LBNL- 46509

Rebuttal to Testimony on 'Kyoto and the Internet: The Energy Implications of the Digital Economy'

Jonathan G. Koomey

Energy Analysis Department
Environmental Energy Technologies Division
Ernest Orlando Lawrence Berkeley National Laboratory
University of California
Berkeley, CA 94720

This report and related documents can be downloaded at
http://enduse.LBL.gov/Projects/InfoTech.html

August 2000

This work was supported by the Office of Atmospheric Programs of the U.S. Environmental Protection Agency. Prepared for the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

Origins of this rebuttal

In the past year and a half, I have been witness to an extraordinary event: an analysis based on demonstrably incorrect data and flawed logic has achieved the status of conventional wisdom, in spite of my and my colleagues' best efforts to refute its assertions. The results continue to be cited by an unsuspecting press, and even by people who ought to know better.

In May 1999, Mark P. Mills published a report for the Greening Earth Society (summarized in an article in Forbes Magazine) that attempted to calculate the "Internet related" portion of electricity use. This report claimed that electricity use associated with the Internet totaled about 8 percent of all U.S. electricity use in 1998, that the entire "digital economy" accounted for 13 percent, and that this sector would grow to consume half of all electricity in the next decade. Subsequent to the publication of this report, there was an exchange of technical emails between Amory Lovins, Joe Romm, Mark Mills, and others (http://www.rmi.org/images/other/E-MMABLInternet.pdf).

My colleagues and I at Lawrence Berkeley National Laboratory examined Mills' calculations of Internet-related electricity use in a technical memo dated December 9, 1999 (Koomey et al. 1999). We found that Mills significantly overestimated electricity use, in some cases by more than an order of magnitude. We adjusted his estimates to reflect measured data and more accurate assumptions, bringing Mills' overall estimate of total Internet-related electricity use down by about a factor of eight.

On February 2, 2000, the House Subcommittee on National Economic Growth, Natural Resources, and Regulatory Affairs held hearings on "Kyoto and the Internet: the Energy Implications of the Digital Economy". At that hearing, Mark Mills, Jay Hakes, and Joseph Romm testified. Mills' testimony was an attempt to defend his earlier analysis. Hakes' testimony dealt mainly with the data used to understand recent events in the overall economy, but he concluded that Mills had overestimated electricity use substantially. Romm's testimony focused mainly on the indirect effects of information technology (IT) on resource use, and cited many examples of how IT can reduce overall resource use.

In June 2000, my colleagues and I completed our first comprehensive assessment of office equipment energy use since 1995 (Kawamoto et al. 2000). This report includes electricity used by network equipment, and estimates total energy used by office and network equipment for residential, commercial, and industrial sectors. Note that this report does not focus just on the "Internet related" portion of electricity use like the Forbes article does, so care must be used in comparing the findings of those two studies. This report confirms the results of our earlier technical memo and our 1995 report (Koomey et al. 1995). The explosive growth in electricity demand that Mills alleges simply does not show up in the data.

Because Mills made specific allegations about our work in his testimony, I felt it necessary to craft a response, and the annotations below are that response. The testimony and responses delve into technical detail that will be superfluous for many readers. For their benefit, I summarize the key points in our rebuttal as follows:

In his Congressional testimony, Mills attempted to refute our critique of his analysis and defend his earlier results. His effort fails on several counts.

1) In spite of larger than normal growth in GDP, recent historical data on electricity use show a slowing of electricity demand growth for the U.S. as a whole, not the acceleration of growth that would be required if Mills' thesis was correct.
2) Mills has yet to offer a detailed point-by-point rebuttal of our critique. The few specific points he raises in the testimony are based on inconsistent comparisons or incorrect assumptions.
3) Mills repeats incorrect statements he has made about physical quantities about which we have many measurements (like the active power use of a PC plus monitor being 1000 W, when it is 150-200 W).
4) Our detailed calculations show that electricity used for all office, telecommunications, and network equipment (including electricity used to manufacture the equipment) is about 3% of total electricity use in the U.S. This estimate is about a factor of four lower than the 13% claimed by Mills.
5) While office equipment is clearly a large enough end use to warrant further study, it is not large enough to support the sweeping policy statements Mills continues to make.
6) Mills incorrectly implies that growth in information technology *must* lead to an increase in electricity use. This link is not demonstrated, and there are many factors that could also lead to a decrease in electricity use from application of information technology.

In summary, Mills' thesis has been decisively falsified, and his assumptions shown to be erroneous. It is inevitable that small errors or differences in interpretation will arise in any complex calculations. However, in this case I have identified selectively misleading examples, unsupported conclusions, and a consistent and willful series of large errors (factors of four to twenty five) all in the direction of overestimating electricity use. It is therefore time to put this urban legend to rest.

Note: Koomey's annotations are in the following font size and style (**Times, 12 pt, Bold**), to distinguish them from the Mills testimony (**Courier, 10 pt, sometimes bold** sometimes not).

"Kyoto and the Internet: The Energy Implications of the Digital Economy"

Testimony of Mark P. Mills
Science Advisor, The Greening Earth Society
Senior Fellow, The Competitive Enterprise Institute
President, Mills-McCarthy & Associates, Inc.

before the
Subcommittee on National Economic Growth, Natural Resources, and Regulatory Affairs
U.S. House of Representatives

February 2, 2000

Thank you, Mr. Chairman and distinguished Members of the Subcommittee for inviting me to speak about the energy implications of the Digital Economy. We live in a special time. It is perhaps not a totally unique time in historical terms, but it is a rare one. Times of major inflections in technology, infrastructure and the economy occur only episodically in history. I am not alone in the belief that we are only at the beginning of one of those powerful inflections, driven by what has been broadly termed the Information Revolution. The Internet is a central part of that revolution and it has only just begun to effect profound changes in our economy.

There have been many attempts to attach numbers to chronicle the growth of the Internet over this remarkable past decade. The number of people accessing the Web has grown from thousands to tens of millions. Web sites have grown from practically none to millions. Computers sold annually have risen from tens of thousands to tens of millions. Digital traffic is measured by prefixes formerly reserved for astronomers; not megabytes, or gigabytes, but petabytes. Still, traffic on the Web is doubling every several months. The entire telecommunications industry as been upended, rebuilt and expanded by the digital revolution. Commerce on the Web has exploded from nothing to tens of billions. New companies, new kinds of equipment, new services appear in a continual flow. Employment in Information Economy jobs has risen from thousands to millions. The real growth by any of these measures has been so astonishing that even the hyperbolic language of headline writers appears understated by comparison.

Against this backdrop, last year I put forth a simple proposition with a colleague that has created some controversy. The proposition is really quite simple. The Internet is using a lot of electricity, and it will use even more in the future.

The currency of the Information Economy, digital bits, are themselves simply bundles of electrons. Every single one of the hundreds of millions of devices, PCs, routers, servers, transmitters and so on, have exactly two kinds of connections: one for bits and one for kilowatt-hours. Just how much electricity does the Internet use? We think something like 8% of the nation's electric supply is absorbed by the sprawling and deeply penetrating hardware of the Internet. And when the broader array of all computers and related equipment are considered, in other words the heart of our new Information Economy, the total probably reaches 13% of all U.S. electricity consumption.

These ideas have been previously submitted to this Committee for the record. The basic concepts are set forth in my report for the Greening Earth Society, "The Internet Begins with Coal," (available at www.fossilfuels.org) and an article published in *Forbes* magazine (5/31/99) with my colleague Peter Huber, a Senior Fellow at the Manhattan Institute.

Subsequently, two respected research organizations and a number of environmental activists have exhibited alarm at the proposition that the Internet uses large and rising amounts of electricity. Before addressing the counter claims, and their deep flaws, I should like to consider the broad context for my analysis to lend perspective to the energy requirements of the Internet.

Let's be quite clear: The alarm my colleagues and I raised is at the incorrect data and flawed assumptions used in Mills' analysis, not at whether "the Internet uses large and rising amounts of electricity". I am concerned that many people and institutions are treating his analysis as technically credible when it so blatantly conflicts with measured data.

In addition, by claiming here that he will show the "deep flaws" in our analysis, he implies that these deep flaws exist and that our calculations are not correct. As shown below, his attempt to refute our work is anecdotal, incorrect, and incomplete, and he has not shown that our calculations were flawed in any way.

Although Mills' analysis is not technically credible, his numbers live on. He continues to make incorrect statements even after they have been conclusively refuted. His numbers continue to be cited by an unsuspecting press, and by some people who ought to know better. In the face of such a refutation, a responsible scientist would correct his analysis and seek to clear the record of his errors. Mills has not done so.

The Internet's Energy Transformation

If the U.S. Department of Commerce is correct, and I believe it is, in concluding that the Information Technology (IT) sector accounts for at least one-third of all GDP growth, then any policy issue that impacts IT must be considered with great caution. Energy policy is just such an issue. Because the explicit and implicit provisions of the Kyoto Protocol would directly impact every aspect of the nations' energy supply, it is appropriate, in fact critical, to consider the energy implications of our emerging Digital Economy.

Energy underpins any economy, in effect because of the laws of physics. Put simplistically, you can't get something for nothing. The Internet has not changed the laws of physics. Even cyberspace has an energy cost. Energy will continue to underpin our economy in the 21st century, just as it did in the 20th. But there will be one difference. In energy terms, the last century belonged to oil. This one belongs to electricity. Oil will not lose its prominent role, but it will take – and indeed, already has taken – second place to kilowatt-hours.

The dawn of the last century saw an explosion of economic activity in the creation of the automobile age. Investors and Wall Street rode chaotic markets investing in new companies. For technology historians, and Wall Street speculators, the dawn of the auto age has important analogs to the dawn of the Digital Age. One consequence of the rise of the automobile was the creation of an enormous and complex oil-related industrial infrastructure to fuel engines in all kinds of vehicles. The engine of the Digital Age is the microprocessor. Its fuel is electricity. Digital bits are bundles of electrons. The billions and even trillions of bits of data created and routed are, perforce, supported and energized by billions of watts. There's no getting around it. Cyberspace, far from virtual, is very real and anchored in electrons. Thus, the Internet, the central driving force of the Digital Age, is both driving and reshaping the electric infrastructure.

The transformation is already in evidence. Our economy today spends four times as much purchasing electricity as oil. This is a profound reversal of the economic positions of oil and electricity 25 years ago. The only basic energy policy that makes sense in this new Digital Economy is to ensure an expanding supply of ever lower cost and ever more reliable electricity, especially considering the trends of the past decade which have been characterized by a dominance of the tools of the Information Economy.

During this past Digital Decade, consumption of electricity has risen by 650 billion kilowatt-hours. For perspective, this growth alone required more new U.S. electric supply than exists in all of Central and South America.

The increase in kilowatt-hour use occurred despite billions spent by federal and state governments and utilities to reduce electricity growth, and despite dramatic improvements in the efficiency of electric appliances, lights and motors. It occurred, I submit, in large part because of the new tools of the Digital Age.

Considering that coal supplied about one-half of the additional electricity over the past decade (about 10% from natural gas), it is easy to see the collision course this trend has with Kyoto-inspired energy policies which are explicitly and implicitly directed at reducing coal use as well as electric consumption.

The implications of this section are ostensibly why Mills did his calculations in the first place: to show that electricity is critical to the growth in the economy, and that any attempt to restrain electricity use would hurt the economy. Conveniently for his sponsor, an offshoot of the Western Fuels Association, he further implies that any restrictions on coal use would prevent us from meeting electricity demand, thereby damaging the economy.

The claim that reducing electricity use would hurt the economy is just a variant of the belief in the iron-clad link between energy and GDP that was decisively falsified

in the mid 1980s. That belief was based on an incorrect understanding of how GDP is generated. Energy is just one factor of production, and other factors of production can be substituted for it with little or no effect on GDP, given enough time for the economy to adjust to a new equilibrium. The same is true for electricity and GDP. In many cases, capital or information can be substituted for energy or electricity and result in a net increase in GDP (for details, see Krause et al. (1993) and http://www.ipsep.org)

The implication that restrictions on coal use must prevent us from meeting electricity demand is also incorrect. There are many studies showing that coal use can be substantially reduced while simultaneously meeting electricity demands, reducing criteria air pollutant emissions, and saving money for society as a whole. There are sufficient untapped energy efficiency, natural gas-fired cogeneration, distributed generation, and renewable power options to displace large amounts of coal in the medium to longer term (for details, see Interlaboratory Working Group (1997) and http://enduse.LBL.gov/Projects/GHGcosts.html).

While there is no dispute that the Internet's equipment uses electricity, the key issue is how much. Mills claims the amount is large and rapidly rising. The factual evidence contradicts this claim.

The first piece of evidence that Mills' claims are not borne out in the data was cited by the Administrator of the Energy Information Administration (Jay Hakes) in his testimony, also delivered at that same hearing on 2 February 2000:

> **"From 1985 to 1995, retail electricity sales grew at a rate of 2.6 percent per year…Since 1995, the use of the Internet has increased dramatically, yet retail electricity sales have grown by 2.1 percent per year, 0.5 percentage points less than the previous 10 years."**

If the Internet were causing the amount of electricity demand growth claimed by Mills, we would expect this growth to show up in the aggregate statistics. In reality, electricity demand growth has slowed in recent years, in spite of the fact that GDP has increased at a greater average annual rate than in the preceding period. This disparity between the aggregate energy data and Mills' claims provides evidence that his thesis is invalid, but it will take several more years of energy data to provide positive proof. Fortunately, our dissection of Mills' assumptions provides such positive proof, without having to wait for the aggregate statistics to roll in.

The Internet & Electricity Demand

Just how much of the nation's electricity demand is a direct result of equipment in the Digital Economy, and more specifically, the Internet? Truth be told, it is hard to draw a bright line between many devices used for the Internet and those that are part of the broader Digital Economy. Nonetheless, we made just such an attempt, precisely because the Internet is at the epicenter of the Digital Revolution.

It is not just hard to draw such a "bright line," it is not possible, nor is it worth doing. The correct approach is to estimate the stocks, power levels, and usage patterns for all office and network equipment in the U.S., regardless of their function. As our detailed technical response to Mills describes (Koomey et al. 1999), trying to focus just on the Internet-related portion of this electricity use virtually guarantees large calculational errors.

It would be exceptionally challenging to catalog all the wide array of devices that comprise the Internet and Digital Economy. Instead, we chose a technique known as sequential approximation. This well-established technique permits one to gain a reasonable order-of-magnitude estimate of a complex factor without a detailed inventory. One can, for example, use sequential approximation to estimate the number of people in a stadium by considering an inventory of hot dogs and soft drinks. Some approximations are required, but the outcome will be in the right ballpark.

Ballpark estimates are accurate provided that the underlying assumptions are accurate. Our analyses demonstrate conclusively that Mills' assumptions are not accurate, so his calculations cannot be correct.

The ballpark estimate: the Internet in all its facets, likely consumes 290 billion kilowatt-hours, or about 8% of the U.S. electric supply system. The broader category, the entire array of all types of computers and computing-related devices (such as storage systems), in homes, businesses and factories which fuel our Digital Economy likely uses 13% of all the nation's electricity.

This last paragraph contains the numbers that have been cited in many places. As I discuss below, they are incorrect even after adding the energy needed to manufacture chips and computers.

These numbers encompass much more than PCs on desktops. One must include for example all the hardware behind-the-wall in the telecommunications and Internet networks which includes, but is far from limited to, such things as routers, the hardware of the dot-coms such as servers, and even the silicon and PC factories. Determining Internet and Digital Age electricity use requires collecting and assessing data across many sectors and boundaries.

It is clear that traditional data sources and methodologies are not adequate to the task of clearly tracking the electric needs of the Information Age. For example, most of the necessary data for the commercial sector is invisible in traditional Energy Information Administration energy accounting. EIA does report on PC electric use in commercial buildings, but all of the other types of information technology hardware (which comprise over three-fourths of Internet energy use) are thrown into a general grab bag category called "other." EIA notes cryptically that "other" includes telecommunications equipment. The data lost in "other" was irrelevant two decades ago at the dawn of the Digital Age. Today, the "other" category of commercial electric use is over 300 billion kilowatt-hours and is greater than all other categories except lighting – and will soon overtake lighting.

Mills incorrectly implies that electricity use in the "other" category is largely driven by information technology. The "other" category in the Annual Energy Outlook forecast contains far more than telecommunications equipment. It also includes service station equipment, district heating and cooling, medical equipment, and automated teller machines. In addition, it includes a correction factor that makes the AEO's estimate of total electricity use for the commercial sector match exactly to the total electricity use reported in the EIA's State Energy Data report for historical years. This correction factor represents just under two-thirds of the commercial "other" category, and it is an accounting factor that has nothing at all to do with information technology.

Before addressing a few points of contention regarding my estimate of 290 billion kilowatt-hours for the entire Internet, it is useful to ask first, is such a result in the ballpark? Much of the confusion and controversy surrounding the issue arises from a key question. In effect, how much of the electric use of a PC (or any IT equipment) is directly attributable to the Internet? Since the Internet is an integral subset of the Digital Economy, the easiest sanity check would be to evaluate the electric needs of all Information Technology equipment. For example, how do you count computers used to develop software for the Internet if those PCs were not directly plugged into the Web? Clearly they are part of the bigger picture, the entire Internet-driven Digital Economy.

A useful starting point for a ballpark check is in the simple fact that the U.S. Information Technology industry sold over $400 billion worth of hardware last year. Over the past three years alone, more than $1 trillion of IT hardware has been installed. This hardware represents the engine of the new Digital Economy. Much of it becomes part of the Internet, most is driven by the Internet. Every single piece of this $1 trillion in hardware gets plugged into a wall somewhere.

The data on dollars spent on computer hardware is at best peripheral to the question at hand. For example, the money spent on mainframe computers last year is larger than ten years ago, but the power use per mainframe computer has gone down by more than a factor of two, because of the shift to CMOS technology in these machines. Cost has nothing to do with the per unit energy use for these machines, to first order. Of course, if more computers are operating, aggregate electricity use may increase or decrease, depending on whether the increase in numbers of units is sufficient to offset the decline in energy use per unit. Our comprehensive analysis accounts for this effect, and I describe those results below.

There's another more specific ballpark check available from the year 1993, the Jurassic Era of the Internet. A 1995 Lawrence Berkeley Labs (LBL) study (the most recent on the subject) reported about 50 billion kWh in 1993 for commercial sector use by PCs, computers and directly related equipment such as monitors and printers.

This 50 TWh includes PCs, monitors, laser printers, serial printers, mainframes, and minicomputers used in the commercial sector.

This 50 billion kWh figure for the commercial sector from seven years ago is a good starting point for the Digital Decade. Let's consider what's happened since then.

1) the number of PCs and related equipment in offices has exploded
2) the number of PCs in homes, schools, everywhere, has also exploded
3) the Internet has burst on to the scene, with all its back-office Web and telecommunications hardware
4) an entirely new class of businesses has been created; the dot-coms
5) the usage level for all computing and IT equipment is up everywhere

These statements need to be quantified. What precisely constitutes an explosion? What if the number of PCs "exploded" but the electricity used by PCs per unit went down significantly? Does the network equipment for the Internet use significantly more power than the network equipment that existed in 1993? What does it mean when Mills says that the "usage level of all computing and IT equipment is up everywhere"? Without quantitative estimates for each of these assertions, they are just rhetoric.

I am quite confident that these factors collectively have brought the 50 billion kWh starting point in 1993 up to my estimates for the broader Internet (i.e., beyond the commercial sector alone) and the Digital Economy today.

It is incorrect and misleading to compare the 50 TWh estimate for the commercial sector to the estimate for the current-day Internet. A correct comparison would tally the electricity used by office equipment and network equipment in all sectors in 1993, and compare it to electricity used by equipment in the same categories in 1999. Without such consistent numbers, the comparison Mills attempts to make is meaningless.

And if we're not quite there yet, just wait a few more months.

This statement is misleading. Mills wants to demonstrate that electricity demand from these devices is growing at a furious pace, and has not done so. He uses these vague statements to create an impression that sounds plausible, yet this impression is not based on credible data or sound analysis.

There are some other useful ballpark indicators. The Information Technology Industry Council's (ITI's) tracking shows the total inventory of computers and computer-type equipment has jumped by at least 100 million units since 1993. The inventory is growing at over 40 million a year now. And their data set specifically does not include such Internet equipment as routers, which are functionally computers. Cisco sells about a million routers a year. Nor does the official data track the number of wireless base stations, amplifiers, ports, hubs, information appliances, and so on. All of these have grown rapidly over the past Digital Decade. All of these devices use electricity. Many are already part of the Internet, and those that are not will soon be.

The number of computers *sold* as reported in ITI's data book is correct, but these estimates appear not to account for retirements of old machines at the levels that I would expect given the rapid obsolescence of computers. Our own stock estimates

are significantly lower than ITI's because our analyses correctly account for retirements. The U.S. inventory is not "growing at over 40 million a year now". That's the statistic for annual sales, which does not account for retirements.

Yes, these devices all use electricity, but almost without exception, they use far less than what Mills asserts they use. For example, most Cisco routers use about 100 W (Mills assumes 1000 W for each and every router, which is a major overestimate).

And this is only part of the story. One must also add the electric needs of the semiconductor, PC and IT manufacturing industries. Semiconductor manufacturing alone has grown in the past half-decade to become the nation's largest manufacturing industry. Silicon plants are the steel mills of the 21st Century. Their fuel of choice; kilowatt-hours.

When you think about it, it is inconceivable that the Digital Age and the Internet, do not already account for a significant and growing share of the nation's electric supply.

Mills uses anecdotes to make his case. Instead, he needs to demonstrate the complete accounting for all office equipment and network equipment in the U.S. I describe the results of such a complete analysis below, based not on arbitrary assumptions but on careful, empirical, peer-reviewed data and measurements.

<u>The Case for 1%</u>

Two organizations have offered rebuttals to the 8% estimate for the Internet's share of national electric use. I believe it important to address these ostensible rebuttals given the importance of this issue to federal energy and economic policy. There is insufficient time here to address all of the details, but a few observations are instructive.

The researchers at Lawrence Berkeley National Laboratory (LBL) have published a superficially **[sic]** analysis of my study "The Internet Begins with Coal." Before addressing a couple of representative examples of the inherent failures in the LBL rebuttal, there are two over arching points that should be made. The first relates to the failure of LBL to step up and take an honest crack at estimating an answer to the core question. The second relates to the strange failure of the LBL team to seek information to clarify their misunderstandings.

First then is the fact that LBL team and others seem preoccupied with rebutting details of my analysis, but are quite unwilling to make their own independent estimate to answer the central and critical question: how much electricity does the Internet use? My recommendation to the LBL team then and now: please undertake a detailed and intellectually honest ground-up analysis of the Internet's electric needs.

I am unwilling to assess how much energy the Internet uses because that is fundamentally the wrong question: it is futile, and irrelevant to policy, to try to estimate how much of the energy and time used by data-processing equipment is Net-related and how much is not. The correct way to analyze the situation is to

examine ALL computers and related equipment as in our most recent analysis, finished in June 2000 (Kawamoto et al. 2000). The total amount of electricity used by all that equipment represents an upper bound for how much the Internet could use if the equipment were devoted entirely to Internet-related activities, and none to any other purpose.

Since Mills is making a set of policy conclusions, it is appropriate to assess the details of his analysis. I am "preoccupied with rebutting details" in his analysis because those details reflect upon Mills' core technical credibility. If he has done analysis using a large number of assumptions that are incorrect (and that all seem to push the analysis towards overestimating electricity use), the people listening to his policy pronouncements ought to find that relevant when assessing his arguments.

The central conclusion of the LBL paper is that 8% is an overestimate of the Internet's use of U.S. electricity by "a factor of eight."

On learning this, I asked the LBL team the obvious question, if you say 8% is an overestimate by a factor of eight:

> "May I quote LBL as claiming/believing/estimating that the Internet uses 1% of the nation's electricity supply?"

Their answer, in full:

> "You may NOT quote LBNL 'as claiming/believing/estimating that the Internet uses 1% of the nation's electricity supply' because your estimate just focuses on direct electricity use, and not the overall effects on the U.S. economy that result from structural changes and substitution effects due to the Internet. You may quote me as believing that your estimate of the direct electricity use associated with the Internet is too high by a factor of eight, but that the NET effect of the Internet on electricity and energy use (which is really what matters) cannot be estimated accurately without assessing the associated indirect effects of the internet on resource use in the economy."

Given what I've outlined earlier, and what practically everyone who reads the news knows, the explosion in Internet equipment is quite unlikely to have led to a reduction in the use of electricity. Data contained in LBL's own research on PCs and computers yields a figure of 2% way back in 1993 and just for the commercial sector.

The statement about "what practically everyone who reads the news knows" contains an incorrect implication, namely that the increase in the numbers of operating electronic devices must lead to an increase in the electricity used by that equipment. There are two ways this implication may be incorrect:

1) ***The power use per unit of certain types of equipment can be declining more quickly than the number of units is increasing.*** **Mainframe computers are one example where power use per computer has dropped so significantly over the past five to ten years that total electricity by these devices has declined. The LCD screen is one where it will occur over the next five years (as low-power LCDs displace the more power-hungry cathode ray tubes).**

2) ***The systemic effects of E-commerce may reduce overall electricity use.* One example is telecommuting. If employees work at home in large numbers, that translates into a reduction in commercial building floor area that needs to be built. Typical office buildings use 15-20 kWh per square foot per year–they are much more electricity intensive than typical homes on a per square foot basis. A large increase in telecommuting will likely reduce electricity used in buildings because it will shift economic activity from more electricity intensive commercial buildings to less electricity intensive residential buildings. Reduced need for warehouse space brought about by "just-in-time" manufacturing (enabled by information technology) could also result in electricity savings. Nobody knows right now whether these effects are large enough to matter because the data do not exist to track them, but they could well be large enough, and they cannot be dismissed out of hand (Romm et al. 1999).**

Both of these effects are potentially important, but Mills incorrectly assesses the first one, and ignores the second.

Furthermore it is disingenuous for the LBL team to state that what really matters is the "NET" effect of the Internet. Certainly it's an interesting issue (more about this in a minute). Fax machines use electricity and displace jet fuel by replacing overnight mail. I believe I may have been the first to publish detailed analyses of this effect of electrification in 1991, and to describe this effect I coined the term "ecowatts" at that time, documenting and publishing widely to extol this important efficiency trend.

But here's a simple arithmetical fact; estimating the net savings from faxing requires, *a priori*, knowing the amount of electricity used by faxes. Accurately calculating the net savings is actually much more difficult than accounting for the electricity used. (Consider, for example, that faxing should have been expected to reduce use of overnight mail; in fact overnight mail has grown.) But LBL suggests that one should not study the use of electricity from PCs, or by inference, faxes or any office equipment "without assessing the associated indirect effects." LBL's own EPA-funded 1995 research on electricity used by all manner of office equipment in commercial buildings does not meet this test – nor should it have to.

The idea that we can or should only study and publish the "NET" effect is the equivalent of claiming that you can figure out the change from dinner without knowing how much money you gave the waiter.

The LBL team dodged the issue.

The LBNL study in 1995 reports direct use of the equipment only. It did not attempt to estimate these systemic effects, because its sole purpose was to estimate the direct use. Our focus in these analyses has always been relatively narrow: to summarize the state of our knowledge of office equipment energy use in the U.S. Our arguments only depend on our assumptions and data about office equipment energy use being accurate and complete, and so that's what we've spent our time to ensure. Our report did not speculate about the overall net impacts of office

equipment on resource use, because we did not then and do not now have the data to assess those impacts in a precise way.

The difference is that we were not making large-scale policy pronouncements in the way that Mills is doing. His policy conclusions *depend* on the net effects of office equipment not reducing or eliminating any direct electricity used by the equipment. Since he has not conducted such an analysis of the indirect effects, he is not justified in making the sweeping policy statements contained in his writings. We do not need to conduct the analysis of systemic effects, we need merely show that his assumptions and conclusions are wrong. We have done so.

Mills says "LBL suggests that one should not study the use of electricity from PCs, or by inference, faxes or any office equipment 'without assessing the associated indirect effects.'" I would never make such a statement. In fact, I've devoted years of my professional life to just such assessments of direct electricity consumption of energy-using devices, without knowing all their indirect effects. I would amend that statement to read "one should not reach sweeping policy conclusions dependent on the total net electricity use associated with any type of office equipment without assessing the associated indirect effects of that equipment on fuel and electricity use. Assessments of the direct electricity use of that equipment can be useful, but cannot by themselves be used to justify sweeping policy pronouncements about climate change policy and electricity demand growth."

In short, Mills leaps from erroneously high estimates of electricity used by data-processing equipment to broad conclusions about the need for more electricity and more coal, without assessing how the Internet may affect other uses of electricity. My colleagues and I correctly assessed the electricity used by data-processing equipment (and found it to be eight times smaller than Mills claimed); did not attempt the futile task of estimating what fraction of that usage is related to the Internet; and noted qualitatively that other Internet-related effects may dampen or reverse growth in electricity used for certain other purposes, making the net effect of Internet traffic growth and E-commerce on electricity demand indeterminate based on present evidence.

I responded to Mills' question in the way that I did because 1) I think the question of how much electricity the Internet uses is the wrong question (as described above) and 2) I had uncovered so many errors in the initial examination of Mills' analysis that I did not want to report my conclusions except relative to his results. Reporting an absolute number (like 1%) implies a greater level of precision than is justified when the analysis itself was so shaky. I was (and am) quite confident that Mills' analysis overestimates the electricity used by a certain segment of office equipment by at least a factor of eight. This is the only relevant result from our analysis.

The second generic point I should like to make arise from the following statement from the LBL paper:

> "Mills' report does not contain enough detailed documentation to assess the reasonableness of <u>many</u> assumptions." (emphasis added)

This is a fair complaint. I note for the record that the LBL team, in full possession of my e-mail, phone number and address, and despite a couple of very general e-mail exchanges with them, made absolutely no attempt to contact me to obtain clarification or expansion on specifics for any assumptions. Considering that clarification was and is necessary for "many assumptions," their failure to do so leaves one wondering if they did not want clarification, and that the rebuttal was motivated by something other than the requirements of technical scholarship.

We analyzed those key assumptions that were sufficiently well defined that they did not require clarification from Mills. In addition, we had the benefit of being copied email exchanges among Mills, Joe Romm, and Amory Lovins, which clarified some of the issues, but left many more unexplained. The assumptions that were not well documented enough in Mills' report we largely ignored. Addressing the errors in just the assumptions that *were* well documented kept us quite busy. We had limited time, and our purpose was to conduct a technical critique of the key assumptions supporting Mills' thesis as quickly as we could.

That the LBL team has, so far, dodged the central question is clear. Thus far their only contribution to this debate has been an attempt to cast doubt on my analysis.

Our critique of Mills' analysis (Koomey et al. 1999) stands on its own. The more detailed and complete analysis (Kawamoto et al. 2000) took another six months, and was not undertaken until we had secured enough funding to do a careful job. Once completed (four months after the testimony was given), it conclusively supported the findings of our technical memo of December 1999, and refuted his claims.

The LBL rebuttal contains numerous serious errors. Let me briefly outline two that are representative.

Mills has to date failed to offer a point-by-point response to our criticisms. Until he identifies our alleged errors, this point is unsubstantiated. If he (or any other reviewer) identifies any errors, we will be happy to correct them.

The first technical point: In the LBL paper, the authors take issue with the claim that the desktop for an Internet-configured PC (i.e., including necessary peripherals) is about a 1,000 Watt device. Setting aside the question of whether it is 1,000 Watts (it is), the LBL researchers know full well that the relevant number used in the calculation is NOT the peak watts, but the quantity of kilowatt-hours used in a year. In analogous terms, what really matters is how much gasoline you use in a year, not the horsepower of your engine.

The statement about the power use of a PC being 1000 W has no empirical basis. As documented in our memo (Koomey et al. 1999), the active power use of a typical computer plus network equipment is in the range of 150 to 200 W, but no more than that. When power management is included (as it is for many PCs sold today) the computer reduces its power use to 60 W or less in standby mode.

In this regard, my analysis for an Internet-configured PC is based on **750 kWh/yr** and is consistent with many other analyses, including their own at LBL.

The footnotes in Mills' Greening Earth Society report make it clear that he used 1000 W for each PC, plus different usage assumptions for home and office PCs. If he based his estimate solely on 750 kWh and not the power level and the hourly usage, then he should not have stated in his documentation that he used 1000 W.

In any case, the energy use per PC figures shown below are for computers used for all purposes, not just those "associated with the Internet". Mills explicitly says that he assumes 12 hours/week for Internet related usage of office PCs, which averages out to be about one-third of the work week. All of the estimates below must therefore be reduced to one-third of their level stated below to make them consistent with Mills' estimates.

• In their 1995 study, LBL finds that a PC and printer uses **650** kWh/yr
("Efficiency Improvements in U.S. Office Equipment: Expected Policy Impacts and Uncertainties," LBL, December 1995, p. 15.).

In the 1995 study, we estimated that a typical PC in 1990 with a monitor PLUS one laser printer uses 650 kWh per year in an office environment. The problem, of course, is that laser printers are shared among many computers, so that forty percent of that 650 kWh (i.e. 250 kWh) must be split among 10-20 computers, yielding 12.5-25 kWh per year per PC for the electricity used by the laser printer.

The computer and monitor together use about 400 kWh, which also happens to match up with our current best estimate for electricity used by current PCs and monitors in an office environment (the comparable number for a home PC is a little over 100 kWh/year). This total is the electricity used by the computer and monitor for all purposes, and is based on the latest measured data. Now we need to reduce this usage to reflect the amount "associated with the Internet" in Mills' calculation.

How much printer use is "associated with the Internet"? It isn't apparent that this question has any meaning, but we'll use one-third to make it match up with Mills' usage of "PCs in offices". Multiplying the PC, monitor, and printer total usage of 425 kWh by one-third, we get 140 kWh, which is more than a factor of five lower than Mills' estimate of 750 kWh. Our work does not in any way support his estimate.

• In an unrelated 1995 EPA study, annual PC electric use was estimated to range from **450 to 2,000 kWh/yr.**
("The Green PC," S. Anzovin, Windcrest, 1994, p. 5).

• A more recent National Academy of Sciences (NAS) report put annual PC/workstation electric use at **1,000** to 1,800 kWh/yr.
(IEEE Spectrum, January 2000).

The upper end of these ranges (1800 to 2000 kWh/year) is for computers and monitors left on 24 hours per day, 7 days per week. There are relatively few

computers that operate this many hours, so these estimates are irrelevant to this comparison.

The lower end of the range from the EPA study is consistent with measured data for total electricity use of PCs and monitors, but we still must divide by three to make these numbers comparable to Mills' estimate of 750 kWh. This calculation yields an estimate of 150 kWh, a number that is consistent with our own estimate above.

The lower end of the range from the NAS report is not consistent with measured data (it is too large), but we ignore that for now. Dividing by three to make that number comparable to Mills' "Internet-related" estimate of 750 kWh yields 333 kWh/year. Even this estimate, which is demonstrably higher than the measured data, is a factor of two lower than Mills' estimate.

Despite the readily verifiable above noted facts, the LBL paper nonetheless concludes that "With these corrections [to Mills' assumptions], PCs in offices use about 7.2 TWh, a reduction of 84% from Mills' estimate."

Surely the LBL team noticed the bizarre inconsistency in this conclusion. Their own 1995 seminal study showed collective commercial sector PC electric use at 50 TWh more than five years ago. How could their "correction" to my analysis yield 7 TWh today?

Our 1995 study showed electricity used by PCs and monitors in the commercial sector to be 17 TWh. The 50 TWh is not comparable to the 7.2 TWh, because it includes printers, mainframes, and minicomputers, not just "PCs in offices" (which is the category of equipment for which I estimated 7.2 TWh in our technical memo). Mills assumes that "45% of new/purchased business PCs are connected directly to the Internet". Multiply 17 TWh by 45% and you get 7.7 TWh. Multiply again by one-third to account for only the "Internet related" portion of the hourly use. This calculation yields 2.5 TWh for Internet-related usage (based on Mills definition of this) for commercial-sector PCs and monitors from our earlier study. The 1995 analysis is perfectly consistent with the later assessment.

Our more detailed study for 2000 (conducted after the testimony was completed) shows 22 TWh for PCs and monitors in the commercial sector. When multiplied by 45% and 33%, we get 3.3 TWh. We can add in the energy used by industrial PCs, but when multiplied by the same factors, we get to 4 TWh for all PCs in commercial and industrial businesses (as distinct from servers and minicomputers). I conclude from this calculation that my original estimate of 7.2 TWh for Mills' category of "PCs in offices" (Koomey et al. 1999) was too generous by 3 TWh. It should have been 4 TWh instead of 7.2 TWh, a reduction of a factor of eleven from Mills' initial estimate.

Let me turn now to a second example of poor analysis in the LBL paper, but of a slightly different 'flavor' of error.

One entirely new category of computer use since 1995 is in Web servers. Servers are really computers ranging in type from PCs, to workstations, and up to mainframes that host the Web sites. Servers run 24-7 and are frequently arranged by the hundreds in enormous banks of racks creating a "server farm" for mid-sized to large Internet Service Providers. LBL claims that we need to adjust downwards both the power used by servers and the total number of servers. The power use issue for servers is essentially the same as I've just outlined for PCs.

There are no measured data indicating that Mills' estimate of 1.5 kW (Mills 1999) is anywhere near the true power level of these machines. A four processor Pentium III machine we measured used about 270 W, and single and dual processor Pentium machines we measured (which are more typical of servers in use) consume much less than that (50-100 W). In addition, most servers do not have dedicated monitors, so it is impossible that 1.5 kW is an accurate assessment of average server power use.

At the time of writing my report, I used an estimate of 4 million servers for 1999 based on an extrapolation from data for the number of Web sites. The LBL team 'adjusted' my estimate arbitrarily to conclude that the "correct" number of servers should have a downward correction of "80%" to 1 million. LBL could have undertaken some modest additional research, as I did subsequently, to learn that there is hard data on the number of servers in operation in 1999 that does not require any extrapolations. The actual number of servers last year was 4 million (Netcraft Internet Survey, www.netcraft.com/Survey/Reports/). Clearly my methodology was more accurate than theirs. As a point of interest; there were fewer than 20,000 servers in 1995. Servers are only one piece of a very big digital pie, but quite indicative of the electric trends.

The source Mills cited for web servers in 1999 (and the one he cites above from his "modest additional research") contains numbers for web sites, and he therefore incorrectly equated web sites with computers acting as web servers. One server computer can serve many sites (at one ISP we visited, each server machine hosted hundreds of web sites).

Here's what the Netcraft methodology page says about this issue:

> **Many servers support facilities *to enable a single computer to run a server for multiple domain names on different ip addresses (italics added)*. Additionally, some service providers offer a more crude domain aliasing facility with multiple hostnames resolving to the same ip address, and customers home pages being referenced by a trailing pathname. *This survey counts each of these domain names as separate servers (italics added)*. (http://www.netcraft.com/Survey/mechanics.html)**

Mills' mistake arose because of confusing terminology in the world of the Web: each time a particular type of server software is used for a web site, it is tallied by the source he cited as one instance of this type of "server". This definition of the word "server" applies to a single instance of the software, not to the hardware

"server" that hosts many different web sites (and hence hosts many different "servers" in the software sense of the term).

In general LBL sought to ignore the basic methodology I used, sequential approximation, and instead clearly sought to undermine the integrity of my work, without attempting their own honest analysis.

We analyzed Mills' assumptions and clearly documented where we thought they were in error, but we did not need to create a complete analysis of all electricity used by office equipment in the U.S. to show that Mills' analysis contained serious analytical flaws. That goal we accomplished with our technical memo.

We subsequently conducted a complete, detailed, and peer-reviewed analysis for direct electricity used by office equipment (Kawamoto et al. 2000), and it validates our initial critique of Mills' work.

They also failed to note the explicit mention in my report that we did not count the electric use of a wide variety of other relevant Internet related devices, totaling in the millions.

Mills' suggests that millions of other devices should also be included, but fails to estimate their electricity use. In a recent speech he held aloft a wireless Palm VII-type handheld computer and claimed that it used as much electricity as a refrigerator—because of the large amount of unspecified Internet-related equipment to which it was allegedly linked. I have not yet seen the detailed accounting upon which this claim is based, but am unable to reproduce it using the measured data we have on hand. Without quantification and documentation, such claims should carry little weight in this discussion.

The LBL researchers are right about one claim, that it is difficult to cleanly separate Internet equipment from all information technology equipment. Thus, I asked the LBL team to consider the conclusion offered in the study and the *Forbes* magazine article, that microprocessors of the Digital Age, in all categories including the Internet, consume about 13% of the nation's electricity. We have yet to receive a response.

Doing these calculations correctly and accurately takes time. Our latest work completed in June 2000 (and emailed to Mills on 7 June 2000) shows total electricity used by all office equipment (including copiers and fax machines) in all sectors (residential, commercial, and industrial) to be 71 TWh. Routers and switches add another 3 TWh. The Koomey et al. (1999) memo estimated total electricity used by telephone central offices (including cooling in these facilities) to be 12 TWh per year. Together, these numbers yield an estimate of 86 TWh for the direct use of all office, telecommunications, and network equipment, not just that "associated with the Internet".

This estimate does not include power used for manufacturing the equipment. The Annual Survey of Manufacturers from the U.S. Census (http://www.census.gov/prod/www/abs/industry.html) reports on electricity use for

1992 to 1996 by industry group, identified by SIC codes, and we summed up electricity use for the relevant industry groups. We were able to use these data to create an upper bound to the amount of electricity used to manufacture office equipment and network equipment in the U.S. for all purposes.[1]

Table 1: Electricity used to manufacture equipment for the digital economy (TWh)

Sector description	*SIC Code*	*1992*	*1993*	*1994*	*1995*	*1996*	*% of sector applicable*
Electronic components and accessories	367	14.6	15.5	15.7	15.8	16.7	75%
Computers and office equipment	357	5.1	4.5	5.0	4.5	4.2	100%
Switchgear and switchboard apparatus	3613	0.6	0.7	0.6	0.7	0.6	50%
Telephone and telegraph apparatus	3661	1.4	1.4	1.3	1.3	1.1	100%
Communications equipment (not identified)	3669	0.2	0.2	0.2	0.2	0.2	100%
Storage Batteries	3691	1.2	1.3	1.5	1.5	1.5	50%
Primary batteries dry and wet	3692	0.3	0.3	0.2	0.2	0.2	50%
Magnetic and optical recording media	3695	0.9	0.9	1.0	1.0	1.0	100%
Total		24.3	24.9	25.5	25.3	25.6	
Total of applicable energy use		19.6	19.8	20.5	20.2	20.3	
Index (1992 = 1.0)		*1.00*	*1.01*	*1.04*	*1.03*	*1.03*	
% average annual growth since 1992		*0.0%*	*1.0%*	*2.1%*	*0.9%*	*0.9%*	

Source: Annual survey of Manufacturers 1992-1996 (US Bureau of the Census 1998).

We included only 75% of the energy for SIC 367 (Electronic components and accessories) and 50% of the energy for SIC 3613 (Switch gear and switchboard apparatus), SIC 3691 (Storage Batteries), and SIC 3692 (Primary batteries dry and wet) because there are many other products for which these electronic components are used. For example, all new cars have a storage battery and at least one computer in them for engine control. Many household appliances have batteries and circuit boards as well. These estimates are rough, but they allow us to create an upper bound.

After summing the electricity use in the categories described above for the 1992 to 1996 period, we found electricity demand growth of 0.9% per year during this period, and total electricity use of about 20.3 TWh in 1996. If we extrapolate this growth to 1999, the upper bound for total electricity used for manufacturing the relevant electronic equipment is roughly 21 TWh. The actual total is almost certainly lower than this, because our assessment of the percentage of each SIC code that is applicable is overly generous, as a conservatism.

---

[1] There is a substantial body of literature on life-cycle assessment (LCA), which attempts to calculate the total electricity and other resources used to create (and use) industrial products. I relied on some of that literature in the 1999 technical memo, but because LCA studies only exist for a few types of products, I felt more comfortable using the aggregate electricity statistics from the Survey of Manufacturers for the calculation above. We will attempt to reconcile the two approaches as more data become available.

Based on these calculations, the direct use of office, telecommunications, and network equipment for all sectors in the U.S. is responsible for about 2.6% of all electricity use (86 TWh/3300 TWh). Electricity needed to run electronics manufacturing facilities to manufacture that equipment adds (at most) another 21 TWh, or 0.6%, which brings the total to 3.2%. This result is about a factor of four lower than Mills' estimate for the electricity demand of "the digital economy." We have thus demonstrated that Mills vastly overestimated electricity associated with computers and other office equipment in the U.S.

The Case for Zero

While the LBL team dodged the specific question of how much electricity the Internet or even the Digital Economy uses, a Cool Companies study led by Joseph Romm was braver. The Cool study has two central contentions that merit brief discussion. One contention, incredibly enough, is that the Internet's electric use is zero. And the other central contention is that the efficiency gains from the Internet offset any putative energy needs. Let me briefly address these two contentions.

The Cool study conclusion about the Internet simply and astoundingly concludes:

> "The authors found that the Internet itself is not a major energy user, largely because it draws heavily on existing communications and computing infrastructure."

This observation reflects such a deep misunderstanding of the telecommunications revolution that it is difficult to know how to respond. Just what exactly do the authors think the past half decade of over several trillion dollars in new investment in telecommunications and computing equipment has been for and driven by, if not the Internet?

The exponential growth in equipment (and related Wall Street valuation) constitutes the electric-intensive infrastructure of the Internet. None of it was "existing." Equally important, it is still rapidly expanding. The entire telecommunications industry has been visibly up-ended and expanded by the Internet. The purchase and installation of hundreds of thousands of miles of fiber optics, and the entire attendant infrastructure has been almost entirely driven by the Internet. Digital traffic now dwarfs voice traffic on the telecommunications networks. And every telecom expert forecasts traffic to grow, and for the growth to be utterly dominated by bits, not voice. The driving force for bits is the Internet.

Equipment installations and stock valuations do not necessarily imply more electricity use for society as a whole, because new equipment may be replacing more energy intensive older equipment, and the new services delivered may displace other more energy intensive services. It is misleading to suggest that the growth in the use of computers and the Internet will necessarily result in a growth in electricity use.

Growth in data traffic does not necessarily imply growth in electricity use. We visited a central office of a major phone company in California, and I asked the

manager of the facility how fast his data traffic was growing every year. He responded "Doubling or quadrupling". Then I asked him about the facility's electricity use over the past five years, and he replied "It has not grown at all".

The Cool study authors would have us believe the Digital Economy is some kind of virtual overlay on existing infrastructure. This is the equivalent of asserting, in 1950, that the several decade build-out of the nation's Interstate Highway system, to support all the new cars and trucks moving into the economy, would not entail any investment (in dollars, materials or energy) since drivers would be using an existing highway infrastructure. It is 1950 for the digital highways.

But the Internet Improves Efficiency

It is widely recognized that the Internet is improving economic efficiency, sometimes astonishingly so. Indeed, this central fact is the very reason that the market is so rapidly consuming digital bandwidth and all of the equipment to create and serve that bandwidth. But economic and energy efficiency are not the same thing. Indeed, economic efficiency can fuel increased energy demand.

There are two aspects to the efficiency argument. One is macro-economic; is the general, overall effect of the Internet to reduce energy and material use? The second, micro-engineering; does the Internet reduce material and energy use in specific applications?

The Internet serves as a kind of economic lubricant. According to the Department of Commerce (Digital Economy II), information technologies drive at least one-third of the GDP growth, and further two-thirds of ALL investment in capital equipment. These results suggest the answer to an oft-posed question from economists and digital skeptics, "when will we see the putative economic effects of the massive investments in computers?" We're seeing them now. Indeed, Chairman Greenspan appears to believe that the reaction is even a little overheated.

Regardless, so far the net effect of the Digital Age at the national level has been to increase energy use. In the last digital decade, total air miles flown have risen from 4.3 to 5.8 billion a year. People are flying more than ever. Planes use fuel. People are driving more than ever, and in bigger vehicles. SUV and light trucks account for one-half of all vehicle sales – doubling in the past Digital Decade. Transportation fuel use is up 12%. Similarly, the digitally-accelerated economy has driven up the size of homes and the spending on home improvements. Whether you think this is good or bad is not relevant to the fact; so far, it has all generated greater energy use.

A robust economy tends to use more energy. To be sure, we're more efficient. But there is no evidence yet in human history, much less the past few decades, of rising economies with sustained declines in energy use. Obviously, improvements in efficiency moderates a growth rate; but the operative word is "growth."

What about the application-specific efficiency argument? The idea, in a nutshell: the Internet is so powerful that it will improve efficiency faster than the energy consumed by the hardware on the Network.

The energy used per dollar of GDP is the favorite efficiency metric of both environmentalists and business leaders seeking environmental coverage. By this measure, the U.S. is incredibly more efficient than just a decade ago. Total U.S. energy use per dollar of GDP has dropped 16% since 1990. Today's economy, operating at the energy efficiency of 1990, would need 15 Quads more fuel – in oil terms, that would be a 40% increase in total U.S. oil use. Interesting, but largely meaningless. The nation still uses more energy today than a decade ago. And more importantly for a Digital Economy, we use a lot more electricity. Increased supply to meet electric growth in just the past five years is equal to the total generating capacity of Italy.

There is a real problem with the dollars per BTU metric for energy efficiency, and easily illustrated. Considering, with this metric, who is the most energy efficient person in America? Bill Gates. Despite enormous energy use in his legendary home, personal jet and so forth, Mr. Gate's wealth yields an efficiency measured as energy per $ that would shame a Sudanese hunter gatherer – only because his wealth is so great. The economic path the U.S. is on, with the Digital Era accelerating economic gains out of proportion to relatively modest energy growth, means that the U.S. economy is following the Bill Gates method for energy efficiency: increase wealth faster than you increase energy use.

What then of the specific energy efficiency gains of the Internet, especially the efficiencies of buying "on line" via e-commerce – what might be termed Amazon-dot-com effect? The jury's still out on whether more or less energy/material infrastructure is used to warehouse and deliver e-commerce products. Books from Amazon via 747 and trucks may use less or more energy than driving an SUV to the book and grocery store. It is far from clear, however, what the final, overall effect will be in retail e-commerce, especially since it is still only a tiny fraction of total retail. The 24-7, send-it-overnight e-commerce economy could increase energy use if aircraft begin to substitute for trucks and trains for product delivery. Many analysts believe that competition in e-commerce will drive business increasingly to delivery overnight. Developers are already building new, dedicated airport hubs that can handle multiple 747s loading-unloading specifically as for e-commerce. ("Developers Rush to Meet the Demands of E-Commerce," 1/23/00, New York Times.)

Even if the net overall impact of the Amazon-effect is improved energy efficiency (it probably is in many cases, if not specifically the Amazon case)– reduced transportation oil use still comes with increased use of electricity. This in fact has been the general macro-energy trend of the past decade.

The Cool study continues also with one long-sought goal of environmental activists. The paperless society. The Cool study sees the Internet saving "2.7 million tons of paper every year by 2003, as it reduces the need to print newspapers, catalogs, direct mail, and the like." ("The Internet may give a boost to energy efficiency," J. Romm, Yahoo.com, 01.24.00) Perhaps. But this sounds eerily like the paperless office touted as the result of word processors a decade ago. So far, paper use is up.

Then too there is the long-promised energy savings from telecommuting. Certainly telecommuting uses less fuel than driving your car. But auto and air travel is up even with the rise of telecommuting. The reasons are complex, but even the co-inventor of the Internet himself has concluded:

> "The Internet has the funny effect of increasing the amount of travel."
>
> (Vinton Cerf, Senior VP of Internet Architecture, MCI WorldCom, actual co-inventor of the Internet, Engineering Tomorrow, IEEE Press, 2000, p.10.)

The jury is still out on this, but I have avoided several trips to DC in the past couple of years because of our videoconferencing facilities. Major companies like AT&T are beginning to measure the impacts of their telecommuting programs on energy use and pollution emissions. Whether the Internet will increase or decrease travel in the aggregate is anyone's guess at this point.

Where does Internet Electric Demand Go From Here?

Up.

Not necessarily, as discussed above and below.

Will there be continued growth in the hardware of the Digital Age? Is the Digital Age fully formed, IT appliance invention, production and utilization fully saturated? All indicators point to the fact that we're just at the beginning. The number of applications, and the range of microprocessor-based devices, the magnitude and extent of the communication networks needed to integrate all the devices is still at that so-called knee in the hockey-stick curve.

One hundred million computers today will become hundreds of millions in a few short years; globally, billions. As the Internet moves increasingly into a wireless mode, power use will grow disproportionately because it is inherently less efficient to broadcast than pipe information. The Palm VII and similar handheld devices and their wireless access to the Internet are only the beginning of an explosive trend. Add to this the ever expanding appetite for faster Internet access, and more broadband services. This is just the beginning.

As computers become more widespread, they will also become more power efficient (mobile applications demand this). Even the PC will give way to Internet access terminals with flat screens, using low power CPUs that are not state of the art, but are perfectly adequate for browsing the web. In addition, the latest CPU chips from Intel and Motorola have vastly more sophisticated power management capabilities than their predecessors. Just because there are billions more devices doesn't mean that electricity use will grow proportionally.

Thus will the Information Economy keep driving demand for electricity? Or will the market's use of new electric devices reach saturation, and efficiency gains combine to flatten out load growth? These two key questions have been posed repeatedly over the past two decades with regard to electric use in general, and are even more critical to understand today, at the dawn of Internet era.

The old conventional wisdom was that PCs and their kin would follow the efficiency trend of all other electric appliances. In one sense they have. Certainly PC monitors are more efficient today, as are many PCs. But unlike lights, chillers and refrigerators, the number of PCs and PC type devices has grown geometrically in a few short years.

In the past, some prominent forecasters have been confident that demand for electricity would stop growing because of efficiency gains and market saturation. We hear much the same language today, with much the same reasoning.

In 1980, a study from the Union of Concerned Scientists predicted:

> "Because saturation levels for most major appliances are achieved, only minor increases in electricity consumption [will] occur."
> (Energy Strategy, Union of Concerned Scientists, 1980)

In 1981, a study from the then Solar Energy Research Institute, since re-named the National Renewable Energy Lab concluded:

> "It appears that the demand for electricity is unlikely to increase significantly during the next two decades."
> (A New Prosperity: Building a Sustainable Energy Future, Solar Energy Research Institute, 1981)

What happened since 1980? Electric demand grew nearly 60%. What went 'wrong'? The analysts completely misunderstood the technology trends of ever greater applications for electricity, uses that more than offset improved efficiency. The same mindset is in place once again with regard to the information age and the Internet.

Incorrect forecasts from 1980-81 are irrelevant. Mills seems convinced that electricity demand must always grow, but his numbers do not in any way demonstrate that it will.

More recently, researchers at LBL concluded in 1995, just five Internet years ago:

> "While total energy use for office equipment has grown rapidly in recent years, this growth is likely to slow in the next decade because the US commercial sector market is becoming saturated (especially for PC CPUs and monitors)."
> ("Efficiency Improvements in U.S. Office Equipment," LBL, December 1995.)

I stand by this statement. Most of the commercial sector is approaching one PC per person, and the declines in energy used by mainframes and minicomputers were greater than we forecast in 1995. There are still five years to go before the decade starting in 1995 will be complete, and I will happily assess the accuracy of this forecast at that time.

To be charitable, forecasters even five years ago could hardly have forecast the growth in electricity-consuming IT-type equipment. But this has not stopped the refrain from continuing. The indicators for future trends are nothing less than amazing.

There are literally trillion of objects manufactured each year. We are rapidly approaching a time when everything will be manufactured with a silicon device of some kind, and where virtually all of them will communicate. Even if the energy needs of this trillion chip industry, and trillion petabyte bandwidth are trivial in per-chip terms – the aggregate electric needs will no doubt be astonishing.

This last conclusion does not follow. New routers have vastly higher data throughput than their predecessors, yet use less power. Distributed and mobile applications of microprocessors require the use of chips that are relatively low power, because batteries have limited life. If new devices are much more efficient than their predecessors, electricity demand growth from these devices could be modest. Without a detailed analysis, there's no way to be sure.

As bandwidth demand rises, power use rises, as does the market's use of the services. Yes efficiency will rise too. But for some time, as we build out the new infrastructure of the Digital Age, efficiency gains will be overwhelmed by sheer growth. Electricity is the fuel of the Digital Age, and the Internet at the heart of this revolution.

Mills' insistence that electricity demand must always grow in conjunction with demand for bandwidth or increased use of computers has not been substantiated. To validate his belief, he must supply data and analysis that stand up to independent peer review by the relevant expert community. Thus far, his work on this topic has crumbled in the face of our independent peer review.

No energy policy, including and perhaps especially the anti-electricity aspects of the Kyoto Protocol, should be considered without passing it first through a Digital sanity test. The integrity, reliability and low cost of the national electric infrastructure will be more, not less important in the future. A juxtaposition of key facts illustrates a policy collision course. Kyoto Protocol advocates call explicitly for the reduction, even elimination of fossil fuels and especially coal from the nation's energy infrastructure. Yet the nation gets 70% of its electricity from fossil fuels (three-fourths of that from coal). EIA forecasts that more fossil fuels will be needed to support economic growth. And while EIA forecasts natural gas will dominate the growth, they also forecast coal use will rise to support the economy.

The EIA's Business-As-Usual (BAU) forecast is just that. It says nothing about how electricity demand and supply could be affected if the U.S. implemented a set of programs and policies to improve efficiency, promote renewable energy, and prevent pollution. A BAU forecast only describes the likely outcome if we do nothing to change the course we're on.

Clearly energy policy and the Digital Economy are tightly linked.

They are related, though there are many possible energy futures that would result in more information technology and vastly lower carbon emissions (in fact, having more information technology may help achieve that goal, because of the autonomous system control it allows). It is misleading to imply that information technology *must* increase electricity use without a detailed analysis of those systemic effects. No one knows exactly how these systemic effects will evolve, but they are likely to be large enough to matter, and they could go either way.

Conclusions

In his Congressional testimony, Mills attempted to refute our critique of his analysis and defend his earlier results. His effort fails on several counts.

1) **In spite of larger than normal growth in GDP, recent historical data on electricity use show a slowing of electricity demand growth for the U.S. as a whole, not the acceleration of growth that would be required if Mills' thesis was correct.**
2) **Mills has yet to offer a detailed point-by-point rebuttal of our critique. The few specific points he raises in the testimony are based on inconsistent comparisons or incorrect assumptions, as shown above.**
3) **Mills repeats incorrect statements he has made about physical quantities about which we have many measurements (like the active power use of a PC plus monitor being 1000 W, when it is 150-200 W).**
4) **Our detailed calculations show that electricity used for all office, telecommunications, and network equipment (including electricity used to manufacture the equipment) is about 3% of total electricity use in the U.S. This estimate is about a factor of four lower than the 13% claimed by Mills.**
5) **While office equipment is clearly a large enough end use to warrant further study, it is not large enough to support the sweeping policy statements Mills continues to make.**
6) **Mills incorrectly implies that growth in information technology *must* lead to an increase in electricity use. This link is not demonstrated, and there are many factors that could also lead to a decrease in electricity use from application of information technology.**

In summary, Mills' thesis has been decisively falsified, and his assumptions shown to be erroneous. It is inevitable that small errors or differences in interpretation will arise in any complex calculations. However, in this case I have identified selectively misleading examples, unsupported conclusions, and a consistent and willful series of large errors (factors of four to twenty five) all in the direction of overestimating electricity use. It is therefore time to put this urban legend to rest.

Acknowledgements

In preparing this rebuttal, I was greatly aided by review comments from Amory Lovins, Bruce Nordman, and Alan Meier. I am also grateful to Skip Laitner of EPA's Office of Atmospheric Programs for funding this work.

This work was supported by the Office of Atmospheric Programs of the U.S. Environmental Protection Agency. Prepared for the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

References

These references (as well as the congressional testimony of Mills, Romm, and Hakes) are accessible at http://enduse.LBL.gov/Projects/InfoTech.html, with the exception of the Interlab Working Group 1997 report, which can be downloaded at http://enduse.LBL.gov/Projects/GHGcosts.html, and the Krause et al. report, which is available at http://www.ipsep.org

Interlaboratory Working Group. 1997. *Scenarios of U.S. Carbon Reductions: Potential Impacts of Energy-Efficient and Low-Carbon Technologies by 2010 and Beyond*. Oak Ridge, TN and Berkeley, CA: Oak Ridge National Laboratory and Lawrence Berkeley National Laboratory. ORNL/CON-444 and LBNL-40533. September.

Kawamoto, Kaoru, Jonathan Koomey, Bruce Nordman, Richard E. Brown, Maryann Piette, and Alan Meier. 2000. *Electricity Used by Office Equipment and Network Equipment in the U.S.* 2000 ACEEE Summer Study on Energy Efficiency in Buildings. Asilomar, CA. August.

Koomey, Jonathan G., Mike Cramer, Mary Ann Piette, and Joe Eto. 1995. *Efficiency Improvements in U.S. Office Equipment: Expected Policy Impacts and Uncertainties*. Lawrence Berkeley Laboratory. LBL-37383. December.

Koomey, Jonathan, Kaoru Kawamoto, Bruce Nordman, Mary Ann Piette, and Richard E. Brown. 1999. *Initial comments on 'The Internet Begins with Coal'*. Berkeley, CA: Lawrence Berkeley National Laboratory. LBNL-44698. December 9.

Krause, Florentin, Eric Haites, Richard Howarth, and Jonathan Koomey. 1993. *Cutting Carbon Emissions–Burden or Benefit?: The Economics of Energy-Tax and Non-Price Policies*. Energy Policy in the Greenhouse. Volume II, Part 1. El Cerrito, CA: International Project for Sustainable Energy Paths.

Mills, Mark P. 1999. *The Internet Begins with Coal: A Preliminary Exploration of the Impact of the Internet on Electricity Consumption*. Arlington, VA: The Greening Earth Society. May.

Romm, Joe, Arthur Rosenfeld, and Susan Herrmann. 1999. *The Internet Economy and Global Warming*. Washington, DC: Center for Energy & Climate Solutions.

US Bureau of the Census. 1998. *1996 Annual Survey of Manufacturers*. Washington DC: US Department of Commerce. M96(AS)-1. February.