

A regression approach to infer electricity consumption of legacy telecom equipment

[Extended Abstract]

Steven Phillips
AT&T Labs–Research
180 Park Ave
Florham Park, NJ 07932
phillips@research.att.com

Sheryl L. Woodward
AT&T Labs–Research
200 Laurel Ave S
Middletown, NJ 07748
sheri@research.att.com

Mark D. Feuer
AT&T Labs–Research
200 Laurel Ave S
Middletown, NJ 07748
mdfeuer@research.att.com

Peter D. Magill
AT&T Labs–Research
200 Laurel Ave S
Middletown, NJ 07748
pete@research.att.com

ABSTRACT

Communications technology has great potential for helping other industries reduce their greenhouse gas emissions. Nevertheless, given the relentless growth of demand for communications services, telecommunications providers will need to transition to more energy-efficient technology in order to limit their own environmental footprint. Here we focus on priority-setting for the transition process. We introduce a method for statistically inferring the electricity consumption of different components of the installed base of telecommunications equipment, while avoiding the high cost of performing direct measurements. Our method relies only on databases of installed equipment in central offices, together with aggregate electricity consumption per office. It takes advantage of inter-office variation in installed equipment to partition per-office electricity consumption by major equipment type. When applied to a collection of 3,918 central offices of a major U.S. telecommunications provider, our approach reveals the (previously unknown) network-wide energy consumption of each major type of equipment. In particular, we find that electricity consumption is dominated by Class-5 telephone switches, which account for 43% of aggregate consumption, and which should therefore be a primary target of central office electricity conservation efforts.

1. INTRODUCTION

While telecommunications networks offer great potential to help other industries reduce greenhouse gas emissions, there is a need to reduce the substantial electricity consumption of networks themselves [11]. In order to reduce or limit their

electricity consumption, telecommunications networks need to manage a transition from installed equipment to newer, more energy efficient technologies. Most research on sustainable information and communications technology has focused on new and emerging technologies, for example optical IP networks [1, 10], efficient data center architecture and design [6] and pico-cell mobile networks [5]. In contrast, there has been little research on improving the energy efficiency of existing, installed telecommunications infrastructure (but see [3]). The approach we take here is to help prioritize the transition out of existing energy-inefficient technologies. We do this by introducing a new method for inferring the electricity consumption of component parts of a major legacy telecommunications network, while avoiding the expense of direct measurements. Our approach allows a telecommunication provider to better understand its own electricity consumption, complementing published efforts to quantify global consumption by servers [8], the internet [2] and the global information technology industry [9]. Note that our focus here is on power that is directly consumed by telecommunications network elements; we do not address the substantial savings that can be achieved by optimizing cooling, power regulation and other equipment support systems [3].

In order to efficiently manage electricity consumption, it is essential to understand how and where electricity is being consumed. For example, power consumption in modern data centers is measured and modeled at multiple scales, from individual servers to larger server groups, and at very fine temporal scales [6]. In contrast, telecommunications networks include many generations of legacy equipment, many installed before fine-scaled monitoring was technically feasible, and before electricity consumption was as pressing a concern. Legacy equipment could simply be retrofitted with power meters, but this is infeasible at scale for two reasons:

- Attaching current measurement devices to installed equipment is invasive, with some risk of power interruptions which would result in unacceptable interruptions of service.

- There would be a very high labor cost for extensive retrofitting of measurement devices in the heterogeneous equipment environment of a central office.

Some natural alternatives exist for avoiding widespread deployment of direct measuring devices, but each has its limitations. While rated power consumption of some network elements is available, the rating may be higher than real-world consumption [6]. In addition, rated consumption is not available for all network elements. Alternatively, individual network elements could be analyzed in a laboratory setting. However, this would be a laborious process, due to the high diversity of network elements in legacy telecom networks (37,398 in the network we study below).

We therefore present a novel statistical approach that bypasses the limitations and expense of existing approaches, in order to give accurate estimates of the power consumption attributable to the various major equipment types in telecommunications central offices.

2. METHODS

2.1 Available data

The data available for this study consisted of two data sets, covering 3,918 central offices of a major telecommunications provider in the United States. We assume these data sets are accurate and exhaustive, but note below some anomalies that are best explained by small errors in the available data.

- DC power plant peak power readings for each central office. This equals the peak total power consumption of all telecom equipment in the office. These readings exclude power consumed by building systems (e.g., cooling, lighting) which run on AC power.
- An inventory of installed equipment in each central office, consisting of the number of installed units of each of 37,398 types of network element identified by HECI codes.

Network service providers use Human (Readable) Equipment Catalog Identifier (HECI) codes in order to inventory, provision and maintain operations. These codes are also referred to as COMMON LANGUAGE[®] Equipment Identification (CLEI) codes, and are an integral part of the Telcordia[®] TIRKS[®] system (an industry standard inventory and provisioning system used for many telecommunication networks). HECI codes are hierarchical, with network elements organized into “families”, “groups” and “types”. They provide the basic building block of the inventory of installed equipment. Units with identical HECI codes should also have identical attributes, including physical form, capabilities, and power requirements.

2.2 Regression estimation of power usage

Most legacy telecommunications equipment has roughly constant power consumption: the equipment is always on, and power consumption is independent of intra-day or other variation in carried traffic. Indeed, inspection of 15-minute AC consumption data for a small number of well-instrumented offices revealed intra-day variation of less than 10%, despite

much larger variation in traffic volumes. We therefore represent the power consumption of each network element by a single time-invariant number. This means that the peak power reading of an office equals the sum of power consumed by each network element. Under the reasonable assumption that all network elements with the same HECI code have the same power consumption, we conclude that the peak power reading of an office should be a linear function of the number of installed network elements of each type. It is therefore reasonable to apply model-based statistical methods [4] to analyze this data. In particular, we will infer the power consumption of each network element using linear regression.

Unfortunately, a standard linear regression model using all network elements as predictor variables would be woefully under-constrained. There are 37,398 unknowns (the power consumption of each type of network element), which is much greater than the number of data points (3,918, the number of central offices). We therefore present two approaches to reduce dimensionality: the first uses the hierarchical grouping of HECI codes into 131 “equipment groups”, and the second does a form of forward stepwise regression, adding HECI codes to the regression one by one in an attempt to model only the most important network elements.

2.3 Regression by equipment group

Each HECI code is a member of a specific equipment “group”, which generally corresponds to a switch type. Each of the different modules that might be incorporated within a specific switching system has its own HECI code within the group. For example, the 5ESS switch group includes all the various line cards developed over the years for this switch, and this group alone contains over 2500 HECI codes. To report the results of our study, we further aggregated similar equipment groups, so for example, we present Lucent 5E and Nortel DMS telephone switches together in a “Switch” group, although they have different HECI groups.

For each central office, we aggregated the counts of all network elements by group. This reduced the equipment counts to 131 numbers per central office. The number of unknowns was therefore reduced to 131, namely the per-element power consumption of each group.

Of course, power consumption varies between network elements in the same group, so the power consumption per element of a particular group may vary from central office to central office, depending on the particular network elements present. However, for many groups, the aggregate count was dominated by the number of line cards, and in particular those at the lowest multiplexing level; this accords well with the fact that for many switch types, power consumption is dominated by the consumption of the lowest level line cards. In addition, among line cards that terminate multiple lines, older line cards that terminate smaller numbers of lines were treated as equivalent to newer line cards; we believe this is reasonable, since line card power efficiency is likely to have improved in tandem with density of lines per line card.

The “POWER” group required special attention. Network elements in this group have power supply / regulation functions, and a regression could represent at least some of the

consumption of each office as a function of the number of power supply or regulation elements. However, these network elements are not themselves consuming much power. Rather, they pass it along to other network elements. We prefer to assign the power passing through these elements to the “end-users”, i.e., the network elements that directly consume power. We therefore omitted the “POWER” group from the regression entirely.

2.4 Forward stepwise regression by HECI code

Our second regression sought to attribute power consumption to individual network elements, rather than groups. To avoid the problems of under-specification described above, we used a forward stepwise regression approach [7] to automatically select the most important network elements to include as predictor variables in the regression; regularization methods such as ridge regression could be used instead.

As a preprocessing step, we first eliminated rarely-used HECI codes, retaining only 1,800 HECI codes out of the available 37,398 to use as candidate predictor variables. The remaining 1,800 consisted of the most numerous HECI codes in their group: 10-60 HECI codes were retained per group, with more being retained for the groups with larger aggregate numbers of network elements. In addition, we omitted HECI codes from the “POWER” group, for the reasons described in Section 2.3 above, along with HECI codes in other groups whose description suggested that their function is power-handling (containing strings such as “48v dc”, “supply”, “pwr” or “power”).

The forward stepwise regression was performed as follows. We start with no variables in the regression, and add a small number (5) of variables until a specified number of variables (375) was included in the model. The variables included at any step are the ones that, when considered individually, yield the greatest decrease in model error, as measured by residual sum of squares.

The resulting 375 variables are used in a linear regression model, which assigns a power consumption value to each of the corresponding network elements. We re-aggregated these values by equipment group, in order to obtain an estimate of aggregate power consumption by equipment group.

2.5 Direct measurements in a training facility

In order to verify results of the regression analyses, direct measurements were made for a small number of network elements. The measurements were made in an industrial training facility, designed to train network operations personnel. We note that the equipment configurations and traffic load of switches in such a facility are likely to be substantially different from switches in a functioning central office. Nevertheless, we relied on the general rule that most legacy network elements have constant power draw, independent of load, and we assume that the direct measurements represent the true power consumption of these network elements.

Measurements were taken on each of 14 switches. In each case, the baseline power consumption of the switch was measured using a clamp-on ammeter. All installed network elements (such as line cards, each identified by a HECI code) were then unplugged and reinstalled one at a time, and the

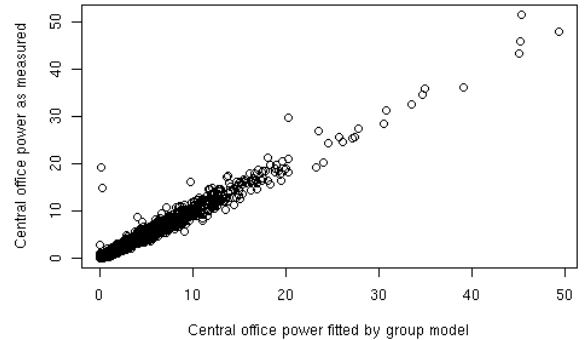


Figure 1: Quality of fit of the regression by equipment group. Measured power (normalized so that median is 1.0) is plotted against normalized fitted power for each of 3,918 central offices.

resulting change in power used by the switch was recorded for each network element.

3. RESULTS

3.1 Results of regression by equipment group

Table 1 shows the results of the regression by equipment group. Class-5 TDM telephone switches (Switch) are easily the largest contributor to power consumption, accounting for 43.0% of the total consumption. The next largest contributors are SONET/MUX systems (14.4%), D4 Carrier systems (8.5%) and digital cross-connect systems (DESX / DAX, 6.6%). In the central offices studied here, newer access technologies such as DSL appear to consume only a small fraction of total power (4.4% DSL).

The group labeled “(Intercept)” is the constant in the linear regression model. Each central office is thus assigned a constant power consumption per year, in addition to the consumption assigned to its equipment. This term may correspond to power-handling and other equipment that were excluded from the regression analysis, and accounts for 2.2% of the aggregate power consumption.

The “Testsets” group contains network elements involved in testing and maintenance of the equipment in many of the other groups. We expect that such testing equipment truly accounts for only a small fraction of total electricity consumption, and we do not have an explanation of why the regression assigned a fairly large fraction (4.3%) of total power to them.

Applying the fitted per-network element consumption to the inventory of equipment in each office gives its fitted electricity consumption, which can be plotted against the measured value (Figure 1). We observe a reasonably close correspondence between fitted and measured consumption. Two outliers lying close to the y-axis probably correspond to data errors, perhaps offices incorrectly matched between the database of DC power plant readings and the database of equipment inventories.

Table 1: Results of regression by equipment group. Groups contributing less than 1% to total power consumption are not shown. Total power and standard error are given in percentage units. All estimated coefficients are highly significant: $p < 0.001$ (*) or $p < 0.01$ (**).**

Equipment group	Number of items	Power percent	Standard error	Significance
Switch	40397721	42.55	1.08	***
SONET/MUX	4677187	14.39	1.35	***
D4 Carrier	2776633	8.51	1.44	***
DESX.DAX	3221307	6.56	0.43	***
DSL	1353256	4.45	0.70	***
Testsets	98911	4.33	0.50	***
SMDI	47210	3.21	0.29	***
Loop Access	3477897	3.02	0.82	***
(Intercept)	3918	2.21	0.76	**
Ethernet	8033	1.89	0.43	***
EBOC	4275	1.63	0.26	***
PGPLUS	24917	1.31	0.16	***
Repeaters	596527	1.13	0.15	***
Broadband	12974	1.03	0.12	***

Table 2: Results of forward stepwise regression by HECI code. Groups contributing less than 1% to total power are not shown.

Equipment group	Power percent
Switch	42.66
SONET/MUX	12.72
DSL	6.75
D4 Carrier	6.6
DESX/DAX	6.04
Loop Access	4.48
Ethernet	1.93
Testsets	1.77
Alarms	1.63
Repeaters	1.59

3.2 Results of stepwise regression by HECI

Table 2 shows the results of the regression by HECI code, aggregated by equipment group. Each term shown in the table corresponds to a sum over many HECI codes, with varying levels of statistical significance, so we do not report significance values for the sums. These results are similar to those of Table 1. Notable differences include:

- Less power assigned to Testsets (1.8% vs. 4.3%)
- Less power for D4 Carrier systems (6.6% vs. 8.5%)
- More power assigned to DSL (6.8% vs. 4.5%)
- No power assigned to constant intercept (vs. 2.2%)

Plotting per-office fitted electricity consumption against measured consumption (Figure 2) shows a better fit than for the regression by group (Figure 1). However, the better fit in the second regression may simply be due to the higher number of predictor variables used (375 vs 131).

Table 3 shows example measurements at the training facility, involving network elements installed in a Tellabs Titan 5500 digital cross-connect system. This system showed the highest level of agreement between the aggregate power consumption predicted by the regression by HECI code and the

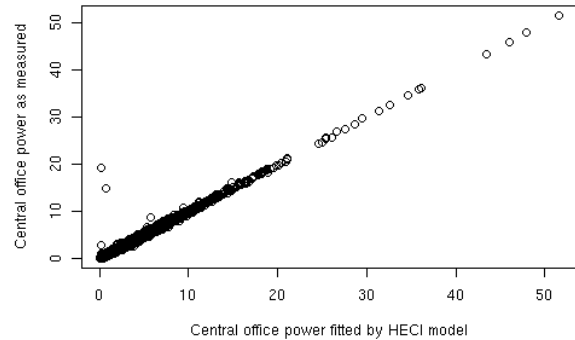


Figure 2: Quality of fit of the forward stepwise regression by HECI code. Measured power is plotted against fitted power for each of 3,918 central offices. Values are normalized so that median measured power is 1.0.

measured consumption (within 2%). The level of agreement in the 14 measured systems was as follows:

- 4 systems: agreement within 2-15%
- 3 systems: little agreement
- 7 systems: no HECI codes in common

It is worth noting that the power consumption assigned to each network element included in the stepwise regression cannot simply be interpreted as an estimate of the power consumed by that network element. The reason is that many network elements are commonly installed in standard configurations, such as a shelf with a common timing unit and multiple line cards. Once one element of such a configuration is included, the regression is likely to assign to it the power consumption of the whole configuration. Other elements in the configuration are not likely to get included in the regression, as doing so would not significantly reduce the residual

Table 3: Verification of forward stepwise regression by HECI code using measurements of a Titan 5500 at a training facility. The line cards in four bays of the switch were analyzed, and the change in load (amps) was measured on both A and B supplies when each bay was unplugged. HECI codes in gray are those that were included in the regression. The sum of load allocated by the regression to the grey HECI codes was 5.04 amps, within 2% of the 4.94 amps total usage of all bays.

Bay 1 Line Cards	Qty	Bay 2 Line Cards	Qty	Bay 3 Line Cards	Qty	Bay 4 Line Cards	Qty	Total load (amps)
T3C1AA8	1	T3C1BB1	4	T3D1AA3	1	T3C1AJ0	10	
T3C1AAG	7	T3C1D04	1	T3D1AAD	1			
		T3C1DCR	1	T3D1AAE	7			
		T3C1DCX	5					
		T3DAKZ0	1					
		T3PQALB	1					
Load A	0.22	Load A	2.57	Load A	0.02	Load A	0.37	4.94
Load B	-0.04	Load B	0.08	Load B	1.44	Load B	0.28	

sum of squares of the regression. This effect can be seen in Table 3, where the consumption of 12 network elements is assigned to only four elements included in the regression.

3.3 Discussion of results

We used two regression analyses to partition the electricity consumption by equipment group. Despite major differences in the two regression analyses, they assign very similar amounts of electricity to each equipment group, with the difference being less than 2% for most groups. This suggests that our analyses have a reasonably high level of accuracy, and we estimate that we have identified the contribution of each equipment group to within two percentage points.

There was less agreement between the predictions for the stepwise regression by HECI code and measurements taken in an industrial training facility: of the seven systems that contained at least one HECI code chosen in the stepwise regression, only four showed reasonable agreement. This is likely due to differences in equipment configurations between the training facility and typical central offices.

4. CONCLUSIONS

The purpose of this work was to analyze central office electricity consumption, which is a major component of electricity consumption for large telecommunications companies. In the 3,918 central offices analyzed here, newer technologies such as DSL still account for a fairly small fraction of total energy consumption. Current research on sustainable information and communications technology focuses on new networks and technologies, and therefore on limiting future growth in energy consumption. In contrast, TDM switches, especially Class-5 telephone switches, are the largest contributors to current central office consumption. Our study highlights the need for strategies to transition from older to newer technologies. In particular, in order to significantly reduce total central office electricity consumption in the near term, we suggest prioritizing efficiency measures and/or technology evolution for TDM switches.

5. ACKNOWLEDGEMENTS

We thank Marshal Tarbet for taking measurements at the training facility. We also thank all those who helped provide access to the databases used in this study.

6. REFERENCES

- [1] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, , and R. S. Tucker. Energy consumption in optical IP networks. *Journal of Lightwave Technology*, 27(13):2391–2403, 2009.
- [2] J. Baliga, K. Hinton, R. Ayre, and R. S. Tucker. Carbon footprint of the internet. *Telecommunications Journal of Australia*, 59(1):5.1–5.14, 2009.
- [3] C. Bianco, F. Cucchiatti, and G. Griffa. Energy consumption trends in the Next Generation Access Network - a Telco perspective. In *29th International Telecommunications Energy Conference*, pages 737–742, 2007.
- [4] L. Breiman. Statistical modeling: The two cultures. *Statistical Science*, 16(3):199–215, 2001.
- [5] H. Claussen, L. T. W. Ho, and F. Pivitt. Effects of joint macrocell and residential picocell deployment on the network energy efficiency. In *IEEE 19th International Symposium on Personal, Indoor and Mobile Radio Communications*, pages 1–6, 2008.
- [6] X. Fan, W.-D. Weber, and L. A. Barroso. Power provisioning for a warehouse-sized computer. In *ISCA '07: Proceedings of the 34th annual international symposium on Computer architecture*, pages 13–23, New York, NY, USA, 2007. ACM.
- [7] R. R. Hocking. The analysis and selection of variables in linear regression. *Biometrics*, 32, 1976.
- [8] J. Koomey. Estimating total power consumption by servers in the U.S. and the world. Technical report, Lawrence Berkely National Laboratory, 2007.
- [9] K. Roth, F. Goldstein, and J. Kleinman. Energy consumption by office and telecommunications equipment in commercial buildings—volume I: Energy consumption baseline. Prepared by Arthur D. Little for the U.S. Department of Energy. A.D. Little Reference no. 72895-00., 2002.
- [10] R. S. Tucker, R. Parthiban, J. Baliga, K. Hinton, R. W. A. Ayre, and W. V. Sorin. Evolution of WDM optical IP networks: A cost and energy perspective. *Journal of Lightwave Technology*, 27(3):243–252, 2009.
- [11] M. Webb. Smart 2020: Enabling the low carbon economy in the information age. <http://www.theclimategroup.org>, 2008.