

Extending the OpenTelemetry Java Auto-Instrumentation Agent to Publish Green Software Metrics

Andreas Brunnert (brunnert@hm.edu)
Ferdinand Gutzy (ferdinand.gutzy@hm.edu)
Munich University of Applied Sciences HM, Munich, Germany

Abstract

The increasing focus on the carbon emissions of the IT sector has led to the creation of green software metrics that attempt to quantify the environmental impact of software in terms of carbon emissions. One of the most prominent examples is the Software Carbon Intensity (SCI) metric, recently standardized as ISO/IEC 21031:2024, which is defined as the rate of carbon emissions for a software system. Despite the merits of these metrics, development teams often lack the tools to capture them for their software systems. This paper attempts to address this challenge by presenting an extension to the OpenTelemetry Java Auto-Instrumentation Agent to collect not only traditional performance metrics such as response time, throughput, and resource utilization, but also carbon emission data for an instrumented software system.

1 Introduction

At the time of writing there is no commonly agreed upon standard for quantifying the carbon emissions of software systems[2]. Researchers and industry practitioners alike have proposed green software metrics to fill this gap[3]. One of the green software metrics that is already standardized as ISO/IEC 21031:2024¹ is the Software Carbon Intensity (SCI) metric². SCI is defined as follows:

$$SCI = ((E * I) + M)perR$$

E: energy consumed by a software system (kWh)

I: carbon intensity of electricity (gCO₂eq/kWh)

M: embodied carbon, the amount of carbon emitted during the creation and disposal of a hardware device (gCO₂eq)

R: functional unit that describes how the application scales (e.g, transactions per minute or simultaneous users).

This paper describes an extension of the OpenTelemetry (OTel) Java Auto-Instrumentation Agent³

to collect the relevant data to calculate the SCI metric for individual transactions of an instrumented software system. The OTel Java Auto-Instrumentation Agent can be attached to any Java application running on a Java version higher than 8 and automatically collects traces, metrics and logs for commonly used libraries and frameworks automatically. It does this by injecting measurement probes before and after each method invocation to capture so-called spans. In this work we extend the probes of this agent to collect the data to calculate the SCI value for all transactions (e.g., API calls) executed by the instrumented application.

2 Architectural Overview

The general architecture of such an extension has been outlined in our previous work [2] and is shown in Figure 1. We use OpenTelemetry⁴ as standard for capturing and transferring the metrics from the application to Prometheus⁵ as time series database. Grafana⁶ is used to visualize and calculate the carbon emissions per transaction as well as per server, container, or virtual machine involved in processing a transaction over time.

Compared to our previous work in [2], we have simplified the overall architecture by eliminating the need to integrate with an external service to capture carbon emission data. In the current implementation the application agent sends the carbon emission data for each transaction directly along with the other transaction data such as resource demands, response time and throughput.

The reason for this change is that we are now using freely available data⁷ from the Cloud Carbon Footprint (CCF)⁸ project to calculate the emissions, which are bundled with the application and do not change over time.

⁴<https://opentelemetry.io>

⁵<https://prometheus.io>

⁶<https://grafana.com>

⁷<https://github.com/cloud-carbon-footprint/ccf-coefficients>

⁸<https://www.cloudcarbonfootprint.org/>

¹<https://www.iso.org/standard/86612.html>

²<https://sci.greensoftware.foundation>

³<https://github.com/open-telemetry/opentelemetry-java-instrumentation>




Figure 1: Architectural Overview [2]




Figure 2: Resource Profile [1]

3 Collecting Resource Demands

The basis for assessing the carbon emissions of a software system is the amount of hardware resources (i.e., CPU, memory, storage, network) used during the transaction processing as this drives the energy consumption and helps us to calculate the E value of the SCI formula.

In [1] and [2] we have already shown and evaluated how to measure resource demand for a software system. We use the same concepts to capture the resource demand (d) values for CPU (d_{CPU}), memory (d_{MEM}), storage (d_{STO}), and network (d_{NET}) for each method (or transaction) by extending the existing measurement probes injected by the OTEL Java Auto-Instrumentation Agent using a SpanProcessor⁹. The resource demand for CPU is measured as CPU time in ms, while the other demands for memory, storage and network are measured as bytes allocated in memory or read (r) from respectively written (w) to storage or network. We collect this data in so-called resource profiles as shown in Figure 2 in which we store a vector of resource demands per transaction of the software system.

This resource demand data now needs to be translated into carbon emission data as explained in the following section.

4 Calculating Carbon Emissions

Since we cannot use the resource demand data as a direct input for the SCI calculation, we need to translate it into energy consumption and carbon emission data. In order to do this, we follow the recommendations provided in the CCF methodology¹⁰.

For CPU resources, the CCF methodology and our

⁹<https://opentelemetry.io/docs/specs/otel/trace/sdk/#span-processor>

¹⁰<https://www.cloudcarbonfootprint.org/docs/methodology>

implementation rely on a simple model originally published by Etsy as part of their Cloud Jewels conversion factors¹¹, which uses data on the minimum (idle, P_{CPUmin}) and maximum (100% utilization, P_{CPUmax}) power consumption of a processor (e.g., taken from SPECpower results¹²) to calculate the actual power consumption (P_{CPU}) based on the current CPU utilization (CPU_{util}) as follows:

$$P_{CPU} = P_{CPUmin} + (CPU_{util} * (P_{CPUmax} - P_{CPUmin}))$$

The CPU power model also takes into account the data center's Power Usage Effectiveness (PUE), a metric that describes how much power is consumed by the computing equipment in a data center relative to the total power (including cooling and other energy consumers). PUE is not directly measured, but is configured based on the cloud provider or data center in which the software is running.

$$P_{CPUinclPUE} = P_{CPU} * PUE$$

As the total CPU utilization (CPU_{util}) is not directly related to the CPU demands of a single transaction of a system, we need to measure CPU_{util} in order to calculate the current power consumption of the CPU, and then we can allocate the power consumption proportionally to single transactions of the system. To do this, the CPU demand data needs to be converted into a utilization value (CPU_{utilT}) by summing up all CPU demand values (measured in ms) of a transaction in a given time period (where t_0 is the start timestamp of the time period in ms and t_n is the end of the end timestamp of the time period in ms) and dividing this sum by the total time the CPU was busy in that time period as follows:

$$CPU_{utilT} = \left(\sum_{t_0}^{t_n} d_{CPU} \right) \div (CPU_{util} * (t_n - t_0))$$

Using this data we can calculate the power usage of the transaction (P_{TCPU}) as follows:

$$P_{TCPU} = P_{CPUinclPUE} * CPU_{utilT}$$

Using this power consumption value, we can calculate the energy (E) in kilowatt-hours (kWh) by considering the time interval (t_0 to t_n) in hours as follows:

$$E_{CPU} = (P_{TCPU} * (t_n - t_0)) / 1000$$

To determine the carbon emissions (C_{CPU} in gCO₂eq) required to produce the energy, we need to know the grid emissions factor (GEF). The GEF denotes the carbon intensity of producing one kWh in a given region and is typically expressed in grams of CO₂ equivalents per kWh (gCO₂eq/kWh). The GEF

¹¹<https://www.etsy.com/codeascraft/cloud-jewels-estimating-kwh-in-the-cloud>

¹²https://www.spec.org/power_ssj2008/results/power_ssj2008.html

is derived by the region in which the software is operated and this region needs to be configured for the agent. For cloud providers there is a list of grid emissions for their regions provided by CCF¹⁰. Alternative sources for this value would be ElectricityMaps¹³ or Carbon Aware Computing¹⁴.

$$c_{CPU} = E_{CPU} * GEF$$

For storage, memory, and network the CCF methodology¹⁰ provides a power consumption per GB (P_{GB}) depending on the type (RAM/HDD/SSD), again following the initial publication by Etsy¹¹ for storage and network. Therefore, it is important to know how much data is consumed in a given time period, which can be calculated based on the resource demand measurements. Since the resource demand measurement values are in bytes, we need to multiply them by 10^{-9} to convert them to GB values. Therefore, the calculation of energy consumption (E) and carbon emissions (c) is the same for memory, network, and storage:

$$E_{GB} = (P_{GB} * PUE * (t_n - t_0))/1000$$

$$c_{MEM} = \sum_{t_0}^{t_n} d_{MEM} * 10^{-9} * E_{GB} * GEF$$

$$c_{NET} = \sum_{t_0}^{t_n} d_{NET} * 10^{-9} * E_{GB} * GEF$$

$$c_{STO} = \sum_{t_0}^{t_n} d_{STO} * 10^{-9} * E_{GB} * GEF$$

Based on the previous calculations we can now calculate the total carbon emissions of a transaction as shown in the following section.

5 SCI Calculation

For any given transaction we can now use the the sum of the carbon emissions (c_T) of each resource as the E*I part of the SCI specification.

$$E * I = c_T = c_{CPU} + c_{MEM} + c_{STO} + c_{NET}$$

What is still missing is the embodied carbon (M) part of the SCI calculation. Since the total embodied emissions (TEE) of a hardware server cannot be used for modern runtime environments such as virtual machines (VMs) or containers, CCF¹⁰ suggests using the relative amount of resources used by a VM or container as M (e.g., CPU count of the instance (CPU_{count_i}) divided by the CPU count of the entire server (CPU_{count_s})). In addition, the time that the server is in use must be taken into account. To do this, the coefficient 0.0289¹⁵ is used to scale down the

TEE value, typically given in kgCO₂eq, to gCO₂eq/h based on four years of usage. Using the above ratio and our time interval the value for M is calculated as follows:

$$M = TEE * 0.0289 * (t_n - t_0) * (CPU_{count_i} / CPU_{count_s})$$

As this value is not yet scaled to the individual transaction, we cannot add it directly to the SCI value of a single transaction (SCI_T). We use the previously introduced CPU_{utilT} to properly scale this value to the various transactions in a given time frame as follows:

$$M_T = M * CPU_{utilT}$$

Finally, we can calculate the SCI value for any given transaction (SCI_T) by adding c_T and M_T . The R value can be omitted because all values are already scoped to a single functional unit, namely the transaction for which the resource demand values were measured.

$$SCI_T = c_T + M_T$$

6 Conclusion and Future Work

This work presented an extension to the OpenTelemetry Java Auto-Instrumentation agent that collects the relevant data to automatically calculate an SCI score per transaction. The current implementation state of the emission calculation relies on many data sources that are based on assumptions and configurations. In addition, the data sources are static which is not realistic considering the GEF differences during day and night for renewable energy sources. Therefore, we are working on integrating other data sources that are more up-to-date with the current grid system on the one hand and the actual energy consumption of the system on the other hand. In addition, we plan to perform evaluations comparing our calculations with real measurements to see how well they represent the reality.

References

- [1] A. Brunnert and H. Krcmar. “Continuous performance evaluation and capacity planning using resource profiles for enterprise applications”. In: *Journal of Systems and Software* 123 (2017), pp. 239–262.
- [2] A. Brunnert. “Green Software Metrics”. In: *Companion of the 15th ACM/SPEC International Conference on Performance Engineering*. ICPE ’24 Companion. London, United Kingdom: Association for Computing Machinery, 2024, pp. 287–288.
- [3] A. Guldner et al. “Development and evaluation of a reference measurement model for assessing the resource and energy efficiency of software products and components—Green Software Measurement Model (GSMM)”. In: *Future Generation Computer Systems* 155 (2024), pp. 402–418.

¹³<https://www.electricitymaps.com/>

¹⁴<https://www.carbon-aware-computing.com/>

¹⁵0.0289 = 1000 (kg to g) ÷ 4 (years of server usage) ÷ 12 (months per year) ÷ 30 (days per month) ÷ 24 (hours per day)