HPCT difference Simulation -- Inverse Kinematics

As the "Little Man" (See ARM1 and ARM2 on PCTDEMOS diskette) demonstrates so well, HPCT suggests a working model of behavior which can be developed further to provide simulation of virtual figures on the computer screen. In August of 1994, I began to examine commercially available animation software, prompting Bill to re-write a proposal he had sent me. The idea was - and still is as of June 1995 - to approach animation software companies to organize a joint development effort -- Dag

Date: Sat Sep 10, 1994 11:23 am PST
From: William T. Powers
To: Dag Forssell
Subject: Animation

Dag: I've rewritten the original piece in the light of what you said about current animation programs. You will probably want to add more language along those lines.

FROM ANIMATION TO SIMULATION

Animation techniques have progressed considerably over the years. However, the ultimate goal of animation, which (in at least some cases) is to produce realistic-looking movements of living organisms and physical objects, is very difficult to reach. The human eye and brain are finely calibrated for detecting slight deviations from naturalness. Current methods still lack a way of producing natural movements, but a new approach to the analysis of real behavior may offer yet another step forward. We will discuss first the difficulties in achieving naturalism by present techniques, and then outline an approach that may overcome these difficulties. This approach, beside having considerably commerical potential, will also advance the sciences of behavior by a significant amount.

The limitations of present methods

The most advanced animation technique available today for depicting the movements of organisms appears to involve a mathematical principle called "inverse kinematics and dynamics." This impressive term means specifying a desired form of action, and then calculating the joint angles (and perhaps even the simulated forces) needed to create that action. This method derives from robotics and from certain theoretical approaches to explaining the behavior of real organisms.

Before this technique can be applied, it is necessary for the animator to specify a "trajectory" of movement for specific points on the body of the person or animal to be pictured. For example, if a human figure is to reach out and touch an object with a hand, the animator must decide on the path by which the hand will move from one position to another and the speed profile of the hand as it moves along the path, as well as the orientations of the hand during the transition. Given this specification and the desired positions of other critical points such as the shoulder, an inverse-calculation program can derive the specific joint angles needed to produce that result, and generate the basic armatures for key frames of an animation. Other techniques of computer graphics such as "morphing" can then be used to fill in transitions between fully- rendered key frames.

The disadvantage of this approach is that the animator must specify the movements desired, and do so in a way that is convincing to the human eye. This is no easier to do through inverse calculations and morphing than it is through manually drawing the frames -- although of course the actual drawing is far easier.

"Naturalness" means behavior that is exactly like that of a living system in a real world governed by the laws of physics. When a ball bounces, it follows curves dictated by mass, elasticity, gravity, and air resistance. On impact it distorts and compresses, picking up rotary motion if the trajectory is not vertical, and decompresses, its shape changing and going through oscillations as it rises again at a new speed in a new direction, following a new curve. To put together a sequence of animated frames making this bounce look true to life requires an unusual ability to visualize realistically, tracings of an actual ball while it really bounces, or calculations based on physical principles.

Doing this with a moving figure such as a human being by the method of inverse calculations involves keeping track of many times the amount of detail needed to make a bouncing ball behave naturally. As a foot of a walking figure touches the ground, the shape of the foot flattens out, the joints compress, the knee flexes under the weight of the body, the muscles react to the forces by changing their tension and spring constants, the body sinks below the hip and then rises again, and the body responds to all these forces by accelerating up, down, sideways, and around various axes of rotation. The entire body reacts to the forces, both body and arms making adjustments against them for the purpose of maintaining balance. And the human eye and brain, watching an animation of walking, mercilessly detects every deviation from the proper relationships among the body parts, between the body and the ground, and among the temporal and inertial properties of moving masses. If any of these are wrong, the movement simply won't look right.

Using the method of inverse calculations, the animator must still specify movements having spatial and temporal patterns that look convincing; the inverse calculations do not help with this. The technique of morphing fills in intervening frames not on the basis of the laws governing the natural world, but simply as a linear interpolation. Thus the properties of the natural world, and of the living creatures being pictured, still do not come naturally out of the basic method.

An alternative: simulations using control theory

There is another approach which promises to work far better than the methods outlined above. It is based on an entirely different conception of how the behavior of organisms works: control theory.

The theory behind the inverse-calculation method says that organisms calculate the forces required to generate specific trajectories of movement, and then produce those forces in a pre-programmed way. Control theory says that they do not do this at all. Instead, organisms are equipped with kinesthetic sensors that detect the positions of all limbs, and with vision which detects the current relationships of the body to the environment, and uses this information all during movements to control the positions of limbs and the relationships of limbs to external objects. The trajectories of movement which come out of behavior controlled in this way simply follow from the properties of the nervous system, the muscles, and the masses of the body. Of course it is still possible for organisms to produce arbitrary patterns of movement, such as moving the hand from one position to another along an arc, but now the actual properties of the real organism's nervous system and muscles, and the laws of physics, come into play to determine the details of these movements, and how the body must react to such movements.

Achieving naturalness of movement using control theory now depends on having an accurate model of the control systems in the body as well as a model of the physics of movement. However, control system models are far simpler than models involving inverse calculations. In a computer model of pointing behavior, in which a stick figure can reach out and touch an arbitrarily moving target in three dimensions, the computer code that describes the actual control systems governing a joint angle can be written in less than 30 lines of code in a language called "C". An inverse-calculation model that produces equivalent movements, written in the same computer language, would require many hundreds of lines of code. A control-system model that accomplishes pointing behavior in a graceful and realistic way runs on a desktop computer in real time. An inverse-calculation program for accomplishing the same result would require far more computer power and would run far more slowly.

But the main difference is that when people look at the control-system model, they are struck immediately by the naturalness of its movements -- which are not specified by the user at all. All the user needs to do is move the target around; the Little Man follows it with its finger without any external instruction.

The Little Man is a true simulation, not an animation. Its movements arise from properties built into it; they are not specified in advance. This principle can, with enough development, produce animations an order of magnitude better than any that are now possible.

Developing the control-system method

The control system model was developed in the context of psychology as a way of explaining many kinds of behavior. Its major advances have been along lines of predicting simple motor behaviors, and modeling more complex behaviors in a qualitative way. This theoretical approach, being new, has not had many resources devoted to it; most of the large efforts going on now use the inverse-calculation model, which dates back to neurological thinking at the beginning of the 20th Century. As a result, the Little Man demonstration cited above represents the best that has been done so far with very limited resources and few people working on it in an unofficial way.

The next steps in the development are easy to describe, but will require considerable effort to accomplish. What is needed is a mathematical description of a whole organism instead of a single jointed arm. In control theory models, the most difficult part of developing the model is also the simplest in principle: calculating how a physical body with jointed arms and legs attached to a torso will move when torques are applied at its joints in all the possible directions. This involves only "forward dynamical" calculations, not the inverse calculations in the approach described at the beginning.

The most direct way of making these calculations would be not to do them at all, but to build an actual body with actuators at all its joints. This would then be the test bed on which the control systems for generating behavior would be built. Unfortunately, this would be an extremely expensive way to do it, and the required kinds of sensors and actuators probably do not exist. The most practical way is to use a computer to simulate the physics of the body by using well-known equations of physical dynamics.

To do this simulation, however, requires mathematics of a level not available to the developers of the Little Man. The required equations are not to be found in the literature, for the simple reason that most current efforts have been devoted to the inverse-calculation approach, which requires equations in a very different form. The basic principles required to produce this simulation are well-known; what is needed is a physicist-mathematician-engineer (in one body or several) who can go through the labor of producing the very complex set of dynamical equations needed to define the simulated physical organism.

Once this simulated physical organism is available, we will have something like a marionette physically capable of producing all the movements that a real organism produces, given only something that can pull the strings in the right way. Actually, any proper development of an inverse-calculation model would require the same test-bed marionette, but that approach, because of other great difficulties attendant on implementing it, has not gone this far.

Given the simulated marionette, the next step is to work out the control systems that are to operate it. This has already been done for a single simplified limb, with excellent and realistic results. A start has been made toward extending the simple two-joint limb to an arm with 14 degrees of freedom, including a hand that can move as real hands do. The same model, using different parameters, would apply to legs. So the development of the control system model would not start completely from scratch.

There are different levels of realism which can be sought. For example, one version of the Little Man includes a model of muscles and of the stretch and tendon reflexes, closely following the architecture of real human bodies. But it is not necessary to model the muscles; the equivalent behavior can be programmed into the model, with only higher levels of control being explicitly modeled.

Similarly, in the most advanced version of the Little Man, the visual systems actually use optical ray-tracing, and depth perception is accomplished by comparing images in a left and right eye. The result is that the model contains perceptions much like those in a real system. A short-cut, however, can eliminate this level of detail: it can be assumed that in the brain of the simulated organism, there are signals representing the actual x, y, and z positions of objects in the environment (including the simulated organism's own hand). Since we will have calculated all those positions in the process of setting up the behavior of the marionette, we do not need to try to reproduce the method by which the brain extracts this sort of information from optical inputs.

So by judicious use of information which we, as modelers, have about the environment of the organism, we can by-pass detailed problems of stimulation without sacrificing overall realism. Once the overall behavior of the model has been brought to an acceptable level, we can, for scientific reasons, go back to fill in the black boxes which have been left unopened in the model. That can be treated as a separate effort.

The scope of the project

Carrying this project to a successful end will, we estimate, require the full-time effort of a team of six people for three years. Computing equipment and a place to work will be needed. The total cost will be in the neighborhood of two million dollars, assuming an average annual salary of $60,000 per person for six people, $200,000 annually for administration, and about $400,000 total for computer graphics equipment, workstations, physical plant, and logistics.

A realistic goal would be to simulate walking of a human or animal figure, with simple interactions with environmental objects such as pushing them, grasping and lifting, and orienting. The control system model behind this effort is inherently modular; if this level of performance can be reached, building even more complex organization on it will then be a matter of creating more modules to use the ones already developed. By the end of the proposed project, we will have a very good picture of the next steps to be taken. It is quite probable that usable results will be obtained before the end of the project. However, since we are in effect launching the practical development of a new theory of human behavior, the true end-point can't be defined: there is more to learn than a handful of people, even very talented people, can accomplish in a lifetime.

A practical approach

There is already a world-wide organization of scientists and other professionals interested in the development of control theory as a model of human behavior. The American contingent formed The Control Systems Group ten years ago, and a European Control Systems Group held its first meeting in Wales in the spring of 1994. The American group incorporated as a 501-3-c not-for-profit corporation in 1994, and one of its members is currently organizing a fund-raising effort for the establishment of a Center for the Study of Living Control Systems.

The project described here could be set up as one of the initial efforts to be carried out by this Center, under contract. An advantage to the client would be that much of the cost of setting up this project, and part of the operating expenses, could be treated as a contribution to this not-for-profit Center. The result would in fact amount to a major contribution to science, even though the immediate goal is a commercial development of a method for creating better animations for entertainment and instruction.

William T. Powers Sept. 10, 1994 Durango, CO

Date: Thu Mar 30, 1995 2:19 am PST
DAG: Copy of Lucas letter for your info

March 26, 1995

Dear Mr. Lucas;

I am writing to enlist your interest and support in connection with a project that is (from my viewpoint) of great scientific importance, and (from yours) of potentially great artistic and commercial value. In a nutshell, the objective would be to develop a method of creating simulated organisms that behave in a natural way, in the form of either physical models or computer simulations.

While I don't know everything about the current state of your art in this field, everything I have learned indicates that all approaches to generating realistic behavior are based on the concept of generating movements by means of stored programs or by using human operators attempting to control models so they show realistic movements.

The problem with this approach is that it depends on human artistic judgment to decide what is a natural trajectory of movement. Even the latest gimmick in animation (of which I know), inverse kinematics, requires that someone decide on the detailed trajectory to be followed by hands, feet, head, and so forth. While artists can be very good at this, they are still only approximating the movements of real people or animals (or Things). The human eye is unforgiving of imperfections in these approximations, however willing the suspension of disbelief.

The ideal approach would involve not animation but simulation. Simulation involves giving a computed object the same physical properties as the real one, and letting the object itself interact with its changing environment in the way that is natural to it. If the object is an organism, we should carry this one step further: we should give the object the same internal means of creating behavior that the real organism has, or would have if it were real. And therein lies the contribution to both science and art that I think we can make by working together.

The current accepted wisdom in mainstream science is that behavior is controlled by commands from the brain. These commands start with the general and finish with the particular -- signals sent to muscles. Behavior is thought of as the creation of movements. When models are constructed on this basis, they become very complex and slow, because the simulated brain has to do the very same thing that the present techniques of animation do: plan out the path, every acceleration and every velocity, all during a movement from its beginning to its end. One result of this approach is that the behavior of these models looks a great deal like an inexpert animation. A graphic illustration of the problems was seen recently in videos of Dante, the robot that descended into a volcanic crater. It could take only about 5 steps per hour, spending the rest of its time and vast remote computing power on planning the next step. It ended up on its back and had to be airlifted out.

There is another approach which began way back when cybernetics was born: an approach based on the principles of control theory. For reasons which have remained a puzzle to me for 40 years, this approach has never been accepted by mainstream behavioral scientists. It is basically very simple. The idea is that organisms form internal perceptual representations of the world they live in, and learn to produce actions that will make the perceived world match reference specifications inside the behaving system. Such systems, basically, are organized to produce any amount and direction of action that will produce and maintain the correct perceptual result. Regulators and servomechanisms are precisely of this nature; there is little mystery in how they work or how to design them for specific purposes. These principles offer a natural basis for a theory of goal-centered purposive behavior.

I've been working on such a theory for most of my professional life and on into retirement. The going has been slow, but over the past twenty years increasing numbers of scientists and other professionals, now in the hundreds, have begun to realize that this theory holds great promise as the next step in understanding behavior. That explains my interest in proposing this project: I am still trying to get mainstream scientists to understand what this model can do, not so much (at this late stage of my life) for my own sake as for the sake of many people who have joined me in knocking at well-defended doors. I think you could help us amplify the volume of the knocking.

This theoretic approach is called "Perception Control Theory", or PCT. I will not shower you with details. What matters here is that in trying to communicate how it works, I have developed many interactive computer demonstrations and simulations which show phenomena of kinds that conventional theories have great trouble in explaining. One of these simulations, the most advanced that I have been able to put together with my limited abilities, is a model of pointing behavior. This model was intended to show how much simpler the PCT approach was than the command-driven approach then and still being touted as an explanation of how people can reach out and touch targets. It was this model that gave me the idea of proposing PCT as a replacement for animation.

The Little Man model is a stick figure on a screen. It has a head with two eyes, and one arm that can swivel up and down and side to side at the shoulder and bend at the elbow.

A triangular target floats in space near the figure. The user can use a mouse or the keyboard to move this target around in three-dimensional space. As the target moves, the figure tracks it with its "fingertip." The user gives no commands to the Little Man; all that is needed is to move the target; the Little Man will track it, wherever it is and however it moves. The movements of the arm arise out of the internal organization of the Little Man; there is no planning of trajectories. There are (for your technical kibitzers) no inverse calculations at all.

One result of the simplicity of the model is that it runs in real time; in fact it has to be slowed down to avoid acting unnaturally fast on my desktop computer. But for our purposes here, an even more important result is that people seeing it for the first time remark on how graceful and natural it looks. It is clearly not an animation. As more than one person has commented, the damned thing looks alive. I was pleased with that comment because I am trying to model living organisms.

So I have a beginning for the project. What is needed now are resources: intellectual and practical. I know what needs to be done next, but I don't know how to do it. You have people who have immense skills at programming and model-making, and you can hire consultants to carry the math to the next level that is required, one step past where I stop. There are many stages to this project, from making small improvements and aids to present techniques on to vast improvements down the line. I can teach your people what they need to know to go beyond where I have gone.

What about it? If you're interested, I will be on the West Coast in the latter part of April, taking part in a presentation at the American Educational Research Association annual meeting in San Francisco. My wife and I are planning visits to friends and colleagues in LA, La Jolla, and San Diego, and have no particular schedule. If that's too rapid (I understand that you have one or two other things to do), we can set up a meeting later. I can bring along my simulation programs and run them for you if you can scare up a 486 computer with a VGA screen. Or you could shunt me off to some of your technical people (although I would really like to start with you and see your face when you grasp what I've been up to).

Let me know.

Best regards,

William T. Powers (Bill)

73 Ridge Place, CR 510
Durango, CO 81301-8136
(Internet) powers_w%flc@vaxf.colorado.edu
(303) 247-7986

[Bill received a polite no thanks from Mr. Lucas' secretary - Dag]