SUPERCOMPUTERS

HEARINGS

BEFORE THE

COMMITTEE ON SCIENCE AND TECHNOLOGY U.S. HOUSE OF REPRESENTATIVES

NINETY-EIGHTH CONGRESS

FIRST SESSION

NOVEMBER 15, 16, 1983

[No. 47]

Printed for the use of the
Committee on Science and Technology

U.S. GOVERNMENT PRINTING OFFICE

28-293 O WASHINGTON : 1984

COMMITTEE ON SCIENCE AND TECHNOLOGY

DON FUQUA, Florida, *Chairman*

ROBERT A. ROE, New Jersey
GEORGE E. BROWN, Jr., California
JAMES H. SCHEUER, New York
RICHARD L. OTTINGER, New York
TOM HARKIN, Iowa
MARILYN LLOYD, Tennessee
DOUG WALGREN, Pennsylvania
DAN GLICKMAN, Kansas
ALBERT GORE, Jr., Tennessee
ROBERT A. YOUNG, Missouri
HAROLD L. VOLKMER, Missouri
BILL NELSON, Florida
STAN LUNDINE, New York
RALPH M. HALL, Texas
DAVE McCURDY, Oklahoma
MERVYN M. DYMALLY, California
PAUL SIMON, Illinois
NORMAN Y. MINETA, California
RICHARD J. DURBIN, Illinois
MICHAEL A. ANDREWS, Texas
BUDDY MacKAY, Florida
TIM VALENTINE, North Carolina
HARRY M. REID, Nevada
ROBERT G. TORRICELLI, New Jersey
FREDERICK C. BOUCHER, Virginia

LARRY WINN, Jr., Kansas
MANUEL LUJAN, Jr., New Mexico
ROBERT S. WALKER, Pennsylvania
WILLIAM CARNEY, New York
F. JAMES SENSENBRENNER, Jr., Wisconsin
JUDD GREGG, New Hampshire
RAYMOND J. McGRATH, New York
JOE SKEEN, New Mexico
CLAUDINE SCHNEIDER, Rhode Island
BILL LOWERY, California
ROD CHANDLER, Washington
HERBERT H. BATEMAN, Virginia
SHERWOOD L. BOEHLERT, New York
ALFRED A. McCANDLESS, California
TOM LEWIS, Florida

J. H. Poore, *Executive Director*
Robert C. Ketcham, *General Counsel*
Regina A. Davis, *Administrator*
David S. Jeffery, *Minority Staff Director*

CONTENTS

WITNESSES

APPENDIX

Additional Testimony Submitted for the Record

Additional Material Submitted for the Record

SHORT BIOGRAPHY

Alvy Ray Smith
Computer Graphics Project Leader
Lucasfilm Ltd.

Alvy was introduced to computer graphics at Xerox Palo Alto Research Center by Dick Shoup, who designed the first sophisticated computer-assisted "paint" machine (1972-1973). Here he united his artistic interests and computer science education (PhD, Stanford, Cellular Automata Theory) to create the videotape *Vidbits* which served as a passport into the newly formed Computer Graphics Laboratory at the New York Institute of Technology in 1975, where for five years he assisted in the development of this facility into one of the premier computer graphics installations in the world. While there he generated stills, videotapes, and films which have appeared in a variety of magazines, books, newspapers, museums, and televsion shows. In particular, just before leaving NYIT, he collaborated with Ed Emshwiller on the videotape *Sunstone*. He joined Ed Catmull in the Computer Division of Lucasfilm Ltd. in early 1980. This division is involved in a major effort to digitize the filmmaking process with special emphasis, of course, on the use of computer graphics in theatrical release motion pictures. Alvy is Project Leader of the division's Computer Graphics Project, consisting of thirteen people engaged in both software and hardware development. The Graphics Project generated a state-of-the-art production, under Alvy's direction, of a sequence (the "Genesis Demo") for the 1982 Paramount movie *Star Trek II - The Wrath of Khan*. Most recently, the Project contributed a sequence to Lucasfilm's *Return of the Jedi*.

COMPUTER POWER FOR FILM AND FLIGHT †
Technical Memo No. 95

Alvy Ray Smith
Computer Graphics Project Leader

Computer Division
Lucasfilm Ltd.
October 21, 1983

INTRODUCTION

Recent incorporation of computer graphics into feature-length films - *Return of the Jedi, Star Trek II: The Wrath of Khan,* and *Tron,* to name a few - has inspired a belief that entire films might soon be generated by computer. It is shown below that even the so-called "supercomputers" of today fall quite short of the power required for this goal. True supercomputers are needed with capabilities just now being conceived and with cost commensurate with typical filmmaking practice.

The discussion here will focus on the computation of images, but supercomputers are general-purpose computers, which means they can be used for any kind of computation. In particular, supercomputers meeting the demands of movie-making computation would further image-making computation in general. Flight simulators are a case. Cost-effective supercomputers of sufficient power could replace expensive special-purpose flight simulators.

HOW TO COUNT PIXELS

A *pixel* is the smallest part of a picture computed by a computer (it is a corruption of "picture element"). A single frame of film is assumed to be divided into a great many scanlines - just as a television screen is divided into scanlines - and each of these is further divided into many tiny spots, the pixels. So a frame is divided into, say, 1000 scanlines each consisting of, say, 1000 pixels.

The *resolution* in this case would be said to be 1000×1000 ("×" is read "by"), or 1K×1K, where K is shorthand for a thousand of anything ‡. The resolution required by film is not yet known with great certainty. Since 500×500 is about the resolution of American television, this is undoubtedly too low for film. 8K×8K is surely too high. We shall assume it is 3K×1.5K pixels. The Japanese have experimented with resolution in a series of tests for establishing their *hi-vi,* or high-definition television system. They found that above 1600 scanlines their subjects did not notice substantial improvements in image quality. For ease of computations and conservatism, we will use 1500 scanlines for our estimates here. Since a common movie film format is Panavision, or Cinemascope, which has a width more than twice its height, we shall assume 3000 pixels along each scanline.

The number of pixels per frame is simply the number of scanlines times the number of pixels on a scanline. So a typical frame is assumed to have 4.5 million pixels; a computer would have to compute 4.5 million spots to make such a frame.

---

† Prepared for testimony before the House of Representatives Committee on Science and Technology, Congressman Don Fuqua Chairman, Washington, D.C., October 26, 1983.

‡ In computer science, K usually abbreviates 1024, the tenth power of 2; we use K = 1000 throughout this article, however.

Another abbreviation we will use is 1M for one million; so 4.5M pixels is 4.5 million pixels. Another common way to say the same thing is "4.5 megapixels". So M and mega- mean one million. There are similar abbreviations for larger and larger numbers. Since we will use very large numbers in this article, it will help to abbreviate them as shown in the following table, where each succeeding entry is three orders of magnitude larger than the one above (a factor of 10 is one order of magnitude):

Large Number Abbreviations			
Abbreviation	Prefix	Meaning	No. of 0's
K	kilo-	thousand	3
M	mega-	million	6
G	giga-	billion	9
T	tera-	trillion	12
P	peta-	quadrillion	15

HOW MANY BYTES IN A MOVIE?

A color film has a red, green, and blue primary component at each pixel to give full color. To "compute a frame" means that a computer must compute one red, one green, and one blue item for every pixel in a frame. The typical smallest unit a computer can generate is called a *byte* (eight bits). So at the very minimum, computing a frame requires the generation of 3 bytes at every pixel, or 13.5M bytes for an entire frame.

It is fairly well known that film is projected at 24 frames per second; so it is not difficult to calculate that there are 24 times 60 times 90, or 129,600, frames in a typical 90-minute movie. Using the abbreviations, there are 129.6K frames in a typical movie. So computing a movie requires computing 13.5M times 129.6K bytes, or 1.7496T bytes. *A typical 90-minute movie requires the computation of about 1.75 terabytes.* This very large number is the amount of memory storage space required to store a movie digitally. Next we estimate the computer power necessary to compute this many bytes of finished film.

HOW LONG WOULD IT TAKE TO COMPUTE A MOVIE?

The calculation above shows that a typical feature-length film at a reasonable resolution has over a trillion bytes. This is the size of the finished film; mistakes, reshoots, and cutting-room floor footage are not counted. Each byte of finished film requires the execution of a series of operations by a computer. The operations of a computer are controlled by a list of instructions, its *program.* Thus the speed of a computer is sometimes measured in *instructions per second.* This works for today's slower general-purpose computers which tend to perform one operation per instruction, but supercomputer speeds are measured in *operations per second* since they tend to to perform many operations for each program instruction.

We estimate that a computer must execute from 1000 to 10,000 operations for each final output byte to film. This is based on an image of a richness we believe filmgoers will find satisfactory, one which approaches reality in complexity (but certainly not restricted to reality in content). We shall use the average of 5000 operations per byte for our calculations. This is believed to be an underestimate since we used about 4000 operations per byte for the Genesis Demo sequence in *Star Trek II: The Wrath of Khan* and are aiming for an even higher complexity in future projects. 5K operations per byte times about 1.75T operations implies 8.75P operations, or 8.75 peta-operations. Thus *a typical movie would require the computation of approximately 8.75 quadrillion operations.*

A common general-purpose computer in 1983 (not a supercomputer) is the VAX 11/780 manufactured by Digital Equipment Corporation. It is rated as a 1Mip machine. This is pronounced "one mip" and means 1M instructions per second. For a VAX, 1 Mip is the same as one million operations per second; so 8.75P operations would take 8.75G seconds to compute. Since

$$1G \text{ second} = 11.574K \text{ days} = 31.797 \text{ years},$$

a VAX would require more than 278 years to compute a movie! And this is being kind. As we will point out several times in this memo, it is inaccurate to use the rated speed of a computer for realistic results. The VAX computes at much less than 1 Mips for actual programs; so 278 years is an unrealizable best case.

Since it is out of the question to compute an entire 3K×1.5K movie at the present, what if video resolution, 500×500, were used instead? A best-case calculation similar to that above yields the result that over 15 years of VAX time would be required; so even a highly degraded image is also out of the question.

It should be emphasized here that we are calculating *raw compute time*. We have assumed that the machine never stops computing once it starts. This is unrealistic because a computer never operates around the clock, day after day, without having a hardware failure or being stopped for some other reason, such as machine maintenance. Furthermore, time must be allotted for system security backups, recomputations for the inevitable mistakes, and program development time.

A typical major motion picture requires three stages called preproduction, production, and postproduction. Preproduction comprises such tasks as design, scripting, storyboarding, and financing; it takes on the order of six months. Postproduction covers editing, audio mixing, film printing, foreign language adaptation, special effects, advertising, and distribution; it requires on the order of a year. Our raw compute time takes neither of these two stages into account; it is to be compared to the actual production time of a film which is about four to six months.

THE EFFECT OF SUPERCOMPUTERS

The numbers derived above are discouraging. Would the use of modern supercomputers change the outlook substantially? The two most well-known supercomputers today are the Cyber 205 by Control Data Corporation and the Cray 1S by Cray Research Inc. The speeds of these machines - and of supercomputers in general - are expressed in "megaflops" instead of, say, "mega-ops" because the operations performed by supercomputers which most interest their users are the so-called "floating-point operations", a particularly difficult set of operations for computers. A supercomputer megaflop is computationally more powerful than a VAX Mip, but the difference is difficult to quantify; so we will take them as equal here. The effect is to underestimate the time needed to compute with a VAX rather than to overestimate the time needed with a supercomputer.

A recent paper [1] argues that in actual practice at Los Alamos National Laboratory, the speeds of these two machines are more like 10-20 megaflops. We shall use 50 megaflops here in an attempt to use a number both benign and realistic. It is important to note that this is substantially lower than the advertised speeds. For example, both American supercomputers have advertised speeds of several hundred megaflops. Parts of some programs can indeed compute this rapidly on them but real programs cannot effectively sustain this speed. As reported in [1], both machines run at less than 20 megaflops on a real mix of programs.

With the assumptions made above, a Cray or Cyber is about 50 times more powerful than a VAX. Thus, *a modern "supercomputer" would take about 6 years to compute a feature-length film*. This is still entirely too long.

WHAT ABOUT JAPANESE SUPERCOMPUTERS?

The Japanese firms of Fujitsu and Hitachi plan to have the supercomputers FACOM VP-200 at 500 megaflops and HITAC S-810 at 630 megaflops, respectively, operational by the end of 1983 [2]. If the top speeds are assumed, our hypothetical movie would take about 29 or 23 weeks to compute, respectively, but, as argued above, top speeds are not to be trusted for realistic speed measurements of supercomputers. In actuality, these machines will probably be comparable in speed with the Cray and Cyber supercomputers and hence too slow for full-length movies.

The Japanese have announced their intention of building a 10 gigaflop computer by 1989. The Super-Speed Scientific Computer Project has a total budget of about $100 million [2]. A 10G flop machine is about 20 times as powerful as a current supercomputer. The raw compute time for a film thus reduces to about 100 days, which is approaching a reasonable number for a feature film. *A supercomputer of the power planned by the Japanese would be sufficiently powerful to compute a feature-length film.* Both Cray Research and ETA (the recent supercomputer spinoff company of Control Data) have announced intentions of producing supercomputers of similar power by the late 1980s or 1990[3]. ETA, in fact, is planning a 30 gigaflop machine.

THE COST OF COMPUTING MOVIES

A conventional movie costs on the order of $10M with about $8M for production and $2M for postproduction and preproduction. Each one takes about two years. Only one of every ten movies succeeds on the average, but we shall not incorporate this unfortunate fact into our calculations here.

A current supercomputer installation costs $1-2M for initial site preparation and peripheral equipment plus about $3-4M per year for staff, building lease, supplies, utilities, and supercomputer maintenance. The machine itself costs about $10M; current prices range from $4M to $12M. Spreading the one-time costs of about $12M over three years means a per year cost of $4M. Thus the annual cost is $7-8M for a current supercomputer facility.

Comparing the production costs of a typical film with this annual supercomputer cost shows them approximately equal, if the machine is used for computing a new (successful) film every year. Of course, as we have seen above, current supercomputers cannot compute a film in a year. *So what we are awaiting are supercomputers of the power planned for 1990 at current "supercomputer" prices or less.*

FILM VERSUS FLIGHT

The production of complex graphics for films and realistic graphics for flight simulators - simulators for commercial aircraft, military aircraft, space shuttles, and even oil tankers - obviously have much in common. There are major differences, however. In film, a single path is taken through a rich terrain where the terrain changes from film to film; the time taken to compute a frame is of secondary importance to the richness of the image, several minutes per frame being tolerable. In flight, multiple paths are flown through a terrain which does not change from flight to flight; the time per frame is strictly that of *real time,* 24 frames per second for film or 30 frames per second for video, and image complexity is sacrificed to meet this schedule. Film implies great complexity at a cost of time, while flight implies real time at a cost of complexity. The data complexity for a film is concentrated along the one predetermined path through the database; Hollywood is famous for its false fronts. Flight simulators must spread data complexity over the wide range of potential paths which are not known in advance (except to within restricted "corridors" through the data).

COMPUTATIONAL REQUIREMENTS OF A FLIGHT SIMULATION

If a flight simulator were to display at the resolution used above for films, 13.5M bytes per frame, and at the complexity used above, 5000 operations per byte, then it would have to compute at a rate of 2.025 teraflops to satisfy the demands of real time (30 frames per second). This is four orders of magnitude faster than current "supercomputers" and two orders of magnitude faster than the projected 10 gigaflop supercomputers. So it is not feasible to even consider the use of a general-purpose supercomputer, current or projected, for flight simulations at movie resolution and complexity.

As already acknowledged, flight simulators typically sacrifice image complexity and resolution for speed. A calculation similar to that above but at video resolution (750K bytes per frame) and one-tenth the complexity (500 operations per byte) yields a computation rate of 11.25 gigaflops, in the ballpark for a 10-30 gigaflop machine. Therefore, *it is conceivable that a supercomputer of the power projected for the 1990s could compute a flight simulation in real time.*

This is important. Current flight simulators are special-purpose machines, generally one of a kind or a few of a kind. We are suggesting that a general-purpose supercomputer could conceivably replace these in the coming decade or two. A change would be a programming change only, not a new hardware development requiring great cost in time and money.

CONCLUSIONS

We have argued that the goal of completely computing a feature-length film (90 minutes) at a reasonable resolution (3K by 1.5K pixels) in a reasonable amount of time (2-3 months, nonstop) is not feasible with today's "supercomputers" but would be with supercomputers of the power planned by the Japanese and Americans for 1990. This assumes that the cost of such a machine would be about $5 to $10 million. This means that, without a special boost, we are unlikely to realize this particular goal in the US until about 1995-2000.

We have further argued that any general-purpose approach that improves movie-making computations also improves all image-making computations. Flight simulation, a case in point, could conceivably be computed on general-purpose supercomputers of the variety planned for the 1990s, although at a resolution and complexity less than that for film.

The main point is that it is not at all difficult to imagine what we would do with the supercomputers apparently still a decade away. We in the image-making professions are preparing ourselves for cost-effective supercomputing; the hunger has begun.

REFERENCES

[1] "The Computational Speed of Supercomputers", Ingrid Y. Bucher, *SIGMETRICS 83 Conference Proceedings,* ACM, August, 1983, Minneapolis, pp. 151-165.

[2] "Overview of Scientific Computing System - Next Generation Japanese Supercomputers", Hiroshi Kashiwagi, *Digest of Papers: Spring COMPCON 83,* IEEE, February 28-March 1, 1983, San Francisco, p. 305.

[3] "The Battle of the Supercomputers: Japan's All-Out Challenge to the U.S.", *Business Week,* October 17, 1983, pp. 156-166.