

From the Look of Things

Alain Fournier

Department of Computer Science
University of British Columbia
Vancouver, BC, V6T 1Z4, Canada
fournier@cs.ubc.ca

Abstract

We can't help wonder occasionally about what we do. The following is the result of such wondering, using a unique opportunity to get a paper in without much scrutiny.

Résumé

On ne peut pas s'empêcher de se demander de temps en temps ce qu'on fait. Le présent document est le résultat de telles questions, profitant de l'occasion unique de pouvoir publier un article sans trop d'examen.

Keywords: what, why.

1. Introduction

1.1. So You Are in Computer Graphics, Eh?

I guess the easy answer to that kind of question is yes. The trouble is that it is followed, explicitly or implicitly by a more difficult one: what are you doing? The standard answer is that I am making pictures with computers. This is only a paraphrase of the expression "computer graphics", but logicians tell us that true statements are all tautologies. A related question, more important for acceptance of papers and career advancement is whether what you do is computer science, or electrical engineering, or whatever is the name of the department which hired you. That is where my definition of computer graphics as *computer science you can see* becomes especially handy. All these attempts at definitions have a serious side, because they help you choose your research directions, the venues you submit papers to, as well as research and thesis topics. A unique feature of computer graphics within computer science is that we synthesize (I carefully avoid that word in verbal communication), and the product of our synthesis is visible (sometimes, rarely, palpable). It would put us definitively into engineering, except that we do not create real objects (if we design graphics hardware, of course then

we are engineers).

1.2. Basic Principles

There are a few principles that every computer graphics person¹ should adhere to.

- A picture **is not** worth a thousand words.

This ought to be obvious enough to prevent the positive version of this aphorism from being printed ever again. There are words, such as "beauty", "love", or "standing committee on curriculum re-assessment" which cannot be expressed by a thousand pictures (calling a bitmap of the letters making the words a picture does not count; you should be ashamed for suggesting that). A few examples of the puzzling results when one insists on replacing words with pictures can be seen in Figure 1 [1]. There are of course pictures which a thousand words cannot begin to describe. Figure 2 is one. We have both words and pictures, and we should use both.

- A model **is** worth a thousand pictures.

I heard that first said by Frank Crow. I do not know if he originated it, but he deserves the credit. We can show a lower bound on that number. Assume we have a 3D model of some object. From it we can compute samples of all the possible views from all possible directions. If we consider only orthographic projections, and using polar angles for the directions, everybody will agree that 5 degree increments is the largest we can get away with, and therefore we need 37 samples in θ and 73 samples in ϕ , for a total of 2522 samples (I grant you only one ϕ value is needed for $\theta = 0$ and $\theta = \pi$). Therefore this model is worth *at least* 2522 pictures. Of course I assumed that the model is good enough so that somebody would want to see 2522 pictures rendered from it,

¹ Some time ago Jules Bloomenthal asked the community for suggestions to replace this awkward expression by something pithier; I am not aware of any results, so I will use *cgp*.





Figure 1. Some pictures not worth a thousand words (guess what words they are meant to replace).

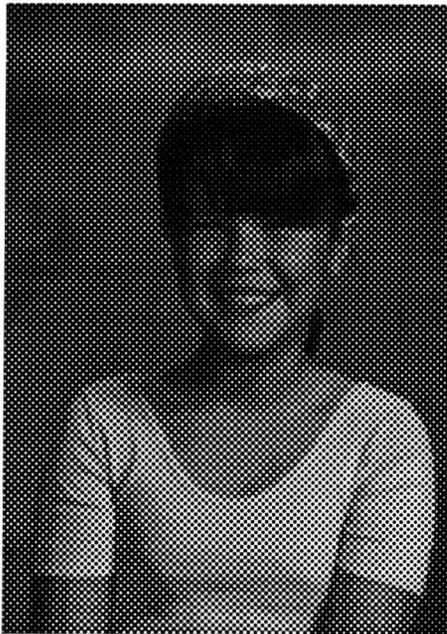


Figure 2. Wordless.

and that all these pictures are different.

- There are things that exist and cannot be seen.

I will not even begin to discuss whether mathematical objects exist. I will just mention in passing that Quine

(maybe Eugene Fiume's favourite philosopher) [10] is less ready to admit existence of ideal objects of physics (such as mass points and frictionless surfaces) than of geometric objects. Would that make "physical modelling" more suspect than geometry and kinematic? You be the judge.

Some things are very hard to model and/or render. A black hole is hard to see (some ray-tracers I know can do a reasonable job, though, they send a ray and you never hear from them again). A single photon is hard to render (though doing them in bulk is our livelihood). Most fractal objects are easy to render if we assume their surface is fully realized. Let us take as a simple example a fractal surface as a sample of two-variable fractional Brownian motion $fBm(x, y) = z$. For any point P on the surface to be visible from the eye (or from a point on a light source), a line segment from that point to the eye cannot intersect the surface. (see Figure 3).

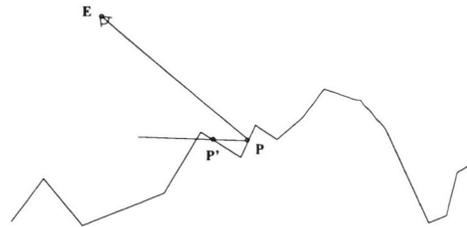


Figure 3. A fractal surface cannot be seen.

It means that for any point P' on the surface the segment PP' has a slope less than the slope of PE . As P' gets closer to P , however, that means that there is a limit to the slope of PP' , which contradicts the property of fBm that the limit slope goes to ∞ with probability 1 [7]. Therefore no point on the surface can be seen from the eye or from the lights (exercise left to the reader, what if the line PE is parallel to the Z axis?). This frees us from the worry of designing an illumination model for such surfaces². Also note that it is not in contradiction with the fact that one standing on the surface would see a finite horizon, as long as the observer is a finite height above the surface.

2. Should you think it does not apply to the real world, recall that one can make approximations of a blackbody by stacking tightly together several old-fashioned double-edged razor blades.



- There are things that can be seen and do not exist.

In fact everything we display in computer graphics not from real images (digitized video or film) falls under that category. Our models are not real, even if the objects they are "inspired" from are real. It is interesting that some programming languages insist on calling "real" some class of numbers they manipulate. We should not make the same mistake. The initial charm of computer graphics is that one can display lines, cubes, spheres, etc., and make them obey our bidding. We should not forget, however, that not only what we see, but even what we represent internally, are only approximations of some ideal. This was a matter of contention when using fractals as models, because our algorithm produced only "approximations" to fBm. But of course any algorithm implemented on a computer can only produce approximations to a sine curve (because of the need for transcendental functions) or even to a straight line (because of possible truncation of slope, among other reasons). The real question is how close we are *vs* how close we want or need to be.

- We can display anything anybody wants us to display.

True³, as long as they do not want the exact shape and the real colours, and we will have to insist that what they see is what they wanted.

1.3. Seeing is Believing

How well do we see

Pretty well, thank you. No matter how many pixels we throw at the viewer, how many bits per pixel we use, how often we update the image, there will be cases where the user can see through our fakery. We commonly think 24 bits/pixel is enough, but it is enough only because we hide behind the pitiful dynamic range of the monitors. Even at the common range of about 25, 9 bits of luminance is needed if we assume we can notice differences of 0.5% across the range. If the dynamic range approached the 13 decimal magnitudes of normal human vision then we would need at least 11 bits/pixel for luminance (assuming equal relative steps of 2%), which means 33 bits/pixel colour if an *RGB* like colour space is used. Even with higher estimates this is not too hard to provide at today's memory cost. Spatial resolution, you say. Just a few interesting data points. Your basic photoreceptors are spaced by some 30 seconds of arc (these are actually the best, not the basic

ones, but we are talking rough numbers here). To achieve the same spacing when viewing pixels on a standard monitor, we should stand about 4 meters away. Another way to look at this is that for the normal viewing distance we should increase the number of pixels by a linear factor of 7 or 8. But, you say, we do not really have to hit each last cone. There are only about a million optical nerve fibres, and that is the number of pixels we have on ordinary displays now. Besides these spacings are for the fovea, the rest is easy. Well, think again. In some tasks, such as using binocular parallax to judge the depth of a line, we can detect a difference of about 5 seconds of arc. During a saccade the eye can travel at speed of 500 degrees/s. That means that the point whose image is on the fovea can move by about 200 pixels in a 1/60th of a second. Not all of this is gloomy, because if we can make "sub-receptor" discriminations, we can also produce sub-pixel changes, thanks to the miracle of filtering.

We really see well when there is a point. As a young child (is there any other kind) I accompanied men who hunted "bizets" (rock doves) and "palombes" (wood pigeons), two birds of the *Colombidae* family that I could not tell from city pigeons even up close. They could tell them apart in flight when to me they were just "v"s in the sky. Some of us can also, at a glance, tell silk from rayon, rag paper from pulp paper and fresh lettuce from about to be wilted one.

How badly do we see

Our vision is notoriously poor at absolute judgements (and of course there are good reasons for this). That is why we have lived so long with the very limited dynamic range of displays. That is also why so many rendering systems limit users to the 0-1 range for light "intensity", and (almost) get away with it, and why many cgps cannot tell you what the luminance of the light from a 100 watt bulb at 2 meters is.

We are limited in bandwidth within the electromagnetic spectrum (one octave *vs* more than 8 octaves in sound; again there are good physiological reasons for this), and even more limited in sensing spectra, since we project an infinite dimensional space to three dimensions. Finally we time sample rather niggardly. All of that means that we have boxes that show images a few hundred pixels across, getting the job halfway done in one sixtieth of a second and where all the colours are generated by three phosphors incapable of getting even close to the rainbow. It would be amusing if it were not the kind of display we use as a standard of realism in computer graphics.

3. Around Imager this is known as the Kelly Booth principle.



We really cannot see very well when we do not want to see. I am not referring here to the standard battery of optical illusions, but to the cases when we do not pay attention, or when we fall for the ploy so well summarized by Chico Marx: "Who you gonna believe, me or your own eyes?". Ask yourselves those questions. Does the full moon look like a diffuse sphere (Bob Woodham asked me that one)? What is the shading model for images produced by scanning electron microscope? Is this paper really white (or how would I know)? And what is the colour of my eyes, anyway?

For many years ray-tracing has been producing what we call "realistic" images, when most of them suffered from the slight problem that the darkest part of the image should have been the brightest. Even pictures produced by cone-tracing John Amanatides, which at the time represented a definite advance for the "ray-tracing" paradigm still suffers from that problem.

What do we see

Do we see pixels when awake, are there jaggies in our dreams? A quote stolen from Eugene Fiume's thesis [4] shows that philosophers are not always wrong:

Let us imagine a white surface with irregular black spots on it. We then say that whatever kind of picture these make, I can always approximate as closely as I wish to the description of it by covering the surface with a sufficiently fine square mesh, and then saying of every square whether it is black or white. In this way I shall have imposed a unified form on the description of the surface. The form is optional, since I could have achieved the same result using a net with a triangular or hexagonal mesh.

Lutwig Wittgenstein (1889-1951)

This is implicitly the Weierstrass theorem of images, and whatever we see it means that we can use pixels, which neatly bypasses much cognitive psychology. How many pixels and what kind is still a big question (see above and below). The shape of pixels is occasionally discussed, more from the point of view of what it is rather than what it should be. I am partial myself to hexagonal tessellations, maybe with jittering, as a practical alternative for displays, but then again I also believe patents should be abolished.

Some fundamental facts about the pixels we do use are often neglected. It is taken for granted, for instance, that there is a trade-off between size and intensity, and in fact newspaper empires are based on this. We have been blithely computing "anti-aliased" lines with intensities of partially covered pixels determined to give the illusion of sub-pixel motion, and similarly with sub-

pixel positioning of characters. In what is now known as the A-buffer (we should have called it the α -buffer, or better still the \bullet -buffer) [3] we were blithe indeed. There were however many simple questions with no answer: how small have the pixels to be for the trade-off to be real, does half the intensity put the boundary halfway between pixels, is the effect the same for all observers, is it linear, how many bits of position can I convert to intensity? It is only with the recent work of Avi Naiman and Walter Makous [8] that these questions have answers based on experiments.

I often thought that the famous WYSIWYG principle is more a threat than a promise. I certainly do not want to see on paper, or anywhere else, the unkered, unfiltered and uncouth characters I see on my screens. I had rather entertain a little longer the illusion that what I write will come out as if set by Bodoni or Fournier (Le Jeune). On the other hand the concept has merit where it has not been used much, in rendering. As surface definitions and illumination models become more sophisticated, the user is less and less able to predict what the result will actually look like. It makes sense, therefore to let the user specify the end instead of the mean. That is the prime motivation of the PhD work of Pierre Poulin [9].

1.4. Beware of the Meta-view

The *meta-view* is the concept that what we should show is what the eye should see. First the eye does not see. The viewer sees⁴. What we should show is such that the viewer will see the same thing that if she/he looks at the modelled object imaged through the same medium. That is assuming that the modelled object exists. If it does not we would like to "evoke" existing visible objects (see below about realism).

The meta-view fallacy is quite widespread and persistent. It appears in almost all popular science descriptions of the eye, where it is said that the image on the retina is upside-down, and the brain (among its many other jobs) "puts it right side up". This is nothing but the survivance of "explanation" of vision by an *homunculus* posted behind our eyes looking at the retina. Some art historians tried to explain El Greco's elongated characters by his astigmatism (of course he would have seen his painted characters even more elongated, astigmatism is not idempotent), or J. W. Turner's sea of red by some degeneracy of the cornea (this one more believable if still suspicious).

4. That is why *scientific visualization* should be called *scientific display*, or more simply *data display*. But then it would be clear that it is what we have been doing all along.



Closer to us, it has been argued [6] that

"The human eye senses intensity; it perceives projected areas, and receives energy within a solid angle $d\omega$ defined by pupil size. Intensity is thus an appropriate quantity for use in the construction of computer generated images."

(intensity here means *radiance*). The conclusion is right (radiance should be used), but the reasons are wrong. Whatever the human eye "senses" is not what we want to use. If we want to create a video or film which looks like real objects were imaged, then what we want to compute is whatever affects the medium, not the eye or the visual system which will later see it. Of course knowing the characteristics of the eye might help us cut corners, since simulating the medium accurately is a sufficient condition, but not always a necessary one.

Another way the meta-view caused us some harm is in the use of *RGB* colour spaces in illumination models. Even though demonstrably wrong, it is used because the human colour space is three-dimensional, but of course reflecting materials do not know that and do not care. It is intriguing that there is a possibility reflection can be expressed in a low dimensional space, but it takes a heavy dose of the *anthropic* principle to make something out of it. A similar effect is at work when we try to justify the use of an opponent colour space, because signals to the brain may be encoded that way. Our frame buffers and screens are not plugged (yet) directly to the optic nerve.

1.5. Realism

I admit it: I and some of my acquaintances have used the word "realism", mostly to claim our pictures are more realistic, without defining it. Quoting Webster (easy because available on-line):

re-al-ism 're[^]-e-,liz-em, 'ri-e-n
(1817)

1: concern for fact or reality and rejection of the impractical and visionary

2a: a doctrine that universals exist outside the mind; specifi: the conception that an abstract term names an independent and unitary reality

2b: the conception that objects of sense perception or cognition exist independently of the mind
-- compare *NOMINALISM*

3: fidelity in art and literature to nature or to real life and to accurate representation without idealization

Definition #1 is for engineers. Definition #2a is for Platonicians. A counter-proposition is found in *The search for intelligent life in the Universe*, by Jane Wagner and Lily Tomlin: "What is reality, after all, but a collective hunch". Definition #2b is right for everybody but me. Definition #3 is what we more or less mean in the context of computer graphics. Now if only we could define "accurate"...

One often hears criticism of the "realistic" school of modellers and renderers. The argument is usually "why imitate reality slavishly when we can communicate more effectively with symbolic or more abstracted representations". First achieving realism is a challenge in itself, a proof that we master the medium and that we model accurately under very demanding conditions. It can be an end in itself. It is a little bit (analogy alert) like telling a runner that she would get there faster with a cab. Second in many applications, such as *Computer Augmented Reality (CAR)* where the goal is to merge seamlessly real and computer generated images, there is no alternative. If the goal is indeed to communicate, then of course we should select the most effective and efficient method, and that might very well not be "realism". One more note: looking real and eliciting real responses is quite different. One recent experiment in Virtual Reality consisted in simulating standing in high places with people suffering from acrophobia. They experienced real fear, some to the point of sickness and panic (it was not done to scare them, but to help treat them). Does that mean that the patients thought they were "really" there. Of course not, just that the response was real. One always hears mentioned that when the Lumière brothers showed their first films (1895), in particular of a train entering the station, and coming right at the camera, spectators ducked for cover. Did they temporarily forget that reality is not in black-and white, and does not flicker badly? No, but the experience was new and powerful enough to make them react, not necessarily to make them believe there was a real train coming at them.

There is not just one axis of "realism". First there is the abstract-representational axis (to borrow the terms used in visual arts). Do I want to show directed acyclic graphs or weeping willow trees? There is the subject matter axis. That is the most important in painting, where schools of realism have been identified by this axis (the best known associated with Gustave Courbet in the 19th century). Do I paint Greek gods floating in the clouds or peasants cutting their toenails? In computer graphics it is less important for now, but as we get better in modelling and rendering this will become more of an issue. A landmark work in this respect is the *pen & ink*



illustration system out of the University of Washington [11]. Then there is the visual complexity axis. Do I want flat shaded spheres (*aka* discs) or chrome balls? I can certainly render abstract models in a highly realistic manner. Everybody knows now that carbon atoms are black, oxygen atoms red, hydrogen atoms white.

This is a central issue in volume rendering, as the structures that have to be shown, although usually "real" enough, are not normally seen in this form. The reason to render them realistically is that we want to put to good use all the training our visual system has undergone these past two million years. At the same time we don't want to "create" objects that are not there, and we want to respect the mind-set and training of the skilled people who use the display for research and/or diagnosis. A flexible yet powerful approach is offered by treating volume rendering as an exercise in filtering. That prevents having any specific visual model built-in into the rendering (not all models fits within the paradigm, actually). This is the basis for John Buchanan's PhD work [2]. His goal was to explore the representations, algorithms and data structures associated with the filtering approach, not the ways people would actually use it. The latter is also important and has yet to be done.

1.6. The Shape of the World

I still don't know what "shape" is. In dictionaries "shape" is always defined by words synonymous with shape, or by totally opaque words. In computer graphics, however, shape is modelled through a variety of primitives, ranging from the point to the forest. This is an age-old activity. Cultures see primitive shapes in nature the way kids see bunnies in clouds. To wit:

Everything in Nature is modeled after the sphere, the cone and the cylinder. One must learn to paint from these simple figures.

Paul Cézanne (1839-1906)

We take that seriously, and in fact Vishwa Ranjan is busy using spheres to represent shapes (and we mean shape) in the context of volumetric data. You would think from that quote that Cézanne believes firmly in the power of geometry. On the other hand:

I seek to render perspective only through colour
The same Cézanne

Anyway one could go on with such examples. A more recent version is:

Clouds are not spheres, mountains are not cones, coastlines are not circles, and bark is not smooth, nor does lightning travel in a straight line.

Benoit Mandelbrot

Notwithstanding that nobody seriously claimed that clouds are spheres or that lightning travels in a straight line, Mandelbrot's eventual point is extremely important, and he gave us a brand new collection of primitives to play with.

To quickly summarize where we cgps are in terms of shape modelling:

- rigid objects are a snap;
- complex objects are coming along;
- articulated objects are fine as long as they have no skin;
- moving rigid objects is down to a fine art;
- moving complex objects is dragging behind;
- fluids are unruly and turbulent.

Examples to justify these statements and references will be provided on request. Let me single out the work of David Forsey for free-form design of complex shapes, and Przemyslaw Prusinkiewicz for stunning objects from plants to shells (and recently even plants shaped as shells). The main weakness we still have is lack of integration. Motion has to be built-in, looks have to be built-in.

1.7. The Look of the World

Is it possible to even separate "look" from "shape"? This is one case, I think, where computer graphics actually helps clarify our concepts. Texture mapping, local and global illumination models all are concerned with the look. The shape has been determined by other means (even in the case of displacement maps). Indeed we use these to infer shape in our visual system, but that's the meta-view reappearing.

While the aforementioned techniques are at the forefront of research in computer graphics, there are however still aspects of looks that have been somewhat neglected but are powerful effects in our visual universe. The lustre of materials (especially linked to binocular vision), the change of colour of materials (not necessarily from interference, but from shadowing and blocking, as with denim), the influence of the microstructure of material (the stochastic approach to these has been successful, but many microstructures are quite regular, as in most woven materials, woods, skin, furs).

The story so far:

- texture mapping is ubiquitous and powerful;
- local illumination is in rebirth;
- global illumination is still almost exclusively with radiosity;
- light models are still quite primitive;
- and beware the participating medium (as they say in



seances).

I and associates are currently working on most aspects mentioned above. Results to follow. Again the main problem here is integration. Just one example: if we want to model and render a tiger moving, we need a good articulated, moveable tiger shape, but we need also to be able to define the tiger stripes, to map it smoothly without deformation on the shape (or to grow it "in place"), to render the fur at various scales, and to make the whole thing move convincingly and cooperatively (the skin should move slightly with respect to the body, but not the texture with respect to the skin, except inasmuch as it is compatible with fact that the pattern is actually carried by the fur, whose elements can move slightly with respect to the skin). One step in that direction has been the use of reaction diffusion processes in 1D to model patterns on shells, where the texture and the shape were beautifully integrated [5].

2. Further Work and Conclusion

There is no definite conclusion because there is no stopping point. If one gets the feeling that "it all comes together now" it is probably because the universe is collapsing. There are interesting folds in the fabric of the world, though. Recently I found myself musing about a technique to generate stochastic models that involves fractional Brownian motion, wavelets, recursive modelling, filtering of normal distributions and *NIL* maps. Those are topics I have dealt with in works that span my whole active research life. It is rather disquieting to have to consider that maybe there was a plan.

Computer graphics is not as different from other analytic sciences in the way it builds models: it has to make them as simple and economical as possible, while trying to make them useful. In our case, usefulness is verified in pictures. We should not be ashamed of the simplicity of our models, neither should we be too proud of them. Speaking of something else (which is what poets do; in this case it was about human love) Louis Aragon (in *Le roman inachevé*) said:

*Mais tout ceci n'est qu'un côté de cette histoire
La mécanique la plus simple et qui se voit
Une musique réduite au chant d'une voix*

Indeed we want the simplest mechanism that can be seen, but we should not think it is the whole story. By summing up computer graphics in a few idiosyncratic pages I have reduced a beautiful piece of music to the sound of a single voice. Thank you for listening to it anyway.

Acknowledgements

This ought to be the longest section, as I have more debts than thoughts.

It all started in Dallas, when I met Henry Fuchs and his Genisco frame buffer (let's not mention the Tektronix storage tube; I still hate green screens). Sam Uselton was first introduced to me with "he knows everything about computer graphics". That was the first time I considered cg as a possibility. Then influenced by Henry, Don Fussell, Bruce Naylor and I soon followed Sam's footsteps and started making pictures with computers. Zvi Kedem, my real supervisor, managed to keep me honest and graduated in the process.

I did go much farther, as Newton would say, because I was standing on the shoulders of taxpayers. Through NSERC, ITRC, ASI and the funding of universities, they provided continuous support for these endeavours. Private funding has been important too, especially here at UBC from IBM Canada, who helped create and support GraFiC. At the University of Toronto Xerox and Apple played an important role.

I have found the universities themselves always supportive, especially here at UBC, where under the leadership of Maria Klawe in the Department of Computer Science I (and we collectively) have been given even more than promised. The various levels of administration (for Dean of Science up to President) have been remarkable in their commitment, even under difficult financial conditions. It is still amazing to me that I could trust, rely on and like all the department heads I knew, Alan Borodin, Derek Corneil and Maria Klawe.

I have been especially lucky to be in two exceptional departments. Exceptional not only for the quality of people, but for an atmosphere quite free of petty power struggles and maneuvering. At Toronto I must single out the DGP crowd, Ron Baecker (who created DGP with Les Mezei and trusted me to join), Bill Buxton (too unique for words), Marilyn Mantei (who convinced me CHI work can be done honestly) and Eugene Fiume. Eugene of course is special there, not only a former student and a friend, but an inspiration, and the one who gave me a "second wind" in research when I badly needed it.

When it comes down to it (and it does soon enough), the students make all the difference. In my case the difference was positive, very positive. I have to single out everybody, because every one gave me something unique. The PhD students who have finished, in chronological order, are Delfin Montuno, Eugene Fiume, John Amanatides, Avi Naiman, Pierre Poulin and John Buchanan. They are all unique and special to me, and I



am happy to have all of them as friends. It is worth noting here that they collectively speak Chinese, French, Greek, Hebrew, Italian, Japanese, Spanish and Tagalog, and for most of them English is not their native language. I think it means something good for Canada. Pierre Poulin and John Buchanan deserve special mention because they took a great risk by following me when I moved from Toronto to Vancouver. It paid off for me, I hope it did for them. I owe them a lot. The PhD students who did not finish do not look as good on your *cv*, but life does not always follow plans (in fact does not get even a copy of the plans). Elsa Campuzano, Atjeng Gunawan and Lili Liu contributed a lot, and accomplished a lot. Current PhD students are the usual exciting, confusing, frustrating bunch. Vishwa Ranjan, Bob Lewis, Marcelo Walter, Paul Lalonde and Bill Gates (the last two shared with Dave Forsey) are all working on aspects of computer graphics discussed here. They will soon turn into polished, omniscient and graduated PhDs. I always claimed it is easy to formulate a good MSc topic in computer graphics, much harder for a good PhD topic. That's one reason MScs in computer graphics are exciting and valuable. The students certainly are. Again in chronological order, they are John Amanatides, Tim Piper, Dave Grindal, Avi Naiman, Colin Hui, Galia Diker-Pildush, Tom Nadas, Pierre Poulin, Andrew Woo, Maria Raso, Michael Penn, David van Blankenstein, Gang Huo, Russ Krywolt and Bill Gates. Tim Piper died tragically in Vancouver a few years after graduation. Current students are Raza Khan and Lili Liu, about to graduate, Chris Romanzin and Paul Fearing. Here again not everything is wrapped up. Peter Schoeler, Tom Milligan and Chuan Chee all have published results, but no thesis (yet). Not students as such, but wonderful visitors and collaborators, were Mikio Shinya, Marie-Claire Forgue and Frederic Taillefer.

The current Imager and GraFiC labs, in addition to the students listed above have of course David Forsey and Kelly Booth. That was quite unexpected to have Kelly arrive one year after I did, but I could not have been luckier. In addition to all I knew about him from a distance, I discovered many other sides of him, in particular his incredible generosity of spirit and deeds. Dave is holding up the "junior" side of the lab, meaning he has to do real research and get students excited (Kelly and I tell them about the time we had a PDP 11/45 driving a 256x256 frame buffer). Many more notable are in the lab. Two special individuals are Peter Cahoon, who always knows the right tool, whether mathematical or commercial, and has done everything at least once (truly), and Ron Lane-Smith, with the Taj Mahal as his beautiful obsession.

Again all of these are friends. In the middle of the night, however, when all seems futile and hopeless, or in the middle of the day, when I think I have found something clever, I first turn to Adrienne. And I last turn to Adrienne. Adrienne and Ariel love me. What else is there to say.

References

1. *Répertoire des Pictogrammes*, Les Publications du Québec.
2. J. Buchanan, "Filtering Volumetric Data," *Ph.D. Thesis, Department of Computer Science, University of British Columbia*, 1993.
3. E. Fiume, A. Fournier, and L. Rudolph, "A Parallel Scan Conversion Algorithm with Anti-Aliasing for a General-Purpose Ultracomputer: Preliminary Report," *Proceedings of Graphics Interface '83*, pp. 11-22, 1983.
4. E.L. Fiume, "A Mathematical Semantics And Theory Of Raster Graphics," *Ph.D. Thesis, Department of Computer Science, University of Toronto*, 1986.
5. D. R. Fowler, H. Meinhardt, and P. Prusinkiewicz, "Modeling seashells," *Computer Graphics (SIGGRAPH '92 Proceedings)*, vol. 26, no. 2, pp. 379-387, 1992.
6. C.M. Goral, K.E. Torrance, D.P. Greenberg, and B. Battaile, "Modeling the Interaction of Light Between Diffuse Surfaces," *Computer Graphics*, vol. 18, no. 3, pp. 213-222, July 1984.
7. B. B. Mandelbrot and J. W. Van Ness, "Fractional Brownian Motion, Fractional Noises and Applications," *SIAM Review*, vol. 10, no. 4, pp. 422-437, October 1968.
8. A. C. Naiman, "The Use of Grayscale for Improved Character Presentation," *Ph.D. Thesis, Department of Computer Science, University of Toronto*, 1991.
9. P. Poulin, "Shading and Inverse Shading from Direct Illumination," *Ph.D. Thesis, Department of Computer Science, University of British Columbia*, 1993.
10. Willard Quine, *Word & Object*, The MIT Press, 1960.
11. G. Winkenbach and D. H. Salesin, "Computer-Generated Pen and Ink Illustration," *Computer Graphics (SIGGRAPH '94 Proceedings)*, vol. 28, to appear.

