

The World is Not Enough: Growing Waste in HPC-enabled Academic Practice

Carolynne Lord
UK Centre for Ecology & Hydrology
Lancaster, United Kingdom
School of Computing and
Communications
Lancaster University
Lancaster, United Kingdom
carlor@ceh.ac.uk

Adrian Friday
School of Computing and
Communications
Lancaster University
Lancaster, United Kingdom
a.friday@lancaster.ac.uk

Adrian Jackson
EPCC
The University of Edinburgh
Edinburgh, United Kingdom
a.jackson@epcc.ed.ac.uk

Caroline Bird
School of Computer Science
University of Bristol
Bristol, United Kingdom
caroline.bird@bristol.ac.uk

Chris Preist
School of Computer Science
University of Bristol
Bristol, United Kingdom
chris.preist@bristol.ac.uk

Simon Lambert
Scientific Computing Department
Science and Technology Facilities
Council (STFC)
Didcot, United Kingdom
simon.lambert@stfc.ac.uk

Gabin Kayumbi
The Alan Turing Institute
London, United Kingdom
gkayumbi@turing.ac.uk

Kelly Widdicks
UK Centre for Ecology & Hydrology
Lancaster, United Kingdom
kelwid@ceh.ac.uk

Abstract

Most research depends to some extent on technologies and computational infrastructures including, and perhaps especially, HCI. Despite the noted environmental impacts associated with information communication technology (ICT) globally, to date little consideration has been given as to how to limit the impact of research and innovation processes themselves. Working to understand the technical and cultural drivers of this impact within the specific but resource-intensive domain of High Performance Computing (HPC), we conducted 25 interviews with academic researchers, providers, funders, and commissioners of HPC. We find intersecting socio-cultural and technical dimensions that link to research institutions like conferences, funders, and universities that reinforce and embed, rather than challenge, expectations of growth and waste. At a time when large scale cloud systems, generative AI and ever larger models are multiplying, we argue to de-escalate demand for computing, aiming for more moderate, responsible and meaningful use of computational infrastructures—including within HCI itself.

CCS Concepts

• **Social and professional topics** → **Sustainability**; • **Human-centered computing** → **Human computer interaction (HCI)**.

Keywords

sustainable HCI, HPC, growth, waste, socio-cultural drivers

ACM Reference Format:

Carolynne Lord, Adrian Friday, Adrian Jackson, Caroline Bird, Chris Preist, Simon Lambert, Gabin Kayumbi, and Kelly Widdicks. 2025. The World is Not Enough: Growing Waste in HPC-enabled Academic Practice. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26–May 01, 2025, Yokohama, Japan. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3706598.3713919>

1 Introduction

The Paris Agreement, ratified by 196 countries in 2016, committed nations to emission cuts of 50% on 1990 levels by 2030 to stay below 1.5°C of global warming. The first global stocktake of the Paris Agreement concluded at the UN Climate Change Conference (COP28) in December 2023, resulting in the decision that we are *not on track* to meet the required levels [53]. The relatively modest sounding 1.5°C detracts from the already severe impacts that are manifest from climate change. These include, but are not limited to, more frequent forest fires, desertification, extreme weather events, threats to biodiversity, heat related threats to health, and countries already seeing major impacts from sea level rise [72]. More alarming still, each degree of warming puts us in ever greater risk of hitting likely irreversible tipping points that would accelerate these effects.

While direct and indirect impacts of ICTs are only some of the many intertwined factors bound up in this problem, there is little question that digitalisation, demand for data and computational services, and the materials and resources that undergird this, continue to accelerate [23]. The impacts of this are of course not equally felt, raising questions of fairness, and social and environmental justice—not least globally, since the populations largely benefiting from such technologies are not those directly dealing with the extraction of minerals, the growing mountains of e-waste, and the toxification of environment associated with it [67].



This work is licensed under a Creative Commons Attribution 4.0 International License.
CHI '25, Yokohama, Japan
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1394-1/25/04
<https://doi.org/10.1145/3706598.3713919>

The depth and scale of this problem is of no surprise to many working in HCI [21]. Since the birth of ‘Sustainable HCI’ or SHCI in 2000s [29], researchers have pointed to systemic issues raised by the widespread adoption of ICT: arguing for ‘Sustainable Interaction Design’ in 2007 [5]; articulating the need for limits [52]; changing the ‘cornucopian’ design paradigm [58]; and the need to avoid a limited techno-solutionism framing of sustainability [24]. This includes the need to recognise the political agency of design [15], and even the explicit targeting of design to influence policy and governance [9] towards sustainability and post-growth HCI [67].

Despite these critical perspectives, there has been relatively little attention paid within HCI to the (un)sustainable growth and waste of computation within specific research communities. Digital methods such as very large simulations, modelling, deep learning/generative AI, and visual rendering are increasingly found, but the sustainability of relying upon them often goes unconsidered. This has even been evident even within HCI with, for instance, the use of AI for sketching systems [43], the application of machine learning to Twitter communities [37], and the use of Large Language Models (LLMs) to construct user personas [63] being just some examples from the 2024 CHI proceedings. Computationally intensive research requires the means to conduct it—the tools of High Performance Computing (HPC). These range from personal workstations equipped with powerful GPUs, through to very large national supercomputer clusters such as the UK examples of the Advanced Research Computing High-End Resource (ARCHER) (hosted at Edinburgh University) and Isambard-AI (at Bristol University), as well as international facilities like those at CERN or US National Laboratories.

HPC has unquestionably enabled a new generation of science, transforming the potentials in exploring scientific questions. However, its environmental costs are significant. The issue with HPC stems not just from the environmental impacts that facilities are having now, but their potential future impacts as they embed and change what it is to do research. Over time, the process of embedding shapes what this infrastructure is, the culture that it is becoming a part of, and also the expectations associated with both [16, 76]. Where growing HPC use is meaningful, this has the potential to generate scientific value. But when forms of growth are wasteful, this will lead to additional unnecessary environmental impacts. It is crucial for forms of waste to be identified and intervened in as early as possible, before they become a part of the nature and expectations of academic practice and culture.

In this paper, we use interviews with academic researchers, providers, funders, and commissioners of HPC facilities, to map out the technical, systemic, and cultural drivers of growth and waste. We mainly use HPC and shared computing as our case study as its magnitude is especially visible and carbon intensive, with high quality historical accounting of its use. In doing this, we position academic practice as a crucial site of enquiry arguing that, often, research institutions like conferences, funders, and universities reinforce, rather than challenge, growth paradigms and escalating expectations for the level of compute needed to conduct research. We then offer a number of strategies for intervention in this space. Finally, we make the argument that research more widely is becoming increasingly interrelated with HPC: with research changing expectations and need for HPC, and HPC changing research. We

further argue the importance of understanding these challenges and how they are constructed, so that all disciplines depending on HPC can address their growing footprint whilst it is still possible to do so.

2 Background

The overall carbon footprint of ICT is globally significant [23]. Large-scale computing infrastructures accounted for an estimated 2% of the US electricity supply in 2016 [68], with a further 2% of emissions potentially attributable to their embodied carbon costs (e.g., creating and disposing of the hardware) [31]. Very large scale HPCs, such as the Fugaku system in Japan, require upwards of 30 megawatts (MW) to power and cool [47]. The latest Exascale system in the US, Aurora, already exceeds 38 MW [71]. As these differentiated figures suggest, the size of HPC systems matters. Size is linked to how much energy they consume; how much hardware, and thus, their embodied energy and materials cost; and for how continuously and cost effectively they can be used.

Analogously, recent figures suggest that data centres, another site of concentrated computing, are growing in number, scale, and resource use. The worldwide energy use of data centres is estimated to have grown from 153 terawatt-hours (TWh) in 2005 to 205–270 TWh by 2010; they account for 1% of global electricity consumption, with a 30% increase in a little over half a decade despite significant improvements in their efficiency [46]. More recent estimates place this at nearer 460 TWh in 2022, with demand estimated to grow to 1,000 TWh by 2026 [1]. The number of data centres has not only increased, but they are also getting larger and more energy dense—containing significantly more compute capability, with more hardware, and demanding more power [10]. This is partly driven by the collection and processing of ever more data. Predictions suggest that 175 trillion gigabytes of new data will be created globally by 2025 [11] requiring further networks for transmission, data storage and compute. The shift to cloud, and software and infrastructure as a service is no doubt a contributing factor.

Such infrastructures have enabled the emergence of large-scale machine learning models, which are noted to have significant environmental impacts. Figures from 2020 suggest that training a single LLM can be responsible for over 300,000kg of CO₂ equivalent emissions [12]. More recent estimations suggest that over 500 tons of CO₂ were used in training ChatGPT3 [61]. While this figure could be argued as not that significant given the large number of users it supports (using a model is less computationally demanding than training it), ChatGPT3 was quickly obsoleted. Plus, this says nothing of the wider market competition and race to develop other and more powerful models, and the hardware and energy this requires.¹ Some models have been noted to be increasing in scale faster than generational gains in hardware capability, raising concerns as to their rates of growth, and the potential future demand for suitably capable hardware [82]. Compute demand thus, *has rapidly grown and continues to grow* [42]. Whilst this growth clearly cannot be solely attributed to research alone, computing of

¹ AI data center energy consumption is forecast a compound annual growth rate of 44.7% until 2027, by data centre specialists IDC, <https://digitalinfrastructure.com/news/ai-driven-data-center-energy-consumption/>, accessed March 7, 2025.

all kinds has become an essential tool in many fields of scientific research and engineering [42] including within HCI [17, 39, 43].

Most attention to computing's environmental impacts focuses on energy demand and efficiency gain. For example, power usage effectiveness (PUE) is a ratio and important key performance indicator (KPI) for data centres [46]. But, overly focusing on energy and how efficiently it enables compute can be misleading. The complexities of where in the world the compute happens, where the energy comes from (and how sustainable that is), how it is transmitted and generated, and the growth in its total demand remain hidden. There is also the question of the ongoing demand for material and 'rare Earth' minerals [6], which all come with hidden social and environmental costs [14, 57]—especially when they are disposed of [54]. Collectively, ICT may directly or indirectly displace populations, contribute to mountains of toxifying e-waste and place significant burdens on energy infrastructures and natural resources [74]. This is especially troubling when it occurs in vulnerable environments where resources like water are already scarce. This growing impact has not escaped regulatory and industry attention. Sustainability related performance indicators are already leading to new metrics relating to water use efficiency, and several others under debate [2]—new regulations in Europe will soon mandate that large scale data centres in the EU disclose their efficiency and energy use, among other possible KPIs.

In this paper, we focus on the specific context of computationally intensive research, especially as it relates to HPC. The disciplines using HPC have expanded as computational power and capacity have increased and offered potential for doing virtually what was only possible previously through laborious lab-based studies and experiments. Centred in core physical sciences such as physics and chemistry, weather forecasting and climate science, we are now reaching the point where *most* research and innovation disciplines are not only increasingly digitalised, but also seeing the suffusion of sometimes highly computationally intensive tools, such as automated translation and transcription [22].

Like cloud data centres, HPC facilities exist at various scales. HPC systems are often classified into 'tiers' of super-compute capability; with 'Tier 0' being the highest capacity national or transnational systems (e.g., Fugaku in Japan, Frontier in the USA etc.). The tiers range from Tier 0 to Tier 3 university scale supercomputers (i.e., systems run purely for local access for researchers and students at one institution, or even one department within an institution). Tier 1 systems are nationally-leading centres and there tend to be a few per country. ARCHER2 in the UK is one such example. Tier 2 are traditionally regionally or specialised facilities (e.g., Isambard-AI for AI only applications), restricting the user base to defined institutions or subject areas. Whilst these categories are loosely defined and somewhat fuzzy, they do highlight the range of HPC types that researchers requiring computing resource have access to. They also demonstrate how resources can be general and for a wide user base, or restricted to local and/or discipline specific groups of users.

Our data for the ARCHER series of HPC illustrates the nature of workloads and growing demand for compute. ARCHER was among the most powerful HPC facilities in the UK whilst it was


active (2014–2020).² ARCHER had around 250,000 jobs run on it in its last full quarter of operation (quarter 4 2019). ARCHER2, its successor, saw a *near doubling* of jobs in the last quarter of 2022 to approximately 450,000.³ Whereas ARCHER required around 1.2 MW of power to operate, ARCHER2 uses between 3 and 4 MW, depending on the jobs being run. The proposed future UK Exascale system will require upwards of 30 MW to power. Figure 1a shows the number of daily active users for our example (data is available at [35]). ARCHER2 was commissioned in two phases: (1) around a sixth of the overall compute and storage capacity was installed for early access users (1,024 compute nodes, 131,072 CPU cores), followed by (2) an upgrade to 5,860 compute nodes (750,080 CPU cores). The *step increase* in usage (number of users) from late 2021 (green dashed line) onwards when a larger set of users were given early access, and from the full system opening (dashed red line) is clearly evident; reaching 200–300 daily active users. The demand for this new, larger, HPC system is proportional to the number of jobs run on the system (Figure 1b). These data suggest that as the capacity and CPU core count of ARCHER2 increased, so too *the scale of jobs also increased* to exploit more parallel cores at each node and also in terms of total availability of cores. Utilisation remains close to 100% on shared batch computing systems of this kind, irrespective of the efficiency of the executing code—*supply always grows to exceed demand*.

As these figures illustrate, there is always demand for more compute. New hardware generations offer more compute per Watt of direct energy [71]; but also increase the number of jobs run, as well as their individual complexity, scale, and resolution. New capacity offers the potential to 'go larger' in terms of data analysis within the same time, increasing expectations and practice in turn. It is always possible for the demand for computation to ratchet up to meet or exceed supply. This is known as the rebound effect, where improvements in efficiency lead to an overall growth in demand, rather than reducing consumption as might be intuitively expected. As Widdicks et al. [79] note, emissions from the ICT industry have risen from almost zero to their current level in the last 50 years due to the huge increases in demand and investment in new infrastructure, despite the incredible increases in capability and rates of efficiency of CPUs. Similar trends have been evident in many spaces where digitalisation has promised efficiency gains but has instead led to increases in carbon emissions [9]. The rebound effect has also been shown to apply specifically to HPC [41, 44, 81]; where increased efficiency has been accompanied by an increase in demand, with HPC use expanding into ever-wider areas of research, and involving the use of more complex models.


HCI has seen a body of work exploring energy use in everyday life and sustainability more generally [32]. Behaviour choice around energy use in the home was one of the early foci, but it was quickly recognised as being too limited a perspective [13, 70]. The focus in recent years has moved from the environmental impact of individuals' behavioural 'choices' [70], design without limits, or the 'cornucopian' paradigm [58], toward broader perspectives of inclusion, social and climate justice [3, 45], challenging HCI's close relationship to consumerism [28], its implication in global social

²<https://www.archer.ac.uk>, accessed March 7, 2025.

³<https://www.archer2.ac.uk>, accessed March 7, 2025.



(a) Daily active users on the ARCHER2 during the early access, upgrade, and production phases of the system



(b) Total compute job duration on the ARCHER2 over the same period

Figure 1: Active users and job duration on the ARCHER2 national supercomputer

injustice [67], and wider misplaced faith in limited technological solutions [24]. HCI’s so called ‘third wave’ recognises that computers exist at such scale and mediate so much of society’s actions that they should be studied contextually [7, 29], and indeed the wider power of HCI design for sustainability may be in recognising this [9, 15].

Whilst ICT’s use has been studied in other HCI contexts [3, 78], the use of HPC in academic research has yet to receive much needed attention. This is somewhat surprising as computing usage has

grown in all areas of academic research and industry. 7 out of 10 researchers report that computing was essential to their research as far back as 2014 [33]. The computing infrastructure used in academic practice can easily come to include services at global scales—massive data sets, large machine learning models, data centres and networks, and other forms of HPC. They represent a growing carbon and environmental intensity to academia, and to the products and services we might collectively influence or go on to design [5]. Beyond this, our proximity to, and potential power in, the practices of our peers, and the structures we are a part of, have great potential as a site of study.

The largely unquestioned scientific value of HPC means that it is crucial to find ways in which to distinguish between meaningful [80] and not-so meaningful forms of computation. Here, we leverage the concept of waste to better explore forms of growth whose value does not balance with its costs. The idea of e-waste has been leveraged at CHI previously [38, 75] to refer to the “[...] various forms of electric and electronic equipment that is loosely discarded, surplus, obsolete, or broken” [38] which include ICTs, but is not specific to them. Our interest in waste, however, goes further than the end-of-life of electronic equipment. In our conceptualisation, we build on Wyuts’ [83] definition of wasteful consumption as the disparity “[...] between the amount of resources required and the resources used for a consumption activity; it may manifest as taking more resources than needed or as [the] inefficient use of resources”. The conceptualisation of waste that guides our inquiry is, therefore, subjective and requires our participants’ judgements to balance the relative values associated with HPC with its environmental costs. Rather than focusing on utilisation efficiency, we ask, in what ways is HPC use and growth justified and justifiable in current and future scientific practice, given the wider context of the climate crisis? Though our starting position on waste is quite simplistic, we will build on this throughout the remainder of the paper.

A focus on understanding waste is of special relevance because of the power of embedding over time. Social science ideas of integration and embedding emphasise the concurrent and interacting influence of ‘materiality’ upon human activity [76]; each redefines the other, and changes the meaning of what it is to ‘do’, in this context, academic practice. We seek to understand the redefinition of academic practice and HPC today, to better understand the path of development that we are currently heading down. In doing so, we identify entry points to intervene in forms of unnecessary or unjustifiable HPC growth within research that must be tackled before they become locked-in.

3 Methodology

We set out to understand the landscape of compute intensive research for academic communities in the UK with a particular focus on HPC. This project was reviewed and approved by the Lancaster University ethics review board. All participants were provided with information regarding the nature of the study, and informed consent was collected before starting the research.

There are three parts to our enquiry: (1) we recruited participants for semi-structured interviews using a combination of convenience and snowball sampling; ensuring coverage of 4 groups of stakeholders: academic researchers using HPC, providers, funders, and

commissioners of HPC. (2) We conducted a workshop in which the authors read overlapping subsets of the coded interview transcripts to reach agreement on the important themes. We drew these together (3) for each stakeholder group, and developed causal loop diagrams (CLD) to illustrate the interrelationships between drivers of HPC growth and waste.

The research team identified the key stakeholder groups in UK academic HPC, drawing up an initial spreadsheet of known contacts covering three higher education institutions, and three research funding agencies. We contacted these for interviews, and also asked for recommendations for further participants (snowball sampling). We recruited academic researchers using HPC, HPC managers, IT network managers, research development representatives, research software engineers (RSEs), management/culture representatives (e.g., ‘meta-research’, senior academic management), and commissioners (e.g., research councils, university procurement).

Interview schedules were developed for each group but the broad areas of questioning were: (A) HPC use and provision; (B) sources of waste and inefficiency; (C) acquisition and upgrade, and; (D) sustainability and net zero. As captured by our definition above, defining sources of waste is a valiative judgement—we followed our participants’ interpretations. Interviews were conducted online or in person, and lasted from 30 to 120 minutes. Recordings were transcribed through Otter.ai. before being reviewed for accuracy, and then anonymised.⁴

25 interviews were conducted in total (Table 1). These were comprised of 12 academic researchers drawn from computer science, computational chemistry, climate science, plasma physics and radio astronomy; 8 interviewees from provider and support roles including HPC managers, RSEs and procurement staff; 4 stakeholders from research funding and commissioning bodies, and; 4 in management positions (with some overlap in roles). RSEs are a relatively new software developer role which aims to support the development and use of scientific computing. The sampling strategy used resulted in two institutional deep dives, A and B in Table 1, with the remainder of the participants being spread across seven others, providing insight to contrasting and comparative dynamics at other institutions.

The interview transcripts were thematically coded, with themes relating to the project’s aims and questions being identified. The two interviewing researchers broke these broader themes down through analysis, meeting to share and agree upon codes. To develop the codes further, the project workshop brought together the wider team, with expertise in HPC provision, research, and funding. Each member read a selection of data, and points of confluence and convergence were discussed. The insights were then collaboratively organised under different categories—for instance, charting all the forms of waste described. The themes and sub-themes were developed through this iterative process.

Once the important themes were identified, the team used CLDs [48] to capture the interrelations between the systemic drivers of growth and waste. CLDs are often used in systems thinking to explore the relationship of different parts of an interacting system [30] with ‘arcs’ showing reinforcing and inhibiting elements.

⁴We are aware of the irony of making use of AI software to transcribe our interviews, yet signal this as further evidence for the growth and embedding of potentially computationally expensive, yet clearly useful, tools in HCI.

The approach used to develop the CLDs was as follows:

- (1) Two of the researchers with expertise in System Dynamics modelling reviewed the findings, and identified phrases from the transcripts which either explicitly or implicitly described relevant ‘stock’ variables (e.g., ‘Level of HPC Demand’) and ‘flow’ relationships within the system (e.g., ‘increasing student use of HPC increases ‘HPC demand’’).⁵
- (2) They documented variables as nodes, and identified the reinforcing (+) or inhibiting (-) relationships between these as arcs. This was initially done on a whiteboard to allow for rapid merging, and cleaning of structure, before moving to the loopy online tool.⁶
- (3) Drafts were reviewed to identify missing relationships based on the interviews and wider engagement with stakeholders. Some new nodes were introduced, some existing arcs were ‘split’, to provide a more nuanced understanding of the dynamics.
- (4) Nodes were colour coded to ease readability. Orange represents the core variables under investigation, whilst the other colours denote factors that emerged when we organised the analytical themes. Red nodes are socio-technological; green are technical and provision; purple are socio-organisational, and yellow are cultural.
- (5) The entire team then reviewed and iterated the CLDs further to ensure validity of findings, and the incorporation of all stakeholder perspectives.

Our findings are presented in two parts, first the interview data, before moving to the CLD representation of the wider system of growth and waste of HPC.

4 Findings

Drawing on our interviews across the four stakeholder groups, we zoom in on themes related to growth in demand for compute, and perceptions of HPC ‘waste’ demonstrating how these are embedded in wider academic culture and practice. To demonstrate the spread of perspectives, we identify the stakeholder group for each participant quote—U for user; P for provider/support; M for management, and; C for commissioning.

4.1 Growth: Escalating Demands for Compute

HPC facilities are varied, and both researchers and providers referred to different scales of facilities (from local to shared national resources) in common use. What is involved in their use and provision differed widely, and also within disciplines and communities of practice. In terms of funding, at the national and globally competitive scale, large infrastructure equipment bids and rolling investments provide large-scale facilities for whole communities of researchers to share. This is analogous to large-scale investments into hyperscale data centres being made by global technology companies [73]. At the institutional level, more modest scale shared research group or institutional facilities typically rely on per project and university level funding. These can be shaped and expanded

⁵‘Stock’ and ‘Flow’, and ‘Causal’ are terminology that exist because CLD originally was used to analyse technical systems. As we rely on qualitative data, we depict influential rather than causal relationships.

⁶<https://ncase.me/loopy/>

Table 1: Participants (ID, Role, Stakeholder Group and Institution)

ID	Role and Domain	Stakeholder Group	Institution
P1	Computer Scientist	user	A
P2	Climate Modeller	user	A
P3	Director of Advanced Computing	management/provider	A
P4	Lecturer, Computer Vision and AI	user	B
P5	Plasma Physicist	user	C
P6	Research Software Engineer	provider/support	C
P7	Scientific Computing/Particle Physicist	user	A
P8	Radio Astronomy/Computational Imaging	user	D
P9	HPC Chief Technical Officer	management/commissioning	E
P10	IT Procurement Manager	provider	A
P11	Computational Chemistry (RSE)	user/support	C
P12	Research Development Officer	provider/support	B
P13	Head of Storage and Virtualisation	provider/support	B
P14	Professor of Particle Physics	management/user	A
P15	Lecturer, Natural Language Processing	user	B
P16	HPC Manager	provider/support	B
P17	Senior Data Scientist	user	B
P18	Research Council HPC Senior Programme Manager	commissioning	F
P19	Head of Research Computing	provider/support	G
P20	Senior Research Associate, Computational Chemistry	user	A
P21	IT Network Manager	provider/support	A
P22	Research Council Data and Infrastructures Team	commissioning	H
P23	Professor, Metaresearch	management	A
P24	Senior Portfolio Manager, Digital Research Infrastructure	commissioning	I
P25	Associate Professor, Systems and Synthetic Biology	user	A

by the the start of new projects or the arrival of new researchers, and their associated expectations.

4.1.1 Availability and Changing Expectations. The computational scale of the problem being investigated matters. At different points, computation exceeds that which can be conveniently processed locally (i.e., workstation and laptop GPUs), or even through institutional level facilities. Some participants, like P6 (P), leveraged different facilities for different aspects of their work, saying “*I do some work on office computers, we have small servers in a server farm that I use for slightly bigger things, we have a Tier 3 machine here at [our institution...] and then I use ARCHER2*”.

Many participants recognised that the need for computation in research had increased or “[...] was] *ramping up*” (P7: U) in recent years. This was not always due to the scale of the problem being investigated—with social factors changing the idea of how much compute was sufficient, and when it should be used. This included domains where compute had not been so typical: “[...] *there are lots of good reasons to think that compute usage will increase and that the social sciences will be dragged into it*” (P23: M).

One such social influence was an ‘escalator of expectation’ over time. With, for example, a direct relationship being evident between research and teaching. P15 (U) for instance, told us “*Students read papers, right? That’s part of the work. So, they investigate the field, [...] and the majority of those papers are results from like, big models. And when I say big models, that’s models that require a lot of computational power*”. The increasing expectation of computation in

academic practice shaped students’ perceptions about what they would be taught, and what their own research projects should include; with them wanting to embrace, or being fascinated by, ‘the latest techniques’ or those most likely to be of value to their future careers—embedding further demand for computation.

The integration of computationally-intensive facilities to support different degree schemes, university courses, or research projects, was seen to have grown the demand for computation. This even led to ‘demand peaks’ and congestion effects, especially around deadlines. For instance, the timing of coursework or paper deadlines led to a temporal patterning to HPC use. Commercial cloud, also known as elastic compute because of the ability of the system to automatically provision additional virtual resources, was seen as a way to mitigate this. P20 (U) said: “[...] *you can simply increase or decrease the ‘number of machines’ you have available depending on how many jobs you want to run*”. ‘Pay per use’ services, however, come at a cost. One which may be hard for organisations to anticipate and control, especially in academic settings. Expanding the capacity of facilities was another means of attending to this issue; though this could also be problematic: “[...] *it’s more a case of the more we add, the more it would be used, there’s basically more demand than supply from that point of view*” (P16: P). Peaks in demand are especially challenging to navigate, then, as they drive the need for ‘burst capacity’ and facility expansion which ultimately further fuels patterns and levels of HPC demand, as seen with ARCHER2 (in Section 2).

4.1.2 Embedded Provision. More ‘micro’ scaled configurations were also in regular use, and often embedded in other processes, like the hiring of staff, and expectations around what is a ‘standard’ workstation. P4 (U): “So I have two new postdocs, and they will each need a workstation and GPU access [...] I’ll be buying a new workstation [...] And I’m guessing, an NVIDIA 4090 or possibly two”. Some institutions were trying to avoid this ‘wild west’ of equipment (e.g., ‘servers under desks’), but also were constrained by funding frameworks and eligible costs of research funders. “[To buy into institutional HPC] the rate starts off about 10,000 pounds [GBP] I think. And so that’s what we request that as equipment in a research proposal, and then that buys the node [...] when you ask for things over 10,000 pounds, especially in research councils, it’s classified as equipment, and then University has to match fund the shortfall of equipment” (P12: P). Different funding streams were therefore more or less desirable to the institution, due to the cost implications and especially overheads or matched costs liable as a result.

The academic researchers interviewed noted the wide range of applications which necessitated access to such facilities. Increasingly, this included the training of machine learning models. Given the wide availability of facilities, the decision of which scale often was linked to the circumstances and practices they would be involved in. The use of different facilities would, however, simultaneously enable and constrain ways of working: shaping what could be done, and what would be required to do it. Issues of control and availability were manifest, with some departments apparently keeping their own facilities running ‘under the radar’; P12 (P) told us “people used to cost in their own server with their personal grants. But they tried to stop that and said you had to go through [central procurement]. [Department] have their own like mini [HPC cluster] version [...] and I don’t know how well known that it is. It’s a bit of a_ they try and keep it secret”.

4.1.3 Barriers for Sharing Infrastructures. Shared institutional or higher level facilities were seen as more complex environments to set up and use. Barriers included permissions, funding policies (e.g., data privacy), user accounts, and the skills and knowledge required to navigate these constraints. P25 (U) told us of the importance of the skills and knowledge of the individual researcher in determining which facility would be used. “If your student or postdoc can do things, then it’s fine [...] various students have problems, [they] just install it on their laptop because of dependencies on see, various bits and if there it is a big piece of software, yeah [...] If they have to wait for someone to help them to set up their code and progress studies, it’s just time that is wasted”. P15 (U) reinforced this, noting “only the really good programmers [use the institutional HPC], they do complain that it’s a bit fiddly, but they’ll use it. Because it’s powerful in the end”.

4.1.4 Materiality and Cost. With particularly larger scale shared facilities (e.g., data centres, supercomputers), access to physical resources such as energy and water, competition for real estate, coupled with organisational and regulatory considerations (e.g., risk, security, data sovereignty, ownership, compliance) constrained where facilities could be put. This could lead to resource duplication. At different scales, however, similar complexes of funding and upgrade cycles, needs, and expectations fed into how much computation can be supported, and how it can be used. P13 (P) told us “we built the data centre with cutting edge technology at the

time. [...] that served us really, really well for 13 years, and we’ve just refreshed it. So it should last another 5–10 years [...] And the price. I mean, we do really, really well on our deals. If you’re in the centre of London, [...] and it’s a tiny little building, and somebody’s going to offer you 10 million pounds just for the land [...] it makes a very compelling argument to move to cloud. [...] Power [and] cooling in the middle of a city is much more difficult and complicated than it is when you’re [...] rural”. Here, the context in which the facility is being built is influenced by a whole range of factors (e.g., spatial, financial), few of which are technical.

4.2 Waste: Computational and Cultural

Our participants were asked to explain how they conceptualised waste in computation and facilities use. There are clear advantages to leveraging HPC resources, and in certain cases, participants reported that their work and use was meaningful and not wasteful. For instance, P8 (U) felt that she “was personally usually [...] mindful about how [she] used [resources]”, and that her use was appropriate. However, when we asked about ever reaching ‘net zero’, P4 (U) said: “[that’s] something that I will never achieve. There’s really no question about that”. It is clear that there is a carbon cost to advancing science. Participants also remarked, however, that the exact costs of their HPC use was not clear to them, with no information available to evaluate their work’s carbon emissions (P1 (U), P11 (U/P), P15 (U), P17 (U), P20 (U)).

4.2.1 Views on Efficiency. Others did feel there were ways, however, to make ‘better’ use of resources, and to leverage computational power more efficiently. Improving the efficiency of each computational task, and maximising the use of shared rather than individual resources were both seen as being appropriate ways of reducing waste. Shared facilities, for instance, were thought to ensure that power, heating, and ventilation were also shared and their costs amortised for more users—though P2 (U) noted that non-technical factors often shaped when, and in what ways, resources could be shared in practice. Projects, funders, and certain constraints on data (e.g., sensitive health or personally identifiable data) were seen as possible barriers to the use of shared facilities. Fears over security and privacy of sensitive data (P2 (U), P4 (U), P17 (U)) were especially important in this regard. Other forms of inefficiencies were also noted like the “duplication of platforms or resources [...]” (P23: M) with some facilities being underutilised given the inability—currently—to join up different platforms to enable tasks to freely migrate tasks to underutilised servers.

4.2.2 Providing Efficient Utilisation. For end users, power demand is often seen as a function of job runtime. For infrastructure providers, the concern extends beyond the dynamic power demands of a particular job; they must also consider the base server and cooling energy demand. “There are energy saving things we do try and do with the cluster as a whole [...] we say to the vendors, we have a fixed sum of money, what amount of compute can you supply for us. We don’t necessarily go for the fastest, because that’s not necessarily the most power effective. So we take into account the amount of compute, the price obviously, and the amount of power that’s consumed, assuming that we can fill the system 80 to 90% capacity” (P16: P). Supporting a certain ‘efficient’ level of utilisation in some sense

here drives expansion: “*the analogy I tend to use is [...] the motorway, what’s the average gap between cars. So if it’s late at night, there may be half a mile between cars. And then as the utilisation increases, [eventually] the[re’s] no space between them [...] bumper to bumper traffic jam, which is what you want to avoid [...] 80 to 90% tends to keep things flowing*” (P16: P).

4.2.3 Profligate Cultures. Though some of the forms of waste were technical, others more socio-cultural. These emanated from expectations from academic culture, with our participants questioning whether the computation needed to be done, or done at that scale, in the first place. The ‘publish or perish’ paradigm was repeatedly mentioned as requiring researchers to publish intensively throughout their careers, and ‘to be competitive’ in terms of their publication profiles. Research that was computationally intensive, even if redundant, was seen part and parcel of academic practice: this could be research which had already been done (P1: U), of little actual scientific value in terms of new knowledge (P23: M), or even ‘junk science’ (P4: U). P23 (M), for instance, remarked that researchers sometimes “*[did] research that doesn’t need doing because we already know the answer, and [they’re] just interested in publishing a paper*”. There is a ‘performativity’ associated to academic publishing which not only drives demand but escalates wasteful practices regarding compute. “*You do a heck of a lot more experimentation [for publishing] than for a real-world application [...] all of the experimentation that you’re doing [...] some of it is essential. A lot of it isn’t. But then again, if I want to keep my job, I better get published*” (P4: U). P20 (U) added that reviews provided a temporal pressure: “*the reviews come back with some questions [...] and you only have a certain amount of time to do it. Sometimes you really need the data next week*”. This could lead to researchers using bigger machines and more resources to meet these deadlines. Once resources had been funded, too, it was unusual that they would not be used: “*You go okay, so I want to use these resources, I’m going to pay for them upfront and then I’m going to use them until the quota is done. You’re probably going to use the whole lot*” (P9: M/C).

4.2.4 Externality of Expedience. When it came to waste reduction strategies, our participants noted that it was not just a simple case of avoiding resource or research duplication. The nature of research (i.e., as a mode of systematic inquiry) means that “*[sometimes] you need to redo [a simulation or run]. And that’s science, isn’t it?*” (P20: U). There is a clear relationship between experimentation and more simulations; though code can, sometimes, be made more efficient. This is not, however, always seen as a quick, or easy process. “*It’s always going to be the case with our research, some things might be quite obvious [...] but most of the time, it’s stuff you could spend three or four weeks looking for, and still not find*” (P2: U). The time pressures of academic practice often limited the time that could be spent in producing more efficient code: “*It’s kind of a time to solution that we’re interested in. You know, how quickly can we get that back?*” (P5: U). Beyond having the time needed to improve code, participants noted not necessarily having access to sufficient in house programming capability, nor the training in algorithm design or software engineering required. It is here that the increased importance of RSEs was noted, as specialist support that could be factored into project funding to aid researchers in their efforts to improve efficiency. These stakeholders could even be

leveraged to design more generalised tools. Current levels of skills and training means that it becomes important to, or it is assumed that researchers should, not ‘touch’ software toolkits. This is to avoid breaking ‘the established benchmark’, even if that benchmark is known to be outdated or inefficient in terms of its utilisation of modern hardware capabilities. “*Oh, this is what I got from, you know, the PhD student before me and this is what I’m going to use, [this] kind of way of thinking, it’s a lot of repetition of mistakes*” (P7: U).

4.2.5 Governance Vacuum. The dynamics uncovered suggest that the individual alone cannot attend to the problems of computational waste. However, participants noted little external pressure from academic conferences, journals, projects, funders, and institutions to work in less environmentally wasteful ways. The academic structures that mediated our researchers’ work even appeared to motivate a lack of sustainability, rather than to encourage it: “*we get very frustrated here, with the university about how much they demand we replace things [...] as soon as the server is out of warranty, they’re like you have to have a new one. And it’s like this computer works absolutely fine [...] often the university says, ‘Well, we’re striving to be a Top 50 World University. So, we always had the latest tech’. It doesn’t work very well with sustainability*” (P2: U). Provisioners, also, noted the endemic cycles of upgrading and refreshing, though this was seen more positively given their focus upon maximising value in the given space of a facility: “*It is one of the brilliant things where every sort of four or five years there’s an old saying called Moore’s Law where everything doubles in computing etc. So what tends to happen in IT speak is we grow to fill a room and then we refresh it, and shrink right back again to a corner and then we grow to fill the room [...] So I’ve been doing this for a very long time. And I’ve seen that happen numerous times*” (P13: P). A combination of factors, only a few of which are technical, combine and interact here to create a system of replacing over extending lifetime of facilities. Concerns over the security of data, facilities and devices, combined with priorities to ensure ‘cutting-edge’ technologies and global competitiveness, means that institutions are caught in a cycle of planned obsolescence [59]. This currently undermines strategies to limit waste and increase the lifespan of facilities and their attendant technologies.

Following climate emergency declarations and a growing awareness of the need to act, our participants noted how HPC policies stood apart and were increasingly inconsistent with other institutional sustainability efforts. For instance, for conferences: “*there’s lots of sort of schemes and policies around being conscious of our use of resources, and trying to be sustainable in our activities [like travel]. Generally, there’s nothing that’s sort of specifically directed at [...] resource usage*” (P5: U). When it comes to universities: “*we think about the plastic we use right? [...] we don’t think what happens when we press the button to launch a big simulation*” (P25: U). Given that the impacts of ICTs and related services are comparable, globally, to the aviation industry [23], this is a missed opportunity. Our participants noted that strategies could be developed to attend to and influence this issue: “*like a best [sustainable] paper award [...] would motivate a lot of people [...] or if one of the companies created something with net zero in mind [that could] motivate researchers*

to do something similar. It will become the norm” (P15: U). Leveraging the relation of academia to publishing and industry, and their relative power, was seen here as a possible driver of change.

4.3 A Systems Perspective

Stepping back from the context to consider the totality of the system, it is important to note that we have necessarily separated out the factors in order to describe them. However, it is clear that they are all highly, and tightly, interrelated, as illustrated in the CLD. Figure 2 is a CLD consisting of nodes representing concepts, linked by directed arcs representing the relationships between them.⁷ A ‘plus’ on an arc denotes a positively influencing relationship (reinforces), whilst a ‘minus’ depicts one that detracts (diminishes). While the CLD can be said to be somewhat reductive, we believe it illustrates well how HPC is embedded in a wider system of growth and waste—locking in the environmental, energy, and material footprints into the ongoing process and support of academic practice.

Nodes and arcs on the left of the diagram show advances within Computer Science result in new computational techniques which in turn drive advances within other disciplines. This results in a usage increase within a discipline, leading to increased expectation of use, including among students of that discipline (4.1.1), and in turn drives further advances. This is a classic reinforcing feedback loop, driving up demand for HPC over time. The right of the diagram show the dynamics driving the different sources of waste and potential mitigations identified in the interviews. The most straightforward conceptually is inefficient code (4.2.4) but the lack of researcher time and expertise means such code often persists unless investment in expertise such as RSEs takes place. Increased runs of simulations can be from increased experimentation: some is unnecessary and driven by strong cultural expectations enforced by, for instance, academic review practices (4.2.3). At the most extreme (4.2.3 P23), academic pressures and poor practices can lead to redundant and unnecessary research taking place.

Behind the scenes, and to some extent present in our findings, competition, digital products and services, as well as broader socio-economic factors drive global research, epistemology and the global economy further [50, 52, 67]. To attend to the issue of growth and waste in HPC-enabled academic practice, we must see beyond the context to the system, so that we can know how these factors relate and shape one another. The CLD demonstrates where some potential intervention levers for this emergent pattern of growth are. The project’s full report delves into further socio-cultural details, and has a more extensive collection of CLDs [4].

5 Intervention Points and Barriers to Net Zero

We have demonstrated how the growing use of HPC is shaped by various socio-cultural factors leading to, at times, wasteful forms of use and provision. Ultimately, growing compute is not always linked to scientific progress. According to our analysis, waste can emerge through code and infrastructures (i.e., a lack of efficiency, wasted or underutilised infrastructures), academic structures (i.e., projects, conferences, journals, institutions) and scientific culture

that drive and expand HPC facilities. We now consider some intervention points that could help to stem these forms of waste. Some of these measures can be taken on by the individual, but the more meaningful interventions we argue, are actions at institution, funder, and scientific culture scales [49]. We identify stakeholder groups and specific actions in italics to make our recommendations more evident. While we ground our discussion in the fairly specific context of HPC and scientific computing in the UK, we suspect similar dynamics exist in many domains given the prevalence of computational methods (as noted previously in HCI [8, 51, 65]).

5.1 Wasteful Code and Infrastructures

Our participants were clear that technical measures were needed to attend to growth and waste of existing HPC. A cornerstone of ‘green software engineering’ is the reduction of energy demand through improved program code efficiency [18]. The drive for ‘more efficient’ HPC hardware and software is deeply ingrained [27]. Researchers and RSEs in particular can *enhance code bases* to make better use of specialist hardware leading to higher utilisation by using these features to full potential. This was thought to reduce overall energy demand. RSEs, with their expertise in coding, could also be leveraged in the *development of shared tools, training and techniques* to help plug capability gaps to exploit this within research teams.

It was generally accepted that shared facilities were more efficient than individual or group HPC servers. Local bespoke resources were anticipated to be more lightly utilised, leading to wasted compute capacity, and weaker amortisation of the costs of facilities’ creation—similar efficiency arguments are made for moving from local private clouds to shared data centres [46]. Batch compute jobs could also be placed, potentially automatically, to lightly loaded systems [36]. Recent findings suggest that *HPC providers can run large multi core systems at lower than maximum CPU clock rates* to make a modest energy saving at the expense of slightly increased job runtime [66]. Given the changing marginal carbon intensity of energy generation, as seen increasingly for energy tariffs, Greenhouse Gas (GHG) emissions could also be reduced for non-time critical jobs by *shifting computation to ‘off peak times’ offered via ‘green job queues’ or to countries with cleaner energy mixes* [46]. Although, we have yet to see consideration of congestion effects, and other wider emergent implications of such strategies be systematically considered in such proposals.

It was repeatedly suggested that there was a lack of technical support and established practice for evaluating the environmental footprint of HPC use. This prevented meaningful comparisons between different computational techniques, whilst also preventing estimation of the lifetime costs of conducting particular research. Tools are being developed by the *HPC community* that could *make computational cost information routinely available and open to inspection*.⁸ While past S/HCI work challenges characterisations of users as rational actors lacking energy related information [55, 56], we believe feedback in this context may be more akin to software performance profiling, through which externality costs of computation could be instrumental in a process of software improvement or to evaluate alternative approaches as discussed above. Without

⁷The diagram is best used interactively to illustrate and animate how each factor links to, and influences, the others—see, <https://bit.ly/44QFhD9>. Note: we have redrawn the figure from the original loopy version for readability.

⁸c.f., <https://github.com/GreenAlgorithms/GreenAlgorithms4HPC>, accessed March 7, 2025.




Figure 2: Causal Loop Diagram (CLD): An overview of socio-cultural factors that influence HPC growth and waste

knowledge and accountability, it is hardly surprising that software is not being optimised for energy performance, and computational runs are not as parsimonious as they could be. It is crucial that any information provided is *complemented by effective training and agreed standardised assessment methods* to avoid potential for creative accounting.⁹ Funders and broader academic structures need to acknowledge and resource specialist programs and career paths for RSEs in HPC too, otherwise the important role and knowledge that RSEs have will not be available, especially to smaller projects that cannot afford them directly. Open source code and models were noted as being a way of increasing energy efficiency of code more routinely. Projects must be able to *commit funds to RSEs and time needs to be committed within projects* to develop more generalised code and models that live beyond any single project.

Strategies centred on promoting resource sharing, of course, have to be understood in context of what this means for project data and wider socio-cultural factors. Shared facilities can have their own disadvantages like long wait times and being less responsive or controllable. Also, ownership is linked to wider considerations of status and ownership such as control, ease, competition, and prestige. Institutional or funder policy, may also inadvertently constrain use of shared facilities, pushing researchers toward establishing their own. For example, through sometimes unwarranted or addressable concerns over data security, data management plans and institutional and funder policy.

Focusing on efficiency of computation alone is—ultimately—not a solution to the overall increase in demand [9, 69]. Especially if it is not accompanied by limits to growth through a more structural, and then cultural transformation [62]. Whilst increasing efficiencies lowers the computational (and probably indirect financial) costs of HPC use, it could also create possible rebound effects [26, 79]. This is evident in our ARCHER2 example (Figure 1b) whereby the increased efficiency of the system is outpaced by the step in growth in both users and the ambition of workloads. This rebound effect is also evident in HPC more generally [41, 44, 81]; with HPC use

expanding into ever-wider areas of research whilst also involving more complex computational models. It is hardly a leap of the imagination, then, that if researchers can do more computation in the time available, they will likely do so.

5.2 Wasteful Academic Structures

The need for addressing climate change in all aspects of society is clear [34]. Yet, the academic structures described by our participants did not recognise the need to include HPC, or the growth in demand for computation and computational facilities in their net zero strategies. Our participants already recognised that a technological, or efficiency-based, solution alone were not sufficient to deal with all forms of waste they perceived. And so, efforts must be made to go beyond this.

Academic structures (i.e., projects, funders, journals, conferences) must address the wider systemic factors driving growing waste in HPC use and provision. The embodied emissions from manufacturing (as well as wider social and environmental impacts of resource use like rare earth metal mining) are a significant contributor to HPC's many impacts [8], and these are not actively considered in current research practice, nor at the level of *academic policymaking*. This suggests the need to work within computational limits [52] and significantly *prolong the life of HPC*—scaling back on drivers of premature retirement and acceleration in HPC infrastructures.

Researchers and research managers could contribute towards an evidence base that allows the evaluation of runtime energy and carbon performance when studies are proposed or designed, to aid in addressing the 'informational deficit' around the impacts of computational work. Ultimately this would require a strategic focus on HPC's footprint at earlier research stages when proposing or even funding research. *Academic structures* would, however, need to *lift the significance of this information* by mandating inclusion as a *performance indicator* in the justification of resources during funding applications, or when papers are considered within peer-review processes. The technological approaches of efficiency, sharing, and information will have limited value without genuine efforts to embed sustainability considerations throughout all practices.

⁹Note several industries are already working on agreeing standards for assessing their own footprints, e.g., ITU <https://www.itu.int/rec/T-REC-L.1450/en> and <https://www.itu.int/rec/T-REC-L.1480/en>, accessed March 7, 2025.

There were times where we observed current structures encouraging growth and waste, rather than opposing it. For instance, financial systems were actually forcing spending at specific times of the year (e.g., end of financial year; regular HPC funding cycles) that could result in wasteful IT spend due to the ease, high value, and the normally fast supply chains of computing equipment. *Further research is needed to understand how policy and practice at different scales is driving growth in infrastructures* and encouraging these profligate practices. Whilst we have argued that care needs to be taken to understand the impact of already-existing policies, an important first step for *academic structures* is to *include HPC in their sustainability policies* [77]. Another is to begin to *incentivise meaningful use and value in research*. This could take the form of *conference awards for low-carbon papers* (as suggested by one of our participants), or *recognition being given when science is done within its limits*. With its background in understanding sustainability in different empirical contexts too, the *CHI community* could be working to *argue the importance and impact of not considering the growing embedding of HPC and computation in research*, in all forums we have access to.

5.3 Wasteful Culture

Ultimately, a wider cultural shift is required across science if we are to tackle the growing demand for HPC and computation [64]. Though we have intentionally separated wasteful code and academic structures from the culture in which they are embedded for presentation here, smaller scale interventions do not go far enough. HPC use is complicated by a range of socio-cultural factors (e.g., practices, expectations, institutions, policies) that shape and hold growth and waste in HPC in place. Interventions must go beyond individual researchers ‘caring’, ‘being conscious’ or ‘making good choices’; changes are required to make space, and embed, more environmentally conscious ways of ‘doing science’ at a structural level. We can encourage researchers to use shared resources or make better utilisation of HPC, but this will have little impact if funding bodies obligate the use of private facilities, or research institutions escalate its expansion and growth. The significance of environmental costs of computation must be raised systematically, *funders, conferences, journals and institutions must embed mandatory reporting of this* for it to have any meaningful effect. It is not enough, either, that only some communities or structures do this. For larger scale impacts to happen, *cross-cutting interventions* must happen *across science and academia*. Pressure to grow research disciplines requiring unsustainable techniques, and for institutions to remain globally competitive with one another infrastructurally, will lead to overall growth in environmental impact [25].

Shifting the culture of science is far more challenging and will require *cross-cutting efforts from policymakers* of different types and across scales. There is little policy oversight or buy-in currently in terms of net zero, especially in terms of HPC and computation use, from specific academic structures. In fact, there were times when key stakeholders were conversely encouraging more, not less, wasteful HPC practices. Close deadlines could lead to researchers intensively using HPC resources, the use of shared facilities was sometimes forbidden, and the pace of equipment turnover was being actively promoted. Dynamics of this type require changing

the way in which science is done, accountability, and top-down as well as bottom-up pressure, for the *academic community* to better consider and *balance the values of its policies and work with the potential harms*.

Addressing digital waste increasingly challenges what it is to do research, holding up a lens to uncomfortable truths such as unnecessary duplication of research and re-computation. Whilst we recognise that science itself depends on some necessary forms of duplication (e.g., replicability to validate results), the instances described by our participants call into question the sustainability of some methods that have already become deeply enculturated (e.g., redundant science). But conversations and appreciation of this is starting in many communities—our participants noted ongoing initiatives such as the *Findability, Accessibility, Interoperability, and Reusability* (FAIR) data principles [20]—with the goal to transform researchers’ data practices to reduce the likelihood of duplicating data, and the efforts involved in (re-)producing it. Though these policies are predominantly about the pursuit of an Open Science agenda, they reflect also an understanding that every instance of HPC use should ‘count’ in the pursuit of new knowledge, and that ‘wasted science’ should be avoided. Within scientific computing, voluntary environmental sustainability commitments are starting to emerge that seek to address this through voluntary processes and consultation.¹⁰ The Heidelberg Agreement specifically asks funders to consider the environmental sustainability of what they fund [77]. This cultural shift must go further than just reducing such redundancy. *Strategies* are yet to emerge *that question the value and importance* of the science that is being conducted against its related harms and externalities.

There are many already-known and existing pressures in the academic sphere [40]. Survival within academia often drives researchers to publish as much as they can to remain competitive, and to retain their ongoing role. A culture that values researchers by output volume, research funding, and conference attendance, with little oversight will drive wasteful and damaging practice. Efforts must be made to develop a *culture of science* that provides space for researchers to conduct meaningful work, whilst *not incentivising that which is wasteful* (i.e., considered inefficient, duplicated, poorly conducted). In fact, *academic communities must consider how they incentivise sustainable science*, and care in methods—and they must do this together, if impact is to be meaningful. *Institutions, institutional governance and wider bodies such as funders and infrastructure providers* have a significant role to play in *driving a better and more moderate scientific culture that resists waste* [77].

Imagining strategies and drivers of this type raises important epistemological questions as to how academic and scientific communities judge work to have sufficient value that their impacts are outweighed. In our view, it is crucial that *each community questions how much computation is really needed* to answer certain research questions. Some avenues of exploration may have to be limited or even precluded on sustainability grounds, much as danger to life is already considered in medical research. *Further research* is needed to *develop methods that enable policymakers and researchers to answer these kinds of questions* and anticipate digital-sustainability impacts

¹⁰<https://wellcome.org/who-we-are/positions-and-statements/environmental-sustainability-concordat>, accessed March 7, 2025.

in advance [60]. Specifically, as *academic communities*, what kinds of *policies* could we enact to put pressure on wasteful, profligate, unnecessary, or duplicated research? What care is being taken to understand how already existing policies promote rather than discourage HPC waste in scientific culture? And crucially, how do we ensure an even and consistent approach across the culture of science, to ensure meaningful action?

As a digitally literate discipline, HCI has powerful methods for *exploring this embedding of computation* in practice across science. We can help to *communicate these impacts and growth trajectories* to other domains. We call on HCI researchers to *engage research communities* to develop and inform practical and practicable methods. Further, we call on HCI to *look within itself* given its especially close relationship with digital fabrics and tools, to see what unsustainable trajectories and HPC methods are already being laid down, and to *urgently attend* to the lock-in of unsustainability in HCI [79].

5.4 Limitations and Future Work

While we call for a wide-ranging cultural shift in HPC and scientific computing, we draw on a specific set of data and participants. We recruited HPC users from a number of disciplines, but these were all STEM with no representation from social sciences and humanities. Although computationally intensive research is rarer in these disciplines, it is emerging [19]. This is an important upcoming area and will likely increase demand. Further work should explore the similarities, and differences, of drivers of expansion and waste of specific disciplines, so more cohesive cross-disciplinary (and therefore, cultural) strategies can be developed. Further to this, we have identified different forms of computational waste but our work is qualitative and hence does not provide quantitative evidence as to the scale and wider impact of this waste. Gathering quantitative data would be necessary to determine the relative magnitude of the different wasteful practices we identify, and to explore future trajectories as demand changes.

Our CLD process necessarily makes decisions about the granularity of modelling. We note that many of the nodes and arcs included could be explored to a finer level, identifying differences in the finer-grained interactions between technologies and practices. We see this as a high-level map of the terrain, enabling further research to explore specific issues in more detail. For example, insights from meta-research could be used to identify the ways in which the easy provision of HPC might lead to unnecessary research taking place.

We caution here, however, that our limitations are provided to guide future research, and *should not* be read as reasons why interventions must not be implemented sooner rather than later.

6 Conclusion

We have explored growth and waste in HPC as a case study of computational embedding in academic practice. We have seen how in use and provisioning, technical and socio-cultural factors create several forms of waste from inefficient use and redundant hardware, to needless science. HPC growth is far from just a technical phenomenon, and is shaped by—and shapes—numerous interacting socio-cultural factors stemming from academic structures, and scientific culture more broadly. The pervasive and entrenched nature of ICTs in all aspects of society make it particularly hard to own

and govern. Yet, as we have witnessed for HPC, the embedding of computational methods is evident in most disciplines and is colonising further afield: changing the very nature of what constitutes research, and the expectations associated with it—even within traditionally qualitative disciplines. As infrastructures and the expectations around computation are further embedded into academic practice, it becomes increasingly difficult to reverse forms of waste that become locked-in.

Some communities are already leveraging their attendant structures to tackle this growth. More needs to be done urgently across the scientific culture to work together to attend to this issue. This will take time, and so it is necessary for each academic structure (including HCI) to question the impact of their own implicit decision-making, include HPC as a part of their sustainability policies, and develop ways to incentivise value over waste in HPC-dependent contributions. It is perhaps easy to see HPC and scientific computing as a special case, and hold HCI apart. This, however, is to take a narrow view of what constitutes large-scale computing and its costs. The growing demand for computation is happening not just in research, but in everyday life, and the pervading culture that embeds digital systems in tools and systems at a huge scale.

We draw attention to the developing HPC-dependent literature in HCI, the domain's reliance on large-scale data and increasingly AI, and the computational services we are complicit in designing and endorsing. As the demand for computation increases, infrastructures are upgraded and supply is expanded, further increasing demand in a cycle of growth. This cycle drives escalating material, environmental, and social costs whose impacts are felt unevenly across communities worldwide. We *cannot* continue to ignore the environmental and social justice issues underlying our work. Each community must leverage its own academic structures to address cultural drivers of growth and waste of computation.

The challenge should not be underestimated. But unfettered compute is no longer justifiable. HCI clearly does not own all the structural and cultural leverage points; but we can act as a catalyst to help research communities understand their relationship to computational growth and waste. We must also urgently take the lead in shifting our own culture and practice toward responsible and meaningful use of computation and HPC.

Acknowledgments

Our sincere thanks to our interviewees for their time and insights. ARINZRIT and IRISCAST were supported by NetZero Scoping project [36] (NE/W007134/1), we gratefully acknowledge our funder UKRI NERC Digital Research Programme.

References

- [1] International Energy Agency. 2024. Electricity 2024. <https://iea.blob.core.windows.net/assets/6b2fd954-2017-408e-bf08-952fdd62118a/Electricity2024-Analysisandforecastto2026.pdf>. Accessed: 18th April 2024.
- [2] Gianmarco Baldini, Isabella Cerutti, and Chrysanthi Chountala. 2023. *Identifying Common Indicators for Measuring the Environmental Footprint of Electronic Communications Networks (ECNs) for the Provision of Electronic Communications Services (ECSS)*. Technical Report JRC136475. Publications Office of the European Union, Luxembourg. <https://doi.org/10.2760/093662>
- [3] Oliver Bates, Mike Hazas, Adrian Friday, Janine Morley, and Adrian K. Clear. 2014. Towards an Holistic View of the Energy and Environmental Impacts of Domestic Media and IT. In *CHI '14: Proceedings of the 2014 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 1173–1182. <https://doi.org/10.1145/2556288.2556968>

- [4] Caroline Bird, Chris Preist, Carolynne Lord, Adrian Friday, Kelly Widdicks, Gabin Kayumbi, Simon Lambert, and Adrian Jackson. 2023. *Learning from the Big Picture: Applying Responsible Innovation to the Net Zero Research Infrastructure Transformation (ARINZRI)*. Technical Report. Net Zero DRI Scoping Project. <https://doi.org/10.5281/zenodo.7948731>
- [5] Eli Blevis. 2007. Sustainable Interaction Design: Invention & disposal, renewal & reuse. In *CHI '07: Proceedings of the 2007 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 503–512. <https://doi.org/10.1145/1240624.1240705>
- [6] Jonathan Blumenthal and Miriam L. Diamond. 2022. Sustainability of the Internet of Things Requires Understanding of Mineral Demands and Supplies. *Environmental Science and Technology* 56 (2022), 9835–9837. Issue 14. <https://doi.org/10.1021/acs.est.2c03124>
- [7] Susanne Bødker. 2006. When Second Wave HCI meets Third Wave Challenges. In *Proceedings of the 4th Nordic Conference on Human-Computer Interaction: Changing Roles*. ACM, New York, 1–8. <https://doi.org/10.1145/1182475.1182476>
- [8] Alan Borning, Batya Friedman, and Nick Logler. 2020. The 'Invisible' Materiality of Information Technology. *Commun. ACM* 63, 6 (May 2020), 57–64. <https://doi.org/10.1145/3360647>
- [9] Christina Bremer, Bran Knowles, and Adrian Friday. 2022. Have We Taken On Too Much?: A Critical Review of the Sustainable HCI Landscape. In *CHI '22: Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 1–11.
- [10] CBRE. 2024. *Global Data Center Trends 2024*. CBRE. <https://www.cbre.com/insights/reports/global-data-center-trends-2024>
- [11] United States International Trade Commission. 2021. *Data Centers around the World: A quick look*. Technical Report. United States International Trade Commission. https://www.usitc.gov/publications/332/executive_briefings/ebot_data_centers_around_the_world.pdf
- [12] Payal Dhar. 2020. The Carbon Impact of Artificial Intelligence. *Nature Machine Intelligence* 2 (2020), 423–425. <https://doi.org/10.1038/s42256-020-0219-9>
- [13] Carl DiSalvo, Phoebe Sengers, and Hrönn Hrynjarsdóttir. 2010. Mapping the Landscape of Sustainable HCI. In *CHI '10: Proceedings of the 2010 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 1975–1984. <https://doi.org/10.1145/1753326.1753625>
- [14] Nina Djukanovic. 2022. *"Green as Fields, not Mines": The case of lithium mining and resistance in Serbia*. Master's thesis. University of Oxford.
- [15] Paul Dourish. 2010. HCI and Environmental Sustainability: The politics of design and the design of politics. In *Proceedings of the 8th ACM conference on Designing Interactive Systems*. ACM, New York, 1–10. <https://doi.org/10.1145/2207676.2208539>
- [16] Paul Dourish. 2022. *The Stuff of Bits: An essay on the materialities of information*. MIT Press, Cambridge, Massachusetts. <https://doi.org/10.7551/mitpress/10999.001.0001>
- [17] Kevin Dunnell, Gauri Agarwal, Pat Pataranutaporn, Andrew Lippman, and Pattie Maes. 2024. AI-Generated Media for Exploring Alternate Realities. In *CHI '24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 1–8. <https://doi.org/10.1145/3613905.3650861>
- [18] Briag Dupont, Nesryne Mejri, and Georges Da Costa. 2020. Energy-aware Scheduling of Malleable HPC Applications using a Particle Swarm Optimised Greedy Algorithm. *Sustainable Computing: Informatics and Systems* 28 (2020), 100447. <https://doi.org/10.1016/j.suscom.2020.100447>
- [19] Achim Edelmann, Tom Wolff, Danielle Montagne, and Christopher A. Bail. 2020. Computational Social Science and Sociology. *Annual Review of Sociology* 46 (2020), 61–81. <https://doi.org/10.1146/annurev-soc-121919-054621>
- [20] GO Fair. 2024. *Fair Principles*. GO Fair. <https://www.go-fair.org/fair-principles/>
- [21] Daniel Fallman. 2010. A Different Way of Seeing: Albert Borgmann's philosophy of technology and human-computer interaction. *AI & Soc* 25, 1 (April 2010), 53–60. <https://doi.org/10.1007/s00146-009-0234-1>
- [22] Marcin Frąckiewicz. 2023. *High Performance Computing for Social Sciences: Unraveling the Complexities of Human Behavior*. TS2 Space. <https://ts2.space/en/high-performance-computing-for-social-sciences-unraveling-the-complexities-of-human-behavior/>
- [23] Charlotte Freitag, Mike Berners-Lee, Kelly Widdicks, Bran Knowles, Gordon Blair, and Adrian Friday. 2021. The Real Climate and Transformative Impact of ICT: A critique of estimates, trends and regulations. *Patterns* 2, 9 (10 Sept. 2021), 100576. <https://doi.org/10.1016/j.patter.2021.100340>
- [24] Tony Fry. 2005. The Scenario of Design. *Design Philosophy Papers* 3, 1 (2005), 19–27. <https://doi.org/10.2752/144871305X13966254124158>
- [25] Alexandra M Gormally, Kirstie O'Neill, Michael D Hazas, Oliver EG Bates, and Adrian J Friday. 2019. 'Doing Good Science': The impact of invisible energy policies on laboratory energy demand in higher education. *Energy Research & Social Science* 52 (2019), 123–131. <https://doi.org/10.1016/j.erss.2019.02.012>
- [26] Cédric Gossart, Bernard Aeberscher, Lorenz M Hilty, Bernard Aeberscher, and Lorenz M. Hilty. 2015. Advances in Intelligent Systems and Computing. In *Advances in Intelligent Systems and Computing*, Lorenz Hilty and Bernard Aeberscher (Eds.). Springer, New York, Chapter Rebound Effects and ICT: A review of the literature, 435–448. https://doi.org/10.1007/978-3-319-09228-7_26
- [27] Udit Gupta, Young Geun Kim, Sylvia Lee, Jordan Tse, Hsien-Hsin S. Lee, Gu-Yeon Wei, David Brooks, and Carole-Jean Wu. 2021. Chasing Carbon: The Elusive Environmental Footprint of Computing. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, New York, 854–867. <https://doi.org/10.1109/HPCA51647.2021.00076>
- [28] Maria Häkansson and Phoebe Sengers. 2013. Beyond Being Green: Simple living families and ICT. In *CHI '13: Proceedings of the 2013 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 2725–2734. <https://doi.org/10.1145/2470654.2481378>
- [29] Lon Åje Erni Johannes Hansson, Teresa Cerratto Pargman, and Daniel Sapiens Pargman. 2021. A Decade of Sustainable HCI: Connecting SCHI to the Sustainable Development Goals. In *CHI '21: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 1–19. <https://doi.org/10.1145/3411764.3445069>
- [30] Hörður V. Haralldsson. 2004. *Introduction to System Thinking and Causal Loop Diagrams*. Lund University, Sweden.
- [31] Jonathan Hays, Nicola Walton, Adrian Jackson, A. Mudaraddi, A. Packer, and R.A. Owen. 2023. *IRISCAS: Final Report*. Technical Report. IRISCAS Project. <https://doi.org/10.5281/zenodo.7692450>
- [32] Mike Hazas, Adrian Friday, and James Scott. 2010. Look Back Before Leaping Forward: Four decades of domestic energy inquiry. *IEEE Pervasive Computing* 10, 1 (2010), 13–19. <https://doi.org/10.1109/MPRV.2010.89>
- [33] Simon Hettrick. 2018. *UK Research Software Survey 2014*. Technical Report. Software Sustainability Institute. <https://doi.org/10.5281/zenodo.1183562>
- [34] IPCC. 2022. *Climate Change 2022: Mitigation of Climate Change*. Technical Report. Intergovernmental Panel on Climate Change. 2042 pages. https://www.ipcc.ch/report/ar6/wg3/downloads/report/IPCC_AR6_WGIII_FullReport.pdf
- [35] Adrian Jackson and Ioanna Lampaki. 2024. *HPC System Usage Data*. EPCC. <https://doi.org/10.5281/zenodo.14284188>
- [36] Martin Jukes, Michael Bane, Jennifer Bulpitt, Katie Cartmell, Miranda MacFarlane, Molly MacRae, Alex Owen, Charlotte Pascoe, and Poppy Townsend. 2023. *Sustainability in Digital Research Infrastructure UKRI Net Zero DRI Scoping Project Final Technical Report*. Technical Report. UKRI Net Zero DRI Project. <https://doi.org/10.5281/zenodo.8199983>
- [37] Hsien-Te Kao, Isabel Erickson, Minh Duc Hoang Chu, Zihao He, Kristina Lerman, and Svitlana Volkova. 2024. Machine Learning Insights Into Eating Disorder Twitter Communities. In *CHI EA '24: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, New York, 1–8. <https://doi.org/10.1145/3613905.3651116>
- [38] Sunyoung Kim and Eric Paulos. 2011. Practices in the Creative Reuse of E-Waste. In *CHI '11: Proceedings of the 2011 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 2395–2240. <https://doi.org/10.1145/1978942.1979292>
- [39] Bailey Kimmel, Austin Lee Geisert, Lily Yaro, Brendan Gipson, Ronald Taylor Hotchkiss, Sidney Kwame Osae-Asante, Hunter Vaught, Grant Winger, and Chase Yamaguchi. 2024. Enhancing Programming Error Messages in Real Time with Generative AI. In *CHI EA '24: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, New York, 1–7. <https://doi.org/10.48550/arXiv.2402.08072>
- [40] Gail Kinman. 2001. Pressure Points: A review of research on stressors and strains in UK academics. *Educational Psychology* 21 (2001), 473–492. Issue 4. <https://doi.org/10.1080/01443410120090849>
- [41] Loïc Lannelogue, Hans-Erik G. Aronzon, Alex Bateman, Ewan Birney, Talia Caplan, Martin Jukes, Johanna McEntyre, Andrew D. Morris, Gerry Reilly, and Michael Inouye. 2023. GREENER Principles for Environmentally Sustainable Computational Science. *Nature Computational Science* 6, 3 (2023), 514–521. <https://doi.org/10.1038/s43588-023-00461-y>
- [42] Baolin Li, Rohan Basu Roy, Daniel Wang, Siddharth Sami Vijay Gadepally, and Devesh Tiwari. 2023. Toward Sustainable HPC: Carbon Footprint Estimation and Environmental Implications of HPC Systems. In *IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis*. IEEE, New York, 1–15. <https://doi.org/10.1145/3581784.3607035>
- [43] Zhihao Liu, Yu Li, Fangyan Tu, Ruiyuan Zhang, Zhanglin Cheng, and Naoto Yokoya. 2024. DeepTreeSketch: Neural Graph Prediction for Faithful 3D Tree Modeling from Sketches. In *CHI '24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 1–19. <https://doi.org/10.1145/3613904.3642125>
- [44] Maël Madon, Georges Da Costa, and Jean-Marc Pierson. 2024. Replay with Feedback: How does the performance of HPC system impact user submission behaviour? *Future Generation Computer Systems* 155, C (2024), 66–79. <https://doi.org/10.1016/j.future.2024.01.024>
- [45] Jennifer Mankoff. 2012. HCI and Sustainability: A tale of two motivations. *Interactions* 19, 3 (2012), 16–19.
- [46] Eric Masanet, Arman Shahabi, Sarah Smith, and Jonathan Koomey. 2020. *Re-calibrating Global Data Center Energy-Use Estimates*. Technical Report. Science. <https://doi.org/10.1126/science.aba3758>
- [47] Satoshi Matsuoaka and Naoki Shinjo. 2020. *Fugaku and Associated Technologies Help to Solve Major Challenges in the World Today*. Fujitsu. <https://www.fujitsu.com/global/about/innovation/fugaku/pickup/interview01/>

- [48] Donella Meadows. 2008. *Thinking in Systems: A primer*. Sustainability Institute.
- [49] Donella Meadows. 2015. *Leverage Points: Places to intervene in a system*. Sustainability Institute.
- [50] Srinjoy Mitra, Suvobrata Sarkar, and Agomoni Ganguli-Mitra. 2023. On the Need for an Anticolonial Perspective in Engineering Education and Practice. *Nature Communications* 14, 1 (2023), 8453. <https://doi.org/10.1038/s41467-023-43952-2>
- [51] Bonnie Nardi. 2019. Design in the Age of Climate Change. *She Ji: The Journal of Design, Economics, and Innovation* 5, 1 (2019), 5–14. <https://doi.org/10.1016/j.sheji.2019.01.001>
- [52] Bonnie Nardi, Bill Tomlinson, Donald J. Patterson, Jay Chen, Daniel Pargman, Barath Raghavan, and Birgit Penzenstadler. 2018. Computing Within Limits. *Commun. ACM* 61, 10 (Sept. 2018), 86–93. <https://doi.org/10.1145/3183582>
- [53] United Nations. 2023. *Global Stocktake*. United Nations Climate Change. <https://unfccc.int/topics/global-stocktake>
- [54] Sarker M Parvez, Farjana Jahan, Marie-Noel Brune, Julia F Gorman, Musarrat J Rahman, David Carpenter, Zahir Islam, Mahbubur Rahman, Nirupam Aich, Luke D Knibbs, et al. 2021. Health Consequences of Exposure to E-Waste: An updated systematic review. *The Lancet Planetary Health* 5, 12 (2021), e905–e920. [https://doi.org/10.1016/S2542-5196\(21\)00263-1](https://doi.org/10.1016/S2542-5196(21)00263-1)
- [55] Lucas Pereira, Filipe Quintal, Mary Barreto, and Nuno Nunes. 2013. Understanding the Limitations of Eco-Feedback: A one-year long-term study. In *International Workshop on Human-Computer Interaction and Knowledge Discovery in Complex, Unstructured Big Data*. Springer, New York, 237–255. https://doi.org/10.1007/978-3-642-39146-0_21
- [56] James Pierce, Diana J. Schiano, and Eric Paulos. 2010. Home, Habits, and Energy: Examining Domestic Interactions and Energy Consumption. In *CHI '10: Proceedings of the 2010 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 1985–1994. <https://doi.org/10.1145/1753326.1753627>
- [57] Thibault Pirson, Thibault P. Delhay, Alex G. Pip, Grégoire Le Brun, Jean-Pierre Raskin, and David Bol. 2023. The Environmental Footprint of IC Production: Review, Analysis, and Lessons From Historical Trends. *IEEE Transactions on Semiconductor Manufacturing* 36, 1 (2023), 56–67. <https://doi.org/10.1109/TSM.2022.3228311>
- [58] Chris Preist, Daniel Schien, and Eli Blevis. 2016. Understanding and Mitigating the Effects of Device and Cloud Service Design Decisions on the Environmental Footprint of Digital Infrastructure. In *CHI '16: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 1324–1337. <https://doi.org/10.1145/2858036.2858378>
- [59] Christian Remy. 2015. Addressing Obsolescence of Consumer Electronics through Sustainable Interaction Design. In *CHI '15: Proceedings of the 2015 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 227–230. <https://doi.org/10.1145/2470654.2481378>
- [60] Christian Remy, Oliver Bates, Alan Dix, Vanessa Thomas, Mike Hazas, Adrian Friday, and Elaine M. Huang. 2018. Evaluation beyond Usability: Validating sustainable HCI research. In *CHI '18: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 1–14. <https://doi.org/10.1145/3173574.3173790>
- [61] Kate Saenko. 2023. *A Computer Scientist Breaks Down Generative AI's Hefty Carbon Footprint*. Scientific American. <https://www.scientificamerican.com/article/a-computer-scientist-breaks-down-generative-ais-hefty-carbon-footprint/>
- [62] Tilman Santarius, Jan C. T. Bieser, Vivian Frick, Mattias Höjer, Maike Gossen, Lorenz M. Hilty, Eva Kern, Johanna Pohl, Friederike Rohde, and Steffen Lange. 2023. Digital Sufficiency: Conceptual considerations for ICTs on a finite planet. *Annals of Communication* 78 (2023), 277–295. <https://doi.org/10.1007/s12243-022-00914-x>
- [63] Andreas Schuler, Doris Janssen, Julian Blumenröther, Theresa Maria Probst, Michael Schmidt, and Chandan Kumar. 2024. Generating Personas using LLMs and Assessing their Viability. In *CHI EA '24: Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, Vol. 67. ACM, New York, 1–7. <https://doi.org/10.1145/3613905.3650860>
- [64] Douglas Schuler. 2024. Tools of Our Tools? Exploring the Cybercene Conjecture. *Commun. ACM* 67, 9 (2024), 33–36. <https://doi.org/10.1145/3654437>
- [65] Sabrina Scuri, Marta Ferreira, Nuno Jardim Nunes, Valentina Nisi, and Cathy Mulligan. 2022. Hitting the Triple Bottom Line: Widening the HCI approach to sustainability. In *CHI '22: Proceedings of the 2022 CHI conference on Human Factors in Computing Systems*. ACM, New York, 1–19. <https://doi.org/10.1145/3491102.3517518>
- [66] Euisong Seo, Jinkyu Jeong, Seonyeong Park, and Joonwon Lee. 2008. Energy Efficient Scheduling of Real-Time Tasks on Multicore Processors. In *IEEE Transactions on Parallel and Distributed Systems*, Vol. 19. IEEE, New York, 1540–1552. Issue 11. <https://doi.org/10.1016/j.future.2015.06.003>
- [67] Vishal Sharma, Neha Kumar, and Bonnie Nardi. 2023. Post-growth Human–Computer Interaction. *ACM Trans. Comput.-Hum. Interact.* 31, 1, Article 9 (Nov. 2023), 37 pages. <https://doi.org/10.1145/3624981>
- [68] Arman Shehabi, Sarah Josephine Smith, Dale Sartor, Richard Brown, Magnus Herrlin, Jonathan Koomey, Eric Masanet, Nathaniel Horner, Ines Lima Azevedo, and William Litner. 2016. *United States Data Center Energy Usage Report*. Technical Report. Lawrence Berkeley National Laboratory. <https://eta.lbl.gov/publications/united-states-data-center-energy>
- [69] Elizabeth Shove. 2017. What is Wrong with Energy Efficiency? *Building Research & Information* 46, 7 (2017), 779–789. <https://doi.org/10.1080/09613218.2017.1361746>
- [70] Yolande Strengers. 2014. Smart Energy in Everyday Life: Are you designing for resource man? *interactions* 21, 4 (2014), 24–31.
- [71] Top500. 2024. *TOP500 List - June 2024*. Top500. <https://top500.org/lists/top500/2024/06/>
- [72] Ravi Kant Upadhyay. 2020. Markers for Global Climate Change and its Impact on Social, Biological and Ecological Systems: A Review. *American Journal of Climate Change* 9 (2020), 159–203. Issue 3. <https://doi.org/10.4236/ajcc.2020.93012>
- [73] Sardar Khaliq Uzaman, Atta Khan, Junaid Shuja, Tahir Maqsood, Faisal Rehman, and Saad Mustafa. 2019. A Systems Overview of Commercial Data Centers: Initial Energy and Cost Analysis. *International Journal of Information Technology and Web Engineering (IJITWE)* 14, 1 (2019), 42–65. <https://doi.org/10.4018/IJITWE.2019010103>
- [74] Julia Velkova. 2024. *Dismantling Public Values, One Data Center at the Time*. NordMedia Network. <https://nordmedianetwork.org/latest/news/dismantling-public-values-one-data-center-at-the-time/>
- [75] Dhaval Vyas, Awais Hameed Khan, and Anabelle Cooper. 2023. Democratizing Making: Scaffolding Participation Using e-Waste to Engage Under-resourced Communities in Technology Design. In *CHI '23: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 1–16. <https://doi.org/10.48550/arXiv.2302.10402>
- [76] Matt Watson and Elizabeth Shove. 2023. How Infrastructures and Practices Shape Each Other: Aggregation, Integration and the Introduction of Gas Central Heating. *Sociological Research Online* 28, 2 (2023), 373–388. <https://doi.org/10.1177/13607804211055495>
- [77] Philipp M. Weber, Sandra Bendiscioli, Gerlind Wallon, Uwe von Ahlsen, Anne Marie de Beaufort, Marion Boland, Florijn Dekkers, Dominique Dunon-Buteau, Martin Farley, Alyson Fox, Stéphane Guillot, Loïc Lannelongue, Marta Łazarowicz-Kowalik, Brendan Rouse, Gabrielle Samuel, Teresa Sanchis, Susan Simon, Mathew Tata, Gisou van der Goot, and Fiona M. Watt. 2024. The Heidelberg Agreement on Environmental Sustainability in Research Funding. <https://doi.org/10.5281/zenodo.13938809>
- [78] Kelly Widdicks, Oliver Bates, Mike Hazas, Adrian Friday, and Alistair R. Beresford. 2017. Demand around the Clock: Time Use and Data Demand of Mobile Devices in Everyday Life. In *CHI '17: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 5361–5372. <https://doi.org/10.1145/3025453.3025730>
- [79] Kelly Widdicks, Federica Lucivero, Gabrielle Samuel, Lucas Somavilla Croxatto, Marcia Tavares Smith, Carolyn Ten Holter, Mike Berners-Lee, Gordon S. Blair, Marina Jirotk, Bran Knowles, Steven Sorrell, Miriam Börjesson Rivera, Caroline Cook, Vlad C. Coroamă, Timothy J. Foxon, Jeffrey Hardy, Lorenz M. Hilty, Simon Hinterholzer, and Birgit Penzenstadler. 2023. Systems Thinking and Efficiency under Emissions Constraints: Addressing rebound effects in digital innovation and policy. *Patterns* 4, 2 (2023), 100679. <https://doi.org/10.1016/j.patter.2023.100679>
- [80] Kelly Widdicks, Christian Remy, Oliver Bates, Adrian Friday, and Mike Hazas. 2022. Escaping Unsustainable Digital Interactions: Toward ‘more meaningful’ and ‘moderate’ online experiences. *International Journal of Human-Computer Studies* 165, C (2022), 102853. <https://doi.org/10.1016/j.ijhcs.2022.102853>
- [81] Jackson Woodruff, David Schall, Michael F.P. O’Boyle, and Christopher Woodruff. 2023. When Does Saving Power Save the Planet?. In *HotCarbon '23: Proceedings of the 2nd Workshop on Sustainable Computer Systems*. ACM, New York, 1–6. <https://doi.org/10.1145/3604930.3605719>
- [82] Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga Behram, James Huang, Charles Bai, Michael Gschwind, Anurag Gupta, Myle Ott, Anastasia Melnikov, Salvatore Candido, David Brooks, Geeta Chauhan, Benjamin Lee, Hsien-Hsin S. Lee, Bugra Akylidiz, Maximilian Balandat, Joe Spisak, Ravi Jain, Mike Rabbat, and Kim M. Hazelwood. 2021. Sustainable AI: Environmental Implications, Challenges and Opportunities. *CoRR* 4 (2021), 795–813. <https://doi.org/10.48550/arXiv.2111.00364>
- [83] Wendy Wyuts. 2020. Market Distortions Encouraging Wasteful Consumption. In *Encyclopedia of the UN Sustainable Development Goals*. Springer, New York, 443–453. https://doi.org/10.1007/978-3-319-95726-5_45