Paul Kaiser of Riverbed Interviews Michael Girard and Susan Amkraut

Paul Kaiser of Riverbed Interviews Michael Girard and Susan Amkraut

Ideas on Character Studio

2002-04-13

When I first met Michael Girard and Susan Amkraut in 1993, I asked them whether it was possible to make a 3D world out of children's drawings. That question led to our close friendship and ongoing collaboration. We recorded this conversation in March 1988, and then revised it the following month.

PK:
How did the two of you start trying to simulate human motion on the computer? What was behind this particular obsession?

SA:
Actually, we didn't start with moving figures. We started with moving abstract shapes.

MG:

Yes, only gradually did we zero in on human movement in all its complexity.

But to go all the way back to the beginning, when I first met Susan, she was a printmaker. She took gestural drawing techniques and put them into an algorithmic framework.

What she was doing with prints was akin to what interested me in computers, which was to work from very primitive sources - tactile, visceral sources that become the primitives in the construction of a piece. We were both interested in things that really hit you in the gut but that aren't well defined formally. We wanted to build an architecture around those primitives that would be extremely baroque and complex.

Now, human figures have very strong emotional and psychological qualities. There is nothing more visceral and immediate than a body moving. When we see the human body move, we feel it -- we don't just see it as a visual image, we also feel it as a muscular experience.

Try to choreograph a dance without invoking any emotions at all, try to make it as robotic as possible, and still -- since you're using the human body as a primitive, you can't get away from its expressive power.

PK:
So you began not by considering the human figure but instead by addressing more fundamental questions of computer animation. How did you begin? And how did your approach to computer animation differ conceptually from traditional techniques of animation?

SA:
I had made a print suggesting simple parallax, which is just the perceptual principle that when things are moving far away, they not only look smaller but also they seem to move more slowly.

By this time, I was already programming software. I started wondering how you could vary these parallax relationships on the computer -- what would that do to your perception? With this in mind, I wrote a primitive particle-system graphics program, in which I could have small things moving quickly and big things moving slowly -- just to see what this reversal of parallax would feel like. We were working on pretty primitive machines back then (this was around 1978) so the results weren't very interesting -- but this was the kind of idea we started exploring.

PK:
These basic perceptual issues of distance and parallax - aren't they the concerns of Chuck Jones and others doing traditional animation?

MG:
The Disney and Warner Brothers cartoonists exaggerated movement in order to convey meaning. They perfected a kind of moving caricature.

By contrast, what we've sought in computer animation is to open the door to a new type of animation -- one in which we can focus on the subtleties and the micro-structure of motion. Slow it down, put it under a microscope.

To do all this, however, we had to begin thinking about programming new systems. So it was at that point that we focused on software development. We studied the whole problem of system design, and then we began designing tools.

As we did so, we realized that the boundary between art technology and art itself was blurring. To clarify, I think most people working in fine arts today take their tools for granted: they have pencils, paint, canvas, cameras, etc. Stable technology. But with computers, the technology changes rapidly and fundamentally from year to year. And as you have more intelligent kinds of behavioral systems, the tools themselves suggest new kinds of artistic content.

PK:
But to return to your story. As you started to build new systems and software architectures for creating new works, what were your exact steps? By this time you'd moved from UC Santa Cruz to Ohio State, right?

SA:
Well, I got there a year after Michael, in fall '84, and we were just excited about having the equipment to be able to actually create an animation, which we'd never been able to do before. We got to work right away on a piece called Hidden Agenda. It explored the use of splining algorithms, which were relatively new back then.

A spline is a set of points that defines a curve or a surface. We were interested in deforming shapes, so we were warping 3D objects to the curve or the surface of a spline. We were also playing with the idea of subjective contours, which are lines that the mind fills in from a visual composition. We wanted to create this illusion in 3D space, having seen it already in 2D illustrations.

PK:
So what was the hidden agenda of Hidden Agenda?

SA:
Well, that the splines to which the shapes were being warped were invisible. It gave the feeling of a hidden surface that was only defined by implication, by the visible influence it had on the shapes around it.

PK:
The basic principle of subjective contours is a well-established Gestalt rule of perception, which you translated from the plane to the cube. What interested you about an implied rather than a visible representation?

MG:
The goal in our field at the time was to create photo-realistic imagery. We were after something else. What interested us was that the mind has degrees of freedom, and that there are many different possible interpretations of a sensory event.

When you read a book, your mind reconstructs a very elaborate scheme internally. What makes the artistic experience compelling is that it's a very personal and sophisticated interpretation. We realized that with computer animation we could make renderings that suggested much more sophisticated spaces than could be seen, and in the process gave the user more degrees of freedom.

A general problem with computer animation is that if too much of it is filled in for you, you see it once and don't bother with it again. It's over-determined. It's not like a piece of music that you can listen to over and over again to deepen your experience each time.

SA:
It was this realization that helped us change direction. We lost interest in visual illusions per se, because they were too gimmicky and also too precise. There was only a fixed way of reading them, and they didn't have soul.

So Michael began playing with the human figure, which did.

PK:
Animating human figures seems about the hardest thing you could have tried.

MG:
I didn't start with humans, actually. The first system we developed let us build creatures with any number of legs.

How to animate such creatures convincingly? It's not enough to show a form moving, we had to go beyond what's visual to what's physical. One interprets what one sees in terms of acceleration and gravity, so we realized we had to depict that play of forces. This meant we had to understand physical movement from the inside out.

So for many years, we had to put aside our interest in building baroque architectures and processes, and instead focus on building a system that has gravitational dynamics. This was a huge problem in computer animation. We started by looking into robotics and physically-based dynamics.

We were fortunate that Ohio State University had one of the premier programs on walking and running machines in the country. Marc Raibert's work on running machines was particularly amazing.

Some of the same algorithms used for making robots walk and run could also be used in computer animation, especially the crucial notion of inverse kinematics.Up to that point in the computer animation field, programmers had relied upon the concept of joint angle rotations. Their idea was to imitate what traditional animators do ' that is, to come up with a set of keyframes and then interpolate what comes in between. In cartoon animation, you have an army of animators to take the keyframe drawings and fill in all the missing frames in between. So the idea was that the computer could take over that task.

It didn't work out very well. The computer interpolation techniques used mathematical functions that didn't capture the physically-based dynamics that even a traditional animator would fill in. The animator, with an instinctive grasp of real-world physics, would do much more than the computer program, which simply calculated continuity. These early interpolation routines didn't even incorporate basic Newtonian mechanics.

So back in the 80's I may have been one of the first to introduce gravitational dynamics into the interpolation process. Since then, almost all of the research has been devoted to questions of physically-based dynamics -- which I have to say I now believe to be an over-correction. The human being has such a complex control mechanism -- the human brain and the nervous system -- that simply modeling it as a physically-based object controlled with forces doesn't fully capture its real intelligence and sophistication.

SA:
I remember that back at Ohio State, as we were starting to address these problems, Michael became obsessed with studying movements in extreme slow motion, trying to see exactly how they worked. We'd watch hours of this stuff ' I was always saying, Speed it up!

MG:
We were starting to build simulations -- at first not of two-legged figures, but of multi-legged creatures, which were a little easier because at least one foot would always be on the ground.

We were posing crazy questions to ourselves, like this one: Suppose you had a three-legged animal and it was rotating constantly while it was moving forward, what would that gait look like? We'd developed gait pattern algorithms that were partly based on the walking and running robots that had been developed at Ohio State and MIT, among other places. Our algorithms were somewhat similar, although I had developed a gait-shifting algorithm that was more advanced than any others.

I enjoyed the gait-shifting problem because it was almost like a musical question. How do you go from one rhythmic pattern of leg movements to another? Add to that the fact that you have a physical entity -- a walking or running multi-legged figure ' which you have to move in a physically credible way.

For gait-shifting, the concrete problem was: Given any two patterns of movement, figure out what the phase shift should be in each of the legs, such that it prevents a leg from going past its kinematic limits -- past the point where it would rip off if it had to stay on the ground. A creature has to move fairly quickly from one gait pattern to another in a natural way. Now there's an infinite number of phase-shifting strategies that will take you from one gait pattern to the other. I solved for the one that minimizes the dynamic fluctuation of the body as a whole, taking the shortest distance in the smoothest way.

I became fascinated with the whole idea of optimization, that is, how do you go from one state to another in such a way that you optimize a given function? So even with something as simple as an arm motion that starts in one position and ends in another, what are the sets of positions of that arm that will minimize the jerk?

This leads to an interesting concept. At least in these simple terms, I think our notion of grace can be formally described. What is it about when the ice skater hits the ground cleanly, why do we think it looks more graceful to minimize the jerk? We don't want to see the wobble. So in animation you can take any set of motions and redefine them so that they satisfy some optimization criteria.

SA:
At one point at Ohio State, Michael had a button that we called the Animation Beautifier. We'd have different animators work on an animation in the human animation system, and then they'd push the button and wait for it to churn away. And then, in optimizing the motion, it would remove the jerk.

PK:
Is your gait-shifting algorithm such that you can define any two gaits and automatically interpolate between them?

MG:
Yes. Given any two periodic (repeating) gaits, you can shift from one to the other. This works for even the most eccentric kind ' you can go from a complicated hop- hop- skip gait, for example, to simplewalking.

We discovered that even with an artificial system like ours the phase-shifts from a walk to a trot to a canter to a gallop to a bound are minimal phase-shifts. Whereas if you go from a walk directly to a bound, it takes more phase-shifting, which is much harder for a creature to do -- certainly much harder to do gracefully.

PK:
I recall the huge controversy in the nineteenth century about horses galloping -- did all four feet go off the ground or not? -- that couldn't be solved until the invention of photography let people freeze that moment, something the naked eye couldn't do. Were you conscious of the history of motion research as a backdrop for your investigations?

SA:
Oh, yes, we studied Muybridge and Marey. And Ohio State had a great library of motion on film. For example, they had a slow motion film of a kangaroo mouse on a treadmill. You know, it hops like a kangaroo and it's just incredible.

MG:
Watch a rhinoceros run in slow motion! I remember being just awestruck by its grace. That's where slow motion became fascinating to me. You could watch something so big and heavy and bulky, and see it move so gracefully.

All the motions that human beings do are optimized by the human nervous system. If you take even the most mundane motions and slow them down and look at them ' which is what's been done in film -- you're immediately struck by the beauty and grace of the body.

But to answer your earlier question, the problem we had before us was formally defined by what we were trying to achieve with computer animation. So it's not as if we were originally inspired by the 19th century investigations, though that work became important as reference material to understand more deeply what was going on. It was a wonderful experience to probe deeply into that, which helped uncover this hidden world of what happens when you move. We also worked with Ohio State's dance department, particularly with a dancer named George Karl who was in both the computer animation and dance programs. He opened our eyes further to the subtleties of human movement.

The more we work with human movement, the more we realize how enormously complex the dependencies are between all the parts of the body. Even something as simple as walking is really hard to capture in a formal model. You know, I think even to this day no one has a really complete understanding of how we walk.

PK:
Now as Michael was delving into locomotion, you were addressing an entirely different question, weren't you, Susan?

SA:
Yes, I'd started working on the problem of flocking. Whereas Michael's project was to look at the human body in motion, mine was to take a mathematical algorithm and to see where it could lead.

I'd begun by animating particles using force-fields in 3D. These force-fields would attract, or repel, or shape the movement of the particles. So, for example, I could have a sink that drew all the particles in, or a source they'd funnel out of, or even a spiral they'd fly around.

I soon saw how this could lead to an elegant solution for flocking. The problem posed by flocking is this: you have multiple creatures who don't want to run into each other, but also want to stay very close together -- and they have to avoid hitting any external obstacles as well.

Algorithms had been developed in which a lead bird guided the flock. But real flocks behave in a more interesting fashion: they have no leader. So, neither did my algorithm, which worked like this. I put a little force-field around every bird, so that if any other bird got near, it was told to go away. And of course each bird had a corresponding attraction field, so that if the other bird got too far away, it was told to come closer. So every bird at every frame of the animation considers every force-field around it, and moves accordingly. It's a difficult algorithm to work with because you can't tell where you are at any given point in time unless you know where you started and have computed all the way back up from there.

My interest in this went beyond wanting to simulate actual flocks. I wanted to create a flock of birds all flying realistically as individuals, but flying collectively in patterns that could never happen in the real world.

PK:
When did your investigations begin to merge with Michael's?

SA:
In a film animation project called Eurythmy, which featured flocks of birds and moving human and animal figures.

MG:
Eurythmy posed an important question for us, which still concerns us today. How much is the aesthetic experience dependent on the complexity or organization of the process that created it?

Eurythmy gave us the chance to see what sort of patterns could be created with our software processes. We both shared very profound experiences of seeing patterns in nature, such as the interplay of light on the ocean surface. We could see that natural systems had complex patterns embedded in them, and Eurythmy was our way of exploring that.

For example, we played with the audience's dawning perception that the animated flock of birds is an abstraction, a complexity of natural motion transformed into algorithmic and synthetic motion.

SA:
The patterns of flight in the beginning of the film are relatively normal and look realistic. But later you end up with this sea of birds at knee level and then a huge pillar of birds rising up into the sky -- two things you would never see in reality.

MG:
Another project we did was called Menagerie, which Scott Fisher was producing for the Pompidou Centre in Paris. Scott's interest was in virtual reality, so the piece had to work in real-time. That meant that we had to be able to control all the moving creatures without pre-animating them. The challenge for us was to take the gait-shifting algorithms we'd written and figure out how they could be computed on the fly.

Menagerie was the marriage of Susan's flocking system with my locomotion gait system. We realized that if you have running animals, you want them to herd in some way, so the flocking algorithm that Susan had applied to birds could be applied to animals as well.

So we had animals running around, responding to the vector force fields that surrounded them in the herd -- and also that surrounded you, the virtual viewer, who was exploring the space.

PK:
Yes, it's a beautiful interaction. I remember that as I approached the birds, for example, they would fly away from me and land a bit further on. More interestingly, as I approached an animal, first it would walk slowly away from me, but if I kept following it, its gait would shift to a canter and then to a gallop, leaving me in the dust, as it were.

SA:
That was our first sophisticated rule-based behavioral system. That is, a system in which you have a set of conditionals, of rules that say if such and such circumstance happens, then the system will behave this way. In the case of the animals, if you are within a given distance of them for a given period of time, they would shift gaits and run away from you. Or with a different set of rules, they might approach you as if attacking. We could develop different scenarios by making different sets of rules to drive the animation.

It's well known that if you have enough rules in a system, you can quickly arrive at a point where you can't predict what's going to happen. Menagerie was actually something of a chaotic system in that you could not predict the location or the gait of the animals twenty seconds from the time you did something. So it was always a different experience. As a simulation it was very much alive.

PK:
What also impressed me deeply in Menagerie was that even when the creatures were moving very far away, I could still discern how they were moving. I would see these few dots in the distance, but could tell that that cluster of pixels was galloping. For me, that was the real proof that your gait algorithms worked -- even with a figure that was practically indistinguishable, I could still read it clearly.

I think human beings are genetically programmed to read certain kinds of things in the real world, even given extremely minimal sensory imput. One can see a figure far, far away and still know something about it. You can tell if it is a man or a woman, for example, or whether it's approaching or retreating. We have very sharp eyes for this.

But to continue. Did the programs you wrote to generate Eurhythmy and Menagerie eventually evolve into your figure animation software Character Studio?

MG:
Eventually. But only after I came up with the idea of using footsteps and footprints in time and space as the key control mechanism. That solved a lot of conceptual problems for us.

PK:
You've patented this technique of footstep-driven animation. What made it such a great solution?

MG:
Before artists can start working with motion on the computer, they'll need fine control over every element of the motion. The keyframe techniques of the past just don't give them that. Just getting a figure to walk is so tedious! Assuming you know what you're doing, you still have to specify every detail by hand, and it takes forever.

I wanted to come up with a much higher-level -- and visible -- control mechanism. I was very interested in rhythm, and so the idea of footsteps -- something that dancers think of automatically -- seemed like a natural starting-point. You'd start your animation by specifying the footsteps, a set of patterns in time and space.

PK:
But now you've gone further than footstep control of single figures. With the Motion Flow Editor, you've created an even higher level means for controlling and choreographing multiple motions for multiple figures. Can you describe how motion flow networks operate?

MG:
A general strategy for the architecture of complexity is to build components first, then combine those components into larger wholes. A major problem in animation systems, and of motion design in general, is that there has been no building block other than the keyframe, which is a single posture of a single element at a point in time. Now we have building blocks or components that are entire motion fragments or motion clips.

With the Motion Flow Editor, you can construct complex sequences of motions from motion fragments or "clips".

Let's say you set up a motion network of a clumsy monkey who keeps tripping and falling down in the forest. In the Motion Flow Editor, you can set up any number of paths (called "scripts") through different clips of his motion, so that you might have him walk, run, trip on a rock, fall down, stand up -- then walk, run, trip on a rock, fall down, stand up again, and then throw a tantrum.

One can make a cybernetic set of possibilities given any set of motions -- of how you might connect them. We developed the Motion Flow Editor for you and Shelley to use with Merce Cunningham and Bill T. Jones. But we also did it with an eye towards video game developers. It was interesting to see how their needs converged with those of choreographers.

Game software often contains a figure animation engine, which lets the user control the motion of any particular animal or figure in the game. Depending on what buttons the user pushes, the game engine will branch to a different motion, blending the first motion into the second. Character Studio allows game developers to explicitly represent that architecture of their game engine.

Now the choreographer faces a challenge similar to that of a games developer. Given a set of motions, how do you build an architecture for a larger ensemble of motions? How do you string them together? Starting from one motion, you could consider any number of possible movements to join it to.

PK:
As Merce, Shelley, and I have worked with the Motion Flow Editor in making Hand-drawn Spaces, it's astonishing how well the software interpolates between one motion and another -- a capability that builds on your earlier work in gait-shifting algorithms. If we connect a clip of a dancer hopping in one direction to a clip of him pirouetting in another, your program almost seamlessly joins the two together. So, just to make the power of this clear: Normally in video or film editing when you splice two clips together, they just meet head on in the cut. But when you splice movement together in your new software there is a complex gait shift in which the movements are interpolated much more smoothly over a range of frames that you can specify.

Now let's talk about Brouhaha, a project that I think points to the future. You were working on this when we first met in 1993, but to my knowledge you've never shown it publicly. Can you describe what it was all about?

SA:
With Brouhaha, we built on the idea of subjective contours. We wanted to suggest an element that the viewer would perceive only by its absence. More exactly, we wanted to convey the sense that behavioral motion was embedded in space itself.

PK:
How? As zones of influence?

MG:
Yes, but it was a little more complicated than that. Imagine a field of very simple limb-like structures, a field of robotic limbs or manipulators. A sphere of influence moved through the field, with a behavior associated with the center of that sphere. As the sphere passed over the limbs, the limbs would begin to act according to that behavior.

So the vector force field no longer was just a force in space, a direction -- it became a behavior. We wanted to know if a viewer could perceive the contour of that behavioral space moving through the field of objects.

We used limbs to convey the sense of a real material, something palpable. When one limb poked another, the fat and muscle would deform in response to the trajectory of that stroke as it touched. We were somewhat successful on the algorithmic level, but what we didn't anticipate was that the actual interaction of limbs looked grotesque. We had a difficult time coping with that unintended visceral result.

SA:
It gave us nightmares.

PK:
What behaviors were borne along by these moving spheres of influence?

MG:
There were several. One type of behavior was stroking: a limb would stroke another limb in different ways -- a tickle stroke, or a long, broad, gentle stroke. Another was circling: the limbs would circle one another without touching. We made these behaviors into a single cluster, and then assigned them to the same sphere of influence.

Then there was wind, which brought hitting. The limbs hit each other very hard.

In some cases we moved the forces through the space. For example, we took a force-field that started very small in the center of this group of limbs but then rippled out.

In the last scene of Brouhaha, we had spheres of influence that were moving in orbits through the fields, and you could see the chain reactions of different behaviors propagating through the field.

PK:
Let's turn to the present -- and the future. One thing we've spoken about before is the idea of emergent structure -- for example, a choreographed dance that emerges from a set of interacting rules. Could you give an example of a rule-based system for generating choreography?

MG:
One way would be to classify different kinds of motion, and then sequence the motions according to their classifications. For instance, let's say you have just two classes of movement: walking and running. Imagine a simple rule that says that after moving through a certain number of walking phrases, the virtual figure must shift to running. Then a similar rule would say how many running phrases the figure must perform before shifting back to the walk. Of course, such rules could be considerably more elaborate. At the very least, you'd want the rules to be probabilistic rather than fixed.

However, a more sophisticated approach would use the software to solve for certain constraints. If you're trying to complete a certain rhythmic pattern, for example, you can ask the software to search for matching pieces among thousands of samples. Let's say you want a hopping motion to maintain rhythmic integrity over a specific length of time in order to counterpoint what the other dancers are doing. The program gives you the chance to specify what global rules can generate such a sequence. You can then have the computer construct the network, selecting the connections and transitions over time.

If you have a number of dancers on the screen, you can split the density of movement at any given point. In such a case, you can have the system choose only those phrases whose footpaths won't cause the dancers to collide with one another.

MG:
You can also extend the notion of Brouhaha to Character Studio by creating spheres of behavioral motion that flow through a particular network. Imagine a sphere of influence that describes extremely frenetic motions, and then a network that consists of more languid movements. If you set the frenetic sphere in motion through the network of languid movements, you'll begin to see these invisible behavioral forces move through the dancers

themselves, whose slow, graceful movements can suddenly be transformed into fierce, powerful movements, and vice versa. You could even have this transformation follow the path of a spline, even something as simple as a sine wave. The movements would still look very realistic, for there would be no distortion of the individual movements. But the over-all pattern of movements would be strikingly unusual, unlike anything you'd see in the real world.

PK:
You can also move a field through just the arms or legs of a dancer, as opposed to his or her entire body, correct?

MG:
Yes. You can create different cycles of motion-control networks for different body parts. For example, you can have a cyclic loop of six or seven movements running through the upper portion of the body, and another set running through the lower portion. And you could set put these loops slightly out of phase with each other. Bear in mind that the motions of the upper body will affect those of the lower, and vice versa, so the interaction of these two loops has complex effects on the body as a whole.

PK:
If you were a choreographer, how do you think you might apply these ideas?

MG:
One possibility that fascinates me is to capture the uncontrolled movements of an ordinary person, even of a young child, rather than those of a highly trained dancer. Spastic, frenetic movements are far more complex than controlled ones, and so they're much harder to replicate. However, with motion capture we now can replicate them on the computer. And so I can imagine arranging these more ordinary and spontaneous movements in an extremely baroque architecture.

Incidentally, I think that choreographers have avoided such movements for good reason. The refined language of ballet, and of dance in general, allows choreographers to work with movements that are simplified and repeatable. They can have their dancers perform the same movements, or vary those movements very precisely, to make a dance that can be danced exactly the same way again and again.

PK:

This notion of a refined language of dance -- or "motion alphabet" -- goes even further. Besides being a simplification of elements that can be controlled over time, it illustrates perfectly the notion of grace as you defined it earlier…