

An Overview of the Mimesis Architecture: Integrating Intelligent Narrative Control into an Existing Gaming Environment

R. Michael Young

Liquid Narrative Research Group
Department of Computer Science
NC State University
Raleigh, NC 27695
young@csc.ncsu.edu

Introduction

The Liquid Narrative research group at North Carolina State University, a multidisciplinary group of faculty and students seeking to create interactive narrative environments. To do this, we are designing and building intelligent systems capable of creating structured interaction within virtual worlds that achieves the same kind of cognitive and affective responses to interactive stories as that seen in the participants of conventional narrative media such as the film or the novel. Our approach exploits a well-founded, declarative model of action and intention (Young, Pollack and Moore 1994, Young, Moore and Pollack, 1994) applied to virtual world contexts, in combination with new computational models of narrative structure. The intended result is the production of engaging, story-based interactive applications for education, training and entertainment.

The LN group's current focus is the Mimesis (mim-E-sis) system, an implementation of an intelligent controller for virtual worlds that generates and maintains a coherent, narrative-based storyline. The system combines current research work in AI planning and natural language discourse generation with the real-time control of an existing commercial game system, Epic Games' Unreal Tournament (Epic Games, 1999). The remainder of the paper gives a short overview of the Mimesis system components. Many of the components of Mimesis have already been completed, and we anticipate end-to-end system integration shortly.

Background: Unreal Tournament

Epic Games' Unreal Tournament (UT) is currently one of the most popular titles in the first-person gaming genre. The game's architecture is a client-server design (with single-user mode simply running both client and server on the same machine). UT client processes are responsible for user-side input and output (e.g., keyboard and/or joystick for input, graphics rendering and audio generation as output); the UT server maintains a consistent world model across all users by acting as the centralized controller where the execution of all functions that change the state of the world are coordinated.

The functionality provided by the UT server executes in one of two forms, one modifiable by game licensees (i.e., anyone that purchases a copy of the game CD), the other available only to developer licensees (e.g., those who enter into special contractual agreements with the game developer). Execution of the later, *called native code execution*, involves running compiled C++ executables. The former involves running functions written in *UnrealScript*, a specialized object-oriented scripting language developed by Epic Games and used to define UT object classes and control the behavior of their instances within a UT game world. UnrealScript compiles into byte code and so is platform independent, allowing the same compiled UnrealScript modules to execute on any hardware and OS combination for which UT runs (currently this includes Macintosh, Windows and Linux systems). The UnrealScript process model simulates its own thread-based architecture, providing the Mimesis developer with the abstraction that each UnrealScript object runs its own code within its own thread. Fortunately, game licensees have access to the complete UnrealScript SDK as well as a class browser, a script editor and a level editor (i.e., a 3D world modeler). Built-in support for third-party development of add-on software

(via linking DLLs on Windows platforms)¹ is provided as well, making UT one of the gaming environments more accessible to AI researchers seeking to integrate their work into game environments.

Mimesis Architectural Overview

The Mimesis system integrates AI control with off-the-shelf commercial games. The benefit of this approach for AI researchers is immediate; use of systems like UT provide readily accessible, stable and high-quality graphics, networking, database and process execution support for virtual environments, eliminating the need for time consuming development of these components in a research project.

As I describe below, Mimesis is composed of a number of distinct modules -- some which are integrated directly into the UT clients and servers, others that reside on remote machines and communicate with the UT server via socket-based connections. The UT server, extended to include Mimesis software components, is called a *Mimesis Unreal Tournament Server* (MUTS) and the intelligent controller operating on a remote machine is called the *Mimesis Controller* (MC). The role of the MUTS is to provide low-level access to the UT environment (process invocation, monitoring and control) while the role of the MC is to serve as a centralized intelligent source for the design and (high-level) control of coherent, compelling narrative-based interaction over time.

The principal factor that allows the integration of AI research tools and technology with the existing UT server engine is the sharing between the two of a declarative representation of action and of the conditions within the UT server's virtual world. On the MC side, an HTN planner-style action representation (Erol, Hendler and Nau, 1994, Young, Pollack and Moore, 1994) encodes all actions that can be taken by characters in the virtual environment. On the MUTS side, each of the primitive actions in this representation is mirrored by a procedural definition of UnrealScript code that directly implements the action it mirrors. Currently, the mirroring of functionality between the MUTS and the MC action representations is achieved manually by system designers, although, as I describe in the final section of this paper, much of the process of creating procedural definitions of declarative plan operators is performed automatically.

¹ For example, we have implemented a Windows DLL that links Microsofts Text-to-Speech Libraries with the UT client. The DLL is described in more detail below.

Mimesis Components

In addition to the standard UnrealScript and native code objects provided in the UT server, the MUTS consists of four main Mimesis components, implemented as instances of UnrealScript classes that are integrated directly into the UT run-time environment, running essentially as independent threads within the standard UT server configuration. In contrast, the MC uses one central process, spawning special-purpose intelligent modules as independent threads when needed to respond to specific conditions arising in the virtual world. The various Mimesis components, shown in the diagram in Figure 1, are described below.

Mimesis Controllers are currently implemented in Lisp and run on dedicated high-end machines distinct from the system on which a MUTS executes. Each MC is composed of its own socket-based communications module (for connection to remote MUTSes), a central controller process and a number of individual intelligent support modules, spawned as needed by the MC during the creation and monitoring of an unfolding Mimesis storyline. Additional modules include, for instance, *the narrative planner*, a variant of the Longbow planner (Young, Moore and Pollack, 1994, Young and Moore, 1994) originally developed for natural language discourse generation. Further, the controller creates a world map of the game environment's virtual space, a knowledge base of information about the current state of the environment, etc. The modular design of the controller allows the straightforward integration of new modules as the system's functional requirements are expanded.

When a MUTS server is started, it contacts an MC via a socket connection, and a handshaking protocol ensues in which the MUTS describes the current (pre-game) state of its virtual world. At this time The MC generates a plan-based storyline for the action sequence that it intends to play out during the user's session within the game environment. As gameplay begins on the MUTS, the MC sends commands to the MUTS environment to drive the state changes dictated by the storyline actions present in the plan. These commands may include directives to control character action and communication, commands that manipulate the state of the virtual world directly (i.e., without the use of characters to manipulate the world) or meta-level messages that send information to one or more of the MUTS modules.

The first (and simplest) MUTS-side Mimesis component is the communications module, named the *AILink*. The AILink module handles all communication with the MC, routing messages to the MUTS through its socket-based connection and distributing incoming messages from the MC to the appropriate MUTS modules.

When the MC sends commands for action to the MUTS, these commands are routed by the AILink to an execution module called the *Funcaller*. The Funcaller is responsible for parsing the commands --- arriving via the socket connection as text strings --- and translating them into UnrealScript function calls (i.e., determining the function to call and invoking it using correct references to the data objects named in the command string's parameters). To manage this translation process, the Funcaller uses a pre-defined hash table to link argument string names with references to the appropriate run-time data objects. Once all object references are resolved, an automatically generated dispatcher function is called, the function dispatches based on the action's command name string and calls the appropriate UnrealScript code with the correct arguments to UT data objects.

One complication in the management of action execution on the MUTS is the mismatch between the temporal representation of primitive actions used by the MC and the temporal extent of corresponding procedures executing on the MUTS. The MC's planning system represents primitive actions as occurring instantaneously, while the procedural nature of the code that runs on the MUTS to implement the corresponding primitive actions requires that the actual execution of an action may occur over a substantial period of time. Because users may alter the state of the world during the execution of a primitive action's procedural

Because typical planning-based representations for action assume that primitive actions happen instantaneously, the semantics of their actions' preconditions are intended to describe the required state of the environment immediately prior to the action's successful execution. The representation used by Mimesis distinguishes between preconditions of this type and *persistent preconditions*, conditions in the world that must hold from the beginning of the action execution through to the point of its completion.

The UnrealScript procedures that implement each primitive action are responsible for spawning *sentinel threads*, processes that monitor the action's persistent preconditions during the life span of the procedure's execution. When a sentinel thread detects a change in the state of a monitored state variable, it reports the change to the MC, indicating that the procedural implementation of the currently executing action may have failed. Monitoring these conditions via MUTS modules is necessary because many conditions such as the collision between character and door in the parenthetical example above is handled by UT at a low level; when these events occur, the UT server does not signal an exception that can readily be detected by Mimesis UnrealScript code.

One central research question for the Liquid Narrative

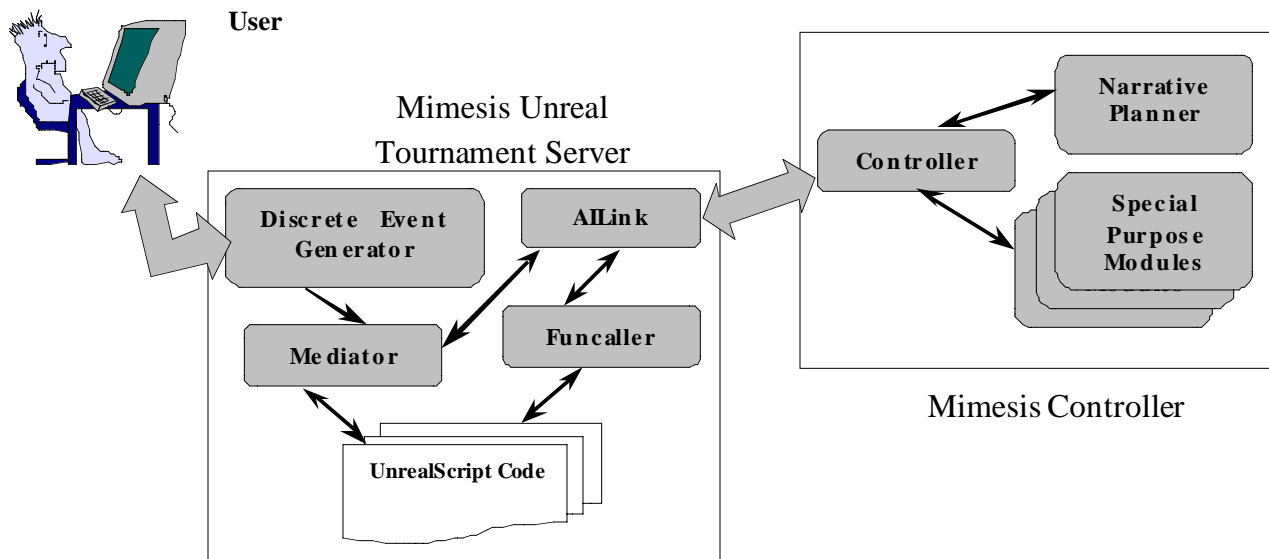


Figure 1. Mimesis Architecture

implementation, it is necessary to monitor each individual system-driven action to make certain that user-initiated state changes do not cause the system's action to fail (e.g., the MUTS must detect when a user closes a door in the face of a character that is executing a **walk-through-door** procedure).

group is how to detect and respond to situations where a user performs an action that does not directly interfere with currently executing procedures, but instead seriously conflicts with the storyline that the MC narrative planner has created. While the Mimesis system can create compelling action sequences for a UT world and then issue

commands that drive the UT world according to the intended storyline, users in the environment may attempt to make changes to the world that can substantively affect the executability of future actions in the plan that drives the action sequence.

To address this issue, we exploit the structure inherent in narrative plans that we create to mediate between the user's commands and the actual execution of those commands in the virtual world. To manage potential conflicts between a user's actions and the MC's storyline, we intervene between the commands issued by the user and their execution by the underlying UT environment. This mediation involves a two-stage process. First, a *Discrete Event Generator* maps relevant input from the user's continuous event stream (mouse motion and keyboard events, for instance) into discrete action commands that match up with the MC's declarative representation of the actions available in the world.

Next, we compare the action that the user is attempting to execute with the plan structure of the storyline that the MC has created, looking for possible conflicts. The plans we create contain a rich causal structure: all causal relationships between steps in the plans are specifically marked by data structures called *causal links*. To ensure that the plans are functionally correct, these links are originally added to a plan at the time that the planning system constructs it. In contrast, we put them to use at execution time; in our system, when a user attempts to perform an action, the declarative representation of that action is checked against the causal links present in the plan. Each action's representation explicitly lists all of its effects – the changes to the state of the world that will occur whenever the action successfully executes. If the successful completion of the user's action poses a threat to any of the causal links (i.e., one of the action's effects conflicts with the condition associated with the causal relationship indicated by the link), an *exception* is raised.

Exceptions are dealt with in one of two ways. The most straightforward is via *intervention*. Because all of a user's actions in the environment pass through the mediator prior to execution, Mimesis itself can determine whether an action succeeds or fails. Typically, the success or failure of an action within a virtual environment is determined by software that approximates the rules of the underlying virtual world (e.g., setting a nuclear reactor's control dial to a particular setting may cause the reactor to overload). However, when a user's action would violate one of the narrative plan's constraints, Mimesis can intervene, causing the action to fail even when its execution in the current world state would normally succeed. In the nuclear reactor example, this might be achieved by surreptitiously substituting an alternate action for execution, one in which the "natural" outcome is consistent with the existing plan's constraints. A control dial momentarily jamming, for instance, will preserve the apparent consistency of the

user's interaction while also maintaining safe energy levels in the storyworld's reactor system.

The second response to an exception is to adjust the narrative structure of the plan to *accommodate* the new activity of the user. When accommodating an exception, the system allows the user's command to execute in the virtual world, then reconsiders the narrative context to look for changes that can be made "behind the scenes" to avoid the original exception's conflict. It may be possible that the conflict caused by the exception can be eliminated by a minor restructuring of the narrative plan, for instance, selecting a different but compatible location for an event when the user takes an unexpected turn down a new path. Alternatively, this may involve more substantive changes to the plan, for instance, should a user stumble upon the key to a mystery early in a narrative or unintentionally destroy a device required to rescue a narrative's central character.

Clearly, the computational process of detecting exceptions is a time-critical task, particularly when intervention is a potential response. There are several techniques that we employ in Mimesis to make exception detection more straightforward. First, the Mediator holds a cache of causal links that are currently relevant to the unfolding plan (i.e., all conditions in the current world state that future steps in the plan are dependent upon). As the storyline unfolds, the MC adds and deletes elements in this cache as appropriate, so that exceptions can be quickly detected by comparing the effects of a user's current action to the cache's contents.

Second, responses to potential exceptions are pre-computed as time permits, allowing the Mediator to select from a list of responses to an exception without having to wait for the MC to generate an appropriate response. Responses are pre-computed based on a dynamically updated list of actions that a user can potentially carry out at the current moment. At any given time, the game's world state will be such that a subset of all possible actions will have all their preconditions satisfied; these actions are precisely the set of actions that are currently available to the user for immediate execution. At times when the MC's processing load is low, it considers each of the actions in this list in turn, determining the appropriate response to potential exceptions raised by the action's execution.

Current Implementation Status and Near-Term Extensions

The Mimesis system is still under development, however, substantial parts of the implementation have been completed. We are currently able to develop plans for interaction using the MC, to send commands for activity based on those plans to the Funcaller and to have the appropriate UnrealScript procedure calls executed. However, the implementation of execution-monitoring is

not yet complete, and the Mediator's functionality is still only partially implemented.

Currently, our group is also working on additional MUTS and MC modules. For example, we have completed a translator module that takes as input primitive action operators and generates UnrealScript class files containing procedure stubs for each action procedure that will need to be implemented for a given domain. The module, named *LBUS* (LongBow-UnrealScript) is used to generate a first-pass automatic class definition; designers are then responsible for filling out the stub functions to implement the low-level behaviors for the various actions in a given world.

Mimesis project members are also implementing a virtual camera controller that will integrate pre-designed idioms for shot composition taken from film with discourse-level direction included in the storyline plans generated by the MC. This work is described in more detail in a separate submission to this symposium (Amerson and Kime, 2001).

Much of the Mimesis design has intentionally been limited to either the MC or the *server*-side Unreal Tournament environment. In this manner, no specialized UT client modifications are necessary to connect to and participate in a Mimesis world. However, client-side extensions (called *mods* by the gaming community) are also easily created for UT, and we have experimented with the addition of text-to-speech capability in UT clients by creating a mod that makes system-level calls to Microsoft's test-to-speech API. The mod allows our system to generate custom text to communicate with each user, then to speak the text as a character or narrator within the Mimesis storyline.

Acknowledgements

The work of the Liquid Narrative group has been supported by a number of sources, including a Faculty Research and Development grant from the NC State Office of the Provost, and equipment grants and gifts from the NC State Department of Computer Science and Microsoft Research's University Grants Program. The Mimesis system could not have been designed or constructed without the constant help of the students that have contributed many hours learning by doing. These include Daniel Amerson, D'arcey Carrol, Shannon Garlick, Shaun Kime, Brian McLaughlin, and C.J. Saretto.

References

Amerson, D. and Kime, S., 2001. Real-time Cinematic Camera Control for Interactive Narratives. In Working Notes of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment, Stanford, CA.

Erol, K., Hendler, J. and Nau, D., 1994. UMCP: A Sound and Complete Procedure for Hierarchical Task-Network Planning. In Proceedings of the Second International Conference on Artificial Intelligence and Planning Systems, 556-560. Chicago, IL.

Epic Games, 1999. Unreal Tournament [Computer Software].

Young, R. M., Pollack, M. and Moore, J. D., 1994. Decomposition and Causality in Partial-Order Planning. In Proceedings of the Second International Conference on Artificial Intelligence and Planning Systems, 556-560. Chicago, IL.

Young, R. M., Moore, J. D., and Pollack, M. E., 1994. Towards a Principled Representation of Discourse Plans. In Proceedings of the Sixteenth Annual Meeting of the Cognitive Science Society, 946-951, Atlanta, GA.

Young, R. M. and Moore, J. D., 1994. DPOCL: A Principled Approach to Discourse Planning, in Proceedings of the Seventh International Workshop on Natural Language Generation, 13-20, Kennebunkport, ME.