ERNEST ORLANDO LAWRENCE
BERKELEY NATIONAL LABORATORY

Jonathan G. Koomey, Staff Scientist 1 Cyclotron Road, MS 90-4000 Berkeley, California 94720
Energy Analysis Department Tel: 510/486-5974 Fax: 510/486- 4247 e-mail: JGKoomey@lbl.gov

MEMORANDUM (LBNL-44698)

December 9, 1999

To: Skip Laitner, EPA Office of Atmospheric Programs
From: Jonathan Koomey, Kaoru Kawamoto, Bruce Nordman, Mary Ann Piette, and Richard E. Brown
RE: Initial comments on "The Internet Begins with Coal"
cc: Mark P. Mills, Rob Bradley, Amory Lovins, Joe Romm, Alan Meier, Alan Sanstad, and Erik Brynjolfsson
Download this memo and related data at: http://enduse.lbl.gov/projects/infotech.html

SHORT SUMMARY

This memo explores the assumptions in Mark P. Mills' report titled *The Internet Begins with Coal* that relate to current electricity use "associated with the Internet". We find that Mills has significantly overestimated electricity use, in some cases by more than an order of magnitude. We adjust his estimates to reflect measured data and more accurate assumptions, which reduces Mills' overall estimate of total Internet-related electricity use by about a factor of eight.

INTRODUCTION

At your request, we have begun to explore some of the assumptions in Mark P. Mills' report titled *The Internet Begins with Coal* (Mills 1999). In this memo, we restrict our comments to a few key assumptions in Part 7 of Mills' report, where he estimated total current electricity demand associated with the Internet. We do not address in this memo any of Mills' assertions about future growth of Internet electricity use, nor do we address any comments he made about the types of electricity supply technologies that would support any such increases in electricity demand. As more data become available, we expect refine the estimates in this memo, which must at this time be treated as preliminary.

The existence of the Mills report highlights the critical need for comprehensive data on electricity used by office equipment and associated network related-hardware. The last time such a comprehensive report was done (Koomey et al. 1995) was prior to the Internet becoming such an important force in the U.S. economy. That report did not address energy used by network hardware, but it did deal explicitly with stocks, energy use per unit, and operating hours to estimate total electricity used by commercial sector office equipment in the U.S. It compiled measured data on many of these parameters, which guided the creation of the scenarios generated in that report. This kind of comprehensive analysis, updated to reflect recent market developments and encompassing a broader scope, would resolve many lingering questions on this issue.

From a methodological perspective, it is problematic to assess only one portion (e.g., the Internet portion) of electricity used by office equipment in the U.S. In the absence of a complete accounting for all office equipment (as found in the Koomey et al. report), the accuracy of the calculations cannot rely on the checks and balances that such a complete accounting would enforce. For example, the total number of personal computers (PCs) is known with much more precision than the number of PCs associated with the Internet, so trying to estimate the latter without first estimating the former will yield a much less certain result.

There are difficult boundary issues in this assessment as well.[1] Mills chose to estimate the electricity used by the Internet and associated equipment, but he did not attempt to assess the effects of *structural changes* in the economy that are enabled by the existence of the Internet. These structural changes will almost certainly affect electricity and energy use. Without assessing the effect of these changes, the *net* effect of the Internet cannot be calculated, yet that is really what we care about. Given the large productivity benefits induced by computer hardware when properly used, it is plausible to speculate that these changes will be large enough to matter.

Mills also makes the assumption that all usage associated with the Internet is incremental. Instead it is actually more likely that at least some of this usage is *substituting* for other energy consuming functions that preceded the Internet (the Internet is expanding uses for the PC at the expense of other energy-using devices). Private computer networks and fax machines, for example, are increasingly being displaced by the Internet. Computer use is substituting for other forms of entertainment, like TV. Even voice communications (formerly the province of the telephone network) are being carried over the net. Such displacement effects represent another difficult boundary issue not treated in Mills' analysis.

In addition, the definition of which hardware is "associated with the Internet" is at best an imprecise one. Is a home computer "associated with the Internet"? People might use it for writing, for doing calculations, for analyzing personal finances, for creating party invitations, or for accessing the net. Does that mean that ALL of its energy use can be attributed to the Internet, or just a part? If just a portion, how much should be allocated to Internet use? Many of the reasons for owning a computer are independent of the Internet, and taken together justify the purchase of a computer. The same conclusion holds even more strongly for PCs in offices, since there are many reasons beyond Internet access for companies to invest in PCs. This kind of arbitrary allocation makes for calculations that are at best limited in usefulness.

In some sense, Mills is asking the wrong question by focusing on the Internet-related portion of electricity use by office equipment. Future studies should analyze total electricity used by this equipment, and not focus on what is Internet related, because these boundary issues are so difficult to resolve.

We turn now to specific assumptions that Mills made in his analysis. There are only a small number of assumptions that drive his results. Table 1 shows that of the 295 TWh that Mills calculates for Internet-related electricity use, more than half is in just three categories: Mainframe computers that serve "Major dot com company" web sites, Web sites using

---

[1] The ISO 14000 standards documents (particularly ISO 14040 and 14041) deal with the methodological issues surrounding such boundary issues. See http://www.iso.ch.

smaller servers, and telephone central offices. An additional 25% of Mills' total energy use is associated with use of PCs in offices and homes, and another 8% of Mills' total energy use is associated with routers. The rest (10%) is associated with the embodied energy to manufacture the equipment. We treat each of these categories in turn.

Table 1: Summary of Mills' estimates of current electricity use associated with the Internet circa 1998

	# of units Millions	*Elect. Used TWh/year*	*% of total*	*Cumulative %*
1) Major dot-com companies	0.033	72	24%	24%
2) Web sites	4	52	18%	42%
3) Telephone central offices	0.01	43	15%	57%
4) PCs in offices for all purposes	40	44	15%	71%
5) PCs at home for all purposes	41	31	10%	82%
6) Routers on Internet	2	16	5%	87%
7) Routers in LANS and WANs	1	8	3%	90%
8) Energy to manufacture equipment	19.5	29	10%	100%
Total		295	100%	

1) MAINFRAME COMPUTERS FOR 'MAJOR DOT COM COMPANIES'

Mills takes the number of mainframe computers in the U.S. from the ITI data book, which is the industry source for such numbers. He assumes that 10% of all mainframe computers in companies other than the major Internet companies are devoted solely to serving web sites. We have no way to judge the plausibility of this assumption, but we note that many such computers serve multiple functions (it is their multitasking abilities that make them so useful). Mills' choice to add the 10% of total mainframe installations to the number of mainframes/web farms in "Major dot-com" companies is an arbitrary one, but one with which we do not have the data to quibble.

Definitions of mainframe computers are not well established, and it appears that Mills did not use the same definitions for his stock and power estimates. The stock estimates rely on the ITI data book numbers, which count any computer costing more than $350,000 as a mainframe. The power estimates he used appear to be inconsistent with this definition.

Mills assumes that each mainframe uses 250 kW, 8760 hours per year. Half of this is assumed to be direct electricity used by the computer, and half for cooling. If the computer's direct consumption is 125 kW, this would place it in the ballpark of LBNL's Phase I supercomputer, installed in July 1999, which draws 150 kW (actual, not rated). It has about 600 processors, and is one of the most powerful in the world. The LBNL supercomputer cost tens of millions of dollars, but such supercomputers number only in the hundreds in the U.S. The bulk of mainframe installations are nowhere near the computing power of a supercomputer, yet that is the power use Mills chose for the typical mainframe.

For LBNL's Phase I supercomputer, the actual power use is about 0.25 kW per CPU.[2] If we use this consumption per CPU, the 125 kW Mills assumes is equivalent to a supercomputer

[2] Note that the LBNL Phase II Supercomputer, now under construction, will use about 0.1 kW per CPU. Source: Howard Walter at LBNL, who is designing the power systems for the new LBNL NERSC building. He generously provided numbers on the power requirements of supercomputers and their associated cooling loads.

with 520 processors. This represents far more processing power than a typical mainframe computer.

The IBM S/390 Enterprise server, which Mills' report cites as an example of the latest mainframe technology, has a rated (maximum) power of up to 6.4 kVA (roughly equivalent to 6.4 kW), depending on the number of processors. For the reasons described in Nordman (1999), the *actual* power use of almost all types of electronic equipment is typically one-half to one-third of the rated power (the rated power is the maximum power that the power supply will consume under fully loaded, worst case conditions). If the actual power is half of the rated power, this machine would use 3.2 kW for typical installations. Of course, IBM's server is relatively new and it relies on CMOS technology to reduce power use, so it probably uses less power than an older mainframe. The rated power may also not include peripheral equipment that would be included in a typical mainframe installation.

It is clear that 125 kW is a much larger power number than has been used in such analyses in the past. The Koomey et al. (1995) report estimates power used by older (1985-1990) mainframes at 25 kW, declining to 10 kW by 1999 (an estimate for 1999 which is validated by the S/390 data described in the previous paragraph). The recent Swiss study by Meyer and Schaltegger (1999) used 30 kW for the average power of each of the roughly 1000 mainframes in Switzerland.

We checked the price of the S/390 on the IBM web site and found that its cost is well above ITI's $350k cutoff for mainframes (S/390s cost millions of dollars). We believe, as Mills also does, that this machine is representative of mainframe computers now being installed. If we accept the 3.2 kW direct power use of the S/390 and quadruple it to account for peripheral equipment, that still leaves our estimate of power used per mainframe (12.8 kW) at about one tenth of what Mills assumes.

Cooling is at most 50% of direct power consumption, not 100% as Mills assumes. This result follows from the compressor-based cooling technologies commonly used in commercial buildings and computer rooms, which have Coefficients Of Performance (COPs) of 2.0 or better. A COP of 2.0 implies that 1 unit of electricity is consumed to move 2 units of heat out of the conditioned space. We consulted with the supercomputer team at LBNL, who use 50% additional power for cooling as their best guess for maximum cooling loads when designing a new supercomputer (although in actual practice, 30% is more typical in the Bay Area).

Using our 12.8 kW direct power use estimate, combined with a 50% multiplier for cooling energy, leaves us at 19.2 kW per mainframe. If we replace Mills' assumption of 250 kW with this new estimate, total electricity used by the Major dot-com companies becomes 5.5 TWh, a reduction of 66.5 TWh or about a factor of thirteen. By itself, this correction reduces Mills' estimate of electricity used by the Internet by 22%.

2) WEB SERVERS

Mills refers to an article titled "WWW Hosts 5 Million Web Sites" (http://www.nua.ie/surveys) to justify his assumption of the number of web servers. He takes 70% of 5 million, which rounds to 4 million servers. The problem is that the article to which Mills refers talks about web sites NOT servers. One server can host dozens of web sites (a fact that Mills acknowledges in his report), so the number of servers is much lower than the number of sites. We assume, for purposes of these calculations, that each server hosts 5 sites (although that is likely to be an underestimate). In practice, some servers will have just one site, and others will have many. This correction factor alone reduces Mills' estimate of electricity use for this component by 80%.

The power used by mini-computers and workstations is assumed by Mills to equal 1.5 kW, 1 kW of which is direct power used by the computer, and 0.5 kW is from peripherals, "especially data backup". Data backup only runs once a day, and services many CPUs. It is unlikely that 0.5 kW per CPU is a reasonable estimate for this service.

Based on the discussion of PC power use below, we reduce Mills' 1.5 kW estimate by a factor of 5, to 0.3 kW. With both corrections (for number of units and power per unit), total power used by U.S. web servers is reduced by a factor of 25, to 2.1 TWh.

3) TELEPHONE CENTRAL OFFICES

Telephone central offices are the next most important item in Mills' list, but much more information and documentation is needed to justify the calculations, particularly the number of such offices and the power use per office.

Mills' estimates that central offices each use 500 kW. His table indicates that there are 25,000 such central offices in the U.S. In fact, most of these central offices are significantly smaller than Mills' assumes (between 30 and 50 kW). We are working on getting an accurate distribution of such central offices by power level, but in the absence of those data, we took another tack.

Our contact at a major phone company reports that a central office uses about 3.3 kWh per thousand minutes of so called "dial equipment minutes" or DEQ (a standard measure of phone connect time). FCC (1999) reports total DEQ for the U.S. of 3,612 billion minutes in 1997. These two numbers together imply electricity used by all central offices of 12 TWh/year. To make this number comparable to Mills' estimate, we multiply this figure by 40%, to get 4.8 TWh/year.

With this revised estimate, power used by central offices is reduced by 37 TWh, or about a factor of nine.

4) OFFICE PCs

The power used by most personal computers is assumed by Mills to equal 1 kW. This estimate is assumed to include all peripheral equipment associated with PCs, as well as some unspecified other equipment. Without a detailed accounting of his assumptions about this equipment, it is difficult to determine what he assumed. However, there is a large body of literature on actual power used by such equipment. A recent power measurement of a 500 MHz Pentium III PC that had no power management showed average power over the course of a day of about 40 W for the CPU.[3] A typical 17" monitor uses about 90 W in active mode.

Of course, most PCs and monitors now are capable of power management (which neither of the above measurements include), so that over the course of a day, these power numbers would be reduced. ENERGY STAR PCs power down to less than 30 W, and typical ENERGY STAR monitors power down to less than 10 W. Whether power management is enabled in many cases is an open question (recent surveys found roughly a third of PCs and monitors had power management correctly enabled in offices), but for the sake of argument, we ignore it.

---

[3] Personal communication with Bruce Nordman, LBNL, November 1999.

Peripheral equipment is often shared. In our office, 20 people share a workgroup printer (HP 8000 DN). We metered this printer, and it draws about 163 W in active mode (when printing), and about 120 W in standby mode. In sleep mode it draws about 30 W, and on average, including sleep modes and printing, it uses about 50 W. Even ignoring the power management of the printer, and assuming it is constantly printing, it would add only 8 W per CPU to our estimate of average PC power. For home PCs, most printers will be inkjets, which typically draw less than 30W even in active mode.

It is not clear what other equipment Mills is referring to in his 1 kW estimate, but we feel it is unlikely to push the average power used by PCs and peripherals to greater than 200 W (and with power management, we feel strongly that 200 W average power is an overestimate). For purposes of these calculations, we use 200 W instead of Mills' 1 kW estimate, and recalculate electricity used by PCs to reflect this revised estimate. If he believes that other "behind-the-wall" components account for a significant amount of power use (800 W/CPU in this case), he needs to specify, item by item, the number and power use of all these components. We examined the "behind the wall" components of the LBNL computer network, but we were unable to figure out how these components, most of which serve multiple CPUs, could possibly add up to 800 W per CPU (tens of watts per CPU over and above router power use is more like it. See Nordman 1999 for details).

The power used by high end personal computers is assumed by Mills to equal 2 kW. We reduce this estimate also by a factor of five, to 400 W, even though this is almost certainly an overestimate (doubling the CPU power for the 500 MHz Pentium III above and assuming a 21 inch monitor at about 120 Watts only leads to actual power use of 200 Watts, without considering power management).

For usage of PCs at the office, Mills assumes twelve hours per week, which is the same as that assumed for home PCs (see point 5, below). It is important to note the inherent arbitrariness of attempting to calculate what part of office PC use is "associated with the Internet". We have no data for what portion of office computer use is "associated with the Internet". The only adjustment we make to Mills' usage numbers in this category is to reduce usage for typical office PCs to seven hours per week from 12 hours per week, to reflect our revision in the home PC usage number below (this adjustment preserves consistency between our methodology and that of Mills). We do not change usage assumptions for PCs at offices behind a firewall or PCs used in commercial Internet services.

With these corrections, PCs in offices use about 7.2 TWh, a reduction of 84% from Mills' estimate.

5) PCS AT HOME

Mills' assumption of 1 kW power draw for home PCs is subject to the same issues examined under point 4, and hence we reduce his 1 kW by a factor of 5, to 0.2 kW. Even though peripheral equipment in a home is associated typically with one PC instead of many in a work environment, that equipment is not always on when the computer is on, and it is likely to be lower power versions of that equipment (e.g. ink jets instead of laser printers, low end scanners, etc.).

We now turn to usage of home PCs. Mills assumes twelve hours per week of usage for these computers, based on an Intelliquest study of home users, but he acknowledges that "other surveys show lower averages". He cites the Neilson/NetRatings March 1999 survey at seven hours per month (less than two hours per week). Another quite recent study (7 December 1999) shows usage of five to eight hours per month, which is also about two hours per week (http://www.nua.ie/surveys/?f=VS&art_id=905355453&rel=true). With this great a range in estimates of usage, it is important to be cautious in drawing conclusions. We were unable to locate any studies that indicated that average U.S. Internet users were logged on more than 12 hours per week, so we feel justified in treating this as an upper bound, with the likely average possibly as low as two hours per week. Even choosing seven hours per week (the midrange between those two estimates) would reduce Mills' estimates for electricity associated with home Internet use substantially.

Mills claims he is being conservative by assuming that

> every single PC and all its relevant peripherals accessing the Net is physically turned on and operating only the 12 hours per week from the Forrester Research (IntelliQuest) survey, and otherwise completely off. As a practical matter many (possibly most) are on at least 50 hours per week, many 24 by 7. A 'realistic' weekly 'on' time of 50 hours yields about the same rough kWh for a 200-300 W duty-cycle compensated PC as the conservative 12 hour/wk duty cycle does for a 1,000 peak W device (footnote 53, in Mills 1999).

The claim of conservatism is spurious. People use their computers for many other things besides Internet access, which is why at least some people have their computers on for 50 hours per week or more (though we doubt many home users do). According to the surveys cited below, the Internet-related component of home PC use is between two and 12 hours per week. That some people keep their home PCs on for more hours than that is irrelevant to Mills stated purpose, that of calculating electricity use "associated with the Internet".

Based on the surveys cited above, we choose usage of seven hours per week for typical home PC users, instead of 12 hours per week. We do not change hours of usage for PC power users or PCs in home offices.

With these corrections for power and usage, PCs in homes use about 5 TWh, a reduction of 84% from Mills' estimate.

6) ROUTERS ON THE INTERNET

Mills' assessment of the number of routers seems inconsistent with our review of the market for these products. It is not clear why there would be twice as many high end routers as low end ones, when in fact the low end ones must be more numerous in any network with central nodes feeding dispersed nodes. We did not correct for this observation, but simply note it for future research. It is also not clear if Mills' stock estimates include switches and routers together, or just routers alone. This issue also must await further research.

Cisco's very highest-end router, which is used in the highest throughput applications, has a rated power of 1.5 to 2 kW. The actual power used for this device will then be 0.75 to 1 kW, because rated power is typically two times the actual power (see text under item 1 above). Unfortunately for Mills' argument, there are very few of these large routers sold every year. Based on a review of the high-end routers sold by Cisco systems, we find that more typical high end routers have rated power of 0.3 to 0.8 kW in typical use (actual power of 0.15 to 0.4 kW). We therefore reduce power use to 0.3 kW, from 1 kW.

Once we correct the power use, routers on the Internet show total consumption of 4.8 TWh, a reduction of 70% from Mills' estimate.

7) ROUTERS ON LANS AND WANS

Routers on Local Area Networks and Wide Area networks (LANs and WANs) use much less power than Mills assumes. In the text of his report, he states that he uses 0.5 kW for the smaller routers, and 1 kW for the larger (Internet) routers. The total TWh calculations do not support this assertion--they imply that the 1 kW assumption was also used for the LAN and WAN routers (divide 8 TWh by 1 million routers, and then by 8760 hours, and you get just under 1 kW).

Mills therefore assumes 1 kW average power draw for all routers. Cisco's typical lower end routers (which account for the majority of all routers) range in rated power from 0.04 to 0.2 kW. We therefore reduce Mills' estimate by a factor of twenty, to reflect a rated power of 0.1 kW and an actual average low-end router power of 0.05 kW (this last factor of two correction from rated power to actual is the same as that used under points 1 and 6 above).

Once we correct the power use, routers on LANs and WANs show total consumption of 0.4 TWh, a reduction of 95% from Mills' estimate.

8) MANUFACTURING ENERGY

Manufacturing energy for computers on the Internet is the most difficult of these categories to analyze, because of the lack of data. The life-cycle assessment needed to calculate embodied electricity use of electronic equipment is a complicated exercise, and one that has only rarely been carried out. The most recent data we examined come from NEC, the largest computer manufacturer in Japan (Tekawa 1997).

NEC estimates total greenhouse gas emissions from manufacturing a desktop PC to be 128 kg/CO_2 equivalent (unfortunately, we don't at this time have much detail on the components of this calculation). Some of these emissions are non-CO_2 greenhouse gases, and some are from non-electricity related fuel use. Nevertheless, we can get an estimate for the upper bound to electricity used for manufacturing all parts of the PC by assuming all of these emissions come from electricity (electricity is more carbon intensive per unit of energy consumed than direct use of fuels). The average emissions factor for Japanese electricity production is about 0.42 kg CO_2 per kWh (115 g C per kWh). This factor implies total electricity use of about 300 kWh per desktop PC, which is an upper bound, as described above. NEC states that the electricity used to assemble their PCs is about 120 kWh per unit,[4] so total electricity use is between 120 kWh and 300 kWh per PC. We chose 300 kWh per PC, which is one fifth of Mills' estimate. This is an absolute upper bound. The true number is almost certainly lower than this.

With this factor of five correction, Mills estimate of 29 TWh for manufacturing energy is reduced to 6 TWh.

CONCLUSIONS

Table 2 shows Mills' estimates corrected as described above. In every category, his estimates must be reduced substantially (by factors of 3 to 25) to reflect more accurate assumptions. For all categories taken as a whole, Mills' estimates are reduced by 88%.

Mills' report does not contain enough detailed documentation to assess the reasonableness of many assumptions, but it is clear from the review of assumptions conducted above that he has vastly overestimated electricity use associated with the Internet. In addition, the value of such estimates is questionable, given the difficult boundary and allocation issues described above. It would be more useful to estimate total electricity used for all office equipment and

---

[4] Both Compaq and Dell appear to use significantly less electricity than NEC to assemble their PCs, and we are investigating this difference.

associated network equipment, because that number is inherently more reliable than deriving what fraction of such devices are "associated with the Internet".

Finally, the structural and substitution effects alluded to above are almost certainly large enough to matter. Future estimates of the impacts of the information technology revolution (which are larger in scope than those of just the Internet) should explicitly account for these effects.

Table 2: Corrected estimates of current electricity use associated with the Internet

	# of units Millions	*Elect. Used TWh/year*	*% of total*	*Cumulative %*
1) Major dot-com companies	0.033	5.5	15%	15%
2) Web sites	0.8	2.1	6%	21%
3) Telephone central offices	0.01	4.8	13%	35%
4) PCs in offices for all purposes	40	7.2	20%	55%
5) PCs at home for all purposes	41	5.0	14%	69%
6) Routers on Internet	2	4.8	13%	82%
7) Routers in LANS and WANs	1	0.4	1%	83%
8) Energy to manufacture equipment	19.5	6.0	17%	100%
Total		36	100%	

REFERENCES

FCC, Federal Communications Commission. 1999. *Trends in Telephone Service*. Washington, DC: Industry Analysis Division, Common Carrier Bureau. September. Found at http://www.fcc.gov/Bureaus/Common_Carrier/Reports/FCC-State_Link/trends.html

Huber, Peter, and Mark P. Mills. 1999. "Dig more coal—the PCs are coming." In *Forbes*. May 31, 1999. Read online at http://forbes.com/forbes/99/0531/6311070a.htm.

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. http://enduse.lbl.gov/Info/37383-abstract.html. Also published in 1996 in *Energy Policy*. vol. 24, no. 12. December. pp. 1101-1110.

Meyer & Schaltegger AG. 1999. *Bestimmung des Energieverbrauchs von Unterhaltungselektronikgeräten, Bürogeräten und Automaten in der Schweiz.* St. Gallen, Switzerland: Bundesamts für Energie. March.

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. http://www.fossilfuels.org.

Nordman, Bruce. 1999. Memo to Jonathan Koomey on Electricity requirements for LBNL's Networking Hardware. December 9. Lawrence Berkeley National Laboratory.

Tekawa, Masafumi, Shigeyuki Miyamoto, and Atsushi Inaba. 1997. *Life Cycle Assessment: An Approach to Environmentally Friendly PCs.* Proceedings of the IEEE International Symposium on Electronics and the Environment. pp. 125-130.