

Software Architecture Axiom 1: No Moving Parts

by **Francis Hsu**

[Born in China and educated at Rutgers College, Francis Hsu works on Data-Systems Architecture for large organizations, mainly Fortune 100, and is now working for the U.S. State Department.]

Preamble

In 1900, David Hilbert challenged mathematicians with 23 problems for the new century. Now as software practitioners face a new millennium, it seems appropriate that we too set challenging tasks for ourselves.

In this light, I think trying to define or discover Software Architecture Axioms is a worthy goal. First, let's be clear that software axioms are not necessarily mathematical in nature, as the following essay shows. Second, although software is as abstract as mathematics, it is clearly very practical—even discounting the billions it has made for investors. Third, as technology, software with hardware has unmistakably spread faster and further than any prior technology ever. Fourth, like sharp edge tools, fire and the wheel, software is a permanent companion of our species. Fifth, let's face it, for many of us, software is an addiction and great fun too.

Software axioms should be a) a goal, b) a means of clarifying a goal, or c) a method to attain a goal. The following is a category a) goal essay.

Introduction

In human history, moving parts were identified in two ways. In ancient times, motion was life. No motion meant death. Since at least the Industrial Revolution, moving parts have meant increased wear, promoted heat, required maintenance and replacements, and degraded the accuracy and working of machines. Though unstated, motion was assumed to be local and took time to complete. Instant motion over large distances is a modern idea.

What the computer and communications industries (collectively the Information Technology (IT) industry) have ignored is how data motion has a severe impact on data quality. In short, the less data is moved, the less error there is.

Data (or information) error is one of the most important issues not only within the IT industry, but in the world at large. This is, in part, because we depend so much on the IT infrastructure. But it is also because we demand to share current information with others anywhere on this planet, at anytime. This is a modern demand. What is this all about? This paper attempts to clarify the situation.

The data error problem has two parts: where data resides and how it is moved. Thus data error can be summarized as 'data residency' and 'data motion' problems. Sometimes they are linked. The solution has been growing on us for years. One is in the form of data storage capacity and the other how the Internet is used.

Data Residency

Data residency is simply, in what form is data stored or recorded? The human brain-mind is the first place where data is stored. We all know it's true. And we all know that we don't know how it's done. But outside our skulls, data residency has an extensive history: from the cave walls where hominids drew animals some 30,000 years ago to the first clay tablets of about 5,000 years ago, to the present day. For us, data residency means the sharing of symbols across space and time. In practice, this means the writing or recording of language symbols.

At first, these symbols were recorded on clay tablets. As technology progressed, we learned to record on many other types of materials, such as bones, leather, wood, etc. As civilization and literacy spread, so did recording technology. Thus the following recording methods appeared: calfskin, copper, paper, pencil, printing, pen, typewriter, carbon paper, stencil, embossing, photographic film, magnetic media, optical media, and soon nano-technology. See Table 1 for a list of recording methods. These recording methods are divided into two broad classes: those which are human-scaled, so that we can both record and read without needing special equipment ('Yes' under the Direct Use column), and those which exceed human-scale ('No' under Direct Use column).

Table 1 A non-exhaustive list of some recording methods in human history.

<u>Date (approx)</u>	<u>Direct Use?</u>	<u>Method name</u>
30,000 BCE	Yes	cave wall drawings
3,000 BCE	Yes	1 st inscribed clay tablets
3,300 BCE	Yes	hieroglyphs appear in Egypt
2,500 BCE	Yes	logographs on bones in China
100	Yes	paper
1450	No	Gutenberg printing press
1802	No	photographic nitrate film
1867	No	Sholes' typewriter
1960?	No	.. microfiche

1956	No	IBM's RAMAC disk storage
1971	No	IBM's 8-inch diskette
1976	No	Shugart's 5.25 inch diskette
1981	No	Sony's 3.5 inch diskette
1981	No	Sony-Phillips's Compact Disc (CD)
1995	No	DVD standard worked out
2002	No	Blu-ray DVD standard set
2010?	No	Carbon nano-tube storage?
2020?	No	Atomic storage?
2030?	No	Quantum storage?
2100?	No	String-theory encoding?

In classing a method as 'Yes' or 'No' under the Direct Use column, there's a degree of arbitrariness. With pen(cil) and paper, we can record (write) directly on the medium (the paper). One could argue that pen(cil) and paper is already a level of indirection: it's not our fingers. This is true. With the typewriter it becomes obvious that its entire purpose is to remove (direct) human hands from the final result. This ensures its consistency, legibility, evenness, cleanliness of final product. In all the cases where Direct Use is 'No,' either because of size, speed, consistency, accuracy, etc. technology has taken humans further and further away from the recording process. In essence, this is what technology does. It replaces the inconsistent and inaccurate motion of human hands with the mechanical motion of machines, where neither size nor speed is a constraint.

As we consider where data errors occurs, the form of data residency or storage does not seem high on the list of suspects. After all, once data is written, most recording media are relatively stable. (A side note: Ever since book publishers have noticed that some books were actually disintegrating within a few years, there has been a concerted effort to publish books on acid-free paper. Such books usually identify the paper as 'acid-free.' Such books are rated to last 500 years or so. That's not too bad for a technology to last several times the lifetime of an average human.) Of course integral to this is the fact that writing is an act for literate peoples. In human history, mass literacy (say greater than 50% of a population) was not widely achieved until the 20th century. Prior to that, published writing was mainly by the elite of any society.

[Sidebar: Why data 'residency' as opposed to 'storage?' In science or technology the use of terms can either help or hinder understanding. While 'storage' is the correct concept, the word also implies a stability of static structure which may not always be true. It is known for example, that human civilization has been broadcasting radio and TV signals outward toward galactic space since the 1920s. Those speed of light signals still encodes the data broadcasted then. Thus, data 'resides' in those signals. Until we can travel faster than light, we'll never catch up with any of those signals. It is awkward to say data is 'stored' there. Although it certainly resides there.]

Data Motion

The second important concept is data motion. Data motion is simply: how is data transmitted?

A complete study of data motion involving all human senses of sight, sound, smell, taste and touch, is not possible in this short paper. Here the focus will be primarily on visual and audio data.

There are three broad categories of data motion. All of them span space and time. This aspect is subsumed in all the others. First, data moves between technologies. For example, when the Rosetta Stone was found, it was stone-based technology. It was then copied onto paper-based technology. It was later photographed onto chemical-film technology. Now it is also on digital-based technology. This latter technology is more inclusive, in that it may have magnetic and optical modes for recording. In general, data motion between technologies is semantically-neutral, that is, the content's meaning does not change. Second, data moves between or across cultures. When the Rosetta Stone was discovered 1799 it moved from stone age culture (actually dated to 196 BCE) into a culture undergoing the Industrial Revolution. Since that time, with colonialism and imperialism spreading worldwide, the data and values of different cultures mixed. This category also includes the big divide between the oral (non-recorded) and the written (recorded) cultures. Both oral and written types of data motions are semantically volatile, that is, the meaning changes radically thus creating both understanding and mis-understanding between cultures. The third and final category of data motion is that between people and computers. This category is in some ways a mix of category 1 technology and category 2 culture. This partial mixture is worthy of a third category, as shown below.

In practice, humans notice the second category, data motion across cultures, because it is the most socially disruptive. The other categories are usually the province of specialists.

In a simplified form, we can view the last (third) category as a grid, as depicted in Figure 1. There are three types of exchanges, labeled A, B and C. The exchange between Human and Human is box 'A.'

	Human	Computer
Human	A	B
Computer	B	C

Figure 1. Data Exchange between Humans and Computers.

Our foci of interest are the exchanges between Human and Computers (B) and between Computers (C). These in fact are where our analysis will reveal the most interesting conditions.

But before analyzing that, we should understand what happens within a single un-connected computer. A personal computer is certainly the most exemplary computer in terms of quantity and extent of distribution. It is also exemplary in the numbers and types of devices that can attach to it. In 2003, the PC has at least these parts: micro-processor, cache memory (some have more than 1 level), on-chip bus, main memory (various flavors of RAMs), systems bus (connects micro-processor to other devices), graphics adapter, graphics RAM, various device controllers, hard disk(s), CD-RW, DVD-RW, diskette, USB ports, modem, network card, keyboard, mouse, soundcard, etc. The prior list contains all the places where data can move into, out of and through. It is well known that these parts are not all equally reliable. So the data error rates in using these parts vary. Data error rates are measured as the number of read or write failures per n attempts. Reliability of micro-electronics has improved so much in past decades that n is usually in the billions (10^9) or higher. These are excellent numbers. The ratings are specified for individual parts. For example, CD-RW devices are reported to have error rates of 1 in 10^{12} . However, when all the listed parts are assembled together to form a computer, their total reliability drops. This is the proverbial 'the chain is only as strong as its weakest link.' The assembly process, connecting the parts, seems to introduce flaws--usually at the junction points.

Table 2 lists the error rates of different parts of an un-connected computer.

Table 2 Error rates for individual computer parts (author's estimates).

1 in	Part name
10^{18}	micro-processor (CPU)
10^{18}	cache
10^{18}	on-chip bus
10^{16}	S/DRAM
10^{16}	systems bus
10^{13}	graphics adapter
10^{13}	graphics RAM
10^{13}	device controllers
10^{13}	*hard disk
10^{12}	CD-RW
10^{12}	DVD
10^{11}	USB
10^{10}	diskette
10^{11}	network card
10^{10}	modem
10^{10}	soundcard
10^7	keyboard (the device itself)
10^7	mouse (the device itself)
10^3	humans using a keyboard (optimistic)

* = (Computer Architecture A Quantitative Approach, 3rd Ed. By John L Hennessy & David A Patterson, San Francisco: Morgan-Kauffman, 2003 p762.)

In general, the more distant a part is from the CPU, the less reliable it seems to be. For both the keyboard and the mouse, the error rate is for the device itself. In both cases, when humans use them and either hit wrong keys or click the wrong button, that should not be counted against the devices. Note that technically, the last item, 'human using a keyboard' is not a device. It is included for comparison. Also, the number of 10^3 is very optimistic: it implies that out of 1,000 key strokes, humans make only one mistake. In typing this paper, I know I frequently hit the 'U' key when I intended to hit the 'I' key, for example.

One more thing: these error rates are not content specific. That is, it does not matter whether the data moved are text files, binary data, computer instructions, audio, video or binary large objects (BLOBs). In short, they are data content normalized.

With the normalized error rates within a single un-connected computer, we are now ready to consider the exchanges between Human and Computers (B) and between Computers (C). To better see where these stand, consider Figure 2 the Analog-Digital and Human-Computer Data Exchange diagram. [There is an attached PowerPoint chart to show this more clearly.] This diagram depicts the interactions among humans and computers in two different worlds. The upper half of diagram is the world where humans live: it's analog. The lower half is where computers exist: it's digital. There are four (4) types of interactions possible:

(1) is between humans which takes place entirely in analog realm. This is simply human conversation or communication. We know from talking in loud environments that we are extremely tolerant of noise: somehow our minds can still figure out what's being said. When we talk, the listener 'gets a copy of what we said.' In most instances, we get an in-exact copy: we extract the meaning, and that mostly is what we remember. Most humans do not remember word for word what another person said, unless the utterance is short. Our memory does a 'data compression' routine on our conversations. We extract the meaning and leave the mass of exactness behind. This is nature's efficiency at work.

(2) occurs when humans enter data into computers; this crosses the analog-digital boundary, from an amorphous analog world to the strict digital one. Prior to personal computers, only large organizations had a multitude of clerks typing data into computers. Today with PCs and connections to the Internet, wireless phones and other devices, anyone can enter data into

computers. As Table 2 data has shown, human data entry has the lowest reliability in terms of error rates.

(3) occurs when computers provides data to humans. This is the reverse of interaction #2. Here data is crossing the boundary from digital back to analog. This is a more subtle interaction because, unless the data is 'erased,' there is no loss of fidelity to the data. To be sure, the quality of data coming out of a computer depends on the quality of data going in. Given the reliability of computer-connected devices, we can say with confidence that if some data is in error (does not reflect reality) it is most likely not because any computer did it. It is because the data was entered incorrectly.

(4) The exchanges take place in the digital world between computers. This can be between two supercomputers using channel-to-channel (CTC) direct links within a datacenter. Or generically, this can be any computer linked to the Internet. For accuracy, this interaction requires low noise level connections between computers. Over the decades, entire industries have grown up providing the higher quality connections. Although this infrastructure is not as reliable as the parts within a microprocessor chip, with redundancy, error correction code (ECC) and standard protocols such as TCP/IP, this has become very reliable indeed.

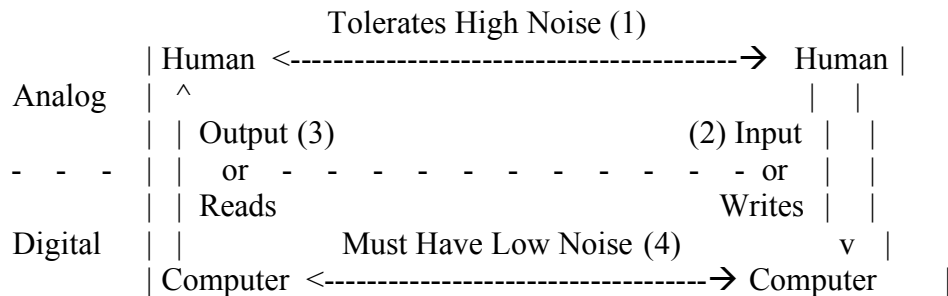


Figure 2 The Analog-Digital and Human-Computer Data Exchange Diagram

There is another perspective to error rates. If we plot the data from Table 2 onto a Sigma graph, (see Figure 3 [Attached Powerpoint slide]) the results are dramatic. The horizontal line is the sigma values, going from left to right, 9 to 1. The famous 6-sigma near center means an error rate of 3.4 per 1,000,000 occurrences. This translates to an accuracy of 99.99966%, sometimes called 'five nines.' While none of the Devices have consistent 'point' sources—the lines show ranges—it is obvious that 'human on -' has the widest range. This means simply that humans are inconsistent. The 80/20 rule is well known: 20% of one's effort should result in 80% of the results. The Sigma graph shows that one source (human on keyboard or mouse) produces close to 100% of the errors. Crossing the analog-digital divide by humans is highly error prone.

Figure 3 6-Sigma and Data Error Rates

	.999999999999669331	.6914	
CPU	-						
D cache	-						
on-chip bus	-						
S/DRAM	-						
e systems bus	-						
graphics adapter	-						
graphics RAM	-						
v device controllers	-						
hard disk	-						
CD-RW	-						
i DVD	-						
USB	-						
diskette	-						
c network card	-						
modem	-						
soundcard	-						
e keyboard		-					
mouse		-					
humans on keyboard						-----	>
	9	8	7	6	5	4	3
		S	I	G	M	A	

Conclusion

Motion used to identify life. With the rise of industrial technology, motion represented the slow decay of friction, heat, and constant monitoring. With the advent of information technology we've gone another step forward. This paper defined two data forms, 'residency' and 'motion.' Of these two forms, 'data motion' is a worse source of data errors. Within 'data motion,' there are three categories: technological, cultural and that between people and computers. This last category focused on the weakest link of the 'data motion' chain: human input (writing) into computers. It seems logical then, that if the IT industry wants to improve data reliability and lower data error rates further, whatever future systems and applications architecture we implement must consider lowering the frequency and necessity of human data inputs. In this regard, no moving parts (NMP) is the ultimate goal, but less moving parts (LMP) is the way to get there.

[End]

