Developing a Framework for Evaluating the Sustainability of Computing Projects

Anton Lundström[1]
KTH Royal Institute of Technology
School of Computer Science and Communication
Stockholm, Sweden
antlund@kth.se

Daniel Pargman[1,2]
KTH Royal Institute of Technology
School of Computer Science and Communication.
Centre for Sustainable Communications
Stockholm, Sweden
pargman@kth.se

ABSTRACT

Toyama [19] has proposed a "preliminary taxonomy" for classifying computing projects as a way of separating sustainable computing efforts from unsustainable ones. In this paper we explore the feasibility of Toyama's taxonomy. We begin by describing how we revised and developed his taxonomy to make it more practically useful and then conducted a pilot study where we used the revised version to evaluate four computing projects. The pilot study was then used as a foundation for further discussing and developing the revised taxonomy into yet another, third and final version which we have chosen to call the Sustainable Computing Evaluation Framework (SCEF). While our proposed framework (SCEF) is more practically useful than Toyama's "preliminary taxonomy", there are still challenges that need to be addressed and we end the paper by suggesting where future efforts could be focused.

CCS CONCEPTS

• Social and professional topics~Computing industry • Social and professional topics~Sustainability • Social and professional topics • Social and professional topics~Professional topics

KEYWORDS

Sustainability; Computing; Taxonomy; Sustainable Computing; Framework

1 INTRODUCTION

There is no doubt we nowadays face substantial challenges in terms of sustainability and radical measures are needed to reach the 2 degree climate goal and achieve a sustainable society. Since computing has become an integral part of modern society, it is natural to ask what computing could do to help create more sustainable societies. Several computer-related areas do ask that question, for example ICT for Sustainability (ICT4S), Sustainable Human-Computer Interaction (S-HCI), Environmental SE4S), Computing within Limits (LIMITS) and so on. These areas all explore how computer systems could be used to increase sustainability and they all by default assume that computing can have a net positive impact in terms of sustainability. While it is well-known that computing also can, and does have negative sustainability impacts due to the production, use and disposal of components (so-called first-order effects, see further [10]), it is implicitly assumed that these negative effects will be outweighed by the positive effects and that efficiency improvements and service substitution can and will result in increased sustainability [9]. But how can we really know if that is the case? To make an assessment of what constitutes a sustainable (or unsustainable) computing project, an evaluative framework is needed.

Toyama [19] proposed a taxonomy for sustainable computing in the hope that it would be able to help shed light on which computing efforts contribute to sustainability and which don't. His basic suggestion was that computing projects should be classified and evaluated according to three different dimensions. Toyama was however careful to present his proposed taxonomy as *preliminary*. The taxonomy was not robust enough to be of direct practical use and could thus not immediately be used to evaluate if a system X in fact could be regarded as sustainable or not, nor did Toyama [19] provide any direct suggestions for how the taxonomy could be improved. This paper represents an attempt to move the taxonomy from *preliminary* to *operational*.

2 A FRAMEWORK FOR SUSTAINABLE COMPUTING

Toyama [19] proposed a "Taxonomy of value for sustainable computing" in the hope that it would help the computing community direct its attention towards sustainability by providing a way to evaluate solutions that claimed to be, or claimed to contribute to sustainability.

The taxonomy (see Figure 1) consists of three dimensions according to which computing systems can be classified in terms of sustainability, namely "*Impact*", "*Intention*" and "*Effort requirements for impact*" (see figure 1 below). Toyama's definition of sustainability builds on a focus on physical resource use, e.g. an "*equilibrium where the amount of resources being*

LIMITS'17, June 22-24, 2017, Santa Barbara, CA, USA.

ACM ISBN 978-1-4503-4950-5/17/06...$15.00.
DOI: http://dx.doi.org/10.1145/3080556.3080562

replenished is equal to the amount of resources being depleted". This criterion is very hard to attain in any practice or industry that depends on non-renewable materials such as minerals or fossil energy sources. It is however semi-compatible with ideas about a circular economy [13] a steady-state economy [3] and of calculating a sustainable ecological footprint that we should keep within [20, 14].

Toyama never directly discusses the use of non-renewable resources such as various metals including rare earth minerals and it could be argued that computing by definition is unsustainable. It does however seem that Toyama subscribes to some version of Daly's "quasi-sustainability" criterion [4], i.e. a "*smartphone app that helps users conserve water may do so at the expense of greater fossil fuel consumption*" [19]. That implies that it can be fine to use (some) fossil fuels (and presumably other non-renewable resources) as long as it saves (a lot of) water. We will not further delve into the concept of sustainability (and how to operationalise it) here, but will instead discuss the proposed taxonomy and the criteria for evaluating whether a computing system is sustainable or not.

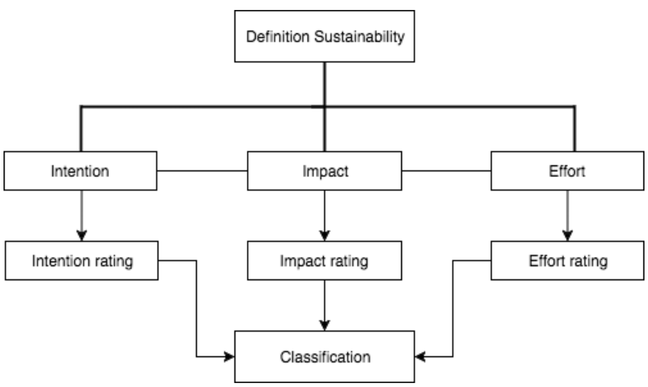

Figure 1: Preliminary Taxonomy

Toyama sketched out a classification system (we will refer to this as rating) for the three proposed dimensions where:

- **Impact** (on sustainability) varies between -3 ("*Adversely affects sustainability*") to +1 ("*Contributes to movement toward a globally sustainable equilibrium*")

- **Intention** (towards sustainability) varies between A ("*Genuine intention to move things toward increasing sustainability*") and C ("*Intention to move things in a direction that runs counter to sustainability, or negligence toward an incidental effect that runs counter to sustainability*")

- **Effort** (for achieving impact) varies between "unlikely" ("*Requires significant or sustained effortful activity that people are unlikely to take up without a considerable external impetus*") and "effortless" ("*Requires almost no significant change in behavior among people or societies*").

Each proposed dimension comes with its own set of problems. Some of these can be anticipated, while others are unknown and are likely to surface only when attempting to put these ideas to use. Anticipated challenges include how to decide on system boundaries and measurements for assessing the Impact, how to understand what constitutes "*intention toward sustainability*" in assessing the Intention, and that Effort will likely vary depending on the stage of implementation etc.

We (as well as Toyama himself[1]) acknowledge that the taxonomy was framed as "preliminary" and that it sufferers from a number of weaknesses. In that vein, we here treat the proposed taxonomy not as the *final* step, but as the *first* step since we believe that the idea of *creating* a taxonomy for evaluating which computer systems are more and which are less sustainable is a worthy task. While the term "taxonomy" might have been useful for Toyama's attempt at classifying and categorising the (three) dimensions that makes up "sustainable computing", we think that the term "*framework*" better fits the outcome of the work we report upon here. We have thus used Toyama's taxonomy as a foundation for developing a framework for classifying computing projects and our work has progressed in three steps:

First by developing Toyama's taxonomy so that it can be used for classifying computing projects and ascertain whether they are sustainable or not (or to what extent they are (un-)sustainable). We have done so by elaborating on the three dimensions so as to make them more operationally useful, and this has resulted in a framework that we from now on will refer to as a "preliminary framework" or just "the framework".

Second by putting the framework to use. We conducted interviews with representatives from four computing projects and used the framework to analyse the answers.

Third and based on the results of our study, we have revised the framework and developed yet another version of it which we from now on will call the Sustainable Computing Evaluation Framework, or for short SCEF.3 METHOD

We started with Toyama's own recommendation that the taxonomy needed a clear protocol to become practically useful and we operationalized this by adding a protocol (a procedure for making a classification). This protocol is temporarily represented as three empty boxes in Figure 2.

[1] Toyama, personal communication (May 13, 2016).

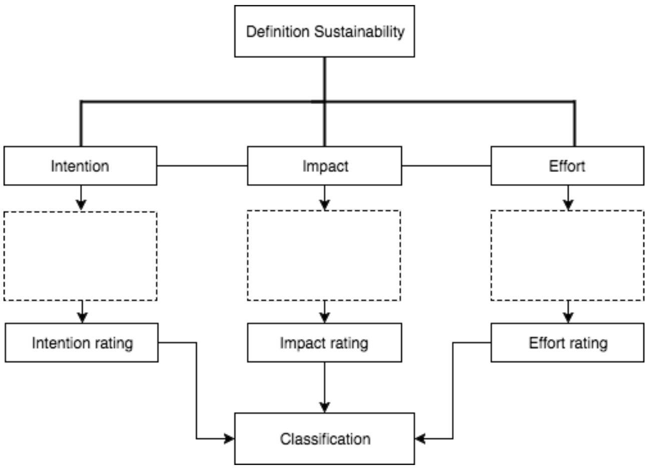

Figure 2: Framework development process

To develop the taxonomy (e.g. to 'fill' the empty boxes) we conducted a literature review and were inspired by [16, 5, 2 11, 12] and others. The result is a preliminary framework (see further Figure 3) and the added protocol for each dimension are explained below:

1. The protocol for assessing Intention includes the questions *What* (is to be sustained?), *How* (is it being sustained?) and *Why* (is it being sustained?). The protocol further requires a description of *How* those responsible for a computing system intend to achieve sustainability, *Why* they intend to achieve it, and by *What* means they intend to achieve it. The system for rating (between "A" and "C") the intention of those responsible for a computing systems has been left unchanged.

2. The protocol for Impact includes *Direct Impact* (impact from the production to the disposal in terms of use of raw materials and energy), *Enabling Impact* (behavior change and process optimization) and *Structural Impact* (institutional, macroeconomic, and societal impact) [12]. An impact assessment needs to consider these three types of impacts so as to make a classification of a computing system. The system for rating of the impact of computing systems between -3 to 1 has been left unchanged.

3. In the process of designing the framework, a decision was made to change *Effort* to *Likelihood of Impact (Likelihood),* drawing on Rogers' [16] theory about the diffusion of Innovations, where we used the attributes of innovations as a protocol. The attributes are: *Relative Advantage* (compared to other products/services of similar kind), *Compatibility* (is it consistent with the need, values, and experiences of potential users?), *Complexity* (is a product/service perceived as difficulty to understand or use?), *Trialability* (is it available for testing before committing?) and *Observability* (is the result of using a computing system visible to others?). Accordingly the rating system was also changed and now covers a scale from *Unlikely* to *Likely*. We have interpreted the function of the dimension "Effort" as an attempt to understand the likelihood that a given project will have an impact or not. We find that what is perceived as effortful or effortless is liable to depend on various contingencies, such as a computing system's stage of implementation or the degree to which it is already being used, where for example using a smartphone will be perceived as effortless given that it is already widely adopted in society. Switching to Likelihood enables the framework to address the same issue but without the risk of getting tangled up in exactly what effort amounts to.

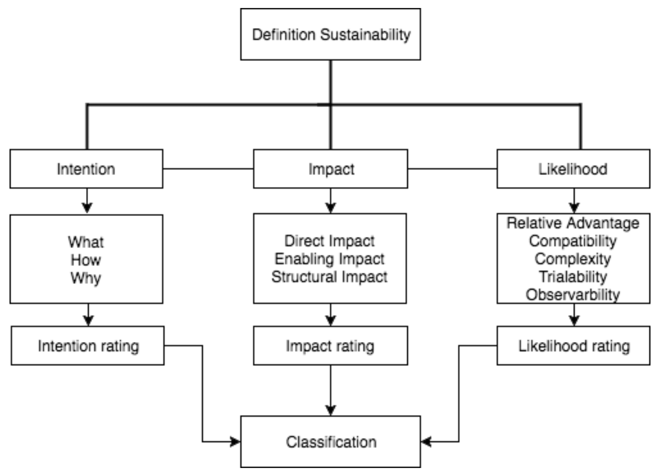

Figure 3: Preliminary framework

3.1 Interviews

We interviewed four representatives of different computing projects and our goal was to 1) connect theory to practice and 2) to get feedback on how the preliminary framework performed when it was applied, and by extension information about how it could be further improved. It should be noted that we were not particularly interested in actually evaluating the computing projects per se, but were rather primarily interested in these projects to the extent that they could help us further revise and improve the preliminary framework. If our intention had been to actually classify the sustainability (or not) of these projects, then we would have conducted the study in another manner (for example by not confining ourselves to interviewing only one representative of each project, see for example [21]).

We used two criteria for finding suitable projects and suitable persons to interview, e.g. the projects 1) *should address sustainability* and 2) *have a clear connection to computing*. There does of course exist a large variety of computing projects, and different types of projects can pose various difficulties in a classification process. This was considered less important at this particular point in time as our purpose was primarily to explore the framework itself. Examining projects that differ from each other is in this case useful, due to the fact that it can draw attention to problems with the framework from various angles. The four projects in question were:

- A start-up company providing a mobile app that allows users (households) to get an overview of their energy consumption by gathering data from smart meters in order to analyze and visualize the data, *EnergyViz.*

- A start-up company providing an electrical vehicle charging system. The system consists of smart power sockets that sends relevant information to the user's mobile app, *ChargeCar*.

- A research project that resulted in a web-based interactive design prototype presenting scenarios for a sustainable future society, *FutureLife*.

- A non profit organization owning/offering a standard for providing third party sustainability certification and labeling for IT hardware products, *HardwareLabel*.

4 SUSTAINABLE COMPUTING?

We will here present selected results from the interviews. These results also provide input to the process of further developing the framework.

4.1 Sustainability

The answers pertaining to sustainability and to how it was connected to the informants' projects varied greatly in depth and in detail. None of the informants considered sustainability in only ecological terms and most instead referred to the three pillar approach [8] for understanding sustainability, e.g. putting ecological sustainability at par with social and economic sustainability. The informants did this either directly or indirectly by highlighting (also) economic and social factors that were important to their projects. None of the informants discussed sustainability in terms of the computing hardware itself (e.g. "the greening of IT", see further [11]). The one exception was *HardwareLabel* (e.g. the informant who represented the project *HardwareLabel*) who worked with certifying and labeling computing (hardware) products and who professionally was mainly concerned with such aspects (e.g. work environment, radiation, toxicity etc.).

EnergyViz was hesitant about using the term "sustainability" and said that it could mean several different things, referring to economic, social, and ecological sustainability and that "*it is perhaps possible to speak of sustainability in a clear way if you can make clear demarcations and connect it to specific actions"*.

4.2 Intention

It is hard to know if informants' stated intentions are intentions toward the goal of their projects/businesses, intentions toward sustainability, or some combination thereof. This is best exemplified by *EnergyViz* who partly took a broad climate oriented stance saying that, *"If you're looking at the development with the climate and the questions we are facing today, it is extremely hard to solve these problems. Now the development is heading in a positive direction and we are trying to be one of the actors who continue the global development around the climate, climate questions, and so forth."* and partly a narrower stance saying that,*"We really just want to supply more information regarding households energy consumption, which can hopefully lead to it* [energy consumption] *decreasing"* and also, *"Our thought is to use the technology, policies and information* [e.g. data gathered through their service] *available today and empower the end-consumer."*

4.3 Impact

The informants had a hard time describing and specifying the potential enabling impact of their respective projects. Even though it seemed this was not something they had thought about explicitly, it was to some extent part of their more general answers in regards to the potential impact of their respective projects. *ChargeCar* referred to how their system would enable users to charge their vehicles when parked (at home or work) instead of having to spend time at a charging station. *ChargeCar* stated that *"Because the existing infrastructure is way to inefficient and demands too much of the user, it's important that there are more user friendly solutions available to make a greater diffusion of electrical vehicles possible and to help make the option of buying an electrical car more attractive". EnergyViz* alluded to the potential of changing users' behaviors through the use of digital tools providing feedback, saying that *"We use visualization, but we also use techniques from social psychology to get people to change their habits and behaviors and enable them to actually lower their energy consumption"*

Questions such as *"What is the goal of the project?"* and *"What effects could the project have?"* did not generate answers that went much further in terms of illuminating what the hoped-for impacts were and most answers did in fact not extend particularly far beyond what any curious reader could learn from each project's website.

4.4 Likelihood of Impact

When the informants was asked about the likelihood of having an impact, external factors was identified as the main hindrance. *EnergyViz* explained how they operated in a *"conservative and undigitized industry"* and that this limited their potential reach. *HardwareLabel* mentioned law and policy frameworks as limiting what they could achieve, explaining how the industry they operated in were reluctant to do more than what was demanded in terms of policy in regards to sustainability. All informants considered their respective services easy to use with one exception; *HardwareLabel* stated that their "product" (labeling) was better than competing labeling schemes in terms of sustainability, but that this also exerted higher demands on their user/customers (e.g. hardware manufacturers).

5 IMPACT, INTENTION AND LIKELIHOOD

The goal of the interviews was not to evaluate particular projects, but to get a feeling for how the framework would perform if applied as-is. Here we will first highlight some findings from the interviews that have implications for the development of the framework and then incorporate these insights into the final version of the framework that we have chosen to call Sustainable Computing Evaluation Framework (SCEF).

5.1 Sustainability and Intentions

Sustainability played a role for all four projects but was of varying importance in the 'day-to-day'-business. Broadly speaking, the projects can be divided into three categories in terms of how important sustainability was:

1. **Central**: Both *HardwareLabel* and *FutureLife* had a clearly formulated idea in terms of sustainability, and used these ideas as a foundation for their respective services.

2. **Peripheral**: *ChargeCar*'s main agenda was to increase the diffusion of electrical vehicles and success of the business. The main agenda was not necessarily to work towards enabling more efficient charging solutions or towards attaining a sustainable society. They did market themselves toward housing cooperatives that had a sustainability agenda, but we judge that sustainability was more of a means than an end.

3. **In-between**: *EnergyViz* gave ambiguous answers to what they were actually trying to achieve, but they were well versed in problems surrounding sustainability. It could be that *EnergyViz* found it hard to reconcile strong sustainability [6] with the daily challenges of running a business.

5.2 Understanding Impact and Likelihood of Impact

Talking about impact proved to be hard. The projects of course had goals for what they hoped to achieve, but these were hard to interpret in terms of concrete effects/impact. When asked about the impact or effect of their projects, the informants did not differentiate between *direct, enabling* or *structural* impact, and most were concerned mainly with enabling impact, for example the hope that information would lead to altered behaviors, which would increase the diffusion of electrical vehicles, which would decrease energy consumption. This makes an impact assessment substantially more complicated not only because it will be necessary to take more effects into consideration, but also because the system boundaries get fuzzy when enabling and structural impact are included in an assessment. Another problem, which affects the assessment of Likelihood of Impact, is that different projects are at different stages of implementation and are therefore, to varying degrees, subject to various external factors. *EnergyViz* was for example easy to use from an end-user perspective, but the biggest barrier for *EnergyViz* to affect large-scale change (e.g. to have an actual measurable impact) was that it was a small player in a conservative industry with giant actors.

6 A FRAMEWORK OF VALUE FOR SUSTAINABLE COMPUTING

The results from the interviews in the end turned out to be of limited importance for the development of the framework. What the projects thought or hoped their impact could or would be may or may not be correct, and a solid impact assessment could not be based on informants' statements. That said, some things still came to light that are worth addressing.

1) An assessment of Intention relies heavily on the statements of insiders and it is (as has been mentioned) an intrinsically hard problem to get at 'the true intention'. One way to address this issue could be to change this dimension to instead reflect how well a project can *account* for what it is trying to achieve in terms of sustainability, rather than just trusting statements and claims about intentions. This means that an assessment could be made according to how well a project can account for *What* they are doing, *How* they are doing it, and *Why* they are doing it. If a project is seeking to "save the planet", it needs to be able to back this up with an accurate and fact-based account of *how* the project can save the planet and by *what means*. The hard-to-get-at intention is, for the purpose of the framework, less important than being able to account for what is done (and why). We thus changed the dimension hitherto called *Intention* to *Credibility*. A good accounting *will be* credible and a bad will not be credible and so the rating system need to be changed accordingly, e.g. credibility can be *high*, *medium* or *low*. It should be noted that it is not unlikely that a credible project is more likely to actually *have* an impact, but a credible project in terms of sustainability does not have to be probable. Rather we think of credibility as "believable" and possible to put your faith in in terms of sustainability.

2) The Likelihood of Impact as presented in the framework is a bit redundant and the essence of what we are trying to achieve is captured by just using three easier-to-measure criteria, namely *Relative Advantage* (compared to other products/services of similar kinds), *Compatibility* (is it consistent with the need, values, and experiences of potential users?) and *Complexity* (is a product/service perceived as difficulty to understand or use?). We want the framework to be useful in practice and therefore rejected additional criteria such as *Trialability* (is it available for testing before committing?) and *Observability* (is the result of using it visible to others?).

3) Taking this into consideration we end up with the final design of a Sustainable Computing Evaluation Framework (SCEF) as shown in Figure 4.

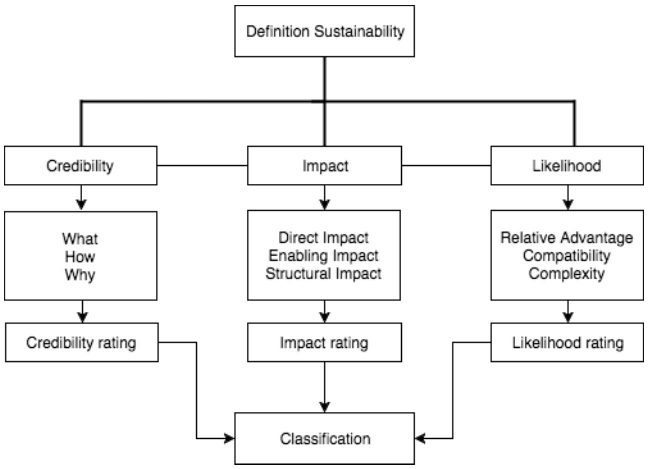

Figure 4: Sustainable Computing Evaluation Framework (SCEF)

7 DISCUSSION

We have attempted to make Toyama's taxonomy [19] more practically useful and have developed it through two revisions. First we added a protocol to each of the three dimensions and this resulted in a revised version that we chose to refer to as a "framework" instead of taxonomy. The protocol for the framework makes the process of classifying a computing system more precise. We then shaped a study around the framework and interviewed representatives from four computing projects pertaining to sustainability. The results of the interviews were used to revise the framework and we have chosen to call the revised framework the Sustainable Computing Evaluation Framework (SCEF, see Figure 4 above). We agree with [19] that intentions are important, but actually assessing intentions is not feasible and reframing this criterion in terms of *Credibility* seems to be the better choice. Several new concepts have been introduced that were not present in Toyamas taxonomy but we have yet to in greater detail clarify how the SCEF is to be used. We will now do so using *ChargeCar* as a test case.

7.1 Is ChargeCar a Sustainable Computing Project?

Impact: The Direct Impact of the larger ChargeCar system comes from the use of raw materials for producing the smart power socket and the energy needed to run the service (e.g. to charge cars). ChargeCar's service enables process optimization since it allows six cars to connect to one power socket where previously only one car could do so. ChargeCar argues that this will enable a greater diffusion of electrical vehicles, something that would count as Structural Impact, yet it is hard to say if ChargeCar will be part of important factors that would cause such a diffusion. Assuming that ChargeCar will contribute to a diffusion of electrical vehicles, the positive effect would be a decrease in the use of fossil fuels (assuming that each new electrical vehicle correlates with one less vehicle using fossil fuels). However, the same amount of resources are required for producing electrical cars as fossil fuel cars. All things considered, ChargeCar might decrease the rate of depletion of (non-renewables) resources, but does not thereby move things towards a sustainable equilibrium (where resource consumption is equal to resource replenishment). This means a *Level (-1)* rating.

Credibility: ChargeCar mainly provided an economic argument for their service when explaining what they were doing. ChargeCar did, from an economic point of view, give a good account of what they were doing and why their service could be attractive for users. They did not however give a credible account in terms of sustainability and our assessment indicates that sustainability was mainly used as a means for attracting potential customers rather then being central to the service they were offering. ChargeCar's Credibility (in terms of sustainability) is therefore *low*.

Likelihood of Impact: The ChargeCar system requires additional hardware (a new power socket) to be used, but their solution enabled users to bypass the inconvenience of going to a charging station, saving both time and effort. In terms of compatibility with current practices, most people do not yet drive electrical vehicles and this currently decreases the likelihood of impact. In terms of Likelihood of Impact, it is *possible* that ChargeCar will have an Impact.

Classification: This adds up to the classification: *Level(-1)/low/possible*. ChargeCar decreases resource use but does not move things toward an "equilibrium where the amount of resources being replenished is equal to the amount of resources being depleted" [19]. ChargeCar's credibility in terms of sustainability was low but it is possible that they will have an impact.

7.2 How is the SCEF Useful?

SCEF adds needed depth to the evaluation process, but some problems still remain. Any assessment will depend on where system boundaries are drawn and this is not necessarily an easy task [15]. *EnergyViz* hoped that people would change their habits and behaviors by providing correct and useful feedback to end users. Should the system boundaries be drawn around behavior specifically connected to energy consumption or also include more general behavior? A user might decrease the amount of energy used, but then use the money saved in other less sustainable ways. This phenomenon not specific to this framework but rather constitutes a problem for any impact assessment that has to deal with "rebound" and "backfire" effects [7], but it is nevertheless something that future work with the SCEF should take into account.

It is furthermore a tough criterion to discuss Impact in terms of resource consumption and replenishment. Even if resource use is what sustainability comes down to, no computing project will ever "replenish" resources in the way it has been defined in the taxonomy. Computing systems will (at least in a foreseeable future) depend on the use of non-renewable resources and *no* computing system can be considered sustainable in this respect. It is however possible to make distinctions between more and less unsustainable systems with the help of the SCEF framework. We have here provided a protocol for evaluating the sustainability (or not) of computing systems, but any assessment will still depend on judgment calls, and a clearly defined method for assessing each of the dimensions is needed.

7.3 Future Work

We have not explicitly discussed whether an SCEF assessment should be (primarily) quantitative or qualitative. In describing how the SCEF can be used, we propose two ways (that possibly can be combined) to proceed working with the framework. The first is to give a narrative account that leads up to a grade. This could act as a "legal" record that can be referred to by future evaluations to determine what the precedent is. The second is to develop a quantitative SCEF "index", for example by rating each of the three dimensions (Impact, Credibility, Likelihood) on a scale from 1 (poor) to 7 (excellent). These values could then be weighed together and form a SCEF Sustainability Index according to SI = x * Impact + y * Credibility + z * Likelihood.

Future work on developing the SCEF should furthermore consider how the scope of a project (in space and time) affects assessment and how this can be integrated into the evaluation process. Questions such as "*who is the system is designed for?*" [17], "*at what cost?*" [18] and *"is the technological intervention at all needed?*" [1] can be useful in guiding such considerations.

7.4 Are Frameworks Useful?

Both Toyama [19] and we have put effort into thinking about and shaping a taxonomy (or a framework) for determining if, or to what extent computing projects can be deemed sustainable (or not). But is our framework - or indeed any framework - useful? Is it important to have frameworks? Do they make a difference? We have assumed so, but developing a framework has also led to some unexpected complications. Building a framework is complicated and the right categories for evaluation do not come naturally. It should also be kept in mind that creating a framework is only a proximal goal where the ultimate goal is to achieve a sustainable society. We hope that the proposed Sustainable Computing Evaluation Framework (SCEF) can be useful and hence contribute to the development of sustainable computing (or at least to less unsustainable computing).

8 CONCLUSIONS

Toyama [19] proposed a taxonomy for evaluating which computing projects are sustainable and which are not. In this paper we have used Toyamas taxonomy as a foundation for the development of a framework that is more operationally useful. We have revised the framework twice and described the process of doing so in detail. The final result of this work is the development of a Sustainable Computing Evaluation Framework (SCEF, see Figure 4). We conclude that although we believe that our developed version is an improvement, there is still room for further adjustments and we suggest that future work with the SCEF should consider:

- Revising the definition of sustainability that is used in the SCEF into make it more tractable for computing projects.
- Add a quantitative aspect to the assessment process and/or consider how this can be combined with a qualitative approach.
- Elaborate on the Credibility dimension and consider how the scope of a computing project (in space and time) affects the assessment and how this can be integrated into the evaluation process.

ACKNOWLEDGMENTS

This research has been partially funded by the Swedish Energy Agency. We would furthermore like to thank Kentaro Toyama for his input early in the process and the reviewers of this paper.

REFERENCES

[1] Baumer, E. P., & Silberman, M. (2011, May). When the implication is not to design (technology). In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2271-2274). ACM.
[2] Brown, B. J., Hanson, M. E., Liverman, D. M., & Merideth, R. W. (1987). Global sustainability: toward definition. *Environmental management*, 11(6), 713-719.
[3] Daly, H. E. (1977). *Steady-state economics*. San Francisco: W. H. Freeman.
[4] Daly, H. E. (1990). Toward some operational principles of sustainable development. *Ecological economics*, 2(1), 1-6.
[5] Dobson, A. (1996). Environment sustainabilities: An analysis and a typology. *Environmental Politics* 5(3): 401-428.
[6] Dobson, A. (2007). *Green political thought* (4th edition). Routledge
[7] Druckman, A., Chitnis, M., Sorrell, S., & Jackson, T. (2011). Missing carbon reductions? Exploring rebound and backfire effects in UK households. *Energy Policy*, *39*(6), 3572-3581.
[8] Elkington, J. (1998). Partnerships from cannibals with forks: The triple bottom line of 21st-century business. *Environmental Quality Management*, *8*(1), 37-51.
[9] Global e-sustainability initiative. (2015). *"SMARTer 2030: ICT Solutions for the 21st Century"*
[10] Hilty, L. M. (2008). *Information technology and sustainability. Essays on the Relationship between ICT and Sustainable Development. Books on Demand GmbH, Norderstedt.*
[11] Hilty, L., Lohmann, W., & Huang, E. (2011). Sustainability and ICT—an overview of the field. *POLITEIA*, 27(104), 13-28
[12] Hilty, L., & Aebischer, B. (2015). ICT for Sustainability: An Emerging Research Field. In *ICT Innovations for Sustainability* (pp. 3-37). Springer International Publishing.
[13] MacArthur, E. (2012). *Towards the Circular Economy: An economic and business rationale for an accelerated transition*. Ellen MacArthur Foundation.
[14] Pargman, D., & Raghavan, B. (2014). Rethinking sustainability in computing: From buzzword to non-negotiable limits. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction* (pp. 638-647). ACM.
[15] Pargman, D., Ahlsén, E., & Engelbert, C. (2016). Designing for sustainability: Breakthrough or suboptimisation?. In *4th International Conference ICT for Sustainability (ICT4S)*. Atlantis Press.
[16] Rogers, E. M. (2003). *Diffusion of innovations*. Simon and Schuster.
[17] Strengers, Y. (2014). Smart energy in everyday life: Are you designing for resource man?. *interactions*, *21*(4), 24-31.
[18] Tainter, J. A. (2006). Social complexity and sustainability. *ecological complexity*, *3*(2), 91-103.
[19] Toyama, K. (2015). Preliminary thoughts on a taxonomy of value for sustainable computing. *First Monday*, 20(8).
[20] Wackernagel, M., & Rees, W. (1998). *Our ecological footprint: reducing human impact on the earth* (No. 9). New Society Publishers.
[21] Walldius, Å., Gulliksen, J., & Sundblad, Y. (2015). Revisiting the UsersAward programme from a value sensitive design perspective. In *Proceedings of The Fifth Decennial Aarhus Conference on Critical Alternatives* (pp. 1-4). Aarhus University Press.