

The Green Practitioner: A Decision-Making Tool for Green ICT

Karl Lundfall
Department of Computer Science
VU University Amsterdam
The Netherlands
Email: kalle.lundfall@gmail.com

Paola Grosso
Institute for Informatics
University of Amsterdam
The Netherlands
Email: p.grosso@uva.nl

Patricia Lago, Giuseppe Procaccianti
Department of Computer Science
VU University Amsterdam
The Netherlands
Email: {p.lago, g.procaccianti}@vu.nl

Abstract—The environmental impact of ICT represents an increasing concern for modern society. Despite that, we still witness a lack of awareness from the ICT industry regarding their carbon footprint. In particular, practitioners lack knowledge and tools to perform informed decision-making in the area of Green ICT. Moreover, best practices for Green ICT are usually conceived and applied in-house. This prevents them from being generalized and shared with the community. In this paper, we present a web-based tool that tackles both needs: on one hand, the tool enables practitioners (i.e. *agents*) to browse a number of Green ICT solutions and calculate their estimated impact if applied to their organization, through customized parameters. On the other hand, the tool also provides a creation interface for *contributors* who want to generalize and share their own Green ICT practices. In addition, we foresee that the data generated by the usage of the tool will provide useful insights and trends regarding the adoption of Green ICT and its effectiveness on reducing CO₂ emissions.

I. INTRODUCTION

The energy consumption of ICT is increasing rapidly towards substantial figures. As of today, data centers account for more than 1.5% of the world's total consumption [12]. On the other hand, the potential of reducing energy consumption and carbon dioxide footprint in this sector is promising, thanks to the numerous technological developments in the area [3]. Therefore, from an environmental perspective, it is crucial to promote and motivate the use of “green” practices in the ICT industry. An important milestone towards a more widespread adoption of such practices is the estimation of the potential impact of a green practice in advance, prior to its adoption. This estimation allows relevant stakeholders (e.g. data center administrators, infrastructure managers) to get concrete figures of the implications of certain green practices. In addition, this increases the awareness and incentive towards *going green*, by showing the positive economic impact of energy savings.

Our aim is to formalize green practices beyond textual descriptions in a standardized way that allows for easy sharing and customization of the practice to different situations. Furthermore we realize that an important gap to fill is knowledge transfer between developers of green practices and their adopters.

The work presented in this paper took as starting point a previous project [4] that had compiled some known green

practices into a public database¹. Every green practice currently available in the databases is summarized with textual descriptions and accompanied by source references for further reading. The authors also developed two prototypes of calculators for the economic and environmental impact of green ICT [2] and green networking [8] practices that also provided guidance in the development phases of our work.

The final result of our work, which is presented here, is *The Green Practitioner*², an online tool available via the Amsterdam Data Science web portal that allows to easily calculate the impact of Green ICT practices. Our tool supports two types of users: practitioners (*agents*) select practices/solutions and apply them to their own environment. *Contributors* are able to define new practices and solutions. In this manner we address the knowledge sharing gap we mentioned before, and move away from purely textual descriptions.

The remainder of this paper is structured as follows. We first present our definition of a green ICT practice and the related work in this area. Then we present the Green Practitioner tool. An illustrative example is presented to show both cases and guide the reader through the use of the tool. Then, we describe the conceptual model, i.e. the building blocks that are used to define practices. We then focus on how one can combine a set of selected practices into one reusable solution. The paper concludes with some discussion and suggestions for future work.

II. RELATED WORK

We define a *Green ICT practice* [4] as a method of making a process more environmentally sustainable within an ICT context.

Different kinds of green ICT practices pursue different goals, such as improving environmental awareness, promoting a sustainable economy, or saving energy in buildings and utilities. Some practices focus on behavioral and organizational changes: one could accommodate more flexible working schedules that enable employees to consume less energy in transit back and forth from work [1]; or one could enforce double-sided printing to save both paper and energy [9].

¹<http://greenpractice.few.vu.nl>

²<http://thegreenpractitioner.amsterdamdatascience.nl>

Likewise one could focus on more technical, ICT-related approaches such as introducing parallelization of operations using the GPU in order to make calculations more efficient [5], or applying cloud computing technology to significantly reduce hardware and software resources needed for individuals [7]. The body of work discovering and documenting these practices is continuously growing.

Reducing costs is regarded as the second major reason to adopt Green ICT practices [10] after reducing power consumption. Especially in times of economic crisis, cost reduction becomes the most important economic objective [13] of many companies. For this reason, making the economic impact of green practices explicit becomes crucial to promote their adoption. In addition, green ICT may save energy consumption and reduce costs but it often requires initial investments, business process changes, and extra effort from both companies and individuals. A clear indication of Return On Investment (ROI) and payback periods becomes necessary to justify the introduction of green ICT practices at management level.

III. THE GREEN PRACTITIONER

The Green Practitioner is a web-based tool to calculate the impact of Green ICT practices. It has been developed as the main outcome of a project supported by the Amsterdam Data Science institute. The tool has been programmed in PHP and uses the Bootstrap framework.

As we said before the tool supports both *agents* in choosing solutions and practices appropriate for their role, and *practioners* in defining them. In sectionIII-A we guide the reader through the interface of our tool that is seen and used by the agents; in sectionIII-B we show how a new solution can be defined in the database that is exposed by the tool. At the moment of writing, the Green Practitioner comes with one predefined solution for immediate use called *Revise current setting with greener technology*. We use this solution to illustrate both the usage and creation processes.

A. Using solutions

The Green Practitioner allows selection of a solution in two ways: either the user selects his/her role within an organization (i.e. the type of actor) and only then browses through the list of solutions that are related to and relevant for this role, or the user knows what she looks for and directly browses through solutions. The tool provides a short description of the solutions available to facilitate the selection.

As it will become clear from sectionIV where we will introduce the conceptual model that supports our tool, solutions are effectively packages of practices; so once an agent has identified the solution of interest it is still possible for her to select the complete package of practices or a subset thereof, as well as adopt them in a certain order.

The selected practices will request a number of inputs from the agent, in order to finally provide the calculation of the impact of the chosen green solution. The tools will provide the agents with customizable graphs where variables can be

included and excluded in the plot depending on what the user is interested in highlighting.

It is easy to understand the whole process by seeing how an user of our tool would navigate through the various phases, namely:

- role selection;
- solution selection;
- practices selection and ordering;
- practices customization;
- impact calculation.

As said, we have currently one example practice defined and available. Starting from the main page, the user will select the agent *data center administrator* as shown in figure 1. Based on this selection, he will get to choose from different practices that are potentially interesting to that agent.




Figure 1. Selecting an agent from the initial screen.

When choosing the presented solution, the user can select practices and their order of application, within the limitations defined in the solution itself (see Figure 2). By hovering over each practice, the user can get a short description of what the practice is about. A current limitation of the tool that we plan to address in future releases is the capability of applying practices within a solution in an incremental manner; this will allow to disentangle and evaluate their individual contribution.




Figure 2. Selecting practices of the solution.

In our example, the solution includes the following three practices that we select in the given order:

- 1) **Run servers as normal**
- 2) **Switch to a greener energy source**
- 3) **Implement efficient server technology**

Afterwards, the user is prompted to customize the solution to his own situation, by specifying the values of all input variables for each chosen practice. In our example, as shown in Figure 3, we first have to specify the variables relevant to the “Run servers as normal” practice.

When we move on to the next practice, in the example case “Switch to a greener energy source”, the user only has to specify the value of one variable, namely

1 Running servers as normal
 2 Switch to a greener energy

Consumption per server: w / Server

Number of servers in use: Servers

Energy cost: € / kWh

Time passed: Months

CO2 emission constant: g CO2 / kWh

Figure 3. Value specification for the first practice.

co2ConstantNew; the values associated to other variables can be re-used as they appeared already in the previously chosen practice. The screen looks like in Figure 4.

1 Running servers as normal
 2 Switch to a greener energy source

CO2 emission constant, new: g CO2 / kWh

Cost of new energy source: €

Electricity price for new energy source: € / kWh

Figure 4. Value specification for the second practice.

We fill a full specification of values in all of the three practices. To show a realistic situation, we let the new electricity price for the green energy source be more expensive than the old one. After doing so, we obtain the results screen shown in Figure 5 (where we chose to specify a time span of 3 months since the implementation of the practices). We observe that: (1) Expenses (initial vs. final) show that implementing the two green practices is more expensive for us; however, (2) operational expenses (Opex) are lower despite the new electricity price being more expensive. This is because of the power savings we obtained through the last practice. Also, (3) capital expenses (Capex) are causing the costs to rise. This suggests that we will eventually reach a point of break-even, after which we will start saving money by implementing these practices. Using the tool to calculate how our expenses change over a period of 24 months (instead of 3 months only), we see that (as shown in Figure 6) somewhere around 13 months the practices start paying off.

Opex normal: 1080 € / Month
 Opex final: 918 € / Month
 Expenses normal: 3240 €
 Expenses final: 4554 €
 Capex final: 1800 €
 Capex normal: 0 €
 Energy cost: 0.25 € / kWh

Figure 5. The results from projecting green practices 3 months ahead.

More importantly, from an environmental perspective the green practices make a substantial difference. Figure 7 shows




Figure 6. The monetary expenses (in €) projected on a 24-months time span.

that our recurring carbon footprint is roughly cut in half. This demo run is logged in our system and available for further exploration ³.




Figure 7. The carbon dioxide emissions (g CO₂) when seeing the time span range from 0 to 24 months.

B. Creating Solutions

Our tool relies on the solutions that have been defined by practitioners and that are stored in a back-end database. The practitioners rely on a creation interface available to them, which guides through the definition of all the components (see section IV). To illustrate the operation of the creation interface, we provide a concrete example for a solution called 'Revise current settings with greener technology'. For the sake of simplicity, the chosen associated practices are relatively simple.

1) *Units*: First all the units necessary for the example solution have to be defined. Table I shows them all. Note that we have not kept the notion of unit in a strict physical sense, as it might seem strange to consider months and servers as units. This is simply done in order to make things as clear as possible for the user. W is an abbreviation for watts.

³<http://goo.gl/67ua0Q>

2) *Variables*: Table II lists the *Variables* we defined for our example usecase. Notice that some of the variables are defined in pairs, where one declares the value under normal circumstances, and the other the value after applying one more practice. The pairs will come in handy for the user when being able to compare the situation of going green rather than continuing as normal. There are also pairs of variables that mimic other variables by adding "new". These are input variables where the "new" variable is referring to a future state. This will become more clear in the practice definition.

Unit	Unit type
Months	Integer
Servers	Integer
kWh	Floating point
W	Floating point
W / Server	Floating point
€	Floating point
€/ kWh	Floating point
€/ Month	Floating point
€/ Server	Floating point
g CO ₂	Floating point
g CO ₂ / kWh	Floating point

Table I

LIST OF *Units* DEFINED FOR THE EXAMPLE *Practice*

3) *Practices and Formulas*: Our example includes three available practices.

- **Run servers as normal.** In this practice, we calculate the environmental and economic implications of running the data center in a normal state, in order to get a benchmark for other practices. Observe that this is not really a green practice but it serves as a means of comparison, initializing the "...Normal" variables, in order to clarify things for the user. The formulas defined for the practice are:

```

- opexNormal = (consumptionServer * 24 * 30 *
  serversInUse * electricityPrice)/1000
- opexFinal = opexNormal
- exNormal = capexNormal + opexNormal *
  timePassed
- capexNormal = 0
- capexFinal = capexNormal
- exFinal = exNormal
- co2emissionNormal = consumptionServer*24*
  30 * serversInUse * timeFrame * co2Constant
- co2emissionFinal = co2emissionNormal

```

- **Implement efficient server technology.** Here we consider an investment in reducing the server consumption in order to achieve an overall lower power consumption. In this simplified model, we assume the consumption of the servers to be the total consumption of the entire database. Observe the similarity of formulas in comparison to the previous practice. Here the formulas are:

```

- consumptionServer = consumptionServer -
  consumptionServerRed

```

```

- opexFinal = (consumptionServer * 24 * 30 *
  electricityPrice * serversInUse)/1000
- capexFinal = capexFinal + costPerServer *
  serversInUse
- exFinal = capexFinal + opexFinal *
  timePassed
- co2emissionFinal = consumptionServer * 24 *
  30 * serversInUse * timePassed * co2Constant

```

- **Switch to a greener energy source.** By switching energy source, the carbon dioxide footprint is reduced and there is also likely to be (positive or negative) economic implications. When we change the emission factor and electricity price, the variables have to be adjusted proportionally to the old values. The formulas look like this:

```

- co2emissionFinal = co2emissionFinal *
  (co2ConstantNew/co2Constant
- co2Constant = co2ConstantNew
- capexFinal = capexFinal +
  costSwitchEnergySrc
- opexFinal = opexFinal *
  (electricityPriceNew/(electricityPrice)
- exFinal = capexFinal + opexFinal *
  timePassed
- electricityPrice = electricityPriceNew

```

Observe that all variables appearing to the right of the equality in each formula are used, except those which also appear as on the left side of the equality in an earlier formula. For example, `consumptionServer` appears to the right side of the equality in the first formula, so we are asked to specify it. `capexNormal` is at the left side of the equality in the formula `capexFinal = 0`, so it is re-used in the subsequent formula (`capexFinal = capexNormal`) and does not function as an input variable in the practice.

4) *Solutions*: We capture the three practices in one solution:

- **Revise current setting with greener technology.** We set a constraint on the practice selection, so that the user has to select the first practice ("Run servers as normal"), and it has to run as the first one. Then the user has to select at least one of the other two, or both, where the order does not matter.

5) *Agents*: The example has only one agent, the data center administrator, related to a single solution.

IV. CONCEPTUAL MODEL

Figure 8 shows the UML conceptual model behind the Green Practitioner. Its concepts are:

- **Unit:** The unit describes how something is measured (e.g. *watts*). A property of the unit is the type of the value that would be expected (e.g. *floating point number*).
- **Variable:** The variable represent a number representing something relevant to the practice (e.g. *server consumption*). Some variables might have an initial value that is assigned by default.

Variable name	Unit	Variable identifier	Description
Servers in use	Servers	serversInUse	The total number of servers used normally in the datacenter.
Time	Months	timePassed	The quantity of time that should be accounted for when accumulating costs. In this example we considered months of 30-days.
Consumption server reduction	W / Server	consumptionServerRed	The consumption reduction per server that can be made by applying a new, more efficient technology.
Electricity consumption of a server	W / Server	consumptionServer	The energy consumption of a server.
Operational expenses, normal	€/ Month	opexNormal	The monetary operational cost under normal circumstances.
Operational expenses, final	€/ Month	opexFinal	The monetary operational cost after implementing some practice(s).
Capital expenses, normal (This defaults to 0)	€/ Month	capexNormal	The initial monetary investment under normal circumstances.
Capital expenses, final	€/ Month	capexFinal	The initial monetary investment or income acquired when implementing practice(s).
Total expenses, normal	€/ Month	exNormal	The total amount of monetary expenses under normal circumstances.
Total expenses, final	€/ Month	exFinal	The total amount of monetary expenses after applying some practice(s).
Cost of new energy source	€	costSwitchNewEnergySrc	The one-time monetary cost of switching to a new energy source.
Electricity price	€/ kWh	electricityPrice	The monetary cost that the supplier charges per electricity quantity.
Electricity price, new	€/ kWh	electricityPriceNew	The monetary cost that the supplier charges per electricity quantity for another power source.
CO ₂ emission, normal	g CO ₂	co2emissionNormal	The amount of CO ₂ footprint that the data center accounts for under current circumstances.
CO ₂ emission, final	g CO ₂	co2emissionFinal	The amount of CO ₂ footprint that the data center accounts after applying some practice(s).
Upgrade cost	€/ Server	costPerServer	The cost of the investment per server to employ another more efficient technique.
CO ₂ emission factor	g CO ₂ / kWh	co2Constant	The weight of CO ₂ emitted per energy unit. Normal values could be 640 for crude oil, or 380 for natural gas.
CO ₂ emission factor for new energy source	g CO ₂ / kWh	co2ConstantNew	The weight of CO ₂ per energy unit for another energy source.

Table II

LIST OF *Variables* NEEDED FOR THE DEFINITION OF THE EXAMPLE *Practices*

- **Formula:** The formula sets the value of a variable, based on current states of the variable(s) used for input. (e.g. *total consumption=server consumption*number of servers*)
- **Practice:** A practice consists of a set of formulas that in combination represent the effect of the practice (e.g. *virtualization of servers*).
- **Solution:** A solution comprises multiple practices which could be selected together (e.g. *virtualization and consolidation in combination*).
- **Practice combination:** Within the solution, it may not be possible to order and select the practices in any form or way, therefore this component sets a limitation as to what practices can be selected in order for the solution to be of value (e.g. *no virtualization without consolidation*).
- **Agent:** A solution can belong to a number of agents, which categorizes solutions for whom the solution is relevant to (e.g. *data center administrator*).

V. COMBINING PRACTICES INTO SOLUTIONS

As mentioned earlier, a novelty of our approach is defining solutions as groups of practices. In doing so, we allow the user to select a subset of the available practices, if desired, as well as reorder them according to a specific temporal application. While powerful, this feature requires particular care in the computation as reordering the practices differently from the default order can result in different dependencies among them. For example, in isolation one practice can define total power consumption as being the sum of power consumption of the servers, and a second practice can take other sources into account. Combining the two practices will most likely not have the desired effect and would be confusing for the user. We illustrate in the following the algorithm that we implemented to handle the proper application of cascading practices.

There are in total n number of practices in the database $\mathbf{P} = p_1, p_2, p_3, \dots, p_n$, and every solution $\mathbf{S}_k = p_1, p_2, p_3, \dots, p_m$ is defined as a subset of these, so $\mathbf{S}_k \subseteq \mathbf{P}$ and $m < n$. In a specific experiment, the user will select an ordered set (a tuple) of practices $\mathbf{U} \subseteq \mathbf{S}_k$, where \mathbf{U} is referred to in figure 8 as *selected practices*.

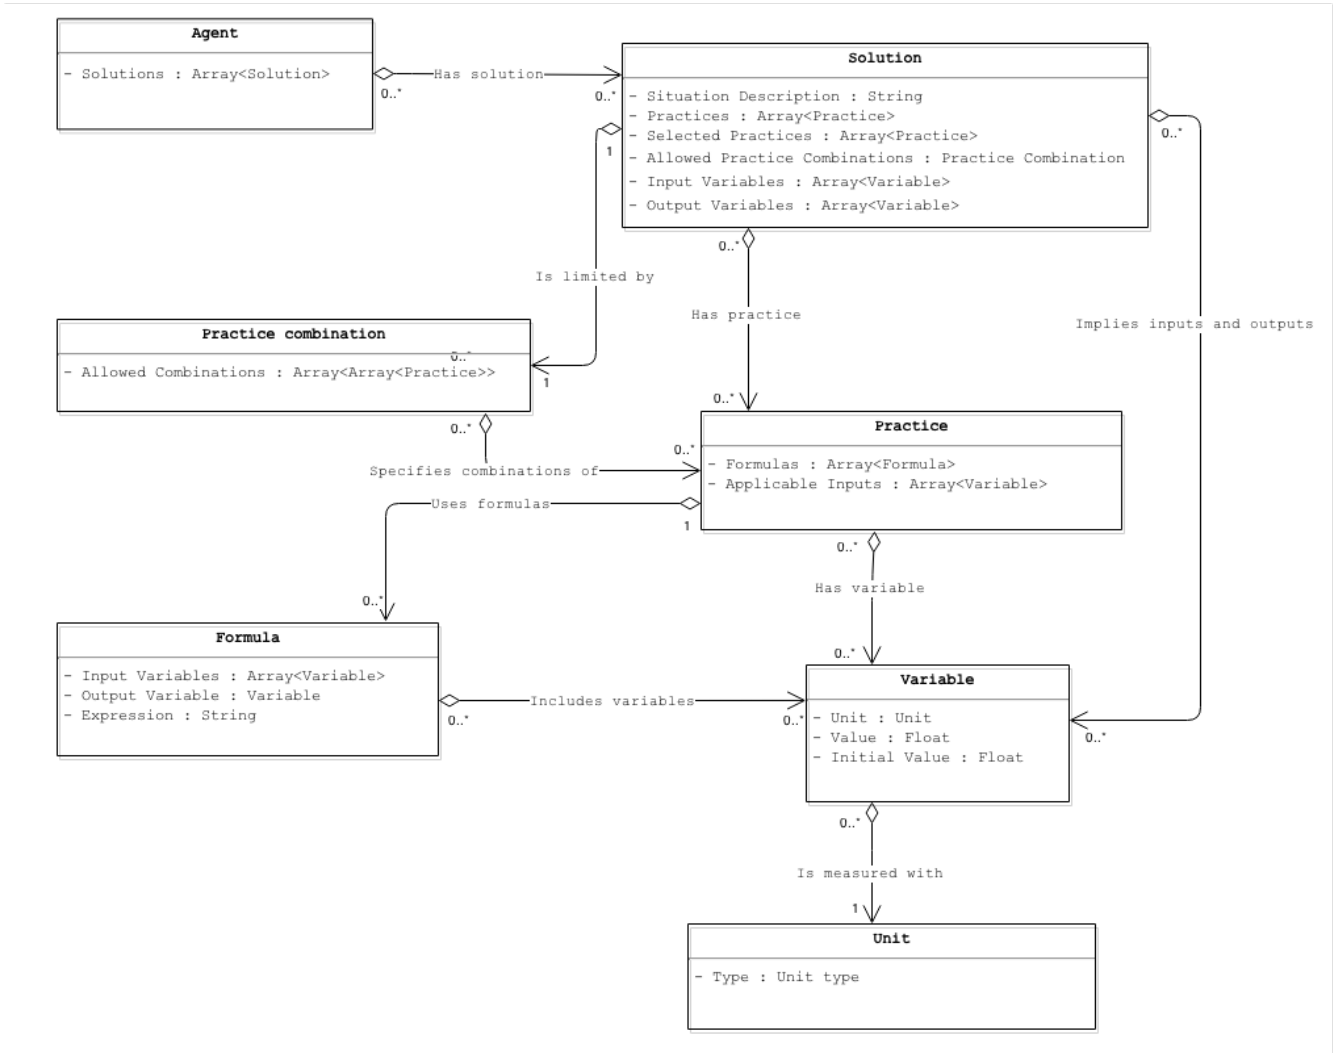


Figure 8. A conceptual model of The Green Practitioner components.

For each selected practice $p_u \in \mathbf{U}$, there belongs a number of formulas $\mathbf{F}_{p_u} = f_1, f_2, f_3, \dots, f_l$, which are executed in the specified order. All variables with unknown values have to be specified by the user in order for the execution to take place. These variables are the *input variables of the solution*.

There might be overlapping input variables, or some variables might be considered input variable in one practice and output variable in the other. It is therefore important to decide what the actual inputs and outputs for the entire selection of solutions will be. Here follows an algorithm describing how

to determine this as well as how to conduct the calculations of the formulas.

- 1) We start from a specified legal combination of practices: $U = p_1, p_2, p_3, \dots, p_n$. Let V_i be the input variables and V_o be the output variables.
- 2) For every practice p_i in U :
 - a) For every formula f_j in F_{p_i} :
 - i) For every variable v_k which appears on the right side of the equality representing an assignment in the formula f_j :
 - A) If v_k has not been assigned to a value yet:
 - **If v_k has a pre-specified initial value val_{v_k} :**
 - Assign: $v_k = val_{v_k}$
 - **Else:**
 - v_k needs to be specified by the user. This variable is now an input variable of the solution, so $v_k \in V_i$. Ask for user-input val_{user} and then assign $v_k = val_{user}$.
 - ii) Let the variable on the left side of the equality be v_o , which is set by calculating the expression. This is an output of the solution, so $v_o \in V_o$.
(Note that if a variable appears at both sides of the equality it will be both an input variable and an output variable).
- 3) Provide the user with all output variables in V_o .

VI. CONCLUSION

The Green Practitioner is a web-based tool relying on a strong conceptual model that offers a generalized method to define green practices and solutions. The tool gives to its users an estimate of the potential impact of adopting a solution in a specific context. At the same time the tool provides an easy to use interface for the input of new solutions and practices.

In this sense the Green Practitioner is the first tool of its kind and it brings two clear novelties. Firstly, it bridges the knowledge gap between ICT professionals that want to adopt new (greener) modes of operations and ICT professionals that have already identified and formalized such practices and want to share them broadly. Secondly, the tool itself and the methodology it implements can be used as a platform for further exploration of green practice-oriented research. To encourage both goals, the tool is openly accessible, as well as the solutions and practices that will be shared in it.

In our opinion, a first step toward making openly accessible solutions already available is to formalize practices put forward in the public sector and by standardization bodies, like CEPIS' Code of Best Practices for Green ICT, DCA's certification of data centre facilities, the Dutch Milieukeur's certification criteria for data centres, or the German blauer engel. If available via The Green Practitioner, many of such practices could be easily accessed, used for fast and cost-less

estimation of potential effects on the own organization and possibly accelerate adoption.

The tool is flexible and meant to be generally applicable. However, defining practices requires some effort. Care has to be taken in correctly defining the semantic meaning of each variable, and how practices are intended to be combined. Further research is needed for gaining experience in formalizing solutions to better appreciate the inter-dependencies of the combined practices. Also, the energy impact of documented green practices or solutions (e.g. [6, 11]) must be first measured. To this aim, our research group inaugurated in year 2015 the new Green Lab where students and researchers design and execute empirical studies on energy-aware software solutions. Relevant results will eventually be turned into solutions for The Green Practitioner.

ACKNOWLEDGMENTS

This research has been funded by the Amsterdam Data Science institute under the SEED project BigData2Nets. It received partial funding by project SIA RAAK MKB Greening the Cloud.

REFERENCES

- [1] Marie-Claude Boudreau, Adela Chen, and Mark Huber. Green IS: Building sustainable business practices. *Information Systems: A Global Text*, 2008.
- [2] Hector Fernandez, Giuseppe Procaccianti, and Patricia Lago. Economic aspects of green ICT. In *Green in Software Engineering*. Springer, 2014.
- [3] GeSI. GeSI SMARTer2020: The role of ICT in driving a sustainable future. Technical report, Global e-Sustainability Initiative, December 2012.
- [4] Qing Gu, Patricia Lago, Henry Muccini, and Simone Potenza. A categorization of green practices used by Dutch data centers. In *3rd International Conference on Sustainable Energy Information Technology*, page 7. Elsevier, jun 2013.
- [5] Song Huang, Shucai Xiao, and Wu-chun Feng. On the energy efficiency of graphics processing units for scientific computing. In *Parallel & Distributed Processing, 2009. IPDPS 2009. IEEE International Symposium on*, pages 1–8. IEEE, 2009.
- [6] Grace Lewis and Patricia Lago. A catalog of architectural tactics for cyber-foraging. In *International Conference on the Quality of Software Architectures*, pages 53–62. ACM, May 2015.
- [7] Liang Liu, Hao Wang, Xue Liu, Xing Jin, Wen Bo He, Qing Bo Wang, and Ying Chen. GreenCloud: a new architecture for green data center. In *Proceedings of the 6th international conference industry session on Autonomic computing and communications industry session*, pages 29–38. ACM, 2009.
- [8] Marc X. Makkes, Dexter Drupsteen, Arie Taal, and Paola Grosso. Storage to energy: modeling the carbon emission of storage task offloading between data centers.

Proceeding of: IEEE CCNC, IEEE 7th International Symposium(4):56–63, January 2014.

- [9] Robert L. Mitchell. Get up to speed on green IT. Technical report, Computerworld, 2008.
- [10] San Murugesan. Harnessing green IT: Principles and practices. *IT Professional*, 10:24–33, January 2008.
- [11] Giuseppe Procaccianti, Patricia Lago, and Grace A. Lewis. Green architectural tactics for the cloud. In *Working IEEE/IFIP Conference on Software Architecture*, pages 41–44. IEEE Computer Society, 2014.
- [12] Jonathan Ruddy and Mark O’Malley. Global shifting of data centers demand. *Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*, 2014 IEEE PES, IEEE 7th International Symposium(4):1–5, October 2014.
- [13] Pradipta Sarkar and Leslie Young. Managerial attitudes towards green IT: An explorative study of policy drivers. In *Proceedings of PACIS*, pages 1–14, 2009.