

An abridged version of this article appears in: Buxton, W. (2001). Less is More (More or Less), in P. Denning (Ed.), *The Invisible Future: The seamless integration of technology in everyday life*. New York: McGraw Hill, 145 - 179. Some minor corrections and improvements of wording have also been made, without changing meaning.

Less is More (More or Less)

William Buxton
Buxton Design
Toronto, Ontario
bill@billbuxton.com
<http://www.billbuxton.com>

*Die Zukunft war früher auch besser.
(The future was better in the past, as well.)*

Karl Valentin

Abstract

In April 1981 Xerox introduced the Star 8010 workstation, the first commercial system with a *Graphical User Interface* (GUI) and the first to use the “desktop” metaphor to organize a user’s interactions with the computer. Despite the perception of huge progress, from the perspective of design and usage models, there has been precious little progress in the intervening years. In the tradition of Rip van Winkle, a Macintosh user who awoke today, after having fallen asleep in 1984, would have no more trouble operating a “modern” PC than operating a modern car.

The “desktop” is still the dominant user interface paradigm. Equally durable is the “general purpose” nature of PC design, which assumes that we channel all our transactions (a diverse lot) through a single interface on a single computer.

While common discourse about digital media is dominated by the concept of *convergence*, we argue that from the perspective of the usage model, just the opposite concept, *divergence*, should be the dominant model. We argue that the diversity of web browsers tomorrow will match the diversity of “ink browsers” (a.k.a. paper) today.

Systems will have to be tailored to dynamically connect the user with artefacts relevant to the user’s current actions -- and do so in a way, form, place, and cost appropriate to the user. This change is as inevitable as it is essential. In the end, it will leave us with new concepts of computers, communications, and computing -- and of computer science itself. Few computer science departments or computer firms appreciate the magnitude of the impending change. It is time for them to wake up.

Introduction: Rushing Slowly into the Future

As someone in the profession of designing computer systems, I have to confess to being torn between two conflicting sentiments concerning technology. One is a sense of excitement

about its potential benefits and what might be. The other is a sense of disappointment, bordering on embarrassment, at the state of what is.

Despite the hyperbole surrounding new media and technology, I firmly believe that we are far behind where we might have otherwise been, and consequently, our society is all the poorer. Furthermore, my view is that our current path has very little hope of getting us to where we should and could be anytime soon, thereby prolonging the missed potential. For this to alter, our approach must change.

Despite the increasing reliance on technology in our society, in my view, the key to designing a different future is to focus less on technology and engineering, and far more on the humanities and the design arts. This is not a paradox. Technology certainly is a catalyst and will play an important role in what is to come. However, the deep issues holding back progress are more social and behavioural than technological. The skills of the engineer alone are simply not adequate to anticipate, much less address the relevant issues facing us today. Hence, fields such as sociology, anthropology, psychology, and industrial design, must be *at least* equal partners with engineering and technology in framing how we think about, design and manage our future.

While the growth of technology is certain, the inevitability of any specific "future" is not. Like mathematics, perhaps we should start to use the word "future" in the plural, *futures*, to reinforce the fact that there are a multitude of different possible futures. The specific future that we build, therefore, will be more easily seen to be a consequence of our own decisions and therefore demand more attention to its design.

What follows is an attempt to establish a conceptual framework from which we can better understand the past and make more informed decisions about the future.

Dimensions of Change

I have begun by expressing disappointment at the state of the art, and at the slow rate of progress, in human terms, in these emerging information technologies and "new media."



Figure 1: *The Xerox Star 8010 Workstation. Introduced in 1981. This is the first commercial system to utilise a "windows, icon, menus, pointer" (WIMP), or "graphical user interface" (GUI).*

But given the general perception that technology is changing at such a breakneck speed that even experts have trouble keeping up, how valid is the suggestion that our progress *is* too slow, and that we could have done better? It all depends on the dimension along which we measure change. What is the relevant metric?

In the areas of microelectronics, telecommunications, and materials science, for example, there is no question that there has been staggering change over the past few decades. But if we shift from the dimension of technology to the dimension of the user, we see something very different. Despite all of the technological changes, I would argue that there has been no significant progress in the conceptual design of the personal computer since 1981. To support this claim, look at the computer shown in the photograph in Figure 1, which dates from that year. My experience is that most computer users, including professionals, cannot identify the *decade*, much less the year, in which the photograph was taken! For how many other "fast changing" products is that true?

The computer shown is a Xerox Star 8010 workstation (Smith, Irby, Kimball, Verplank & Harslem, 1983). It incorporated all the design and user interface characteristics of a contemporary personal computer: windows, icons, a mouse, and CRT¹. In fact, there is an argument to be made that this 1981 machine was better designed from an ease-of-use perspective than most "modern" computers (so rather than progress, we may even have gone backwards in the intervening years!).²



Figure 2: *The iMac from Apple Computer. The iMac typifies both the best and the worst in current computer design. On the one hand, it illustrates how departure from the status quo and investment in design can have a strong impact on the success of a product. On the other hand, it illustrates how superficial our investment in design has been. If you look past the candy colours and the sleek curves, what you see on the screen is the essentially the same old GUI and conceptual model that was there on the Xerox Star in 1981.*

Now I have the greatest respect for the innovators that made this machine possible. But I must ask, "Did they get it so right that no further design or refinement was required?" I think not. What I feel is missing is the next wave of innovation - innovation that does to the Xerox Star what the Xerox Star did to its predecessors. This is something that I believe we have been capable of, yet have failed to do, for several years. This is also something that I feel has to be done before we can achieve the benefits that are so often offered, but so infrequently delivered, by this emerging technology.³

¹ The system also incorporated email, a local area network and a networked laser printer.

² It is significant that the launch of the Xerox Star was in 1981. This preceded by a year the first *Conference on Human Factors in Computer Systems* which took place in Gaithersburg, Maryland. This meeting became the start of the ACM Special Interest Group on Human-Computer Interaction (SIGCHI), and the establishment of HCI as a distinct discipline. Observe that the landmark innovations of the Xerox Star were, therefore, made without the benefit of any significant literature in the field of HCI. And yet, despite a huge body of literature being generated in the intervening years, none of it has had any impact on the design of personal computers, relative to the impact of the Star. What does that tell us about the discipline and the industry? For me, it is cause for serious concern.

³ The Xerox Star is a data point supporting what I call "The Law of Inertia of Good Ideas." This says: *The better an idea in its time, the more it holds back progress in the future.* It is relatively easy to displace bad ideas and bad design. Really good ideas, such as the QWERTY keyboard, or the GUI, however, take hold and are extremely difficult to replace.

One of the motivations for this essay is to put forward a view on how we can bring the progress of the design and benefits of computational devices more in line with the progress of the underlying technologies and its unfulfilled potential. To accomplish this, we need to delve a little deeper into the nature of the changes that *have* been taking place.

A little practical fieldwork will help us here. The exercise is this: ask 10 people what they think the most significant changes have been in computers over the past 15 years. If the list that you thus obtain is like mine, it will look something like:

1. *Smaller*: computers are much smaller than they were, making them portable, among other things
2. *Faster*: we can do things on small machines that used to take huge machines
3. *Cheaper*: the cost of computation is falling dramatically
4. *More of them*: the number of computers has exploded. The population of microprocessors in the world is now about three times the population of humans.
5. *Networked*: our machines can communicate with one another, with faster speed and increasingly using wireless means
6. *Location/Motion Sensing*: our devices are starting to have the capacity to know where they are, both geographically (such as the GPS-equipped navigation computer in some cars, for example), and "socially". Social awareness comes in two forms: technological and human. Devices are developing the capacity to have an awareness of what other devices are in (what Microsoft calls) the "Network Neighbourhood" (what I call the "society of appliances"), and in terms of the human social context.
7. *Input/Output (I/O) Transducers*: the input output devices available are changing dramatically, offering a range of options from printers, scanners, etc. to the user, thereby opening the opportunity to redefine the nature of what constitutes a "computer terminal".

When I have done this exercise, I typically get the same basic results whether it is a layperson or technology professional that I poll. The greatest consistency is with the first three or four items. Further down the list the responses are fuzzier and less consistent.

This is significant, since I would argue that the items are listed *in inverse order of importance*. The things that come first to mind are the least important, and the things that come last to mind and are most vague, are the most important. I see this discrepancy between *consciousness* and *importance* as pointing to the root of the stagnation in computer design.

Am I suggesting that the improvements of size, speed and cost of microelectronics are not important? No. Rather, my argument is that there are so many resources allocated to solving the underlying problems along these dimensions that the improvements will happen regardless of what you or I do. They have momentum, and verge on inevitable. What is not inevitable, at least in the short term, are some of the opportunities that they afford when coupled with the things at the bottom of the list - things that do not have adequate resources or attention being paid to them.

This brings us to the trap inherent in the above list of changes.

The models and language that we use to articulate, or discuss, things, frame our perceptions and ways of thinking. If we discuss design in terms of this list, our perspective, like the list itself, will have a *technocentric* bias. To break out of our current rut, we need to recast our list of changes using a *human centric* perspective, which reflects the importance of usage and activity rather than technology:

- *Who* is using the computer?
- *What* are they doing?

- *Where* are they doing it?
- *When* can they do it?
- *Why* are they doing it?
- *How* they do it?

These are the questions that matter most and can guide design towards the *right solution* in the *right form* for the *right person* in the *right location* at the *right time* and at the *right cost*. They prompt a concern for design that reflects respect for human skill at all three levels: motor-sensory, cognitive, and social (Buxton, 1994).

Bridging the Two Solitudes of the Physical and the Virtual



Figure 3: *Bridging between the two solitudes of the physical and the virtual domains. (Image courtesy of Gray Holland)*

In the previous section, I argued that changes in input and output (I/O) technologies constituted perhaps the most important dimension of change in terms of defining the nature of future technologies. Why is that?

One of the most significant issues confronting computer users, as illustrated Figure 3, is the problem of bridging the gap between the physical and virtual worlds. For most activities, most current systems make it too difficult to move the artefacts back and forth between these two worlds, the physical and virtual. Hence, the relevant documents, designs, etc. are isolated in one or the other, or split between the two.

With appropriate design, however, the I/O technologies of future systems will be designed so as to *absorb* and *desorb* the artefacts relevant to the intended activity, thereby providing a much more seamless bridge between these two solitudes.

Tuning the I/O for specific activities is contrary to most current design, which more follows what might be called, *The Henry Ford School of Design*. Just as Ford is reputed to have said about his automobiles, "You can have it in any colour you want as long as it is black," so today's computer designers say, "You can have it in any form you want as long as it has a keyboard, a display and a mouse, regardless of what you are doing." Hence, temporarily

assuming the role of an anthropologist examining the tools of a society, we notice that there is little significant difference in the tools used by the physician, accountant, photographer, secretary, stockbroker architect, or ticket agent. Instead of specialisation, a *one size fits all* approach to design prevails.

As we shall see going forward, I/O devices are key to our ability to tailor computer systems to the specific needs of various users. Hence, their place at the top of my list in importance, at least insofar as technology is concerned. But technology is only of secondary importance. It is the human's capability, intent and need that should be (but too often is not) the driving function in all of this.

On Complexity, Skill, and Human Limitations

If the human should be the centre of focus in the design of technology, how can we get a better understanding of the relevant, but too often neglected issues, especially as they relate to complexity and design?

Let us begin with a series of examples that I have used over the past few years, beginning with the graph shown in Figure 4. This is an approximation of the often-cited *Moore's Law*, which states that the number of transistors that can fit on a chip will double every 18 months. The graph simplifies this to simply state that there will be more technology tomorrow than there is today. So far, so good.




Figure 4: *Moore's Law: the growth of technology as a function of time. The simple interpretation is that there will be more technology tomorrow than there is today.*

The next graph in the sequence is shown in Figure 5. It illustrates what I immodestly will call *Buxton's Law of Promised Functionality*, which states that the functionality promised by technology will grow proportionally with Moore's Law. In layperson's terms, this simply means there is going to be more functionality promised/offered tomorrow, than there is today.

At this point, readers will be excused if they are wondering why I am wasting time and space stating the seemingly obvious, and using quasi-scientific graphs in the process?




Figure 5: *Buxton's Law: Promised functionality and benefits of technology will increase directly with Moore's Law. The simple interpretation is simply that there will be more functionality promised/offered tomorrow than there is today.*

Many may see the previous two graphs as obvious and banal. But to my way of thinking, they are less obvious than the next graph that we will look at, which nevertheless, seems lost on most computer designers. This graph, shown in Figure 6 illustrates the "growth" of human capacity over time. This represents what I call *God's Law*, which can be expressed as follows: the capacity of human beings is limited and does not increase over time. Stated more simply, the law says that our neurons do not fire any faster, our memory doesn't increase in capacity, and we do not learn or think faster as time progresses. If there is a problem with the graph, it is that it is too generous, since my personal experience is that, if anything, my capacity is decreasing, rather than increasing.




Figure 6: *God's Law: Humans' capacity is limited and does not increase over time. Stated more simply, our neurons do not fire faster, our memory doesn't increase in capacity, and we do not learn or think faster as time progresses.*

"God's Law" has been obviously true much longer than either of the previous two laws discussed. Nevertheless, while virtually every engineer can quote Moore's Law, their designs show little, if any, understanding of God's Law. Yet which is the more important law in terms of good engineering? Which one should be the primary concern in designing technology intended for use by human beings?




Figure 7: *Relating functionality to human capability. The observation to make here is that access to the functionality today is already at the threshold of human capability.*

Figure 7 attempts to illustrate the root of our current design problems by relating the growth of functionality to human capability. As a good scientist, I have normalised my data relative to the Y-axes. If I take the Y intercept of the Functionality graph (that is, where we are today) and, using the dotted line, project it across over to the Capability graph, we start to see some interesting relationships. Most significantly, we see that we are already right at the threshold of human capability, a threshold that we are about to cross (if we have not already done so.)

Using what in scientific visualisation is called "a level of detail zoom", we can look at this a little closer. Figure 8 shows the growth of functionality with the limit of human performance superimposed. This limit we have labelled the *Threshold of Frustration*, but it is also known as the *Complexity Barrier*.

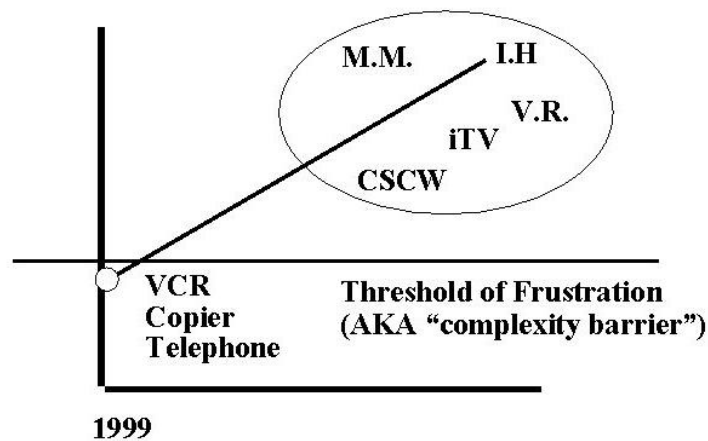


Figure 8: *The Threshold of Frustration as it relates to access to functionality.*

The first point that that I want to make with this figure is that any functionality that lies above the Threshold of Frustration does not exist in human terms. This is regardless of whether you can demonstrate that functionality to me, show me the code, or show me the documentation. The functionality may even be the very reason that the system was acquired. Nevertheless, statistically speaking, it simply does not exist in human terms. If it lies above the threshold, it will not be used.⁴

⁴ At this point I can already hear a chorus of voices protesting that *they* can use such functionality, or that they know someone who can. That does not disprove my point. Every transaction has a cost, regardless if it is performed on a computer or shopping. If the transaction cost is too high, the "business" will fail *even if some people can afford it*.



Figure 9: The computer section of a typical bookstore. What dominates the shelves would more properly be called "documentation" than books. The preponderance of such documentation, which has a half-life of about six months, is only one of the many symptoms of the overall malaise in design. (Photo: Tom Wujec)

Second, based on their track record, engineering and design practice have very little credibility in terms of being able to deliver the promised functionality (such as e-commerce, multimedia, interactive television, computer supported collaborative work, or virtual reality) below this threshold. As proof, just consider the extent to which documentation has replaced books on the shelves of our bookstores, as illustrated in Figure 9.⁵

Those of us in academia and industry are quick to talk about the importance of education and the training gap. I agree. However, more often than not what is meant is the need to teach people more about technology: to make them "computer literate." The existence of all of this documentation is a testament to a training gap, *but the gap is on the side of the engineer and computer scientist, the perpetrators of these systems, not on the part of the users* (who I

⁵ Something is "documentation" if it has a half-life of about 6 months. Books have content that persists beyond the limited lifespan of some product. This has prompted me to introduce the concept of a *Design Index*. This is a design quality indicator for the tools used in some application domain d . Ideally a layperson should be able to determine the index for a given application without having to suffer through it first hand. The bookstore provides the mechanism, as follows:

$$Design\ Index_d = \text{documentation shelf space}_d : \text{content shelf space}_d$$

In the domain of accounting, for example, you could compare how many books there are on how to use spreadsheets and other related programs, compared to the number of books on accounting theory or practice.

Now I recognise that this formula will be dismissed as absurd, and perhaps quite rightly. But what is more absurd, my formula, or the situation that prompted it in the first place, as reflected in the bookstore? The problem is real and requires a change in how we do things. On the one hand we have a skills shortage, and on the other, we are wasting a broad range of existing skills that users bring to the system.

would rather characterise as victims). To overcome this situation, the technological expertise of the computer scientists and engineers who make these systems must be matched or exceeded by their knowledge of people and their capabilities.

Currently that is nowhere near the case. To the best of my knowledge, there are virtually no computer science or engineering schools where to graduate you must develop an application which has been used by another human being, much less be marked on your ability to do so.⁶

Yes, it is hard to fit all the technological basics into the engineering and computer science curriculum. But the result is that the human factor gets squeezed out, which is unacceptable given the growing importance of technology-based tools in our society.

The problem is systemic and must be addressed. But the good news is, just think about the competitive advantage that will be gained by the first jurisdiction to address this situation with conviction!

“Super Appliances”: The Behavioural Implications of Artefacts

Most popular discussions seem to assume that the personal computer and the interactive television set⁷ will be the primary appliances through which the public will access the promised wonders of the digital revolution. Yet, if we are to take the preceding section seriously, perhaps this assumption is worth looking into a little more deeply.

Let us proceed by looking at three everyday appliances, since experience with them, coupled with a bit of common sense, sheds light on how we think about future information appliances.



Figure 10: *The Cuisinart Food Processor*

The first of our three examples is a food processor, the *Cuisinart*,⁸ illustrated in Figure 10. This is a kitchen appliance that performs several functions such as slicing, dicing, mixing, blending, etc. Each of these different functions is accomplished by reconfiguring the device using the various components, such as blades and containers that come with it. Thus, this single multi-purpose appliance, with its accessories, can perform the function of several simpler appliances, such as a mixer and a blender.

⁶ Considering this, perhaps we should stop complaining about how hard things are to use, how error prone they are, and how poorly they are designed. Rather, perhaps we should be amazed at how good they are under the circumstances!

⁷ That is, a television set connected to the Internet via a “set-top box.”

⁸ <http://www.cuisinart.com>



Figure 11: *The Shopsmith: a 5-in-1 Powertool. The unit is shown along with each of its five different ways of being used.*

Our second example could be considered a Cuisinart for the woodworking shop. It is a power tool for carpentry called *the Shopsmith*,⁹ and is illustrated in Figure 11. This is a multifunction appliance that, in a few minutes, can convert to any one of the following five functions:

- Table saw
- Disk sander
- Horizontal drill
- Drill press
- Lathe

The benefit is that one gets a suite of tools for the cost of slightly more than one, and the resulting appliance takes up little more space in the shop than a bicycle.

Our third example is illustrated in Figure 12. It is the common *Swiss Army Knife*¹⁰. By virtue of having a range of tools such as a nail file, bottle opener, cork screw, scissors and yes, a knife blade, embedded in it, the Swiss Army Knife joins the Cuisinart and the Shopsmith as a single device that can perform a wide range of functions. Nevertheless, one thing that distinguishes the Swiss Army Knife from the Cuisinart and the Shopsmith is that it is more portable. However, as the number of tools in the knife increases, so does its weight and bulk.

⁹ <http://www.shopsmith.com>

¹⁰ <http://www.airtime.co.uk/shop/SwissArmyKnives/> and <http://www.swissarmy.com>

The convenience, and thus the likelihood of carrying it around, are therefore reduced proportionally.



Figure 12: *The Swiss Army Knife*

From a design perspective, each of our examples, the Cuisinart, Shopsmith and Swiss Army Knife, are representatives of what I call “*super appliances*”. This class of appliance is important to understand since both the personal computer and the interactive television fall into the same category. Hence, we can potentially use past experience to inform future expectations.

Collectively, appliances of this class all share the following properties:

- *Multiple Function:* The appliance is a single device that encompasses the functionality otherwise delivered through a number of simpler special purpose devices.
- *Single Location:* The functionality offered by the appliance can only be accessed where the appliance is located. The higher the overhead or inconvenience in moving the appliance, the more constraining this is.
- *Single User:* Only one person can use the appliance at time.¹¹
- *Single Function at a Time:* While offering multiple functions, these functions are time multiplexed, meaning that only one function at a time can be used.
- *Space/Complexity Trade-off:* These multifunction appliances are inherently more complex than any of the single-function tools that they overlap with. This is largely due to the overhead in switching the modality from one function to the other. On the other hand, they

¹¹ I acknowledge that there are exceptions, such as interactive video games, played either on a PC or TV game consol. Another obvious exception would be watching television or something similar on a TV equipped with a set-top box. But consider the battles that rage nightly in homes around the world over who has the TV remote control. Now imagine this amplified ten times over as the choice expands beyond selecting TV channels, to selecting web pages. Finally consider reading your personal email on the television, in full view of the rest of the family. I think that even the most rudimentary analysis, much less common sense, suggests that for the most part these are services tailored for individual, not communal interaction.

typically occupy less space, and cost less, than the equivalent set of single-function tools.

Sometimes the benefits of these tools warrant the inherent increase in their complexity. If this were not true, the companies that make them would not have survived as long as they have. But on the other hand, everyday experience says that this class of tool represents only a small part of our tool usage. Consequently, I think that it is fair to ask the following question:

If you don't do the bulk of your food preparation using a Cuisinart, you don't do the bulk of your carpentry with a Shopsmith, and you don't do the bulk of your eating, can opening, cork screwing, etc. with a Swiss Army Knife, then why would you find the same type of appliance acceptable as the primary tool that you use in your work or recreation?

Some may argue that the new digital super appliances are a new breed and will foster a different set of behaviours. It would be foolish to argue that this is not partially true. Having said that, in my opinion, it is equally foolish to assume that our experience and expectations with tools will not transfer to the next generation of appliances. Likewise, it would be short sighted to assume that there are not viable design alternatives to the digital super appliance.

A look into some of the implications of the *Single Location* attribute of super appliances can shed some insight onto why we need to explore alternative approaches to design.

One of the foundations of the discipline of architecture is to design each physical space to best accommodate the specific activities and tasks which it is intended to support within. Each room, itself, is a key component influencing how well (or poorly) such activities and tasks are performed. This notion is nicely captured in the following quote from the architect Louis I. Kahn:

Thoughts exchanged by one and another are not the same in one room as in another.

My view is that the relationship between function and space is just as important to the designer of digital tools as to an architect. We choose our tools according to the task, just as we do so with rooms (which are "just" another level of tool. A sense of the tool-location-activity relationship can be demonstrated through the following exercise.

Make a list of the tools in your Swiss Army Knife. Then opposite each, list the room most associated with the function that that tool most normally serves. An example of such a list is shown in Table 1.¹²

¹² Not only is there a fairly distinct location associated with each function, with conventional specialised tools, each activity can take place independent from, and simultaneously with, any other activity, and as many people can work concurrently as there are tools.

Tool	Location
Saw	Workshop
Spoon	Kitchen
Fork	Kitchen
Scissors	Sewing Room
Leather Punch	Stable
Nail File	Bathroom
Corkscrew	Dining Room

Table 1: *Function/Location Relationships with a sample Swiss Army Knife*

The implications of this become clearer and more immediate if you do the same exercise for your family PC. The home computer promises that it will let your children do their homework, write or practice music, and play games. It claims to let you do your accounting and correspondence, another person watch a movie, and someone else plan a vacation. But just like the Swiss Army Knife, there are specific places in the house (the bedroom, music room, den, living room, games room, etc.) associated with each of these activities. Funnelling all of these activities through a single appliance in a single location is, again, inconsistent with these task-location associations.

One could argue that this is a temporary problem, and that computers (or interactive televisions) will soon be inexpensive enough that one could simply have several of them distributed around the house, one in each location associated with a particular task. On the one hand I agree, as will be seen later. But I do so only if there are significant changes to what constitutes a computer or TV — changes that are not implicit when most people make this argument.

The essence of these changes is the notion that once we associate a particular activity with a tool and place it in a location appropriate for that activity, then we can also tailor the design of the tool so that it is optimised for that specific purpose. In other words, we can break out of the one-size-fits-all approach of the super appliance and evolve towards purpose-built tools.

In the process, we will discover that the three rules of computer design are the same as those of real estate: location, location, and location, which we interpret as:¹³

1. Location: of the appliance, i.e., physical context
2. Location: relative to user(s), i.e. social context
3. Location: relative to other appliances, i.e., context in the society of appliances

In biological systems, there is a tendency for specialised organisms to win out over generalised ones. My argument is that the evolution of technology will likely be no different. Rather than *converging* towards ever more complex multifunction tools, my claim is that going forward we must *diverge* towards a set of simpler more specialised tools. Underlying this is what I call my law of the inverse relationship between usefulness and functionality, expressed as:

$$\text{Usefulness} \sim 1 / \text{Functionality}^n$$

¹³ The influence of tools and location go far deeper than this in terms of how they impact human performance and cognition. See Hutchins (1990 & 1996) for an excellent and penetrating treatment.

Where n relates to the number of functions of the device, such as the number of tools in the Swiss Army Knife. These consequent notions of divergence and specialisation are themes that we will explore in more detail in following sections.

Plumbing, the Internet and the Waternet

In the previous section we used everyday appliances to highlight some of the limitations of personal computers and televisions as the primary means to access digital media and functionality. Luckily, there are also everyday appliances that can help lead us to a better alternative. These are the appliances that attach to another ubiquitous network, the *Waternet*.



Figure 13: Three "appliances" on the "Waternet". Despite standardization and "convergence" in the underlying network, note the diversity of the devices attached to it, despite only looking at the limited subset of sinks.

The Waternet is just what you think it is: that great network of pipes and reservoirs that brings water to and from your house or office, run by a relatively small number of very large companies. It is a lot like the Internet. Like with our computers, we didn't start connected. We had our own wells or cisterns, and we used things like septic tanks for our waste. Just as people have connected their PCs to the Internet via their local Internet Service Provider (ISP), so have most of us moved away from wells, cisterns and septic tanks and taken our plumbing online by hooking up with our local WSP (Waternet Service Provider).

Going forward, I think that there will be some other similarities. For example, the Waternet is essentially invisible. You only notice it when it breaks, and what is nice about it is that it does not break very often. Furthermore, many of us (including me) do not know who their WSP is, especially at my place of work. As it matures, the same will be true of the Internet. It will be the services and content, not the network that will be visible. As it is with copper, PVC,

or other plumbing materials, so will it be with the Internet: as long as it works, nobody will care if it is delivered by coaxial cable, twisted pair telephone wire or fibre optic cable. Nobody will care if it is delivered by the cable TV or the telephone company. It will just have to work.

With the Internet in the future, as with the Waternet today, what *will* be noticed are the appliances that are attached. Here is where we see the difference in the maturity of the two networks today. With the Internet, there is very little differentiation among these. As we have already discussed, there is little significant difference among personal computers that hang on the Internet. On the more mature Waternet, this is not the case. Attached to it are diverse appliances such as sinks, lawn sprinklers, toilets, baths, showers, fire hydrants, etc.

Even within these categories, such as with the sinks shown in Figure 13, there is sufficient diversity that one can easily distinguish which belongs in which context. We even know (due to its ornate nature) that the hand sink shown is intended for guests and would be found in the downstairs bathroom rather than intended for everyday use such as one that would be found in the upstairs bathroom.

The diversity and fine granularity of differentiation in these Waternet appliances hints at the diversity of what we will see in future Internet appliances. Also, note that the companies that make these appliances are not those that own or manage the network. They are specialized companies whose products and competence lie in niche areas.

The Pilot: Success Can Be In Hand

There is a significant implication to this last point. The industry today is dominated by large companies who, like the WSPs, exercise near monopolistic control. This is at the same time many jurisdictions, investors, and small companies are attempting to understand how they can benefit economically through the ongoing “information technology revolution.” The question frequently asked is, “How can we get a foothold in the face of the domination of the established large companies?”

The key to the answer lies in what we have just been discussing. A key area of future growth will be in the “terminals” or “appliances”. Fundamental to success in this class of technology will be one’s expertise and insights in the application domain, not just engineering. More to the point, and as supported by the precedent of the Waternet, no quasi monopoly can possibly have the depth of knowledge of the full gamut of application domains. Hence, there remains a healthy and growing opportunity for those who focus on the human/application centred approach that we are discussing.

A hint at the scale of the opportunity can be seen in the example of Palm Computing, and the success of the Palm Pilot (an example of which is shown in Figure 14.) The decade preceding the introduction of the Pilot was littered with the corpses of companies that had tried to introduce a pen-based personal digital assistant (PDA).

There is a compelling argument that the reason that the Palm Pilot succeeded when its predecessors had failed, is that the designers specified the product in human, rather than technological terms. For example, these included the need to fit into a jacket pocket, to be able to find an address faster than one could in a traditional address book, and to find “When did I had dinner with you?”, or “Am I free next Thursday?” faster than one could with a traditional date book. They also included the specification that one should be able to back-up the contents of the device in one button push, so that if the device was lost, one would only lose the cost of the device, not the information.



Figure 14: The Palm VII Handheld. The Palm Pilot is an example of “less-is-more” design. Despite the failure of previous pen-based PDA’s, the Palm devices won market dominance by formulating their product specifications in human, rather than technological terms.

What is significant to me in all of this is *when* Palm achieved its success: exactly the same time Microsoft was being found guilty of monopolistic practice and preventing other companies from competing in the personal computer space!

If the Waternet is an appropriate example, convergence is something that exists in the plumbing, not in what is visible to the public. Here, diversity and specialization is the rule. This implies that if there was an opportunity for a company like Palm to benefit from understanding human needs, and developing a new class of device that addressed them, then the same is true for others.

The conclusion to me is that the human-centric approach is not only good design, it leads to good business as well.

On Strength vs. Generality

In the preceding, we have been building the case for a move to diverse specialized tools rather than the *status quo* of one-size-fits-all super-appliances. It would be nice if it was that simple, and we could just start building a broad assortment of such devices. But it is not. Again, our design options are constrained by the amount of complexity that humans can readily manage.

In design, there is a trade-off between *weak general* and *strong specific* systems. Tools, like people, seem to be either a jack-of-all-trades, or specialized. Inherently “super-appliances” fall into the jack-of-all-trades weak-general category.

This weak general vs. strong specific trade-off plays a role in the design of all tools, including chisels, screwdrivers, bicycles, and supercomputers. For reasons that we shall see, *strong general* systems have never been an option. One has always had to make a choice. One of the most significant aspects of combining emerging technologies with appropriate design is the potential to change this for the very first time.




Figure 15: A Graphical Representation of a Weak General System. The area under the red line, shaded in red, represents the cognitive load required to have control or facility with the system.

For a number of years, I have used a series of simple graphics to illustrate the nature of this new potential. We begin in Figure 15 with an illustration of a classic *weak general system*, such as a good kitchen knife. It can do a lot of things, hence its generality, but for many of these there are other tools that can do better, hence its relative lack of strength.

The system illustrated in Figure 16 is at the opposite extreme of the tradeoff. It represents a prototypical *strong specific system*, such as a potato peeler. Compared to the kitchen knife, it can only do a few things, hence its specificity. However, what it does do, it does better than the general purpose knife. Hence its strength.

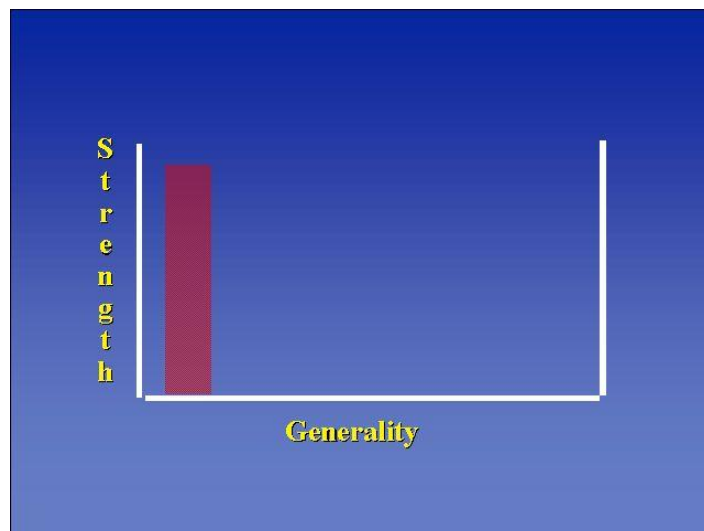


Figure 16: A Strong Specific System. This system can do few things, but it can do them very well.

At this stage, one could reasonably ask, “Why not achieve both strength and generality by having a suite of strong specific tools?” This approach is represented graphically in Figure 17.




Figure 17: *Achieving Strength and Generality Through a Suite of Strong Specific Tool. The problem with this approach is the aggregate complexity of the individual tools. Considered individually, each is manageable, however, considered together, their aggregate complexity is beyond the human capacity to manage.*

Where is the problem here? If the domain and the tools are relatively simple, then there is no problem. But then, if things are so simple, why would we be applying computers and other digital technology in the first place? It is where the domain is broad and complex that the problems arise.

In such cases, even though each tool may be individually manageable, their collective complexity rapidly exceeds a human's ability to cope. Just think of all of the electronic appliances around you. Even if you can work some of them, collectively they begin to overwhelm. This aggregate complexity is represented in these figures as being proportional to the area shaded in red. Going back to our earlier discussion, consider this area to represent the *cognitive load* imposed on the user in order to take advantage of the tools.

Referring back to *God's Law* from Figure 6, in Figure 17, we see that the apparent potential of the tools far exceed the behavioural limits imposed by the *Threshold of Frustration*¹⁴. While over simplified in its telling, this story is nevertheless true. The accumulation of complexity limits the number of specialized tools that we can manage. At least up until now.

As stated at the beginning of this section, my belief is that the technologies emerging today begin to enable us to, for the first time, change this. Understanding how and why this is so is closely linked to my earlier contention that I/O, networking, and location/motion sensing are the most important aspects of technological change to understand and exploit.

My argument is that when the strong specific tools are digital *and* networked, they have (for the first time) the capacity to *communicate* and *cooperate*, thereby assuming much of the load that would otherwise burden the user. This I call the *Net Benefit*, illustrated graphically in Figure 18.

¹⁴ c.f. Figure 8.




Figure 18: The “Net Benefit”. The effective load of exercising control over the set of strong specific tools is made manageable by the benefit of the networked tools communicating and assuming some of the burden.

As an example, consider some of the specialized digital appliances that are found in a modern automobile: the cell phone, the stereo, and the GPS system. Each has an independent function and may likely be manufactured by a different company. Operating each also imposes a degree of load upon the user. In the context of a moving car, this load can become dangerous.

Imagine driving down the freeway listening to the car stereo and having the phone ring. Especially if the music is playing loudly, before answering the phone, your first reaction will likely be to turn down the stereo. All of this can take your concentration, not to mention your eyes and hand, away from driving. This is dangerous in most circumstances.

Next, imagine that these appliances can communicate over a wireless network in the vehicle. If this is the case, they can all know that they are in the car, and the GPS, for example, can let the phone and the stereo know that the car is moving. Thus, when a call comes in, the phone can notify the stereo to turn down the music. In order to let the driver further concentrate on the safe operation of the vehicle, the phone could also use the car network to divert the voice of the caller away from the handset to the stereo speaker closest to the driver. The phone then “listens to” the driver’s voice through a microphone built into the steering wheel it, rather than the one built into the handset. As a result, when the phone rings, the driver need only say the word “answer”, and the overhead in speaking to the remote person is reduced to about that of speaking to a person sitting in the passenger’s seat.¹⁵

This example not only provides an example of “net benefit”, but how a *society of appliances* can leverage their potential benefit through a knowledge of where they are spatially as well as socially (relative to both other appliances and the user).¹⁶

¹⁵ Some of the benefits that I describe in my example are already available in some vehicles. The automobile manufacturers *Lexus*, *Land Rover* and *BMW*, for example, offer car phones that turn down the stereo when there is an incoming call. This is a hint of things to come, but only a start.

¹⁶ The use of an automotive example in this section is significant. There is a strong argument to be made for the case that it is in automotive electronics (which make up about 30% of the cost of a modern car) where these types of systems will first reach the public. The reason is that cars constitute environments where behaviours are well understood, involve a range of electronic systems, and where (unlike the home or office) the environment is under the control of a single company. (What I mean by this last point is that GM has more control over what goes inside your car than the architect Frank Gehry has over what goes into his buildings.) These systems involve real-time, process control, information sharing, and networked cooperating parallel processors from diverse vendors. If I am right, this type of system is what will soon appear in our office and home systems. Yet, given that most schools still focus

Devices as Notation

One of the first things that one studies in Philosophy 101 is the intimate relationship between thought and language. We all have heard the axioms:

- Thought beget language.
- Language beget thought.
- There can be no language without thought.
- There can be no thought without language.
- Notation is a tool of thought.¹⁷

To any sceptics, a good exercise is to see how a switch from Roman to Arabic numerals affects the complexity of performing "simple" long division. Compare the following two representations of a problem:

$$478 \div 72 = ? \quad (1)$$

$$CDLXXVIII \div LXXII = ? \quad (2)$$

At the semantic level, both are identical. Nevertheless, (1) is clearly simpler than (2), and this is not simply due to a lack of familiarity with Roman Numerals. Before the Phoenicians brought us Arabic numerals, long division (from a question of cognitive load) was about as hard as 2nd year calculus is today, even for Romans. But with the switch to the Arabic numerals, with the introduction of the decimal, long division became something that it can be performed by a child.

The significance of this to the design of information appliances derives from a belief that computers are notational instruments, *par excellence*, which have the potential to reduce the complexity of today's world, much as the introduction of the decimal did for mathematics in the past.




Figure 19: *Scientific Visualization.* Here, 3D computer graphics are used to render complex data in a form that is manageable by a human. (UNIRAS/Pacific Visualization Systems)

on teaching traditional functional programming, how well are we preparing our students for this inevitable change in computational paradigm?

¹⁷ A cogent example illustrating the notion of *Notation as a tool of thought* is the Turing Award lecture by Iverson (1980).

Scientific visualisation, as illustrated in Figure 19 is one example where this is manifest. Of course, science is not the only source of complexity in the world. Others have applied similar techniques to visualizing other types of information. One example is illustrated in Figure 20.




Figure 20: *Information Visualization.* Here, visualization techniques are used to render complex information structures manageable. In this case, what is displayed is the organizational structure of a business unit. (Card, Robertson & Mackinlay, 1991).

This is an example on information visualization undertaken at Xerox PARC (Card, Robertson & Mackinlay, 1991). In the example, a 3D graphical device called a *cam tree* is used to effectively represent a complex organizational structure in a form that can easily be assimilated and searched.

But it is not just what is on the screen that can serve as a notational aid facilitating thought or understanding. The device itself can be a component in the representation of the problem. An example of this is seen in the contents of Figure 21.



Figure 21: The Portable Accurate Timepiece. In the spirit of the television game, Jeopardy, what question was this technology developed to answer? Hint: The answer is not "What time is it?"

The way that I like to use this example is in the tradition of the popular television show *Jeopardy*. In a reverse of what most quiz shows do, Jeopardy gives the answer, and asks

contestants “What is the question?” In keeping with this, after I show the image in Figure 21, I ask people, “What question was this device designed to answer?”

Typically, the response is, “What time is it?” But, of course, that answer is too obvious. The response that I am looking for is, “Where am I?” This normally causes confused looks until I point out that the accurate portable time piece was initially developed as an navigational aid, as a means to help calculate longitude (Sobel, 1996).

Prior to the invention of the chronometer, it took three hours of manual calculation to determine longitude. These calculations were typically learned when the navigator was a midshipman, a young boy who left school to go to sea at about 12-14 years old.

Given how far a ship could travel in three hours, and the consequences of error, it is not surprising that some other mechanism was desired.

With the introduction of the chronometer, the calculations were reduced to simple arithmetic, and could be done in minutes. In the sense of causing a reduction in complexity, the chronometer was arguably as important a notational device as the decimal.

Increasing complexity is the enemy, and well-designed devices, appropriately deployed, can play a significant role in turning this tide.

Design in Front of the Glass: What You See (Hear and Touch) is What You *Think* You Get

If, therefore, we think of an information appliance itself as a notational device, how informative is that notation when virtually all devices have, as we have argued, looked the same since 1981? How can the form-factor of a computational device help reduce complexity and increase its value and effectiveness?

Up until now, user interface design has mainly focussed on design issues that lie “behind the glass” of the display. The physical computer, itself, has largely taken as a given, and the energy has been focussed on how to then maximize human performance with that device.

If we want to explore the notion of “device as notation,” the first thing on our agenda should be to change this and focus far more of our efforts on the other side of the glass: on the physical device itself.

I have two simple exercises that I use with audiences to illustrate the potential of going down this path.

First, I ask them to make a quick sketch of a computer. This I ask them to do in less than 15 seconds. What I want is their first response, much like you do in a word association test.

Second, I ask them to perform the same task, this time imagining that it is about 1969.

Over the past few years, I have had over 3,000 people perform these exercises. The drawings illustrated in Figure 22 are typical of the responses that I have received. In response to the first task, for example, the overwhelming majority draw a monitor and a keyboard. A smaller but significant number, as in Figure 22(a), draw a mouse as well.




Figure 22: Two 15 second sketches of "computers": (a) circa 1999, and (b) circa 1969. These are representative of thousands of sketches collected by the author. Similar to about 80% of the sketches collected, neither sketch contains a computer. Rather, what are shown are the input/output devices: the terminal!

The responses that I get to the second task are more varied. (I suspect that this is due to the likelihood that the majority of respondents were not alive in 1960, much less using computers!) What people typically draw looks much like a collection of refrigerators and washing machines, intended to represent the keypunch machines, card readers, tape drives and line printers of the era, such as illustrated in Figure 22(b).

While the consistency of what people draw is interesting, the heart of the exercise lies in noticing the consistency of what they do *not* draw: *almost nobody draws the computer!*

What they draw are the input/output devices, which brings us back to the point made earlier concerning the importance of I/O devices. The exercise highlights the power of what users see and touch (the I/O devices, or "terminal") have in shaping their mental model of the system. Furthermore, the differences in perception that these same input/output transducers elicit illustrate that they are "accidents of history" *and therefore candidates for change.*

These are two of the most powerful observations that a designer of computers could have. What they say is:

1. You can change the input/output devices.
2. By your choice, you can have a huge influence on shaping the end user's mental model of the system.

Consequently, if you know understand your users, their skills, and the context, you can shape their mental model through the affordances of the devices that you use to implement the system. You need not stick to the *status quo*, where every system has the same devices, say the same thing about their purpose, and therefore say nothing.

From these simple drawing tasks emerges much of the power to fundamentally change the perception of computation forever.

In the next two sections we will discuss two examples that hint at where this approach could take us.

Example 1: Tuning Into Browsers

The mechanism whereby most people access the Internet is via a *browser*, which is typically a specialised piece of software that lets one find and read requested documents and data. In this case, the word "read" is used broadly, since reading might involve watching a video or listening to audio. For most, a browser is defined by the example of the two most used

commercial products, Netscape's *Navigator*¹⁸ and Microsoft's *Internet Explorer*¹⁹. This perception is very narrow since conceptually and functionally these two products are virtually identical. Here again, we have an example of a weak general approach to design in that one form of browser is used to access all electronic documents, regardless of type (text, graphics, video or audio). This situation cannot and will not continue. My prediction about the future is this:

The diversity of web browsers tomorrow will match the diversity of ink browsers today.

In everyday English 'ink browser' is just another word for "paper." I used the pseudo technospeak both to show how language can make simple things confusing, as well as to compare the past and the future in common terminology. In today's terminology, books, for example, are a technology for "browsing" ideas written in ink. Given this observation, now think about how diverse, and specialised, the range of "ink browsers" is. We have "Post-it" notes, posters, notepads, books, newspapers, and other forms too numerous to list. There is a context and purpose for each. In the future, this is how diverse Internet browsers will be. And just as we associate certain kinds of information, location, and activity with each form of "ink browser", so will it be with Internet browsers.

Let me give you an example. Ask yourself, "What is a radio?" In computerese, a radio is a browser. It browses the AM, FM, middle wave or long wave spectrum instead of the world-wide web, but it is a browser, nevertheless. And what are the buttons on your car radio? They're clearly bookmarks, just like your Internet browser. Having established this, it is interesting to contemplate the fact that whereas my eighty-year-old mother would claim on the one hand to have no idea how to work a browser (much less know what one is), on the other hand, she has been operating one effectively for decades. As it turns out, it is only poorly designed and inappropriately deployed browsers that she has trouble with. But that is less her fault than the fault of bad design.

Now let us consider that notion of a "radio as browser" in the context of the much-celebrated ability of the internet to deliver audio content, either as a downloaded file (such as in the popular MP3 format), or streamed over the network (using a technology such as Real Audio²⁰).

The technocentric approach to delivering this capability has mainly centred on integrating this capability into existing PC Internet browsers.

The human-centric approach would be to begin by asking questions such as, "Where do you listen to music?", "What appliances do you use and where are they?", and "Where is your record collection?"

Anyone pursuing this line of questioning would rapidly find out that most people have their stereo in a different room than their computer and do most of their spoken word radio listening in the car. From this, they would quickly see the limited benefit of using the PC as the audio player.

On the other hand, they would also quickly see the benefit of delivering this content in a form and at a cost appropriate to where one *does* listen to audio. Hence, one might very well arrive at the design of something that looks like a conventional audio tuner that hooks up to your existing stereo, has radio buttons on it, and gives you immediate and simple access to the available audio.

¹⁸ <http://www.netscape.com/>

¹⁹ <http://www.microsoft.com/catalog/display.asp?site=808&subid=22&pg=1>

²⁰ <http://www.realaudio.com/>



Figure 23: *Two Web Radios: Two Internet appliances for browsing streaming audio are shown. On the left is the iRad²¹ and to the right the Kerbango²². Each appliance bypasses the PC and accesses audio over the Internet directly.*

Since it is a strong-specific device, it will be great for audio and useless for word processing. More to the point, like a digital camera²³, it won't be called a "computer." Rather, it will be called something like a "web radio" or "tuner." In keeping with our "draw a computer" exercise, if it looks like a duck, walks like a duck, and quacks like a duck, it's a duck. It will be defined by its form, location, and the function that it provides, not its internal componentry or logic. Now, after using this example for about four years, this type of device is now starting to appear. Two examples that resemble traditional radios are shown in Figure 23. Another example uses wireless low-power radio to enable your conventional FM radio in one room present audio captured on your PC in another.²⁴



Figure 24: *Taking the Web With You: The Rio portable digital audio MP3 player²⁵.*

²¹ <http://www.audioramp.com>

²² <http://www.kerbango.com>

²³ While most people do not think of their camera as a computer, it is interesting to note that most modern 35 mm cameras have more computational power than many early PCs. This includes not only includes digital cameras, but film cameras as well. These can be thought of as a specialised computers that instead of a mouse and keyboard for input, take in light through the lens, and instead of a CRT, output chemical emulsion on film. Great for taking pictures. Useless for doing your income tax.

²⁴ <http://www.radiowebcaster.com/index.html>

<http://www.sonicbox.com>

²⁵ <http://www.riohome.com/>

And, in keeping with the trend towards divergence, "digital audio browsers" are beginning to come in forms that are increasingly matching the diversity of more conventional audio listening devices, such as the portable MP3 players, such as illustrated in Figure 24 which follow in the tradition of portable audio cassette and CD players. With these technologies it is now possible to download not only music but a large range of spoken word books²⁶.

Example 2: e-Commerce Beyond the Amazon

Our second example concerns one of the most discussed and written about aspect of the new economy: e-commerce. Given the amount of money at stake, this seems a good place to make our last point.

Again, my approach is by way of a word-association game. Typically, what I do is ask audiences to tell me, as quickly as possible, what company name comes to mind when I say the word "e-commerce."

More than 90% of the time the answer is "Amazon.com".

I then tell them what my answer is, which is "Symbol Technologies,"²⁷ which generally results in confused looks and the question, "Who are they?" and "Why did you choose them?"

To answer the second question first, I would estimate that the amount of e-commerce that goes through browsers driven by their technology likely exceeds the total e-commerce being transacted on all Netscape and Internet Explorer browsers combined (including Amazon.com) by about five million times!



Figure 25: A Barcode Reader used in a store checkout counter. (Image courtesy of Symbol Technologies).

²⁶ <http://www.audible.com>

²⁷ <http://www.symbol.com>

How can this be, especially if most people have never heard of them? The answer is that Symbol Technologies is the largest manufacturer of bar-code readers in the world, and therefore their technology, such as that shown in Figure 25, is used in supermarket checkout counters all over the world.

Now most people would claim that checking out groceries is not e-commerce, and even if it was, the barcode-driven technology is not a “browser” in any meaningful sense. But then, most people have not read this article, and therefore do not appreciate the power of I/O devices to shape perceptions. For, on deeper consideration, this clearly is a browser, and e-commerce, and the transactions run over the same wires, routers and servers as most other e-commerce. But, in the tradition of the Waternet, you only notice what doesn't work. Successful design is transparent.

We don't notice the effective use of e-commerce at the checkout counter precisely because it *does* work and is trusted.

But if this is true, then what does this say about the depth of analysis of all those e-commerce experts that we have read, none of whom even considered this class of system?²⁸

The Renaissance is Over – Long Live the Renaissance

This essay has been largely about missed potential, and the need for change. The emphasis has been on moving outside of the relatively narrow confines of technology and informing it with a more human-centric approach.

In the process, I could be criticised for being too critical of engineering and computer science. I have spent time describing things that they should know and why they should know them, but I have been almost silent on how they might achieve this knowledge, especially given the pressures of an already strained curriculum.

It is interesting that the design principles that we can apply to the social engineering that addresses this issue are the same as those that we have already discussed in terms of the engineering of future information appliances. In this, I am referring to our earlier discussion of weak-general vs strong-specific systems.

The exact same issues that we saw in this discussion are evident in the tension between the need for discipline specialization vs general holistic knowledge. Given the much-discussed constraints on human ability, how can we expect an individual to maintain the requisite specialist knowledge in their technological discipline, while at the same time have the needed competence in industrial design, sociology, anthropology, psychology, etc., which this essay implies are required to do one's job?

In short, just as we discussed the notion of a networked society of appliances, so does the solution in this case lie in a social network of specialized individuals. Likewise, we get a *net benefit*; however, in this case, it is due to a *social network* of people rather than a network of computers.

In 1959, Sir Charles P. Snow presented a landmark lecture on the relationship between the science community and that of the arts and humanities, or “literary intellectuals” (Snow, 1964). Snow characterized these two communities as having lost even the pretence of a common culture. That is, they had lost the ability to communicate on any plane of serious intellectual endeavour. This he argued from the perspective of creativity, “intellectual life,” and normal day-to-day living. He coined the term “the two cultures,” to characterize the polarization of these two communities.

²⁸ This is doubly worrying, since in my opinion, it is precisely this type of approach that is going to dominate in the long run.

Today, our academic institutions and funding agencies (among other things), are set up to reinforce the silo mentality of these separate cultures. Industrial designers go to art college, psychologists, sociologists, and anthropologists are in the Faculty of Arts and Sciences at a liberal arts school while the computer scientists and engineers are in a separate faculty and perhaps at a separate institute of technology. And all are funded by different agencies.

Yet, if you are to believe what I am saying, the skill sets of all these disciplines must be effectively applied in concert to achieve the potential that is sitting there waiting to be exploited. But the cultural biases identified by Snow are working strongly against this.

One of the consequences of Snow's initial lecture and the accompanying essay was a spate of over 100 articles discussing the educational implications of the "two cultures." A commonly expressed view expressed was the need to teach engineers more about the arts and social scientists and artists more about math and science.

A superficial reading of my own essay might conclude that such a simplistic approach to rethinking the education of engineers and computer scientists is what I would advocate. But that would be wrong.

Essentially, this school of thought can be reduced to "Let us create a community of renaissance men and women, instead of these specialists." I would suggest that the result of this would be a culture of mediocre generalists, which is not what we need. Renaissance man and woman have not been viable for the past 300-400 years. The world has simply become too complex.

On the other hand, the notion of *renaissance team* is entirely viable: a social network of specialists from different disciplines working as a team with a common language. But while viable, the systemic biases of language, funding and institutional barriers make this type of team the exception, rather than the norm.

The problems begin right from the start of our educational system, with its emphasis and reward system is based on the performance of the individual rather than the team. And, cases where team performance is encouraged, from the perspective of skills and culture – in the sense articulated by C.P. Snow – their composition risks skewing towards the homogenous, rather than heterogeneity.

If, as claimed by the psychologist Jean Piaget (1947/2001), intelligence is defined by the ability to *adapt* to a changing environment *via accommodation* and *assimilation*, then given the societal changes being driven by these new technologies, our policies and educational institutions must behave in an intelligent way and adapt their structure and curricula to reflect these changes.

Conclusions

I have long described myself as a *sceptomist*: one who strives to achieve a balance between their inner sceptic and optimist. What I have seen of the evolution of technology over the years has been most discouraging in terms of meeting the social and human benefits that might have been achieved. On the other hand, I cannot help but marvel at the technological advances that have been made.

Yet, frustration continues due to the degree to which the potential of the latter to positively impact the former is left unfulfilled. Yet, the latent potential is real and realizable.

However, to make it so is going to require a significant rethink of where the human fits into the disciplines of engineering and computer science. Computer scientists are going to have to realize that "primary memory" is the human brain, not RAM, and that the effectiveness and integrity of the transfer of data between primary and secondary memory is as important as that between RAM and a hard disk.

Technology is a critical component in the future of systems design, but it is not sufficient. It is time that this fact was reflected in our educational institutions, industry, and professional societies.

Hopefully this essay has made some small progress towards making that happen, even if 20 years late.

References

- Buxton, W. (1994). [Human skills in interface design](#). In L.W. MacDonald & J. Vince (Eds.). *Interacting with virtual environments*. New York: Wiley, 1-12.
- Card, S.K., Robertson, G.G., and Mackinlay, J.D. (1991). [The information visualizer, an information workspace](#). *Proceedings of the 1991 Human factors in Computing Systems Conference, (CHI'91)*, 181 – 186.
- Hutchins E. (1990/2014). [The technology of Team Navigation](#). In J. Galegher, R. Kraut & C. Egido (Eds.), *Intellectual teamwork: Social and technical bases of cooperative work*. Hillsdale, NJ: Lawrence Erlbaum, pp 191-220.
- Hutchins, E. (1996). *Cognition in the Wild*. Cambridge MA: MIT Press.
- Iverson, Kenneth E. (1980). [Notation as a Tool of Thought](#). *Communications of the ACM*, 23(8), 444-465.
- Smith, D.C., Irby, C., Kimball, R., Verplank, W. & Harslem, E. (1983). [Designing the Star User Interface](#). In P. Degano & E. Sandewall (Eds), *Integrated Interactive Computing Systems*. Amsterdam: North-Holland, 297-313. Originally appeared in *Byte*, 7(4), April 1982, 242-282.
- Snow, C.P. (1964). [The Two Cultures: and a Second Look](#). Cambridge: Cambridge University Press.
- Sobel, D. (1996). [Longitude: The True Story of a Lone Genius Who Solved the Greatest Scientific Problem of His Time](#). New York: Penguin.