



## **The future of sustainable digital infrastructures: A landscape of solutions, adoption factors, impediments, open problems, and scenarios**

Downloaded from: <https://research.chalmers.se>, 2025-07-01 23:29 UTC

Citation for the original published paper (version of record):

Verdecchia, R., Lago, P., de Vries, C. (2022). The future of sustainable digital infrastructures: A landscape of solutions, adoption factors, impediments, open problems, and scenarios. *Sustainable Computing: Informatics and Systems*, 35. <http://dx.doi.org/10.1016/j.suscom.2022.100767>

N.B. When citing this work, cite the original published paper.



# The future of sustainable digital infrastructures: A landscape of solutions, adoption factors, impediments, open problems, and scenarios

Roberto Verdecchia<sup>a,\*</sup>, Patricia Lago<sup>a,b,\*</sup>, Carol de Vries<sup>c</sup>

<sup>a</sup> Vrije Universiteit Amsterdam, The Netherlands

<sup>b</sup> Chalmers University of Technology, Sweden

<sup>c</sup> PhotonDelta, Eindhoven, The Netherlands

## ARTICLE INFO

### Keywords:

Sustainability  
Green IT  
Energy efficiency  
Digital infrastructures  
Data centers  
Cloud  
Landscape  
Qualitative research

## ABSTRACT

**Background:** Digital infrastructures, *i.e.*, ICT systems, or system-of-systems, providing digital capabilities, such as storage and computational services, are experiencing an ever-growing demand for data consumption, which is only expected to increase in the future. This trend leads to a question we need to answer: How can we evolve digital infrastructures to keep up with the increasing data demand in a sustainable way?

**Objective:** The goal of this study is to understand what is the future of sustainable digital infrastructures, in terms of: which solutions are, or will be, available to sustainably evolve digital infrastructures, and which are the related adoption factors, impediments, and open problems.

**Method:** We carried out a 3-phase mixed-method qualitative empirical study, comprising semi-structured interviews, followed by focus groups, and a plenary session with parallel working groups. In total, we conducted 13 sessions involving 48 digital infrastructure practitioners and researchers.

**Results:** From our investigation emerges a landscape for sustainable digital infrastructures, composed of 30 solutions, 5 adoption factors, 4 impediments, and 13 open problems. We further synthesized our results in 4 incremental scenarios, which outline the future evolution of sustainable digital infrastructures.

**Conclusions:** From an initial shift from on-premise to the cloud, as time progresses, digital infrastructures are expected to become increasingly distributed, till it will be possible to dynamically allocate resources by following time, space, and energy. Numerous solutions will support this change, but digital infrastructures are envisaged to be able to evolve sustainably only by (i) gaining a wider awareness of digital sustainability, (ii) holding every party accountable for their sustainability throughout value chains, and (iii) establishing cross-domain collaborations.

## 1. Introduction

With the introduction of high bandwidth data transfers, affordable data plans, the generalized migration to the cloud of software applications and data management, and the popularization of streaming services, digital infrastructures are experiencing an ever-growing demand of data consumption [1]. As expected, the related energy consumption is steadily increasing over time. This motivated sector leaders like Microsoft, Google and Amazon, to increasingly adopt renewable energy resources, *e.g.*, solar and wind farms, as a means to lower the environmental impact of their hyperscale data centers. Nevertheless, adopting renewable energy can be considered only as part of the solution, as (i) such adoption does not tackle the need to optimize the use of cloud resources, and (ii) the production of renewable energy will not meet its demands already in the near future. With the global transition toward the adoption of renewable energy resources, the whole society

will need them. Therefore, exploiting renewables for the future digital infrastructures will not, as such, make them sustainable, but rather they need to become energy efficient, too, not to compete with the other industrial sectors. This is especially true when considering Europe, where prominent “data hubs” such as the Netherlands have to be situated in relatively small geographic areas.

For the last decades the digital infrastructure industry has been able to maintain a relentless pace of introducing new generations of faster and more energy efficient computing hardware approximately every two years [2]. Data consumption is, however, rising faster than the improvement in energy efficiency and now also the so-called “Dennard scaling” [3], that allowed lower power consumption with each new semiconductor generation, is irremediably coming to an end. In addition, with the ever-growing increase of data transport speed, the power consumed in wiring and communication is rising more than linearly. To

\* Corresponding authors.

E-mail addresses: [r.verdecchia@vu.nl](mailto:r.verdecchia@vu.nl) (R. Verdecchia), [p.lago@vu.nl](mailto:p.lago@vu.nl) (P. Lago), [Carol@photondelta.com](mailto:Carol@photondelta.com) (C. de Vries).

<https://doi.org/10.1016/j.suscom.2022.100767>

Received 10 January 2022; Received in revised form 23 March 2022; Accepted 21 May 2022

Available online 31 May 2022

2210-5379/© 2022 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

maintain the increase in data processing power, innovative solutions are needed.

Our work was carried out in the context of the *Lower Energy Acceleration Program*<sup>1</sup> (LEAP) exploring alternative solutions towards a sustainable growth of the data center industry. The aim of LEAP is to accelerate the transition towards sustainable digital infrastructures by integrating innovative developments at the heart of our energy and digital infrastructures. One of the goals of the LEAP, and focus of our work, is to develop a technology landscape for energy efficient digital infrastructures. The landscape focuses on three different temporal horizons, namely:

- Horizon 1 (H1): State of the art (today)
- Horizon 2 (H2): Within the next 4–6 years (near future)
- Horizon 3 (H3): Beyond 6 years (future)

In our previous study [4], we conducted a set of exploratory interviews involving 11 participants, in order to gain a preliminary understanding of the topic of sustainable digital infrastructures. Such study constituted the first of the three research phases (Phase 1, see Section 3.2.2) on which the study reported in this paper is based. In addition to conducting two additional research phases (Phase 2 and Phase 3, see Sections 3.2.3 and 3.2.4), and extending our preliminary findings, in this paper we discuss the solutions that according to a diverse set of experts can be adopted to sustainably evolve digital infrastructures, and their related adoption factors, impediments, open problems, and evolution scenarios.

The results reported in this paper are based on a mixed-method empirical process, born from the combination of semi-structured interviews (Phase 1, see Section 3.2.2), focus groups (Phase 2, see Section 3.2.3), and working groups involving a larger number of participants (Phase 3, see Section 3.2.4). In total, 48 unique participants, belonging to heterogeneous sectors relevant to the domain of digital infrastructures, took part in this study.

The contribution of this paper is a *landscape for sustainable digital infrastructures*, comprising:

- *Solutions* that are or will be available in the future to sustainably evolve digital infrastructures;
- *Adoption factors* that will guide the deployment and support the success of solutions;
- *Impediments* that may inhibit the adoption of solutions;
- *Open problems* that need to be tackled to sustainably evolve digital infrastructures;
- *Four future scenarios* of incremental nature, for the sustainable evolution of digital infrastructures.

The *target audience* of this paper includes (i) researchers studying the sustainability of digital infrastructures, who can profit from this investigation by deepening their knowledge of the state-of-the-art of sustainable digital infrastructures, and the key concepts, many of which open for future research, that characterize their present and future evolution; (ii) practitioners interested in improving the sustainability of their digital infrastructures, who can gain insights into the current and future solutions they may adopt and/or invest in, the trends that will shape the future of their infrastructures, and the related concepts to be considered while reasoning about the sustainability of digital infrastructures; and (iii) any readers interested in gaining awareness of sustainable digital infrastructures, who can gather an overview of the key related concepts, and the factors that will characterize the phenomenon of sustainable digital infrastructure evolution.

This paper is structured as follows. In Section 2 we introduce the essential background knowledge required to understand our study. In Section 3 we describe the study design and execution. Section 4

presents the results of our research, in terms of identified sustainable digital infrastructure solutions, adoption factors, impediments, and open problems. In Section 5 we discuss our findings, by explicitly answering our research questions, presenting four scenarios for the future of digital infrastructures that emerge from our results, and providing further considerations on the results we obtained. Section 6 presents the threats to validity that, despite our best efforts, may have influenced the results of our study. In Section 7, we document a review of the related work, by discussing the main commonalities and differences w.r.t. our study. Finally, Section 8 closes the paper by presenting the conclusions and directions for future work.

## 2. Background

Over the years, numerous studies focusing on various aspects of the sustainability of digital infrastructures have been presented in the literature [5]. As this research field progressed, the concept of “sustainability dimension” [6] has been introduced in order to systematically reason about the multifaceted nature of (digital) sustainability. Based on this concept, among others, Lago et al. [6] reworked the definition into four core dimensions of sustainability, namely *economic*, *technical*, *social*, and *environmental*. The *technical dimension* concerns the long-term use of digital infrastructures and their appropriate evolution in a continuously changing execution environment. The *economic dimension* instead regards the preservation of capital and economic value of digital infrastructures. The *social dimension* regards the support of current and future generations to possess the same or greater access to social resources by pursuing generational equity. For digital infrastructures, this dimension includes from the direct support of social communities, to the promotion of indirect activities creating benefits for such communities. Finally, the *environmental dimension* focuses on the improvement of welfare while protecting natural resources. For digital infrastructures, this dimension aims at addressing ecologic concerns, including energy efficiency and ecologic awareness creation. In this study, we leverage the four sustainability dimensions as defined by Condori et al. [7] in order to systematically reason about and categorize the concepts emerging from our research.

As further clarification on the concept of the sustainability of digital infrastructures, as utilized in this investigation, our primary focus is on the *energy efficiency* of digital infrastructures. Hence, while some concepts related to energy efficiency (e.g., *strategy for awareness creation*, see Section 4.2.1) may be touched upon, other sustainability aspects of digital infrastructures (e.g., life cycle assessment, carbon footprint, circularity and waste management, etc.) fall outside the scope of the study reported in this paper.

Regarding the concept of “digital infrastructure”, in this paper we utilize such term to denote ICT systems, or system-of-systems, providing digital capabilities, such as storage and computational services. For example a digital infrastructure, as defined in this paper, can range from the composition of a single hyperscale server and a thin client, to a completely distributed network of edge devices. Instead of considering different dimensions, e.g., socio-technical [8,9] or socio-economic [10,11], our definition considers digital infrastructures as a primarily technical concept [12]. Adopting such definition allows us to study the topic from a well-defined and focused angle, scope-down our research design (e.g., in terms of stakeholders to be considered), and gain more depth in our research design and related results.

## 3. Study design and execution

In order to develop our landscape of sustainable digital infrastructures, we designed and conducted a qualitative mixed-method empirical study [13]. In this section, we describe the main aspects of our study design and execution. Specifically, in Section 3.1 we report the research goal and research questions of the study, while in Section 3 we describe the research process followed to obtain our results.

<sup>1</sup> <https://amsterdameconomicboard.com/en/initiative/leap>.




Fig. 1. Research process overview (S: Number of distinct sessions; P: Total number of participants).

### 3.1. Research goal and research questions

The overarching goal of this research is to understand how we can sustainably progress the development of digital infrastructures in the future, what are the key solutions, impediments, adoption factors, and related open problems. More formally, by applying the research goal formulation proposed by Basili et al. [14], the objective of our investigation can be defined as follows:

**Analyze digital infrastructures**

**For the purpose of** identifying solutions, impediments, adoption factors, and open problems

**With respect to sustainability**

**From the viewpoint of** practitioners and researchers

**In the context of** current and future digital infrastructure development.

By achieving our goal, we aim at providing concrete insights on how to transition towards sustainable digital infrastructures, by making explicit the determining concepts of this phenomenon. The concepts, reflected by four sub-research questions ( $RQ_1 - RQ_4$ ), are (i) sustainable infrastructure solutions, (ii) adoption factors, (iii) impediments, and (iv) open problems, as further articulated in the reminder of this section. By combining the specific concepts emerging from our study we are able to derive four scenarios which answer our main research question (RQ). An overview of the concepts, their dependencies, and mapping to the research questions is provided in Fig. 7.b.

In order to achieve our goal, we have to answer the main RQ of our study, namely:

**RQ** What is the future of sustainable digital infrastructures?

With our RQ, we aim at studying the state-of-the-art of digital infrastructure sustainability, and understand how their ever-growing requirements can be satisfied in a sustainable fashion in the future. More specifically, we want to investigate what solutions for sustainable digital infrastructures can be utilized, and in addition gain further knowledge of the related impediments, adoption factors, and open problems. In order to systematically answer our main research question, we decompose our RQ in four sub-RQ as follows.

**RQ<sub>1</sub>** What are the solutions to develop sustainable digital infrastructures?

With  $RQ_1$ , we aim at identifying what present and future solutions can be utilized in order to create and maintain sustainable digital infrastructures. With our subsequent sub-RQs instead, we aim at further investigating concepts related to the identified solutions.

**RQ<sub>2</sub>** What drives the adoption of sustainable digital infrastructure solutions?

With  $RQ_2$ , we aim at identifying what can drive the adoption of sustainable digital infrastructure solutions. More specifically, we aim at identifying what technological, social, environmental, and economic factors can stimulate and accelerate the adoption of solutions.

**RQ<sub>3</sub>** What hinders the adoption of sustainable digital infrastructure solutions?

As a mirror question w.r.t.  $RQ_2$ , after we identify the key challenges of employing sustainable digital infrastructure solutions, we aim at understanding which are instead what can hinder or impede the adoption of such solutions. This is expressed in our third sub-RQ ( $RQ_3$ ).

**RQ<sub>4</sub>** What are the open problems related to sustainable digital infrastructure solutions?

With  $RQ_4$ , we intend to investigate what are the open problems and challenges related to sustainable digital infrastructures. Answering this last sub-RQ allows us to understand what problems need to be tackled in order to ensure a sustainable future development of digital infrastructures.

### 3.2. Research process

An overview of the research process followed in this study is depicted in Fig. 1. During Phase 1, we conducted a series of exploratory interviews with targeted participants. This phase was adopted to gather the initial insights required to develop our landscape of sustainable digital infrastructures (for more information on Phase 1, refer to Section 3.2.2). In Phase 2, we conducted a set of focus groups [15] with targeted participants. This phase was adopted to deepen and refine our preliminary findings, leading to a preliminary outline of our landscape. A detailed documentation of how we applied the focus group method in Phase 2 is reported in Section 3.2.3. Finally in Phase 3, we started by presenting our findings to an audience of targeted participants. The presentation served as an introduction for a set of subsequent working groups, which were carried out in parallel after the presentation. Phase 3 was used to gather feedback from a larger audience of experts, in order to assess the correctness and completeness of our findings. The results presented in this paper are the results of the incremental findings collected during the three phases of our research, which are further detailed below.

#### 3.2.1. Participant sampling

Regarding the sampling strategy adopted to select participants for our investigation, for all three phases we leveraged a mixed-method non-probability sampling, consisting of accidental sampling guided by *ad hoc* quotas [16]. This allowed us to involve participants belonging to heterogeneous sectors relevant for the development and maintenance of digital infrastructures (see Fig. 2). In addition to accidental sampling supported by *ad hoc* quotas, we utilized purposive sampling [16] to further refine the selection of participants. Specifically, our investigation focuses on the *future evolution* of sustainable digital infrastructures. Therefore, the participants were required to possess not only knowledge of the state-of-the-art and/or state-of-the-practice about digital infrastructures, but also sufficient insights to carry out the speculative forecasting process implied by our RQs. To ensure the participants possessed the required knowledge, we selected representatives of the various organizations in the digital infrastructure value chain. The selection of participants was conducted by ensuring, through the analysis of their current role and past experiences, that they possessed both technical expertise and a broad strategic overview of the topic within their area of expertise.




Fig. 2. Distribution of study participants across sectors.<sup>2</sup>

Participants were involved in one or more of our research phases, as further detailed below. More demographic information regarding our participants, in terms of sectors, unique companies, and research phase they were involved in, is reported in Table 4 of Appendix.

Regarding the sampling of Phase 1 participants, the selection process started by gathering an initial selection of diverse companies working in the field of digital infrastructures, by considering as starting set the companies involved in the LEAP Initiative and/or connected to the Amsterdam Economic Board. Subsequently, employees of such companies, who worked in roles relevant for our investigation (e.g., digital sustainability managers, CTOs, and sustainability researchers) were contacted to inquire about their availability to partake in our interviews. After gathering an initial set of available participants, we selected our final set of 11 Phase 1 participants by ensuring that different stakeholders of the digital infrastructure domain, e.g., hardware manufacturers, cloud providers, software service providers, consultancy firms, academic/research institutes, and digital infrastructure customers, would participate in Phase 1.

Similarly, for Phase 2, we sampled participants from the ones suggested during our research Phase 1 (see also Q4, Table 1), and added additional participants, by considering the participants discarded from Phase 1 (due to satisfied quotas). As for the sampling process of Phase 1, *ad hoc* quotas were utilized to ensure that digital infrastructure stakeholders of different sectors were equally represented during Phase 2. In total 19 participants took part in this phase.

Finally, for Phase 3 we re-invited the participants who were involved in either Phase 1 or Phase 2. This allowed us to gain feedback on our findings, and assess the quality of our data analysis and synthesis processes by presenting our results directly to the people who provided the data. In total, 30 participants were involved in this phase, with 13 participants who were involved also in a previous phase (5 Phase 1 participants, and 8 Phase 2 participants, see also Fig. 3). The remaining 17 participants of Phase 3, (i.e., the participants who were involved only in this research phase) were selected by considering the potential participants of Phase 1–2 who were discarded due to met quotas, integrated with further participants suggested as relevant and potentially missing by the participants of Phase 2.

In total, our study involved 48 unique participants, employed by 33 different organizations belonging to heterogeneous sectors (see Fig. 2 and Table 4).




Fig. 3. Distribution of study participants across research Phases sectors.

### 3.2.2. Phase 1: Exploratory interviews

During our first research phase, we carried out a set of semi-structured interviews with targeted practitioners and researchers. The goal of these interviews was to gain a preliminary understanding of the topic under investigation, and gather the groundwork of findings to be refined in the subsequent research phases.

Regarding the structure of the interviews, each interview started with a general exposition of our planned research process and goal (as presented in this section). Subsequently, 5 interview questions of open-ended nature were posed to the participants. Follow-up questions were posed after the answers of the participants, if these required further clarification or enrichment.

The interview questions (see Table 1) were purposely designed as wide-ranging, in order to let participants express their expert opinion on sustainable digital infrastructures without biasing them towards focusing on a specific topic. Specifically, with Q1 we aimed at gaining an initial understanding of the areas of development deemed by the participants as concrete options to build sustainable digital infrastructures. With Q2 instead, we explored concepts related to sustainable digital infrastructure solutions, in terms of the factors relevant to adopt the solutions. While with Q1 we focused on the solutions available in the near future, we used Q3 to investigate the solutions which will be available in the longer-term, by exploring development areas showcasing high potential in the future. Q4 focused on preparing for our next research phase (Phase 2), by deploying a chain-referral sampling strategy [17], executed by asking the participants who, in their opinion, are the key stakeholders/experts/decision makers of the area under investigation. Finally, with Q5, the participants were provided the possibility to add further insights or remarks, in order to avoid the accidental omission of potentially relevant data collection during this phase.

<sup>2</sup> As some participants could be mapped to more than one sector (cf. Table 4), the total sum of participants depicted in this figure exceeds the total sum of unique participants of our study (48).



**Table 1**  
Interview questions (Phase 1).

ID	Question
Q1	What will be the next-generation energy efficient digital infrastructure solutions?
Q2	What are the key factors which are important to adopt energy efficient digital infrastructure solutions?
Q3	Are you aware of any interesting developments in energy efficient digital infrastructure solutions, which show high potential to become successful in the future?
Q4	Are you aware of other key stakeholders/experts/decision makers on energy efficient digital infrastructure solutions to contact?
Q5	Anything to add?

Each interview was conducted and moderated by the three authors of this paper. Seven interview sessions were conducted during this phase, by considering a total of 11 participants. Each interview involved between one and three participants, was conducted via video conference, and lasted between 42 and 58 min. The interviews were video-recorded and transcribed manually by following the denaturalism approach, *i.e.*, by correcting grammar, removing interview noise (*e.g.*, stutters) and unifying nonstandard accents (*i.e.*, non-majority), while ensuring a full and faithful transcription [18].

The final output of Phase 1 consisted in a confidential technical report summarizing our preliminary findings gathered during this phase, in terms of identified solutions, impediments, adoption factors, and open problems. For more information regarding the data analysis and synthesis process followed, refer to the dedicated section (Section 3.2.5).

### 3.2.3. Phase 2: Focus groups

In order to expand and refine the findings of Phase 1, in Phase 2 we applied the focus group method [15]. Specifically, Phase 2 consisted of presenting to groups of participants 3 thinking scenarios, followed by the findings of Phase 1. This research phase was conducted to gather feedback based on the discussion among participants to extend and refine our preliminary results. As an introductory step to Phase 2, the technical report produced as output of Phase 1 was shared with the participants a week prior to the focus group session. This process ensured that participants possessed sufficient background knowledge on the topics to be covered, and ensured that participants assimilated our preliminary results before the session.

Similar to Phase 1, each focus group session started with a general exposition of our research process and goal (as presented in this section). Subsequently, participants were presented 3 thinking scenarios (summarized in Table 2) they had to jointly reason about. The first thinking scenario (T1), was utilized to explore, from the point of view of the participants, what would be the future evolution of digital infrastructures, in terms of software/hardware technologies, deployment strategies, and other related concepts. After discussing the future of digital infrastructures, with T2 we introduced to participants the concept of limitations on the energy to be consumed in the future by digital infrastructures. This second thinking scenarios allowed us to investigate what solutions, according to the participant, would progress the development of digital infrastructure in a sustainable fashion, as dictated by limitations on the energy resources to be consumed. Finally, with thinking scenario T3, we further pushed the boundary of the limitations imposed on the energy consumed by digital infrastructures, by envisioning a long-term future where digital infrastructures have to grow in a sustainable manner by consuming the same amount of energy as today. This last thinking scenario was used to explore future breakthrough solutions which will influence the long-term future of sustainable digital infrastructures, and what concepts, such as open problems to be solved and adoption impediments, will arise.

After the discussion of the 3 thinking scenarios, the preliminary solutions, adoption factors, impediments, and open problems identified

**Table 2**  
Thinking scenarios (part of Phase 2).

ID	Question
T1	From your perspective, think of your investments in digital infrastructures of the next 5+ years: what would you invest in?
T2	Imagine now a future where you have to operate with limited energy resources. What would you invest in?
T3	In 5 or 10 years time you have to deliver your customer service within the same energy budget of today. What breakthrough solutions would you need? And how can this be a success?

during Phase 1 were presented to the participants. This step allowed us to gather feedback on our preliminary results, expand them, and refine them according to the discussions within the participants of the focus groups.

As conclusion of the focus group sessions, participants were given the possibility to add further insights or comments on the topic under investigation, to ensure that no relevant information was omitted from the data gathering process of this phase.

The 19 participants of this phase were distributed over 5 focus groups, with the participation ranging from 2 up to 6 participants per focus group. The sessions were moderated by the first two authors of the study, with the exception of one focus group, which was moderated by the third author. The sessions were conducted via video conference, and lasted between 55 and 108 min. The interviews were video-recorded and transcribed manually by following the denaturalism approach (as described in Section 3.2.3).

The final output of Phase 2 consisted in a confidential technical report summarizing the findings of this phase into a preliminary landscape of solutions, adoption factors, impediments, and open problems. As for Phase 1, the data analysis and synthesis process of this phase is reported in Section 3.2.5.

### 3.2.4. Phase 3: Preliminary landscape presentation & working groups

In Phase 3, we evaluated and reviewed our findings by involving a larger audience of participants. Specifically, this phase was carried out as a single session, and consisted of a plenary presentation of our preliminary landscape (*i.e.*, the findings of Phase 2), followed by 5 working groups, carried out in parallel, where participants could express their opinion and feedback regarding our results. The plenary presentation consisted in presenting the overarching goal and research process of this study, followed by the exposition of the results obtained in Phase 2, *i.e.*, the identified solutions, adoption factors, impediments, and open problems. The presentation lasted a total of 25 min.

The subsequent working groups consisted of a discussion between participants, each supervised by a moderator, and a research support member, in charge of transcribing the discussion and supporting the interaction among participants. In total, 5 working group were conducted, by involving 5 participants in each working group. Each working group was conducted with the support of a Miro board,<sup>3</sup> *i.e.*, an online whiteboard where participants could use virtual “sticky notes” to map their comments, feedback, and ideas on a summary graphical representation of our findings. An example of Miro board structure, as utilized in the working group sessions, is presented in Fig. 4. Each Miro board was composed of (A) an overview of the working group participants, (B) two general question to kick-start the feedback session and discussion, namely: “What would you like to add to each of the categories?” and “What is most important in your view for each of the categories?”, (D)–(G) summary of the results reported in the preliminary landscape report (obtained via Phase 2), and (H) feedback notes created by participants during the session. In total, the parallel working groups lasted 45 min, and were followed by a plenary *post mortem* phase,

<sup>3</sup> <https://miro.com/>.

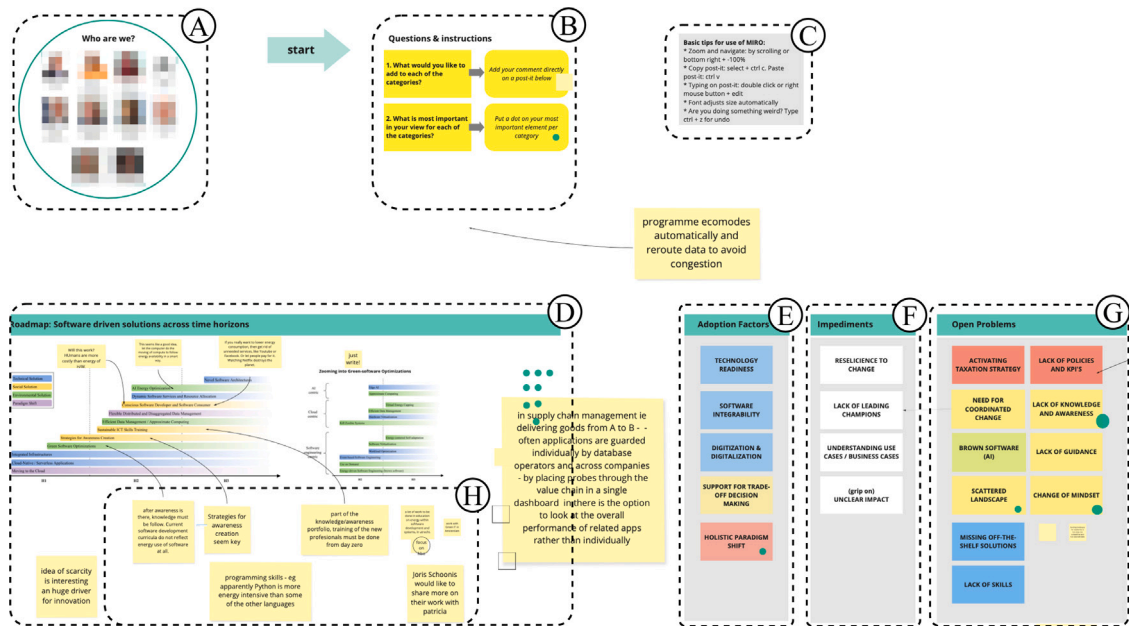


Fig. 4. Example of Miro board structure, as utilized in the working groups (part of Phase 3), where (A) summarizes participants, (B) general process instructions, (C) basic Miro tips, (D)–(G) results of Phase 2, and (H) exemplary positioning of feedback notes.

lasting 15 min, where the key takeaways of the working groups were presented and discussed.

One week prior to the execution of Phase 3, the technical report produced as output of Phase 2 was shared with the participants, to assure they possessed sufficient background knowledge, and did not have to assimilate the results of our preliminary landscape exclusively during the introductory plenary presentation of Phase 3.

After Phase 3, participants were provided the possibility to integrate their comments and feedback by contacting the authors of this study via mail. After all comments were assessed and integrated, the final results of the Phase 3 data analysis and synthesis process was shared with the participants, in order to validate our findings and further refine them according to their final feedback. The output of this last step resulted in this paper.

### 3.2.5. Data analysis and synthesis

The data analysis and synthesis process followed in this research was defined *a priori*, and consisted of 4 core sub-processes, namely (i) incident identification, (ii) coding, (iii) constant comparison, and (iv) memoing. The data analysis and collection was carried out 3 times, after each of our research phases (see Fig. 1). The data analysis was executed after the raw output of each phase, i.e., video-recordings (Phase 1–2) and feedback notes/transcriptions (Phase 3), were reported in textual form (e.g., via transcription for Phase 1–2, see Sections 3.2.2–3.2.3) and organized for analysis.

A detailed description of each of the research sub-processes followed in our data analysis and synthesis steps is described in the following:

- i. **Incident identification:** As a first step, in order to systematically analyze our data, we pre-processed it by subdividing the textual output of each phase into *incidents* [19], i.e., bits of data, such as sentences or paragraphs, which were related to the phenomenon investigated in our research;
- ii. **Coding:** The core of our data analysis relied on a coding process [20]. Specifically, in a initial coding phase, we adopted a mix of provisional and open coding to classify our data into different categories. More in detail, via provisional coding we mapped each incident to one of the four concepts at the basis of our study and RQs, namely *solutions*, *impediments*, *adoption factors*, and *open problems*. Further provisional coding was then used

to map the four aforementioned concepts to their sustainability dimension (i.e., *technical*, *social*, *economic*, or *environmental*), and the solutions to their respective time horizon (H1–H3). Following this preliminary provisional coding step, open coding was used in order to develop substantial codes describing our data. This process was conducted by attaching to each incident a keyword summarizing its content, in order to swiftly analyze and compare the content of different incidents in the next coding step. Subsequently, we adopted selective coding in order to identify the core concepts of our data, leading to the establishment of the concepts and categories, i.e., the solutions, impediments, adoption factors, and open problems, which constitute the results of our research (see Section 4). Finally, in order to draw our main conclusions, and answer our main research question (RQ), we adopted selective coding to abstract and synthesize our findings into four scenarios, presented in the discussion of this work (Section 5,) with which the energy efficiency of digital infrastructures can be addressed in the future.

- iii. **Constant Comparison:** In the execution of each data analysis and synthesis processes, all research artifacts (i.e., incidents, codes, categories, and memos) were constantly compared and revised. This procedure was conducted to ensure that our preliminary and final results constituted a faithful representation of the data collected during Phases 1–3.
- iv. **Memoing:** Throughout our data analysis and synthesis steps extensive *memos*, i.e., textual notes, were recorded in order to keep track of research observation which fell outside our predefined coding process. This process, referred to as *memoing* [19], was used to annotate, among others, coding trends, emerging concepts, research actions to be taken (e.g., required homogenization of codes), and potential relations between incidents.

### 3.2.6. On the selection of the research process

For the interested reader, in this section we describe the rationale behind the adoption of our research process.

The very nature of this research is inductive, as it aims at understanding a phenomenon, namely the future of sustainable digital infrastructures, rather than testing a preexisting theory. Therefore, the data gathered from the participants can be seen as the fundamental “building blocks” which, throughout our three research phases,

contributed in a bottom-up incremental fashion to the results of this study.

Among the various qualitative research methods available to achieve our goal [21], few can be considered as potentially fit to answer our RQs. For example, grounded theory (GT) [19] might be deemed as a potential candidate research method, as it aims at building a theory on a phenomenon based on incremental findings. However, as presented in Section 3, our study aims to answer an *a priori* defined goal and RQs. Therefore, grounded theory could not be applied, as an *a priori* defined goal and RQs would violate numerous fundamental assumptions of GT, such as the unknown nature of the emerging theory, theoretical sampling, and theoretical sensitivity [19].

A method which could be deemed as a better fit to answer our RQs is case study research [22]. However, given the broad scope of our research goal, adopting such method would constitute a major threat to external validity of the study, making the results hard to generalize, and ultimately defying the very goal of our exploratory investigation. To mitigate threats to external validity, a *multi* case study design could be contemplated. However, given the heterogeneous nature of industrial sectors considered, which ranged from academic institutes to funding agencies and hardware manufacturers (see Fig. 2), the effort required to conduct such type of research would make the study unfeasible. In a similar vein we deemed other qualitative methods, such as narrative research, phenomenology, or ethnography [21], too narrow focused: while potentially providing more depth to the findings, they are not suitable to answer our RQs, which are broad in nature.

Given the above, for this study we chose to design an *ad hoc* 3-phase mixed-method inductive qualitative study, which allowed us to tackle the encompassing nature of our goal, while guaranteeing a satisfactory level of depth in the results. Specifically, in the first phase, exploratory interviews were conducted to delineate the general shape of our findings, in terms of relevant stakeholders, topics to be considered, and preliminary results. Subsequently, the preparatory results needed to be refined and consolidated. This task was achieved by leveraging focus groups, which allowed participants to jointly reason, discuss, and review findings which could have been hard to digest individually. Focus groups also allowed us to gain more depth in the findings, by empowering experts to make a “deep-dive” into the sustainability of their sector during each targeted focus group. Finally, as closure to our research process, we wanted to conduct an as encompassing as possible validation of our achieved results. Therefore, we re-invited all participants to take part to the evaluation and refinement session. Given the significant number of participants (30, see also Section 3.2.4), conducting exclusively a plenary research phase involving all participants at the same time would have been unfeasible. Therefore, we leveraged a set of separate working groups, where a smaller number of participants could provided in a structured and moderated environment feedback on the results we achieved.

## 4. Results

In this section, we report the result of our research. We open the section with a documentation of the landscape’s *key of reading*, i.e., a description of the landscape elements categorization, supporting the reader in the interpretation of the landscape (see Section 4.1). In Section 4.2 we detail the solutions for energy efficient digital infrastructures emerging from our study. In Section 4.3 instead we document the emerging adoption factors that would facilitate, or even accelerate, the adoption of the solutions. Section 4.4 reports the impediments elicited by participants, i.e., the factors that could hinder the adoption of solutions. Finally, the current open problems that need to be solved in order to advance and accelerate the sustainability of digital infrastructures are presented in Section 4.5.

### 4.1. Landscape elements categorization.

Throughout the Results section (Section 4), concepts such as solutions and open problems, are categorized *via* acronyms. The general structure of the acronyms is “[sustainability dimension]-[concept type]”. *Sustainability dimensions* can be technical (T), social (S), environmental (E), economic (Ec), paradigm shifts (PS), or general (G) when not tight to any particular one. *Concept types* can be solutions (S), adoption factors (AF), impediments (I), or open problems (OP). For example, “Domain-specific Hardware (T-S)” is a technological solution. It must be noted that our main focus is on novel technologies/solutions; as such, existing technologies (e.g., *software virtualization* like virtual machines and containerization) are left implicit even if they will undergo continuous improvement and optimization, over time, to contribute to energy efficiency.

The solutions are temporally ordered throughout the landscape horizons, showcasing the technology readiness in terms of widespread adoption and full impact, as perceived by the participants of this study.

### 4.2. Solutions for sustainable digital infrastructures

Fig. 5 gives an overview of the solutions described in this section. It is important to note that, while in Fig. 6 software-centric solutions are reported, it is in general hard to distinguish between hardware- and software-centric solutions for most concepts presented in this section (e.g., *non-Von Neumann architectures* are interlaced with *novel software architectures*). Hence, in the reminder of this section, solutions are presented without making a distinction between hardware- and software-centric ones for the sake of clarity.

#### 4.2.1. H1: Solutions for today

Solutions belonging to the H1 are characterized as being readily available for adoption. While the focus of this study is on horizons H2 and H3, we captured also the solutions in H1 that have been mentioned by the participants, or well-known solutions that are being adopted in H1 but are expected to reach further maturity or full potential in H2. Examples may include software virtualization solutions (e.g., virtualization and containerization also mentioned in the H1 report [23]) which are instrumental to maximize energy efficiency in the H1 scenario of cloud centralization (see Section 5.2.5.1), and implement more innovative scenarios like energy-driven dynamic consolidation, and the flexible geolocation highlighted in Section 5.2.5.2. Although well known, there is (still) significant room for optimization to maximize the energy saving potential [24].

**Moving to the Cloud (PS):** During H1 we see a first paradigm shift, already occurring in present time, namely *moving to the cloud*. This paradigm shift entails moving data, computational, and software capabilities from on premise to the cloud. In other words, rather than owning resources locally, resources are accessed on-demand, as provided by renowned cloud computing services (e.g., Amazon Web Services, Microsoft Azure, and the Google Cloud Platform). This paradigm shift influences the rise in popularity of software applications specifically designed to be deployed on the cloud, such as *cloud-native and serverless applications*. For instance, a study from Eclipse about cloud computing growth between 2008 and 2014 [25], found that 86% of companies were already using between 1 and 4 different types of cloud computing services, and predicted that 50% of all IT would be cloud-based between 2019 and 2025. Figures from 2020 report 90% of companies being cloud based [26], hence far exceeding the expectations. As another example, while cloud-stored data witnessed a yearly growth of 20%, about 99% is waste as hardly ever used [27].

**Heuristics for Hyperscale Hardware Management (T-S):** Moving to the cloud entails a growing centralization of software and hardware



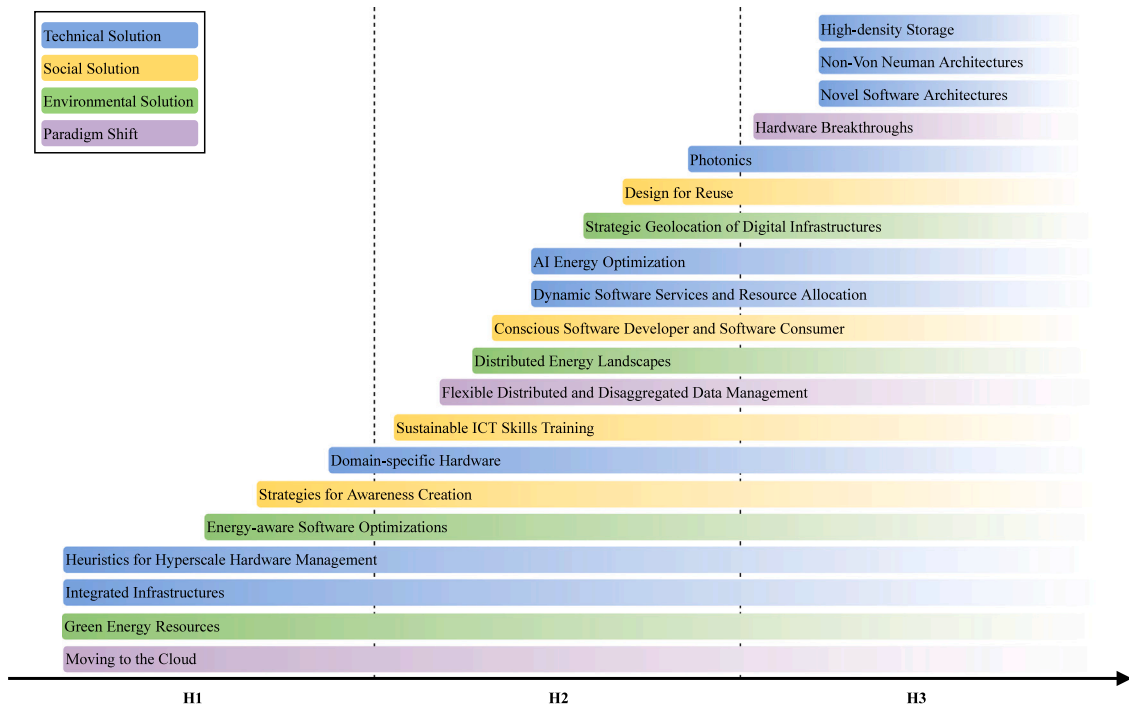


Fig. 5. Overview of the temporal distribution of solutions (H1–H3).<sup>4</sup>

resources in hyperscale digital infrastructures. Therefore solutions relative to this paradigm shift are prominently characterized by the energy optimization of this type of digital infrastructures, *i.e.*, *heuristics for hyperscale hardware management*. A frequent adopted solution regards *heat management*, such as *efficient cooling strategies* (*e.g.*, *immersion cooling*) and the *reuse of dissipated thermal heat*. In addition, energy consumption of hyperscale digital infrastructures can be lowered by adopting *energy-aware storage optimization*. This solution often entails moving the data that requires high transfer speeds to solid-state drive storages (SSDs), while archiving less-frequently accessed data via long-term backup storage solutions, *e.g.*, Amazon Glacier,<sup>5</sup> that are far less performant, but also more energy efficient.

**Green Energy Resources (E-S):** Another prominent category of solutions regards the adoption of *green energy resources*, *e.g.*, solar and wind farms, which often envisions the proximity of the future hyperscale digital infrastructures to green energy resources.

**Energy-aware Software Optimizations (E-S, H1):** *Energy-aware software optimizations* of the applications running on the digital infrastructures is another category of solutions that start to arise in H1. An overview of this type of solutions, distributed over the three horizons, is reported in Fig. 6. Related to this horizon, a cloud-centric specific solution is what is referred to as *kill zombie systems*, *i.e.*, the detection and shut down of idle servers to ensure that no energy is wasted to keep unused hardware resources running. In addition, the transition to the cloud encourages the *use on demand* of resources, enabled by *event-based software engineering*, and allowing to timely consume cloud resources only when certain triggers appear in the event stream. Related to the *moving to the cloud* paradigm shift, H1 sees the rise of *cloud-native* and *serverless applications*.

**Integrated Infrastructures (T-S):** Another solution, which spans also across H2 and H3, regards the creation of *integrated infrastructures*

born from the tight collaboration between software and hardware companies, that develop dedicated hardware built to satisfy the needs of software companies in a sustainable fashion.

**Domain-specific Hardware (T-S):** A solution specific to hardware regards instead the appearance during H1 of *domain-specific hardware*. Such hardware components are designed to efficiently solve specific problems, *e.g.*, the appearance of new graphics processing unit (GPU) types developed and optimized specifically for deep learning.

**Strategies for Awareness Creation (S-S):** A solution addressing the social dimension of the sector and much broader in scope, regards *strategies for awareness creation*, *e.g.*, monitoring and communicating about the environmental impact of data production, manipulation, and usage, hence striving towards a conscious use of energy, and a behavioral change in the data consumption patterns. *Design for reuse* implies a possible implicit tradeoff between energy consumption optimization and hardware waste, as utilizing longer lived technologies can imply a delay in using the newer more efficient technologies.

#### 4.2.2. H2: Solutions for the near future

**Flexible Distributed and Disaggregated Data Management (PS):** Occurring in H2, this paradigm shift foresees a transition from hyperscale data centers to *flexible distributed and disaggregated data management*. With the steady advancements in communication technologies, and the growing affordability of computational power, *edge computing* is expected to gain a widespread popularity in the near future. At the same time the “edge” will take different shapes from what we first thought, varying from static mini-clouds on premise to flexible “follow-the-need services”. This will imply to elastically move a vast number of computational tasks as near as possible to both the consumer premises and the increasingly decentralized energy production, *e.g.*, via *onboard computing*, and *distributed networks of computational nodes*.

**Strategic Geolocation of Digital Infrastructures (E-S):** Transitioning towards a flexible distributed and disaggregated data management will allow for the *strategic geolocation of digital infrastructures*. With this solution, digital infrastructures can be strategically positioned close to

<sup>4</sup> Solution specific to *energy-aware software optimizations* are depicted in Fig. 6.

<sup>5</sup> <https://aws.amazon.com/glacier/>.

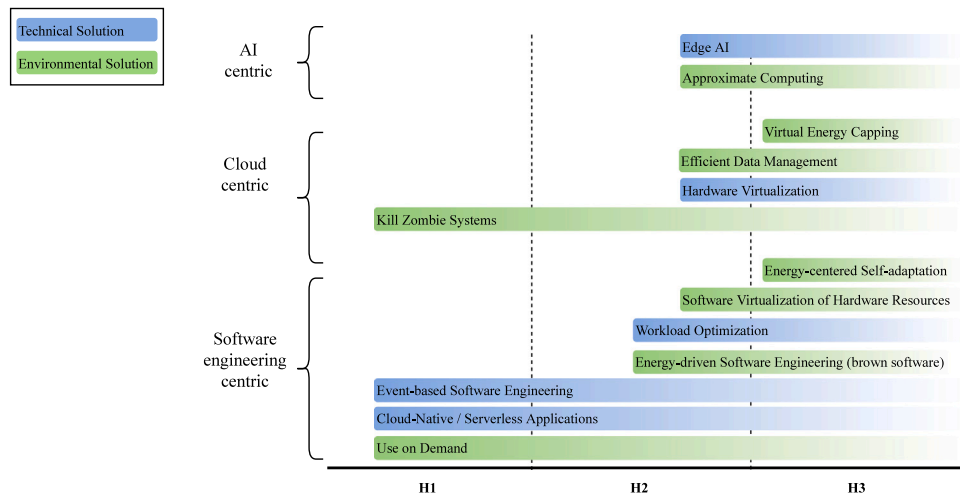


Fig. 6. Zoom-in into “Energy-Aware Software Optimization” solutions.

their end-users, in order to ensure high bandwidth, and keeping low the energy consumption of data flow and related communication.

**Dynamic Software Services and Resource Allocation (T-S):** In addition, distribution and disaggregation supports better profiling of energy consumption patterns, allowing for *dynamic software services and resource allocation*. At a coarser level of distribution and disaggregation, *mini-clouds* can be seen as intermediate steps supporting the transition towards a completely distributed paradigm. This solution enables to profile data usage patterns, and dynamically allocate services and resources by considering also the specific action performed on the data, e.g., data transport, storage, or manipulation. For example, the energy efficiency of data staging can be optimized via profiling by analyzing its frequency of use, and subsequently allocating the best fitted resources to reduce data traffic and improving performance at the same time.

**Distributed Energy Landscapes (E-S):** In addition to mitigating the energy waste of both, or either, hyperscale and co-location data centers (e.g., due to idle times or suboptimal virtualization practices), the shift towards a distributed paradigm enables also the use of *distributed energy landscapes*, supported by *smart energy grids*, to locally produce and consume energy, avoiding the inevitable energy dissipation characteristic of a centralized monolithic system.

**AI Energy Optimization (T-S/E-S):** Another characteristic aspect of H2 is the prominent role that artificial intelligence (AI) will play. As for other computational tasks, AI is expected in the near future to shift further towards distribution and disaggregation, enabled via (i) novel *federated learning algorithms*, supported by the appearance of *edge AI*, (ii) *data optimization/compression strategies* allowing to transfer high volumes of curated information rather than raw data, and (iii) *approximate computing*, i.e., the provisioning of results of acceptable quality, rather than optimal, in order to reduce energy consumption. As AI training/serving are known as particularly energy greedy computational tasks [28], future developments of AI require applying energy efficiency software engineering to AI-based systems, a field which is currently rapidly gaining traction [29]. In addition, promising results showcase how energy consumption of AI-based systems can be reduced by utilizing *data-centric green AI techniques* [30] and *AI-dedicated hardware components*, i.e., *AI on chip*.

**Energy-aware Software Optimizations (E-S, H2):** While in H1 *energy-aware software optimizations* apply software engineering practices for energy efficiency, in H2 it will be instrumental to create innovation as *energy efficient distributed software*, e.g., flexible distribution and disaggregation, smart virtualization, and sustainability-aware self-adaptation [31]. While current advancements towards stable and

reliable edge computing are promising, more significant research will be required to systematically shift towards a distributed adaptation paradigm. This will soon require to consolidate aspects such as *serverless architectures*, and *optimized service orchestration strategies*. In addition, the shift towards integrated infrastructures, spawn from tight collaborations between software and hardware manufacturers, will require further advancements in fields such as *infrastructure partitioning*, i.e., the partition of hyperscale data centers to optimize and sustain different tasks and workloads. Current trends predict the widespread popularization of innovative software optimizations, such as *fine-grained dynamic load balancing*, and *AI-enabled optimization of software energy consumption* (e.g., to manage virtualization and scheduling tasks). Other solutions specific to energy-aware software optimizations (cf. Fig. 6) regard (i) the *software virtualization of hardware resources*, virtualizing pools of hardware resources in order to ensure the seamless allocation and use of heterogeneous hardware components available on the cloud, (ii) *workload optimizations*, carried out to dynamically tune hardware and software resources to best fit the task at hand, and further advancements in the field of *energy-driven software engineering*, allowing to refactor software applications to make them more energy efficient, while maintaining unvaried their delivered functionality.

**Sustainable ICT Skills Training (S-S):** H2 is characterized by the popularization of sustainable ICT skills training, carried out both in academic and industrial settings. Such educational paths are necessary to systematically create profiles able to reason about and address the energy sustainability of digital infrastructures in all types of industrial sectors.

**Conscious Software Developers and Consumers (S-S):** In addition to educational training of sustainable ICT skills, advancements in energy efficiency measurement and monitoring allow for the rise during H2 of *conscious software developers and software consumers*, i.e., people who develop a renewed sense of responsibility regarding the sustainability of the ICT solutions they implement and use. The rise of responsible software developers and software consumers on the one hand leads to more sustainable software solutions “by design”, and on the other hand increases the (quality of) sustainability requirements of digital infrastructures demanded by consumers.

**Design for Reuse (S-S):** H1 is characterized by an average lifecycle expectancy of digital infrastructure hardware components equal to approximately two years. Differently, H2 is expected to witness a growing trend of *design for reuse* and *hardware lifecycle management* practices. The adoption of these circular economy strategies allows to produce longer lasting and maintainable hardware components, amortizing the financial cost and environmental impact of hardware

production over a longer time span. Obviously, this corresponds to a trade-off between energy-efficiency and use of critical material which is not a one-size-fits-all solution.

#### 4.2.3. H3: Solutions further away

**Hardware Breakthroughs (PS):** H3 regards solutions that will appear in the longer term, *i.e.*, beyond 6 years. From this study, H3 appears to mainly consider novel *hardware breakthroughs*. Such hardware breakthroughs, described in the reminder of this section, are expected to drastically change how portions of digital infrastructures are designed and operate. It is important to bear in mind that in the foreseeable future, such solutions will not substitute the hardware technologies currently adopted in digital infrastructures, but will rather co-exist with them.

**Photonics (T-S):** Prominently, H3 is marked by a widespread use of *photonics*. The demand of low latency communication will be steadily growing while the shift towards distributed paradigms, starting in H2, will gain momentum. While photonics is already in use during H1, it is expected to become much faster and ubiquitous during H2, and will be of widespread and consolidated use in H3. In addition to inter-server/client communication (*copackaged optics*), which is already showcasing promising results, the widespread transition towards optic communication is expected to occur also *within* data centers, with the replacement of micro-electronic hardware with optics-based one by adopting *integrated photonics*.

**Non-Von Neumann Architectures (T-S):** In parallel, current advancements in *non-von Neumann* research showcase how, during H3, complex computational tasks will be executed at a fraction of the energy consumed by the hardware of today. Prominently, the evolution of *neuromorphic computing* and similar solutions can lead to groundbreaking energy savings required to carry out computational-intensive tasks. This notable change of hardware technologies can lead to drastic changes in the underlying hardware structures of data centers. In addition, current progress in *high-density storage solutions* showcase promising digital infrastructure storage optimizations, *e.g.*, by making use of electron spins or relations between protons and electrons, and enabling to store high volumes of data at a negligible energy cost. A consideration can also be made regarding *quantum computing*: the successful implementation of quantum computers can find applications in a dedicated class of computing, fundamentally out of reach for conventional computers, *e.g.*, quantum physics, molecular chemistry of logistics. Nevertheless, such computational tasks will with high probability fall outside the domain of conventional digital infrastructures. In addition, the low temperature at which quantum computers may have to operate ( $-272^{\circ}\text{C}$ ), can pose a serious concern regarding their energy efficiency.

**Novel Software Architectures (T-S):** As hardware is expected to present some considerable breakthroughs in H3, software will be needed to evolve and adapt to the new underlying hardware. This will require the creation of *novel software architectures*, in order to evolve software systems to best fit the drastic technology changes implied by the hardware breakthroughs of H3.

#### 4.3. Adoption factors

In this section, we present the elicited key adoption factors that, if present, would facilitate and even accelerate the adoption of the solutions presented in Section 4.2. An overview of the adoption factors (as well as impediments and open problems described further on) are depicted in Fig. 7.

**Technology Readiness (T-AF):** The most-frequently mentioned adoption factor is *technology readiness*, especially relevant for H3. The technology readiness results to be a key adoption factor, as developing and on-boarding a preliminary solution with not clearly understood

benefits and drawbacks implies a great risk. While solutions presented in H1 are production ready, the technology readiness of solutions appearing in H2 and H3 is hard to define. While the positioning of solutions throughout the landscape depicted in Fig. 5 orders the solutions temporally, it is important to bear in mind that such positioning is not definitive, and may change as future developments of the solutions take place.

**Ease of Integration (T-AF):** Another prominent adoption factor, specific to the software solutions, entails the *ease of integration*, allowing to integrate solutions without any drastic change in the normal functioning of data centers. This results to be another key adoption factor, as the normal functioning of digital infrastructures cannot be interrupted while integrating a new solution, and the return of investment of adopting a certain solutions cannot be hindered by the cost of integrating the solution.

**Digitization and Digitalization (T-AF):** The adoption of solutions also highly depends on the future maturity of *digitization* (*i.e.*, the conversion of information into digital form) and *digitalization* (*i.e.*, the adoption of digital technologies in business processes). Digitization and digitalization processes can either pave the way, or inhibit, the development of sustainable digital infrastructures. This depends on the progress that digitization and digitalization will make in the future, and the extent to which their advancements will consider sustainability aspects.

**Support for Trade-off Decision Making (S-AF):** General to all solutions is a clear understanding of potential tradeoffs and hence the *support for trade-off decision making*, as energy savings should not deteriorate the quality of provided services. This entails also a systematic analysis of implementation and deployment costs involved, in order to understand the economic implications of proposed solutions, *i.e.*, other business cases.

**Holistic Paradigm Shift (Ec-AF):** Specific to the movement towards a distributed paradigm is instead the requirement of a holistic paradigm shift, allowing to distribute the cost of research and development across a wide range of stakeholders, instead of burdening with the implied risk a single party. Additionally, the widespread paradigm shift allows stakeholders ensuring that the undertaken change will be used and supported by other parties, hence mitigating the potential risk of developing silos technologies, *i.e.*, technologies that are hard to interface with others present on the market.

#### 4.4. Impediments

In this section, we report the four key impediments, identified with our research, which can hinder the adoption of sustainable digital infrastructure solutions.

**Unclear Impact (G-I):** Related to the key adoption factors, are a list of impediments, which might hinder the adoption of the solutions reported in Section 4.2. The first and most important impediment is the unclear impact of the solutions on service provision quality (*e.g.*, performance), energy savings, and evolution of technology ecosystems. This impediment is related to the *technology readiness* adoption factor, and can be mitigated only by conducting systematic experimentation to evaluate/measure the impact of the landscape solutions.

**Adversity to Change (G-I):** The *unclear impact* impediment can lead to *adversity to change* of certain parties (*e.g.*, telecommunication and cloud providers), as reshaping currently consolidated technologies may lead to uncharted situations, that have to be clearly analyzed and understood before undertaking major investments.

**Lack of Leading Champions (G-I):** The study uncovered a relation between the impediment *resilience to change* and the *lack of leading*

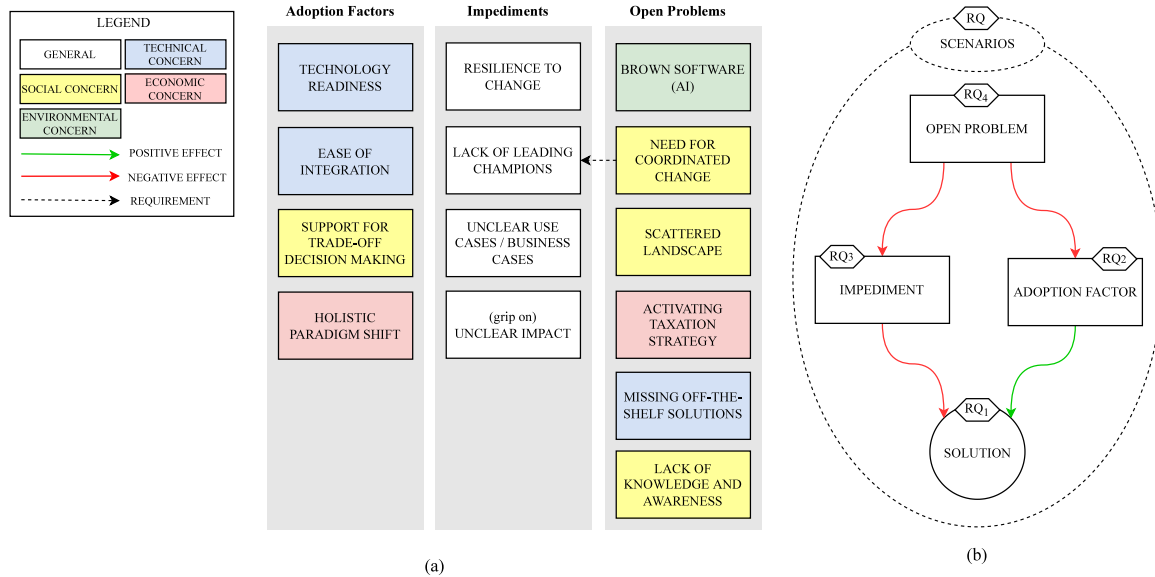


Fig. 7. (a) Summary of adoption factors, impediments, and open problems. (b) Network of dependencies between concepts/RQs.

*champions*, i.e., leading figures in organizations, or influential organizations themselves, that take initiative and steer the change towards the next generation of sustainable technologies. This impediment is further discussed in the related open problem *lack of guidance* reported in Section 4.5.

**Unclear Use Cases and Business Cases (G-I):** An impediment characteristic to research oriented solutions (e.g., quantum computing) is the yet *unclear use cases and business cases*, which have to be understood before they can be successfully adopted in industrial contexts.

#### 4.5. Open problems

In this section, we report the encompassing open problems of the next generation of energy efficient digital infrastructure solutions.

**Need for a Coordinated Change & Scattered Landscape (S-OP):** Related to the *lack of leading champions* is the perceived *need for a coordinated change*, enabling stakeholders to jointly progress, while sharing costs/risks involved, and avoid a *scattered landscape*, characterized by compartmentalized technology silos adopted only by few companies. In addition, real progress needs shared responsibility and shared accountability throughout the whole value chain, which makes all parties responsible for their actions towards the sustainability of digital infrastructures, and empowers them.

**Lack of Activating Taxation Strategy (Ec-OP):** A key driver towards a communal paradigm shift, and a current open legislation problem related to energy usage, is the *lack of activating taxation strategy* that should both activate all stakeholders, and target other factors than just electricity or carbon emissions. Energy bills are among the highest costs of data centers, nevertheless the current taxation strategies do not drive disruptive changes in energy consumption patterns of data centers<sup>6</sup>. In a foreseeable future, smart taxation strategies may be formulated, such as higher taxation of fossil energy sources, and dynamic pricing based

on real-time energy demands. Similarly, e-waste production is currently only marginally supervised, but more stringent legislation could lead in the future to the popularization of “design for reuse” and other optimizations of data center hardware life cycle strategies.

**Lack of Policies and KPIs (Ec-OP):** An open problem related to missing activating taxation strategies is the *lack of policies and KPIs*.<sup>7</sup> Specifically, this open problems regards the current lack of regulations, standards, and policies regarding sustainability requirements of digital infrastructures. The complexity of this problem is also due to the current absence of KPIs, such as sustainability labels for software and hardware components, that can track and guide the progress of stakeholders developing sustainable digital infrastructures. The introduction of KPIs throughout value chains can also support an enhanced monitoring, regulation, and resolution of sustainability concerns of complex ICT business cases. The *lack of policies* should be addressed at both the national- and the international level, as numerous hardware and software companies are nowadays characterized by an international nature, both in terms of multi-national companies, and companies operating at a global scale.

**Lack of Guidance (S-OP):** Another open problem of social nature is the perceived *lack of guidance*, i.e., guidance supporting companies in becoming more environmental sustainable, and supporting them in the systematic adoption of energy efficient solutions. This guidance can from either from a governmental institution, a research consortium, or even a private company, that champions and supervises the common endeavor of parties towards more sustainable digital infrastructures.

**Change of Mindset/Sense of Urgency (S-OP):** An open problem that can inhibit a shift towards energy efficient solutions is the need of a *change of mindset*, in order to give a higher priority to sustainability of digital infrastructures, which often is neglected in favor of other goals such as business targets and customer satisfaction.

**Brown Software (T-OP):** Another open problem regards applying energy-efficiency practices to the software for and in data centers which is energy inefficient, so-called *brown software*. This problem is becoming more prominent due to the ever-growing adoption of AI, which is at the moment characterized by severe software energy inefficiencies. As

<sup>6</sup> A quite successful activation strategy of the Dutch Government concluded in 2020, was the MJA (*Meerjarenspraken*, or in English multi-annual agreements), a nationwide initiative intended to improve energy efficiency of ICT products, services and processes by granting tax exemptions when the efficiency targets were met. Results and details are online at [www.rvo.nl/onderwerpen/duurzaam-ondernemen/energie-besparen/mja3-mee](http://www.rvo.nl/onderwerpen/duurzaam-ondernemen/energie-besparen/mja3-mee). Some of the resulting practices are reported in [32].

<sup>7</sup> Key Performance Indicators.



AI-based systems become more and more adopted, this open problem may lead to serious energy supply shortcomings if left unaddressed.

**Missing off-the-shelf-solutions (T-OP):** As a general trend, the current stall to move towards the next generation of energy efficient digital infrastructures, implies a current perception of *missing off-the-shelf-solutions* to optimize the energy consumption of data centers. This translates in the current need of creating or adapting *ad hoc* solutions, instead of having the possibility to efficiently and effectively applying solutions that are available.

**Energy Availability:** During H1, we observe a reliable and satisfactory energy supply. Nevertheless current energy consumption trends display the rapid transitions towards electric of applications that used to run on fossil fuel. A prominent example is transportation, e.g., electric cars. To support this transition, it is estimated that the electric grid in the Netherlands will need to at least double in capacity during H2. This combined with the abandonment of brown energy resources, poses two urgent and still open energy-related problems: (i) new challenges in the capacity planning of the energy infrastructure, and (ii) the impending scarce supply of renewable energy, due to the usage from other sectors and the need for space and materials. In addition, this transition will also lead to the inclusion of new notable consumers of electric energy, such as the chemical industry and hydrogen producers, who will require the implementation of dedicated energy lines and energy buffers. Nevertheless transitioning towards electric energy entails some slow processes (from both a technical and regulatory point of view), e.g., installing new cables and transformers. This slow installation processes may require years in order to set up a large power connection in a geographical region, putting practical limits on the growth of certain areas. In addition, as digital infrastructures are moving towards the adoption of green energy resources, they will start to compete with industries of other nature, and even the public sector. This poses both a practical and political problem, that is currently still open.

**Lack of Knowledge and Awareness (S-OP):** Related to this problem, industrial contexts often suffer of a lack of knowledge and awareness regarding what ICT sustainability really means, which could be mitigated with dedicated education and training programs. This is also reflected in a current *lack of skills* regarding energy efficiency of digital infrastructures, that often leads to a lack of knowledge on how to address sustainability in practice, and possibly the adoption of suboptimal solutions.

In the following, we report two open problems that, given their different and encompassing nature, are reported separately.

**Inherent Time Required for a Paradigm Shift (G-OP):** A general problem regarding any new technology, which is present also in the context of energy efficient digital infrastructure solutions, is the *inherent time required for a paradigm shift*. From historical data, we know that technology leaps often take between 10–20 years to gain traction, as time is required in order to clearly understand its pros and cons of the technology, and to be adopted by a wider audience. While the information and communication technology domain is characterized by an extremely fast innovation cycle, such consideration holds to a large extent also for this domain.

**Culmination of Micro-electronics Computational Advancements (G-OP):** The last open problem, laying at the root of the growing concerns of data center energy consumption, is the *culmination of micro-electronics computational advancements*, notably displayed by the culmination of Moore's Law [2]. As reducing further power consumption of micro-electronics is becoming physically impossible, there is no degree of freedom left to further optimize the energy consumption of such technology. Hence, in order to improve the energy efficiency of digital infrastructures, two different possibilities are available, paving the landscape for H2 and H3, namely (i) transforming the current data processing paradigms and data volumes, and (ii) systematically transitioning towards the next generation of hardware technologies.

## 5. Discussion

In this section, we report an in-depth discussion of the results of our research. Specifically, in Section 5.1, we outline the general purpose we envision for the landscape which emerged from our results. In Section 5.2 instead, we revisit our research questions, and explicitly answer them by considering the results we obtained. Finally, in Section 5.3, we draw further considerations on our results and the future of sustainable digital infrastructures.

### 5.1. A landscape open for inspiration

In Section 4.2 we describe the solutions we uncovered in our investigation, as well as the time horizons in which our participants placed them. This suggests a strategy about when each concept is expected to be available on the market provided we invest in research and development to create them—hence a landscape. However, this landscape can (and should) be read in various ways depending on the reader's perspective. In general, it should be an inspiration to think out-of-the-box about solutions and innovations needed to develop a future-proof and energy efficient digital infrastructure. In particular it can act as a strategy to achieve a goal over time, e.g., to position technological solutions over time, so that they incrementally build upon one another (see the sketch in Fig. 8.(a)); or to incrementally realize target scenarios that help creating socio-technical solutions which contribute to a systemic way of thinking (see Fig. 8.(b)). The aim of the landscape is to inspire to take action. Researchers, practitioners, funding agencies, and policy makers (among others) should assess which innovations to support and take forward in their next research and development steps, to further accelerate the transition towards sustainable digital infrastructures development and maintenance.

### 5.2. Research questions revisited

In this section, we revisit our research questions and answer them based on the results we obtained with our investigation. Specifically, we start by answering our sub-research questions (RQ<sub>1</sub>–RQ<sub>4</sub>), and leverage such answers to subsequently draw our main conclusions, i.e., answer our main research question (RQ).

#### 5.2.1. Answer to RQ<sub>1</sub> (solutions)

With RQ<sub>1</sub>, which states “What are the solutions to develop sustainable digital infrastructures?”, we aimed at identifying solutions which can be utilized in order to create and maintain sustainable digital infrastructures. In Section 4.2, we documented the set of solutions we identified through our research (see also Fig. 5). Identified solutions belong either to the technical, social, or environmental sustainability dimension. By considering the solutions we identified, we can draw the following observations: (i) future solutions will not only be technical, but also of social, and environmental nature, (ii) with the passing of time, the sustainability of digital infrastructures will be supported by an increasing distribution of computational and energy resources, facilitating presumption, and dynamic allocation of assets, (iii) solution will reflect this increasing distribution, with a shift of focus from hyperscale/hardware optimizations towards distributed/software-related solutions in the future, (iii) paradigm shifts, i.e., systematic changes of the digital infrastructure domain, will deeply affect the development of sustainable infrastructures, starting from an initial shift to the cloud, to flexible distributed and disaggregated data management, and finally revolutionary hardware breakthroughs, (iv) energy-aware software optimizations will play a key role in achieving the sustainability of digital infrastructures.

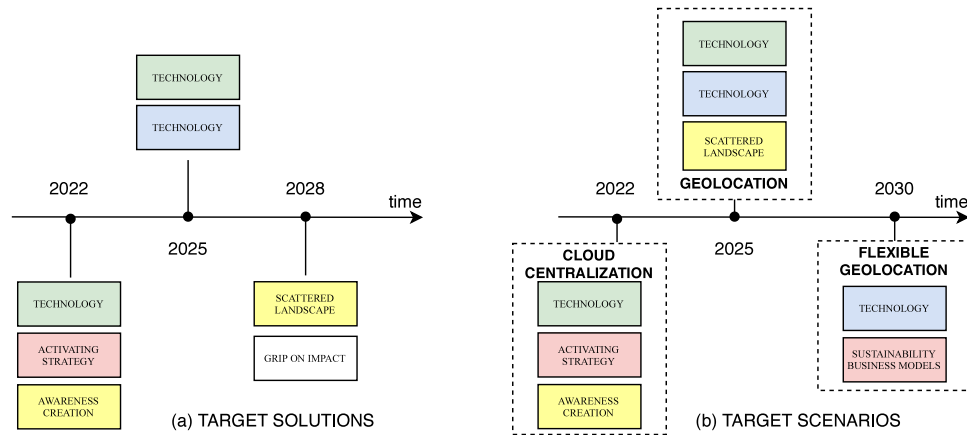


Fig. 8. Landscape illustrations (a) per target solutions and (b) per target scenarios.

**Summary of findings RQ<sub>1</sub> (Solutions):** We identified 30 solutions for sustainable development of digital infrastructure of heterogeneous nature, ranging from technical (14 solutions), to social (4 solutions) and environmental ones (12 solutions). Solutions are expected to focus increasingly on distribution of resources, be supported by numerous energy-aware software optimizations (13 solutions), and culminate in radical hardware breakthroughs in the long-term future. In addition we identified 3 paradigm shifts, which will guide the present and future of sustainable digital infrastructures.

#### 5.2.2. Answer to RQ<sub>2</sub> (adoption factors)

With our second sub-RQ, namely “What drives the adoption of sustainable digital infrastructure solutions?”, we aimed at identifying which factors stimulate the adoption of sustainable digital infrastructure solutions. In Section 4.3 we reported the five key adoption factors identified in our research, three of which related to the technical sustainability dimension, one the social dimension, and one to the economic dimension. Regarding technical-related adoption factors, participants highlighted the characteristics solutions would be required to have, in terms of *technology readiness*, *ease of integration* within existing digital infrastructures, and the support of future *digitization and digitalization* advancements, which will either enforce or inhibit the adoption of solutions. In addition, participants expressed the need of *support for trade-off decision making*, required to assess, reason about, and discuss how adopting a certain sustainability solutions would impact other digital infrastructure properties. Finally, from an economic point of view, the requirement of an *holistic paradigm shift* emerged, i.e., the need to distribute sustainability research and development efforts across multiple stakeholders, in order to mitigate the cost related to the adoption of solutions.

**Summary of findings RQ<sub>2</sub> (Adoption Factors):** Five key adoption factors emerged from our study, namely *technology readiness*, *ease of integration*, *digitization and digitalization*, *support for trade-off decision making*, and *holistic paradigm shift*. The identified adoption factor express key properties of the solutions, and the context they are deployed in, required for the adoption of solutions.

#### 5.2.3. Answer to RQ<sub>3</sub> (impediments)

Our third sub-RQ states “What hinders the adoption of sustainable digital infrastructure solutions?”. This research question was designed to understand what factors discourage or impede the adoption of

sustainable digital infrastructure solutions. In Section 4.4, we reported the 4 key impediments identified with our research. Such impediments, rather than being related to a specific sustainability dimension (e.g., technical or economic), are all of generic nature. One of the most prominent impediments, related to the emphasis of this study on the future, is the *unclear impact* of the solutions on sustainability. As stated by the participants, before adopting a solution, solutions have to be proven as effective, and allow for transparent trade-off analysis with respect to other properties of digital infrastructures (e.g., performance and maintainability). This can be deemed to constitute a vicious cycle, where the impact of solutions has to be proven before deployed, but only by deploying solutions *in vivo* the exact impact of solutions can be studied. As a related impediment, participants noted the *lack of a leading champion*, i.e., entities who take it upon themselves to develop, experiment, research, and steer sustainable digital infrastructure advancements. Similarly, the unclear impact of solutions and the *unclear use cases and business cases* lead to an *adversity to change*, that may influence digital infrastructure stakeholders in adopting well-known yet unsustainable solutions, instead of sustainable ones that may showcase unknown properties.

**Summary of findings RQ<sub>3</sub> (Impediments):** We identified 4 key impediments for the adoption of sustainable digital infrastructure solutions, namely *unclear impact*, *unclear use cases and business cases*, *adversity to change*, and *lack of a leading champions*. All impediments exhibit different facets of the uncertainty bound to future solutions, and call for *joint* effort in progressing the field of sustainable digital infrastructures.

#### 5.2.4. Answer to RQ<sub>4</sub> (open problems)

Our last sub-RQ states “What are the open problems related to sustainable digital infrastructure solutions?”, and was formulated to identify the open problems and challenges related to sustainable digital infrastructures solutions. In Section 4.5, we identified a total of 13 open problems, 2 of which of economic nature, 5 social, 3 technical, 2 environmental, and 1 general. From an economic point of view, open problems highlight a current lack of legislative overview, standards, and administration of digital infrastructure sustainability, expressed in terms of a *lack of policies*, *KPI*, and *activating taxation strategies*. From a social perspective, on one hand open problems demonstrate the need of joint action to progress the sustainability of digital infrastructures discussed also in Section 5.2.3, and is embodied in open problems such as the presence of a *scattered landscape*, the *lack of guidance*, and the *need for a coordinated change*. On the other hand, social problems reflect the need of a social change in digital infrastructure providers and consumers, embodied in our results as the open problems of *lack*




Fig. 9. Scenarios of sustainable digital infrastructures across temporal horizons.

of knowledge and awareness, and change of mindset/sense of urgency. Technical open problems instead showcase the technical need for novel solutions (*culmination of micro-electronics advancements*), the perceived immaturity of the field (*missing off-the-shelf solutions*), and the current lack of technical sustainability skills in the job market (*lack of skills*). The environmental problems identified in our research regard the current neglect of *brown software* issues, and problematics related to the future *energy availability* and supply requirements. The last open problem, of general nature, regards the *inherent time required for a paradigm shift*, i.e., the time needed by a technology to gain traction and be adopted by a wide audience.

**Summary of findings RQ<sub>4</sub> (Open Problems):** In total, we identified 13 open problems related to sustainable digital infrastructure solutions, ranging from economic problems (2 open problems), to social (5 open problems), technical (3 open problems), environmental (2 open problems), and general ones (1 open problem). Economic problems display the lack of legislation and standards for sustainable digital infrastructures, while social problems highlight the need for a coordinated change, and changes in the mindset of digital infrastructure providers/consumers. Technical, environmental, and general open problems result to be of more heterogeneous nature, displaying different facets of the challenges related to sustainable digital infrastructures ahead.

#### 5.2.5. Answer to main RQ (the future of sustainable digital infrastructures)

In order to answer our main research question, we analyzed the results of our study (see Section 4), and derived 4 incremental scenarios which address the sustainability of digital infrastructures. The scenarios, ordered temporally from the one that is currently taking place, to the one that can only be achieved in a long time horizon, constitute the answer to our main research question, namely: “What is the future of sustainable digital infrastructures?”. As the scenarios are incremental, they can be ordered temporally throughout the three horizons of the landscape. An overview of such ordering is provided in Fig. 9.

The documentation of each scenario, mapped to the involved stakeholders (see also Table 3), is provided below.

**5.2.5.1. Scenario 1 (S1): Cloud centralization.** This first scenario entails the migration of software and hardware resources from on-premise to a centralized remote cloud. This scenario is connected to the first paradigm shift (*moving to the cloud*) presented in Section 4.2.1, that occurs during H1. An overview of this scenario is depicted in Fig. 10.

Intuitively, this scenario strives towards more energy efficient digital infrastructures by delegating sustainability concerns to prominent cloud providers (e.g., Amazon, Microsoft, and Google). The sustainability of this solution is based on the assumption that hyperscale digital infrastructures deploy already some energy efficient solutions,

Table 3

Stakeholders involved in the sustainable digital infrastructure scenarios.

Stakeholder	Scenario			
	S1	S2	S3	S4
Cloud Provider	✓	✓	✓	✓
Costumers and Consumers	✓	✓	✓	✓
Hardware Producers	✓	✓	✓	✓
Telecommunication Providers	✓	✓	✓	✓
Governments	✓	✓	✓	✓
Prosumers		✓	✓	✓
Smart Energy Grid Providers		✓	✓	✓
Municipalities		✓	✓	✓
“Seamless Continuum” supervisors			✓	✓

e.g., *green energy resources* and *energy-aware software optimizations* (see Section 4.2.1), that might be currently hard to be achieved by customers concerned with managing only a fraction of their software and hardware resources. In addition, the sustainability of this scenario is supported by a high level of *use on demand*, as the consumption of cloud resources is directly associated to a economic expense by the consumers.

As presented in Section 4.2.1, this scenario has to be regarded only as a temporary scenario adopted as a transition towards more energy efficient scenarios (presented in the reminder of this section). In fact, solutions adopted by hyperscale digital infrastructures, e.g., *hyperscale hardware management*, should be regarded as heuristics to mitigate the environmental impact of digital infrastructures, rather than making them sustainable in the long term. In addition, geographical space limitations, and the current trends of green energy resources development, pose serious concerns regarding the growth of hyperscale digital infrastructures, as their centralized paradigm can constitute a challenge for grid operators to ensure that the required infrastructure facilitates all consumers in any specific geographical area. In addition, the increasing user mobility (reflected in the pervasive use of mobile devices) is turning remote-cloud data storage and -traffic into important bottlenecks [33].

**Stakeholders of Scenario 1.** The most prominent stakeholders involved in this scenario are:

- *Cloud providers*, who manage their hardware and software resources, and make them available to customers in form of services;
- *Customers and Consumers*, who make use of cloud services, and migrate to various extents their hardware and software capabilities to the cloud; they may include cloud-based applications accessed by both office-workers and home-workers;
- *Hardware Producers*, who supply hardware components to cloud providers, or support them in implementing their own hardware solutions (e.g., via *integrated infrastructures*);




Fig. 10. Cloud centralization (Scenario 1).




Fig. 11. Flexible geolocation (Scenario 2).




Fig. 12. Seamless continuum (Scenario 3).

- *Telecommunication providers*, who provide communications services to connect the different entities of the scenario;
- *Governments*, who supervise the energy taxation and regulation of customers and cloud providers.

**5.2.5.2. Scenario 2 (S2): Flexible geolocation.** The scenario is characterized by a hybrid nature, in which remote clouds and micro-clouds coexist. An overview of this scenario is depicted in Fig. 11.

The “flexible geolocation” scenario is supported by the paradigm shift occurring in H2 *flexible distributed and disaggregated data management* (see Section 4.2.2). Specifically, as *edge computing*, *distributed energy landscapes*, and *dynamic software services and resource allocation* gain traction, it is possible to exploit for the sake of sustainability the energy- and computational resources available at the edge of ICT networks. This allows to distribute both the computational and energy consumption load between different geographical areas, hence mitigating the energy consumption centralized in specific geographical areas characteristic of Scenario 1. In addition, this scenario enables the appearance of *hardware, software, and energy prosumers*, i.e., consumers of hardware or software resources that can not only make use of their local or personal hardware capabilities when possible, but can also make use of local energy smart grids in order to consume energy in more self-sustainable fashion. Flexible geolocation brings a systemic vision of the energy sector and the ICT sector working together, hence enabling both novel economies of scale and stability in terms of energy needs and quality of service, for both sectors. For example, the geolocation of heat production from e.g., data centers and heat consumption from e.g., greenhouses creates mutually-beneficial ecosystems (see also Fig. 13). **Stakeholders of Scenario 2.** In addition the stakeholders of the previous scenario (Scenario 1, see also Table 4), the following new stakeholders emerge in this scenario:

- *Prosumers*, who make use of their hardware, software, and energy resources;
- *Smart Energy Grid Providers*, who supervise and manage access to smart grid services;
- *Municipalities*, who make urban decisions about e.g., spatial planning for data centers, and the support of urban solutions influencing the production and consumption of energy.

This scenario also uncovers the need for centralized government (e.g., ministries) and decentralized government (e.g., municipalities) to synchronize their decisions, and strategies.

**5.2.5.3. Scenario 3 (S3): Seamless continuum.** This scenario is supported, among others, by advancements in *hardware virtualization*, *workload optimization* and *dynamic software services and resource allocation* that take place in H2 (see Section 4.2.2). This scenario is characterized by a pool of shared hardware and software resources, constituted by the resources made available by both micro-clouds and remote clouds. An overview of this scenario is depicted in Fig. 12.

The shared pool of resources constitutes in this scenario a “seamless continuum”, i.e., hardware and software resources are allocated at runtime, in order to select from the resource available from the pool the ones which are best fitted to provide a certain service. This scenario is supported by advancements in communication technologies, both from a software and hardware perspective, allowing higher transmission speeds, and automatically avoiding congestions by profiling and rerouting data traffic as needed. For example, by considering the AI domain, the computational intensive training of an AI model can be delegated to AI accelerators available on the remote cloud, while the subsequent classification based on the trained model can be executed by a device as close as possible to the end-user. This scenario allows to progress towards the sustainability of digital infrastructures by seamlessly selecting the hardware and software resources most fitted to the task at hand, while leveraging the hybrid nature presented in Scenario 2.

**Stakeholders of Scenario 3.** In addition the stakeholders of the previous scenario (Scenario 2, see also Table 4), the following new stakeholders emerge in this scenario:

- “*Seamless Continuum*” *supervisors*, who supervise the distribution of resources available in the pool. This could be a new type of *aggregator*.

**5.2.5.4. Scenario 4 (S4): Follow time, space, and energy.** This scenario builds upon the previous ones, with specific emphasis on the dynamic allocation of resources characteristic of the *seamless continuum* (Scenario 3). More specifically, differently from Scenario 3, in this scenario resources are allocated based on both their software and hardware capabilities, and on the availability of the energy the resources need, the proximity of resources, and the timeliness of the task at hand. An overview of this scenario is depicted in Fig. 13, where any digital infrastructure resource can be used seamlessly, hence conceptually linking anything and everything (illustrated by the gray circle in the figure).

As an example for Scenario 4, we can consider a computational intensive task, requiring eventual consistency, that has to be carried out daily. As the execution of such task can be postponed during the day, it is possible to allocate the task to the high computational






Fig. 13. Follow time, space, and energy (Scenario 4).

power available in the remote cloud, and executing it when the energy demand of the remote cloud is low (e.g., at nighttime). If instead the load of the remote cloud never presents a decrease, it is possible to allocate the task to a network of edge computing nodes available in the proximity of the end-user, in order to carry out the task throughout the day by making use of locally available energy.

In this last scenario, the sustainability of digital infrastructures is achieved by making use of information regarding the task at hand, the energy availability, and proximity of resources, in order to achieve a sustainable service provisioning without any apparent degradation of its quality aspects. In addition, this scenario enables the dynamic prioritization of energy resources used based on energy re-use, presumption, and overall sustainability of the resources.

Interestingly, in line with our results, the Gartner report on the data centers of the future foresees that “by 2025, 85% of infrastructure strategies will integrate on-premises, colocation, cloud and edge delivery options, compared with 20% in 2020” [34].

**Stakeholders of Scenario 4** This scenario involves all stakeholders involved in Scenario 3 (see also Table 3).

**Summary of findings main RQ (The future of sustainable digital infrastructures).** Digital infrastructures are currently experiencing a shift towards cloud centralization, temporarily improving their sustainability by delegating sustainability concerns to cloud providers (Scenario 1). Sustainability limitations of hyperscale digital infrastructures will subsequently be addressed by distributing both the computational and energy consumption load between different geographical areas, via the introduction of *flexible distributed and disaggregated data management* (Scenario 2). As time progresses, a *seamless continuum*, constituted by a pool of shared resources, will allow to conveniently select the hardware and software resources best fitted to achieve the task at hand (Scenario 3). Finally, digital infrastructures will evolve to *follow time, space, and energy*, allowing to fulfill tasks by allocating resources based on software and hardware requirements, energy availability, proximity, and timeliness of tasks (Scenario 4).

### 5.3. Further considerations

Due to the continuous increase of data and digital services, the energy required to operate digital infrastructures is steadily increasing, so much so that it is rapidly reaching its feasibility limits unless new technologies are applied. Open problems, impediments, adoption factors, and solutions create a network of dependencies (see Fig. 7.(b)<sup>8</sup>) framing the concerns relevant for deciding on the future sustainability of digital infrastructures, and the related evolution scenarios. To mitigate environmental repercussions, the current adoption of green energy resources, and hardware/software optimizations for hyperscale infrastructures has to be intended exclusively as part of the solution, as it will not be able to scale with the ever growing data consumption demands. In the near future, this can be mitigated by shifting towards a more distributed architectural paradigm, bringing data and its processing closer to the consumer premises, with a mix of strategic data center geolocation and edge computing. Distribution and disaggregation will dynamically promote energy consumption patterns based on renewable energy supply, flexibility of services, and smart transfer and use of information, hence mitigating the environmental impact characteristic to modern hyperscale data centers. This will allow to progress in a sustainable fashion, till novel hardware breakthroughs will occur, setting new standards of low energy data storage, communication, and processing. Rather than seeing digital infrastructures completely replaced due to the introduction of new technology, the most likely progress will entail building heterogeneous digital infrastructures, where old and new technologies co-exist to provide a seamless service provision, but with a sustainable mindset.

To progress towards a energy-aware future of digital infrastructures, it is paramount that people gain awareness of the sustainability of the digital services they develop and use. This allows to hold every party present throughout value chains accountable for the sustainability of their actions, potentially transitioning towards *designing for less*, i.e., ensuring that only what is really needed is produced and consumed.

In addition, the urgency to tackle the sustainability of digital infrastructures requires to promptly activate all stakeholders involved. This is reflected in the current need to revise and innovate the *modus operandi* of funding agencies, as only timely interventions can resolve the current (un-)sustainability trends of digital infrastructures. The adaptation of funding schemes requires also the strategic focus on significant national-wide sectors in the Netherlands, such as the flower industry, and the digital infrastructure industry itself.

To build the sustainable future of digital infrastructures that might be, however, all stakeholders must act together: cloud providers, cloud customers, technology providers, consumers, government, and researchers - *we are all decision makers*. Further, we *must* take into account, both qualitatively and quantitatively, possible *rebound effects* [36,37] of optimized features of current and future infrastructures, e.g., the more data/processing speed is made available, the more data is being consumed, which in turn causes a further need for speed.

## 6. Threats to validity

In this section we discuss the threats to validity of our study, by following the categorization provided by Wholin et al. [38].

### 6.1. Internal validity

As any qualitative research investigation, the subjective influence of researchers on the data collection, analysis, and synthesis may have influenced our results. In order to mitigate this threat, interviews were conducted or reviewed by at least two researchers, and the coding process, albeit being conducted by a single researcher, was supervised and

<sup>8</sup> The diagrams use the notation defined in [35]

scrutinized independently by two additional researchers. In addition, to mitigate internal validity threats related to the data analysis/synthesis processes (see Section 3.2.5), the final results were presented back to the participants, in order to ensure that the data they provided had been correctly interpreted, processed, and faithfully reported.

## 6.2. External validity

The most prominent threat to external validity of this study lies in the focus on the data provided by the 48 participants who took part in this study. In order to mitigate this threat, *via ad hoc* quotas we ensured that our sampling considered participants of heterogeneous sectors (see Fig. 2 and Table 4). As a further threat to external validity, the topic considered in this research focuses on the future evolution of digital infrastructures, as envisioned by our participants. Hence, our results do not have to be interpreted as absolutely certain and deterministic, but rather as the probable evolution of digital infrastructures, as described by our participants and their understanding of the state-of-the-art. In addition, as any quantitative study, our results constitute a mid-range substantive theory, *i.e.*, our results do not have to be interpreted as absolute or final, and can be refined and extended by follow-up or replication studies.

## 6.3. Construct validity

An inherent threat to construct validity of our investigation may be caused by the research process followed to answer our research questions. In order to reduce this threat, we adopted different mitigation strategies. First, the research process we followed was carefully designed and finalized *a priori*, in order to systematically execute our research plan without altering it to accommodate unexpected incoming data, impediments, etc. In addition, while the data analysis and synthesis process was the same across our 3 research phases, we adopted a mixed-method empirical study (see Section 3 and Fig. 1), in order to mitigate potential threats to validity bound to a specific research method. Finally, potential threats to validity related to our coding strategy (see Section 3.2.5) were mitigated by involving 3 independent researchers in the process, who scrutinized, discussed, and reviewed the coding process and related results.

## 6.4. Conclusion validity

In order to mitigate potential threats to conclusion validity which may have effected our results, we leveraged the participants of our study in order to assess and validate our findings. Specifically, after each of our 3 research phases, the obtained intermediate/final results (*i.e.*, the preliminary findings report, the preliminary landscape report, and the final landscape report, see Section 3 and Fig. 1) were presented back to the participants, who had the opportunity to scrutinize how we analyzed and synthesized their input, and provide feedback on our conclusions. In addition, our intermediate results were presented not only to the participants who provided the data, but also to new participants, *i.e.*, the ones who participated to the subsequent research phases, and were never exposed to our result before (see Table 4 for an overview of new / “recurring” participants). This allowed us to further mitigate threats to conclusion validity, by having independent participants, who were not involved in the study yet, assess our findings.

## 7. Related work

In this section, we review the related literature which we deem the closest to our study in terms of treated research topic, and discuss the main commonalities and differences. Specifically, given the vast body of knowledge available on the topic of ICT sustainability [5], we concentrate our discussion on the related work studying the topic of

future evolution of sustainable digital infrastructures, or the solutions available to support such evolution.

As the impact of digital infrastructures on sustainability became clearer, during the years several studies, position papers, and roadmaps on the topic emerged. This led to the emergence of manifestos, both academic and industrial, outlining the future of sustainable digital infrastructures, and the actions to be taken to sustainably evolve them. In “A Manifesto for Energy-Aware Software” [39], Fonseca et al. describe nine principles of energy awareness, which can guide the way towards establishing sustainable digital infrastructures. By comparing the manifesto to our study, we note that numerous reported principles emerge also in our investigation (*e.g.*, as the solutions *Sustainable ICT Skills Training* and *Strategies for Awareness Creation*), albeit our research considers a much vaster range of concepts (*e.g.*, adoption factors, impediments, *etc.*) and reports energy-aware software optimizations (see Fig. 6) only as minor portion of the total solutions identified.

In an earlier work produced by IBM [40] in 2012, key elements of green digital infrastructures, and a set of energy efficiency techniques for digital infrastructure sustainability are presented. Given the time elapsed from the publication of such work, it comes to no surprise that the paper covers exclusively solutions reported in H1 of our landscape, and does not consider a setting other than a centralized one (Scenario 1, see Section 5.2.5.1). Furthermore, as a difference w.r.t. our study, we further characterize the sustainability of digital infrastructures by focusing not only on available solutions, but also on future ones, with the addition of concepts characterizing the phenomenon (*e.g.*, adoption factors and impediments).

A study which instead focuses also on future temporal horizons is the taxonomy of sustainable cloud computing presented by Gill et al. [41]. In such survey, the authors present a taxonomy of solutions for the sustainable evolution of digital infrastructure. When compared to our work, we note that the study by Gill et al. focuses exclusively on technical solutions, with particular emphasis on virtual machines (*e.g.*, virtual machine load balancing and consolidation), and heuristics for hyperscale hardware management (*e.g.*, thermal-aware scheduling and cooling management). The landscape presented in our study instead considers solutions belonging also to other sustainability dimensions, such as the environmental and social one. In addition, our study aims at a broader understanding of sustainable digital infrastructures, by considering not only solutions but also impediments, adoption factors, and scenarios. Interestingly, similarly to our investigation results, Gill et al. elicit a set of open challenges, often posed as open questions. By comparing these challenges with the open problems of our landscape, we note that the ones presented by Gill et al. lay at a completely different level of abstraction, and result to possess a much finer-grained level of focus (*e.g.*, “How does size of the cloud datacenter affect its energy efficiency?”) w.r.t. the encompassing nature of the problems emerging from our investigation.

The work of Varghese et al. [42] presents an outline of the future generation of digital infrastructures. In line with the results of our study, Varghese et al. note an increasing distribution of computational and storage resources in the future. Albeit presenting sustainability as a key to architecting future digital infrastructures, the study of Varghese et al. only marginally touches upon sustainability aspects of digital infrastructures, by discussing the energy-aware positioning of data centers, and the importance of considering energy efficiency as a quality-of-service metric. In contrast, the primary focus of our investigation is on the sustainability of digital infrastructures, and hence reports a much deeper level of detail and results on the topic.

Harmon et al. [43], in a study dating back to 2012, defined a roadmap for sustainable IT for the upcoming decade (*i.e.*, 2012–2022). The roadmap reports market segments, technologies, compliance requirements, organizational changes, and value migration of the future sustainable IT. The authors identify two “waves” of sustainable IT. The first wave, occurring in the present to near-term future, is characterized

by “green datacenter products and technologies” (e.g., power and workload management, cloud computing, and server design). The second wave instead, occurring between the near-term and intermediate-term future, is characterized by “sustainable IT services and technologies”. By comparing our landscape to the roadmap of Harmon et al. we observe that the first wave described by Harmon et al. is still ongoing, showcasing the slow advancement of sustainable digital infrastructures in practice. In addition, as major differences w.r.t. Harmon et al. [43], our study focuses also on other factors, e.g., open problems and impediments, and considers solutions and their interconnection at a refined level of granularity.

In the work of Uddin et al. [44], a framework leveraging different energy optimization solutions (e.g., virtualization, cloud computing, and green metrics) to improve the sustainability of digital infrastructures is presented. As for the aforementioned paper by IBM [40], also this study focuses exclusively on solutions appearing in the temporal horizon H1 (see Section 4.2.1), and does not consider other concepts emerging in our landscape, e.g., adoption factors, impediments, and open problems.

Finally, there are various studies that compared to ours, tackle only specific aspects within digital infrastructure sustainability. These studies range from the assessment of digital infrastructure sustainability [45–47], to monitoring and visualizing sustainability metrics [48, 49], supporting decisions to improve the sustainability of digital infrastructures [50,51], adopting sustainable networking solutions [52], and utilizing software engineering techniques to improve energy efficiency [53,54].

## 8. Conclusions and future work

Digital infrastructures are experiencing an ever-growing demand for data consumption, which is only expected to increase in the future with the popularization of high bandwidth connections and affordable data plans. This trend opens for a new question that we, as digital infrastructure providers, researchers, and consumers, will have to answer in the near future: *How can we evolve digital infrastructures to keep up with the increasing data demand in a sustainable fashion?* In order to answer this question, in this study we present a mixed-method qualitative empirical research, conducted by involving 48 unique practitioners belonging to heterogeneous digital infrastructure sectors. Our findings shed light on the solutions which are or will be available in the future to sustainability evolve digital infrastructures, the adoption factors that will guide their deployment, the impediments related to their usage, and the open problems that need to be tackled. In addition, from our investigation emerge four subsequent scenarios for the future of sustainable digital infrastructures, starting from an initial shift from on premise to the cloud, and culminating in a distributed scenario, where digital infrastructures will follow time, space, and energy to dynamically allocate resources. While technological, environmental, and economical solutions will support this change, another factor emerged as crucial to achieve the sustainability of digital infrastructures. The shift towards sustainability of digital infrastructures cannot be achieved by a limited number of stakeholders alone. Only by gaining a wider awareness of the sustainability of digital infrastructures, hold every party present throughout value chains accountable for the sustainability of their actions, and establish cross-domain collaborations, we will be able to sustainably evolve digital infrastructure while keeping up with their growing demands.

As documented in Section 2, in this research we studied the future evolution of sustainable digital infrastructures primarily from a technical viewpoint. We hope that our results can inspire and be used as handle by researchers from other (cross-)domains, e.g., socio-technical [8,9] or socio-economic domains [10,11], to further enrich our findings.

Based on this finding, as future work we are currently establishing an interdisciplinary collaboration involving heterogeneous parties,

from hardware manufacturers, to cloud providers, software service providers, and research/academic institutes. The goal of our research plan is to implement, apply, and evaluate the sustainable digital infrastructure solutions reported in this study, with keen interest in the shift towards the distributed paradigm, which is at the basis of the incremental scenarios for sustainable digital infrastructures documented in this research.

## CRedit authorship contribution statement

**Roberto Verdecchia:** Conceptualization, Methodology, Investigation, Data curation, Writing – original draft, Visualization. **Patricia Lago:** Conceptualization, Methodology, Investigation, Writing – review & editing, Supervision. **Carol de Vries:** Investigation, Writing – review & editing.

## Declaration of competing interest

One or more of the authors of this paper have disclosed potential or pertinent conflicts of interest, which may include receipt of payment, either direct or indirect, institutional support, or association with an entity in the biomedical field which may be perceived to have potential conflict of interest with this work. For full disclosure statements refer to <https://doi.org/10.1016/j.suscom.2022.100767>. Roberto Verdecchia reports financial support was provided by Netherlands Enterprise Agency. Patricia Lago reports financial support was provided by Netherlands Enterprise Agency. Carol de Vries reports financial support was provided by Netherlands Enterprise Agency.

## Acknowledgments

This research received funding from the Netherlands Enterprise Agency Project “Energy Efficient Digital Infrastructures” (project number RVO-TSE2200010) and support from the LEAP Initiative of the Amsterdam Economic Board. Our sincere gratitude goes to the 48 participants who took part to this study, for their time, invaluable insights, passion, and support. Types of companies who participated in the study were hardware manufacturers, cloud providers, software service providers, consultancy firms, academic- and research institutes, funding agencies, and digital infrastructure customers. In addition to the participants that chose to remain anonymous, we would like to thank (in alphabetical order): Joris van den Aker (TNO), Orhan Alici (Senior Solutions Architect at Redhat), Nicola Calabretta (Assistant Professor in Electro-Optical Communication Systems at TU Eindhoven), Jeroen Cox (Strategic Lead Energy & Environment at KPN), Aaron Ding (TU Delft), Sagar Dolas (Adviser at SURF Innovation Lab), Niels Hensen (ITB2 Datacenters), Hans Hilgenkamp (University of Twente), Robbert Hoeffnagel (Adviser at SDIA/Green IT Amsterdam), Judith Inberg (University of Twente), Jos Keurentjes (University of Twente), Sjaak Laan (Director Consulting Expert at CGI), Jan-Willem Lammers (Principal Solutions Architect at VMware), Johan Mentink (Assistant Professor of Physics at Radboud University), Erik Negenman (Application Expert Low Voltage at Global Operations, Schneider Electric), Job Oostveen (TNO), Theo Rasing (Professor of Physics at Radboud University), Patty Stabile (Associate Professor on Neuromorphic Photonics at TU Eindhoven), Jeroen van der Tang (Public Policy Manager Duurzaamheid at NL Digital), Loek Wilden (CDCAP Team leader Digital Services & Execution at Schneider Electric).

## Appendix

See Table 4.

**Table 4**

Overview of Study participants, mapped to company, sector, and research phase they took part in.

Participant ID	Company ID	Company sector	Research phase		
			Phase 1	Phase 2	Phase 3
P1	C1	Special interest group	✓		
P2	C1	Special interest group	✓		
P3	C2	Cloud Provider, Software Service Provider, Hardware Manufacturer	✓		
P4	C2	Cloud Provider, Software Service Provider, Hardware Manufacturer	✓		
P5	C2	Cloud Provider, Software Service Provider, Hardware Manufacturer	✓		✓
P6	C3	Academic Institute	✓		
P7	C3	Academic Institute	✓		✓
P8	C4	Academic Institute	✓		✓
P9	C4	Academic Institute	✓		✓
P10	C4	Academic Institute	✓		
P11	C5	Cloud Provider, Software Service Provider, Hardware Manufacturer	✓		
P12	C5	Cloud Provider		✓	
P13	C6	Consultancy Firm		✓	
P14	C7	Research Institute		✓	
P15	C8	Cloud Provider, Software Service Provider, Hardware Manufacturer		✓	
P16	C9	Software Service Provider		✓	
P17	C10	Cloud Provider		✓	
P18	C11	Cloud Provider		✓	
P19	C12	Consultancy Firm		✓	
P20	C13	Funding Agency		✓	
P21	C14	Academic Institute		✓	
P22	C14	Academic Institute		✓	
P23	C15	Academic Institute		✓	✓
P24	C16	Software Service Provider		✓	✓
P25	C17	Cloud Provider		✓	✓
P26	C18	Cloud Provider		✓	✓
P27	C19	Cloud Provider, Software Service Provider, Hardware Manufacturer		✓	✓
P28	C20	Software Service Provider, Consultancy Firm		✓	✓
P29	C21	Cloud provider, Software Service Provider		✓	✓
P30	C22	Consultancy Firm		✓	✓
P31	C1	Academic Institute			✓
P32	C23	Hardware Manufacturer			✓
P33	C23	Hardware Manufacturer			✓
P34	C24	Telecommunication Service Provider			✓
P35	C24	Telecommunication Service Provider			✓
P36	C25	Special interest group			✓
P37	C26	Research Institute			✓
P38	C26	Research Institute			✓
P39	C27	Academic Institute			✓
P40	C27	Academic Institute			✓
P41	C28	Consultancy Firm			✓
P42	C29	Cloud Provider, Software Service Provider, Hardware Manufacturer			✓
P43	C29	Cloud Provider, Software Service Provider, Hardware Manufacturer			✓
P44	C30	Research Institute			✓
P45	C31	Cloud Provider, Software Service Provider, Hardware Manufacturer			✓
P46	C32	Special interest group			✓
P47	C33	Special interest group, Funding Agency			✓
P48	C33	Special interest group, Funding Agency			✓

## References

- [1] S. Podder, A. Burden, S. Kumar Singh, R. Maruca, How green is your software?, 2020, URL: <https://hbr.org/2020/09/how-green-is-your-software>. (Accessed: 11 November 2021).
- [2] G.E. Moore, Cramming more components onto integrated circuits, *Proc. IEEE* 86 (1) (1998) 82–85.
- [3] R.H. Dennard, F.H. Gaensslen, H.-N. Yu, V.L. Rideout, E. Bassous, A.R. LeBlanc, Design of ion-implanted MOSFET's with very small physical dimensions, *IEEE J. Solid-State Circuits* 9 (5) (1974) 256–268.
- [4] R. Verdecchia, P. Lago, C. Ebert, C. De Vries, Green IT and green software, *IEEE Softw.* 38 (6) (2021) 7–15.
- [5] R. Verdecchia, F. Ricchiuti, A. Hankel, P. Lago, G. Procaccianti, Green ICT research and challenges, in: *Advances and New Trends in Environmental Informatics*, Springer, 2017, pp. 37–48.
- [6] P. Lago, S.A. Koçak, I. Crnkovic, B. Penzenstadler, Framing sustainability as a property of software quality, *Commun. ACM* 58 (10) (2015) 70–78.
- [7] N. Condori-Fernández, P. Lago, Characterizing the contribution of quality requirements to software sustainability, *J. Syst. Softw.* 137 (2018) 289–305.
- [8] O. Henfridsson, B. Bygstad, The generative mechanisms of digital infrastructure evolution, *MIS Q.* (2013) 907–931.
- [9] D. Tilson, K. Lyytinen, C. Sørensen, Digital infrastructures: The missing IS research agenda. Research commentary, *Inf. Syst. Res.* 21 (4) (2010) 748–759.
- [10] L. Locatelli, M. Guerrero, B. Russo, E. Martínez-Gomariz, D. Sunyer, M. Martínez, Socio-economic assessment of green infrastructure for climate change adaptation in the context of urban drainage planning, *Sustainability* 12 (9) (2020) 3792.
- [11] A.K. Sarangi, R.P. Pradhan, ICT infrastructure and economic growth: A critical assessment and some policy implications, *Decision* 47 (4) (2020) 363–383.
- [12] E. Brynjolfsson, A. Saunders, *Wired for Innovation: How Information Technology Is Reshaping the Economy*, Mit Press, 2009.
- [13] J.W. Creswell, *Qualitative, Quantitative and Mixed Methods Approaches*, Sage, 2014, pp. 43–45.



- [14] V.R. Basili, G. Caldiera, D. Rombach, The goal question metric approach, in: *Encyclopedia of Software Engineering*, Wiley, 1994, pp. 528–532.
- [15] J. Kontio, J. Bragge, L. Lehtola, The focus group method as an empirical tool in software engineering, in: *Guide to Advanced Empirical Software Engineering*, Springer, 2008, pp. 93–116, <http://dx.doi.org/10.1007/978-1-84800-044-5>.
- [16] V. Vehovar, V. Toepoel, S. Steinmetz, Non-probability sampling, in: *The Sage Handbook of Survey Methods*, Sage Thousand Oaks, 2016, pp. 329–345.
- [17] P. Biernacki, D. Waldorf, Snowball sampling: Problems and techniques of chain referral sampling, *Sociol. Methods Res.* 10 (2) (1981) 141–163.
- [18] D. Oliver, J. Serovich, T. Mason, Constraints and opportunities with interview transcription: Towards reflection in qualitative research, Oxford University Press, 2005, p. 1273, <http://dx.doi.org/10.1353/sof.2006.0023>.
- [19] B.G. Glaser, Strauss, Discovery of grounded theory: Strategies for qualitative research, 1967, p. 39, AldineTransactions.
- [20] J. Saldaña, The coding manual for qualitative researchers, Sage, 2015, pp. 291–298.
- [21] J.W. Creswell, C.N. Poth, Qualitative Inquiry and Research Design: Choosing Among Five Approaches, Sage publications, 2016.
- [22] R.K. Yin, Applications of Case Study Research, sage, 2011.
- [23] D. Harryvan, Savings Options in Datacenters and Server Rooms in 2021 and 2025 (in Dutch), Technical Report, Certios BV, 2021.
- [24] A. Barillas, C. Miller, S. Ramesh, Digital Transformation and a Net Zero Emissions Europe, Technical Report, Aurora Energy Research, 2021.
- [25] J. Desjardins, C. Ang, M. Lu, Cloud computing growth - Visual capitalist, 2014, URL: [www.visualcapitalist.com/cloud-computing-growth](http://www.visualcapitalist.com/cloud-computing-growth), [www.visualcapitalist.com/cloud-computing-growth](http://www.visualcapitalist.com/cloud-computing-growth).
- [26] N. Galov, 25 Cloud computing statistics in 2020 - will AWS domination continue?, 2020, URL: <https://hostingtribuna.com/blog/cloud-computing-statistics/>. (Accessed 1 May 2021).
- [27] G. Cook, T. Dowdall, D. Pomerantz, Y. Wang, Clicking Clean: A Guide to Building the Green Internet, Greenpeace, 2015.
- [28] K. Saenko, Why AI is so power-hungry, 2020, Ars Technica, URL: <https://arstechnica.com/science/2020/12/why-ai-is-so-power-hungry>.
- [29] R. Schwartz, J. Dodge, N.A. Smith, O. Etzioni, Green AI, *Commun. ACM* 63 (12) (2020) 54–63, <http://dx.doi.org/10.1145/338183>.
- [30] R. Verdecchia, L. Cruz, J. Sallou, M. Lin, J. Wickenden, E. Hotellier, Data-Centric Green AI: An Exploratory Empirical Study, *International Conference on ICT for Sustainability (ICT4S)* (2022).
- [31] I. Gerostathopoulos, C. Raibulet, P. Lago, Expressing the adaptation intent as a sustainability goal, in: *IEEE/ACM 44th International Conference on Software Engineering (ICSE)*, IEEE, 2022.
- [32] Q. Gu, P. Lago, H. Muccini, S. Potenza, A categorization of green practices used by dutch data centers, *Procedia Comput. Sci.* 19 (2013) 770–776, URL: <http://www.sciencedirect.com/science/article/pii/S1877050913007096>.
- [33] D. Vellante, A new era of innovation: Moore's law is not dead and AI is ready to explode, 2021, SiliconANGLE, URL: <https://siliconangle.com/2021/04/10/new-era-innovation-moores-law-not-dead-ai-ready-explode/>.
- [34] H. Cecci, D. Cappuccio, Your Data Center May Not Be Dead, but It's Morphing, Technical Report, Gartner, 2020.
- [35] P. Lago, Architecture design decision maps for software sustainability, in: *IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Society, ICSE-SEIS*, 2019, pp. 61–64, URL: <http://dx.doi.org/10.1109/ICSE-SEIS.2019.00015>.
- [36] C. Gossart, Rebound effects and ICT: A review of the literature, in: *ICT Innovations for Sustainability*, Springer International Publishing, 2015, pp. 435–448, URL: [http://dx.doi.org/10.1007/978-3-319-09228-2\\_6](http://dx.doi.org/10.1007/978-3-319-09228-2_6).
- [37] K. Widdicks, T. Ringenson, D. Pargman, V. Kuppasamy, P. Lago, Undesigning the internet: An exploratory study of reducing everyday internet connectivity, in: *International Conference on ICT for Sustainability, ICT4S*, in: *ICT4S*, 2018, pp. 384–397, URL: <https://easychair.org/publications/paper/VjQ1>.
- [38] C. Wohlin, P. Runeson, M. Höst, M.C. Ohlsson, B. Regnell, A. Wesslén, *Experimentation in Software Engineering*, Springer Science & Business Media, 2012.
- [39] A. Fonseca, R. Kazman, P. Lago, A manifesto for energy-aware software, *IEEE Softw.* 36 (6) (2019) 79–82.
- [40] M. Ceron, A green IT approach to data center efficiency, 2012, IBM Redbooks. IBM Academy of Technology, URL: <http://www.redbooks.ibm.com/redpapers/pdfs/redp4946.pdf>. (Accessed 1 May 2021).
- [41] S.S. Gill, R. Buyya, A taxonomy and future directions for sustainable cloud computing: 360 degree view, *ACM Comput. Surv.* 51 (5) (2018) 1–33.
- [42] B. Varghese, R. Buyya, Next generation cloud computing: New trends and research directions, *Future Gener. Comput. Syst.* 79 (2018) 849–861.
- [43] R.R. Harmon, H. Demirkan, D. Raffo, Roadmapping the next wave of sustainable IT, *Foresight* (2012).
- [44] M. Uddin, A.A. Rahman, Energy efficiency and low carbon enabler green IT framework for data centers considering green metrics, *Renew. Sustain. Energy Rev.* 16 (6) (2012) 4078–4094.
- [45] D. Gmach, Y. Chen, A. Shah, J. Rolia, C. Bash, T. Christian, R. Sharma, Profiling sustainability of data centers, in: *Proceedings of the 2010 IEEE International Symposium on Sustainable Systems and Technology*, IEEE, 2010, pp. 1–6.
- [46] M. Marwah, P. Maciel, A. Shah, R. Sharma, T. Christian, V. Almeida, C. Araújo, E. Souza, G. Callou, B. Silva, et al., Quantifying the sustainability impact of data center availability, *ACM SIGMETRICS Perform. Eval. Rev.* 37 (4) (2010) 64–68.
- [47] H.I. Moud, J. Hariharan, H. Hakim, C. Kibert, et al., Sustainability assessment of data centers beyond LEED, in: *2020 IEEE Green Technologies Conference, GreenTech*, IEEE, 2020, pp. 62–64.
- [48] C. Bash, T. Cader, Y. Chen, D. Gmach, R. Kaufman, D. Milojicic, A. Shah, P. Sharma, Cloud sustainability dashboard, dynamically assessing sustainability of data centers and clouds, CiteSeer, 2011.
- [49] M.A. White, Sustainability: I know it when I see it, *Ecol. Econom.* 86 (2013) 213–217.
- [50] G. Lewis, P. Lago, Architectural tactics for cyber-foraging: Results of a systematic literature review, *J. Syst. Softw.* 107 (Sep) (2015) 158–186, URL: <https://www.sciencedirect.com/science/article/pii/S0164121215001211>.
- [51] M. Pawlish, A.S. Varde, S. Robila, Decision support in data centers for sustainability, in: *2013 IEEE 13th International Conference on Data Mining Workshops, IEEE*, 2013, pp. 613–620.
- [52] F.A. Moghaddam, P. Lago, P. Grosso, Energy-efficient networking solutions in cloud-based environments: A systematic literature review, *ACM Comput. Surv.* 47 (4) (2015) 1–32.
- [53] R. Verdecchia, R.A. Saez, G. Procaccianti, P. Lago, Empirical evaluation of the energy impact of refactoring code smells, in: *ICT4S*, 2018, pp. 365–383.
- [54] R. Verdecchia, G. Procaccianti, I. Malavolta, P. Lago, J. Koedijk, Estimating energy impact of software releases and deployment strategies: The KPMG case study, in: *2017 ACM/IEEE International Symposium on Empirical Software Engineering and Measurement, ESEM, IEEE*, 2017, pp. 257–266.



**Roberto Verdecchia** is a Postdoctoral Research Associate at the Software and Sustainability group (S2) of the Vrije Universiteit Amsterdam, the Netherlands. He holds a double Ph.D. in Computer Science, appointed by the Gran Sasso Science Institute, L'Aquila, Italy, and the Vrije Universiteit Amsterdam, the Netherlands. His research interests focus on the adoption of empirical methods to improve software development and system evolution, with particular interest in the fields of technical debt, software architecture, software energy efficiency, and software testing. More information is available at [robertoverdecchia.github.io](http://robertoverdecchia.github.io).



**Patricia Lago** is Full Professor in software engineering at the Vrije Universiteit Amsterdam, the Netherlands, where she leads the Software and Sustainability (S2) research group in the Computer Science Department. Her research is in software architecture and software quality with a special emphasis on sustainability. She has a PhD in Control and Computer Engineering from Politecnico di Torino and a Master in Computer Science from the University of Pisa, both in Italy. She is initiator of the Computer Science Master Track in Software Engineering and Green IT, and co-founder of the Green Lab, a place where researchers, students and companies collaborate to measure the energy footprint of software solutions and the impact on software quality. She is a member of VERSEN and the Steering Committees of IEEE ICSA and ECSCA conference series. She is also Steering Committee Chair of ICT4S. More information is available at <http://patricialago.nl>.



**Carol de Vries** is program and technology manager at Photondelta in Eindhoven, the Netherlands. His focus is on integrated photonics technologies and applications and creating more business opportunities in this emerging field. He has a long history in the electronics industry working in R&D and new business management in the fields of semiconductors, sensors and integrated systems in different markets where has worked a.o. for Philips, NXP and Sensata. He has a MSc in Physics from the Technical University of Eindhoven.