



BY GORDON BELL

BELL'S LAW FOR THE BIRTH AND DEATH OF COMPUTER CLASSES

A theory of the computer's evolution.

In the early 1950s, a person could walk inside a computer and by 2010 a single computer (or “cluster”) with millions of processors will have expanded to the size of a building. More importantly, computers are beginning to “walk” inside of us. These ends of the computing spectrum illustrate the vast dynamic range in computing power, size, cost, and other factors for early 21st century computer classes.

A computer class is a set of computers in a particular price range with unique or similar programming environments (such as Linux, OS/360, Palm, Symbian, Windows) that support a variety of applications that communicate with people and/or other systems. A new computer class forms and approximately doubles each decade, establishing a new industry. A class may be the consequence and combination of a new platform with a new programming environment, a new network, and new interface with people and/or other information processing systems.



Bell's Law accounts for the formation, evolution, and death of computer classes based on logic technology evolution beginning with the invention of the computer and the computer industry in the first-generation, vacuum-tube computers (1950–1960), second-generation, transistor computers (1958–1970), through the invention and evolution of the third-generation Transistor-Transistor Logic (TTL) and Emitter-coupled Logic (ECL) bipolar integrated circuits (ICs) from 1965–1985. The fourth-generation MOS and CMOS ICs enabling the microprocessor (1971) represents a “break point” in the theory because it eliminated the other early, more slowly evolving technologies. Moore's Law [6] is an observation about integrated circuit semiconductor process improvements or evolution since the first IC chips, and in 2007 Moore extended the predic-

one as a wireless sensor network class. Field Programmable Logic Array chips with tens to hundreds of millions of cells exist as truly universal devices for building nearly anything.

BELL'S LAW

A computer class is a set of computers in a particular price range defined by: a programming environment such as Linux or Windows to support a variety of applications; a network; and user interface for communication with other information processing systems and people. A class establishes a horizontally structured industry composed of hardware components through operating systems, languages, application programs and unique content including databases, games, images, songs, and videos that serves a market through various distribution channels.

The universal nature of stored-program computers is such that a computer may be programmed to

$$\text{Transistors per chip} = 2^{(t-1959)} \text{ for } 1959 \leq t \leq 1975; 2^{16} \times 2^{(t-1975)/1.5} \text{ for } t \geq 1975.$$

tion for 10–15 more years, as expressed in Equation 1. The evolutionary characteristics of disks, networks, displays, user interface technologies, and programming environments will not be discussed here. However, for classes to form and evolve, all technologies must evolve in scale, size, and performance at their own—but comparable—rates [5].

In the first period, the mainframe, followed by minimal computers, smaller mainframes, supercomputers, and minicomputers established themselves as classes in the first and second generations and evolved with the third-generation integrated circuits circa 1965–1990. In the second or current period, with the fourth generation, marked by the single processor-on-a-chip, evolving large-scale integrated circuits (1971–present) CMOS became the single, determinant technology for establishing all computer classes. By 2010, scalable CMOS microprocessors combined into powerful, multiple processor clusters of up to one million independent computing streams are likely. Beginning in the mid-1980s, scalable systems have eliminated and replaced the previously established, more slowly evolving classes of the first period that used interconnected bipolar and ECL ICs. Simultaneously smaller, CMOS system-on-a-chip computer evolution has enabled low-cost, small form factor (SFF) or cell-phone-sized devices (CFSD); PDA, cell phone, personal audio (and video) devices (PADs), GPS, and camera convergence into a single platform will become the worldwide personal computer, circa 2010. Dust-sized chips with relatively small numbers of transistors enable the creation of ubiquitous, radio networked, implantable, sensing platforms to be part of everything and every-

replicate function from another class. Hence, over time, one class may subsume or kill off another class. Computers are generally created for one or more basic information processing functions—storage, computation, communication, or control. Market demand for a class and among all classes is fairly elastic. In 2010, the number of units sold in classes will vary from tens for computers costing around \$100 million to billions for SFF devices such as cell phones selling for under \$100. Costs will decline by increasing volume through manufacturing learning curves (doubling the total number of units produced results in cost reduction of 10%–15%). Finally, computing resources including processing, memory, and network are fungible and can be traded off at various levels of a computing hierarchy (for example, data can be held personally or provided globally and held on the Web).

The class creation, evolution, and dissolution process can be seen in the three design styles and price trajectories and one resulting performance trajectory that threatens higher-priced classes: an established class tends to be re-implemented to maintain its price, providing increasing performance; minis or minimal-cost computer designs are created by using the technology improvements to create smaller computers used in more special ways; supercomputer design (the largest computers at a given time come into existence by competing and pushing technology to the limit to meet the unending demand for capability); and the inherent increases in performance at every class, including constant price, threaten and often subsume higher-priced classes.

All of the classes taken together that form the com-

puter and communications industry shown in Figure 1 behave generally as follows:

- Computers are born—classes come into existence through intense, competitive, entrepreneurial action over a period of two to three years to occupy a price range, through the confluence of new hardware, programming environments, networks, interfaces, applications, and distribution channels. During the formation period, two to hundreds of companies compete to establish a market position. After this formative and rapid growth period, two or three, or a dozen primary companies remain as a class reaches maturity depending on the class volume.

- A computer class, determined by a unique price range, evolves in functionality and gradually expanding price range of 10 maintains a stable market. This is followed by a similar lower-priced subclass that expands the range another factor of five to 10. Evolution is similar to Newton's First Law (bodies maintain their motion and direction unless acted on externally). For example, the “main-frame” class was established in the early 1950s using vacuum tube technology by Univac and IBM and functionally bifurcated into commercial and scientific applications. Constant price evolution follows directly from Moore's Law whereby a given collection of chips provide more transistors and hence more performance.

A lower entry price, similar characteristics subclass often follows to increase the class's price range by another factor of five to 10, attracting more usage and extending the market. For example, smaller “mainframes” existed within five years after the first larger computers as sub-classes.

- CMOS semiconductor density and packaging inherently enable performance increase to support a trajectory of increasing price and function.

Moore's Law single-chip evolution, or microprocessor computer evolution after 1971 enabled new, higher performing and more expensive classes. The initial introduction of the microprocessor at a substantially lower cost accounted for formation of the initial microcomputer that was programmed to be a calculator. This was followed by more powerful, more expensive classes forming including the home computer, PC, workstation, the shared microcomputer, and eventually every higher class. Home and personal computers are differentiated

from workstations simply on “buyer”—a person versus an organization.

The supercomputer class circa 1960 was established as the highest performance computer of the day. However, since the mid-1990s supercomputers are created by combining the largest number of high-performance microprocessor-based computers to form a single, clustered computer system in a single facility. In 2010, over a million processors will likely constitute a cluster. Geographically coupled computers including GRID computing, such as SETI@home, are outside the scope.

- Approximately every decade a new computer class forms as a new “minimal” computer either through using fewer components or use of a small fractional part of the state-of-the-art chips. For example, the hundredfold increase in component density per

decade enables smaller chips, disks, and screens at the same functionality of the previous decade especially since powerful microprocessor cores (for example, the ARM) use only a few (less than 100,000) transistors versus over a billion for the largest Itanium derivatives.

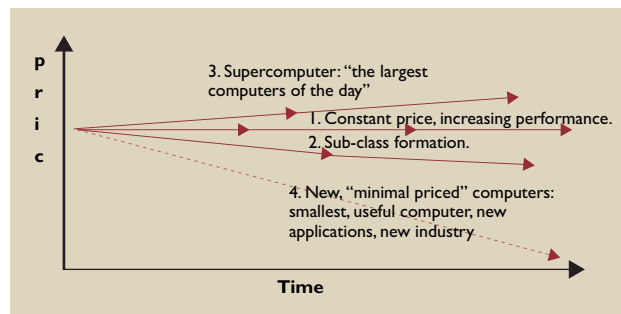


Figure 1. Evolving computer classes based on technology and design styles.

Building the smallest possible computer accounts for the creation of computers that were used by one person at a time and were forerunners of the workstation (for example, the Bendix G-15 and LGP 30 in 1955), but the first truly personal computer was the 1962 Laboratory Instrument Computer (LINC). LINC was a self-contained computer for an individual's sole use with appropriate interfacial hardware (keyboards, displays), program/data filing system, with interactive program creation and execution software. Digital Equipment's PDP-1 circa 1961, followed by the more “minimal” PDP-5 and PDP-8 established the minicomputer class [1] that was predominantly designed for embedded applications.

Systems-on-a-chip (SOCs) use a fraction of a chip for the microprocessor(s) portion or “cores” to create classes and are the basis of fixed-function devices and appliances beginning in the mid-1990s. These include cameras, cell phones, PDAs, PADs, and their convergence into a single CPSD or SFF package. This accounts for the PC's rapidly evolving microprocessor's ability to directly subsume the

1980s workstation class by 1990.

- Computer classes die or are overtaken by lower-priced, more rapidly evolving general-purpose computers as the less-expensive alternatives operating alone, combined into multiple shared memory microprocessors, and multiple computer clusters. Lower-priced platforms result in more use and substantially higher volume manufacture thereby decreasing cost while simultaneously increasing performance more rapidly than higher-priced classes.

Computers can be combined to form a single, shared-memory computer. A “multi” or multiple CMOS microprocessor, shared-memory computer [2] displaced bipolar minicomputers circa 1990 and mainframes circa 1995, and formed the basic component for supercomputers.

Scalable, multiple computers can be networked into arbitrarily large computers to form clusters that replace custom ECL and CMOS vector supercomputers beginning in the mid-1990s simply because arbitrarily large computers can be created. Clusters of multiprocessors were called constellations; clusters using low latency and proprietary networks are MPPs (massively parallel processors).

Generality *always* wins. A computer created for a particular, specialized function, such as word processing or interpreting a language, and used for a particular application is almost certain to be taken over by a faster-evolving, general-purpose computer. The computer’s universality property allows

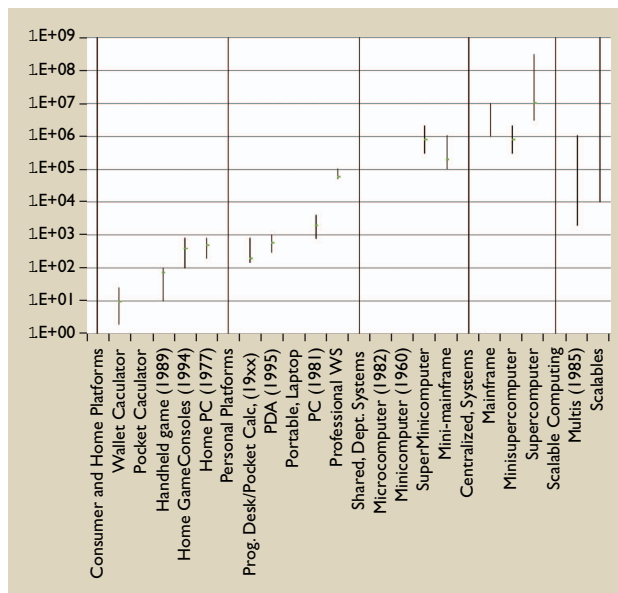


Figure 2a. Computer classes and their price range circa 2005.

any computer to take on the function of another, given sufficient memory and interfaces.

SFF devices subsume personal computing functionality as they take on the communications functions of the PC (email and Web browsing), given sufficient memory and interfaces. SFF devices, TVs, or kiosks accessing supercomputers with large stores, subsume personal computing functionality. The large central stores retain personal information, photos, music, and video.

The specific characteristics of the classes account for the birth, growth, diminution, and demise of various parts of the computer and communications industry.

OVERVIEW OF THE BIRTH AND DEATH OF THE COMPUTER CLASSES 1951–2010

The named classes and their price range circa 2010 are given in Figure

2a. In 1986, David Nelson, the founder of Apollo computer, and I posited that the price of a computer was approximately \$200 per pound [7]. Figure 2b gives the introduction price and date of the first or defining computer of a class.

Here, I will use the aspects of Bell’s Law described previously and follow a timeline of the class formations beginning with the establishment of the first computer classes (mainframe, supercomputer, shared personal professional computers or workstations, and minicomputers) using vacuum tubes, transistors, and bipolar integrated circuits that continue through the mid-1990s in the first period (1951–1990). In the second period beginning in 1971, the MOS microprocessor ultimately overtook bipolar by 1990 to

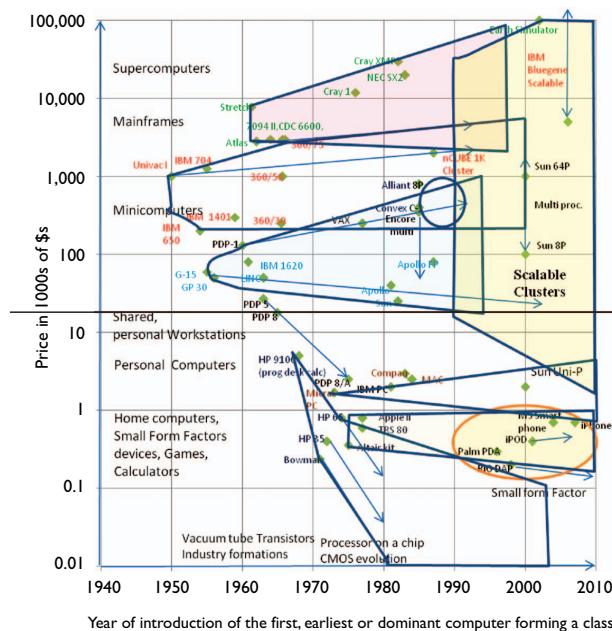


Figure 2b. Introduction price versus date of the first or early platforms to establish a computer class or lower-priced sub-class originating from the same company or industry.

establish a single line based on CMOS technology. The section is followed by the three direct and indirect effects of Moore's Law to determine classes:

- Microprocessor transistor/chip evolution circa 1971–1985 establish: calculators, home computers, personal computers, and workstations, and lower (than minicomputer) priced computers.
- “Minimal” designs establish new classes circa 1990 that use a “fraction” of the Moore number.

Microsystems evolution using fractional Moore's Law-sized SOCs enable small, lower-performing, minimal PC and communication systems including PDAs, PADs, cameras, and cell phones.

- Rapidly evolving microprocessors using CMOS and a simpler RISC architecture appear as the “killer micro” circa 1985 to have the same performance as supercomputers, mainframes, mini-supercomputers, super-minicomputers, and minicomputers built from slowly evolving, low-density, custom ECL and bipolar integrated circuits. ECL survived in supercomputers the longest because of its speed and ability to drive the long transmission lines, inherent in large systems. In the end, CMOS density and faster system clocks overtook ECL by 1990.

The “killer micro” enabled by fast floating-point arithmetic subsumed workstations and minicomputers especially when combined to form the “multi” or multiple microprocessor shared memory computer circa 1985. “Multis” became the component for scalable clusters when interconnected by high-speed, low-latency networks. Clusters allow arbitrarily large computers that are limited only by customer budgets. Thus scalability allows every computer structure from a few thousand dollars to several hundred million dollars to be arranged into clusters built from the same components.

In the same fashion that killer micros subsumed all the computer classes by combining, it can be speculated that much higher volume—on the order of hundreds of millions—of SFF devices, may evolve more rapidly to subsume a large percentage of personal computing. Finally, tens of billions of dust-sized, embeddable wirelessly connected platforms that connect everything are likely to be the largest class of all enabling the state of everything to be sensed, effected, and communicated with.

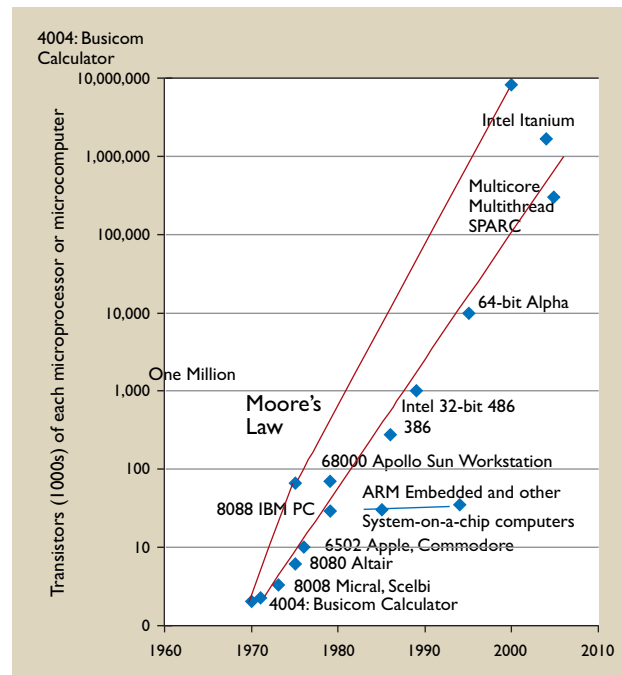


Figure 3. Moore's Law, which provides more transistors per chip per year, has resulted in creating the following computer classes: calculators, home computers, personal computers, workstations, “multis” to overtake minicomputers, and clusters using multiple core, multithreading to overtake mainframes and supercomputers.

MICROPROCESSORS CIRCA 1971: THE EVOLVING FORCE FOR CLASSES IN THE SECOND PERIOD

Figure 3 shows the microprocessors derived directly from the growth of transistors/chips beginning in 1971. It shows the trajectory of microprocessors from a 4-bit data path through, 8-, 16-, 32-, and 64-bit data paths and address sizes. The figure shows a second path—the establishment of “minimal” computers that use less than 50 thousand transistors for the processor, leaving the remainder

of the chip for memory and other functions (for example, radio, sensors, analog I/O) enabling the complete SOC. Increased performance, not shown in the figure, is a third aspect of Moore's Law that allows the “killer micro” formation to subsume all the other, high-performance classes that used more slowly evolving bipolar TTL and ECL ICs. Calculators, home computers, personal computers, and workstations were established as classes as the processor on a chip evolved to have more transistors with wide data paths and large address spaces as shown in Figure 3.

In 1971, Intel's 4004 with a 4-bit data path and ability to address 4KB was developed and programmed to be the Basicom Calculator; instead of developing a special chip as had been customary to implement calculators, a program was written for the 4004 for it to “behave” as or “emulate” a calculator.

In 1972, Intel introduced the 8008 microprocessor coming from the Datapoint terminal requirement, with a 8-bit data path and ability to access 16KB that allowed limited, programmable computers followed by more powerful 8080-based systems MITS used to introduce its Altair personal computer kit in 1975,

Over time, the performance of lesser-performing, faster-evolving products eventually overtakes the established, slowly evolving classes served by sustaining technology.

which incidentally stimulated Gates and Allen to start Microsoft. In 1977, the 16-bit 6502 microprocessor and higher-capacity memory chips enabled personal computers for use in the home or classroom built by Apple, Commodore, and Radio Shack—computers that sold in the tens of millions because people bought them to use at home versus corporate buyers. By 1979, the VisiCalc spreadsheet ran on the Apple II establishing it as a “killer app” for personal computers in a work environment. Thus, the trajectory went from a 4-bit data path and limited address space to a 16-bit data path with the ability to access 64KB of memory. This also demonstrates the importance of physical address as an architectural limit. In the paper on DEC’s VAX [3], we described the importance of address size on architecture: “There is only one mistake that can be made in a computer design that is difficult to recover from—not providing enough address bits for memory addressing and memory management...” The 8086/8088 of the first IBM PCs had a 20-bit, or 1MB address space, the operating system using the remaining 384KB.

Concurrent with the introduction of the IBM PC, professional workstations were being created that used the Motorola 68000 CPU with its 32-bit data and address paths (4GB of maximum possible memory). Apple Computer used the Motorola “68K” in its Lisa and Macintosh machines. IBM’s decision to use the Intel architecture with limited addressing, undoubtedly had the effect of impeding the PC by a decade as the industry waited for Intel to evolve architecture to support a larger address and virtual memory space. Hundreds of companies started up to build personal computers (“PC clones”) based on the IBM PC reference design circa 1981. Dozens of

companies also started to build workstations based on a 68K CPU running the UNIX operating system. This was the era of “JAWS” (Just Another WorkStation) to describe efforts at Apollo, HP, IBM, SGI, Sun Microsystems and others based on 32-bit versus 16-bit. Virtually all of these “workstations” were eliminated by simple economics as the PC—based on massive economies of scale and commoditization of both the operating system and all constituent hardware elements—evolved to have sufficient power and pixels.

“Minimal” CMOS Microsystems on a Chip circa 1990 Establish New Classes using Smaller, Less-Expensive, Chips. In 2007, many systems are composed of microprocessor components or “cores” with less than 50,000 transistors per microprocessor core at a time when the leading-edge microprocessor chips have a billion or more transistors (see Figure 3). Such cores using lower cost, less than the state-of-the-art chips and highly effective, rapid design tools allow new, minimal classes to emerge. PDAs, cameras, cell phones, and PADs have all been established using this minimal computer design style based on small cores. In 1990, the Advanced RISC Machine (ARM) formed from a collaboration between Acorn and Apple as the basis for embedded systems that are used as computing platforms and achieved two billion units per year in 2006. Other higher-volume microsystem platforms using 4-, 8-...64-bit architectures including MIPS exist as core architectures for building such systems as part of the very large embedded market.

Rapidly Evolving Killer CMOS Micros circa 1985 Overtake Bipolar ICs to Eliminate Established Classes. In the early 1980s, the phrase “killer micro” was introduced by members of the technical comput-

ing community as they saw how the more rapidly evolving CMOS micro would overtake bipolar-based minicomputers, mainframes, and supercomputers if they could be harnessed to operate as a single system and operate on a single program or workload.

In the *Innovator's Dilemma*, Christensen describes the death aspect basis of Bell's Law by contrasting two kinds of technologies [4]. Sustaining technology provides increasing performance, enabling improved products at the same price as previous models using slowly evolving technology; disruptive, rapidly evolving technology provides lower-priced products that are non-competitive with higher-priced sustaining class to create a unique market space. Over time, the performance of lesser-performing, faster-evolving products eventually overtakes the established, slowly evolving classes served by sustaining technology.

From the mid-1980s until 2000, over 40 companies were established and went out of business attempting to exploit the rapidly evolving CMOS microprocessors by interconnecting them in various ways. Cray, HP, IBM, SGI, and Sun Microsystems remain in 2008 to exploit massive parallelism through running a single program on a large number of computing nodes.

Two potentially disruptive technologies for new classes include:

- The evolving SFF devices such as cell phones are likely to have the greatest impact on personal computing, effectively creating a class. For perhaps most of the four billion non-PC users, a SFF device becomes their personal computer and communicator, wallet, or map, since the most common and often only use of PCs is for email and Web browsing—both stateless applications.
- The One Laptop Per Child project aimed at a \$100 PC (actual cost \$188 circa November 2007) is possibly disruptive as a “minimal” PC platform with just a factor-of-two cost reduction. This is achieved by substituting 1G of flash memory for rotating-disk-based storage, having a reduced screen size, a small main memory, and built-in mesh networking to reduce infrastructure cost, relying on the Internet for storage. An initial selling price of \$188 for the OLPC XO-1 model—approximately half the price of the least-expensive PCs in 2008—is characteristic of a new sub-class. OLPC will be an interesting development since Microsoft's Vista requires almost an order of magnitude more system resources.

FUTURE CHALLENGES

The Challenge of Constant Price, 10–100 billion Transistors per Chip, for General-Purpose Computing. The future is not at all clear how such large,

leading-edge chips will be used in general-purpose computers. The resilient and creative supercomputing and large-scale service center communities will exploit the largest multiple-core, multithreaded chips. There seems to be no upper bound these systems can utilize. However, without high-volume manufacturing, the virtuous cycle is stopped—in order to get the cost and benefit for clusters, a high-volume personal computer market must drive demand to reduce cost. In 2007, the degree of parallelism for personal computing in current desktop systems such as Linux and Vista is nil, which either indicates the impossibility of the task or the inadequacy of our creativity.

Several approaches for very large transistor count (approximately 10 billion transistor chips) could be:

- System with primary memory on a chip for reduced substantially lower-priced systems and greater demands;
- Graphics processing, currently handled by specialized chips, is perhaps the only well-defined application that is clearly able to exploit or absorb unlimited parallelism in a scalable fashion for the most expensive PCs (such as for gaming and graphical design);
- Multiple-core and multithreaded processor evolution for large systems;
- FPGAs that are programmed using inherently parallel hardware design languages like parallel C or Verilog that could provide universality that we have not previously seen; and
- Interconnected computers treated as software objects, requiring new application architectures.

Independent of how the chips are programmed, the biggest question is whether the high-volume PC market can exploit anything other than the first path in the preceding list. Consider the Carver Mead 11-year rule—the time from discovery and demonstration until use. Perhaps the introduction of a few transactional memory systems has started the clock using a programming methodology that claims to be more easily understood. A simpler methodology that can yield reliable designs by more programmers is essential in order to utilize these multiprocessor chips.

Will SFF Devices Impact Personal Computing?

Users are likely to switch classes when the performance and functionality of a lesser-priced class is able to satisfy their needs and still increase functionality. Since the majority of PC use is for communication and Web access, evolving a SFF device as a single communicator for voice, email, and Web access is quite natural. Two things will happen to accelerate the development of the class: people who have never

used or are without PCs will use the smaller, simpler devices and avoid the PC's complexity; and existing PC users will adopt them for simplicity, mobility, and functionality. We clearly see these small personal devices with annual volumes of several hundred million units becoming the single universal device evolving from the phone, PDA, camera, personal audio/video device, Web browser, GPS and map, wallet, personal identification, and surrogate memory.

With every TV becoming a computer display, a coupled SFF becomes the personal computer for the remaining applications requiring large screens. Cable companies will also provide access via this channel as TV is delivered digitally.

Ubiquitous Wireless: WiFi, Cellular Services, and Wireless Sensor Nets. Unwiring the connection around the computer and peripherals, televisions, and other devices by high-speed radio links is useful but the function is "unwiring," and not platform creation. Near-Field Communication (NFC) using RF or magnetic coupling offers a new interface that can be used to communicate a person's identity that could form a new class for wallets and identity. However, most likely the communication channel and biometric technology taken together just increase the functionality of small devices.

Wireless Sensor Nets: New Platform, Network, and Applications. Combining the platform, wireless network, and interface into one to integrate with other systems by sensing and effecting is clearly a new class that has been forming since 2002 with a number of new companies that are offering unwiring, and hence reduced cost for existing applications, such as process, building, home automation, and control. Standards surrounding the 802.15.4 link that competes in the existing unlicensed RF bands with 802.11xyz, Bluetooth, and phone transmission are being established.

New applications will be needed for wireless sensor nets to become a true class versus just unwiring the world. If, for example, these chips become part of everything that needs to communicate in the whole IT hierarchy, a class will be established. They carry out three functions when part of a fixed environment or a moving object: sense/effect; recording of the state of a person or object (things such as scales, appliances, switches, thermometers, and thermostats) including its location and physical characteristics; and communication to the WiFi or other special infrastructure network for reporting. RFID is part of this potentially very large class of trillions. Just as billions of clients needed millions of servers, a trillion dust-sized wireless sensing devices will be coupled to a billion other computers.

CONCLUSION

Bell's Law explains the history of the computing industry based on the properties of computer classes and their determinants. This article has posited a general theory for the creation, evolution, and death of various priced-based computer classes that have come about through circuit and semiconductor technology evolution from 1951. The exponential transistor density increases forecast by Moore's Law [6] being the principle basis for the rise, dominance, and death of computer classes after the 1971 microprocessor introduction. Classes evolve along three paths: constant price and increasing performance of an established class; supercomputers—a race to build the largest computer of the day; and novel, lower-priced "minimal computers." A class can be subsumed by a more rapidly evolving, powerful, less-expensive class given an interface and functionality. In 2010, the powerful microprocessor will be the basis for nearly all classes from personal computers and servers costing a few thousand dollars to scalable servers costing a few hundred million dollars. Coming rapidly are billions of cell phones for personal computing and the tens of billions of wireless sensor networks to unwire and interconnect everything. As I stated at the outset, in the 1950s a person could walk inside a computer and by 2010 a computer cluster with millions of processors will have expanded to the size of a building. Perhaps more significantly, computers are beginning to "walk" inside of us. ■

REFERENCES

1. Bell, C.G. The mini and micro industries. *Computer* 17, 10 (Oct. 1984), 14–30.
2. Bell, C.G. Multis: A new class of multiprocessor computers. *Science* 228 (Apr. 26, 1985) 462–467.
3. Bell, G. and Strecker, W. Computer structures: What have we learned from the PDP-11. *IEEE Computer Conference Proceedings* (Florida, Nov. 1975).
4. Christensen, C.M. *The Innovator's Dilemma*. Harvard Business School Press, 1997.
5. Gray, J. and Shenoy, P. Rules of thumb in data engineering. In *Proceedings of ICDE2000* (San Diego, Mar. 1–4, 2000). IEEE press.
6. Moore, G.E. Cramming more components onto integrated circuits. *Electronics* 8, 39 (Apr. 19, 1965); revised 1975.
7. Nelson, D.L. and Bell, C.G. The evolution of workstations. *IEEE Circuits and Devices Magazine* (July 1986), 12–15.

GORDON BELL (gbell@microsoft.com) is a principal researcher in Microsoft Research Silicon Valley, working in the San Francisco Laboratory.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.