A Trillion and Change

Loren Rhodes

Convocation Address May 5, 2009

Thank you President Kepple, Provost Lakso, platform party, faculty and staff, parents and friends—and most of all—the class of 2009. Congratulations to the graduating seniors and today's award winners. I'm honored to speak to you today. I will try not to embarrass anyone, especially Casey Chew's friend Aaron Rhodes[*]—we'll see. One opening word, however—the one that the men's tennis team yells in encouragement—"SUT!"

We've been through an amazing year of changes. Some "ups" and a few more "downs" than "ups". Where the country and world is going is anyone's guess. We've just been through a historic presidential election that promised change. I began working on this talk the week before President Obama's inauguration and now his first 100 days have just passed. In January I truly hoped the new administration would be an "up".

Of course the most obvious "down" has been the economy. In January the reports were that the stock market value lost $1.2 trillion over 2008, and when I looked at my end-of-year investment portfolio I certainly know I was a part of that loss. So we have a national debt of eleven trillion dollars, total government assistance since the beginning of the financial crisis of nearly twelve trillion dollars, and the promise that we'll likely have one trillion dollar deficits annually for the near future. [1] These dollar amounts are truly astronomical numbers. Have we seen at least the bottom of this economic crisis? We can only hope.

While politicians and experts debate, it is anyone's guess what changes the world will be facing economically and politically, and now we add in an H1N1 swine flu pandemic. Here I am talking on topics for which I hold deep concern as you do; yet they are beyond my areas of expertise, so I offer no ready answers for the world problems of today. Rather, I'd like to share my passion for computers and technology and some of the advancements therein that will affect you in the future, most of which are an "up" type change. We need to be cognizant of these rapid technology changes and the promises we are experiencing, so that we may influence progress for a positive world. Wait, don't turn on that iPod yet, don't answer that text message, or send that next tweet—all this new technology now in your pocket. This is important stuff.

Let me start with just a little background on my own computing experiences and those at Juniata over the past thirty years to help you appreciate the expanding computing technologies that seem to run our lives and in some cases overtake them. My first computer experience was in high school on a 1960's era IBM mainframe with sixteen kilobytes of memory. That's "kilobytes"— a thousand bytes. We use up sixteen kilobytes of memory for less than one second of music today. My undergraduate years were spent on a third generation IBM mainframe with sixteen megabytes of storage and probably about 600 MB disk space. That's "megabytes"— a million bytes. That storage, shared by all Penn State computer users, would be equivalent to one CD's worth of music today. During my first fifteen years of teaching at Juniata, we had a fourth generation mini-computer that again served all academic computing needs. Memory on that machine was also on the order of megabytes of memory and disk space and the machine was accessible from about sixty terminals. Ah, the politics of who got a terminal in their office and which buildings and floors should house the public labs! (I don't miss those days.) In the mid-nineties—the start of your recollection of computing technology—the PC era began, and ultimately, the connection to the commercial internet that we expect today.

And you have observed, I trust, the truly exponential increases in speeds and capacities with your own PCs, and, the aggravating obsolescence that comes with any technology purchase within months. We have a principle in computer science that describes this exponential increase called Moore's law.[2] It was first observed by Gordon Moore, cofounder of Intel Corp. It states that the number of components placed in the same area of a computer chip, doubles every 18-24 months. He proposed this "law" in 1965 and it has held rather true to today (check out a nice graphic in Wikipedia showing this amazing consistency).[3] This doubling principle is observed in increased memory capacities and has, until about four years ago, also resulted in doubling computing speeds for all types of computers. So today we enjoy literally gigahertz computing speeds and gigabytes of storage. Giga—the prefix is for a billion—with a "B". But the terabyte disk drive is becoming commonplace, which can be bought for a mere $100 for all those gigabyte video and music collections you might have.

A terabyte storage device—that is a trillion bytes. Moore's Law has brought us quickly to that "trillion" number in technology. A trillion dollar annual debt. The twelve trillion dollar bailout cost. Scores of trillions of spam emails sent daily. Google, last summer indexed its one trillionth web page. A trillion—a hard number to fathom. So…how big is it? I recall secondary school math teachers with activities conveying the size of a million to us students: projects of collecting a million aluminum can tabs; or a million pennies (only $10,000—that doesn't cover a semester of tuition!). It's fairly safe to say that you will earn a million dollars over your career, probably several times over—once you find that job in today's economy (but don't forget to give some of that back to good old JC).

This “trillion” number, however, has caught my attention recently as it keeps cropping up. So let’s try to get a sense of a trillion. A trillion is a million million—or a million squared. I am guessing the size of Rosenberger auditorium, excluding the stage area, is 100 by 80 by 35 feet, for a volume of 280,000 cubic feet. If we want to fill this space with a trillion of something, how big would those objects be? Beachballs? Tennis balls? Marbles? A trillion divided by the 280,000 cubic feet means about 3.6 million units of something per cubic foot. The cube root of 3.6 million is 153. Or, 153 of these objects if laid end to end would make a foot, which is about thirteen to the inch. Those little Styrofoam balls (not the peanuts) used in packing or bean-bag chairs are roughly that size. So it would take about a trillion of them to fill Rosenberger. So now I expect that you’ll remember nothing else of what I say today but the image of us buried in one trillion bean bag chair Styrofoam balls filling this auditorium.

Let’s push these big numbers just a little further as to what we might encounter in the future. Storage network servers already have capacities rated in petabytes (a thousand trillion bytes). Google processes at least twenty petabytes of information in images and data every day. That’s 20,000 trillion bytes. The fastest supercomputers are predicted to soon be able to process at twenty petaflops, or twenty thousand trillion calculations per second. The next storage unit after petabytes will be exabytes (million trillion). Petabytes and exabytes will soon be commonplace like megabytes and gigabytes. But be careful; it has become easy to just think of these increasing magnitudes as simply the next larger number like we just added another thousand—instead these changes are truly a 1000 fold factor.

Moore’s Law has given us increased storage capacities and has shrunk the size of devices that we see in Blackberries, iPhones, iPods, and now Kindles. Curiously the doubling of computing speeds has, in practical terms, reached its end, but Moore’s Law continues. Instead, computer manufacturers have begun to pack more computing processing units on the chip rather than a single faster one. Dual core and quad core computers are already the norm. Moore’s Law now might suggest that the number of processors will double in a PC every eighteen months so that PCs in a few years could easily have over 100 processors and servers over a 1000. Wow! But, unfortunately, today’s applications won’t run four times faster simply because there are four processors available. The problem is that billions of lines of software have to be completely rewritten to utilize these multiple processor PCs. Supercomputers with thousands of processors have been around for decades predicting weather, doing DNA and Google searches, mining phone calls for terrorist threats, all by software written by a small cadre of specially trained software professionals. Supercomputing will soon be in your PC. So my own computer science education field is now facing a major upheaval to alter our teaching because multiple processor coding will be mainstream. Not that we haven’t dealt with changes before—we’re used to that, but this is a significant change in what we must incorporate in our teaching.

I've mentioned Google; it is, at the moment, THE technology company to work for. This is where the top new PhDs in computer science go. In my undergraduate days it was IBM and Bell Labs, the research arm of the phone company. In the '80s and '90s it was Microsoft and now it's Google. Google is only ten years old and has become a verb in our lexicon. We say "Google it" and we know what you mean while your grandmother may wonder if she should be offended. We never IBM'ed anything or Microsoft'ed something. But we "Google" daily. It's a fascinating company that started out performing 20,000 Internet searches per day with results taking about three seconds. According to the Internet research company comScore, Google performed 235 million searches per day last July searching its trillion plus web pages with an average response time of less than a second. Randall Stross's book, *Planet Google: One Company's Audacious Plan to Organize Everything We Know,* describes how Google won the search engine wars over the likes of Yahoo and company, by keeping humans out of the indexing process.[4] He discusses Google's attempts to index all digital and printed material of the world as well as video and audio, despite copyright concerns. Keep your eye on Google, whose company motto is, by the way, "Do no evil." Let me describe one of the many interesting projects at Google. They have built a natural language translation engine, much like their searching engine, that involves no human intervention. Why no humans? They understand the explosion of digital information that is occurring. Google's engine does language translation mindlessly by using its huge indexing and text scanning capabilities, and matching one to five word phrases from texts in one language to phrases from parallel texts in a second language. To build its translation engine, over two trillion words were analyzed from its wealth of digitized documents. As of May 2008, Google could translate between any pair of twenty-three languages.

Combining these advances in machine language translation, voice recognition, robotics and other developments in artificial intelligence, the question can become, "How soon will technology rival the human mind?" Did you know that checkers, the game, was solved in July 2007?[5] That is, a computer can play a perfect checkers game and not lose. Fear not, chess, on the other hand, still requires a heuristic approach, but machines have beaten grand masters for years. Futurist Ray Kurzweil, author of *The Age of Spiritual Machines,*[6] is producing a movie for release later in 2009 called *The Singularity is Near* which expands on his 2005 book of the same name.[7] He predicts that singularity will happen in 20 years—this is the point at which computers reach human intelligence and capacity. We should not be surprised!

Other technological changes are on the forefront and soon in the marketplace: flexible displays, implanted chips, wearable computers, Kindles, all tied to ubiquitous communication networks. Terabytes, no…, petabytes, thousand trillion bytes of information of all forms will abound and amass. Devices yet to be on the market with exabyte capacity will record everything you see and do in ever so smaller packages.

The exponential pace of technological change continues unabated. You can't stop it. You can't expect it to slow. Expect change. Change will happen! You've seen this technological change already affect your lives like no previous generation, so embrace it as it involves trillions of bytes and trillions of dollars. Juniata has prepared you with critical thinking skills and sensibilities, and a desire to make the world a better place. My hope is that you can leverage these technology changes and advances for the good into personal success and progress.

Congratulations to today's award winners; good luck to Aaron and his 2009 classmates. And whatever it means—SUT!

NOTES

[1] Graduating son of the author.

[1] Associated Press graphic, published in *Altoona Mirror*, March 24, 2009.

[2] Gordon Moore, "Cramming More Components Onto Integrated Circuits," *Electronics Magazine*, 19 (1965).

[3] http://en.wikipedia.org/wiki/Moore%27s_law.

[4] Randall Stross, *Planet Google: One Company's Audacious Plan to Organize Everything We Know*, (New York: Free Press, 2008).

[5] Jonathan Schaeffer, et al, "Checker is Solved", *Science,* 317 (2007): pp. 1518-1522.

[6] Ray Kurzweil, *The Age of Spiritual Machines: When Computers Exceed Human Intelligence* (New York: Penguin Books, 1999).

[7] Ray Kurzweil, *The Singularity is Near: When Humans Transcend Biology* (New York: Viking, 2005).