

Green Software Engineering: The Curse of Methodology

Abram Hindle¹

¹Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada, abram.hindle@ualberta.ca

ABSTRACT

Computer Science often seems distant from its natural science cousins, especially software engineering which feels closer to sociology and psychology than to physics. Physical measurements are often rare in software engineering, except in a few niches. One such important niche is that of software energy consumption, green mining, green IT, and sustainable computing, which all fall under the umbrella of green software engineering.

With the physical measurement of energy consumption comes all of the limitations of measurement and experimentation that exist in the natural sciences and engineering. Issues abound, from attribution of energy use, isolation of components, to replicable experiments. These get further complicated by cloud computing whereby systems are virtualized and attribution of resource usage is a serious issue.

Thus in this work we discuss the current state of software energy consumption, and where will it go.

Keywords: Green Mining, Energy-aware mining, green software engineering, greenIT, sustainability, sustainable software

1 INTRODUCTION

Fundamentally all computation comes at a cost. It is of no surprise that electrical measurements of work correspond to computation as well. With the availability of smart phones, heavily parallelizable clusters, cloud-mad data centers, software and energy interact more readily than ever before. Energy comes at a cost to generate, to deliver, and to store. Delivery requires infrastructure, storage requires materials for batteries, and the by-product of energy consumption, heat, requires cooling. While hardware primarily consumes energy, it can only be as efficient as the software that commands it.

Software's interaction with energy is split among many contexts. Two important contexts are mobile applications and software services hosted within data centers. Other contexts include embedded sensors, the desktop.

Data Centers are limited by energy in terms of power limits of rack power systems as well as cooling. Typically energy accounts for 50% to 100% of the cost of purchased equipment over the equipment's lifetime [29]. Racks have limited energy hookups. Only so many 600W+ units may be powered. Furthermore for every unit put in, the waste heat must be addressed. A data center with poor cooling will pay even more in energy consumption due to the excessive use of the cooling system of each hosted server. Typically services offered by a data-center are software services and in many cases the services are dynamically provisioned on virtual machines or containers.

Mobile applications are slightly different, their availability is affected by the availability of energy. Without battery energy left, no application could survive. The energy used by mobile devices is negligible, usually less than CFL light bulb while charging – yet the batteries are composed of potentially toxic and costly materials. Reducing mobile energy use leads to longer battery lives, combined with reduced battery replacement, and more availability for the end-user.

These two contexts alone motivate the importance of energy efficiency and the study of software energy consumption. Hardware creators can only do so much until it becomes the responsibility of the software developer to develop software in an energy efficient and sustainable manner. Software engineering researchers have noticed this problem and have taken up the torch, thus accepting their responsibility for some of the energy consumption costs of applications.

In this paper I introduce software energy consumption, discuss past and present challenges, works, and issues relevant to software engineering communities. Then I discuss my predictions for the future of software energy consumption, and where such research will go in the next decade.

2 BACKGROUND

Energy is the effort expended to complete a task. For electricity we typically use joules (J) (and sometimes watt-hours) to indicate the energy that a task takes. Power is the instantaneous rate of energy consumption or the work that is being done. Typically power is measured in Watts, which is the instantaneous amount of work done. The multiplication of Power by time is energy, or energy is the integration of power over time. For long running services power is a common measure (average energy use per second), where as for tasks with a clear beginnings and ends energy is a common measure – the cost of a task.

Software energy consumption is a kind of performance and thus part of the non-functional requirement (NFR) of efficiency. Generally we want software to consume the least amount of energy and have low power use. As a kind of performance software energy consumption testing is typically a kind of performance regression testing. This kind of testing typically is evolutionary [Hindle] and seeks to compare performance between versions on the tasks of the product.

Benchmarking is another kind of regression testing that the allows comparison between products. Benchmarking is less about comparing versions, than it is about comparing different implementations of the same task. Some energy research seeks to benchmark applications for energy efficiency [56, 28].

3 THE PAST AND PRESENT

There are many issues in energy-aware software engineering, green software engineering, and green-mining [Hindle] ranging from the complexity of testing, dependency on hardware, dependency on the environment, or dependency on software. The generalizability of this research is hampered by the complexity, and the lack of availability of tools. All of these issues compound the difficulty of applying static or dynamic analysis to software energy traces.

3.1 Ranking Applications by Energy Efficiency

Consumers tend to lack information about software. When a consumer buys an oven, the oven is ranked by its energy efficiency. What if the same ranking existed for software? Research exists that seeks to rank software in terms of energy efficiency much in the way that energy stars rates and ranks consumer products [56].

3 main challenges that face ranking software by energy efficiency include:

- Software executes more than 1 task
- Fair benchmarks for multiple products
- Efficiency per platform (Software/Hardware)

The challenges that software faces versus ovens is that software does multiple tasks and some of these tasks are quite distinct. Thus without agreement about the shared tasks, not every feature or task can be compared between products. This is complicated by the lack of standardization. Figure 1 demonstrates an example of how application energy rankings could be integrated into an App store: different apps that fulfill the same tasks could be measured on a per task basis, as to allow consumers to see the difference efficiencies each app has to offer. The ranking must be on a per-task basis rather than a holistic basis, as certain apps will focus on specific behaviours and tune themselves for it.

Not all hardware is created equal and not all software works the same on all hardware. Thus if software executes differently on different hardware it must be measured, simulated, or estimated on that hardware. Thus when one ranks software, should it be invariant of the hardware? If so how should one normalize it [56, 25]? Zhang [56] poses a method of normalizing across platform with linear scaling – this technique is used by the Green-Miner [25] to normalize measurements from different smart-phones under test..

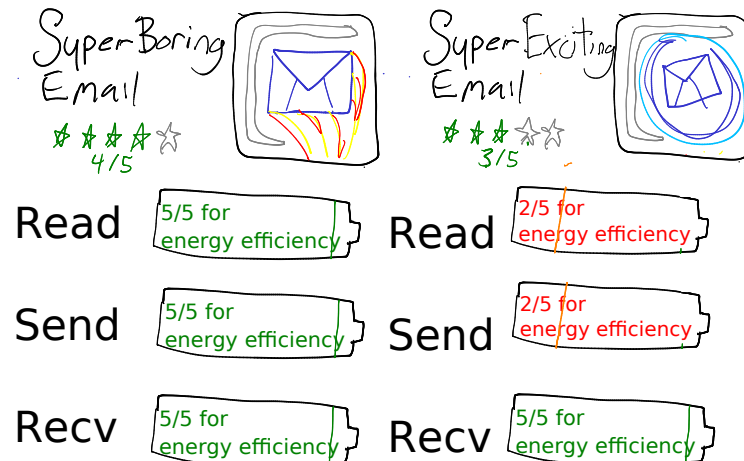


Figure 1. Storyboard mock-up of future App ratings in the App Store

Furthermore software energy consumption is not stable across versions [Hindle, 23, 24, 48]. Testing a single build of the software might not be enough, a partial or entire energy consumption profile should probably be built. Users of an app will probably care if there is change in energy efficiency [53].

There are three main works which attempt to benchmark or pose the software energy problem in a similar to Energy Star star rankings. Amsel et al. [10] discussed green tracker and compared web browsers for energy consumption performance. Zhang et al. [56, 54] describe the differences between applications that do the same tasks yet perform differently in terms of energy use. They propose SACER/Green Star, a method to measure and compare and rank applications much like energy star. Ecodroid [28] employed static and dynamic analysis to automatically rank applications – they did not use task based measurement.

Thus consumers need access to energy performance information and the app-store might be the perfect place to display such details, as shown in Figure 1.

3.2 Generalizable Models

One overarching goal of much of the energy consumption research is to produce a model that generalizes across many applications. The use-case of such a model is that developers do not have access to expensive hardware and cannot accurately measure the energy consumption of their applications – thus they must rely upon estimations based on different kinds of analyses and models. But these generalizable models suffer from the range of hardware, operating systems, environment, software domains, and versions of software.

In the mobile arena the wide-range of screen-sizes, memory sizes, and kinds of processors tends to hamper generalizable models. Furthermore the Android ecosystem is considered fragmented in terms of hardware and software [19]. Server-side, the difference between different x86 manufacturers chips can be significant. Thus there is much hardware variation in terms of primary components. This ignores the range of peripherals and I/O devices that be prevalent on mobile devices: GPS, motion sensor, touch sensors, light sensor, accelerometers, cameras, bluetooth, wifi, etc.

The issue of different operating systems is also relevant, Windows and Linux do not share the same code base and handle energy management differently. Android includes customizations distinct from Linux as well. Furthermore there are different versions and distributions of Linux, Windows, Android, OSX, and iOS. Thus measurements from one environment might not hold for another.

Furthermore generalizable models suffer from a lack of data. Energy traces are not prevalent in the operational data within Github git repositories or other publicly available repositories. Continuous integration tools tend not to measure or estimate energy. One possible repository of energy data, from the Carat project [45], is not publicly accessible to developers and researchers. Thus there is a real lack

of software energy data available to researchers and what is available is not very comprehensive [55].

Too many models are very hardware dependent. For instance the models of Pathak et al. [44] require arduous component modeling. Where as much of the work of Hindle et al. [25] only is tested on a small subset of Android devices and platforms. Karan et al. [8, 7] built upon Pathak et al. [44] work and suggested a rule-of-thumb model based on system calls that works relatively well: if the system call count significantly changes between versions then the energy use between versions changes significantly. This rule-of-thumb specifically avoids mis-classifying many of the 90% of changes which do not affect the energy profile of an application. This model was extended and generalized as a regression problem by Shaiful et al. [14] who estimate not only change, but actual energy usage. The system call based models are general and relative to the products themselves, they model applications that face the user quite well. Yet the models fail to account for CPU use effectively without the use of counters.

While these models range in deploy-ability and usefulness, will programmers want to apply these models?

3.3 How knowledgeable are programmers about energy?

While software engineering researchers are interested in software energy consumption are programmers knowledgeable or aware? Currently in 2015, the answer is a resounding, “not really.”

Pinto et al. [46] studied StackOverflow [2], a question-answer site for programmers. They found that energy related questions were poorly answered and that many questions were asked.

Pang et al. [Pang et al.] followed up and surveyed and interviewed programmers. Pang et al. found that programmers surveyed did not have much experience with software energy consumption. Not only did they lack experience but rarely were they asked to address software energy consumption. The programmers also said they would consider energy consumption when buying a mobile device.

Wilke et al. [53] corroborates the view of these developers. They found that App ratings on Google Play Store suffered when user commented on poor energy consumption behaviour. Khalid et al. [31] have made similar observations.

Many works aim to help developers by finding specific energy bugs [37]. Manotas et al. [38] provide suggestions for energy efficient collections. Others have suggested using genetic programming to optimize already existing programs [13]. Some works discuss the cost of using libraries that provide advertisements [17] and some works describe the costs and benefits of ad-blocking with respect to energy consumption [47]. Some work aims to optimize display usage through color choices [36]. While others help to provide feedback to developers if anything has changed [7].

3.4 Measurement

Software energy consumption needs to be measured. Many researchers use time as a proxy but for idle applications this might not be appropriate [56].

Hindle et al. [25] describes the green miner, a hardware-based continuous regression test framework. The green miner is a software queue for tests that enables deployment of tests onto a series of Android phones, enabling parallel execution of tests. A similar work was presented by Banerjee et al. [11] whereby they use physical instrumentation.

Li et al. [34] tried to use high frequency measurements to attribute energy use to particular source lines, while Hao et al. [20] applied program analysis to estimate energy use. Gupta et al. [18] attempted to correlate measurements with library usage.

LessWatts.org from Intel develops and provides the PowerTop tool to estimate energy use at run-time based on ACPI information [26].

Thus there are many hardware methods of measuring energy but many are too complicated for programmers so they opt for either server hardware with instrumentation or for estimations from ACPI.

4 IMMEDIATE CHALLENGES

The field is currently immature. There is a lack of:

- Shared tools;
- Shared datasets;
- Benchmarks datasets;

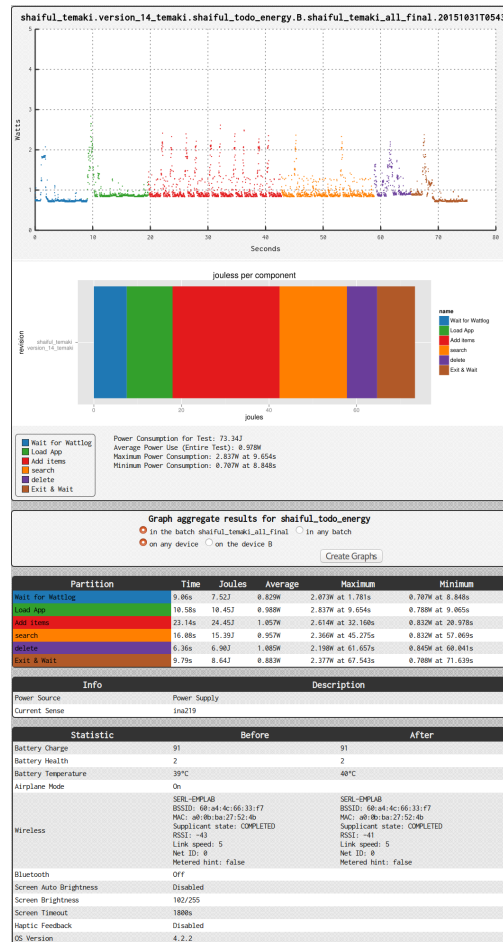


Figure 2. Green Miner Example Test Run

- Agreement on methodology;
- Coherent community;
- Methodological discussion about addressing threats;
- Lack of agreement on methodological threats;

All of these issues pile up into immediate challenges that face current software energy research as well as those works of the future. This section tries to illustrate the potential issues that the field faces.

4.1 Lack of Data

Currently there is very little data for researchers to work with. The current pattern is for researchers to setup a test and measure everything themselves. This is different than a lot of mining software repositories [21, 41] research – there is not a repository available here, unlike other kinds of dynamic analysis such as crash reports.

Each energy trace, such as the one depicted in Figure 2, from 1 test-run, contains many measurements of the same test over time. Typically these tests are re-run to address error in physical measurement and the environment. The re-running of the tests leads researchers to summarize the distribution of measurements leading to a collapse in the amount of data. 5000 kilobytes of energy measurements can be quickly collapsed down into 10 to 40 rows of summary statistics about the test-runs. The data is effectively limited by the researchers time and ability to run all the necessary tests. As dynamic analysis and tracing is typically used it takes a lot of run-time to execute tests.

Thus dynamic analysis takes time, but there is also a limitation in the number of applications that meet the requirements of the research. For instance if one is focusing on Android applications with available source code, the set of testable applications is quite limited. Furthermore given those applications very few come with tests so tests need to be generated.

Thus the field lacks data due to a lack of collecting existing data, a lack of sharing of existing data, a lack of appropriate applications to test, and a lack of available tests for these applications to enable dynamic analysis.

4.2 CPU is not enough

Many works – especially in the area of distributed computing – simply relate CPU time to energy [50, 51, 30, 52]. While this is correct for CPU bound processes, many applications are not CPU bound. Some are not even IO bound, they are event bound. Thus they have an idle cost, but not much in the way of CPU work. They might induce IO when woken up but for the most part most user facing applications are quite idle. Furthermore CPU use does not necessarily represent the activity of peripherals. In the case of GPU clusters, CPU use could be almost irrelevant as the GPU would be the dominant energy consumer.

4.3 Virtualization

If one cares about sustainability [12, 32] and reducing the global energy consumption of computing, virtualization cannot be ignored. Many services online are virtualized, running on a virtual machine in the cloud, or within a container of a container service.

Measuring VMs and containers is quite difficult as resources are not equally shared [40]. Many clouds use over-subscription, whereby resources are over promised to many services with the hope that these services do not need all of these resources at the same time. With virtual machines CPU is often over subscribed while for containers both CPU and memory are oversubscribed.

With most of the worlds services running within data-centers any savings in the resources used by a service has a potential for saving energy: when the CPU, memory or peripherals such as hard-drives. The less use, the less heat, resulting in less cooling and more savings.

Currently there is some work on attributing energy consumption to virtualized machines [39, 40], but it is just the beginning of such research. Little to no work has been done to estimate the energy use of services within containers such as docker. At the moment it is very difficult to estimate the energy cost of a task that is virtualized or container-ized.

4.4 Multi-version analysis necessary

In a software product 1 change can fundamentally change the performance of the software [24]. The same is true to software energy performance. Thus to characterize the performance of a product by only its latest version is unfair. Projects such as Hadoop have had issues with performance regressions – Hadoop 2 on smaller clusters often performs worse on the same task than Hadoop 1 due to resource management [27].

Software is more than the just the current version – most developers exist in a continuously evolving context producing many builds and many versions at once. Versioning is a constant problem within software distributions such as Debian [16] or Ubuntu.

Thus 1 change can have a significant effect on performance of software, and that change might be required to address a raft of issues. This performance changing commit could negate past results [48]. Romansky et al. [48] investigated if every revision needed to be measured or just some, and found that the performance changing commits generally were either immediately corrected or initiated a long plateau of similar performance across subsequent versions.

For instance if one tests if a refactoring was impactful and just look at refactoring commits – what happens afterward [49]? Is the behaviour stable? Was there a bug? Just looking at the immediate before and after commits might not be enough to determine the actual effect of a design pattern or a refactoring.

Thus multi-version analysis adds more software to analyze, enables more data to be collected but also adds robustness against some threats to validity.

4.5 Non-determinism of hardware state

One problem with modern software is that it runs on complicated platforms. Furthermore the realistic scenario of running applications on various hardware can be complex.

One such difficulty is hardware power saving functionality which enables CPUs to use more or less voltage or to change clock rates, as well as enable low power or high latency mode in peripherals. In real world use these options are often on. Experimentally they are often turned off or set to a constant setting, but not always, and it isn't always beneficial to create such an artificial setting [56].

Even if we start the CPU in a certain state for a test, the test input might induce a different CPU state [15].

4.6 Non-determinism of software state

Software can exhibit non-determinism. Mobile platforms are quite adaptive and small changes in the environment can result in different behaviour. Furthermore events within the operating system are not always controllable or deterministic. Network communication is not deterministic as well, thus one serious confound is the non-determinism of software state in the OS alone. When this is combined with long running services non-determinism abounds.

4.7 The need for science

Within this section I have brought up many issues, but how many have empirical evidence to demonstrate the dangers or costs of ignoring these issues. What if the measurements are strong against noise after enough runs are executed? Perhaps after 40 runs the initial state does not matter.

Furthermore in terms of publications regarding energy what we need is more science. Not every energy consumption paper can be a tool paper. Sometimes a result or technique could be integrated into an existing tool or be deployed as tool, but that is a high bar when most bug prediction work never produces a deployable tool. Currently for software energy consumption research the bar is quite high, there is little taste or favour to scientific publications, such as the work of Li et al. [33], Romansky et al. [48], or Linares-Vásquez et al. [35], rather than tool publications such as Green Advisor [7]. Yet communities such as MSR [41] and ESEM [ese] promote this kind of research with other kinds of nonfunctional requirements (NFRs) such a performance or maintainability have much empirical work behind them.

4.8 Community

Not only do we need more science in software energy consumption, we need more community support. There are some industry wide groups that discuss Green IT, sustainable IT infrastructure [42, 9?]. As of writing there is a smattering of specifically sustainability and green IT conferences, none are truly coherent when it comes to software energy consumption, as each venues have different goals [3, 5, 4].

There are workshops such as GREENS [4], but the motivation to publish at GREENS is low when one can submit papers to other venues which garner more recognition. Furthermore in software engineering venues one could perceive there is a lack of a knowledge regarding software energy consumption, which mirrors the current reality of programmers' knowledge [Pang et al.].

All of these factors lead to an software energy research diaspora, where sub-communities are made and results are quietly published but not noticed by other communities. It is almost as if researchers are publishing into a vacuum whereby other researchers do not see each other.

4.9 Impossible Bar to reach, or potential paper?

These limitations should not scare anyone away from the field, in fact for empiricists and experimentalists these are papers in waiting. Many of these issues might not pose as significant effect as we think, or their effect might be avoided or controlled for methodologically. for any of these issues there is an impactful avenue of research found by asking the question, "Do we have to address this potential pitfall?" Potential paper topics that anyone could address:

- How to model the difference in performance between OSX, Windows, and Linux, or iOS and Android.
- The effects of state on repeated tests.
- The effects of differing wifi-state on energy tests.
- Temperature, mobile devices, and energy consumption.
- Effective version test selection.
- Effect of hyper-visors on energy consumption.
- How to control for non-determinism in disk I/O.
- How to control for non-determinism in network I/O.
- How to account for different background cloud utilization.
- How many versions do we really need to measure?
- What is the effect of isolation on our tests?
- Is there difference between software instrumentation and human input?
- What issue affect ACPI energy estimates?
- How many energy measurements do I need per version?
- Does the quality of test cases matter when measuring energy?

Thus these limitations should spur scientific research into the effects and costs of addressing and ignoring the issues brought up in this section. These limitations and proposal for future work segways into what will be expected in the future, next.

5 THE FUTURE

In this section I lay out my prediction for the future of software energy consumption research.

5.1 Multi-version analysis will be expected

In the future researchers will engage in multi-version analysis of performance and energy consumption. They will use multi-version analysis because a primary concern of energy consumption is performance regression e.g., "has performance worsened?" Multi-version analysis will also be used to increase the generality and robustness of their research. Instead of making claims about 1 snapshot of a program's performance researchers will establish the profile [48] of a program's performance. This is especially important in research that engages in factor analysis as it provides more measurement of the system but also protects against spurious factors being reported as significant.

5.2 An end to developer measurement

In the future we will never expect a developer to physically benchmark or measure their software. This will be the realm of technicians and researchers, not developers. Physical measurement will be avoided by easy to access services, better models of services and apps, and better software energy estimation frameworks that appropriately address the needs and limitations of developers.

No hardware Instrumentation: hardware is expensive, and it requires much knowledge and training to address hardware measurement. The developer of the future will not have to rely on expensive testing hardware, or the questionable measurements of their ACPI chip-sets.

Access to Hardware Regression Testing Services: if developers truly need physical measurement they will be able to outsource it. We expect in the future that services will be available that will be like the Green Miner [25] – developers will submit applications, specify the hardware to test on and simply wait for a result back from the framework. No awkward setup, no difficult testing. What physical measurements are made will probably be integrated into even better models.

For verification programmers will have the option to submit their application to a continuous integration, testing, and deployment service that will provide some hardware based measurement. Yet for the most part the future engineers need not worry about actually measuring software energy consumption.

5.3 Online Shared Repository

The future holds promise as large open shared repositories of dynamic traces of energy consumption are curated. Different platforms, different applications, different tests and different runs all aggregated in large online repositories of data. These shared repositories would allow the curation of community tuned models of energy consumption. Much like the PROMISE repository [6].

The repositories would allow the hours and hours that practitioners and researchers spending benchmarking and testing software to be used to develop better models. The variation in available runs alone would be intensely beneficial. Even the tests themselves could be shared, enabling further collaboration.

The future is crowd sourced and open shared traces available to all.

5.4 Cloud and Container Estimation

One of the largest concerns in the future will be the sustainability of software services [12]. There will be pressure from social causes, combined with carbon taxes and worldwide sustainability pressure to reduce carbon emissions. This will affect the software as a service market. Furthermore companies will be asked to estimate their energy use so they can argue if they are green and sustainable. The requirements of sustainable engineering will prompt for developer awareness of the issues and the ability to estimate the impact of services.

Thus all the difficulties mentioned before in Section ?? will conflict with the requirement of energy estimation – programmers of the future will have to estimate or measure the sustainability of distributed software ecosystems. These ecosystems might not be fully subscribed to – many will be relatively idle services – but such estimates of energy consumption will be required.

Programmers will submit usage scenarios, configuration, and their software to a testing service that will estimate the energy usage of their services at different loads. This will require a new kind of continuous integration and deployment software to operate. Furthermore such a system will need measurement instruments to enable measuring, modeling, and estimation of software energy consumption.

5.5 Budgeted Software and Energy Requirements

More managers and customers will explicit request software energy consumption be addressed in their applications. As Lago et al. [32] suggest, sustainability will be perceived as a software quality issue.

This will imply that not only will energy requirements exist, but likely services will be granted energy budgets that they have work within. It is likely that services provided by Amazon AWS and other cloud providers will start explicitly charging for energy rather than just CPU, Memory, IO and network usage. With this change in pricing part of the requirements elicitation process will be to define the energy budgets of a service.

5.6 Education

As Pang et al. [Pang et al.] found, developers are not very aware of software energy consumption and thus if they were asked to act on it, as not very educated on the causes of software energy consumption.

As of writing this software energy consumption is a niche topic rarely taught to computer scientists or software engineers – although somewhat addressed in electrical engineering and computer engineering curriculums. Developers of the future will face the demand for sustainable systems, thus computer science and software engineering curriculums will change to address green software engineering.

6 CONCLUSIONS

Software energy consumption research currently faces many challenges and threats to validity. Among these are attribution of energy use to processes, measurement of virtualized or containerized processes, estimation of energy use, and lack of freely available software energy tools that do not require physical hardware.

Methodologically software energy research is plagued by threats to generalizability regarding OS versions, application versions, environments, the variety of available hardware, and a lack of recorded operational data and measurements. Currently there is a very limited research community who has done little to share data and tools. This is further compounded by a fragmentation of the community across numerous small conferences and workshops.

The future holds much promise for the field of software energy consumption as there are many hard challenges that need to be addressed. The programmers of the future will face sustainability as a requirement and will have to design software with energy efficiency in mind. These programmers not only will receive education, instruction, and training, they will have at their disposal powerful models and tools that are integrated into their development environment ever ready to provide them with software energy awareness when they need it.

REFERENCES

[ese]

- [2] (2008). Stack Overflow. <http://stackoverflow.com>.
- [3] (2012). *GREENCOM '12: Proceedings of the 2012 IEEE International Conference on Green Computing and Communications*, Washington, DC, USA. IEEE Computer Society.
- [4] (2015). Fourth international workshop on green and sustainable software (GREENS 2015).
- [5] (2015). *Proceedings of the Sixth International Green and Sustainable Computing Conference (IGSC'15)*.
- [6] (2015). The promise repository of empirical software engineering data.
- [7] Aggarwal, K., Hindle, A., and Stroulia, E. (2015). Greenadvisor: A tool for analyzing the impact of software evolution on energy consumption. In *31st IEEE International Conference on Software Maintenance and Evolution*. IEEE Computer Society.
- [8] Aggarwal, K., Zhang, C., Campbell, J. C., Hindle, A., and Stroulia, E. (2014). The power of system call traces: Predicting the software energy consumption impact of changes. In *Press of the 2014 Conference of the Center for Advanced Studies on Collaborative Research, IBM Corp.*
- [9] Alliance to Save Energy (2007). PC Energy Report 2007: United States. Technical report, 1E.
- [10] Amsel, N. and Tomlinson, B. (2010). Green tracker: a tool for estimating the energy consumption of software. In *Proceedings, CHI EA*, pages 3337–3342, New York, NY, USA. ACM.
- [11] Banerjee, A., Chong, L. K., Chattopadhyay, S., and Roychoudhury, A. (2014). Detecting energy bugs and hotspots in mobile apps. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 588–598. ACM.
- [12] Becker, C., Chitchyan, R., Duboc, L., Easterbrook, S., Penzenstadler, B., Seyff, N., and Venters, C. C. (2015). Sustainability design and software: The karlskrona manifesto. In *37th IEEE/ACM International Conference on Software Engineering, ICSE 2015, Florence, Italy, May 16-24, 2015, Volume 2*, pages 467–476. IEEE.
- [13] Bruce, B. R., Petke, J., and Harman, M. (2015). Reducing energy consumption using genetic improvement. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation, GECCO '15*, pages 1327–1334, New York, NY, USA. ACM.
- [14] Chowdhury, S. A., Kumar, L. N., Jabbar, M. T. I. M. S. M., Sapra, V., Aggarwal, K., Hindle, A., and Greiner, R. (2015a). A system-call based model of software energy consumption without hardware instrumentation. In *Proceedings of the Sixth International Green and Sustainable Computing Conference (IGSC'15)*.

- [15] Chowdhury, S. A., Sapra, V., and Hindle, A. (2015b). Is HTTP/2 more energy efficient than HTTP/1.1 for mobile users? *PeerJ PrePrints*, 3:e1280.
- [16] Claes, M., Mens, T., Di Cosmo, R., and Vouillon, J. (2015). A historical analysis of debian package incompatibilities. In *Mining Software Repositories (MSR), 2015 IEEE/ACM 12th Working Conference on*, pages 212–223. IEEE.
- [17] Gui, J., Mcilroy, S., Nagappan, M., and Halfond, W. G. J. (2015). Truth in advertising: The hidden cost of mobile ads for software developers. In *37th IEEE/ACM International Conference on Software Engineering, ICSE 2015, Florence, Italy, May 16-24, 2015, Volume 1*, pages 100–110. IEEE.
- [18] Gupta, A., Zimmermann, T., Bird, C., Naggapan, N., Bhat, T., and Emran, S. (2011). Energy Consumption in Windows Phone. Technical Report MSR-TR-2011-106, Microsoft Research.
- [19] Han, D., Zhang, C., Fan, X., Hindle, A., Wong, K., and Stroulia, E. (2012). Understanding android fragmentation with topic analysis of vendor-specific bugs. In *Reverse Engineering (WCRE), 2012 19th Working Conference on*, pages 83–92.
- [20] Hao, S., Li, D., Halfond, W. G. J., and Govindan, R. (2013). Estimating Mobile Application Energy Consumption using Program Analysis. In *Proceedings of the 2013 International Conference on Software Engineering, ICSE '13*, pages 92–101.
- [21] Hassan, A. E. (2008). The Road Ahead for Mining Software Repositories. In *Proceedings of the Future of Software Maintenance (FoSM) at the 24th IEEE International Conference on Software Maintenance*, pages 48–57.
- [Hindle] Hindle, A. Green mining: A methodology of relating software change to power consumption. In Submission to MSR 2012, <http://softwareprocess.es/a/green-change-e.pdf>.
- [23] Hindle, A. (2012). Green mining: Investigating power consumption across versions. In *Proceedings, ICSE: NIER Track*. IEEE Computer Society. <http://url.ca/84vh4>.
- [24] Hindle, A. (2015). Green mining: a methodology of relating software change and configuration to power consumption. *Empirical Software Engineering*, 20(2):374–409.
- [25] Hindle, A., Wilson, A., Rasmussen, K., Barlow, J., Campbell, J., and Romansky, S. (2014). Green-Miner: A Hardware Based Mining Software Repositories Software Energy Consumption Framework. In *Mining Software Repositories (MSR), 2014 11th IEEE Working Conference on*. ACM.
- [26] Intel (2011). LessWatts.org - Saving Power on Intel systems with Linux. <http://www.lesswatts.org>.
- [27] Ivanilton Polato, Denilson Barbosa, A. H. and Kon, F. (2015). Hadoop branching: Architectural impacts on energy and performance. In *Proceedings of the Sixth International Green and Sustainable Computing Conference (IGSC'15)*.
- [28] Jabbarvand, R., Sadeghi, A., Garcia, J., Malek, S., and Ammann, P. (2015). Ecodroid: an approach for energy-based ranking of android apps. In *Proceedings of the Fourth International Workshop on Green and Sustainable Software*, pages 8–14. IEEE Press.
- [29] Jed Scaramella, M. E. (2007). Solutions for the datacenter's thermal challenges. <http://whitepapers.zdnet.com/abstract.aspx?docid=352318>. IDC white paper.
- [30] Joseph, R. and Martonosi, M. (2001). Run-Time Power Estimation in High Performance Microprocessors. In *Proceedings of the 2001 international symposium on Low power electronics and design, ISLPED '01*, pages 135–140, New York, NY, USA. ACM.
- [31] Khalid, H., Shihab, E., Nagappan, M., and Hassan, A. E. (2014). What do mobile app users complain about? A study on free iOS apps. *Accepted to be published in IEEE Software*.
- [32] Lago, P., Koçak, S. A., Crnkovic, I., and Penzenstadler, B. (2015). Framing sustainability as a property of software quality. *Commun. ACM*, 58(10):70–78.
- [33] Li, D., Hao, S., Gui, J., and Halfond, W. G. J. (2014). An empirical study of the energy consumption of android applications. In *30th IEEE International Conference on Software Maintenance and Evolution, Victoria, BC, Canada, September 29 - October 3, 2014*, pages 121–130. IEEE Computer Society.
- [34] Li, D., Hao, S., Halfond, W. G., and Govindan, R. (2013). Calculating source line level energy information for android applications. In *Proceedings of the 2013 International Symposium on Software Testing and Analysis*, pages 78–89. ACM.
- [35] Linares-Vásquez, M., Bavota, G., Bernal-Cárdenas, C., Oliveto, R., Di Penta, M., and Poshyvanyk, D. (2014). Mining energy-greedy api usage patterns in android apps: An empirical study. In *Proceedings of the 11th Working Conference on Mining Software Repositories, MSR 2014*, pages 2–11, New

- York, NY, USA. ACM.
- [36] Linares-Vásquez, M., Bavota, G., Cárdenas, C. E. B., Oliveto, R., Di Penta, M., and Poshyvanyk, D. (2015). Optimizing energy consumption of guis in android apps: A multi-objective approach. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2015*, pages 143–154, New York, NY, USA. ACM.
- [37] Liu, Y., Xu, C., and Cheung, S. (2015). Diagnosing energy efficiency and performance for mobile internetware applications: Challenges and opportunities. *Software, IEEE*, PP(99):1–1.
- [38] Manotas, I., Pollock, L., and Clause, J. (2014). Seeds: A software engineer’s energy-optimization decision support framework. In *Proceedings of the 36th International Conference on Software Engineering, ICSE 2014*, pages 503–514, New York, NY, USA. ACM.
- [39] Moghaddam, F. A., Geenen, T., Lago, P., and Grosso, P. (2015a). A user perspective on energy profiling tools in large scale computing environments. In *2015 Sustainable Internet and ICT for Sustainability, SustainIT 2015, Madrid, Spain, April 14-15, 2015*, pages 1–5. IEEE.
- [40] Moghaddam, F. A., Lago, P., and Grosso, P. (2015b). Energy-efficient networking solutions in cloud-based environments: A systematic literature review. *ACM Comput. Surv.*, 47(4):64.
- [41] MSR (2013). Mining Software Repositories. www.msrconf.org.
- [42] Murugesan, S. (2008). Harnessing Green IT: Principles and Practices. *IT Professional*, 10(1):24–33.
- [Pang et al.] Pang, C., Hindle, A., Adams, B., and Hassan, A. E. What do programmers know about the energy consumption of software? *IEEE Software*.
- [44] Pathak, A., Hu, Y. C., and Zhang, M. (2011). Bootstrapping energy debugging on smartphones: a first look at energy bugs in mobile devices. In *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*, page 5. ACM.
- [45] Peltonen, E., Lagerspetz, E., Nurmi, P., and Tarkoma, S. (2015). Energy modeling of system settings: A crowdsourced approach. In *Pervasive Computing and Communications (PerCom), 2015 IEEE International Conference on*, pages 37–45.
- [46] Pinto, G., Castor, F., and Liu, Y. D. (2014). Mining questions about software energy consumption. In *Proceedings of the 11th Working Conference on Mining Software Repositories*, pages 22–31. ACM.
- [47] Rasmussen, K., Wilson, A., and Hindle, A. (2014). Green mining: energy consumption of advertisement blocking methods. In Müller, H. A., Lago, P., Morisio, M., Meyer, N., and Scanniello, G., editors, *Proceedings of the 3rd International Workshop on Green and Sustainable Software, GREENS 2014, Hyderabad, India, June 1, 2014*, pages 38–45. ACM.
- [48] Romansky, S. and Hindle, A. (2014). On improving green mining for energy-aware software analysis. In *Press of the 2014 Conference of the Center for Advanced Studies on Collaborative Research, IBM Corp.*
- [49] Sahin, C., Pollock, L., and Clause, J. (2014). How do code refactorings affect energy usage? In *Proceedings of the 8th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement, ESEM ’14*, pages 36:1–36:10, New York, NY, USA. ACM.
- [50] Seo, C., Malek, S., and Medvidovic, N. (2007). An Energy Consumption Framework for Distributed Java-Based Systems. In *ASE ’07*, pages 421–424.
- [51] Seo, C., Malek, S., and Medvidovic, N. (2008). Component-level energy consumption estimation for distributed java-based software systems. In *Component-Based Software Engineering*, pages 97–113. Springer.
- [52] Tiwari, V., Malik, S., Wolfe, A., and Tien-Chien Lee, M. (1996). Instruction level power analysis and optimization of software. *The Journal of VLSI Signal Processing*, 13.
- [53] Wilke, C., Richly, S., Gotz, S., Piechnick, C., and Aßmann, U. (2013). Energy consumption and efficiency in mobile applications: A user feedback study. In *Green Computing and Communications (GreenCom), 2013 IEEE and Internet of Things (iThings/CPSCoM), IEEE International Conference on and IEEE Cyber, Physical and Social Computing*, pages 134–141. IEEE.
- [54] Zhang, C. (2013). The Impact of User Choice and Software Change and Energy Consumption. *University of Alberta, Edmonton, Alberta, Canada*.
- [55] Zhang, C. and Hindle, A. (2014). A green miner’s dataset: mining the impact of software change on energy consumption. In *Proceedings of the 11th Working Conference on Mining Software Repositories*, pages 400–403. ACM.
- [56] Zhang, C., Hindle, A., and Germán, D. M. (2014). The impact of user choice on energy consumption. *IEEE Software*, 31(3):69–75.