Computer History Museum

Volunteer Information Exchange

Sharing what we know with those we know

Volume 1 Number 9 June 17, 2011

Contribute To The VIE

We would like to hear from all you ex-HP'ers, Applers (Is that a word?) Intelers, and all the Silicon Valley types who have stories about your company. How about you SRI'ers, Stanford'ers, Berkel'ers?

(It's fun making up words.) How about you Googlers...

And all the staff and long time docents – you are a great repository of stories, facts and factoids that should be shared.

So, please send me your experiences, stories articles and jokes.

Jim Strickland jlstrick@aol.com

Questions and Answers

Q: We currently display four of the original 40 ENIAC panels. Where are the others?

Sue Mikell

A: This turned out to be a super interesting question. The short answer (subject to change -- boy, is it subject to change!) follows.

Note that, including the three portable function tables, there are 43 panels plus the card reader and punch units built especially for ENIAC by IBM.

See page 3 and 4 for additional information.

CHM *	4
Univ of Michigan **	4
Univ of Pennsylvania *	7
Smithsonian Institution	11
Perot Group **	8
Heinz Nixdorf Forum* Padeborn, Germany	1
Total	35

* On loan from Smithsonian
** On loan from the US Army

CONTENTS

VIE	1
Univac – 60 years old	2
Facts and Factoids 1st keyboard input	2
Computers in Space: 4 Pi	2
Where is Eniac?	2
Arthur Burks – Colleague of Mauchly and Eckert	2
Bay Are Travel Writer article on CHM	4
Iphone vs. Cray	7
Coming Events	9

Happy 100th Birthday to IBM

June 16, 1911 – 2011

(of course it was CTR from 1911 - 1924)

FACTS AND FACTOIDS

Factoid (Oxford English Dictionary) "something which becomes accepted as fact, although it may not be true." If you submit an item, please differentiate the facts from the factoids. And if you can verify something, thus changing it from a factoid to a fact, please let us know.

The First Software to Allow a Computer to be Operated by a Keyboard

- 1949 Betty Holbertson at Eckert-Mauchly develops UNIVAC Instructions Code C-10.
- C-10 was the first software to allow a computer to be operated by keyboarded commands rather than dials and switches. It was also the first mnemonic code.

(*From the Smithsonian Magazine; they rerun items from past issues.)*

60 YEARS AGO -- NOT YOUR FATHER'S LAPTOP

UNIVAC, the first commercially successful computer, gets an initial public workout in Philadelphia, June 14,1951. Five years in the making by J. Presper Eckert and John Mauchly (inventors of the earlier ENIAC), UNIVAC (universal automatic computer) is 141/2 feet long, 71/2 feet wide and 8 feet high. Bought by the Census Bureau, it spits out demographic data at 120 facts per second. Forty-six UNIVACs, with prices starting at $600,000, will be sold for inventory payroll, insurance and other business applications.

COMPUTERS IN SPACE

The IBM 4 Pi computer

IBM had been involved in the space program since its inception. The IBM CPC (Card programmed Calculator) was used in the Redstone rocket program ca. 1950 and the first computer in space was built by IBM for the Gemini program. And, several space and defense programs built special purpose computers.

But just as general purpose computers are usually less expensive than special purpose computers on the ground, so they should be in space, so in 1970, IBM announced System /4 Pi for the space and aviation industries.

System/4 Pi was a family of radiation hardened avionics computers used, in various versions, on the B-52 Stratofortress bomber, the F-15 Eagle fighter, E-3 Sentry, NASA's Skylab space station, MOL, and the Space Shuttle, as well as other aircraft.

The top-of-the-line 4 Pi was the AP-101, used in the B-52. The U.S. Navy used a similar variant in the carrier based A-6E/A-6E TRAM medium attack aircraft. The Shuttle is controlled by five AP-101s, four of which are arranged in a redundant configuration, with the fifth as backup. Skylab employed the model TC-1, which had a 16-bit word length in contrast to the AP-101's 32 bits.

The name of the system is derived from the fact that the angular measure of a complete sphere (solid angle) is 4π steradians, while the angular measure of a complete circle is 360 degrees; hence System/4 Pi and System/360. This implies that System/4 Pi is a version of the IBM System/360 for the three-dimensional world of avionics.

Indeed, programs for the 4 Pi could be written, compiled and tested on a mainframe /360 greatly simplifying early phases of development and test.

I'm thinking of getting off Facebook, getting off Twitter. You know why?

I just want to sign up for a service that, every hour, texts me the message 'You're not alone.' I think that would do the trick.

Pete Holmes

Where Is Eniac?

- *The following is from Wikipedia.*

    The School of Engineering and Applied Science at the University of Pennsylvania has four of the original forty panels and one of the three function tables of the ENIAC. The Smithsonian has five panels in the National Museum of American History in Washington DC. The Science Museum in London has a receiver unit on display. The Computer History Museum in Mountain View, California has a single panel on display. The University of Michigan in Ann Arbor has four panels, salvaged by Arthur Burks. The U.S. Army Ordnance Museum at Aberdeen Proving Ground, Maryland, where ENIAC was used, has one of the function tables. There is also a panel on display at Perot Systems in Plano, Texas.

But, we know the part about the CHM is incorrect, how about the rest?

- According to the curator of the Perot collection in Plano Texas, the US Army Military History Museum considers themselves to be the owner of ENIAC. As you know ENIAC was funded by the Army and installed at the Aberdeen Proving Grounds to do firing tables for new artillery. (Note that the Smithsonian Institution also considers themselves to be the owner of ENIAC.) She found that after it was decommissioned, panels were sent to several Army locations where they had no idea what to do with them. Most just got rid of them.
- Fort Sill, Oklahoma (which has an artillery mission) was given 8 panels. She found them and negotiated so that Ross Perot would restore them and display them until Ft. Sill had an appropriate place to display them. (Ft. Sill's Artillery Museum is underway but not ready for ENIAC yet.)
- Perot has restored the 7 panels, (the 8th is beyond help.) They are on display and have lights on the accumulator set to flash based on a motion sensor.
- Further, our research suggests that: in addition to Perot's eight:
    - ✔ U of Pennsylvania has 4 panels.
    - ✔ University of Michigan has a function table and 3 other panels, from Arthur Burks. (See article on page 4)
    - ✔ The Smithsonian has as many as 11 panels. (Kim Harris recently visited and tells us they are not on display.)
    - ✔ An accumulator panel is on display at the Heinz Nixdorf Forum in Padeborn, Germany.
    - ✔ That makes a total of 35 extant. (Various smaller parts are in private and public collections.)
- The picture below shows the seven panels on display at the Perot Group lobby in Plano, Texas. The rightmost large panel is purposely turned so that visitors can see a panel from both sides and see its many vacuum tubes. The small unit at the right is a "decade ring counter."

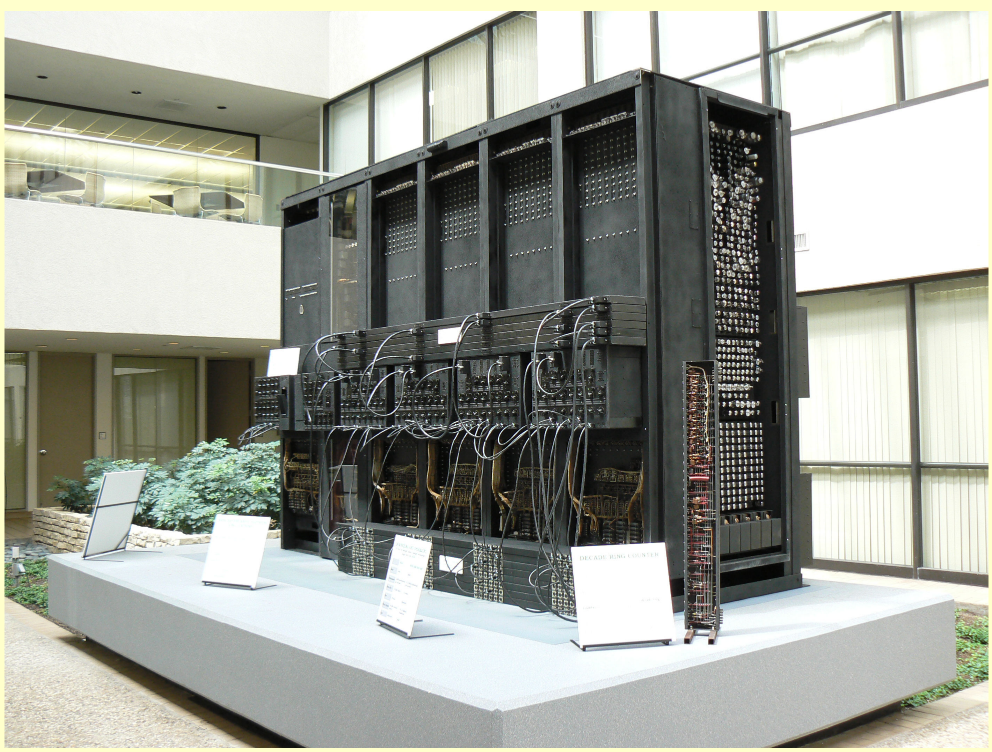

Arthur Burks – Colleague of Mauchly and Ekert

From Wikipedia **Arthur Walter Burks** (October 13, 1915 – May 14, 2008) was an American mathematician who in the 1940s as a senior engineer on the project, contributed to the design of the ENIAC,

When Mauchly and Eckert's proposed concept for an electronic digital computer was funded by the U.S. Army's Ballistics Research Laboratory in June 1943, Burks was added to the design team. Among his principal contributions to the project was the design of the high-speed multiplier unit. (Also during this time, Burks met and married Alice Rowe, a computer employed at the Moore School.)

In April 1945, with John Grist Brainerd, Burks was charged with writing the technical reports on the ENIAC for publication. Also during 1945 Burks assisted with the preliminary logical design of the EDVAC in meetings attended by Mauchly, Eckert, John von Neumann, and others.

On March 8, 1946 Burks accepted an offer by von Neumann to join the computer project at the Institute for Advanced Study in Princeton, New Jersey, and joined full time the following summer. (Already on the project was another member of the ENIAC team, Herman Goldstine. Together, Goldstine and Burks gave nine of the Moore School Lectures in Summer 1946.) During his time at the IAS, Burks worked to expand von Neumann's theory of automata.

In the 1960s he was presented with the opportunity to acquire four units of the original ENIAC, which had been rusting in a storage Quonset hut in Aberdeen, Maryland. He ran the units through a car wash before restoring them and donating them to the University of Michigan. They are currently on display in the entryway of the Computer Science Building.

After working on this project, Burks relocated to Ann Arbor, Michigan in 1946 to join the faculty of the University of Michigan, first as an assistant professor of philosophy, and as a full professor by 1954. He helped found the university's computer science department, first as the Logic of Computers group in 1956, of which he was the director, then as a graduate program in 1957, and then as an undergraduate program within the new Department of Computer and Communication in 1967,

Over the past two months we have been pitching the Museum to travel writers, to help drive local and tourist visitors to Revolution. Below is one of two articles/postings that was just published on a family-centric travel site: Bay Area Travel (The second article will appear in the next issue.)

Computer History Museum: The Famous Past in our own Backyard

Bay Area Family Travel Joyce Kiefer June 8, 2011

http://www.bayareafamilytravel.com/index.php?pr=SoBayMuseum

Once upon a time, back in the middle of the last century, computers were bigger than the palm of your hand. Way bigger. They were large and boxy with banks of tape decks the size of refrigerators and ran on vacuum tubes. They were fed by punched cards and spit out information on loops of wide printer paper. They had flashing lights but no screens. But did you know that computers go farther back than the 1950's? There was the abacus, the Babbage Engine, and . . .

I wanted my granddaughters Gwynne, Age 9, and Amri, 16, to learn the tale of computers before they left Silicon Valley to move to Washington State, so I took them to the newly renovated Computer History Museum in Mountain View, California. What better companions could we bring along with us than two older generations of computer experts - their father, a software developer, and their Grandpa Bill, a retired programmer?

The museum presents "R/evolution, the First 2000 Years of Computing." (The line between the R and the E in "revolution" suggests both the social and the technical points of view.) The location, fittingly, is the former headquarters of a Silicon Valley giant, Silicon Graphics (SGI). The museum reopened early this year after a two-year expansion and remodel. Now you can follow the development of computers through 19 alcoves that each feature a landmark artifact and describe its importance. Accompanying videos show how it works and present interviews with its developers.

WHAT TO EXPECT

We started at the beginning, although we could have followed directions on the Visitor Map and gone directly to whatever was of interest. The museum presents a five-minute orientation film, but our family experts chose to bypass this. However, I think it's helpful to find out how a museum introduces its purpose and the basic theme of its exhibits.

As we began to move through the exhibit alcoves, Bill spotted a huge slide rule. Nostalgia swept over him; he used slide rule in high school trigonometry. I, in turn, was stopped by the Hollerith Census Machine. When I started college in the mid-1950's, I thought my course registration cards filled with numbers and holes meant to be read by computers had to be the latest thing - a symbol of the modern world I was about to enter on my own. But no, Herman Hollerith invented the punched card machine to process the census back in 1890!

I look at computers from a practical aspect: what can they do for me? I don't want to know the process involved to get the result I'm after. Just tell me what steps to follow and I'll write them down. "No, don't write anything down; just listen to me!" said my former boss and my husband. They and my son-in-law focus on the logic, mechanics, and the math that drive the computer to do what it does. The quote on one of the museum walls by inventor/professor Donald Knuth echo the way these techies regard computers: "Computer programming is an art because it applies accumulated knowledge to the world, because it requires skill and ingenuity, and because it produces objects of beauty."

KIDS'S VIEW

Our nine-year-old granddaughter Gwynne has never seen a computer larger than a desktop PC. So the portion of a 1940's ENIAC computer on display was beyond her comprehension. One of these huge computers filled a 1,000 square foot room but had much less storage than a laptop. I recall a New Yorker cartoon of a giant computer spitting out a small message. The technician reads, "It wants a goat sacrificed to it."

She will never get the joke.

But all of us were fascinated by the German Enigma encryption machine from World War II. It looked like a portable typewriter (Gwynne hasn't seen one of those, either) but it converted each letter typed into another letter to create a code. However, the Allies figured out the code before WWII began and were able to keep a step ahead of the enemy. What messages went out from this little machine? Secret codes are the stuff of mystery and fantasy for any generation.

Bill was swept by nostalgia at the display of the colorful IBM 360 with its blue and orange disk drives. He thought of the mainframe computers he used throughout his career from the early '60's to 2004 at Hercules Powder Company, Control Data and Stanford University. And he flashed on a scene from the good old days of mid-'60's, which he described to me.

"I'm standing in a large computer room, chilly because the machine needed air conditioning to keep from overheating. Inside is a Control Data 3800 mainframe, a card reader, a printer and 10 tape decks, each with a reel mounted inside. Several of us stand by expectantly. Then the computer comes alive. The mainframe D-register lights start blinking and the computer begins to perform. Card reader starts ingesting cards. Printer chugs out loops of paper in rhythm. Tape reels jiggle back and forth in time. Carefully programmed, the computer is playing "Stars and Stripes Forever."

Not as spectacular but still fun: the twice-monthly museum demo of the musical abilities of Digital Equipment Corporation's PDP-1, built in 1959. See details at the end of the column.

SUPER COMPUTERS AND MORE

Much more lay ahead as we explored the alcoves: the Cray-1Supercomputer, music and art, robotics, games, networking. But only about 2% of the museum holdings are on display. The website presents photos and descriptions of most of the rest. The Computer History Museum claims to be the world's largest collecting and exhibiting institution of its kind that explores the history of computing and its ongoing impact on society, more than you'll see on the floor.

Even the most dedicated geek may feel overwhelmed and need a break. The café in the lobby offers far more than orange soda. You can get sandwiches, salads, coffee, and sweets and eat them in the fresh air on the patio. When you look down at the lobby floor, you'll notice that the floor tiles are configured as a giant punched card.

We should have allowed about four hours for our tour and included a break. But since the museum was going to close before we'd seen everything, I looked hard to find the first computer I ever used. In Fall of 1984 I began a new job. Sitting on my desk was the envy of everyone in my department: the new Apple MacIntosh that had been announced so dramatically in a Superbowl commercial that January. But I didn't know what to do with it. My boss, a graduate of MIT, couldn't understand why I couldn't teach myself. Eventually I did conquer that little box and I still use a Mac. How ironic to see this once revolutionary invention behind glass as a museum piece!

"It looks like a microwave oven," Gwynne observed.

EARLY GAMES

One of the few hands-on sections is the games section. I would like to see this section expanded with more terminals and information on how games are designed but not at the expense of other displays. However, if you show up on the first or third Saturdays of the month at 3 p.m., you can play Spacewar! in the PDP-1 Room in the wing left of the lobby where there's no admission charge. Digital Equipment created Spacewar! in 1959 to show off its first computer, PDP-1, which did something revolutionary: it focused on interaction with the user. The game has two players, each in command of a spaceship, who each try to destroy the other against an astronomically correct star field.

Also in the free section of the museum you'll find an alcove on Computer Chess with online documents, images, artifacts, oral histories, moving images, and software related to the game from 1945 to 1997. You'll be right next to the Babbage Engine. This was 16-year-old Amri's favorite computer.

In the mid-1800's Charles Babbage designed a "Difference Engine" to mechanically calculate and tabulate mathematical functions called polynomials. An attached printer would produce out the results. He persuaded the British government to fund the project. Twelve years later, the machine was still unfinished. Tired of waiting, the government withdrew its support. But Babbage didn't give up. He designed a sleeker model (weighing only five tons) and called it "Difference Engine No. 2. But he never saw it built. That was done in the early 2000's in London. Former Microsoft guru Nathan Myhrvold ordered a duplicate and has lent it to the Computer History Museum until he can shore up the floor of his living room to support the five-ton machine.

I returned to the museum a few weeks later to watch Bill's life-long co-worker give a presentation on the Babbage Engine and have an assistant actually make it work. First he wheeled out a blackboard and described the kind of calculations it would do. Then his assistant turned the crank and the five-ton machine ticked and clacked to life. Pins in the back rippled in the shapes of DNA spirals. Parts moved sideways and up and down. It was a marvel to behold.

John Hollar, the museum president, describes the ambitious purpose of "R|evolution," The First 2,000 Years of Computing: "It enables us to vastly expand our interpretation of computer evolution, its deep roots in Silicon Valley and beyond, its growth into a singular force in global life, and its ongoing social and historical impact."

After touring the museum I was left with the feeling that the story and use of computers has far to go in directions we can't imagine, just as it has in the past. I wondered how my grandchildren's kids will look at smart phones and I-Pads.

DETAILS: Computer History Museum 1401 North Shoreline Blvd. Mountain View

DIRECTIONS: Head east on Highway 101 at the Shoreline exit.

MORE DETAILS: 650/810-1010
Hours: Wed.-Sun. 10 a.m. - 5 p.m. Admission: Adults - $15, Students/Seniors - $12, children under 12- free.

WEB: Go to http://www.computerhistory.org/visit/

The Computer History Museum offers many online exhibits on a variety of topics related to the history of computing. Some online exhibits like Mastering the Game complement the physical exhibits you see when you visit in person. Others are available only through the Internet. The website also provides information about museum lectures.

SPECIAL DAYS: Babbage Engine demos - These hour long presentations take place every Saturday and Sunday at 1 p.m.

PDP-1 demos - Play Spacewars! with a partner. Listen to the computer make its own music. Go the lst or 3rd Saturday of the month at 3 p.m. for this one-hour presentation. Free.

Docent-led tours - Docents are available for informal one-hour tours that focus on a theme such as communication as it revolves around five or six artifacts. Call to confirm.

-

If Honeywell and Fairchild merged, what would the name be? Fairwell Honeychild!
From Kim Harris

Iphone Vs. Cray

Last fall, during one of our docent training sessions, someone made the remark that todays iPhone is more powerful than the Cray-1. Having worked with supercomputers during the 1980s and therefore aware of the kinds of problems they were designed to solve, this comparison just didn't sound legitimate to me. I followed up with questions to some of the docents and received an excellent analysis of the comparison from Mike Albaugh. My conclusion after reading Mikes analysis was that I should refrain from making the comparison on my tours because they have radically different architectures, they are many generations of technology apart, and each is optimized to run a different set of applications efficiently.

However, just last weekend, I got exactly that question from one of the visitors, so I thought it would be worth sharing Mikes work with the other docents. Mike follows his conclusions below with a section called Random Data Dump which provides the raw data on which he based his analysis.

Sue Mikell

Mike Albaugh's Updated Answer

Very Briefly: "It depends", but judged by the measure that placed the Cray 1 at the top of the list of Supercomputers, in 1979, the iPhone actually beats the Cray-1. That measure is Millions of Floating Point Operations per second (MFLOPS), and the Cray-1 came in at 3.4 MFLOPS on the LINPACK (n=100) benchmark in 1979. This is the first year I found data for [1]. By 1983, the same chart shows the Cray-1 at 12 MFLOPS, probably as a result of both increased memory sizes and better compilers. Meanwhile, I found a site [2] which ran that benchmark on the iPhone (3G, 3GS, and 4), with results of 18.2 to 20.9 MFLOPS.

"We walk from here":

Note that neither of those measured benchmark numbers are anywhere near the best-case (What John Mashey calls "Speed of light") number for the respective machines. The Cray-1 was rated at 135 MFLOPS [3] with an 80MHz clock and the iPhone 4 (making some assumptions about the F.P. unit I have not confirmed) would be rated at 2GFLOPS with a 1GHz clock [4]. Why is that, and why does it matter?

1) Real problems don't typically massage a small data set for hours. A major advantage of modern systems is the sheer amount of memory they contain, driven by Moore's Law which has its greatest effect on memory. Assuming an iPhone4, it has the same amount of "main memory" as the largest Cray-1M (1982) had in Total main memory and "Solid State Storage). Or four times the "main memory" of the Cray-1M (32 times the main memory of the original Cray-1, in 1975).

2) Today's computers depend quite heavily on "the Memory Hierarchy", where the CPU has fast access to a small memory, slightly slower access to a somewhat larger memory, etc. through many layers (e.g. 4 in the iPhone4, as far as I can tell. As many as 7 in a modern desktop [5]).
If the problem at hand fits well into the innermost layers of that model, the CPU can run close to "full speed". If not, things can slow dramatically. Staying in "main memory or closer", the worst-case speed of memory access of the iPhone is a little worse than the Cray-1, with latency 1.2 to 2 times. If, however, "the stars align", the iPhone4 can be up to 12 times as fast as the Cray-1.

3) As the auto ads say: "Your Mileage May Vary". We do not run the same software, or even remotely similar software, on an iPhone as was run on a Cray-1. On the other hand, the rule of thumb for emulating one computer on another, given roughly similar architectures and more memory than the original, is one order of magnitude. Another factor of 1.4 could be needed to account for the narrower memory of the iPhone.

4) So, the iPhone4 is (roughly) 1/2 to 12 times the speed of the Cray-1. Which implies that for some code, you could expect to emulate [6] a Cray-1 on your iPhone at (roughly) the original speed, as long as you do not do a lot of access to "main memory". Tweaking and re-compiling the code to run "native" on the iPhone could speed up those applications to run faster, up to maybe 6x. Of course, people who bought Crays tended to have problems which did have a lot of accesses to main memory, so you would have to cherry-pick your demo.

Note that 2010 - 1975 = 35 years, or 23 Moore's Law iterations, so the base technology (might have) improved 2^{23} times, or a factor of about 8 million)

Side Note: Things get really muddy around software, which always grows to fill the available memory and take time just short of frustrating the user. When I was preparing my talk on the LGP-30, I did a "back of the envelope" calculation, and found that in the approximately one minute that a modern laptop takes to go from power-on to a usable desktop, it executes as many instructions as (best case) an LGP-30 running 24/7 could do in about 9.5 years, the approximate service life of an LGP-30.

Some pre-emptive responses:

One criticism I have gotten about the thought of rough equivalence is "ARM does not have floating point, so it's nothing like a Cray". That was true of the ARMs introduced while the Cray-1 was still king of supercomputers, but not today. There are several (incompatible) floating-point extensions to the ARM instruction set. Most provide IEEE-754 32 and 64-bit F.P., and some provide vector operations. It may very well be that the ARM (thus iPhone) F.P. could not – exactly -- match Cray performance, but it is not true that you would have to fall back to software emulation. (except as noted in [6])

I already mentioned a slug of caveats about the application mix being different for the two machines. YMMV, but while blazing F.P. was indeed one of the main reasons for buying a Cray, [7] the ability to stream large datasets through that F.P. was a nearly equal reason, and the iPhone ARMs have pretty decent memory bandwidth (and latency), compared to a 1975 machine. :-)Note that "programmers these days" are totally hooked on F.P., so even a spelling checker likely depends on it. Javascript definitely does, so the iPhone has to be decent. Also, video games (and other image-processing tasks) can use the 16 and 32 bit floating point that the (later) ARM cpus have and the Cray lacks. This is another case of "horses for courses."

"What can't you do with an iPhone that you could with the Cray-1?" Well, you can't develop software for it, on it. At least not without violating your EULA. You also can't log into it remotely (again, without violating your EULA). And you can't sleep on the power supply.

Notes and further reading

[1]
http://www.netlib.org/utk/people/JackDongarra/faq-linpack.html
1979 is the first year listed in Jack Dongarra's FAQ for the LINPACK benchmark. While this has now been pretty much superceded, it was an important measure of (super)computer performance at the time.

[2] http://www.xyster.net/blog/?p=40
Note that this performance was gotten with the C version of the benchmark (so less amenable to optimization that the Fortran version almost certainly used by the Cray), and with n=1000.
The larger n is (partly) to make the test run longer, for better timing accuracy, but requires more memory. n=1000 numbers are not available at that site for prior to 1986, when the Cray-1 was no longer a "Supercomputer", so not listed. It was widely used as a comparison for folks trying to convince you that their mini or micro was "supercomputer-like".

[3] Some raw Cray-1 specs:
http://en.wikipedia.org/wiki/Cray-1

Cray-1 ran at 80MHz, with 16-way interleaved memory (8MBytes, sounds dubious, as 24 bits of word address would be 16MW, 127MB), 4Wds every 50ns, (So again, memory at 1/4 CPU, but with multiple active banks could run at CPU Speed.) 160MIPS, 134MFlops, 250 Mflops with vector chaining. Again via WikiPedia, 1/4, 1/2, and 1MW available on Cray-1 (1975), 1/2/4 MW on Cray-1S (1979) Vector registers (8 x 64) roughly equivalent to 4KB L1 cache (1/8th of A4 L1 Dcache). Instruction buffer only 16 Words (128B) compared to A4 32KB L1 Icache) Cray-1M(1982) could have up to 32MW (256MB) of MOS memory as Solid State Storage.

[4] Some raw iPhone specs:
http://www.gsmarena.com/apple_iphone-1827.php
says original iPhone has 412MHz ARM 11, 4/8/16GB flash. un-spec'd RAM
http://furbo.org/2007/08/21/what-the-iphone-specs-dont-tell-you/
says about 128MB RAM, CPU @400MHz, Bus-speed 100MHz.
http://vr-zone.com/articles/iphone-3g-s-cpu-and-memory-specs-revealed/7195.html
says iPhone 3GS is 600MHz CPU, 256MB RAM, 32GB Flash, HW OpenGL acceleration
http://www.computerandvideogames.com/article.php?id=249927
says iPhone4 is 1GHz A4, 512MB RAM, 32GB flash
"http://pdadb.net/index.php?m=cpu&id=a40000&c=apple_a4"
has a few few A4 specs, basically: 13-stage pipeline, 64KB (32+32) L1 cache (1 or 2 clocks, probably) 640KB L2 cache (4..10 clocks, probably).
DDR2 64-bit main memory interface. Assuming A4 also has CPU/4 memory clock, main memory latency is (4..6 mem-clocks x 4nSec = 16..25 nSec. About 1.2..2 x Cray latency. (based on DDR2 article at http://en.wikipedia.org/wiki/DDR2_SDRAM) Not

clear if 256MB is "integrated" on chip or in package, so not really sure of speed. Muddying the waters a bit, according to "http://news.cnet.com/8301-31021_3-20008089-260.html"
the iPhone4 has 512MB, so some inside, some out?

5] The Cray-1 memory hierarchy was nearly "flat", with just vector-registers and a tiny icache. main-memory, and a solid-state equivalent of "very fast disk".

[6] As far as I can tell, and remember, the Cray-1 did not use IEEE-754 arithmetic, so it's entirely possible that if you were actually writing a Cray emulator to get bit-for-bit-equal answers, you might indeed have to go to software emulation, and slow dramatically.

[7] Some of Cray's best customers used little or no F.P. The iPhone make lack a "Pop-count" instruction as well. :-).

More on Similar topics coming …

Dag Spicer suggests that we check out Jack Dongarra. Following is an excerpt from:

http://bits.blogs.nytimes.com/2011/05/09/the-ipad-in-your-hand-as-fast-as-a-supercomputer-of-yore/

His [Jack Dongarra's] research group has run the test on Apple's new iPad 2, and it turns out that the legal-pad-size tablet would be a rival for a four-processor version of the Cray 2 supercomputer, which, with eight processors, was the world's fastest computer in 1985.

Dr. Dongarra's researchers also discovered that the new iPad2 is about 10 times as fast as its predecessor, the original iPad. That is likely because of some design changes in the microprocessor used in the new version of the Apple tablet.

To date, the researchers have run the test on only one of the iPad microprocessor's two processing cores. When they finish their project, though, Dr. Dongarra estimates that the iPad 2 will have a Linpack benchmark of between 1.5 and 1.65 gigaflops (billions of floating-point, or mathematical, operations per second). That would have insured that the iPad 2 could have stayed on the list of the world's fastest supercomputers through 1994.

Coming Events

Date	Day	Time	Event
Jun 29	Wed.	12 Noon	CHM Soundbytes - The History of Magnetic Striped Media Technology – A Lecture by Jerome Svigals
Jul 14	Thur.	6:00 PM 7:00 PM	Reception Something Ventured (film)
Aug 24	Wed.	12:00:00 PM	Software patent debate
Nov 05	Sat.	02:00:00 AM	The Challenge and Promise of Artificial Intelligence, a Bay Area Science Festival Wonder Dialog
Nov 08	Tues.	02:00:00 PM	The Technology of Animation DreamWorks Animation's Jeffrey Katzenberg and Ed Leonard will kick off this series, in a conversation moderated by HP's Phil McKinney.

Please contribute to the Computer History Museum Volunteer Information Exchange.

Share your stories, your interesting facts (and factoids) and your knowledge.
Send them to Jim Strickland (Jlstrick@aol.com)

Contribution is voluntary BUT please volunteer to contribute.