Carbon-Aware Process Execution for Green Business Process Management

Philipp Hehnle[1], Maximilian Behrendt[1], Luc Weinbrecht[1] and Carl Corea[2]
[1]*envite consulting GmbH, Stuttgart, Germany*
[2]*Research Group Process Science, University of Koblenz, Koblenz, Germany*

Keywords: Green BPM, Carbon-Aware Process Execution, Green ICT, Camunda.

Abstract: Traditional business process management (BPM) focuses on the improvement of performance dimensions such as time, costs, and quality. Ecological aspects are usually not considered as an equal performance dimension. In this context, Green BPM approaches have been proposed to strengthen the awareness among people and organisations about the impact of business processes on the climate. However, little research in Green BPM covers the runtime of digitised processes, or provides concrete means to reduce carbon emissions during process execution. Therefore, we present an approach for carbon-aware process execution, which allows to automatically postpone energy-intensive activities to times when energy with low CO_2 emissions, e.g. solar energy, is better available. Importantly, our approach considers and complies with external regulations such as Service Level Agreements (SLAs) when postponing activities. Our approach is implemented in Camunda and has been evaluated in interviews with domain experts.

1 INTRODUCTION

Climate change has the potential to cause unprecedented natural disasters that impact, inter alia, food production and thereby threatens livelihoods (Matemilola et al., 2020). In result, governments across the world adopted the Paris Agreement in 2015 (Matemilola et al., 2020), committing to implement measures in order to keep the global temperature increase below 2°C compared to pre-industrial levels. To limit the global temperature increase, it becomes necessary to drastically reduce carbon (CO_2) emissions, so that there are net zero emissions in the long run (Matemilola et al., 2020). To this end, there is a broad consensus that Information and Communication Technology (ICT) can play an important role in reducing CO_2 emissions (Kim et al., 2009; Oloo Ajwang and Nambiro, 2022; Gohar et al., 2020), e.g. web meetings may reduce the need for transportation. However, it is not safe to assume that the use of ICTs will always save emissions. In particular, the sheer execution of the ICTs produces emission itself. According to (Freitag et al., 2021), ICTs produce 2.1%-3.9% of global green house gas emissions. Therefore, the ICT sector itself must also strive to reduce its emissions to contribute towards the overall goal of the Paris Agreement. In this work, we investigate how digitised business processes can contribute towards this goal. In particular, we present an approach for *carbon-aware process execution*, which allows to automatically execute activities of a business process at times at which energy sources with low CO_2 emissions (*green energy*) are available, e.g. solar energy, lowering the CO_2 emissions. Importantly, our approach leverages predictive insights to ensure that the business process execution remains compliant with external requirements such as service level agreements (SLA).

In traditional BPM, the main performance dimensions for which processes are optimised for are *time, cost*, and *quality* (van der Aalst, 2013). In this work, we aim to raise awareness that the dimension of *ecology* should be seen as an equal. To this aim, our approach supports companies in considering carbon emissions as a process performance dimension.

Our approach is implemented in Camunda, and can be used “out-of-the-box” by companies. We show that our approach can be used to reduce CO_2 emissions with a comprehensive case-study and evaluate our approach in interviews with domain experts. For our investigation, we follow a design science oriented approach as proposed in (Peffers et al., 2007), in that we identify important requirements (Section 3.1) and develop/evaluate our approach based on these.

Hehnle, P., Behrendt, M., Weinbrecht, L. and Corea, C.
Carbon-Aware Process Execution for Green Business Process Management.
DOI: 10.5220/0012557100003690

In *Proceedings of the 26th International Conference on Enterprise Information Systems (ICEIS 2024) - Volume 2*, pages 659-666
ISBN: 978-989-758-692-7; ISSN: 2184-4992

In Section 2, fundamentals on Sustainability, Green BPM and Predictive Process Monitoring are stated. In Section 3, we present our approach to reduce CO_2 emission during process execution and a proof of concept implemented for the Workflow Management System (WfMS) Camunda. In Section 4, the presented approach is evaluated. Finally, this paper concludes with a summary and outlook.

2 FUNDAMENTALS

In this section, we introduce concepts to save carbon when using ICTs, SLAs, predictive process monitoring, and Green BPM.

2.1 Sustainability in Information and Communication Technology

There are various approaches to operate ecologically sustainable ICTs, mainly, by trying to improve the efficiency of the software or the hardware. As an example, improving the hardware software is operated on can reduce the energy consumption and thereby the ecological impact (Freitag et al., 2021). Furthermore, the choice of technology (e.g. programming language (Pereira et al., 2017)) and the programming style (e.g. used data structures (Hasan et al., 2016)) affects the energy consumption of software.

Besides these approaches, the *time* when software is executed may be chosen carbon-aware. For instance, (Radovanovic et al., 2023) discusses how data centre computing may be run in time windows when green energy is available. For this purpose, the Green Software Foundation's Carbon Aware SDK[1] may be used to get a forecast about when and where additional energy consumption causes the least CO_2 emissions. The SDK takes into account whether there are enough sustainable energy sources available for the additional energy consumption. Building on this forecast, it therefore becomes possible to execute energy-intensive activities when green energy is available. Consequently, in this work, we focus on this latter line of approaches, i.e. trying to reduce emissions by optimising the execution *time*.

2.2 Service Level Agreements in Business Process Management

An SLA (Frankova et al., 2011) constitutes a contract between a service provider and a consumer describing a non-functional requirement. SLAs of a business process may be the maximal execution time of the entire (Frankova et al., 2011) or parts of the business process (del Río-Ortega et al., 2015). By monitoring the SLAs of a business process, it becomes possible to prevent a violation of an SLA, e.g. by alerting experts or by automatically removing/adding activities of the business process (del Río-Ortega et al., 2015).

In the context of this work, we implement an SLA-based demand shifting to ensure that any alterations in process execution (in favour of less CO_2 emissions) still remain compliant to the SLAs.

2.3 Predictive Process Monitoring

Predictive Process Monitoring (PPM) (Di Francescomarino and Ghidini, 2022) is a discipline aimed at predicting future aspects of running business processes. Here, history data from process executions (event logs) are considered to train machine learning models. These models can then take as input a running instance and predict different aspects such as the next activity, the remaining time, or possible outcomes (Di Francescomarino and Ghidini, 2022).

To train the models, traces from the event log, i.e. a sequence of executed activities, are usually considered to recognise emerging patterns and behaviours. To this end, process variables can be taken into account, such as the assigned employee.

In the approach presented in this work, remaining time predictions (i.e. the time until the end of the *entire* process instance) are leveraged to estimate the potential degree of freedom for postponing current activities (e.g. the postponement of an activity plus the estimated remaining time should not exceed SLAs). As our approach is applied to WfMSs, event logs tracked by a WfMS can be used for model training. For this work, we build on our previous results for training machine learning models in Camunda (Bartmann et al., 2021). The tool presented in that work can be used to train models in a user-friendly way and provides an interface to predictions in Camunda.

2.4 Green Business Process Management

In broad terms, Green BPM is a specialisation of traditional BPM, aimed to lessen the impact of processes on the natural environment (Seidel et al., 2012). Due to an increasing societal and organisational attention on topics such as climate change (cf. e.g. (Couckuyt and van Looy, 2021)), the topicality of "Green BPM" has again gained some recent momentum, with recent works broadly agreeing on potential benefits and the

---

[1] https://github.com/Green-Software-Foundation/carbon-aware-sdk

need for future research in this area (Couckuyt, 2018; Couckuyt and van Looy, 2021). Also, there seems to be a growing awareness that organisations do in fact contribute to environmental degradation through their processes and should therefore aim to implement more environmentally sustainable processes (Seidel et al., 2012; Couckuyt and van Looy, 2021). Following (Opitz et al., 2014), Green BPM can be described to have two main goals, namely 1) the reduction of the environmental impact of business processes, and 2) inducing cultural change w.r.t. more eco-friendly behaviour. This work focuses on the first goal, namely means for reducing the environmental impact, concretely, reducing the carbon footprint.

Regarding the existing works on reducing the environmental impact of business processes, most research can be found in the area of *design-time analysis* (cf. (Maciel, 2017; Couckuyt and van Looy, 2020; Roohy Gohar and Indulska, 2020; Hoesch-Klohe and Ghose, 2010; Ghose et al., 2010)). That is, most works focus on facilitating means for *modelling* processes in a more eco-aware way, for example, by means of extensions of modelling notations to indicate fuel consuming activities (Recker et al., 2011), or design patterns for modelling processes in a more ecologically sustainable way (Alexander Nowak et al., 2014). Works such as (Ghose et al., 2010) and (Hoesch-Klohe et al., 2010) annotate activities in process models with carbon emissions at design time to substitute fragments of the process models with ones that preserve the behaviour but reduce the carbon impact. However, the authors (Ghose et al., 2010) acknowledge the fact that, at runtime, carbon emissions may vary depending on the execution time. Despite the clear benefits of such approaches, we do see a current lack of Green BPM research in the area of *runtime* monitoring, i.e. support for Green BPM once processes are actually running. And in fact, this research gap has also been identified in recent research agendas, i.e. (Maciel, 2017; Couckuyt, 2018) state that more attention should be paid to “[green] process implementation and execution”, and the “development of instruments”. Further, the authors in (Ghose et al., 2010) argue that “process execution management can also contribute to carbon footprint minimization”. To close this existing research gap, in this work, we develop new means for the operational support of reducing the environmental impact of running processes, concretely, by developing eco-aware monitoring techniques in combination with results from prescriptive process monitoring.

The Activity-Based Emissions (ABE) approach presented in (Recker et al., 2011) is developed for analysing the carbon footprint of processes and allows to conceptualise the emissions of individual activities. In this work, we follow the ABE approach and focus on activities with high emissions due to high energy consumption. Our work lifts the ABE approach to a runtime perspective. In particular, our approach allows to reduce the environmental impact of carbon-aware process models at runtime, by orchestrating the execution of individual activities s.t. these activities are executed in time windows where green energy is better available. This is mainly done by postponing activities, as has also been proposed by (Zhu et al., 2015). While other approaches comprise manufacturing business processes (Ghose et al., 2010; Hoesch-Klohe et al., 2010), we focus on digitised business processes only.

Importantly, by leveraging predictive insights, activities are postponed while still ensuring that processes remain compliant, e.g. are in line with SLAs. In such, our approach represents a novel form of “Carbon-Aware” Prescriptive Process Monitoring.

Note that CO_2 emissions are not the only KPI measurable in regard to the environmental impact (cf. e.g. the dimensions mentioned in (Roohy Gohar and Indulska, 2020) like water consumption, waste, etc.). However, reducing CO_2 can be seen as an important part of lowering the environmental impact of processes (Roohy Gohar and Indulska, 2020), which is why we focus on this aspect in this work.

3 CARBON-AWARE PROCESS EXECUTION

In this section, we will present our approach to reduce the environmental impact of business processes. We begin by exploring important requirements for such an approach and then continue with the concrete presentation and discussion.

3.1 Requirements

The main goal of our approach is to reduce the CO_2 emissions of digitised business processes. We propose to shift the execution of energy-intensive activities to time windows when green energy is available. For such a run-time approach, we see the following general requirements:

Typically, a WfMS is used to orchestrate various services. The required energy for these services varies widely. Getting an entry in a database might not use much energy compared to complex tasks such as image processing. Therefore, an approach is required that is simple to use and takes a holistic view of the business process instead of single services separately.

Furthermore, the process instances need to continue to comply with the specified SLAs. In result, when moving the execution of an activity to a later time window, the overall process execution has to remain within these bounds.

Based on these observation, we therefore raise the following requirements.

- *R1: The adjustment effort must be low.* In principle, it would be possible to configure every service orchestrated by a process individually to consider the carbon intensity of the current energy supply. However, that would entail a lot of adjustment efforts in many services. Ideally, only small changes to the original process model are required.
- *R2: The business process must be optimised holistically.* Rather than optimising services individually, the optimisation should aim for the business process as a whole.
- *R3: The process instances need to comply with specified SLAs.* When business processes are executed, the specified SLAs must be adhered to. For example, if the maximal allowed execution time for a process instance is 3 hours, the approach needs to ensure that even if execution is postponed in favour of carbon savings, the process instance is executed within this frame.

3.2 Proposed Approach for Carbon-Aware Process Execution

To reduce CO_2 emissions produced by energy-intensive activities, we propose to shift the execution of these activities to time windows where green energy is available. For instance, before executing an energy-intensive activity (e.g. image processing), it is verified if green energy is available entailing immediate execution or otherwise it may be postponed.

Two main problems in this use-case are 1) getting data on the availability of green energy, and 2) deciding if and how long to postpone activities. For problem 1), we build on the publicly available Carbon Aware SDK by the Green Software Foundation (cf. Section 2.1). Regarding problem 2), a method is needed for deciding how to postpone activities, e.g. SLAs must be kept. In the following, we present our proposed method for this decision task. See process model in Figure 1 as a running example.

Following the ABE approach (Recker et al., 2011), activities causing high emissions have to be identified by domain experts, e.g. based on metrics such as energy consumption. In Figure 1, this is schematically shown as the activity "energy intense work". In case of such an activity, a new model element, denoted a *postponement activity*, can be added before this element (see "postpone execution for greener energy" in Figure 1). The postponement activity decides whether to postpone subsequent activities to time windows when energy is less carbon-intensive. To allow for making this decision, the postponement activity is called with the following data:

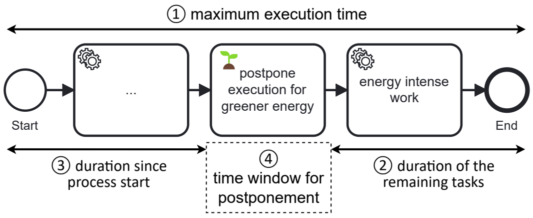

Figure 1: Exemplary process model containing energy-intensive activity adapted from (Hehnle et al., 2023).

- the region in which the process is executed
- the start time of the process instance
- the maximal runtime until the end of the process instance (SLA)
- the (predicted) duration of the remaining activities (cf. Section 2.3)

The region is required to determine the carbon-intensity in the energy forecast. It is assumed that the entire process and its called services are executed in the same region. The time information (start time, maximal execution time according to SLA, predicted remaining time) are needed to calculate the possible time window for postponing the carbon-intensive activity. To prevent SLA violations, the execution may not be postponed arbitrarily. Equation (1) describes how the time window for postponement is calculated:

$$t_P = t_{MET} - (t_{PS} + t_{RA}) \tag{1}$$

- t_P: time window for postponement, i.e. the maximal time subsequent activities may be postponed.
- t_{MET}: maximal execution time of the business process (i.e. the SLA)
- t_{PS}: duration since process start
- t_{RA}: predicted duration of the remaining activities

The calculation of the time window for postponement is visualised in Figure 1. The maximal execution time of a process instance, which must not be exceeded, can be seen at the arrow with the mark ① (e.g. according to an SLA). The (predicted) duration of the remaining activities following the postponement-task (mark ②) as well as the time since the process instance start (mark ③) need to be deducted from this maximal execution time, which results in the time window for the postponement (mark ④). In result, the execution of activities can be postponed within this time window while still considering the specified maximal execution time. This ensures that carbon optimisation does not result in SLA violations.

Finally, the API by the Green Software foundation is used to retrieve a forecast of the carbon-intensity within the time window for the postponement. Based on this forecast, the WfMS postpones the execution of the energy-intensive activity until the optimal time with regard to green energy.

3.3 Proof of Concept

To demonstrate our approach, we implemented a proof of concept (PoC) using Camunda together with the Green Software Foundation's Carbon Aware SDK. In the following, we refer to our PoC as the *Camunda Carbon Reductor*. The Camunda Carbon Reductor is publicly available[2] and can be used "out-of-the-box" with any Camunda (7 or 8) installation.

In the following, the functionality of the Camunda Carbon Reductor is described in interaction with Camunda 8. *Connectors* are a special type of task in Camunda. They act as reusable templates consisting of a model element and a small application that interacts with third party systems. For our PoC, the Camunda Carbon Reductor is implemented via a connector. It can be added to a process model before energy-intensive activities. Configured accordingly, it can automatically decide to postpone the energy-intensive activity to a better time window, while still ensuring the SLAs. Figure 2 highlights the configuration possibilities of the Camunda Carbon Reductor:

1. The location where the energy-intensive activity is executed in, which is required for the forecast.
2. An estimated duration of the remaining activities, which may be specified manually or retrieved via a remaining time prediction. The concrete methods for obtaining such a prediction are beyond the scope of this report. We build on our previous tool presented in (Bartmann et al., 2021), which allows to retrieve such predictions directly in Camunda.
3. The milestone is a timestamp and the starting point for the SLA based duration calculation.
4. The maximal duration of a process instance from the milestone to finish to be SLA compliant.

Once the Camunda Carbon Reductor determined the optimal time window for execution, a timer is set to pause the execution until green energy is available.

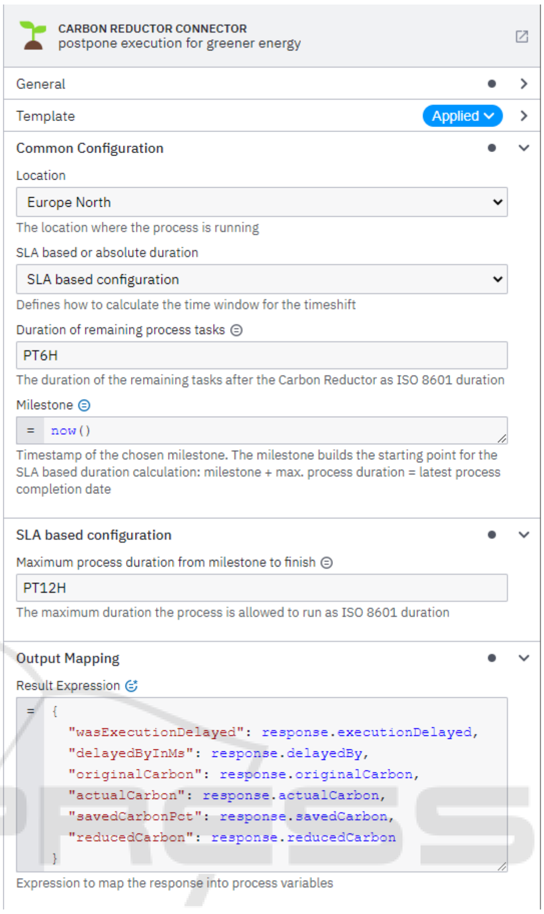
CARBON REDUCTOR CONNECTOR
postpone execution for greener energy

Common Configuration

Location: Europe North
The location where the process is running

SLA based or absolute duration: SLA based configuration
Defines how to calculate the time window for the timeshift

Duration of remaining process tasks: PT6H
The duration of the remaining tasks after the Carbon Reductor as ISO 8601 duration

Milestone: = now()
Timestamp of the chosen milestone. The milestone builds the starting point for the SLA based duration calculation: milestone + max. process duration = latest process completion date

SLA based configuration

Maximum process duration from milestone to finish: PT12H
The maximum duration the process is allowed to run as ISO 8601 duration

Output Mapping

Result Expression:

```
= {
    "wasExecutionDelayed": response.executionDelayed,
    "delayedByInMs": response.delayedBy,
    "originalCarbon": response.originalCarbon,
    "actualCarbon": response.actualCarbon,
    "savedCarbonPct": response.savedCarbon,
    "reducedCarbon": response.reducedCarbon
  }
```

Expression to map the response into process variables

Figure 2: Showcase of configuration possibilities of the Camunda Carbon Reductor adapted from (Hehnle et al., 2023).

For analysis in the monitoring tool Camunda Optimize, results are written as variables into the process instance (also shown in Figure 2)

3.4 Discussion

In the following, we discuss the fulfilment of the requirements (cf. Section 3.1) as well as limitations.

3.4.1 Requirements Fulfilment

The benefit of the proposed approach is that only the process model has to be adjusted in order to save carbon emissions, and no further adjustments are necessary for the called services. Therefore, *R1* is met. The Camunda Carbon Reductor orchestrates the called services and is thus able to achieve the maximum carbon reduction for the business process through carbon-aware optimisation. Here, *R2* is satisfied as well. The developer can specify the SLA of the maximal execution time of the business process, which is considered when postponing energy-

---

[2] https://github.com/envite-consulting/camunda-carbon-reductor

intensive activities. Based on the specified SLA, the predicted remaining time, and the duration since a specified milestone, the time window for postponement is calculated. In result, *R3* is met as well.

3.4.2 Limitations

If a business process is run in a region where the demand of energy is higher than green energy is available postponing activities will not reduce emissions. Furthermore, only activities whose consumption is higher than that of the Camunda Carbon Reductor should be postponed or otherwise more carbon emissions are produced than saved. However, measuring the energy consumption of software is difficult and still a subject of research (Ardito et al., 2019). Integrating the Camunda Carbon Reductor multiple times in one process model may result in SLA violations as all Camunda Carbon Reductors might exploit the maximal execution time. When subsequent activities after the Camunda Carbon Reductor include user tasks it becomes difficult to predict the remaining time of the process. When the Camunda Carbon Reductor postpones the execution of a business process and in addition to that later on in the business process it takes more time than anticipated for a human to complete a user task it may result in an SLA violation as well. Finally, the current approach favours only the process performance dimension of emitted carbon. However, it might be necessary to consider a trade-off among time, cost, quality, and carbon emissions.

4 EVALUATION

The approach as well as the Camunda Carbon Reductor are evaluated in the following. First, a case study explores the fitness in production use. Then, we present feedback from experts gained in interviews.

4.1 Case Study

We conducted a case-study to investigate the potential benefits of our approach. For this, we created an exemplary process model and simulated its execution within a Camunda installation. The process model (shown in Figure 3) consists of a start event, an energy-intensive activity and an end event. The Camunda Carbon Reductor was included as a predecessor activity of the energy-intensive activity. We then used Camunda to run concrete process instances of this process over the course of one week (every 30 minutes, a process instance was started). Importantly, the execution was performed in *real-time*, allowing to also obtain real-time energy forecasts (meaning the data obtained from the case study reflects the actual CO_2 savings!). In total, 173 process instances were executed in this time. For the energy-intensive activity, an execution time of one hour was assumed. Furthermore, an exemplary SLA of 23 hours was assumed. In result, there is a time window of 22 hours to postpone the energy-intensive activity. The activity was assumed to be run in a data centre in the region West US. The results of the case study are depicted in Figure 3.

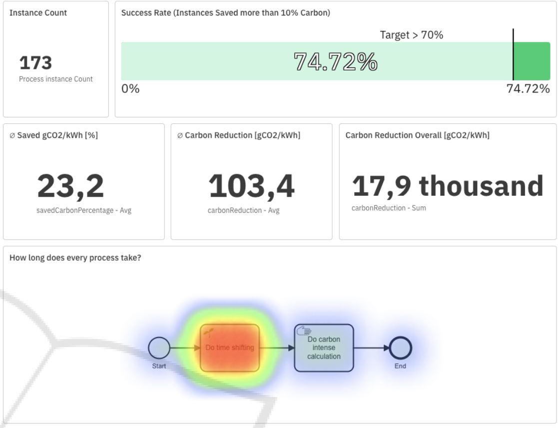

Figure 3: Case study results (Camunda Optimize Dashboard) adapted from (Hehnle et al., 2023).

The data collected by the Camunda Carbon Recutor allow to report KPIs such as the amount of emissions reduced per kWh. As can be seen, 17.9 thousand gCO_2/kWh could be saved over all instances. On average, 103.4 gCO_2/kWh were saved.

Furthermore, the dashboard allows to track the progress of goals, e.g. an instance should ideally save at least 10% of carbon. 74.72% of all (173) instances satisfied this goal. Among all instances on average 23% CO_2 emissions could be saved.

4.2 Expert Interviews

To evaluate the plausibility of our approach, we conducted interviews with eight domain experts who were approached from a pool of industrial partners, but not associated with the affiliations of the authors. An overview of the experts is shown in Table 1.

As can be seen, all experts had a strong background in BPM and Camunda. Especially, inputs from experts A, B, D, E and H, who explicitly work on implementing Camunda in industrial settings, make us confident the interview partners were very suitable to provide feedback. With each expert, we conducted an interview via Microsoft Teams. At least two researchers were present in every interview and an pro-

Table 1: Expert panel overview.

ID	Job Title	Experience with BPM/ Camunda	Company/Sector
A	Developer Advocate	9 years	Software Vendor
B	Developer Advocate	1 years	Software Vendor
C	Senior Consultant	12 years	IT Consulting
D	BPM Consultant	6 years	IT Technology
E	IT Consultant	5 years	IT Consulting
F	Principal Scientist	11 years	Software Vendor
G	Data Scientist	3 years	Tech Company
H	Software Architect Process Development	3 years	Insurance Company

tocol was followed to support internal validity: First, the general approach was motivated. Then, the PoC was shown to the experts, followed by a roughly 20-30 minute discussion. After the interview, the experts were shown four statements (e.g. "I found the tool useful"), to which they should indicate whether they agreed or not (on a 5-point Likert scale from *strongly agree* to *strongly disagree*). The statements were based on the technology acceptance model. Figure 4 shows our survey results.

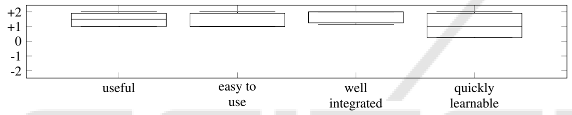

Figure 4: Results of the conducted user survey (Statements: "I found the tool..."), on a 5-point Likert scale from *strongly agree* (+2) to *strongly disagree* (-2).

As can be seen from Figure 4, our approach was uniformly seen as useful and easy to use. It was especially mentioned that the approach was very well integrated, which is in line with requirement R1. The experts did see a slight need for training when using such a tool, but this was seen as expectable.

Also from the open questions/interview, the overall approach was uniformly well-received. The experts were specifically asked whether they could imagine to use the PoC in one of their industrial project, to which five experts answered *yes*. Expert C emphasised that the PoC is especially useful and necessary in sectors that run huge simulations, e.g insurance companies run simulations before many contracts. Expert D indicated that the approach should also be able to transfer the execution of an activity to a different location where green energy is available. Experts F and G proposed a configuration to specify a minimum saving of CO_2 which need to be achieved to postpone the execution or otherwise the execution will be continued immediately. Expert B argued that it might become necessary for a human process participant to manually resume a paused process instance.

5 SUMMARY AND OUTLOOK

Climate change is a challenge for society as a whole. The ICT Sector also needs to contribute to overcome it by reducing carbon emissions (Gohar et al., 2020; Seidel et al., 2012). Research has unveiled various strategies to reduce carbon emissions of ICTs, one of which is them being executed carbon-aware , i.e. time windows can be chosen for running the ICT solutions in which green energy is available.

In this work, we applied the well-established approach of carbon-aware execution from the field of sustainable ICTs to business processes, i.e. energy-intensive activities are executed in time windows in which green energy is available. While traditional BPM focuses on the performance dimensions time, cost, and quality, Green BPM has evolved to reduce the environmental impact of business processes. Previous approaches of Green BPM push for design time optimisation, whereas we focus on the runtime.

The Camunda Carbon Reductor is a proof of concept that implements the presented approach and is publicly available. A case study revealed that the use of the Camunda Carbon Reductor can achieve substantial CO_2 reduction. In this context, the conducted expert interviews underline these findings, where the majority of experts confirmed the need, usefulness and plausibility of our approach.

Despite the achievements in the case study and the good response in the expert survey, the Camunda Carbon Reductor may be improved. Currently, an elements needs to be integrated into the process model, whereas in the future, no element shall be needed. Also, in future work, the Camunda Carbon Reductor shall be able to relocate the execution of business processes to data centres where green energy is available. Currently, it is reasonable to pause the execution of a business process only at one position. However, multiple energy-intensive activities in a process model shall be considered. Also, we plan to integrate results from multi-instance predictive process monitoring to verify whether activities from parallel instances can be postponed in a batch, e.g. image processing tasks from different process instances can be batched, s.t. the hardware used for the image processing does not have multiple (energy-intensive) "cold-starts".

As stated, CO_2 emissions are not the only KPI in regard to the environmental impact (e.g. water consumption, waste (Roohy Gohar and Indulska, 2020)). However, reducing the carbon footprint is uniformly viewed as important in the context of Green BPM (Gohar et al., 2020), which makes us confident that the presented approach is a step towards improving the environmental impact of business processes.

REFERENCES

Alexander Nowak, Uwe Breitenbücher, and Frank Leymann (2014). Automating green patterns to compensate co2 emissions of cloud-based business processes. In *Proc. 8th Int. Conf. ADVCOMP.*, pages 132–139. Xpert Publishing Services.

Ardito, L., Coppola, R., Morisio, M., and Torchiano, M. (2019). Methodological guidelines for measuring energy consumption of software applications. *Scientific Programming*, 2019:1–16.

Bartmann, N., Hill, S., Corea, C., Drodt, C., and Delfmann, P. (2021). Applied predictive process monitoring and hyper parameter optimization in camunda. In *Proc. CAiSE Forum 2021*, pages 129–136. Springer.

Couckuyt, D. (2018). An overview of challenges and research avenues for green business process management. In *On the move to meaningful internet systems*, volume 10697 of *LNCS Inf. sys and applications, incl. internet/web, and HCI*, pages 270–279. Springer.

Couckuyt, D. and van Looy, A. (2020). A systematic review of green business process management. *Business Process Management Journal*, 26(2):421–446.

Couckuyt, D. and van Looy, A. (2021). An empirical study on green bpm adoption: Contextual factors and performance. *J. of Software: Evolution and Process*, 33(3).

del Río-Ortega, A., Gutiérrez, A. M., Durán, A., Resinas, M., and Ruiz-Cortés, A. (2015). Modelling service level agreements for business process outsourcing services. In *Advanced inf. sys. eng.*, volume 9097 of *LNCS*, pages 485–500. Springer.

Di Francescomarino, C. and Ghidini, C. (2022). Predictive process monitoring. In *Process Mining Handbook*, volume 448 of *LNBIP*, pages 320–346. Springer International Publishing.

Frankova, G., Séguran, M., Gilcher, F., Trabelsi, S., Dörflinger, J., and Aiello, M. (2011). Deriving business processes with service level agreements from early requirements. *J. Syst. and Softw.*, 84(8):1351–1363.

Freitag, C., Berners-Lee, M., Widdicks, K., Knowles, B., Blair, G. S., and Friday, A. (2021). The real climate and transformative impact of ict: A critique of estimates, trends, and regulations. *Patterns*, 2(9):100–140.

Ghose, A., Hoesch-Klohe, K., Hinsche, L., and Le, L.-S. (2010). Green business process management: A research agenda. *Australasian J Inf. Sys.*, 16(2).

Gohar, S. R., Indulska, M., et al. (2020). Environmental sustainability through green business process management. *Australasian J. Inf. Sys.*, 24.

Hasan, S., King, Z., Hafiz, M., Sayagh, M., Adams, B., and Hindle, A. (2016). Energy profiles of java collections classes. In *Proc. 38th Int. Conf. Softw. Eng.*, pages 225–236. ACM.

Hehnle, P., Weinbrecht, L., and Behrendt, M. (2023). The camunda 8 connector for carbon-aware process execution. https://camunda.com/blog/2023/07/carbon-aware-process-execution-connector/ (retrieved 16/02/2024).

Hoesch-Klohe, K. and Ghose, A. (2010). Carbon-aware business process design in abnoba. In *Proc. 8th Int. Conf. Service Oriented Comp.*, pages 551–556. Springer.

Hoesch-Klohe, K., Ghose, A., and Lê, L.-S. (2010). Towards green business process management. In *Int. Conf. Services Comp.*, pages 386–393. IEEE.

Kim, S., Kim, H.-K., and Kim, H. J. (2009). Climate change and icts. In *INTELEC 09*, pages 1–4. IEEE.

Maciel, J. C. (2017). The core capabilities of green business process management—a literature review. In *Proc. Int. Conf. Wirtschatsinformatik*, pages 12–15.

Matemilola, S., Fadeyi, O., and Sijuade, T. (2020). Paris agreement. In *Encyclopedia of Sustainable Management*, pages 1–5. Springer International Publishing.

Oloo Ajwang, S. and Nambiro, A. (2022). Climate change adaptation and mitigation using information and communication technology. *Int. J. Comp. Sci. Research*, 6:1046–1063.

Opitz, N., Krup, H., and Kolbe, L. M. (2014). Green business process management – a definition and research framework. In *47th Hawaii Int. Conf. Sys. Sci.*, pages 3808–3817. IEEE.

Peffers, K., Tuunanen, T., Rothenberger, M. A., and Chatterjee, S. (2007). A design science research methodology for information systems research. *J. management inf. sys.*, 24(3):45–77.

Pereira, R., Couto, M., Ribeiro, F., Rua, R., Cunha, J., Fernandes, J. P., and Saraiva, J. (2017). Energy efficiency across programming languages: how do energy, time, and memory relate? In *Proc. 10th ACM SIGPLAN Int. Conf. Softw. Language Eng.*, pages 256–267. ACM.

Radovanovic, A., Koningstein, R., Schneider, I., Chen, B., Duarte, A., Roy, B., Xiao, D., Haridasan, M., Hung, P., Care, N., Talukdar, S., Mullen, E., Smith, K., Cottman, M., and Cirne, W. (2023). Carbon-aware computing for datacenters. *IEEE Transactions on Power Systems*, 38(2):1270–1280.

Recker, J., Rosemann, M., and Gohar, E. R. (2011). Measuring the carbon footprint of business processes. In *BPM 2010 Int. Workshops and Education Track*, pages 511–520. Springer.

Roohy Gohar, S. and Indulska, M. (2020). Environmental sustainability through green business process management. *Australasian J. Inf. Sys.*, 24.

Seidel, S., Recker, J., and vom Brocke, J. (2012). Green business process management. In *Green business process management*, Springer Management/Business for professionals, pages 3–13. Springer.

van der Aalst, W. M. P. (2013). Business process management: A comprehensive survey. *ISRN Software Engineering*, 2013:1–37.

Zhu, X., Zhu, G., vanden Broucke, S., and Recker, J. (2015). On merging business process management and geographic information systems: Modeling and execution of ecological concerns in processes. In *Geoinformatics in resource management and sustainable ecosystem*, volume 482 of *Communications in Computer and Inf. Sci.*, pages 486–496. Springer.