AUTONOMOUS AGENTS FOR REAL-TIME ANIMATION

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF COMPUTER SCIENCE
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

James J. Kuffner, Jr.
December 1999

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Jean-Claude Latombe
(Principal Adviser)

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Leonidas J. Guibas

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Christoph Bregler

Approved for the University Committee on Graduate Studies:

Abstract

Advances in computing hardware, software, and network technology have enabled a new class of interactive applications involving 3D animated characters to become increasingly feasible. Many such applications require algorithms that allow both autonomous and user-controlled animated human figures to move naturally and realistically in response to task-level commands.

This thesis presents a research framework aimed at facilitating the high-level control of animated characters in real-time virtual environments. The framework design is inspired by research in motion planning, control, and sensing for autonomous mobile robots. In particular, the problem of synthesizing motions for animated characters is approached from the standpoint of modelling and controlling a "virtual robot".

Two important classes of task-level motion control are investigated in detail. First, a technique for quickly synthesizing from navigation goals the collision-free motions for animated human figures in dynamic virtual environments is presented. The method combines a fast 2D path planner, a path-following controller, and cyclic motion capture data to generate the underlying animation. The rendering hardware is used to simulate the visual perception of a character, providing a feedback loop to the overall navigation strategy. Second, a method for automatically generating collision-free human arm motions to complete high-level object grasping and manipulation tasks is formulated. Given a target position and orientation in the workspace, a goal configuration for the arm is computed using an inverse kinematics algorithm that attempts to select a collision-free, natural posture. If successful, a randomized path planner is invoked to search the configuration space (C-space) of the arm, modeled as a kinematic chain with seven degrees of freedom (DOF), for a collision-free path connecting the arm initial configuration to the goal configuration.

Results from experiments using these techniques in an interactive application with high-level scripting capabilities are presented and evaluated. Finally, how this research fits into the larger context of automatic motion synthesis for animated agents is discussed.

Acknowledgments

First, I would like to thank my advisor Professor Jean-Claude Latombe for his helpful guidance and support over the past six years. Through him, I have been introduced to a wealth of new ideas ranging from motion planning to mountaineering. I also gratefully acknowledge Professor Leo Guibas and Professor Christoph Bregler for serving on my reading committee, and for their mentorship during my time at Stanford.

While conducting this research, I had the opportunity to work with and learn from many people. I am grateful to Steve LaValle for allowing me to collaborate with him, and for offering his thoughtful advice and friendship. I thank Brian Mirtich of MERL, whom I enjoyed working with to write a software library based on his thesis research. I thank the team of Yotto Koga, Koichi Kondo, and Jonathan Norton, who made creating the ENDGAME computer animated short so much fun.

The students, staff, and visitors of the Stanford Computer Science Robotics Laboratory have enriched my educational experience immeasurably. In particular, I thank David Hsu, Lydia Kavraki, Craig Becker, Tsai-Yen Li, David Lin, Hector Gonzalez, Mark Ruzon, Diego Ruspini, KC Chang, Scott Cohen, Mehran Sahami, Stephen Sorkin, Joel Brown, Serkan Apaydin, Jing Xiao, Frank Hickman, Jutta McCormick, and everyone in the motion-algo group for their friendship and for the many interesting and insightful discussions.

I thank the Department of Computer Science for providing the facilities for conducting this research. I am very grateful to Professor Mark Levoy, Professor Pat Hanrahan, and the Stanford Computer Graphics Laboratory for allowing me to utilize their rendering hardware and video equipment.

Financial support for this research was provided in part by a National Science

Foundation Graduate Fellowship in Engineering, and MURI grant DAAH04-96-1-007, of which I am very grateful.

I thank Professor Hirochika Inoue and Professor Masayuki Inaba of the University of Tokyo for allowing me to join their robotics research laboratory in Japan upon completion of this dissertation.

I am indebted to the many friends who have helped me during my time at Stanford. In particular, I am grateful for the friendship of my former roommates, Apostolos Lerios, Daisuke Terasawa, and Iacovos Koumas, who survived living with me. Thanks also to Michael Svihura, Salli Tazuke, and Phil and Mariko Yang who have been such supportive friends. I owe a special debt of gratitude to Bill Behrman, who generously donated his time and energy to help me print and file the final copy of this dissertation.

I thank my family: Maureen, Andy, Teresa, and Joseph, for filling my childhood with wonderful memories. Special thanks to my parents, James and Kathleen Kuffner, to whom I dedicate this thesis, for their unconditional love and for teaching me the value of education.

Finally, I thank my wife Haruko, for her patience and encouragement throughout.

Contents

List of Figures