



HAL
open science

Energy Efficient Software Development Process Evaluation for MacOS Devices

Shokhista Ergasheva, Dragos Strugar, Artem Kruglov, Giancarlo Succi

► **To cite this version:**

Shokhista Ergasheva, Dragos Strugar, Artem Kruglov, Giancarlo Succi. Energy Efficient Software Development Process Evaluation for MacOS Devices. 16th IFIP International Conference on Open Source Systems (OSS), May 2020, Innopolis, Russia. pp.196-206, 10.1007/978-3-030-47240-5_20 . hal-03647260

HAL Id: hal-03647260

<https://inria.hal.science/hal-03647260v1>

Submitted on 20 Apr 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Energy efficient software development process evaluation for MacOS devices^{*}

Shokhista Ergasheva, Dragos Strugar, Artem Kruglov, and Giancarlo Succi

Innopolis University, Innopolis, Russia
s.ergasheva@innopolis.ru

Abstract. InnoMetrics is the system that aims at collecting software development process metrics in a non-invasive way to access and optimize the development process and its efficiency. This paper demonstrates the development and analysis of energy consumption of MacOS systems based on the software process measurement data. It represents the experience of the development of MacOS system collector and Transfer, in addition to the user interface and early analysis of energy consumption metrics calculations.

Keywords: Energy efficient software · MacOS software development · Software development process metrics.

1 Introduction

The importance of energy efficiency of any system and devices becomes sufficient property of the present Information Technology life. Thus, measuring and what is more further analysis of energy efficiency and its prediction can be managed by monitoring the software development process. Before going deep into the software development process metrics let's define what is the measurement itself. The history of measurement can be divided into two generations. The first generation is the Personal Software Process(PSP) - self improvement process that helps developers to control, manage and improve the way they work. This method can be called as invasive as it requires the direct involvement of participants in the data collection process. Users of the-invasive method should create and print forms in which they log their effort, size and defect information. One obvious downside of this invasive approach is the high overhead cost it entails. The developers should often switch between development tasks and metrics collection tasks, which imposes a high cognitive burden to the developers while the second generation of measurements, non-invasive measurements, do not require manual intervention of the participants during metrics collection [4]. The main aim of the successful use of a non-invasive measurement system is to support the developers, the development process optimization and maintain the data privacy. InnoMetrics is one of the promising software development process data collection

^{*} This research project is carried out under the support of the Russian Science Foundation Grant №19-19-00623.

systems with automatic data collection and transfers to the server for further data analysis. It allows the developers and managers to be aware of the process strength and weaknesses based on the data from the developers performed. The crucial insights are visualized concerning the metrics and analysis. The benefits of such system can be counted as real-time process analysis on a daily basis, the system can be accustomed to different levels of company sizes and audits of the software development process and development itself can be monitored at the same time.

2 Metrics

The first thing in the way to data collection framework was to perform a Systematic Literature Review (SLR) of existing research papers, energy-related metrics and best practices to collect data during the software development process. The Systematic Literature Review surveyed more than 500 studies in the field of software metrics collection and analysis. As a result, at about 170 metrics were derived, and divided into 3 categories as code, product and process metrics. Based on the findings from SLR, we concluded that the mostly unexplored branch of software development metrics is process metrics [8]. Besides, all the tools observed during the research project consider static analysis of the code, nevertheless, the process of developing the software still has no much analysis. The main reason for this insufficient analysis of process measurements is the absence of the tool. Furthermore, the process cost increases, since usually, it requires the developers' participation.

The system of InnoMetrics is developed based on the monitoring of the software development process energy consumption and its efficiency, the developer's productivity. All direct and model-based measurements that involve usage of third-party hardware tools to get energy metrics from various components were considered out of scope of non-invasive software development process analysis approach. The energy consumption of software applications was thoroughly researched and concluded based on the research findings. Battery draining applications result in lower user experience and dissatisfied users [18]. Optimal battery usage (energy usage) is an important aspect that every client must consider.

Application energy consumption is dependent on a wide variety of system resources and conditions. Energy consumption depends on, but is not limited to, the processor, the device uses, memory architecture, the storage technologies used, the display technology used, the size of the display, the network interface that you are connected to, active sensors, and various conditions like the signal strength used for data transfer, user settings like screen brightness levels, and many more user and system settings [18].

For precise energy consumption measurements, one needs specialized hardware [21]. While they provide the best method to accurately measure energy consumption on a particular device, such a methodology is not scalable in practice, especially if such measurements have to be made on multiple devices. Even then, the measurements by themselves will not provide much insight into how

the application contributes to the battery drainage, making it hard to focus on any application optimization efforts.

The InnoMetrics system aims at enabling users to estimate their application's energy consumption without the need for specialized hardware. Such estimation is made possible using a software power model that has been trained on a reference device representative of the low powered devices applications might run on [15, 22, 19, 20, 17]. Based on the findings of the research, metrics like following were investigated [10]:

- Software Energy Consumption (SEC) - the total energy consumed by the software;
- Unit Energy Consumption (UEC) - the energy consumed by a specific unit of the software;

Considering our profiling method and the tools available for us, the ability to attribute the energy consumption was possible only at the process level in coarse granularity. However, the hardware resource usage can fill the gap when it comes to accurately relating Energy Consumption (EC) to individual software elements hence enabling the computation of the UEC.

Profiling the performance requires a basic understanding of hardware components that has to be monitored through "performance counters", which is possible in Windows System. While interpreting performance data for further analysis, the context information has to be taken into account (e.g. hardware-specific details).

To evaluate the Unit Energy Consumption (UEC) the following hardware resources should be monitored:

- Hard disk: disk bytes/sec, disk read bytes/sec, disk write bytes/sec;
- Processor: percentage of processor usage;
- Memory: private bytes, working set, private working set;
- Network: bytes total/sec, bytes sent/sec, bytes received/sec;
- IO: IO data (bytes/sec), IO read (bytes/sec), IO write (bytes/sec).

Attributing some weights to elements of the UEC or by some reliable assumption such as considering the power model to be linear in the nature for each individual component, the SEC Metric is computed.

Besides, the energy usage can also be appraised using Performance Counter, Performance Counter Category and related classes that are available with .NET Framework. To be specific, by analyzing the MSDN documentation to attain the goal of collecting energy related metrics, it was concluded that the information about CPU time, Total Processor Time per process, CPU usage, Memory usage, network usage that Performance Counter provides can be reliable.

The bottleneck in this situation is that it is difficult to match up constantly changing application process IDs and names. The energy consumption of the system depends on a variety of factors that are not limited to those which can be collected using the above mentioned performance classes.

2.1 MacOS energy metrics

In order to obtain energy-related data in MacOS devices, we explored some of the APIs that MacOS provides. The first and most obvious tool was Activity Monitor, an application that comes built-in every MacOS system. One of its aspects is the Energy consumption, shown in Fig. 1.

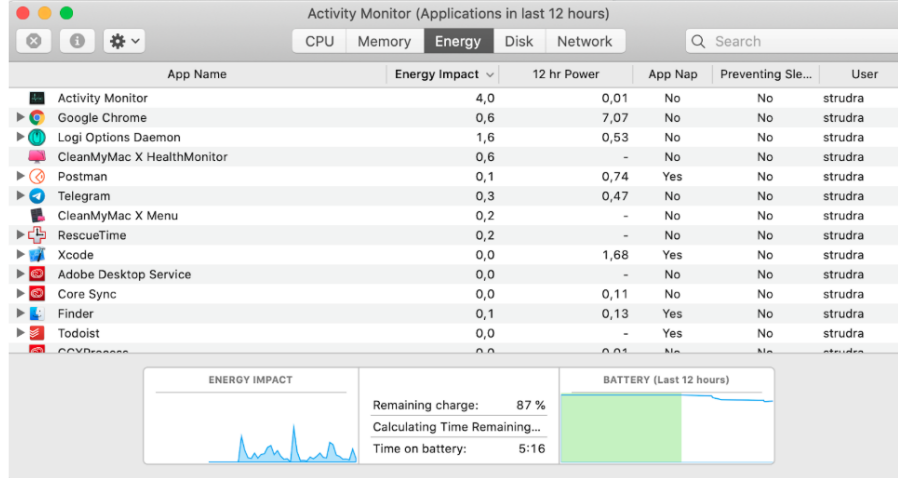


Fig. 1. MacOS energy metrics

The second column, Energy Impact, attracted our attention immediately. We then wanted to figure out what these values (4.0, 0.6, 1.6, etc) represent. It has been then brought to our attention that the definition of Energy Impact is not precisely defined by Apple. According to Activity Monitor’s documentation, the definition of Energy Impact says “Energy Impact: A relative measure of the current energy consumption of the app. Lower numbers are better” [3].

As other documentation says “The Energy tab of Activity Monitor displays the Energy Impact of each open app based on a number of factors including CPU usage, network traffic, disk activity and more. The higher the number, the more impact an app has on battery power” [2]. Both of these are vague, and we needed a concrete way of obtaining this metric’s values.

One article on the Mozilla blog attempted to figure out a formula for this. They indicated that the result of MacOS’s top command line tool which performs periodic measurements of all kinds of metrics; including ones relevant to energy consumption: CPU usage, wakeups and the power measure. The article we mentioned above suggests running the following command to get the above-mentioned information:

```
Yielding the results (trimmed):
top -stats pid,command,cpu,idlew,power -o power -d
```

Table 1. *top* results

| PID | COMMAND | % CPU | IDLE | POWER |
|-------|------------------|-------|------|-------|
| 50300 | firefox | 12.9 | 278 | 26.6 |
| 76256 | Plugin-container | 3.4 | 159 | 11.3 |
| 151 | Coreaudiod | 0.9 | 68 | 4.3 |
| 76505 | Top | 1.5 | 1 | 1.6 |
| 76354 | Activity Monitor | 1.0 | 0 | 1.0 |

They go on to suggest that POWER measure is calculated using a simple formula, and a specific configuration file (tuned for every machine’s architecture). Using the findings from that article, we decided to use some of the most impacting metrics:

- Battery percentage
- Battery status (is charging or not)
- RAM measurement (how much RAM does the active process use)
- vRAM measurement (how much vRAM does the active process occupy)
- % CPU utilized (per process)

All of these metrics were obtained using the macOS command line interface. E.g. to get the current battery percentage we used `pmset -g batt`, and for other measurements we used the `ps -axm -o` command with varying parameters (depending on the use case). It was also possible to use the `top` command, but as we are performing periodic checks anyway, `top` was not necessary. Incorporating these metrics and collection process to the `InnoMetricsCollector` was successful and results are available on GitHub:

<https://github.com/InnopolisUniversity/innometrics-agent-os-mac/>.

3 Collector development

InnoMetrics system contains of the following components [14, 26, 23]:

- Data Collectors - the separate for widespread Operating Systems (Windows, MacOS, Linux) data collecting frameworks;
- Server - which includes analytic module of obtained data referring quantitative and qualitative analysis;
- Dashboard - the component that responsible for visual representation of the analyzed development process data.

The scope of the paper do not consider describing the data collectors for all operating systems, nevertheless MacOS collector will be described thoroughly. MacOS collector includes two separate applications using a service oriented approach[25]:

- InnoMetrics Collector
- InnoMetrics Transfer

InnoMetrics Collector is a daemon process, it runs in the background. The only interaction the user has with the application is that the collector sits in the background and collects data. It displays the user information (Figure 2):

- User Name;
- System version and specifications;
- IP address;
- Mac address of the device;
- as well as the currently open application.

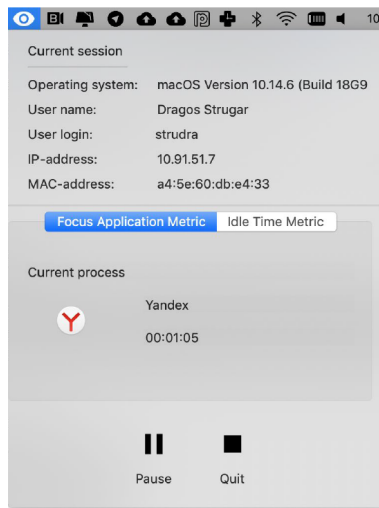


Fig. 2. MacOS Data Collector: Current Process

The Idle Time tab (Figure 3) shows the total idle time and the top 3 application that have been idle for the longest. In addition, the user can stop the process of collection, as well as they can quit the app.

3.1 Metrics being Collected

The data being collected on a per-process basis consists the following types:

- Metric type (app focus/idle);
- Application Name;
- Bundle Domain of the app (e.g.com.apple.finder);
- Bundle Path (where the app is installed, e.g. Applications/Yandex)
- Timestamp of the opening of the app/process;
- Timestamp of the quitting of the app/process;
- Duration of the process open (timestamp end-timestamp start);

It consists of two main areas: the left bar and the process metric list on the right. We have the feature to filter the processes based on keywords, applications, by date. Also, the user can select multiple metrics and hit the Backspace key to delete the metrics before sending them to the server (privacy feature). The right panel presents the data obtained from the collector (connects to a local database where collector was sending data). Finally, user can send the metrics to the server. After selecting the metrics user wants to send to the server, and hitting the “Send Metrics” button, user gets presented with a pop up (Figure 5).

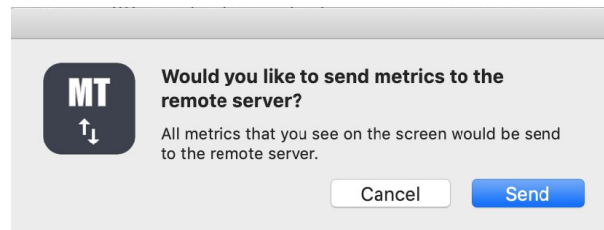


Fig. 5. Popup for sending data to the server

Merging InnoMetrics Collector and InnoMetrics Transfer As mentioned above, there is a dependency between these two applications. One first needs to collect data using the collector and then transfer the data using transfer. The UX would be better if everything was done in the same application, not requiring users to use both of the apps. That is why it was decided to merge the two applications. As a result of this decision, no longer do the users have to click on the submit button to send the data to the server. The data is automatically and periodically being sent as they are using the computer, without any distractions.

Auto Update functionality Desktop applications in general (deployed outside of the Apple App Store) suffer from the inability to update frequently [1]. That is why it was decided to implement the auto update functionality. For this we used the open source framework for auto update of MacOS applications called Sparkle (<https://github.com/sparkle-project/Sparkle>) [13]. It’s a widely known and used framework that allowed to easily update our apps. As the two applications have been merged, there was no need to add the auto update functionality to the Innometrics Transfer. After implementing this functionality, the process of quickly iterating on versions became much quicker and it allows to become more agile when responding to user feedback [16, 24, 5, 11, 6]. Let us stress that the use of an open source application appears particularly suited as it enhances the reliability and the adoption of the system, as evidenced in several existing works [12, 9, 7]

4 Conclusion

To sum up, the system of InnoMetrics was developed for MacOS have been described thoroughly throughout the paper. The energy-related metrics are easily derived from the device with the help of Activity Monitor service in the MacOS. All kinds of metrics were derived and analysed based on the research studies. MacOS framework consists of two basic components: MacOS Collector and Transfer. The details of collecting data and its types, properties and their transfer is integrated in the system. Furthermore, the User Interface developed as a separate system - InnoMetrics Dashboard- for visual representation of the collected data and analysis, was presented. In this report, we have presented: the overall information how to get energy metrics in terms of MacOS systems, a deeper view of how the data collector and transfer operate with the back-end, displayed some results of development phase.

Further improvements of the system will be concentrated on the testing of the system with the functioning industrial companies, adding the external agents for deeper analysis of the software development process like Trello, GitLab, Github, and others. Besides, future work will include exploration of data analysis patterns and best practices for data analysis referring to the data records collected from the Data Collector.

References

1. Mozilla blog: What does the osx energy impact actually measure? <https://blog.mozilla.org/nnethercote/2015/08/26/what-does-the-os-x-activitymonitors-energy-impact-actually-measure/>, accessed: 2020-29-01
2. About mac notebook batteries. <https://support.apple.com/en-au/HT204054> (2020), accessed: 2020-29-01
3. How to use activity monitor on your mac. <https://support.apple.com/en-au/HT201464> (2020), accessed: 2020-21-02
4. Bykov, A., Zaplatnikov, D., Ivanov, V., Rogers, A., Shunovich, A., Sillitti, A., Succi, G., Tormasov, A., Yi, J., Zabirov, A.: A new architecture and implementation strategy for non-invasive software measurement systems. In: Proceedings of the 33rd Annual ACM Symposium on Applied Computing - SAC '18. ACM Press (2018). <https://doi.org/10.1145/3167132.3167327>, <https://doi.org/10.1145/3167132.3167327>
5. Coman, I.D., Robillard, P.N., Sillitti, A., Succi, G.: Cooperation, collaboration and pair-programming: Field studies on backup behavior. *Journal of Systems and Software* **91**, 124–134 (2014). <https://doi.org/10.1016/j.jss.2013.12.037>, <http://dx.doi.org/10.1016/j.jss.2013.12.037>
6. Corral, L., Sillitti, A., Succi, G.: Software Assurance Practices for Mobile Applications. *Computing* **97**(10), 1001–1022 (Oct 2015). <https://doi.org/10.1007/s00607-014-0395-8>, <http://dx.doi.org/10.1007/s00607-014-0395-8>
7. Di Bella, E., Sillitti, A., Succi, G.: A multivariate classification of open source developers. *Information Sciences* **221**, 72–83 (2013)
8. Ergasheva, S., Khomyakov, I., Kruglov, A., Succil, G.: Metrics of energy consumption in software systems: a systematic literature review. *IOP Conference Series: Earth and Environmental Science* **431**, 012051 (Feb 2020).

- <https://doi.org/10.1088/1755-1315/431/1/012051>, <https://doi.org/10.1088/1755-1315/431/1/012051>
9. Fitzgerald, B., Kesan, J.P., Russo, B., Shaikh, M., Succi, G.: Adopting open source software: A practical guide. The MIT Press, Cambridge, MA (2011)
 10. Jagroep, E., Procaccianti, G., van der Werf, J.M., Brinkkemper, S., Blom, L., van Vliet, R.: Energy efficiency on the product roadmap: An empirical study across releases of a software product. *Journal of Software: Evolution and Process* **29**(2), e1852 (Feb 2017). <https://doi.org/10.1002/smr.1852>, <https://doi.org/10.1002/smr.1852>
 11. Janes, A., Succi, G.: *Lean Software Development in Action*. Springer, Heidelberg, Germany (2014). <https://doi.org/10.1007/978-3-642-00503-9>, <http://dx.doi.org/10.1007/978-3-642-00503-9>
 12. Kovács, G.L., Drozdik, S., Zuliani, P., Succi, G.: Open Source Software for the Public Administration. In: *Proceedings of the 6th International Workshop on Computer Science and Information Technologies* (October 2004)
 13. Mathur, A.: *A human-centered approach to improving the user experience of software updates* (2016)
 14. Maurer, F., Succi, G., Holz, H., Kötting, B., Goldmann, S., Dellen, B.: Software Process Support over the Internet. In: *Proceedings of the 21st International Conference on Software Engineering*. pp. 642–645. ICSE '99, ACM (May 1999). <https://doi.org/10.1145/302405.302913>, <http://doi.acm.org/10.1145/302405.302913>
 15. Musílek, P., Pedrycz, W., Sun, N., Succi, G.: On the Sensitivity of CO-COMO II Software Cost Estimation Model. In: *Proceedings of the 8th International Symposium on Software Metrics*. pp. 13–20. METRICS '02, IEEE Computer Society (June 2002). <https://doi.org/10.1109/METRIC.2002.1011321>, <http://dl.acm.org/citation.cfm?id=823457.824044>
 16. Pedrycz, W., Russo, B., Succi, G.: A model of job satisfaction for collaborative development processes. *Journal of Systems and Software* **84**(5), 739–752 (2011). <https://doi.org/10.1016/j.jss.2010.12.018>, <http://dx.doi.org/10.1016/j.jss.2010.12.018>
 17. Pedrycz, W., Russo, B., Succi, G.: Knowledge transfer in system modeling and its realization through an optimal allocation of information granularity. *Appl. Soft Comput.* **12**(8), 1985–1995 (Aug 2012). <https://doi.org/10.1016/j.asoc.2012.02.004>, <http://dx.doi.org/10.1016/j.asoc.2012.02.004>
 18. Rahmati, A., Qian, A., Zhong, L.: Understanding human-battery interaction on mobile phones. In: *Proceedings of the 9th international conference on Human computer interaction with mobile devices and services - MobileHCI '07*. ACM Press (2007). <https://doi.org/10.1145/1377999.1378017>, <https://doi.org/10.1145/1377999.1378017>
 19. Ronchetti, M., Succi, G., Pedrycz, W., Russo, B.: Early estimation of software size in object-oriented environments a case study in a cmm level 3 software firm. *Information Sciences* **176**(5), 475–489 (2006). <https://doi.org/10.1016/j.ins.2004.08.012>, <http://dx.doi.org/10.1016/j.ins.2004.08.012>
 20. Rossi, B., Russo, B., Succi, G.: Modelling Failures Occurrences of Open Source Software with Reliability Growth. In: *Open Source Software: New Horizons - Proceedings of the 6th International IFIP WG 2.13 Conference on Open Source Systems, OSS 2010*. pp. 268–280. Springer, Heidelberg, Notre Dame, IN, USA (May 2010).

- https://doi.org/10.1007/978-3-642-13244-5_21, [http : //dx.doi.org/10.1007/978 - 3 - 642 - 13244 - 5_21](http://dx.doi.org/10.1007/978-3-642-13244-5_21)
21. Saborido, R., Arnaudova, V.V., Beltrame, G., Khomh, F., Antoniol, G.: On the impact of sampling frequency on software energy measurements (Jul 2015). <https://doi.org/10.7287/peerj.preprints.1219v2>, <https://doi.org/10.7287/peerj.preprints.1219v2>
 22. Scotto, M., Sillitti, A., Succi, G., Vernazza, T.: A Relational Approach to Software Metrics. In: Proceedings of the 2004 ACM Symposium on Applied Computing. pp. 1536–1540. SAC '04, ACM (2004). <https://doi.org/10.1145/967900.968207>, <http://doi.acm.org/10.1145/967900.968207>
 23. Sillitti, A., Janes, A., Succi, G., Vernazza, T.: Measures for mobile users: an architecture. Journal of Systems Architecture **50**(7), 393–405 (2004). <https://doi.org/10.1016/j.sysarc.2003.09.005>, <http://dx.doi.org/10.1016/j.sysarc.2003.09.005>
 24. Sillitti, A., Succi, G., Vlasenko, J.: Understanding the Impact of Pair Programming on Developers Attention: A Case Study on a Large Industrial Experimentation. In: Proceedings of the 34th International Conference on Software Engineering. pp. 1094–1101. ICSE '12, IEEE Press, Piscataway, NJ, USA (June 2012). <https://doi.org/10.1109/ICSE.2012.6227110>, <http://dl.acm.org/citation.cfm?id=2337223.2337366>
 25. Sillitti, A., Vernazza, T., Succi, G.: Service Oriented Programming: A New Paradigm of Software Reuse. In: Proceedings of the 7th International Conference on Software Reuse. pp. 269–280. Springer Berlin Heidelberg (April 2002). https://doi.org/10.1007/3-540-46020-9_19, [http : //dx.doi.org/10.1007/3 - 540 - 46020 - 9_19](http://dx.doi.org/10.1007/3-540-46020-9_19)
 26. Vernazza, T., Granatella, G., Succi, G., Benedicenti, L., Mintchev, M.: Defining Metrics for Software Components. In: Proceedings of the World Multiconference on Systemics, Cybernetics and Informatics. vol. XI, pp. 16–23 (July 2000)