PRINCIPAL ACADEMIC TUTOR
Monica Vitali,
Department of Electronics, Information Technology and Bioengineering,
Politecnico di Milano

ACADEMIC TUTOR
Michela Meo,
Department of Electronics and Telecommunications,
Politecnico di Torino

EXTERNAL INSTITUTION
Oliver Wyman

EXTERNAL TUTOR
Anna Maria Rosati, Oliver Wyman

TEAM MEMBERS

Giuliano Bellini,
Computer Science Engineering,
Politecnico di Torino

Luciano Camilli,
Industrial Engineering,
Politecnico di Torino

Francesco De Marco,
Energy Engineering,
Politecnico di Torino

Andrea Restelli,
Computer Science Engineering,
Politecnico di Milano

Filippo Zita,
Management Engineering,
Politecnico di Milano

AppDCO$_2$

Executive summary

Rapid digitalization is causing an increase in energy consumption by IT systems, significant progress has been made in improving the energy efficiency of hardware, data center facilities, and network, while on software, instead, advancements need to be made, in order to develop applications applying sustainability criteria.

The project is very innovative because the literature is at an embryonic stage and the presence of trade-offs between technology, performance, features, and consumption does not allow universal solutions to be identified. The objective of AppDCO2 is to support enterprises with specific applications either in their decision-making process when choosing their tech partner companies or in assessing the impact of their internal applications and finding ways to reduce it.

We developed a prototype, conducting also experiments on two selected applications, capable of monitoring the resource consumption, converting it into energy consumption, and then into an economic cost in real-time. The prototype provides: (i) a sustainability rating calculated from individual factors affecting the software efficiency, including database, frontend, backend, architecture, etc., (ii) indications on which factors / levers to reduce the emissions.

This work establishes a structured methodology for assessing energy consumption that can be replicated to other applications.

Key Words

Software | Energy Efficiency | Green IT | CO2 | Sustainability

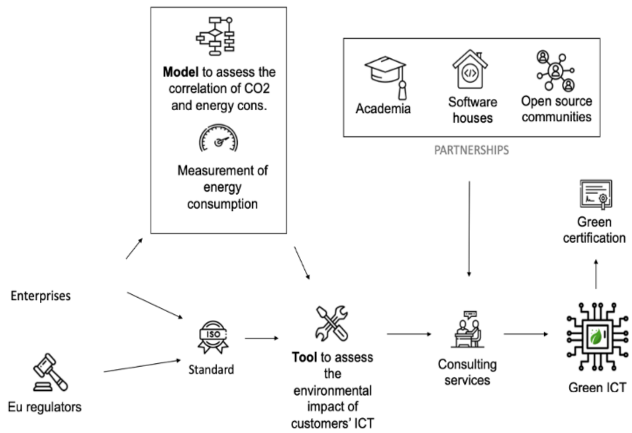

Figure 1: AppDCO2 at a glance

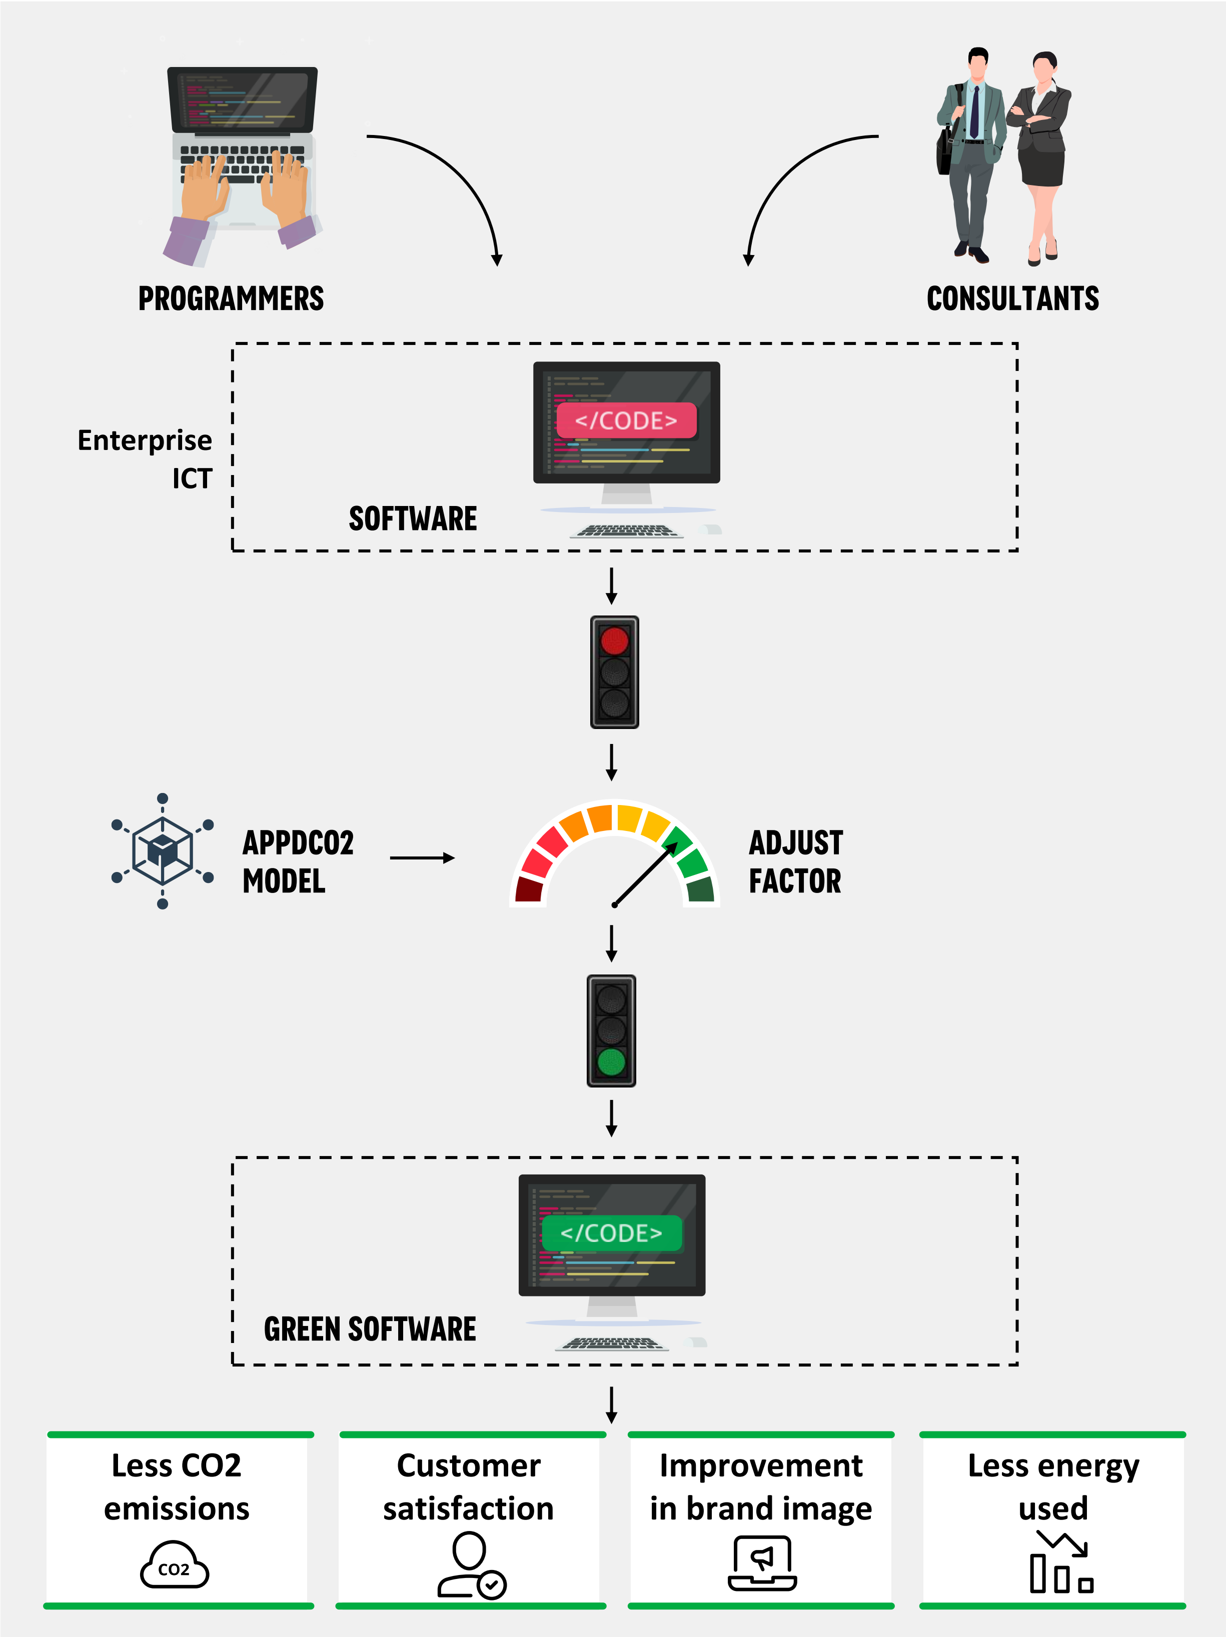
PROGRAMMERS
CONSULTANTS
Enterprise
ICT
</CODE>
SOFTWARE
APPDCO2
MODEL
ADJUST
FACTOR
</CODE>
GREEN SOFTWARE
Less CO2
emissions
CO2
Customer
satisfaction
Improvement
in brand image
Less energy
used

Project description written by the Principal Academic Tutor

Environmental sustainability is getting more relevant every day due to climate change and new governmental regulations aiming to mitigate it. National and international regulations are setting new goals and constraints related to carbon emissions that will affect all companies in the near future. The impact of IT in generating emissions is not negligible and should be taken into consideration. At the same time, IT has a relevant role in supporting the transition towards sustainability in several sectors.
In the near future, organizations will be required to be able to assess the environmental impact of their information systems and to reduce their carbon emissions. An analysis of the state of the art has highlighted the lack of reliable tools for measuring or estimating the energy of the software used to support the daily activities of organizations. Additionally, no guidelines on how to design and manage these applications are available. What an organization looking at carbon neutrality should do to reduce the impact of its software?
APPDCO2 aims at anticipating the path towards the reduction of the impact of IT on environmental sustainability supporting organizations in the transition towards a greener application portfolio. The core part of the project takes an IT perspective to investigate the issue of software sustainability aiming at proposing a method for supporting sustainable choices in software adoption and usage. A classification of different types of applications is performed and for each class, a set of relevant aspects to be considered is highlighted starting from literature and expert opinions. The proposed model assigns a sustainability relevance to each of the identified aspects. This assignment is based on an extended set of experiments aimed at understanding which parameters or design choices are more influential. Along with this, an analysis of the energy model to estimate emissions and costs from energy consumption is presented. Finally, an economical model is proposed to evaluate the advantages of adopting such an approach for organizations, considering technical, economic, and brand image factors.
This project has been developed in collaboration with Oliver Wyman, a leading consulting firm with expertise in Digital and Climate Sustainability.

Team description by skill

Due to the challenging nature of the project, the team conducted extensive research on the state of the art to identify the various optimization levers. The specificity of the project allowed non-technical team members to develop new knowledge about software applications, while computer engineering team members complemented their hard skills with management and energy-related topics. The first subproject team - Luciano, Filippo, and Francesco – worked on state-of-the-art research, developed the tool to translate energy usage into CO2 emissions, and analyzed the economic benefits of the solution. The second subproject team – Andrea and Giuliano – worked on the state of the art and conducted the various experiments in the two case studies respectively on Artificial Intelligence and Web Applications.

Giuliano, Computer Engineer: he was responsible for the examination of possible optimization levers during the research about the state of the art, more specifically concerning the subjects of business logic, application structure, and technical architecture. Moreover, he actively assisted in the software taxonomy phase, conducted extensive tests on the web application under evaluation, and contributed to defining the specifications of the models related to both case studies.

Luciano, Industrial Engineer: he contributed to the research of the state of the art related to database and data management, created the APPDCO2 Tool in Excel, and developed the economic convenience equation quantifying benefits for enterprises.

Francesco, Energy Engineer: he contributed to developing the environmental impact estimation framework, implementing on MATLAB the real-time carbon impact tool that enables users to precisely assess the CO2 emissions related to software deployment both on the cloud and locally.

Andrea, Computer Engineer: he contributed to researching the state of the art, focusing on data management and programming languages. He also played a role in defining the model and conducted experiments on both the web application in the cloud environment and the artificial intelligence application.

Filippo, Management Engineer: he contributed the state-of-the-art research mainly focusing on the AI area, he participated in the definition of the taxonomy for the 2 case studies and was co-responsible for economic benefit analysis.

Goal

This project aims to fill the green software literature gap and provide companies with concrete methods and tools to measure IT sustainability to support them in the green transition. The ambition of the team is to set innovative "green" standards, through which companies can ensure that their digital systems are designed in the most energy-efficient way, thereby reducing energy consumption and carbon emissions.

Reducing the carbon footprint of human technologies is one of the main challenges to tackle to fight global warming and climate change. The ambition of the project is to pave the way for the definition of international standards for energy-aware and sustainable digital applications and the set of methods and tools developed in the project are aimed at supporting organizations in this transition. APPDCO2 is positioned at the heart of the 17 Sustainable Development Goals identified by the United Nations member states. In particular, the main goals targeted are number 13 “Climate action”, reducing the carbon emissions of software artifacts, and number 9 “Industry, innovation and infrastructure”. The model targets private companies in the third sector that want to certify the green approach of their IT systems. Looking at the current state of the market, we can see that industries can have benefits both in terms of return of image as well as savings in taxes.

The ambition is to create a standard and a tool to assess the environmental impact of their customers’ ICT which is currently a complex and resource-intensive activity. Once the actual efficiency, the solution consists of a guideline to help make ICT greener so that the companies are guided in the transition and can adopt the best practices in terms of sustainable IT by exploiting the optimization levers outlined in the APPDCO2 Tool. In particular, the solution aims at providing savings across the entire value chain. The ultimate goal for companies could be to be certified in adherence to ESG standards. As a result, the client's credibility will be boosted in terms of brand image and brand equity, and indirectly impact when the end-users choose their products and services.

Understanding the problem

The challenge was converted into specific goals and tasks that our group addressed:

1. Analysis of the state of the art. The objective was to thoroughly investigate the literature to understand the starting point, i.e., to date what are the best practices and international standards related to green IT. In addition, since any improvement starts with monitoring the state-as-is of companies, the various types of tools available to date for calculating the energy consumption of software were also researched and investigated in depth.
2. Definition of the taxonomy, characteristics, and coding guidelines of software. The goal was to create a framework through which to analyze software and identify various levers for improvement. A breakdown of software into its various fundamental components and levels is necessary to propose the most effective and comprehensive solution possible.
3. Case studies. Identification of two applications and analysis of their architecture using the previously identified framework. Empirically test the effect of software-level changes on the energy consumption of the two applications to verify or extend the state of the art.

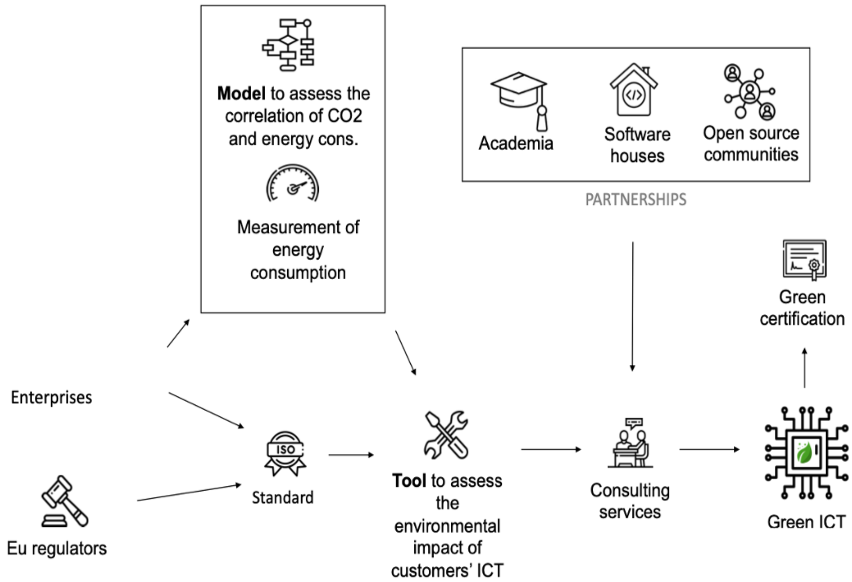

Figure 2: AppDCO2 at a glance

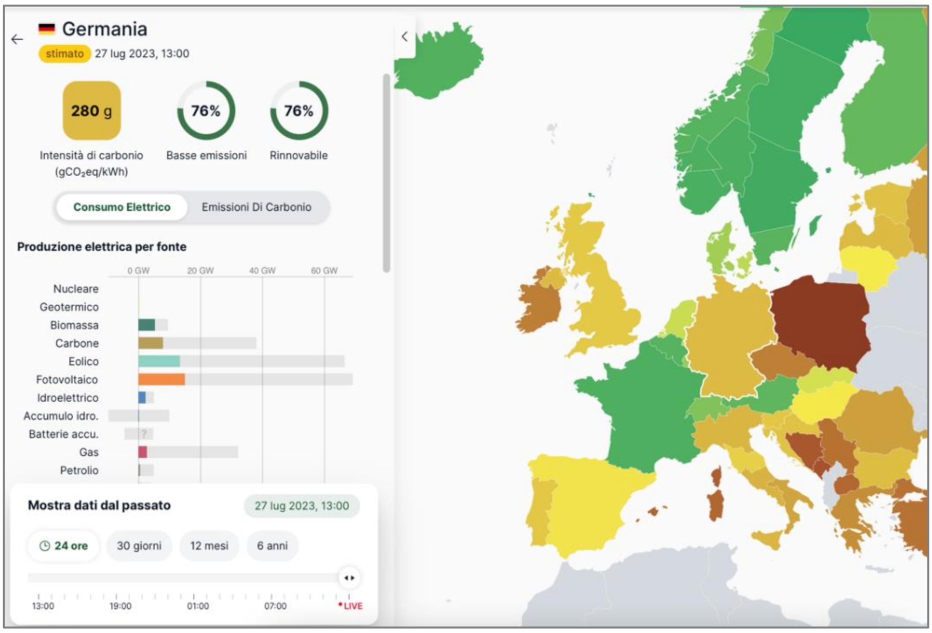

Figure 3: Electricity map in Europe

Exploring the opportunities

When assessing a company's information systems sustainability, we found out that significant research was related to green hardware, while green software literature is still in the early stages. Creating green IT best practices is an arduous challenge because often involves considering trade-offs between different objectives. Striking a balance between performance and sustainability goals requires careful decision-making, considering the specific context and priorities of the organization.

To define the various optimization levers, we carried out preliminary research work related to the state of the art. The findings gave us a complete view of the web applications scenario, which was needed to adequately define possible levers and examine the software considering all the elements it is composed of. A relevant point to underline is that the overall power consumption is not a linear function of CPU utilization. This means that given a defined CPU usage and fixed a specific configuration (number of cores, scheduling policy, cache size, and so on), distinct tasks may differ significantly in energy consumption. Simple models based purely on CPU activity are not suitable for predicting the aggregated impact of a system when it performs database operations, invokes HTTP requests, or loads images to render. Therefore, the challenging nature of the work involved finding additional metrics and diving into more detailed implementation mechanisms.

Continuing from the previous case study, we conducted comprehensive research to comprehend and assess the nature of the Artificial Intelligence scenario. The significance of the carbon footprint within the artificial intelligence case study becomes even more prominent. Considering the statistics about the enormous amount of carbon produced by training an artificial intelligence model using very computationally powerful hardware such as GPUs and TPUs, tuning the artificial intelligence pipeline of execution to keep track of the carbon impact it takes becomes crucial. In this context, our model should address additional factors from the ones we considered for the Web Applications.

Given the vastity of potential applications and techniques within the realm of artificial intelligence, we opted to narrow down the problem's scope to the most prominent field: deep learning. Specifically, we chose to concentrate our investigation on neural network models, while maintaining the overarching versatility of our model. We also decided to focus on supervised learning tasks, which involve training a model using labelled data pairs, consisting of input examples and their corresponding desired outputs. This approach provides a clear framework for the model to learn the relationships between inputs and outputs, enabling it to make accurate predictions or classifications on new, unseen data.

Generating a solution

The APPDCO2 solution measures the energy consumption and sustainability level of applications and allows comparisons between different configurations of the same application. The study also presents each of the factors influencing software efficiency (database, architecture, front-end, etc.) with the associated subfactors. For each subfactor, an optimization lever is identified so that contingency measures can be taken. The solution prototype consists of two tools.

The first, a Matlab-based tool, can monitor CPU usage, energy consumption, and economic costs in real-time. It addresses multiple factors to provide a comprehensive and reliable estimate of the CO2 emissions associated with IT operations. The first step underlying the tool is the monitoring of CPU usage and convert it into an energy consumption profile. The second step involves using the Electricity Map API. Electricity Map is a publicly available platform that provides comprehensive information on the carbon intensity of electricity generation for multiple regions around the world. By integrating with the Electricity Map API, the solution gains access to this vast repository of carbon intensity data. When combined with the energy consumption profile obtained from the monitoring tool, it can accurately estimate the amount of CO2 emissions associated with the energy consumed by the IT infrastructure. An additional capability of the tool is to estimate the cost of energy associated with the energy profile in a particular region.

The second tool is a model developed in Excel and consists of a framework for assessing the energy consumption of software applications, and provides the sustainability rating of an IT infrastructure. Based on the critical factors that influence the carbon footprint of software systems, it provides a reproducible methodology to assess the weight of these factors in real-life scenarios. As the range of applications is very broad, it is almost impossible to provide a valid general model for every type of application. For this reason, two categories of applications have been used as a starting point, namely web applications and artificial intelligence. The factors that may be common to similar types of applications were identified and a reproducible methodology was defined to measure their energy consumption and assess their impact within the application. This tool is composed of both a model with categories, subfactors, and weights to assess the carbon footprint of an application and a set of methodologies to measure those factors. The final model is not immutable and based on perfect weights, but it helps organizations assess the energy consumption and the potential for improvement of their software solution.

Main bibliographic references

Anthony, L.F.W., Kanding, B. and Selvan, R. (2020) Carbontracker: Tracking and predicting the carbon footprint of Training Deep Learning Models, arXiv.org. Available at: https://arxiv.org/abs/2007.03051

Lawrence, A. (2020) Data Center pues flat since 2013, Uptime Institute Blog. Available at: https://journal.uptimeinstitute.com/data-center-pues-flat-since-2013

Strubell, E., Ganesh, A. and McCallum, A. (2019) Energy and policy considerations for Deep Learning in NLP, arXiv.org. Available at: https://arxiv.org/abs/1906.02243