The Stuff of Dreams[1]

Alvy Ray Smith

The Stuff of Revolution

I like to state Moore's Law as "10× in 5"—everything good about computers gets better by an order of magnitude every five years. This beats its traditional wording "2× in 1.5", a doubling every 18 months, because an order of magnitude is more than simply larger—it implies a conceptual leap. So looking back 25 years implies five Moore's Law cycles, five orders of magnitude, or five conceptual leaps! To calibrate yourself, compare concepts of the size of the Earth as we progressed from 50 mph airplanes, as commonplace, to 500 mph airplanes. Now think of 5,000 mph air travel. The "Law" describes revolutionary stuff.

To test the Law, I check it against my own experience. My first patron, Alexander Schure, purchased the first commercial RGB graphics card for us over 20 years ago at NYIT for only one million dollars—at video resolution and in current dollars. The Law tells that the price of a graphics card today should have lowered 10,000×, to about $100. Right on. Incidentally, the volume of that first "card", about the size of three refrigerators, shrank 10,000× also, to an actual card.

Chip architects tell me, and even Gordon Moore has said, that there are 10-15 more years in the Law. (It might slow down to "2× in 2", about "10× in 6.5", but I'll ignore this correction). After that we hit the quantum mechanical wall and invoke various religions to keep the faith—for example, neural networks, massive parallelism, or quantum computers. But we can rather confidently predict another 100× to 1000× improvement in the computer world, and hence computer graphics futures.

We don't exaggerate—given a technology charging through seven or eight orders of magnitude in one or two generations—by describing our era as a bona fide revolution. Moore's Law characterizes its raw material, the stuff or fluid of progress, in place for another decade or more. What have we computer graphicists done with it? And what shall we do with it now?

Realized Dreams

I shall list a couple of my dreams and one of a close friend which have already been realized. You might have a different list, but I think these represent collective desire.

---

[1] Originally titled *25 Years = 100,000×* in draft.

My colleagues and I once dreamed, 20 years ago, of a completely computer generated movie. This dream had two sides—as does computer graphics in general—its pixel and its geometry aspects. *Rescuers Down Under* (and every Disney "cel" animation since) realized our dream on the 2D imaging side, as did *Toy Story* for 3D geometry.

We dreamed that one day computer graphics would be a fundamental tool of filmmakers. To some extent, the films above established this, but with the live-action film *Titanic*, I sense true arrival. I have been embarrassed by the use of "our" tool for simple-minded, gee-whiz effects and am gratified that storytelling is now being accomplished where no other technology suffices and technique is a side dish (though a savory one).

An old girlfriend of mine, a medical student, used to beg me to do something "real" with computer graphics, such as simulate human anatomy for budding physicians more lifelike than flattened, yellowed, fluidless cadavers. Her dream is being realized now with the Visible Human Project.

A Plea to Artists

I have dreamed too that artists would use the new technology ordinarily. This has barely begun. My plea to artists is this: Don't leave this new materia prima—called computation—to us technologists. It is a new creative medium of great profundity, realized on a machine that is the most malleable tool ever invented by human beings. It is necessary that you, the explorers at the edges of our culture—the artists—embrace this stuff and tell the rest of us what it really is. If you don't like what you see now—and there is much to dislike here in the early days—please remember: You can define its development. That's what malleable means. Make it in your image.

Some deep contributions of computation to the arts I've identified: Control of complexity, spatial editing, separation of creative space from display space, modulated evolution, immersion, and convergence of media types. Did you expect this list?

Not the To Do List

Audiences often ask me, "What are the great challenges remaining in computer graphics?" Before I mention them three of them that might surprise you, let me list what are *not* the challenges.

There are certain improvements that we need, but I am sure they will happen. Raw machine speed (à la the Law) will make 3D capability commonplace. I shudder every time I look at my kids' video games. They are thrilled but all I see are the aliasing problems we've spent years eliminating. Makers of the most powerful 3D graphics machines crow about their polygons per second numbers, but I still see crude realtime graphics by my measure. Recall that *Toy Story* frames took an average of seven hours each to compute! On state-of-the-art workstations. And it's a cartoon. But time and the Law will cure these ills.

Today's internet bandwidth is atrociously inadequate. The kind of graphics being shipped around the web makes even video games look good. This too will pass. My favorite candidate for a solution is the new digital TV system. Think of it as a bit pipe—not necessarily a string of video frames. Each channel will spew 20 megabits per second into your home! Compare that to 28.8 or 36.6 kilobits per second. Of course, DTV is only one-way, but it is the right way. Key clicks going the other direction don't bog down phone lines.

Do we need new algorithms? It would be foolish of me to say no, but I do think we have reached a time of diminishing returns here (I'll probably have to eat these words). I am currently fascinated by the resurgence of sampling theory as the active research domain in the very hot "image rendering" field, but to some extent this is what happens when the Law delivers us gravy in the form of cheap cycles. Volume visualization is another cycle eater. So what hard problems (implies no known solution) remain?

The To Do List

First, we have to figure out how to deliver 3D to real people. I use my parents, kids, wife, aunts and uncles when I want to imagine one of them. Frankly, 3D is too hard, even for us so-called experts. Any dummy can drive a 2D interface known as a car, but most of us hire a specialist, called a pilot, to fly a plane in 3-space. I think that instead of delivering a technology ("Here's a complete 3D program") we have to parse it into usable solutions. I don't think real people are ever going to learn perspective and six degrees of freedom etc. The 3D will have to be packaged so they are essentially unaware of it or are only ever required to "drive" it in highly constrained circumstances. One of my favorite examples is Pixar's old Typestry program—3D graphics and RenderMan rendering constrained to the simple task of 3D type. To illuminate a 3D word, the user simply turns on or off preestablished lights in a given array. They don't aim, color, shape, or move them. It isn't a 3D graphics program but a 3D typesetting program, an example of extracting a simple useful application from the general technology.

Second, we've got to create an impedance match between artists and complex models. What I have in mind here are Woody and Buzz of *Toy Story*, with about 700 controls each. But they are cartoons. Human-like characters will probably require thousands of controls. This is simply too much, even for those astonishing artists called animators who have, it seems, built-in physics simulators. A lot of "intelligence" needs to be supplied between them and the characters. I use a filmmaking metaphor here for what is, I believe, a general problem confronting 3D graphics as we further exploit the control-of-complexity dimension of the new medium. This is probably a good place to mention that I do *not* predict actors will be replaced with virtual ones. I don't even know what that means. Something I will predict, however: Actors shall "drive" virtual characters

of satisfying visual complexity. In fact, if you think of animators as actors (that's what Pixar thinks) then this is already happening to some extent.

Third, it's time for what I call The Single Creative App. This is what I want as a creator, so I'm inducting that it's what you want too. Imagine (what appears to the user as) a single application that integrates sound and pictures, 2D and 3D, geometry and imaging, animation and interactivity, and handles all logistics, including intellectual property, asset and project management. Not 15 apps and the conversion filters between them all. This is doable.

Too Soon for History

With two to three orders of magnitude yet to run (and who knows what after that), it's really premature to write the history of computer graphics. I am instructed again by digital TV. A well-known national reporter recently recorded the history of DTV, but his book closes with the proposed technology still using old-fashioned interlaced video and us computer guys just entering the fray. The scene has completely changed since he published his history just a year ago. It was too soon.

We have some very hard problems to solve. I've mentioned a few. Let's get to them. And, artists, once again, please join us in the discovery of this medium. It will take both cultures to understand this one.