

University of Groningen

Sustainability in Software Architecture

Andrikopoulos, Vasilios; Boza, Rares-Dorian; Perales, Carlos; Lago, Patricia

Published in:
2022 48th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)

DOI:
[10.1109/SEAA56994.2022.00073](https://doi.org/10.1109/SEAA56994.2022.00073)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2022

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Andrikopoulos, V., Boza, R.-D., Perales, C., & Lago, P. (2022). Sustainability in Software Architecture: A Systematic Mapping Study. In *2022 48th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)* (pp. 426-433). IEEE. <https://doi.org/10.1109/SEAA56994.2022.00073>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Sustainability in Software Architecture: A Systematic Mapping Study

Vasilios Andrikopoulos
University of Groningen
Groningen, the Netherlands
v.andrikopoulos@rug.nl

Rares-Dorian Boza
University of Groningen
Groningen, the Netherlands
r.d.boza@student.rug.nl

Carlos Perales
Vrije Universiteit Amsterdam
Amsterdam, the Netherlands
c.peraleslinan@student.vu.nl

Patricia Lago
Vrije Universiteit Amsterdam
Amsterdam, the Netherlands
p.lago@vu.nl

Abstract—Sustainability is an increasingly-studied topic in software engineering in general, and in software architecture in particular. There are already a number of secondary studies addressing sustainability in software engineering, but no such study focusing explicitly on software architecture. This work aims to fill this gap by conducting a systematic mapping study on the intersection between sustainability and software architecture research with the intention of (i) reflecting on the current state of the art, and (ii) identifying the needs for further research. Our results show that, overall, existing works have focused disproportionately on specific aspects of sustainability, and in particular on the most technical and “inward facing” ones. This comes at the expense of the holistic perspective required to address a multi-faceted concern such as sustainability. Furthermore, more reflection-oriented research works, and better coverage of the activities in the architecting life cycle are required to further the maturity of the area. Based on our findings we then propose a research agenda for sustainability-aware software architecture.

Index Terms—sustainability, software architecture, systematic mapping study, sustainability-aware architecture

I. INTRODUCTION

The developments of the last decades have generated a raising awareness of sustainability in an increasingly digital world run by software systems [1], [2]. With software systems being deeply embedded in most sectors of our society, a unique opportunity is being offered for a shared intervention across sectors towards a more sustainable world [3]. Reflecting this need, in the recent years sustainability has been acknowledged as an essential software quality across four dimensions [4]: a *technical* one referring to the ability of a software system to evolve and remain used over a long period of time, an *economic* one concerned with the preservation and creation of capital and value, an *environmental* one aiming to minimize the impact of the system on natural resources, and a *social* one focusing on the continuity of communities using it.

The state of the art in sustainable software engineering in general has evolved significantly in the last couple of decades, starting with a focus on ‘green’ software, that is, focusing on the environmental dimension, e.g. [5], [6], [7], [8], but expanding also to the other ones, e.g. [9], [10], [11]. A recent survey of sustainable software research based on the 5Ws formula (why, when, who, where, and what) [12], shows that this is a very active and collaborative area of research, with a good level of maturity. The research community, or more precisely its intersection with that of software engineering, is

also mapped out in terms of major outlets and collaborative efforts by [13] and [14], indicating the existence of active research groups in the area.

However, the software engineering-focused secondary studies such as the ones mentioned above lack by definition the sufficient level of depth and detail when discussing sustainability in software architecture. Focusing specifically on architecture is essential since architecting practices allow us to reason and evaluate sustainability as a system quality throughout the system’s lifecycle, from the elicitation of requirements all the way to managing the maintenance and evolution of the system. This is especially important since there is still a lot of ground to be covered with respect to establishing and exploiting the relation between software architecture and sustainability as a software quality [15].

In comparison to the number of studies on sustainability in software engineering, there are fewer works describing the state of the art in the intersection between sustainability and software architecture, and the secondary studies that do exist on the topic have their own limitations. The survey by Venters et al. [15], for example, provides a good overview of this area but it is not conducted systematically. Some surveys focus on technical sustainability, e.g. the one by Koziolok on sustainability evaluation metrics for software architectures [16], or the ones on the sustainability of reference architectures [17] and their description [18]. Verdecchia et al. [19] discuss technical in combination with economic sustainability in the context of architectural technical debt, while Volpato et al. [20] and Grua et al. [21] look into the implications of software architecture on the social dimension. No systematic survey so far, to the extent of our knowledge, attempts to *organize the state of the art of sustainability in software architecture across all sustainability dimensions and architecting activities*. This is a need that this work aims to address by mapping out the research efforts in the field and thus identifying areas of future research. For this purpose, in this paper we report on our *systematic mapping study* [22], [23] on this topic.

As a consequence, the rest of this paper is organized as follows. Section II discusses the methodological aspects of this study and presents the actions we took as part of conducting it. Section III summarizes the findings of the study, aiming to answer the research questions we have identified in the process. Section IV synthesizes these findings into a research




Fig. 1. The study protocol as a flowchart.

agenda for sustainability-aware software architecture for the coming years. Finally, Section V discusses the threats to the validity of this work, and Section VI concludes it with a summary of its main points.

II. STUDY DESIGN

A. Methodology

To conduct our mapping study in a systematic manner, we are following the guidelines of Petersen et al. [22], [23] for this type of studies in software engineering. The systematic mapping process prescribed by these guidelines can be roughly broken down into the following phases:

- 1) Definition of research questions to scope the study.
- 2) Searching and filtering of relevant papers through well defined inclusion and exclusion criteria.
- 3) Classification of the selected papers through keywording of the abstracts.
- 4) Extraction of relevant data and preparation of the map.

For the purposes of our study we have adopted this process with some amendments as shown in Fig. 1. In the following we discuss each of these phases and our amendments in detail.

B. Definition of Research Questions

As the goal of this work is to provide an overview of the research on sustainability in software architecture we pose questions from three distinct perspectives. First, we ask what are the overall publication trends in this research area in the last years, and how mature is the related literature. Second, from a software architecture perspective, we are interested in which activities of the architecting life cycle [24], [25] are taken into consideration when sustainability is of concern for the systems under study. Third, from a sustainability research perspective, we want to know which of the dimensions of sustainability as a software quality [4] are discussed in the literature. These perspectives translate into the following research questions:

- RQ₁** What are the existing approaches related to sustainability in the field of software architecture as reflected by the literature?
- RQ₂** To what extent are the activities of the architecting life cycle covered?
- RQ₃** To what extent are the dimensions of sustainability covered?

C. Search Process

As per the recommendation of Petersen et al. [23], we focus only on the Population and Intervention dimensions of the

TABLE I
SEARCH RESULTS PER ONLINE BIBLIOGRAPHICAL SOURCE

Source	#Publications
ACM Digital Library	22
IEEE eXplore	138
Science Direct	20
Scopus	187
Wiley Online Library	6
Total:	373

PICO strategy [26] for identifying keywords and formulating search strings from the research questions. More specifically, as the intended Population of our study we define “*all primary studies discussing the architecture of sustainable software systems and relevant activities*” where an Intervention of “*one or more dimensions of sustainability as a software quality is the target of the presented approach*” takes place. Based on this definition and after extensive piloting to calibrate the efficacy of our search we decided on using the query:

```
"software architect*" AND sustainab* AND
(requirements OR design OR implementation OR
evaluation OR decision)
```

which allows searching for publications using variations of the main concepts (sustainable and sustainability, and software architecture, architecting, and architects) in combination with activities that are indicative of actually discussing software architecture (e.g. eliciting requirements, system design, etc.) Five online bibliographical sources were used for the search as shown in Table I. The query was adjusted accordingly for each source based on the syntax and limitations of their respective search engines. The search took place in June 2021. It was additionally constrained to publications after 2000 since previous secondary studies on sustainability in software engineering [9], [11], [12] have not discovered any relevant papers before this year. Table I shows the result of this search per source, for a total of 373 candidate publications identified.

D. Publications Filtering

In alignment with our research questions we defined the following inclusion criteria for this study:

- IC₁** Publication focuses on software architecture.
- IC₂** Publication is related to one or more sustainability dimensions.
- IC₃** Publication is peer reviewed (journal articles, conference and workshop proceedings, magazine articles).

A publication was excluded from our study if one or more of the following criteria were satisfied:

- EC₁ Publication is not published in English.
- EC₂ The publication full text is not publicly accessible.
- EC₃ Publication is a duplicate of another study.
- EC₄ Publication is a secondary study.

Both the last inclusion and the first three exclusion criteria can be checked quite easily automatically. For this purpose we created a set of scripts that took the search result lists from the previous step as an input, retrieved all references and publication metadata in a unified format, merged the lists and removed duplicates, and checked the language and type of the publication based on its bibliographical data. Publications that could not be automatically retrieved were flagged for possible manual retrieval, with 71 publications being found to be completely irretrievable for various reasons. This resulted in a set of 220 publications after this automated filtering for further processing as shown in Fig. 1.

The last exclusion criterion (EC₄), however, and the first two inclusion criteria (IC₁ and IC₂) required reading the whole publication in order to decide if it was admissible to the study. We remove secondary studies because we feel that such works should be discussed separately (e.g. in a tertiary study) as their aggregation and synthesis process inadvertently introduces bias to the mapping process. The latter two criteria proved quite elusive and needed further deliberations for their application. Ultimately, we consider that a publication focuses on software architecture *if we are able to identify which architectural activities it concerns itself with*. This is irrespective of whether the work in question is explicitly contributing to the study of software architectures in general, or discusses a reusable architecting approach to a specific problem across a whole domain. The relation to sustainability, on the other hand, is established through the primary study providing sufficient details on *how its content relates to one or more of the sustainability dimensions as they pertain to software systems*. With respect to the latter point we found quite common for publications to use sustainability as a keyword in their opening and/or closing sections without however providing any connection to the concept throughout their main body. As a result we ended up removing quite a lot of publications from the final sample due to their superficial relation with sustainability. Out of the 220 publications from the automated filtering step, 50 were selected for inclusion in our study through this process.

Since both the identification of the discussed sustainability dimension(s) and the architecting phase(s) is part of the data extraction, as will be discussed further in the following, we ended up doing this second stage of filtering together with the data extraction one. In order to minimize the introduced bias, two co-authors worked on these two tasks independently but at the same time. Disagreements between them concerning the inclusion/exclusion of a primary study were resolved in iterative consensus meetings mediated by a third author who also read through all publications.

TABLE II
DATA EXTRACTION FORM USED IN THIS STUDY; FIELDS MARKED WITH
ASTERISK EXTRACTED AUTOMATICALLY

Field	Values	RQ
Publication Venue*	Full name of venue	
Venue Type*	Journal or Conference	
Publication Date*	Year of publication	
Keywords	Comma-separated list of keywords	RQ ₁
Research Facet	One of the facets identified by Wieringa et al. [27], as per [22]	
Contribution Viewpoint	One of the viewpoints identified in Venters et al. [15]	
Architecting Phases	One or more from the phases identified by Tang et al. [25]	RQ ₂
Sustainability Dimensions	One or more from the dimensions identified by Lago et al. [4]	RQ ₃

E. Data Extraction & Snowballing

Table II summarizes the form used for data extraction from the (primary) studies, following the example of Petersen et al. [22]. Keywords were extracted from the title and abstract of each study, reusing in some, but not all cases the keywords provided in the publication itself. Obvious keywords such as ‘software architecture’ or ‘sustainability’ and derivatives were not included since they were used for the filtering of the study. Publication venue types were bundled into two categories: *journals*, including also magazines, and *conferences*, including also workshop proceedings. Together with the publication date, these venue-related fields were extracted automatically from the metadata of the publications with occasional manual interventions for providing missing data. Research facets use the classification of Wieringa et al. [27] to categorize research efforts into *validation or evaluation research*, *solution proposals*, or *philosophical, opinion, or experience papers*. Following Venters et al. [15] we characterize primary studies depending on whether they discuss *sustainable software*, where the software system and its architecting process is the goal of the study, or *sustainability through software*, where delivering sustainability to the stakeholders is the product of this process.

Instead of the more widely accepted but less comprehensive categorization scheme of Hofmeister et al. [24], we use the one by Tang et al. [25] to identify architecture-related activities in five phases: *analysis*, *synthesis* (design), *evaluation*, *implementation*, and *maintenance* (covering also evolution). To characterize the sustainability dimension(s) addressed by the primary study we use the Lago et al. [4] proposal of looking at the economic, environmental, social, and technical aspects of sustainability as a software quality. More than one phases or dimensions could be addressed by the same study, so both these fields are structured as comma-separated lists of elements from the respective categorization scheme.

As discussed above, extracting data for these latter two fields provided us with the means of verifying IC₁ and IC₂: if no phase or dimension could be extracted for a study then it was treated as a strong indication that it had failed the respective inclusion criterion. This however meant that we could perform a search for additional studies through back-




Fig. 2. Publication year and venue type; proceedings include both conference and workshop proceedings.

TABLE III
PUBLICATION VENUES WITH MORE THAN ONE STUDY; CONFERENCES INCLUDE ALSO CO-LOCATED WORKSHOPS, EXCEPT IF PUBLISHED INDEPENDENTLY FROM THE MAIN CONFERENCE.

Venue	Type	#Studies
European Conference on Software Architecture (ECSA)	Conference	8
Working IEEE/IFIP Conference on Software Architecture (WICSA) — later renamed to: IEEE International Conference on Software Architecture (ICSA)	Conference	7(5 + 2)
IEEE Software	Journal	4
Information and Software Technology	Journal	3
IEEE Access	Journal	2
International Workshop on Requirements Engineering for Sustainable Systems (RE4SuSy)	Conference	2

ward snowballing [28] only after the combined filtering and extraction was finalized. In practice, this turned out not being a major issue; we found 8 additional studies by snowballing through the 50 studies that were filtered by this stage and carried out an additional extraction for them. The results of the mapping process presented in the following section aggregate both sets of 58 primary studies. The final list of primary studies, together with the outcomes of the previous steps of the process are available online¹.

III. FINDINGS

A. RQ_1 — State of the Art overview

To answer our first research question we start by looking at the publication of the selected studies using descriptive statistics. Fig. 2 summarizes the number of publications per year after 2000 and per publication venue type: journal or magazine articles, and workshop or conference proceedings. Publications in proceedings dominate with 67% of the total, indicating space for more mature works to appear in the literature. Table III outlines the publication venues with $n > 1$ from the selected studies. Perhaps unsurprisingly, the

¹<https://doi.org/10.6084/m9.figshare.19077860>




Fig. 3. Top keywords categories after consolidation; multiple categories per study are possible.

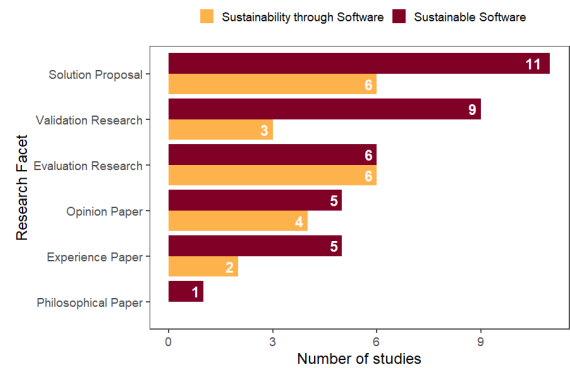


Fig. 4. Research facets of the studies, together with their viewpoint with respect to sustainability

most popular venues for publishing on the topic appear to be the European and International Conferences on Software Architecture (ECSA and ICSA respectively), the latter in both its incarnations as WICSA and later ICSA. The Software and Access IEEE journals, and the Information and Software Technology (IST) journal seem to be also preferred venues for publishing on the topic. Of special mention is the International Workshop on Requirements Engineering for Sustainable Systems (Re4SuSy) run in conjunction with the Requirements Engineering (RE) conference. The fact that the remaining 32 studies are published in 32 different venues is possibly a sign of a need for more concentrated targeting of the research output by the community.

Further, we use the extracted keywords to identify the most popular topics covered by the studies in the field. Since the keywords themselves are defined at different levels of abstraction and exhibit a high degree of variability in the use of terms for the same concepts, we first organize them into categories. Fig. 3 presents the top ten such categories ordered by popularity. The topics that are most frequently discussed fall within the *architectural concepts* category, aggregating such topics as architectural principles, patterns, and tactics, design decisions and the respective principles and patterns, and software product lines and variability-related keywords.




Fig. 5. Architecting phases addressed in the primary studies and their intersections.

Going beyond these simpler descriptive statistics, we next look into how to analyze and synthesize the remaining extracted data into a richer mapping of the state of the art. The “traditional” way of achieving this result is through a bubble plot a la Petersen et al. [22], with e.g. the research facets, architecting phases, and sustainability dimensions as the axes of the plot, and the size of the bubbles signifying the number of studies in the intersection of the axes. However, we have two categorical variables that can take multiple values at the same time (phases and dimensions), and one variable (contribution viewpoint) that is by definition correlated with another variable (sustainable software being usually an expression of technical sustainability [15]). As such, we feel that the bubble plot only diffuses the presented information instead of illustrating it, and we thus opt to present this information distributed across multiple complementary plots.

For the first part we are looking at the research facet of the primary studies, and their contribution viewpoint. Fig. 4 synthesizes the extracted information from these fields to allow us to discuss the type of research conducted in the field. From the figure we can conclude the following:

- 1) While the majority of studies focus on sustainable software (~64%), the ones aiming to deliver sustainability through software is not insignificant (~36%). Given the fact that the former viewpoint is heavily associated with the technical dimension, this is somewhat expected.
- 2) Sustainable software-focused studies dominate only two out of the three “applied” research facets, i.e. *solution proposals* and *validation research*, with *evaluation research* studies being split equally between the two

viewpoints. This can be interpreted as a sign of both viewpoints perceived as the potential target for adoption by practitioners.

- 3) There is a distinct lack of the more reflection-aiming types of studies (*opinion papers*, *experience papers*, and particularly *philosophical papers*). It is our opinion that having more such studies in the future can facilitate further the maturation of the research topic.

Before moving on to also combine phases and dimensions, we first examine each of these extracted fields in isolation as the means of answering RQ₂ and RQ₃, respectively.

B. RQ₂ — Architecting phases coverage

In order to answer RQ₂ we summarize the architecting phases addressed by the primary studies and their intersections using an UpSet plot [29]. UpSet plots visualize set intersections in a matrix layout, allowing for quantitative analysis of data from more than three sets at a time, such the ones here. As seen in Fig. 5, activities related to architecture synthesis are overall addressed the most, appearing in ~64% of the studies. However, maintenance (and evolution) is the most popular phase discussed *individually*, as hinted already by the popularity of the respective keywords. On the other hand, implementation, referring to the low level design of the system, is clearly the least-studied phase, discussed in only ~10% of the studies. Also, we noticed that only 15 studies (~26%) are addressing *more than 3 phases* at the same time, and 10 out of these 15 are due to the prevalence of combining analysis-synthesis-evaluation. Furthermore, the majority of primary studies (37 out of 58, ~64%) are staying within the Hofmeister




Fig. 6. Sustainability dimensions addressed in the primary studies and their intersections

et al. model of architecting activities (i.e. analysis, synthesis, and evaluation, or a combination thereof). Only 3 studies (a meager 5%) address 4 phases or more, and only 1 addresses all five.

C. RQ_3 — Sustainability dimensions coverage

Similarly to the architecting phases, we use the UpSet plot of Fig. 6 to answer RQ_3 . As shown in the figure, the technical dimension is by far the most popular one in the selected primary studies, appearing in a staggering $\sim 86\%$ of all studies. As a matter of fact, it is more popular than all other dimensions put together including their combinations, and it is also discussed in most combinations with any other dimension. Social, in turn, is the most under-addressed dimension (13 studies, $\sim 22\%$), and almost always discussed in combination with other dimensions. Similarly, the economic dimension, while discussed as much as the environmental one, is always treated together with another dimension and usually the technical one. Perhaps more importantly, however, it appears that the focus on the technical dimension outweighs the number of approaches dealing with multiple dimensions: only 10 studies ($\sim 17\%$) address 3 or more dimensions, and only 4 studies ($\sim 7\%$) all four of them.

D. Combining phases and dimensions

Returning now to RQ_1 , we plot the co-occurrences of the phases and dimensions appearing in each primary study as a heatmap in Fig. 7. This provides us with an eagle eye’s view of their interactions. As shown in the figure, for example, *only two dimensions are discussed across all phases of software architecting*: the *technical and economic* ones; however, the

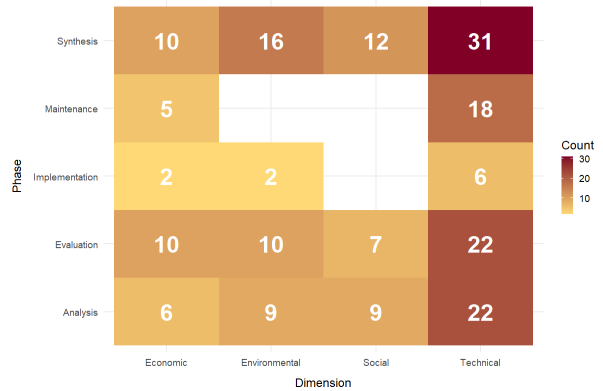


Fig. 7. Co-occurrences of sustainability dimensions and architecting phases addressed by the primary studies

former is much more popular than the latter (50 out of 58 studies, as discussed above). There are also obvious gaps in phases discussed in conjunction with the environmental and social dimensions, and a need for more research addressing the non-technical dimensions in general.

IV. DISCUSSION

In the previous section we presented our analysis of the state of the art of the intersection of sustainability and software architecture. We stayed clearly away from identifying this area as “sustainable software architecture” since it has been historically associated with a very specific inward-facing meaning of sustainability in the literature associated with its resilience

under change. We focused on mapping out the most popular aspects of the state of the art, and more importantly, on identifying gaps. One point that was perhaps lost in this analysis is the set of possibilities created by the accumulated work in the area. To be more specific: the number of available works and their apparent increasing maturity present an opportunity for the establishment of *sustainability-aware architecture* as a separate field of study. Such field would concern itself first and foremost with the interaction between architecting activities from across the system lifecycle and the various dimensions of sustainability as a software-system quality property.

An important step towards realizing this vision is addressing first the diagnosed gaps in our knowledge of the topic. To this aim, we identified a list of research items we find promising for the research community to work on in the coming years:

- 1) **Actively pursuing reflection-oriented studies.** As discussed in Section III-A, the literature seems to favor the more practical types of research at the expense of the more reflection-oriented ones such as opinion, experience, and philosophical papers. This difference is expected to become smaller as the field continues to mature and more knowledge is accumulated by the involved researchers. However, actively pursuing more of these types of studies already offers the opportunity for performing a sanity check through deliberate reflection to a still growing field such as the one discussed here.
- 2) **Wider coverage of the architecting lifecycle.** Figures 5 and 7 paint a picture of inequalities in the current works with respect to their treatment of phases. With the exception of works exclusively focusing on maintenance, there are few works going beyond the core Hofmeister et al. analysis-synthesis-evaluation combination of phases. More works addressing ideally the whole lifecycle are therefore necessary, or at least considering the implementation and maintenance phases, and specifically going beyond the intersection of the latter with the technical dimension.
- 3) **Going beyond the technical dimension.** In a similar manner, and by combining Fig. 6 with 7, we can easily observe how pervasive is the technical dimension in the current discourse. To some extent, this is to be expected: the notion of sustainable software itself has been equated for many publications with that of maintainability, adaptability, resilience, and other concepts under the technical sustainability umbrella [15]. However, and especially in the more recent years, the research community has been stressing the need for a multi-dimensional approach to dealing with sustainability as formulated by the Karlskrona manifesto [1]. As such, there is an urgent need for approaches incorporating as many dimensions as possible, and addressing especially the least popular ones (environmental, economic, and social).
- 4) **Towards a sustainability-aware architecting framework.** There is no primary study in our sample which covers all dimensions and phases at the same time. This might have been understandable and acceptable while the field was still

in its inception, but with its increasing maturity it becomes important to be able to deal with the multi-dimensionality of sustainability in a structured and holistic manner. Assuming that the identified gaps in the state of the art begin to be covered by the previous items in this agenda, an appropriate *sustainability-aware architecting framework* [30] would then be necessary for integrating them into one coherent solution. We plan to pursue this by extending the SAF Toolkit [31].

V. THREATS TO VALIDITY

For the discussion on the threats to the validity of this work, and the mitigating measures we took, we use Ampatzoglou et al. [32] as our guide. Accordingly, we acknowledge potential issues along all three of the there-defined threat categories:

Study Selection Validity: This category aggregates threats rising from the first two phases of secondary studies, i.e. search and filtering. The publications used for the mapping process in this study were retrieved by querying 5 online databases using as broad query terms as possible. The query string used was constructed in alignment with the PI(COC) search strategy, and was piloted over multiple iterations before converging to the form used in this study. Duplicate studies were automatically removed from the final sample, and snowballing was performed to provide additional depth to our search. In addition, selection criteria were defined following the research questions, and their application was performed first independently and then after consensus' building to minimize the bias.

Data Validity: This category includes threats applying to the data extraction and analysis phases of the secondary study. These threats are roughly organized in three groups: limitations of the dataset, data extraction bias, and research-introduced bias. With respect to the first, and by including primary studies that propose (reusable) architectures in addition to works that study explicitly architecture as a topic, we increased both the size and heterogeneity of the dataset. To mitigate the issues that arise from data extraction, and similarly to the filtering process, we performed the extraction first independently and then developed a consensus through discussion among the researchers. The senior researchers involved in this study have provided separate quality assessment to this goal by also reading all the publications to facilitate the consensus building and verify the data extraction. Finally, to minimize research-introduced bias we adopted widely accepted conceptual frameworks (architecting phases defined by Tang et al. [25] and sustainability dimensions by Lago et al. [4]) for classification purposes.

Research Validity: This last category is concerned with the overall research design; it applies throughout the phases of the study and focuses mainly on two aspects: generalizability and repeatability. With respect to the former, our research questions are motivated by both the lack of a similar review in the literature, and the necessity of the topic itself. This means that we cannot use other studies for judging the generalizability of the results. However, the used search string as defined in Section II was piloted in a number of iterations to be as broad as possible while still returning results that are suitable to this

study. With respect to the latter, the replication package for the study has been made available online and is publicly available for replication purposes.

VI. CONCLUSION

Software engineering studies have been increasingly concerned with the topic of sustainability, both as it is delivered by the systems they study, and of the software systems themselves. A number of secondary studies from the last years are attesting to the existence of a vibrant community working on various related research topics. One such topic, that of the intersection of software architecture and sustainability, is the focus of this work. Having no other secondary study specifically addressing this topic, in this work we systematically map existing approaches from the literature on both sustainable architectures and sustainability through architecture.

Our findings show a quickly maturing research community that however needs to address specific deficiencies in its coverage of both architecting activities and sustainability dimensions. Based on our analysis, we propose a list of research items for the following years towards the establishment of sustainability-aware (software) architecture as a new field of study. The items in this research agenda are the focus of our future work.

ACKNOWLEDGEMENTS

This work is partly funded by the project Sustainable-Cloud (project number OCENW.M20.243) of the research programme Open Competition Domain Science by the Dutch Research Council (NWO).

REFERENCES

- [1] C. Becker *et al.*, “The Karlskrona manifesto for sustainability design,” Tech. Rep., oct 2014. [Online]. Available: <http://arxiv.org/abs/1410.6968>
- [2] GeSI, “Digital with Purpose: Delivering a SMARTer 2030,” GeSI, Tech. Rep., 2019.
- [3] B. Penzenstadler *et al.*, “Safety, security, now sustainability: The non-functional requirement for the 21st century,” *IEEE Software*, vol. 31, no. 3, pp. 40–47, 2014.
- [4] P. Lago *et al.*, “Framing sustainability as a property of software quality,” *Communications of the ACM*, vol. 58, no. 10, pp. 70–78, sep 2015.
- [5] S. Naumann *et al.*, “The GREENSOFT Model: A reference model for green and sustainable software and its engineering,” *Sustainable Computing: Informatics and Systems*, vol. 1, no. 4, pp. 294–304, Dec. 2011.
- [6] S. S. Mahmoud and I. Ahmad, “A green model for sustainable software engineering,” *International Journal of Software Engineering and Its Applications*, vol. 7, no. 4, pp. 55–74, 2013.
- [7] E. Kern *et al.*, “Impacts of software and its engineering on the carbon footprint of ICT,” *Environmental Impact Assessment Review*, vol. 52, pp. 53–61, Apr. 2015.
- [8] A. Hindle, “Green Software Engineering: The Curse of Methodology,” in *2016 IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*. Institute of Electrical and Electronics Engineers (IEEE), May 2016, pp. 46–55.
- [9] B. Penzenstadler *et al.*, “Sustainability in software engineering: A systematic literature review,” in *International Conference on Evaluation & Assessment in Software Engineering (EASE)*. IET, 2012, pp. 32–41.
- [10] G. A. García-Mireles *et al.*, “Interactions between environmental sustainability goals and software product quality: A mapping study,” in *Information and Software Technology*, 2018, vol. 95, pp. 108–129.
- [11] N. Wolfram, P. Lago, and F. Osborne, “Sustainability in software engineering,” in *IFIP Conference on Sustainable Internet and ICT for Sustainability (SustainIT)*. IEEE, Jun. 2018, pp. 1–7.
- [12] C. Calero *et al.*, “5Ws of green and sustainable software,” *Tsinghua Science and Technology*, vol. 25, no. 3, pp. 401–414, Jun. 2020.
- [13] C. Calero and M. Piattini, “Puzzling out Software Sustainability,” *Sustainable Computing: Informatics and Systems*, vol. 16, pp. 117–124, Dec. 2017.
- [14] P. Lago and B. Penzenstadler, “Editorial: Reality check for software engineering for sustainability-pragmatism required,” *Journal of Software: Evolution and Process*, vol. 29, no. 2, Feb. 2017.
- [15] C. C. Venters *et al.*, “Software sustainability: Research and practice from a software architecture viewpoint,” *The Journal of Systems and Software*, vol. 138, pp. 174–188, 2018.
- [16] H. Koziolok, “Sustainability evaluation of software architectures: a systematic review,” in *Joint ACM SIGSOFT conference on Quality of software architectures—QoSA and architecting critical systems—ISARCS*, 2011, pp. 3–12.
- [17] T. Volpato *et al.*, “Two perspectives on reference architecture sustainability,” in *Proceedings of the 11th European Conference on Software Architecture: Companion Proceedings*, 2017, pp. 188–194.
- [18] P. H. D. Valle *et al.*, “Towards suitable description of reference architectures,” *PeerJ Computer Science*, vol. 7, p. e392, 2021.
- [19] R. Verdecchia, I. Malavolta, and P. Lago, “Architectural technical debt identification: The research landscape,” in *2018 IEEE/ACM International Conference on Technical Debt (TechDebt)*. IEEE, 2018, pp. 11–20.
- [20] T. Volpato, A. Allian, and E. Y. Nakagawa, “Has social sustainability been addressed in software architectures?” in *ECSA '19: Proceedings of the 13th European Conference on Software Architecture*, vol. 2. New York, New York, USA: ACM, sep 2019, pp. 245–252. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=3344948.3344979>
- [21] E. M. Grua *et al.*, “Social Sustainability in the e-Health Domain via Personalized and Self-Adaptive Mobile Apps,” in *Software Sustainability*, C. Calero, M. Á. Moraga, and M. Piattini, Eds. Springer, 2021, pp. 301–328.
- [22] K. Petersen *et al.*, “Systematic mapping studies in software engineering,” in *12th International Conference on Evaluation and Assessment in Software Engineering (EASE) 12*, 2008, pp. 1–10.
- [23] K. Petersen, S. Vakkalanka, and L. Kuzniarz, “Guidelines for conducting systematic mapping studies in software engineering: An update,” *Information and Software Technology*, vol. 64, pp. 1–18, 2015.
- [24] C. Hofmeister *et al.*, “A general model of software architecture design derived from five industrial approaches,” *Journal of Systems and Software*, vol. 80, no. 1, pp. 106–126, Jan. 2007.
- [25] A. Tang *et al.*, “A comparative study of architecture knowledge management tools,” *Journal of Systems and Software*, vol. 83, no. 3, pp. 352–370, 2010.
- [26] B. Kitchenham and S. Charters, “Guidelines for performing systematic literature reviews in software engineering,” EBSE Technical Report, Tech. Rep., 2007.
- [27] R. Wieringa *et al.*, “Requirements engineering paper classification and evaluation criteria: a proposal and a discussion,” *Requirements engineering*, vol. 11, no. 1, pp. 102–107, 2006.
- [28] S. Jalali and C. Wohlin, “Systematic literature studies: database searches vs. backward snowballing,” in *Proceedings of the 2012 ACM-IEEE international symposium on empirical software engineering and measurement*. IEEE, 2012, pp. 29–38.
- [29] A. Lex *et al.*, “Upset: visualization of intersecting sets,” *IEEE transactions on visualization and computer graphics*, vol. 20, no. 12, pp. 1983–1992, 2014.
- [30] V. Andrikopoulos and P. Lago, “Software sustainability in the age of everything as a service,” in *Next-Generation Digital Services - Essays Dedicated to Mike Papazoglou*, ser. LNCS. Springer.
- [31] P. Lago and N. Condori-Fernandez, “The Sustainability Assessment Framework (SAF) Toolkit: Instruments to help Sustainability-driven Software Architecture Design Decision Making,” Apr. 2022. [Online]. Available: <https://github.com/S2-group/SAF-Toolkit>
- [32] A. Ampatzoglou *et al.*, “Identifying, categorizing and mitigating threats to validity in software engineering secondary studies,” *Information and Software Technology*, vol. 106, pp. 201–230, 2019.