Let us not put all our eggs in one basket

Florence Maraninchi

Let Us Not Put All Our Eggs in One Basket: Towards New Research Directions in Computer Science

Florence Maraninchi, University Grenoble Alpes

March 15, 2022

Our colleagues at IPCC and IPBES have been telling us for years that the situation is serious. 2021 has seen both the publication of the 6th IPPC report, and dramatic illustrations of the impacts of climate change. In all disciplines researchers and teachers face the question: what can you do in your professional life? If you search the internet for occurrences of "carbon neutral university", you will find a long list of declarations by universities worldwide, claiming they will be carbon neutral by 2030 or 2040. I will not discuss here whether carbon-neutrality objectives are feasible or even make sense at all (see [1]). I take this series of declarations as a symptom that the academic world is hopefully starting to take scientific results seriously, at least concerning the impact of our work organizations.

In computer science, several personalities have started questioning our peculiar organization that gives an important role to conferences [2], advocating for a massive change in how research is made and disseminated. Funders also have a significant impact [3]. As far as I am concerned, I stopped flying completely, and that is the least I can do, having flown quite a lot in the past 30 years.

But when I ask myself "what should I do?", when my students ask "are we part of the solution, or part of the problem?", I also look at my research and teaching topics, and I feel compelled to question the contributions of these topics to the development and impacts of the digital world as a whole. It is tempting to look at the positive impacts only. The public discourses tend to present the "digital transition" as a necessary and non-questionable solution to the needed "ecological transition". Our research community has the responsibility to consider several hypotheses, including one in which the digital world is part of the problem.

A tale of three futures

Let me now tell you a short tale meant to let our imagination escape the very pregnant determinism of tech discourses, at least for a few minutes. In 2005, I had a very simple mobile phone allowing to place and receive calls (almost)

everywhere, and which needed to be charged once a week. Telephone booths were still available in urban or rural areas. I am now one of at least one billion people carrying an always-connected always-on portable computer in our pocket, and if we really use all of its functions, we need to charge it twice a day. Telephone booths have disappeared completely. Cafes all over the world advertise the availability of electric plugs and free WiFi to attract a crowd of connection-hungry customers. You can charge your phone by practicing on a static bike while waiting at airports, and you can carry a solar panel on your backpack for a two-day hike. GPS and maps are an example of functions that were already available on dedicated devices anterior to smartphones, and have migrated to smartphones thanks to the versatility of this type of platform. Entirely new functions have appeared thanks to 24/7 connectivity, for instance platforms like Uber.

What happened between 2005 and 2021? There is absolutely no doubt that huge progress has been made on several key points: the technology of batteries has improved; the hardware architecture and the operating systems have been enriched with sophisticated mechanisms to optimize energy consumption; the capacity of memories has increased a lot; new underwater cables and optical fibers have been installed, 4G and 5G have been deployed; etc. But what about the overall environmental impacts of this growing infrastructure and the huge number of short-lived devices connected to it, or the indirect impacts on other sectors?

Let us imagine for one minute that we are back in 2005, doing our job of computer scientists, optimizing hardware and software. What futures did we envision? Future 1, in which our simple phones, functionally unchanged, would need to be charged once a month only, thanks to the improvements of batteries, software and hardware? Or future 2, in which the one-week charging period would be preserved, and as many new functions packed in the device as made possible by those improvements? Could we have imagined future 3, i.e., what we have now? The huge improvements of all aspects of the digital world have been accompanied by massive rebound effects [4], both direct and indirect. The fact that futures 1 and 2 were very unlikely to emerge because there would have been no economic incentives for such massive improvements, without an expected market increase, i.e., a bet on the rebound effects, makes the path followed between 2005 and 2021 a quite slippery slope. This phenomenon cannot be explained by technological arguments only. When we are working on optimizations of digital systems now, are we not in the position of our own selves 15 years ago, believing we were working for futures 1 or 2, but allowing future 3 instead?

Should we try to avoid a new future 3 and if yes, how?

Evaluating the total environmental impacts of the digital world is a complex task. According to the meta study [5], the greenhouse gases emissions of the digital world account for 1.8 to 3.9% of total emissions and are likely to increase. Arguably, compensating those impacts by corresponding cuts in the emissions of other — non-digital — sectors, would require such profound and quick transformations that it might not be feasible.

The moral of the story, put in a provocative form, could be: If there is a single example in the history of computing, where a particular optimization has not been accompanied by massive direct and indirect rebound effects, then we should study it extensively, from various points of view: technological, economical, sociological, etc., in order to try and reproduce it. If there is no such example, then we should stop believing that optimizations always help reducing environmental impacts.

When we start thinking of what it would take to avoid rebound effects and keep the impacts of the digital world within certain *limits*, at least two types of arguments are common: individual ethics and self-limitations, or regulations designed collectively. Both imply *choices and priorities*.

I personally think that, in front of climate change dramatic consequences, 8K videos, connected fridges, cloud-dependent home automation, cashierless retail stores, autonomous vehicles, smart shoes, the metaverse, web3 and NFTs are at best helpless and misdirected innovations, at worst and most probably, harmful. Other technologies, like high-tech medicine, may be useful, but concern the happy few only.

Whatever our personal opinion, as computer scientists we can start exploring the notion of limits even if we do not agree on the moral judgments related to the choice of those limits. We can even explore the notion of limits without being convinced that there should be limits in the first place, just because this is a fascinating territory of undone science [6]. How to stay within *limits* has become a scientific and technical problem that it is little addressed.

Towards new research directions in CS: Limits as First-Class Citizens

Aside green-IT, which deals with optimizations of digital systems, and green-by-IT, in which IT is used to reduce the impact of some non-digital sectors, avoiding the slippery slope of future 3 requires that we also work on an entirely new topic: *limited-by-construction IT*. The recently created series of conferences LIMITS [7] or the notion of Collapse Informatics [8], advocate for a digital world that deals with planetary limits, or may survive collapse scenarios.

When it comes to designing and developing computer systems, thinking in terms of *limits* requires a paradigm shift. We can start by highlighting the

implicit *anti-limits* most of the digital systems of our everyday life are based on. An anti-limit is both a promise and a deliberate hypothesis that resources will grow as needed. For instance, there are obvious anti-limits if a digital system:

1. Requires an increasing amount of resource globally (unlimited number of cryptocurrencies relying on proof-of-work, space, or bandwidth, ...)
2. Promises immediate service delivery, whatever the number of clients and usages (most of the cloud services)
3. Promises unlimited storage in both space and time (Gmail)
4. Assumes availability of some hardware, software and vendor cloud forever (some home automation devices)
5. Is designed to allow for unlimited functional extensions
6. Bets on the availability of a more efficient machine, soon
7. Needs more users or an increased usage per user to be profitable

Most of these examples are clearly rooted in economical choices, but thinking without limits has become so tightly knitted with the very principles of technical solutions, that in some cases it could be hard to continue delivering solutions, should environmental, (geo-) political or social constraints impose restrictions on the development of the digital world.

So what can we do? Having spent most of my 30-year carrier working on critical embedded real-time systems, I am used to languages and tools meant to determine the worst-case execution time, and the amount of memory needed by a program, before deployment: limits are part of the specification, and a very stringent constraint for the implementation. Other sources of inspiration include Gemini [9], which is designed *"to be difficult to extend in the future"*, while in any software engineering course, "extensibility" is presented as a desirable property. According to Wikipedia, it is *"a software engineering and systems design principle that provides for future growth"*. Designing systems that are not scalable, on purpose, is one way to keep limits in mind. Designing for intermittent resources or user quotas is another: *"A solar-powered website, which means it sometimes goes offline"* is presented in [10]. The ultimate limit, as addressed by collapse informatics, is: what if we stopped manufacturing new hardware?

Our discipline may need a radical approach, redesigning the digital world from scratch with specifications based on explicit hardware and software limits. If something is not feasible without assuming that some resource will grow as needed, then it should be considered as infeasible.

Let us take some eggs out of the good old optimizing basket

Even if we are not all convinced that optimizations cannot win over rebound effects and that we should therefore impose limits, even if we do not agree on where the limits should be, it would be a good idea not to put all our eggs in one basket. We should devote some research to the selection and preser-

vation of a somewhat minimal, robust, limited-by-construction digital world, and we should teach it. Asking which computer systems can still be designed and maintained, if we cannot count on the unlimited growth of the hardware and the infrastructure, leads to very intellectually challenging research topics. Moreover it is our responsibility to provide the necessary scientific background to the legitimate questions on technological choices that should be possible in democratic contexts.

1 Acknowledgments

I thank N. Halbwachs, B. Tourancheau, F. Berthoud, J. Combaz, S. Bouveret, M. Périn, and the anonymous reviewers for their thoughtful comments and suggestions.

References

[1] J Dyke, R Watson, and W Knorr. Climate scientists: concept of net zero is a dangerous trap. *The Conversation*, 2021.

[2] Moshe Y Vardi. Publish and perish. *Communications of the ACM*, 63(1):7–7, 2019.

[3] Teun Bousema, Leonard Burtscher, Ronald P van Rij, Didier Barret, and Kate Whitfield. The critical role of funders in shrinking the carbon footprint of research. *The Lancet Planetary Health*, 6(1):e4–e6, 2022.

[4] Wikipedia contributors. Rebound effect (conservation) — Wikipedia, the free encyclopedia. `https://en.wikipedia.org/w/index.php?title=Rebound_effect_(conservation)&oldid=1055996937`, 2021. [Online; accessed 13-January-2022].

[5] Charlotte Freitag, Mike Berners-Lee, Kelly Widdicks, Bran Knowles, Gordon S Blair, and Adrian Friday. The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations. *Patterns*, 2(9):100340, 2021. `https://www.sciencedirect.com/science/article/pii/S2666389921001884`.

[6] David J Hess. *Undone science: Social movements, mobilized publics, and industrial transitions*. MIT Press, 2016.

[7] Computing within limits workshop series. `https://limits.pubpub.org/`.

[8] Bill Tomlinson, Eli Blevis, Bonnie Nardi, Donald J Patterson, M SIX Silberman, and Yue Pan. Collapse informatics and practice: Theory, method, and design. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 20(4):1–26, 2013.

[9] Project gemini. `https://gemini.circumlunar.space/`.

[10] Roel Roscam Abbing. 'this is a solar-powered website, which means it sometimes goes offline': a design inquiry into degrowth and ict. In *Workshop on Computing within Limits*, 6 2021. `https://limits.pubpub.org/pub/lecuxefc`.