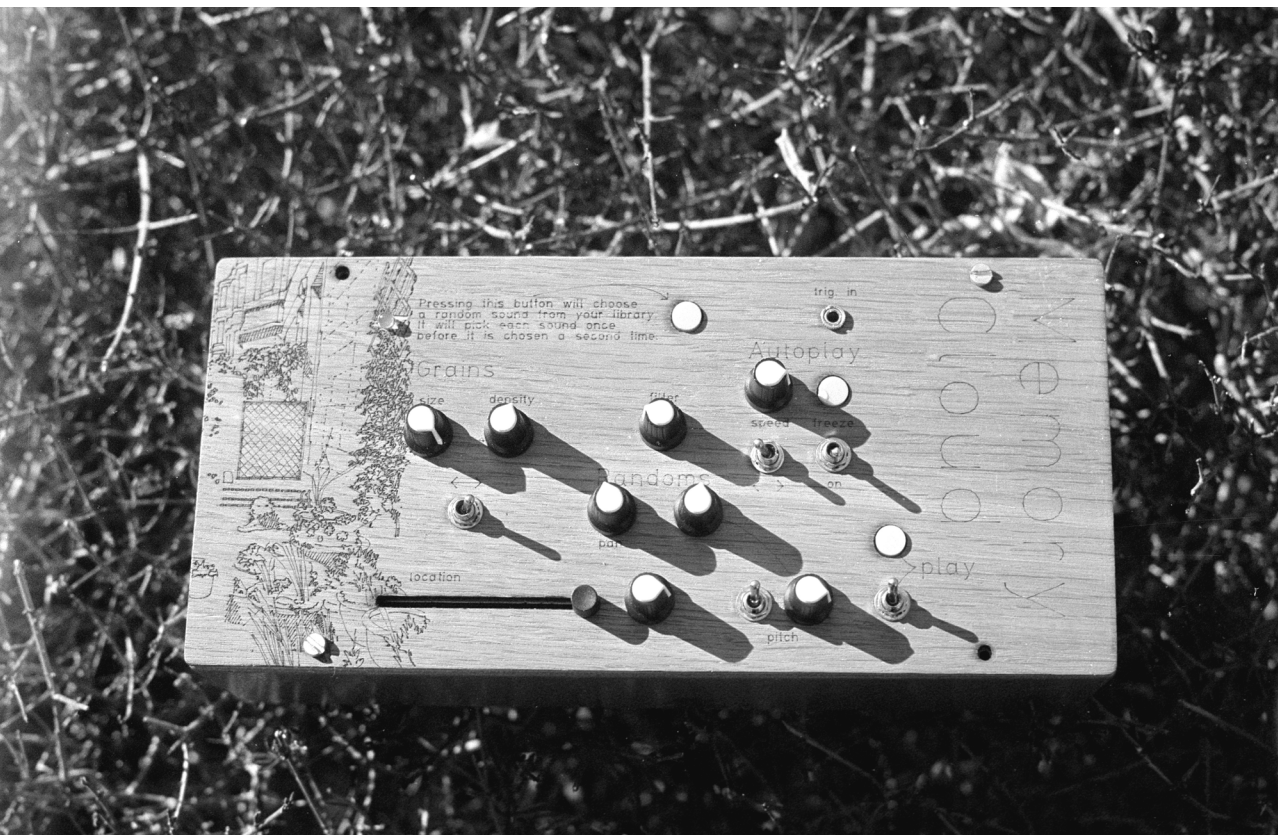
Doctoral Thesis in Media Technology

Configuring Alternative Human-Media Relations

YANN SEZNEC

KTH ROYAL INSTITUTE OF TECHNOLOGY



Configuring Alternative Human-Media Relations

YANN SEZNEC

Academic Dissertation which, with due permission of the KTH Royal Institute of Technology, is submitted for public defence for the Degree of Doctor of Philosophy on Wednesday the 3rd of June 2026, at 2:00 p.m. in F2, Lindstedtsvägen 26, Stockholm.

Doctoral Thesis in Media Technology
KTH Royal Institute of Technology
Stockholm, Sweden 2026

© Yann Seznec

Cover page photo: Yann Seznec

TRITA-EECS-AVL-2026:52
ISBN 978-91-8106-621-0

Printed by: Universitetsservice US-AB, Sweden 2026

To Kristina, Arthur, and Enora

Abstract

This thesis describes the design of a set of sound making devices. These were built as a Research through Design inquiry into how to confront the biases that guide the design of interfaces for digital media. Each of the resulting designed artifacts enacts alternatives to the so-called cornucopian paradigm of computing design, which promises permanence and endless growth.

The work is situated within Sustainable HCI (SHCI), specifically borrowing from limits and constraints-based design principles. The resulting designs include devices that manifest concepts of destruction, degrowth, and a reimagining of how to interact with vast media libraries. An analysis of the devices points towards a set of specific design qualities that emerge from their creation and use.

These qualities suggest a move away from designing for unsustainable growth and permanence and towards something based on the underlying characteristics of the media itself. The portfolio thus represents a reconfiguration of the human-media relationship by designing for and with the limitations of our personal media.

Keywords
media, research through design, sustainability

Sammanfattning

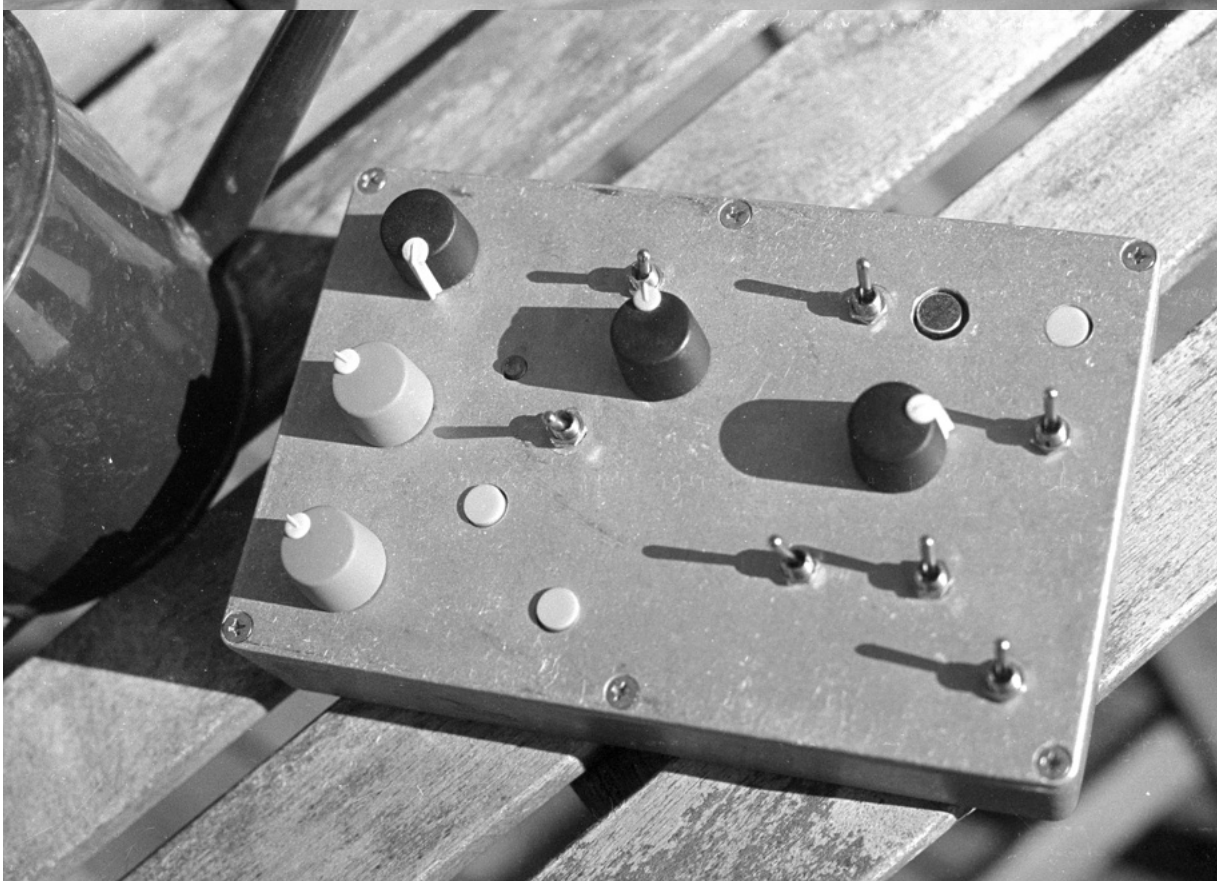
Denna avhandling beskriver framtagandet av en kollektion ljudframställande objekt. En forskning-genom-design-undersökning för att konfrontera de antaganden som ligger till grund för utveckling av gränssnitt för digitala medier. Artefakterna representerar alternativ till det så kallade cornucopianska paradigmet inom datordesign som lovar varaktighet och oändlig tillväxt.

Detta arbete positioneras inom området hållbar människa och datorinteraktion (SHCI), och drar det inspiration från forskning om ramar och designbegränsningar. De resulterande artefakterna inkluderar apparater som manifesterar koncepten destruktion, av-växt (degrowth) och omdefinierar interaktioner med stora mediabibliotek. En analys av kollektionen leder till en uppsättning specifika designkvaliteter, som framkom efter deras tillverkning och användning.

Dessa kvaliteter föreslår ett skifte från design för ohållbar tillväxt och varaktighet till något baserat på de underliggande egenskaperna hos själva mediet. Denna portfolio representerar därigenom en omkonfiguration av förhållandet mellan människa och media, genom att designa för och med begränsningarna i vårt media.

Nyckelord
mediateknik, forskning genom design, hållbarhet


Pressing this button will choose
a random sound from your library
It will pick each sound once
before it is chosen a second time
Grains
size
density
filter
speed
freeze
on
pan
pitch
location



Acknowledgment

It has been an immense privilege to develop the work that I present in this thesis. This is the result of the incredible support of a huge number of people, and I hope the designs and the words contained in this book do justice to the faith that has been placed in me over the past few years. My two supervisors showed particular patience and fortitude on my journey, and it is not any exaggeration to say that I could not have come this far without them. Thank you to Rob Comber for encouraging me to *make weird things*, whilst challenging me to think deeply about what all three of those words really mean. Many thanks also to Elina Eriksson, who slowly, carefully, and thoughtfully guided me towards a profound change in how I view the world. Many others at KTH have also supported me in ways large and small, including Kia Höök, Henrik Artman, Leif Handberg, Jarmo Laaksolahti, Eva-Lotta Sallnäs Pysander, Madeline Balaam, and Patric Dahlqvist. A special thanks to Kjetil Falkenberg for his well-considered and thorough feedback on this thesis.

The designs presented here, along with many other discarded concepts and ridiculous experiments, were built thanks to the unwavering enthusiasm and technical support of Dee Tobin and Laura McHugh in MIDDLA. A special thanks also to Alex Gent for sparking my initial thoughts about sound destruction many years ago.

Among all of the things that came together to make this process successful, perhaps the most surprising and most rewarding was how I formed a part of a truly incredible group of other PhD researchers going through the journey in our little department at around the same time. These include Joseph La Delfa, Nadia Campo Woytuk, Andreas Lindegren, Annika Andreasson, Aksel Biørn-Hansen, Arjun Menon, Minna Thorslund, Kristina Popova, Jooyoung Park, Mafalda Gamboa, Petra (and Furbol) Jääskeläinen, Anna-Kaisa Kaila, Derek Holzer, Fatemeh Bakhshoudeh, Rémy Rey, Kelsey Cotton, and Vincenzo Madaghiele. You have all been an inspiration and I count myself lucky to have spent this time with you.

Outside of KTH my work has benefited enormously from discussions with Noura Howell and Will Odom, and I would like to thank in advance my opponent Thor

Magnusson and my committee Sarah Fdili Alaoui, Palle Dahlstedt, and Stefano Fasciano - I am very proud to be in the company of such fantastic researchers and artists.

To my incredibly supportive family - Mama, Papa, Cory, Emily, Gwen, Peter, Juliet, Cato, Michelle, Pete, and Kylee: your unending encouragement means everything to me.

To my wife Kristina - this acknowledgement could never adequately convey my gratitude and love for you and everything we have built together. You have guided every idea, concept, and design in this thesis, and thanks to you I see this work not as a finished, final thing but instead as a seed that we plant as a promise to the future.

And finally, to our children - you two are an absolute joy, I am a very lucky Papa.

Yann Seznec
Stockholm, May 6, 2026

List of included papers

This dissertation is comprised of five published articles written over the course of my PhD studies, together with a Kappa which presents the work as a whole. These five papers will be referenced throughout the Kappa.

Paper A **Music Within Limits**, Yann Seznec. In 2022 International Conference on ICT for Sustainability (ICT4S), 2022

Paper B **The Memory Cloud: Personal media libraries as affordance and constraint**, Yann Seznec. In Proceedings of the International Conference on New Interfaces for Musical Expression, 2025

Paper C **Three Ways to Destroy a Sound**, Yann Seznec. In Adjunct Proceedings of the 2024 Nordic Conference on Human-Computer Interaction, 2024

Paper D **Play/Destroy**, Yann Seznec, Andreas Lindegren, Frida Lindegren, and Rob Comber. In Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems, 2026

Paper E **All of Your Sounds**, Yann Seznec. Submitted Manuscript 2026

List of included designs

Design	Description
Memory Cloud	Instrument that contains all of the sounds from a user's cloud video library. One sound can be loaded at a time into a granular synthesis engine for manipulation and exploration. You can not return to a previous sound without hearing every other sound first. Technical information and code available here: https://github.com/yannseznec/ys.granular
All of Your Sounds	Device that creates a virtually infinite composition based on all of the sounds from a user's cloud video library. Technical information and code available here: https://github.com/yannseznec/memoryMachine
Sound Destruction Engine	Software designed for loading a digital sound file and destroying it. Various adjustable parameters are available. Technical information and code available here: https://github.com/yannseznec/soundDestroyer
The Sound Destroyer	A standalone instrument containing the Sound Destruction Engine. Sounds are recorded into the instrument and destroyed through manipulation of the knobs and buttons. Technical information and code available on the Sound Destruction Engine repository.
Kopiator	Device with microphone, single button, and built-in speaker. It records a sound while the button is pressed and instantly destroys it on loop (using the Sound Destruction Engine). Technical information and code available on the Sound Destruction Engine repository.
Memory Box	A hexagonal wooden box with a single visible button. A meaningful sound is transferred into the device, and pressing the sound will both play and simultaneously destroy the sound. Technical information and code available on the Sound Destruction Engine repository.
The Period Instrument	A collaboration with Nadia Campo Woytuk. A granular synthesizer instrument designed to respond to menstrual blood. Technical information and code available here: https://github.com/yannseznec/periodInstrument
The Book of Knowledge of Impractical Musical Devices	A set of three instruments designed and built before my PhD studies. Includes the first iteration of sound destruction. Technical information and code available here: https://www.impracticaldevices.com/

Contents

Abstract vii

Sammanfattning ix

Acknowledgment xi

List of included papers xiii

Contents xv

I Kappa 1

1 Foreword 3
1.1 An anecdote 3

2 Introduction 5

3 Background 7
3.1 Sustainability 7
3.2 Considering a Media Ecology 10
3.3 Sound, time, and sustainability 11
3.4 HCI, NIME, and other design work in this area 12
3.5 Summary 14

4 Making Weird Things 15
4.1 Theory 15
4.2 Method 16
4.3 Limits and Constraints 17
4.4 Presentation of the work 17

5 Work 19
5.1 Included papers 19
5.2 Design Portfolio 22
5.3 A note on other related work 26

6 Design Qualities 33
6.1 Intermediate Time 33
6.2 The Firehose 35
6.3 Play/Destroy 36

6.4 In the Box 38
6.5 External Memory 40

7 Discussion: Reconfigurations 41
7.1 Taking Time, Making Time 41
7.2 Embracing Impermanence 42
7.3 Seamful Media 43
7.4 Media Instruments 43
7.5 Summary 45

8 Conclusion 47

Bibliography 51

II Included papers **61**

Kappa

1 Foreword

1.1 An anecdote

Early in my PhD studies I was asked to clear some bookshelves. The books sitting on those shelves had almost certainly not been touched in many years, and space at KTH is at something of a premium. They had to go.

As I worked with my colleagues to laboriously shift these books from the shelves to the large mobile trash bins that had been rented for the occasion, I began to notice that the majority of these books were, in fact, the byproduct of the same kind of work that I was engaged in at that very moment. Thesis dissertation after thesis dissertation, many of them dating back decades or even nearly a century, gathering dust on a shelf until their unceremonious final journey to the bin and eventually out to a recycling center to be burned and converted to steam in the district heating factories of Stockholm.

An academic publication, I have come to learn, is *knowledge*. Precisely how we define the boundaries of that knowledge, how we have arrived at it, and how we justify and support it, forms the bulk of an academic argument. All of this is contained within the words and pages of a written text, whose publication represents the acceptance of this knowledge contribution (within a specific community, at least). This knowledge will then be archived until the day that it is used to heat our homes.

My work has looked into digital media through the particular focus (or tunnel vision) afforded by political ecology, sustainable computing, and research through design, and it is fitting to accept the impermanence of this work as a part of that. I imagine this thesis will end up on a shelf one day, virtual or otherwise, and its destruction and disappearance will go unnoticed. That ending has in many ways already happened - it is built in, just as it is built into all things. Insofar as there is an overarching theme to this dissertation, it might be a call to embrace this impermanence and integrate it into the work.

2 Introduction

This thesis will present a number of designs for sound making machines. These devices were designed as a way of exploring a research space around how the field of Human-Computer Interaction can alternatively configure our relationships with media. But before we even get to the thorny issue of *how* to reconfigure, it is reasonable to ask: *why* do our relationships need reconfiguring?

I will not be the first to point out that we are baffled by our own media libraries. The way we deal with them is confused at best. We casually carry tens of thousands of photographs and videos, a selfie with our grandparents next to a picture of a receipt we sent to Finance. We record six videos of a concert, and never watch them again. The emotional scaffolding we use for these libraries is fascinating to dismantle with thought experiments - when faced with the need to shrink my library, how many random photographs would I be happy to delete? Deleting just one photograph out of 10,000 is probably fine, but what about 10? 20? 100? What if I never knew which ones were deleted? What if every time I looked at a photograph it would also be deleted permanently? At least some of the design work I'll present in this thesis could be seen as real-life instantiations of thought experiments like this.

In my first year as a PhD student I took a course with Elina Eriksson and Daniel Pargman about sustainability and computing. It was a course that profoundly changed my approach to my work in ways that I am still understanding and exploring now. Although I have been making sound machines of varying quality for nearly twenty years, I have found that engaging with sustainability as a deep and systemic concept has given direction and focus to my work. I have learned that while it is easy to point to where we go wrong with our approach to digital media, exploring the underlying reasons why this is the case is more uncomfortable. What do we want from our media? What is enough? What do we need? What can we afford?

The need to reconfigure our relationship stems from the understanding that the issues with digital media are not a problem to fix, but the result of a fundamental set

of characteristics based on an unsustainable approach to computing more generally - an obsession with permanence, seamlessness, and growth. This bias, mirrored throughout society, frames all of the assumptions inherent to human-computer interaction (HCI) and by extension the design of our digital media technology tools.

And so we arrive at the questions framing this thesis: How can we design human-computer interactions that challenge the unsustainable bias towards growth in computing? What qualities would emerge from designs that attempt to do this?

My approach to answering these questions begins with sustainability, looking specifically at Sustainable HCI to contextualize my explorations. I will situate my research within a broad range of work in media ecology and HCI, focusing on sound as a particular form of personal media. A number of designed artifacts that make and manipulate sound will be presented that enact core principles from this background research. These will be presented as a portfolio, alongside five included papers. The papers and portfolio serve as the basis for identifying a set of design qualities, which are then discussed in terms of how they lead to potential reconfigurations of our relationship to media.

3 Background

3.1 Sustainability

The very concept of *sustainability* is to a certain degree a contested space, with a number of competing definitions whose differences may appear subtle but whose ramifications can be significant. The UN Brundtland Report from 1987 defines *sustainable development* as that which "meets the needs of the present generation without compromising the ability of future generations to meet their own needs" [6]. More recently the overarching concept of sustainability is perhaps most visible in the 17 UN Sustainable Development Goals [7], which function as a call to action "for peace and prosperity for people and the planet"[1]. Other definitions explore the relationship between environmental, social, and economic sustainability, either as a Venn diagram or in various different types of overlap [8]. Still others look specifically at resource use [9] or planetary boundaries [10, 11] which take into account 9 ecological subcategories (Ocean acidification, Carbon Dioxide concentration, etc) which, when within a safe operating space, form a stable planet. As I write this, 7 of those boundaries have been crossed [12].

The particular subfield of sustainability within which I am situating this work, sustainable HCI (SHCI), was first named and described by Blevis in 2007 [13][2], when he argued that sustainability should be the primary focus of interaction design. In his paper this is discussed from a material perspective by "linking invention and disposal" and encouraging "renewal and reuse", and also from an experiential standpoint by several principles that embed sustainable thinking into interaction design (de-coupling ownership and identity, using natural models, etc).

[1] `https://sdgs.un.org/goals`

[2] The original Blevis article refers to Sustainable Interaction Design (SID), though SHCI is now more common. "SID" is also an acronym for Sonic Interaction Design, therefore I will use SHCI to avoid confusion as the work here clearly has a sonic component.

In addition to these relatively practical guidelines, Blevis pointed out the issues of situating HCI within sustainability in any capacity, as it is inherently human-focused:

> "The very title Human-Computer Interaction has embedded within it meanings which are problematic from the perspectives of sustainability—it is anthropocentric, and even if the anthropocentrism was not in-and-of-itself a condition of ontological blindness, the sense of human-centeredness in the HCI context is oftentimes construed as a notion of method in which engineering "needs and requirements" follow from cognitive models of "users," rather than a concern for human conditions, particular or global."

This tension takes a number of different forms throughout the SHCI literature. Silberman, for example [14] suggests that computing systems are in no way inherently "pro-sustainability", and indeed the systems that generate them are very often explicitly anti-sustainable. Knowles and others [15] instead suggest that the current global economic and social system is fundamentally incompatible with the types of changes required to avoid catastrophic climate collapse. A definition of sustainability based on development and growth is, perhaps, something of an oxymoron.

In response to traditional concepts of sustainability, which often contain an embrace of growth (or at least a maintenance of the status quo), concepts of post-growth or degrowth have arisen [16, 17, 18]. These theories are based on the assumption of resource scarcity and global instability, suggesting that society will be forced to reckon with these issues, and it is likely in our best interest to accept them as a reality and actively slow our economic and material activities accordingly [19]. Crucially, these researchers often argue that moving away from growth does not necessarily correlate with lower quality of life, but rather that the growth imperative is itself a driver for inequality, poverty, and injustice.

Specifically within SHCI, researchers in collapse informatics [20] have looked into how to account for a "future of scarcity". This area of work argues that the idea of the global crisis as something that *will* happen is fundamentally flawed and ethnocentric. The crisis is, in fact, already happening in much of the world [21]. A more applied approach has looked at exactly what design constraints will be placed upon technology designers after a global infrastructure collapse [22]. As a whole one goal of this line of research is to problematize the notion of technological advancement, and attempt to formulate a methodology for working within limitations that are either strictly placed on the designs or imposed due to outside forces [23, 24, 14]. This "limits" or constraints-based approach [25] represents a move towards a more socially and technologically sustainable future, and forms a critical part of my work.

Other researchers have pointed out that computing technology has a particularly close and direct relationship to notions of growth and speed. Dubbed the *cornucopian paradigm* [26], this reflects a fundamental bias towards unsustainable growth in computing. Preist et al. outline a number of principles which illustrate this paradigm (Instant, Pervasive, Eternal, etc) that are considered to be the ideal in computing system design, and which each require and expect growth in materials and

infrastructure in order to maintain. A crucial aspect of the definition of this concept is how the design decisions that enable it are pervasive and implicit - the expectation of cornucopia is so deeply embedded into the design process that it becomes nearly invisible to the designer (and user). Examples of this include the requirement of always-on broadband or unlimited data packages built into the design of cloud storage systems, or the intensification of features in mobile phone messaging systems [27].

In terms of sustainability studies, this thesis will focus on this cornucopian paradigm - not because it is the most important concept in sustainability (or even in sustainable computing), but because it is particularly striking in digital media design and can be seen as an umbrella concept with a number of different consequences. It is a fundamental computing bias that guides the design of all digital tools, most often to their detriment. It is the pervasive need for computing design to offer *more*: more features, more speed, more space.

As an example - the concept of forced obsolescence is a common point both in general complaints about technology and within SHCI [28]. While there is undeniably pressure from computing companies to encourage people to buy new products and services in order to generate income, it is counterproductive to consider design decisions as deliberate attempts to create poor user experiences in the service of forcing new purchases. Rather, using a cornucopian lens can allow us to see the implicit design biases that actually guide these decisions and understand that they are most often woven into the identity of computing design as whole. These are not deliberate attempts to force people to buy, but rather they are assumptions: that users will have high quality hardware, high-speed broadband, multiple devices, virus-free software. This is arguably worse. While the former can be apparently "fixed" with regulations, the latter is an in-built bias that can not be shifted without fundamental system change. This is much harder to accept and also much harder to do.

Returning to the tensions within the common definitions of sustainability, the corncucopian approach to technology mirrors the general glorification of economic growth as a measure of progress, and is equally unsustainable. We simply can not depend on endless server space or bandwidth or hard drives - while it may seem unimaginable now, resources are never limitless. Limits-based or constraints-based design, guided by sustainability concepts of post-growth, offer an alternative to this cornucopian paradigm.

To a lesser extent, this thesis will also engage with another tension in sustainable HCI: the historic focus on individual behaviour change. As Dourish points out, "HCI has often transformed the problems of sustainability into the costbenefit trade-offs of rational actor economics, promoting sustainability as a matter of personal morality rather than industrial regulation or political mobilization." [29] Arguably an entire field has arisen over the past few decades which involve the development of some technology or design that is positioned as something that will enable a single person to make some small contribution towards a sustainable future. This is perhaps most visible in numerous smart-home "resource man" [30] style projects such as systems that encourage users to lower their heating by one degree, and even in more experimental design works including my own Singing Shower [31]. More recently, this focus within SHCI on designs aimed at individual behaviour change has been

heavily criticized, along with the fundamental belief that users will make informed or rational decisions based on designs [32, 33].

The pressure on the individual has a tendency to dominate narratives on sustainability, not just in technology but more broadly throughout society. The prompts to upgrade our cloud storage are annoying but we blame ourselves, because they are simply the extension of the other external expectations being placed on us every day under the dubious guise of sustainability - fly less, eat less meat, drive less, and crucially: do not consider that these personal changes are only necessary due to the large structural forces encouraging us to have these problems in the first place. Personal media provides a particularly noticeable window into this tension when we are prompted to delete our memories or upgrade or storage. This is portrayed as both our own fault (why are you taking so many pictures?) and easy to solve with money (buy to upgrade!).

Individual behaviour change can be characterized as an approach with a low leverage point, in the Meadows sense [34]. It may make some difference, but because it is not attempting to engage with (let alone change) the underlying systems any results will always be limited. A systems thinking approach [35] can suggest an alternative way of thinking about sustainability, considering inner transition [36] or the concept of hospicing or “composting” [37] our current systems as we work towards building anew.

In summary, the work presented here can be broadly situated within sustainable HCI, and in particular responds to and is contextualized by the SHCI concept of the cornucopian paradigm. This focus overlaps strongly with an embrace of limits or constraints-based design.

3.2 Considering a Media Ecology

Recorded media can be seen as a particular application of the cornucopian paradigm. This arguably has formed a part of the political ecology of media since the advent of recorded technologies, such as with the explosive growth of recorded music being intertwined with colonial extractivist structures [38]. As media has become increasingly associated with computing technology, to the point of being now virtually indistinguishable, it has taken on these characteristics even further. In particular, personal digital media is embedded with three characteristics: permanence, seamlessness, and endless access. Digital media systems are implicitly designed with these characteristics assumed as the ideal.

These biases are reflected in the seemingly unstoppable growth of digital storage mediums. Each generation of digital storage is portrayed as a permanent solution to the problems raised by the previous generation. The 3.5” floppy disk improved the durability of the 5.25” disk while offering more space (as the 5.25” had done in comparison to the 8” disk, and so on). Zip drives, SCSI drives, CD-Rs, SSDs - each of these offered more space or fidelity to accommodate the growing size and quality of our media libraries, and each was portrayed as a permanent solution, most famously with Philips marketing the CD as providing “Pure Perfect Sound Forever”.

The latest step in this cycle is the cloud, whose main distinction from previous

generations of storage is how it promises effectively endless space. The quiet tradeoff to this promise is that the data you store is not physically with you but instead lives in an obscure set of servers located in various data centers around the world. This trend towards outsourced data is conceived, as with every previous generation of media storage, as a method for enhancing the ideal of seamless and endless access of media.

Meanwhile, the general expectation is that a recorded sound, video, or photograph should be the same each time it is viewed - pressing "play" will reproduce the media the same way each time. Not only is this curiously gendered language indicative of how technology can reinforce hegemonic norms [39], this assumption reflects how cornucopian concepts of growth, permanence, and access are built in to the design and experience of recorded media. Recording a personal video or taking a photograph is in effect a promise to your future self that you will revisit this moment. That promise is dependent on a technology that is itself promising that the media will be made permanently available to you, and when you revisit it that moment will be faithfully reproduced.

The relentless push towards growth and the inherent impermanence of computing media are interleaved, with each amplifying and enabling the other. Viewing this media ecology through a lens of sustainable computing and SHCI, particularly with regards to post-growth thinking, clarifies how this cornucopian bias is counterproductive and unsustainable. Our media libraries grow because we deny their impermanence, and we require new forms of media to accommodate them, which will in turn lead to a new cycle of loss. These characteristics - permanence, seamlessness, and endless access, form a media ecology that we engage with, informing a relationship with the media itself. This is a media ecology based on promises that can not possibly be kept. The work presented in this thesis attempts to unpack this relationship, offering a set of designs that engage with alternative forms of media experience and suggest new human-media relations.

3.3 Sound, time, and sustainability

It will become clear in the following pages that my PhD work has focused almost exclusively on the medium of sound. On the one hand, I have a background in sound, music, and creative audio technologies, so there is a certain element of simply continuing to develop my work in the area that I have the most expertise and applying it to a new field. I am digging where I stand [36]. However it is equally true that the area of sustainability, and in particular sustainable computing, attracted me precisely because of my previous experience in sound. Much of my work with sound has focused on its nuanced relationship with time and permanence, both of which are fundamental to a deeper understanding of what sustainability can mean.

The Paris Agreement of 2015 had the stated goal of limiting the rise in average global temperature to a maximum of 2°, and preferably to a maximum of 1.5°. It is generally accepted that achieving these targets will require a massive global effort, with an estimate of 50% reduction in emissions required by 2030 to have any possibility of achieving the goal. Any rise in temperature beyond these targets is likely to cause serious and irreversible harm to the planet.

However these numbers are effectively an illusion - the eco-disaster to which we are responding has already happened, and a growing body of evidence indicates that the "tipping points" towards a vastly changed future have already been surpassed [10], with studies indicating that humanity's ecological demand outpaced the planets capacity starting sometime in the 1980s [40]. Only recently have the consequences of this so-called "overshoot" been felt [34], contributing to a disconnect between scientific reality and felt experience and meaningful action [41].

Whilst the need for action to immediately lower global emissions is undeniably crucial, it is also important to acknowledge that even just the existing carbon in the atmosphere is almost certainly enough for centuries worth of elevated temperatures [42]. This disconnect between our perceived present and the actual present (often framed as a possible future to be avoided) can be seen as an extension of the more general disconnect between human-time and nonhuman-time. Our inability to interpret non-human timescales in effect reduces the perceived agency of nature, leading to a fundamental lack of understanding of our present and potential futures [43, 44] and our present and potential agency.

Nevertheless the Paris Agreement figures, and other similar statistics, are often used as a starting point for design research that attempts to propose a technology or design strategy that will in some way mitigate or slow climate change-related issues and contribute to a more sustainable future. Design is, after all, traditionally defined as intentionally changing an existing situation into a more desirable one [45].

It is perhaps precisely the seemingly intractable nature of these issues that has attracted me to frame my work with a sustainability perspective. A crucial factor of sustainability work is the act of dealing with the past, in the service of the future. This is at least metaphorically, and to a certain extent experientially, exactly the same as working with sound: listening is the act of sonically perceiving the outcome of the release of energy in the recent past. In this sense, sound illustrates the passage of time on a micro scale when listening to something nearby, an unimaginably macro scale when listening to universes collapsing [46], or in a disorienting middle space that Jensenius [47] and others explore - the latency between seeing a flash of lightning and hearing a clap of thunder suddenly making you aware of time and space.

Working with sound therefore reflects the issues faced by the world as a whole - a troubling relationship between the past and the present, and an attempt to control or understand an uncertain future. Just as the current climate disaster is a result of unchangeable actions and decisions already taken, whenever we listen to a sound we are hearing the after effects of a burst of energy that happened sometime in the past. The videos and sounds that we record, store, and forget about are one way that we mediate this uncomfortable relationship, and the technology that we use to do so must be interrogated with that in mind.

3.4 HCI, NIME, and other design work in this area

My work leading up to my PhD has involved building musical devices for exploration, music making, and performance. This has included making a musical pig sty, building a soundtrack from the sounds of a whiskey distillery, generating electro-mechanical

music from mushrooms, and designing musical composition games, among other things. To a certain extent my doctoral studies provided me an opportunity to deepen this work through research and academic contextualization. I have found my work aligning most strongly with research looking into alternative musical instruments as epitomized by the NIME (New Interfaces for Musical Expression) community, which itself has origins in Human-Computer Interaction. One important component of this has been developing an understanding of how musical instruments, much like scientific instruments for measurement, embody and perpetuate knowledge and biases, as explored by Baird [48] in terms of *thing knowledge* and specifically for musical instruments by Magnusson [49, 50] in terms of *epistemic tools*.

The field of HCI has generated a number of approaches to the overwhelming mass of media and the heavy weight of history that goes along with it. As personal media archives have grown in response to advances in technology and the integration of high quality media recording into personal smartphone devices, the issue of how to effectively navigate the resulting libraries has arisen. Much of the HCI work in this area is based on accepted heuristics [51, 52] such as “Minimize user memory load”, “Be consistent”, or “Offer informative feedback”.

While early academic research revealed problems of growing photo archives [53], further work quickly uncovered that “find and retrieve” is not enough to cover the nuance of personal archiving practices [54]. As size of photo and video libraries started to regularly be measured in the thousands, major technology companies such as Apple and Google introduced a number of automatic tools for sorting and presenting the content. Apple’s Photos app on my computer now sorts my media not only by type (videos, selfies, panoramas, etc) and date, but also by “Memories”, an automatically generated set of albums with names such as “Pet Friends”, “By the Water”, and “A Big Day”. There are now 752 memories available to me, which still feels rather overwhelming. The solution has ended up reproducing the very problem it was meant to solve (this is something of a habit for HCI [55]).

An approach which aims to offer a different experience can be seen in HCI research from the “slow technology” movement [56]. This design philosophy explicitly avoids efficiency and speed in favor of reflection and care. This has been applied to visual media with projects including Odom’s Photobox [57] and Capra [58], and Chen’s Chronoscope [59], which all provide users with glimpses of their own photo archives, encouraging reminiscence through tightly constrained interactions. A similar strategy, when applied to location, resulted in White’s Memory Tracer and Memory Compass designs [60], and when applied to domestic vibrations resulted in The Inner Ear by Desjardins et al [61]. Of particular relevance to my thesis work is Odom’s design work around sound with the OLO Radio [62] and Olly [63] , which both empower users to explore the data of their music listening past.

One of the main characteristics uniting this work is how they all are the result of an alternative method of approaching time. Many of these designs push our interaction models for the temporality of our media archives, either through the momentary access to an old photograph or to the granular navigation of a listening past. The revaluation of traditional HCI methods for dealing with time is also seen in research projects looking into the vast array of digital content being left behind when people die [64, 65], and by extension in theoretical discussions around how designers can

consider time [66].

3.5 Summary

This background section has outlined a field of inquiry starting with sustainability and focusing on the SHCI concept of the cornucopian paradigm. Applying this to digital media suggests a particular media ecology that can be explored and enriched through alternative approaches to sound and HCI. In the following section I will outline how my theoretical framing and design methodology has attempted to approach this area through the creation of a set of designed artifacts.

4 Making Weird Things

4.1 Theory

While the motivation and conceptual background for this thesis stemmed primarily from the sustainable computing world, the actual work was undeniably that of design. Thus the theory and method will focus heavily on the act of making.

Early in my PhD I took a course about research methods where, during the introduction, I flippantly described my methodology as “making weird things”. In perhaps one of the most radical moments of my PhD journey, the professor refused to let my casual definition slide but rather engaged with it as a reasonable approach to the exploration of concepts and the generation of knowledge. The legitimisation of my self-declared outsider status had the dual effect of both disarming the performative cynicism that I was clearly using as a defense mechanism and forcing me to question exactly what it could mean. I have made weird things for some time, though if I interrogate that characterization I realize that the “weirdness” is most often not an intention but is a personal reaction to how I think others are perceiving what I have made. By describing what I have made as “weird” I am attempting to preempt how I believe it will be interpreted, even if my original explorations were driven by other intentions entirely. I never set out to make a musical pigsty or menstrual blood instrument *because* it is weird, even if I recognize that the outcomes will likely be interpreted as such.

From a design theory standpoint, this can be placed firmly within what Pierce refers to as *alternative designs* [67]. This definition groups a number of design theories (counterfactual, frictional, adversarial, critical, etc), positioning them all as approaches which resist “progressional” approaches to design which are intended to solve a problem or user need. According to Pierce, alternative designs can be defined by how they “prefigure possible, though improbable productions”, and function precisely because they are non-progressional and not intended for further

development. To drive this home, he provides a useful quote intended to help those attempting to justify this approach in their work:

> Our approach is frictional, our aims are not strictly progressional, and our contribution is transproductional. We prefigure production, but as a way to imaginatively engage with the future. If you take these designs literally as progressional new product development or R&D, then you're missing some, most, or all of the point!

My "making weird things" stance fits rather neatly into this framework, stemming as it does from the desire to create designs that can stay weird (rather than function as a prototype for further development). This framing also allows for an understanding that frictional designs function best as an interruption, an open question, a way of challenging the very assumptions that opened up the design space to begin with.

Within this space my explorations have borrowed heavily from the politically oriented aspects of DiSalvo's adversarial design [68] as well as the more critical/speculative [69, 70] approaches of Dunne and Raby and the counterfactual design of Wakkary and others [71, 72]. As a slight distinction from these, a crucial element that informs my "making weird things" approach is how every one of my designs, with the partial exception of the Period Instrument, is the result of working towards something that I want for my own use. I am creating things that I myself want to play with.

4.2 Method

There is, of course, a clear autoethnographic approach to this work. The design-led nature of this research was informed by my own background as an artist, performer, and "maker" working primarily in sound and technology. The 15 years in between my masters degree, where I first learned to code audio systems and build interactive electronics, and the beginning of my PhD were spent developing and refining the skills (and to a certain extent the concepts) that I then applied to my PhD research. As such I enthusiastically embrace the fact that much of what is portrayed here as "background" or "context" was in fact revealed throughout the design process, and conversely many of the concepts that were used as background to generate the design space from which a project emerged were later discarded. The work therefore is a product of a long and irregular process, spanning many years and different contexts. The outcomes, both in terms of design and knowledge, reflect this temporal irregularity and are enriched by it. The temporal vocabulary developed by Oogjes et al [73] is particularly useful for this.

Thus my PhD has lent heavily towards a Research through Design (RtD) process. This is an approach that prioritizes the practice of design as a method for generating knowledge [74]. In particular this work engages with the Gaver definition of RtD[75] in which the manifestation of the concepts is prioritized and is seen to create and contain knowledge in of itself [76]. Of course, describing a set of designs using words on paper will always be a poor substitute for holding an physically manifested instrument or listening to the sound of your own memories as they disappear.

Working with sound, in particular, is especially challenging as it resists being presented visually at all.

To address this disconnect I have reached most often towards the concept of annotated portfolios by Gaver [77], Löwgren [78], Bowers [79], and others. With this approach a series of designs are described and I attempt to tease out a set of themes and "family resemblances to be reasoned about". The use of this technique was well suited to my PhD work as it is often comprised of a number of interrelated designs that each contribute to a larger set of ideas, and is also a deliberate choice to help distinguish the work from other design or science traditions which might come with embedded statistical, solutionist, or data-driven expectations.

This thesis will therefore take on many of the characteristics of an annotated portfolio. I will present the papers I have included and a set of designed artifacts. These will be analyzed to present a set of particular design qualities that they contain. I will then argue that these design qualities inform a specific set of reconfigurations.

4.3 Limits and Constraints

Rejecting assumptions is a fantastic design exercise. It can be a powerfully generative practice to identify or uncover characteristics of a medium and use that as a limitation. What kind of game would you design for people who can not see, what kind of music can you make with only two notes, what kind of internet can you design with no electrical grid - imposing these types of limitations functions both as a way of revealing the biases in the medium as well as encouraging creative thinking. In my PhD design work I have used the rejection of the cornucopian paradigm as a starting point: what kinds of designs can emerge if we do not accept assumptions of growth and permanence, and how can these new designs point towards new relationships with media?

This approach reveals how assumptions of growth are built into our technologies. These assumptions of growth are intertwined with unreasonable and impossible personal responsibilities, much like the sustainability discourse throughout society today. In the resulting designs, the decentering of growth allows space for a different form of meaning-making and engagement with media and computing technology.

In practice, a rejection of growth means an embrace of limits, as described in my background, and much of the design work presented here is the result of an applied constraints-based design approach [25, 80, 81]. This represents the integration of sustainability (via SHCI) into the practical design method as it was applied in this research.

4.4 Presentation of the work

To summarize, I would like to position the theoretical framing of this research as a form of alternative design, suggesting that by "making weird things" I am proposing (and producing) designs that are overtly non-progressional but nevertheless concrete prefigurations of a desired future. The work followed a Research through Design

methodology, where the creation, presentation, and analysis of design artifacts was prioritized. The designs themselves were the result of a sustainability-led process embracing limits and constraints.

This thesis will now continue with a detailed summary of the included papers, followed by a presentation of each of the designed artifacts that these papers discussed. These will be presented separately, and their descriptions will primarily summarize how they function.

I will then present five specific qualities that emerge from the designs. These are Intermediate Time, Firehose, Play/Destroy, In the Box, and External Memory. These will form the basis of a discussion regarding what these designs *do*. I will frame this as *reconfigurations*: points at which the relationship between human and media is somehow altered, warped, or even broken. The thesis will show how the qualities in my designs can encourage particular reconfigurations of the human-media relationship.

5 Work

5.1 Included papers

5.1.1 Paper A: Music Within Limits

Yann Seznec. 2022 International Conference on ICT for Sustainability (ICT4S)

My first published paper is an examination of how music technology can (or indeed, should) be viewed through the lens of contemporary sustainable computing scholarship. The paper uses as a starting point the example of Alexander Graham Bell and Clarence J Blake's "ear phonautograph", a 19th century tool used for transcribing sound using a disembodied human ear. This instrument formed a fundamental part of Bell's research into sound leading to his invention of the telephone, which was instrumental in the development of electromagnetic transmission (and eventually recording) of audio.

The provenance of the ears that were used in these experiments is vague at best, and it is extremely likely that the body parts were taken from unwilling donors. As with all technologies, audio recording has been shaped by its history: our microphones are listening through dead, un-consenting ears. This forms the basis in the paper of an analysis of the political ecology of digital music making, looking at some of the technical and material sustainability issues in music technology. Following this, the paper speculates on approaches to making musical tools, proposing a set of guidelines that could help frame the design process. These include considering different power sources, decoupling audio and visual stimuli, breaking the myth of neutrality, and using meaningful content as musical material.

More broadly, I include this paper in the thesis as a way of overtly situating the design work it preceded within a context of sustainable computing and political ecology. As my first academic publication I look back on this paper with a mixture of pride and embarrassment, the latter in particular with regards to the final section where I

introduce the design guidelines. In retrospect these were unnecessarily specific, however the resulting artifacts that I built in the years following the publication of this paper did generally reflect those guidelines (apart from the suggestion to consider alternative power sources, which I would still very much like to explore in future work). Meanwhile the political stance of the paper remains relevant to this thesis and is something of a touchstone for the work as a whole. This paper therefore functions both as a stance and as framing for the design process that unfolded throughout my PhD studies.

5.1.2 Paper B: The Memory Cloud: Personal media libraries as affordance and constraint

Yann Seznec. 2025 International Conference on New Interfaces for Musical Expression (NIME)

The Memory Cloud is a digital instrument that transforms a user's cloud media library into musical content. The ramifications of this are twofold, with this paper analyzing how the use of personal media acts as a design constraint as well as an affordance. This is based on Magnusson's analysis of constraints and affordances in musical instrument design [49]. This work is unique among the included papers in the sense that it is the only one that discusses the use of the device in any depth by someone other than myself. The paper describes how I worked with a musician to create a version of the instrument with their own sounds, which was then used creatively by him and another musical collaborator. My discussions with these musicians around how the instrument fit into and expanded their existing musical practice form the basis of the analysis.

On the one hand, providing a user with not only a great deal of their own media, but *every single sound they have ever recorded*, is not exactly a constraint. The design of the Memory Cloud, however, subverts that apparent glut by loading only a single sound from the library at a time and shuffling the sounds so that each sound will only be loaded once before it is loaded a second time.

The affordances provided by personal digital media, when used in a creative way like with the Memory Cloud, gesture towards a broader possibility - where commercial instruments or creative sound products are generally marketed in terms of "realism" or "quality", the approach outlined in this paper encourages the player to build a relationship not only with the enormity of their own media (and, by extension, their own life history), but also engage with it as a temporary, fleeting, and unstable thing. This aligns with Magnusson's explorations of the relationship between constraints and affordances, where working within the limitations of the possible points towards new potential features and experiences.

5.1.3 Paper C: Three Ways to Destroy a Sound

Yann Seznec. 2024 Nordic Conference on Human-Computer Interaction

Drawing upon SHCI research that looks at the design bias towards endless growth, here I present three designed exemplars that each enact destruction in slightly

different ways. These all use a set of code (the “sound destruction engine”) that is reconfigured in different ways depending on the physical controls and intended use case. The designs include a performance-oriented instrument which has all of the controls available, the “Kopiator” which offers playful interaction, and the “Memory Boxes” which are intended to contain (and destroy) the only remaining copy of a personally meaningful digital sound.

This paper frames the design process from a limits-based perspective, looking at how each designed artifact enacts destruction in a different way based on what constraints are placed on the interface and presented to the user. Each resulting device opens up different questions about how destruction can function and what it can accomplish. The Sound Destroyer instrument, for example, is analyzed with regards to how it is manipulated by a knowledgeable player, whereas the Kopiator functions best as a something closer to a philosophical toy.

5.1.4 Paper D: Play/Destroy: A portfolio of sound destruction devices

Yann Seznec, Andreas Lindegren, Frida Lindegren, Rob Comber. 2026 CHI conference on Human Factors in Computing Systems

This paper expands significantly on Three Ways to Destroy a Sound, looking to situate the work of destruction within the context of the myth of media permanence in computing design. Specifically the paper argues that computing technology is built on a cycle of promise of permanence followed by inevitable loss. Each successive generation of media storage is portrayed as a solution to the previous iteration, but the only true guarantee is that it too will one day be replaced. Through the presentation of an annotated portfolio of work exploring the permanent destruction of digital audio, we suggest that HCI would benefit from integrating the inevitable loss of digital media into the design process.

From a design theory perspective, this paper uses the concept of seamfulness [82, 83, 82], arguing that destruction of sound can act as a method of revealing the underlying nature of the media. This is accomplished by bringing listening and destruction together into a single interaction. The paper concludes with a call for HCI to embrace the inevitability of digital decay.

I am one of four co-authors on this paper. As the lead author I designed all of the artifacts in the annotated portfolio, and formulated the main arguments with regard to how the work relates to the cornucopian paradigm.

5.1.5 Paper E: All of Your Sounds

Yann Seznec. Unpublished manuscript

A second approach to engaging with the enormity of modern cloud media libraries resulted in *All of Your Sounds*, a listening device that generates an endless sound composition from a users library. It takes the form of a small domesticated object that can be turned on or off when desired, allowing the user to tune into their memories as they fade in and out and are never repeated. The paper uses the designed device to open up questions about how we engage with our content, in

particular discussing the “conversation with the material” that arises when the media is presented in this way.

Designing for massive personal media archives, whilst attempting to stay true to the framework of sustainable computing, was a particular challenge. Where the destruction-led devices have a fairly obvious (if superficial) entry point for discussion (“burn it all down!”), the Memory Cloud and All of Your Sounds are both somewhat more subtle. They are less confrontational, but nevertheless point towards new ways of configuring our relationships to media.

5.2 Design Portfolio

As this work used a Research through Design process, it is crucial to present the actual design artifacts that were produced, particularly as there is some overlap amongst the papers. The following section will therefore describe each design in some detail, focusing first on the core body of work and finishing with several tangentially related projects.

5.2.1 The Memory Cloud

The Memory Cloud is a standalone synthesizer instrument which contains the sounds from every video in a user’s cloud library. One sound file at a time can be loaded into a synthesizer engine for manipulation. The audio engine itself is a classic granular system [84] that breaks the sound apart into individual “grains” which can be reconstituted according to the player’s desires. When the audio library is loaded into the machine, the order of the samples is shuffled, and a dedicated button selects one sound at a time for manipulation. The player does not know which sound will be selected next, and choosing a new sound will mean leaving the current sound behind until every other sound has been loaded.

The granular engine that drives the instrument can be used with virtually any sound file - I developed the code in order to be flexible and easy to use in a wide variety of situations (such as my custom football-pitch instrument [85]). The Memory Cloud is, therefore, defined not by the sound engine itself but primarily by the content that it contains - a (generally enormous) sound library made from all of the videos in one person’s “cloud”. The preparation of this library is something of an undertaking, involving the exporting and converting of a much larger number of files than media library management systems are generally intended to support. Going through this process is the first stage in building a connection with the instrument itself, as it makes the user deeply aware of the scale and intensely personal nature of the library. When this process is completed, it is a standalone instrument removed from the global techno-social infrastructures and political ecologies that enabled the cloud library to exist, becoming a local and physical manifestation of what is generally an invisible, intangible, and formless archive.

Playing this instrument is an act that blurs the understanding of permanence and certainty that media libraries are generally intended to reinforce. Pressing the button to load the sound - already a loaded act, as it will leave behind the current sound until you have listened to every other sound in the archive - is a leap into the unknown.

Very often the first step of exploration involves adjusting the parameters in a way that will reveal what the new sound *is*: a dog bark? a child's shout? yet another "happy birthday" recording?

The use of granular synthesis in this instrument is based on the unique relationship the audio processing technique has with time. By breaking a sound into many thousands of fragments and processing them each separately before recombining them, time becomes a fluid characteristic that can be explored. You can slow down the laughter of a child indefinitely, to make it last minutes or hours. When you're ready, you can move on to the next recording and leave that laughter behind.

5.2.2 All of Your Sounds

All of Your Sounds (AoYS) is also a device which contains the sounds from a user's entire cloud video library. In this instance, rather than being something controllable, AoYS is a passive listening device. It automatically generates an endlessly variable audio soundtrack from the sound library, layering multiple sounds from disparate parts of a person's life.

This design took as a starting point the same library preparation as The Memory Cloud - it emerged from one of a few experiments with the large folder of sound files that I had created for that instrument. Thus as a standalone device that contains these sounds AoYS explores some of the same issues of scale, locality, and personal media as described above. It is in some ways even more self contained then the Memory Cloud, as it contains an amplifier and built-in speaker, and requires only a power supply to function. The device features only one control, a volume knob. These characteristics generate further distance and independence from external infrastructures, enabling it to be theoretically used by nearly anyone, at the cost of control or listener choice.

The sound that is generated is rather the result of a set of choices hard-coded into the audio generation engine. My goal was to create a system for playing numerous sounds simultaneously in a dynamic and ever-changing way, cycling through all of the files in the library. In service of this some of the choices were imposed technically, for example the number of concurrent voices was dictated primarily by the relatively limited memory of the embedded computing device. However for the most part the parameters were set by me as the designer, a process that most closely resembled algorithmic composition. The two most important parameters available were the range of potential random playback lengths and the shape and timing of the volume envelopes. Changing the values of these resulted in drastically different listening experiences, and choosing a satisfyingly tuned combination of them was a large part of the design process.

The resulting object is in some ways modeled after the concept of a domesticated object [86], which can sit comfortably among the daily habits and routines of a household. It could be characterized as something of a audio equivalent of a digital photo frame, providing sonic glimpses into a semi-curated personal library (whenever you choose to listen). However it also functions to show the limitations of sound in this regard, as it is difficult to see it performing the same ambient function as a photo

frame due to the particular nature of audio and how we listen. It takes energy and time to listen, which is perhaps not conducive to a passive domesticated experience. Instead, AoYS is a device that suggests a different engagement with media, something that allows a listener to engage with the scale of the media library as a fundamental characteristic, allowing a new relationship to emerge with with sonic content.

5.2.3 Sound Destruction Engine

This is the software behind The Sound Destroyer, the Memory Boxes, and a number of other physical designs I have developed throughout my PhD. It is a set of code made in Pure Data for loading a sound, playing it back with a potentially modified wavetable, and overwriting the result onto the original sound. The engine itself is designed to be flexible, with a large number of user-adjustable parameters controlling the glitching of the wavetable and the application of the destruction to the sound.

Its origins date to the Book of Knowledge of Impractical Musical Devices project (described below), when I first became interested in a system that would destroy a sound each time it is heard. My first attempts to implement this concept involved deliberately injecting the audio buffer (i.e., the location in the local memory which has loaded the sound file) with random noise each time the sound was played. This was an ineffective approach, which turned out to be both inefficient and unsatisfying - the noise itself ended up framing the sound more than giving the impression of destruction. It took a great deal of random data before the sound would functionally disappear, and even at that point the disappearance felt unrelated to the original sound.

My second approach was based on the idea of using the sound against itself, making the actual sonic material responsible for the destruction. This involved separating out the playback from the soundfile in the form of a wavetable - a relatively common approach to audio manipulation which enables rapid and drastic manipulation of the time domain. Inserting noise into the wavetable (rather than the sound itself) meant that the resulting audible glitches were coming from the source sound, rather than directly from random numbers. The resulting audio output can then (when desired) be fed back into the audio buffer, overwriting the original sound with any glitches that have been heard.

A parallel can be made here with the historical lineage of sound recording, whose origins focused on the act of making a “record” of the sound, often using smoked glass pulled under a stylus attached to a vibrating membrane, such as with the ear phonautograph described in Paper A. These so-called “recordings” could not be played back and were primarily made for the study of sound waves and vibration, though modern analysis has been able to recreate the sounds from the original drawings [87]. One of the challenges faced by the modern researchers is accommodating the uneven and non-standard speeds at which these recordings would be made. The digital wavetable that is used in the Destruction Engine can thus be compared to this act of pulling the smoked glass under the stylus.

Taking this specific approach results in some specific outcomes. The use of a

wavetable in Pure Data, for example, lends itself strongly to a looping playback system (as the wavetable itself is most easily driven by a looping sawtooth wave). For similar reasons, the glitches that are written to the wavetable tend to be linear, because the code uses random numbers fed into ramp generators. These two factors lend the resulting sonic output a particular sonic aesthetic which is applied in a similar way to different sounds, particularly when the files are of comparable length. The overwriting system, meanwhile, resulted in an unexpected fading out of the playback as the sound is destroyed. The decaying white noise that characterizes the final stages of most sounds being processed by the engine is the result of the particular assemblage of code and design choices that went into the software.

5.2.4 The Sound Destroyer

This is the instrument instantiation of the sound destruction engine, in the sense that nearly every user-controllable variable in the engine in made available with a user interface element (knob, switch, or button). As an instrument the Sound Destroyer is somewhat unwieldy and requires knowledge and practice in order to use. From a music technology standpoint it could be classified as a sampler or looper: it records sound which can then be manipulated instantly. The specific manipulations that are available are based on the wavetable glitching system as set up in the destruction engine.

While the interface of this instrument provides a number of knobs, switches, and buttons, implying a great degree of control, the Sound Destroyer is in practice best characterized as being semi-controllable or semi-responsive. While some binary buttons and switches will react immediately and as expected (record, permanence, etc), some of the main knobs will instead provide indirect control. The glitch controls, for example, set not the exact value of how the sound will be destroyed, but rather set the potential maximum random value that *could* be applied to the wavetable in order to destroy the sound.

The instrument therefore embodies an uneasy relationship with reproduction. Loopers and samplers are generally expected to faithfully play back the sound they just recorded, and instruments are generally expected to react to a user's input - musical instruments are after all waiting for human interaction to complete their existence [50]. The Sound Destroyer makes no such promises. Faithful reproduction is one choice, but is not an assumption.

5.2.5 Kopiator

The Kopiator records a sound and repeats it instantly, on loop, while it is steadily destroyed. Recording a short sound will result in the sound disappearing over the course of perhaps a minute, while a longer sound will take a correspondingly longer time. The microphone and the speaker are both built into the device, so the sonic experience is entirely self contained.

This design began as an exercise in constraints-based user interface design, namely attempting to make satisfying sonic interactions with one button and one microphone. The built interface thus contained a number of different prototype

interactions (a synthesizer that talked back, a sampler that turned any input into danceable techno, etc) before the sound destroyer was implemented. This design process informed the specific interaction mechanic that drives the Kopiator - pressing and holding the button to record a sound, and releasing it to play the sound.

The instrument runs on the Sound Destruction Engine as described above, with most of the parameters fixed. They are tuned specifically to create a relatively rapid interaction where a user will almost certainly hear their sound disappear completely. There is some variability to this, as the number of glitches inserted into the sound is fixed on a per-loop basis, meaning that the time it takes for the sound to disappear will depend on the length of the recording that is made. The Kopiator is a design that takes advantage of this very specific interaction - pressing and holding a button, and then releasing that button - to explore concepts of time and decay.

5.2.6 The Memory Boxes

A Memory Box contains the only copy of a sound. A button on the box, when held down, controls the playback of the sound. Each time the button is held, the sound will be played and also slightly destroyed (using the sound destruction engine described above). Each time the sound is played, therefore, will be the last time that the sound is heard in its current form. Eventually the playback of the sound will result in its disappearance.

This represents a different instantiation of the destruction engine. While the interface is comprised of only one button, similar to the Kopiator, the button is instead used to play (rather than record). More crucially, the user in this case is well aware of the sound that is contained in the device, because they put it there. Similar to the Memory Cloud and AoYS designs, the experience of using the Memory Box begins before the button is ever pressed: choosing, preparing, and transferring the sound into the device. Ensuring that the audio in the Memory Box is the only remaining copy of the sound is a particular challenge - does it remain on a backup disk? Did you message it to your parents? The design intention is only activated when these questions have been addressed, and you can be sure that pressing the button will enable you to hear a sound for the last time.

5.3 A note on other related work

The method section of this thesis makes clear that this thesis was not a linear process. This is not unusual, of course, as most dissertations portray a small slice of the full body of research, giving an impression of a finite and planned journey from problem space to proposed solution. In my case a number of related projects and designs that formed part of my research journey have not been included as papers in this dissertation, but I will mention them at certain points and it is worth clarifying them here.

5.3.1 The Book of Knowledge of Impractical Musical Devices

This was a set of three instruments built before my PhD studies which contain the first iterations of some of the concepts that developed into this thesis. They are

described in detail in a paper for SMC 2025 [88], and one of the instruments appears in the design portfolio presented in Paper D.

This project took as its inspiration the Book of Knowledge of Ingenious Mechanical Devices from the 12th century Arabic scholar and engineer Al-Jazari. It was originally conceived as something of a critical examination of the state of music technology, but it grew into a broader exploration of how limitations can expand our experiences with sonic media.

The instruments took the form of three "Volumes". The first used a library of sounds recorded on a single day. A real time clock in the device changed their order each day that the instrument was used, loading different sounds and randomizing some settings. Thus the original day was represented and remixed differently each time it was used. As my first experiments with large personal media libraries, this instrument foreshadowed the Memory Cloud and AoYS work.

The second Volume was a GPS-dependent synthesizer, which used an early version of the granular synthesis code that ended up being used in several instruments presented in this thesis. This instrument used the data from a built-in GPS sensor to set a number of parameters of the synthesis engine, meaning that the sonic output would be different depending on where it was being used. Apart from the technical audio similarities, this instrument also implemented constraints-based design principles (before I was aware of the term).

The third Volume was a device for listening to and simultaneously destroying a sound. This may sound familiar at this point, and certainly this project was the first place I explored and implemented this concept. During my PhD research I was able to revisit this software and design in a number of different ways, resulting in the variety of implementations seen in this thesis.

5.3.2 The Period Instrument

A side project with my PhD colleague Nadia Campo Woytuk, the Period Instrument is a menstrual blood responsive synthesizer. We constructed and "played" the instrument, albeit with fake blood due to research restrictions with regard to bodily fluids. In our NIME paper [89] we explored how the Period Instrument represents a particular application of time, in the sense that it is dependent on a cyclical and fundamentally human (or at least mammalian) form of time, but one which is generally ignored.

The instrument itself is a fully functional granular synthesizer (built using the same core code as the Memory Cloud), loaded with a set of sounds that were chosen aesthetically to feel organic. An RGB sensor housed in a custom made enclosure detects whether a liquid poured into the top of the system is blood-colored, and triggers grains of sound accordingly, which are heard through a set of speakers attached to the structure.

This system provided the basis to examine some fundamental tensions with regard to constraint-based instrument design, such as how the system can be subverted (by, for example, using wine rather than menstrual blood), and whether that in turn defeats

the intention. The constraint, in this case, was perhaps speculative in the sense that the instrument did not technically require blood in order to function. However the instrument itself was real, responsive, and playable.

The Period Instrument has a number of connections to sustainable computing as well, and in our paper we situate its design firmly within an alternative computing context that draws from feminist design in particular, though in this thesis I will mainly mention this work in relation to how my designs attempt to reconfigure our relationship with time.





The Sound Destroyer *(top)*

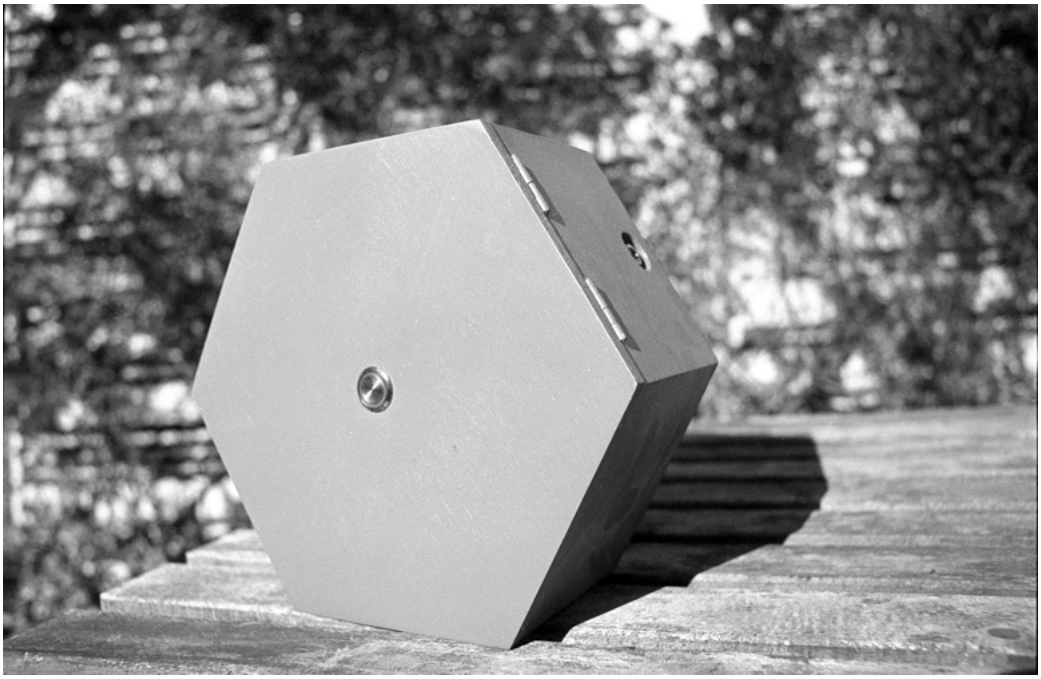
The Sound Destruction Engine *(left)*





Kopiator *(left)*

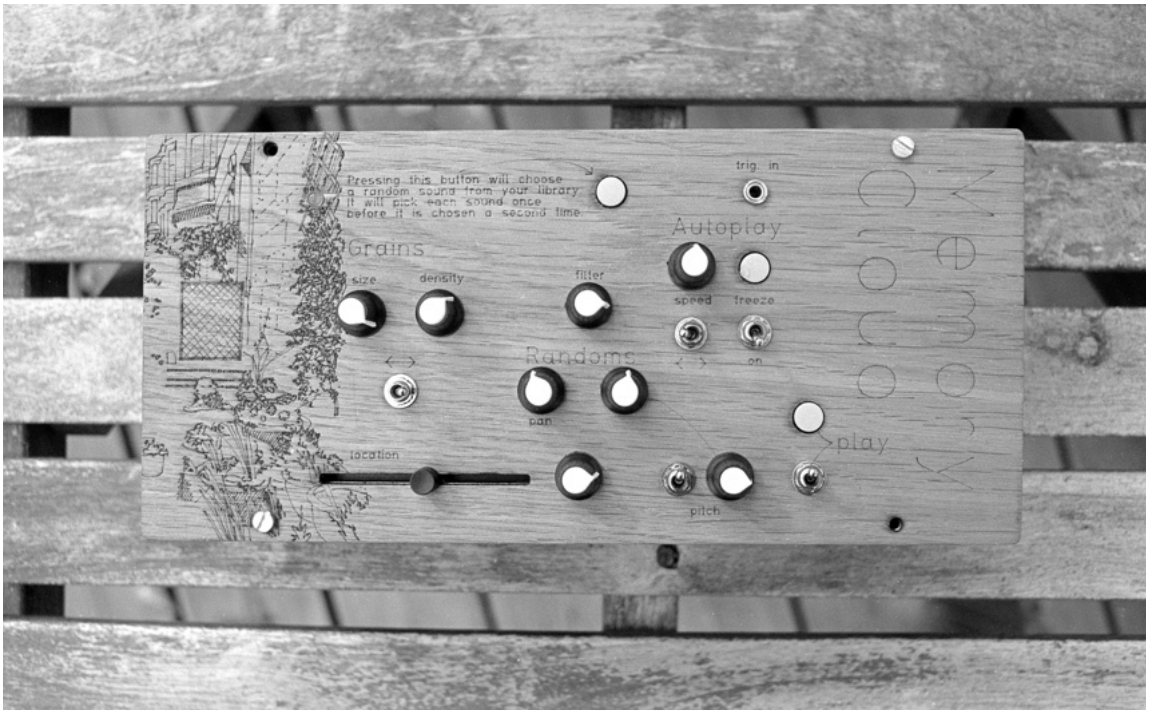
All of Your Sounds *(right)*





Memory Box *(top)*

Volume 3 from The Book of Knowledge of Impractical Musical Devices *(bottom)*





The Memory Cloud

6 Design Qualities

This section will present a set of specific qualities from the designs that I have included in this thesis. As both a designer and a user, these design qualities lead to a reconfiguration of my human-media relationship, from something where the media is at my service to a place where I become keenly aware of the size, weight, power, and meaning contained in my library.

6.1 Intermediate Time

In Victor Borges' micro-short story On Exactitude In Science, an empire becomes so obsessed with cataloging itself that only a 1:1 scale map will suffice - that is, a map the same size as the empire it was representing. Future generations deemed the map useless, and it was left to decay in the desert.

Listening to a sound recording takes time. It takes, in fact, the same amount of time that it took to make the recording. It is a 1:1 scale map of the time that the recording is representing. A 1:1 scale map of a place is absurd, a concept fit only for a paragraph of magical realism, yet a 1:1 scale map of time is considered a reasonable method for capturing and remembering something. The problems inherent in each are identical: like a 1:1 map of time, listening to a sound recording demands the same amount of time as the recording took to make. We have to *make* time.

The Kopiator is an exploration of this concept. A single button on a metal enclosure, with a microphone protruding from the top, looking vaguely reminiscent of a communication module from a midcentury space command center or the talkback mic in a recording studio - the design offers little in terms of instruction or guidance, but then the feature set is also extremely limited. Pressing and holding the button will make a recording, and releasing the button will play back that recording on loop, inserting and overwriting destructive glitches on each loop until the sound disappears.

The interaction can be funny, especially when the recording is very short. Pressing the button for just a moment and saying “Hello” will immediately trigger a looped “Hello” which disintegrates relatively quickly “HelloHellozzzHzelloHeizuiloElllaosEallaizxsszssszsssszsss”. Perhaps because of peculiarly comedic feedback loop it quickly becomes something that can be played with in a rapid-fire way, creating a rhythmic interaction.

Holding down the button for more time generates a longer loop of sound, which invites an alternate form of listening. The glitches are introduced to the sound at a frequency dictated by the length of the recording, which means that an extended sound recording will take correspondingly longer to disappear. With a longer looped sound we move towards an ambient or minimalist listening experience, perhaps closer to Basinki’s long-form destruction [90]. The same system, and the same action, thus offers at least two wildly different experiences - comedy, and poignancy.

Sound recording is a time machine. A recorded sound contains a piece of time, and listening back to that piece of time enables (or forces) us to re-inhabit that time. To a certain extent this is true with all forms of media, but arguably photographs and books and videos represent the transfer of the energy of time to a visual medium that can be looked at and studied at leisure, freeing us from the need to re-inhabit the time they took to create. As sound recordings became something that could also be heard back, rather than simply seen, they took on time-based qualities missing from other forms of media.

The design of The Kopiator took the particular relationship between recorded sound and time as a starting point. I wanted to make a thing that helped underscore the relationship of recording sound and the relentless arrow of time [91], something that perhaps could drive home the inadequacy of recorded media as a attempt to recreate the past within the present.

The button press to record a sound in the Kopiator opens the window of time. Pressing the button down is also a guarantee that the button will at some point be released, and the amount of time elapsed between these two actions dictates the precise amount of time that will then be re-experienced, on loop. The user physically inhabits this time with their finger held on the button.

The release of the button instigates an immediate reflection of the amount of time that was captured. If you recorded three seconds of sound, you’ll hear those three seconds back immediately. This could be characterized in some ways as an echo, though one in which the distance to the theoretical echoing surface is dynamic - a three second sound projected onto a wall 514.5 meters away would repeat the sound back precisely as it ends. A 5 second would need to be projected onto that same wall from 857.5m to achieve the same affect. A time machine is, by definition, also a space machine.

The Kopiator extends this loophole in time by repeating it and, of course, destroying it. Each time the window of time is repeated and modified, it drives home the original moment of time which was captured. Listening to the sound disintegrate is a reminder that the moment that was captured is now in the past, as are all moments apart from right now, the tip of the arrow. Are you recording this moment now, as

you read? If not, it's already in a lost past. And if you are recording this, you'll need to *make the time* to listen back.

I present this quality as Intermediate Time, a characteristic formed from an alternative approach to media reproduction, one which embraces the tension inherent in the desire to recreate time.

The quality of Intermediate Time is enacted in several other designs. While not included officially in this thesis, The Period Instrument [89] used the concept of menstrual time as the starting point for a musical instrument that was designed exclusively to be played with menstrual blood, framing the instrument as dependent on a deeply human form of time imbued with a complex web of social and cultural values. Meanwhile the Memory Box embodies a form of Intermediate Time by linking the decay of the sound it contains to the active choice to listen (or not).

Every sound *decays* in an acoustic sense, reflecting around the space where it is being heard until the energy dissipates. In everyday listening that is nearly instantaneous, perhaps stretching to a few seconds in churches or mausoleums. Meanwhile even the least stable recording formats that we use for capturing sounds can last for years or decades before they fail. Somewhere in between the near-immediate dissipation of sound in a flurry of reflections and the decades-long disappearance of an mp3 file sits Intermediate Time, taking a window of physically embodied time and destroying it before our very ears.

6.2 The Firehose

I have so many sounds. It is likely that you do too. An unanticipated consequence of the development of smartphones is that sound libraries, once the exclusive domain of sound designers and niche hobbyists with reel-to-reel machines, have become oddly prosaic.

It is in some ways easier to ignore the sound on these videos - indeed it is notable that Apple's Photos app on iOS now keeps video files muted by default. Watching them without sound makes it easier to ignore their time-based nature: you simply swipe to where you want to watch, or glance quickly at a set of videos before choosing one to open. The sound for these videos has become a casualty of the need for HCI to present media as efficiently as possible. Recorded sound, as a medium, resists navigation. It is impossible to listen to all of these videos, it would simply take too much time. Instead they because a library of unheard sounds, a forest of unfallen trees.

With All of Your Sounds I made a device for listening to this library, while attempting to avoid a solution-oriented approach. I wanted a thing that provided another way to listen. Once you decide to make a machine that takes all of your sounds and presents them back to you, many of the design decisions are pretty straightforward. The sounds should be randomized - no, *shuffled*, so you hear them all once before you hear them again. The length of playback should also be randomized, the offset into the sound too, because otherwise you'd only hear the beginnings.

Running this system with even just one voice (in the traditional synthesizer design sense) is initially almost like a game, you try and identify every sound as it passes.

That's a baseball stadium, that's my son as a toddler, that's...I'm not sure what that is, now that must be my brother playing guitar. It's fun, in a way, though also tiring. Raising the number of simultaneous voices pushes the experience into something more overwhelming. It's impossible to keep up with the memories, they rush past you and new ones are presented faster than you can identify them. A dog's bark is overlaid over a recording of a brass band, ten years separate those recordings and your brain struggles to accommodate the two concepts at the same time, and meanwhile four more recordings have come and gone.

It's not a nice feeling, necessarily. It can be magical to hear small snippets of my life floating past. Relinquishing control over my experience of them is powerful, humbling, and disarming. But it can also be immensely frustrating when it feels like too many significant moments are passing me by, and I can't just sit with a single memory without another one coming fast on its heels. It's an emotional firehose.

After sharing this design online, I received a number of messages from people expressing a desire to have one for themselves. As one of the few people who has actually tried this device, my immediate reaction to this was: "*really?*" I am constantly tempted to emphasize to people that while it is a powerful experience, it is not always a pleasant one. On the other hand, perhaps that's just what memories are.

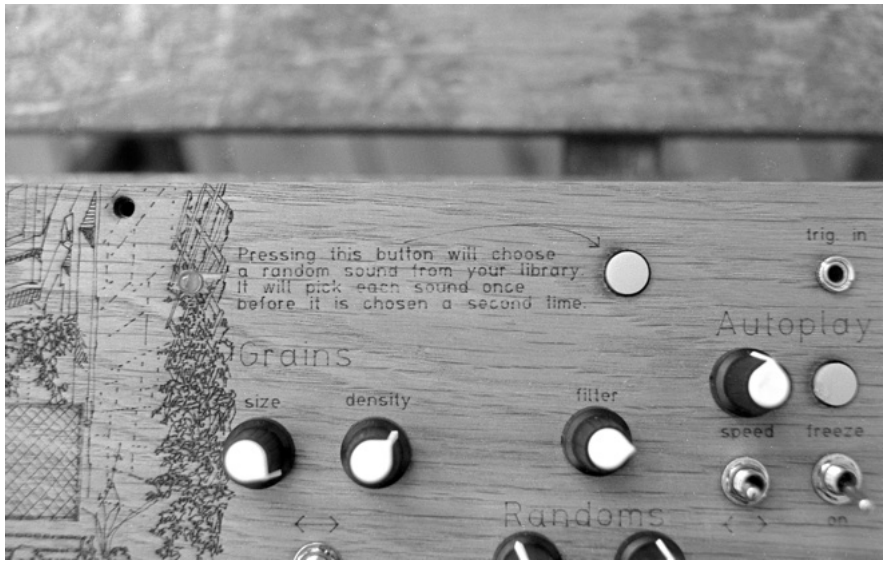
The Memory Cloud device is also based on a user's entire personal video library. Rather than dealing with this massive set of media by taking away control, the Memory Cloud offers the user the ability to play with a sound, taking it apart, stretching or squeezing it in any number of ways. However the user is not given total freedom - they can only play with one sound at a time, and they must load every single sound in the library before they will hear the same one twice. The pressing of the button for loading the next sound is thus a loaded decision: *am I ready to move on?*

This is also an overwhelming feeling, albeit one that inverts the association with time. Rather than being overwhelming because the sounds are rushing past you without end, in the case of the Memory Cloud it is overwhelming because you are stationary, inhabiting a single static moment within an enormous ocean of memories, deciding whether you are ready to start over.

The Firehose is the term I thus use to describe this design quality found in All of Your Sounds and the Memory Cloud. It is the feeling of being overwhelmed by a mass of personal media. Rather than providing the user with the ability to find a specific sound, the user is instead forced to reckon with the entire weight of their library, hearing sounds that they may not recognize or realize they had stored.

6.3 Play/Destroy

Roughly half of the design work presented in this thesis is about destruction. The act of listening and simultaneously permanently destroying a sound has been a running theme in my work for a number of years, dating back to my Book of Knowledge of Impractical Musical Devices project started in 2016 [88]. As described in Paper D,



the origins of this concept date from a conversation with an art conservator, who explained to me that visible light is the enemy of historical paintings. The visible light that touches a painting - allowing us to see it - is simultaneously robbing the painting of its pigment. A painting that is stored where there is no light will last far longer than one that is seen. The work of a conservator is thus finding a balance between a painting being seen and enjoyed by an audience versus its inevitable disappearance.

The promise of digital media is that it will never disappear, and it can be endlessly and losslessly copied. Developing an approach to destroying digital media within this context results in a design quality in which listening and destruction are a single entity. More than being linked, they are one action and one experience, interleaved and inseparable.

From a technical standpoint, the bulk of this quality is achieved in the Sound Destruction Engine, by separating out the destruction from the storage of the sound. In digital sound parlance, the current state of the sound itself is kept in its own buffer, whereas the wave that will drive the playback of that sound is kept separately. That wave (or “wavetable”) is the one that is modified with glitches whenever a listening session is begun. These two buffers - the wavetable and the soundfile - are kept entirely separate until the playback is initiated. The destruction contained within the wavetable is potential which is only activated when the sound is heard.

Playback is the moment that brings these two together. The sound can only be heard when the wavetable is applied to it - the audible energy is generated by the interleaving of these two discrete buffers. A metaphorical comparison can perhaps be made to quantum physics and the observation of a particle collapsing the probabilities



into an observed reality, or more prosaically this process could be compared to the rendering of a scene in a game engine or video processing system. Keeping these two elements separate until the moment when they are heard enables the destruction and the listening to be a single event, both technically and experientially.

When describing the system (and more broadly, the concept behind the sound destruction work), I most commonly explain that the sound can not be heard without being destroyed. In the case of the Memory Boxes, among other devices, a button on the box functions as the interface that enables you to listen at the cost of loss. On several of the design instances this button is overtly labeled with “Play/Destroy”. However it is equally important to acknowledge the reverse of this quality - neither can you destroy the sound without hearing it. These are not two separate things happening at the same time: listening and destroying are interwoven into a single thing.

6.4 In the Box

I am conscious that to a large extent the designs I am discussed here can be summarized as boxes with buttons. Admittedly, nearly half of the images in my Play/Destroy paper are in fact pictures of a box with a button and a hand holding or pressing the box. I don’t feel any great need to defend my work on this level, as I am not claiming to be an industrial (or even graphic) designer, and the aesthetic properties of the objects I’ve made are not really a point of discussion in this thesis. I can point out, though, that my tendency to make boxes stems from some of the restrictions that I placed on my work from the beginning - a desire to move away from screens, mainly, and an attempt to make objects that fully enact singular

concepts.

A consequence of this approach has been the emergence of a quality around what it means to place media *in* a device. This is seen most clearly in the Memory Boxes, which are designed to contain the only remaining copy of a meaningful sound. In that case, the quality was formed as part of the lead-up to the use of the Memory Boxes, as described in Paper D. The placing of the sound into the boxes acted as the final in a set of steps, which involved selecting, listening, discussion, transferring, and deleting the sound. The actual "placement" of the sound took the form of a rather banal USB transfer between my laptop and the embedded computer board inside the Memory Box, but putting the sound *in the box* is better described as encompassing the entire cycle. That selection process generated meaning, understanding, and a new relationship with the sound recording.

The Memory Cloud and AoYS devices contain this quality too, as the user's entire cloud library is, once again, in the box. There are several layers to this. Similar to the Memory Box process, preparing the cloud library to be placed into the box takes time, effort, knowledge, and (in practice) guidance from someone who has done it before. The cloud library media systems that we now rely on are strongly biased against local storage, and by default the videos in particular are stored remotely until they are selected for viewing. The process therefore involves selecting the videos, exporting them to local disk, waiting for the system to catch up, waiting again for significantly longer while they download, and ensuring that there is enough room on the local disk to store the resulting files. To export the sound then requires knowledge about batch processing and audio file types. In some ways this is a political act, shifting control of the media from an invisible remote power towards one that is controlled locally, something that is reflected in the tensions and frictions embedded into the process.

The end result, once the sounds have been transferred into the devices, are objects which both "contain" and are distinct from the cloud libraries of the user. They are offline and can not be updated with the sounds of new videos as they are recorded, but rather those must be added manually if desired. This is both a deliberate choice to ensure that the devices remain stable and independent of external infrastructure, and is also realistically impossible due to the limitations of the commercial cloud library APIs. The boxes become a vessel for a snapshot of a library, a static version of the state of personal media at the moment it was made. Just as with the Memory Boxes, the design quality of placing the sound in the Memory Cloud or AoYS emerges mainly in the process, the preparation, and the friction that is built-in to the experience.

In all of these cases this design quality reflects the curious combination of inaccessibility and tangibility that characterizes our relationship with media. While our media may be easily accessible (given correct circumstances) it is not actually present. Pushing through the frictions that keep this in place, we can put our sounds In the Box, which brings them closer physically at the cost of making them less accessible. We know that the sounds are *in* the thing that we are holding, yet listening to them in a reproducible and normally expected way is challenging and in some cases impossible.

6.5 External Memory

In 2024 I was lucky enough to participate in a symposium in Umeå, Sweden, entitled Hurricanes and Scaffolding, where I introduced my work around sound destroying and showed one of my sound destruction machines. With a small group, I randomly selected a video from my personal cloud media library - a clip of my daughter ice skating. I exported the video to my local drive and deleted the video from the cloud permanently. I then exported the sound from the video and deleted the video file itself. We listened to the sound together. You could hear her voice briefly, perhaps a small giggle, the scraping of skates, and my own voice encouraging her at the end. I then took this soundfile and loaded it into my Sound Destruction Engine on my computer. At each stage, I made extra clear that I was doing my best to ensure that I was fully deleting all instances of the file (this involved checking my cloud library several times, checking to see if I had sent the video to anyone on a messenger platform that may have saved it, etc).

Once I was certain that there remained only this one remaining copy of the sound, loaded into the Sound Destruction Engine, I clicked Play/Destroy and we listened as the sound looped, glitched, was overwritten, and eventually faded to silence. The presentation ended and after a short discussion with some lovely insightful questions we moved onto a coffee break, at which point several people from the session discreetly found me and asked the same question: “You didn’t…really delete it, did you?”

As I had these conversations, nearly everyone also mentioned something specific about the video - how listening to only the sound made them notice different things, how my daughter was cute, how it was emotional to hear the moment disappear. I was struck by this, because I had to admit that this was a video I did not remember recording, and I’m not sure I had watched it even once after taking it. That video had been continually stored on a set of remote servers in unknown places around the world for several years, but it was not in my own memory. In my paper about All of Your Sounds (Paper E) I recount a similar instance where in the creation of my personal sound library I come across a video of my son as an infant - a video that I did not know existed and had not seen in nine years.

These videos were media that was *stored* by the definitions of HCI and computing, in the sense that they could be filed, recalled, edited, and otherwise form a part of my vast personal library. However I was unaware of their existence. AoYS makes me confront my own library, presenting me with videos I did not know existed, moving them into my own memory. Meanwhile, permanently destroying the video in Umeå exported it not just to me, but to everyone in the room who witnessed the process. It lives and exists far more now than it ever did before, in the memories of those who listened to it disappear. In both cases, my designs became a conduit for an alternate form of memory, where the videos moved from being apparently permanent but obscure towards being impermanent but present and even distributed across multiple people. Where the In the Box quality emphasizes the experience that emerges from preparing and “placing” sounds into a device, this External Memory quality illustrates how the devices in use reshape the media.

7 Discussion: Reconfigurations

7.1 Taking Time, Making Time

Sound as a medium is uniquely unsuited to cornucopian design principles, due to its fundamentally time-based nature. The design artifacts presented in this thesis have explored how sound is completely reliant on (and therefore particularly illustrative of) time, as suggested by Chion and others [92], and have looked at how the act of listening to a sound requires inhabiting the same amount of time as it took to record it.

On the one hand this is taking advantage of what is generally considered to be a downside of sound - the particular challenge it presents in terms of display and navigation. A library of photographs or digital videos can be perceived at a glance, leading to the standard "thumbnail" display paradigm, but a hard drive full of sound recordings could take hours to listen to, if not significantly more (the library I built for Paper B totalled over 80 gigabytes, or roughly 140 hours of sound). Rather than considering this a problem to solve, the design qualities here demonstrate what can emerge when this is considered a characteristic to embrace. By suggesting an intermediate timescale of decay, for example, a listener becomes aware of both the window of time that has passed and the particular nature of sound as inhabiting time and physical place. The combination of the qualities of Intermediate Time and the Firehose point out that recording sound is an assumption that we will one day be able to make the time to listen back. The qualities use the perception of sound to illustrate that not only are we storing media, we are also attempting to store time and space.

There is a materiality to this understanding as well, in that by reinforcing this connection between time and space we are also forced to acknowledge the actual physical space that our media takes up. The creation of the massive sound libraries that generate the Firehose and In the Box qualities are generated by converting the media from an abstract, invisible, and formless concept to a tangible, physical, and

unwieldy object. It is a process that stems from considering the particularly time-based nature of sonic media.

These qualities reach beyond solutionism. Rather than considering data as something that needs to be managed and presented in a consumable way, it engages with the fundamental characteristics of the media. When considering sound, this means an acceptance that we can not *make time* any more than we can *make space*. Recorded sound is a 1:1 scale map of time, built upon the fantasy that the moment we recorded has not really passed, we can re-experience it anytime, it's in my pocket right now, along with all of those other times. My designs use the design qualities of the Firehose and Intermediate Time in an attempt to break this relationship by accepting that our media *is* overwhelming. There is no hiding it, it is a feature (not a bug).

7.2 Embracing Impermanence

The anti-cornucopian principles that guided this work can be seen as a political stance that comes as a result of an acceptance of the fragility of our computing systems. Our media ecology is built on a cycle of promise and loss, where new technological solutions are offered to address the problems inherent in the previous (supposedly permanent) solutions. The designed alternatives presented in this thesis relate instead to concepts of degrowth computing. Degrowth literature stresses (admittedly in no small part due to negative associations with the term) that an anti-growth mindset should not be equated with negativity or suffering but rather offers opportunities for better options [18, 16, 93]. These theories posit that we can not rely on growth-based paradigms to solve the climate metacrisis, because the growth is precisely the root cause of the problem. Moving away from growth provides the potential for an altogether better system. As applied specifically to my work here, it's not as if the systems for storing personal media were working beautifully and just need to be tweaked or improved. They are based on deeply embedded assumptions of growth, and as such we will continue to be confronted with their instability until we are forced to reckon with their collapse.

The design qualities outlined here contain an embrace of impermanence. The Sound Destruction Engine which forms the basis of several of the physical artifacts (Sound Destroyer, Memory Box, and Kopiator) had its origins in a thought experiment about linking perception and decay as a reflection of that relationship in visual art. The design process revealed that the most effective way to achieve this was to fully interleave the act of listening with the act of destruction: Play/Destroy.

The quality of External Memory expands on this as well, illustrating the quality that emerges once the destruction has been enacted. This encompasses the full cycle, from selection and preparation of the media through to its eventual loss. When this quality is achieved the media has been effectively exported away from the systems that have kept it apparently accessible and permanent but practically unavailable and unstable, towards an external memory (or memories) in a very human sense.

Impermanence and instability are inherent qualities of computing and media technologies, and some of the standard HCI heuristics can be characterized as attempts to mitigate these characteristics in the service of user friendliness. The

Play/Destroy and External Memory qualities suggest that we could instead be designing *for* and *with* impermanence and instability. Embracing these qualities as an alternative to the assumption of growth can enrich and expand our experience of media.

7.3 Seamful Media

Our media ecology contains the assumption that our media should be seamlessly accessible and available when we choose. This assumption is a cornerstone of the cornucopian paradigm, leading to the creation of tools that require always-on broadband access, syncing across devices, and logarithmically expanding server space. The massive (and growing) digital infrastructure that enables these characteristics is designed to be seamless and invisible.

Two of the qualities outlined here, In the Box and the Firehose, emerge from a repositioning of the concept of access. These qualities come from specifically limiting access to the media by removing the ability to browse, search, or explore the archive. These two qualities are, it must be said, fairly confrontational. When a user experiences the In the Box quality they are certainly confronted with the scale of their media, the thousands of minutes of recordings that must be downloaded and processed in order to be placed into a device. The conversion from invisible and abstract to tangible and local is inherently frictional and challenging. Similarly, the Firehose can in some ways be described as aggressive: *you want to hear your past? Well, here are all of your sounds all at once, try to keep up...*

This confrontational approach to media display forms a type of seam [82, 94, 83], where the underlying qualities of the material come to the surface. I explored this deeply in Paper D, particularly with regards to the various designs that employed sound destruction. That paper describes how a *seamful* design process can lead to new forms of interaction and engagement with media. In particular, here I would like to suggest that In the Box and the Firehose represent a seamful approach to access. These qualities are perhaps confrontational and aggressive precisely because they reflect the issues contained in the media itself - issues that the cornucopian paradigm attempts to paper over. The media *is* enormous, overwhelming, and impossible to engage with. We *do* struggle to relate to it. These are simply characteristics of the media ecology which we ignore at our peril. A seamful approach to the media challenges this approach, even if it isn't always pleasant.

7.4 Media Instruments

The ear phonautographs of the 19th century were many things - groundbreaking, grimly fascinating, illustrative of embedded social hegemony, and to our modern tastes utterly repulsive. They performed an apparently crucial role in the development of our understanding of sound and, perhaps indirectly, the birth of audio technology. Limited to the recording of sound waves, and incapable of playing them back, the ear phonautograph was also undoubtedly a scientific *instrument*. The structure of the ear phonautograph was itself built from another instrument, most probably a dismantled microscope [95]. These machines were built in order to enable

a deeper understanding of the nature of sound. An instrument is, of course, a device for recording or measurement.

In perhaps the least groundbreaking sentence of this thesis, I can also confidently state that an instrument is also something for making music. This linguistic quirk, present not only in English, positions the word as describing an object that creates, rather than measures. Perhaps more accurately, it can be said that the word "instrument" can cover a gradient of meanings between something that measures and something that creates - synthesizers certainly have roots in scientific instruments such as signal generators and oscilloscopes, and navigational instruments provide data to which a pilot can respond in real time.

One common thread between the various forms of "instruments" is how they embody and contain knowledge in of themselves, as explored by Magnusson [50] and Baird [48], building on ideas from Ihde of how instruments enable an embodiment relationship with the world [96]. Acoustic instruments, for example, can carry an implicit knowledge, with strings and and vibrating wood enacting physics and the perception of sound and emotion before they are necessarily described scientifically.

The devices I present in this thesis can certainly be described as "instruments" across the entire gradient of definitions. To a certain extent this stems from my own history as someone who has played, built, and hacked musical instruments for performance and composition as a career. I also see my own instruments as part of a larger political ecology of technological instruments such as the ear phonautograph. But more specifically to this argument, my instruments offer an enactment of a specific set of anti-cornucopian values. They do this by reconsidering some of the core assumptions of how we should be presenting and experiencing media.

As a whole the design qualities presented in the previous chapter represent re-evaluations of what would be considered HCI heuristics. These heuristics emerge from, and are intertwined with, the cornucopian paradigm which biases towards the development of instant, seamless, endless, and permanent designs for media. Play/Destroy, for example, breaks the concept of digital reproducibility - pressing the button twice will not reproduce the sound again in the same way. In the Box forces a user to reckon with the enormity of their library, removing the illusion of seamlessness from the media experience. The Firehose works actively against the standard heuristics of media display, which suggests that media should be presented in a way that provides information and access.

The work of breaking, or at least avoiding, these heuristics, is accomplished by a fundamental shift *towards* the media itself. Rather than relying on abstractions of the media, the designs here are folding together the material and content of the media with the interaction and the interface. The Memory Cloud, for example, is on one level a granular synthesizer, but transforms when the personal sounds are placed In the Box. It takes on depth and risk and meaning when the button to load a new sound is pressed, because the player is being confronted with their own library. There is no "find" function here, you will get another sound from your own past, whether you like it (or even remember it) or not.

The devices that use sound destruction also work directly with the characteristics of the media by interweaving the act of listening and the act of destruction, becoming a single moment both experientially and technically which can not be unwoven - Play/Destroy. You hear to destroy and you destroy to hear. You can not rely on permanence and reproducibility, it simply does not exist in this system. The system is inherently unstable and impermanent, and that's what you hear.

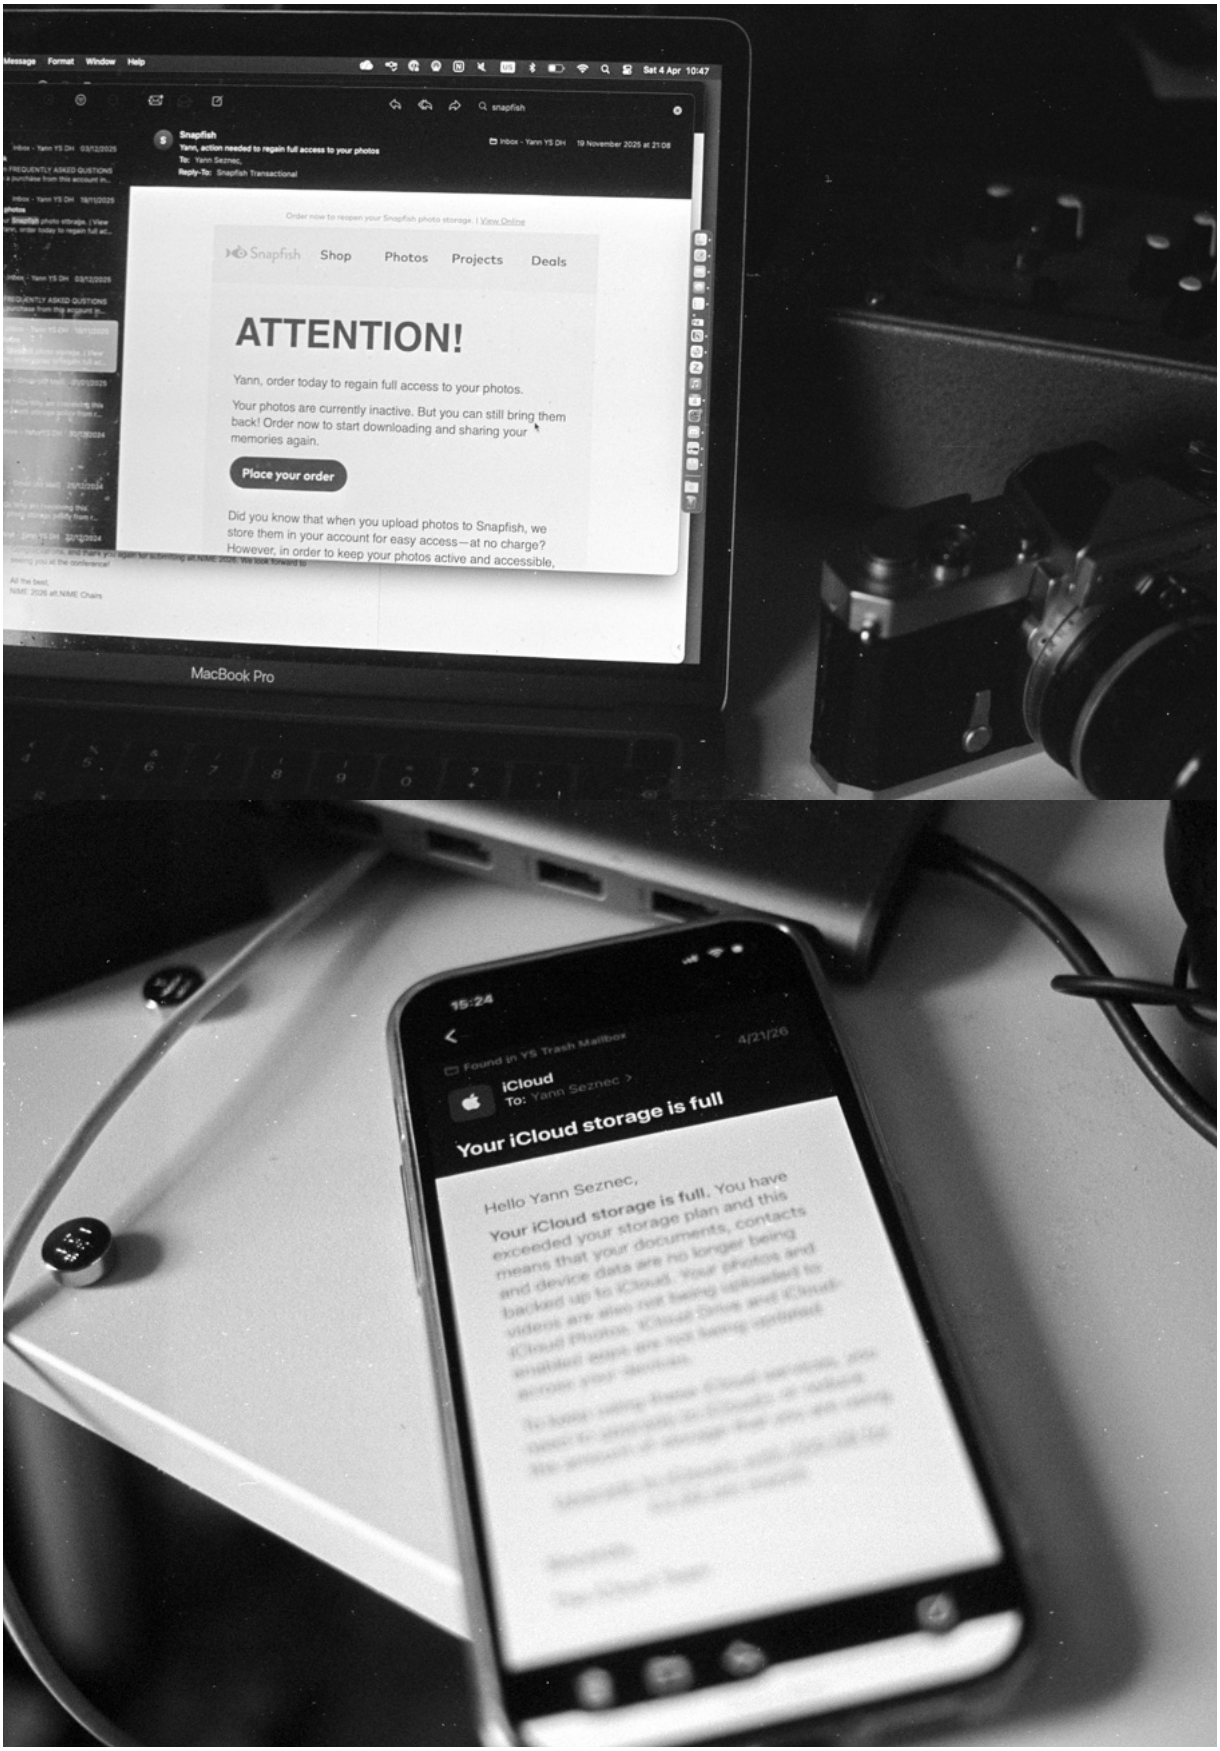
The designs here all engage with fundamental properties of the media that HCI, through the cornucopian paradigm, is primed to avoid. The sheer size of a personal media archive, the amount of time it takes to listen to a sound, the impermanence of digital media - rather than being something to avoid or solve, these become the material that the designs are made from. In each case, new qualities emerge from a folding together of the media and the interface of (or interaction with) the device, to the point where the media is inseparable from the instrument.

I use the word "instrument" deliberately there, as it is precise this folding together of media, content, and interaction that generates the qualities which aligns them with the characteristics described above. These are not devices for exploring, measuring, or presenting media, rather they are instruments that embody and enact a political stance that results in qualities that reconfigure our relationship to the media. The instruments are not separate from the media, they are a single epistemic force.

7.5 Summary

This discussion has looked across the qualities as a way of generating what I refer to as *reconfigurations*: ways that the designs can modify, reimagine, or even break our relationship with our personal media. In particular I have explored reconfigurations of our understanding of time, access, and seamlessness - values which are deeply embedded into our existing relationships with media. The instruments that form the portfolio in this thesis are able to enact these reconfigurations by interweaving the interface and interaction with the core characteristics of the media itself. This approach suggests an alternative approach to designing for media, one which engages with the fundamentally unsustainable nature of computing. In many ways this can result in awkward, challenging, or uncomfortable experiences, but comfort is a luxury we can not afford.

I would like position these reconfigurations as leverage points, in the Meadows sense [97]. Within the current computing and user experience paradigm, when space or permanence becomes an issue we are provided with a solution: buy to upgrade. Purchasing more storage will solve your problem. This is mirrored in many aspects of Western capitalist modern life, of course, and could be seen as a reflection of how these values become embedded into the design of technologies. If, however, we decide to embrace impermanence and limitations, aiming to use their inherent qualities to design otherwise, this is a leverage point that can topple a host of previously assumed requirements.


Snapfish
Yann, action needed to regain full access to your photos
To: Yann Seznec,
Reply-To: Snapfish Transactional
Inbox - Yann YS DH 19 November 2025 at 21:08
Order now to reopen your Snapfish photo storage. | View Online
Snapfish Shop Photos Projects Deals
ATTENTION!
Yann, order today to regain full access to your photos.
Your photos are currently inactive. But you can still bring them back! Order now to start downloading and sharing your memories again.
Place your order
Did you know that when you upload photos to Snapfish, we store them in your account for easy access—at no charge? However, in order to keep your photos active and accessible,
MacBook Pro
15:24
Found in YS Trash Mailbox
4/21/26
iCloud
To: Yann Seznec
Your iCloud storage is full
Hello Yann Seznec,
Your iCloud storage is full. You have exceeded your storage plan and this means that your documents, contacts and device data are no longer being backed up to iCloud. Your photos and videos are also not being uploaded to iCloud Photos. iCloud Drive and iCloud-enabled apps are not being updated across your devices.

8 Conclusion

In 2016, bleary from the birth of my first child and naïve about the long-term strategies and policies of digital media companies, I uploaded a number of photographs to several different photo printing services. I managed to order some mugs and a few books, before the time pressures of fatherhood outweighed my instincts for documentation. In the years that followed, I would periodically receive emails from these companies, encouraging me to relive my past: *"Remember these memories from five years ago? We thought you'd like to press rewind and relive these times."* The emails would conclude with a call to action, inviting me to memorialize that time with a fridge magnet and a mousepad.

More recently, however, these emails took a darker turn. The light, casual language became more formal; they were *reaching out* to me to let me know that their *policies* were changing. Initially, they were still really quite enthusiastic about this! They would really love to help me! I can still keep my photos active!

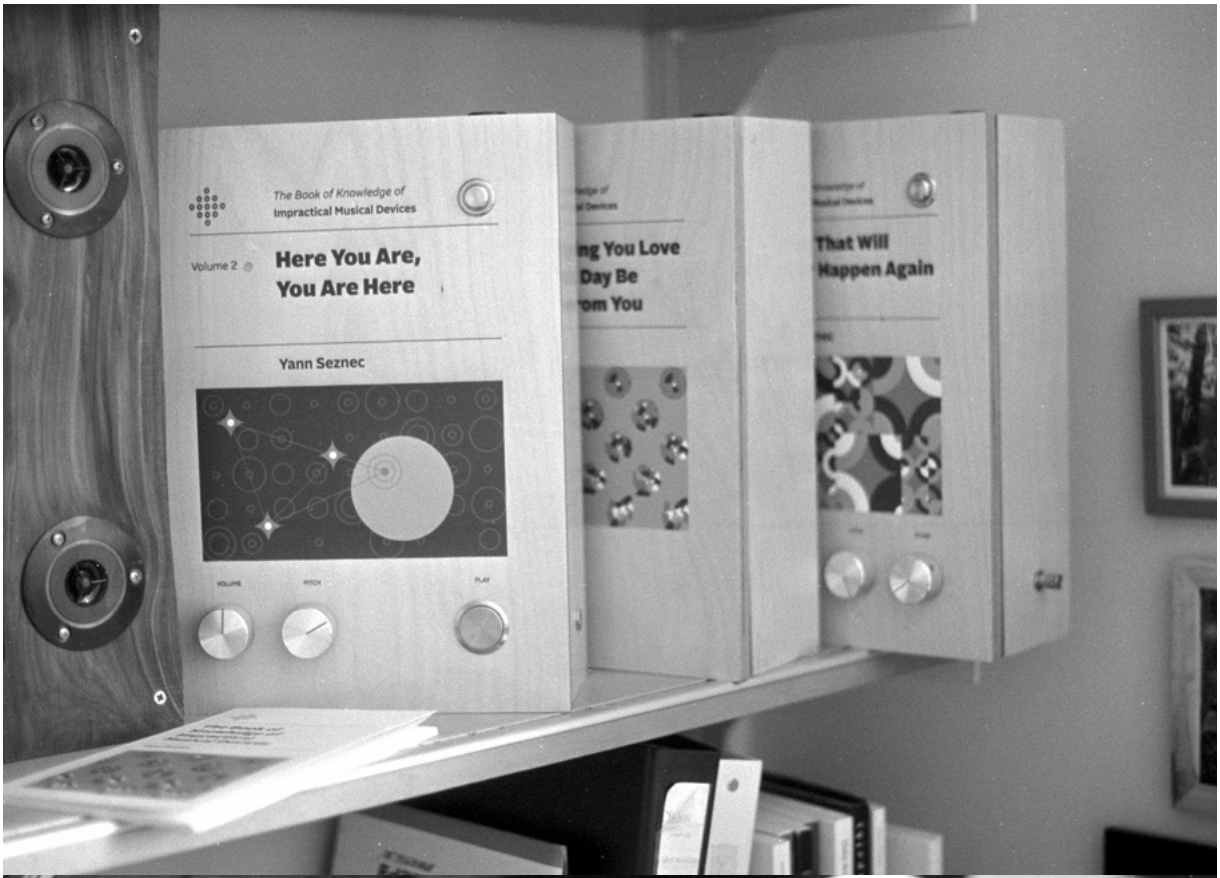
Eventually this enthusiasm was replaced with a "Final Notice" and a menacing all-caps "ATTENTION". My photos were going to be "archived". The actual mechanism of archiving was somewhat obscure, but my cherished digital media was definitely being threatened: buy a mousepad, or the memories die.

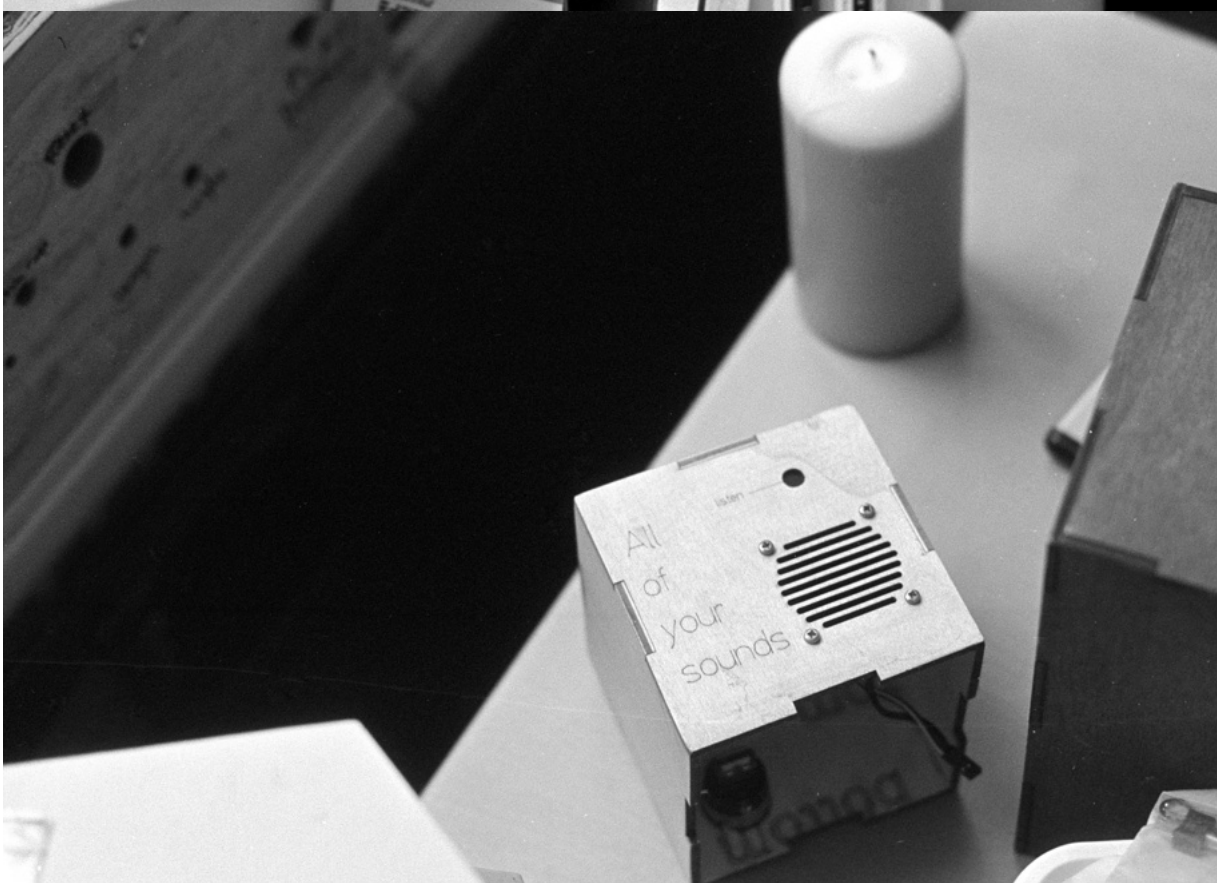
This experience was only striking because it ran counter to the notions of permanence and infinite storage that are built into digital media, which represent a denial of the realities of past, present, and future collapse. Computing technologies are built on promises of permanence, seamlessness, and endless access - concepts that are based on unsustainable assumptions of perpetual growth.

Even today, these promises are not fulfilled, and when they are broken we blame ourselves: our digital lives are currently full of moments where the responsibility of limiting our resource use is suddenly shifted to us personally. Our cloud library fills up, or our email inbox is full, and we are prompted to purchase an upgrade.

Technology encourages us to create and share media without limits, but then punishes us when we do so and expects the individual to make amends.

The designs in this thesis have explored an alternate approach, considering what might emerge if we eschewed the assumption of growth in favor of a strategy that embraces the fundamental characteristics of digital media. This approach results in new configurations of our relationship with media with regards to time, space, access, and permanence. I offer these reconfigurations as a contribution, alongside the design portfolio, and all of the code and resources that went into the creation of the artifacts. I hope that they can provide the basis for more design explorations in this space.


The Book of Knowledge of
Impractical Musical Devices
Volume 2
Here You Are,
You Are Here
Yann Seznec
VOLUME
PITCH
PLAY
That Will
Happen Again


All
of
your
sounds

Bibliography

[1] Yann Seznec. “Music Within Limits”. en. In: *2022 International Conference on ICT for Sustainability (ICT4S)*. Plovdiv, Bulgaria: IEEE, June 2022, pp. 172–178. ISBN: 978-1-6654-8286-8. DOI: 10.1109/ICT4S55073.2022.00028. URL: https://ieeexplore.ieee.org/document/9830084/ (visited on 12/01/2022).

[2] Yann Seznec. “The Memory Cloud: Personal media libraries as affordance and constraint”. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*. Ed. by Doga Cavdir and Florent Berthaut. Canberra, Australia, June 2025, pp. 446–451. DOI: 10.5281/zenodo.15698924. URL: http://nime.org/proceedings/2025/nime2025_65.pdf.

[3] Yann Seznec. “Three Ways to Destroy a Sound”. en. In: *Adjunct Proceedings of the 2024 Nordic Conference on Human-Computer Interaction*. Uppsala Sweden: ACM, Oct. 2024, pp. 1–7. ISBN: 979-8-4007-0965-4. DOI: 10.1145/3677045.3685496. URL: https://dl.acm.org/doi/10.1145/3677045.3685496 (visited on 03/10/2025).

[4] Yann Seznec, Andreas Lindegren, Frida Lindegren, et al. “Play/Destroy: A portfolio of sound destruction devices”. en. In: *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. Barcelona Spain: ACM, Apr. 2026, pp. 1–13. ISBN: 979-8-4007-2278-3. DOI: 10.1145/3772318.3791701. URL: https://dl.acm.org/doi/10.1145/3772318.3791701 (visited on 04/13/2026).

[5] Yann Seznec. “All of Your Sounds”. Submitted Manuscript 2026.

[6] World Commission on Environment and Development, ed. *Our common future*. Oxford paperbacks. Oxford ; New York: Oxford University Press, 1987. ISBN: 978-0-19-282080-8.

[7] Lon Åke Erni Johannes Hansson, Teresa Cerratto Pargman, and Daniel Sapiens Pargman. “A Decade of Sustainable HCI: Connecting SHCI to the Sustainable Development Goals”. en. In: *Proceedings of the 2021 CHI*

Conference on Human Factors in Computing Systems. Yokohama Japan: ACM, May 2021, pp. 1–19. ISBN: 978-1-4503-8096-6. DOI: 10.1145/3411764.3445069. URL: https://dl.acm.org/doi/10.1145/3411764.3445069 (visited on 10/18/2021).

[8] Christoph Becker, Ruzanna Chitchyan, Leticia Duboc, et al. "Sustainability Design and Software: The Karlskrona Manifesto". en. In: *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*. Florence, Italy: IEEE, May 2015, pp. 467–476. ISBN: 978-1-4799-1934-5. DOI: 10.1109/ICSE.2015.179. URL: http://ieeexplore.ieee.org/document/7202997/ (visited on 03/17/2026).

[9] Richard Heinberg and Daniel Lerch, eds. *The post carbon reader: managing the 21st century's sustainability crises*. OCLC: ocn681376694. Healdsburg, Calif. : Santa Rosa, Calif. : Berkeley, Calif: Watershed Media ; Post Carbon Institute ; Distributed by the University of California Press, 2010. ISBN: 978-0-9709500-6-2.

[10] Will Steffen, Katherine Richardson, Johan Rockström, et al. "Planetary boundaries: Guiding human development on a changing planet". en. In: *Science* 347.6223 (Feb. 2015), p. 1259855. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1259855. URL: https://www.science.org/doi/10.1126/science.1259855 (visited on 01/12/2023).

[11] Andrew L. Fanning and Kate Raworth. "Doughnut of social and planetary boundaries monitors a world out of balance". en. In: *Nature* 646.8083 (Oct. 2025), pp. 47–56. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/s41586-025-09385-1. URL: https://www.nature.com/articles/s41586-025-09385-1 (visited on 03/23/2026).

[12] Boris Sakschewski, Levke Caesar, Lauren Andersen, et al. *Planetary Health Check 2025: A Scientific Assessment of the State of the Planet*. en. Tech. rep. Artwork Size: 144 pages, 22 MB Medium: application/pdf Publication Title: Planetary Boundaries Science (PBScience). Potsdam Institute for Climate Impact Research (PIK), Sept. 2025, 144 pages, 22 MB. DOI: 10.48485/PIK.2025.017. URL: https://publications.pik-potsdam.de/pubman/item/item_32589 (visited on 03/23/2026).

[13] Eli Blevis. "Sustainable interaction design: invention & disposal, renewal & reuse". en. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. San Jose California USA: ACM, Apr. 2007, pp. 503–512. ISBN: 978-1-59593-593-9. DOI: 10.1145/1240624.1240705. URL: https://dl.acm.org/doi/10.1145/1240624.1240705 (visited on 10/18/2021).

[14] M. Six Silberman. "Information systems for the age of consequences". en. In: *First Monday* (July 2015). ISSN: 1396-0466. DOI: 10.5210/fm.v20i8.6128. URL: https://firstmonday.org/ojs/index.php/fm/article/view/6128 (visited on 11/16/2021).

[15] Bran Knowles, Oliver Bates, and Maria Håkansson. "This Changes Sustainable HCI". en. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Montreal QC Canada: ACM, Apr. 2018, pp. 1–12. ISBN: 978-1-4503-5620-6. DOI: 10.1145/3173574.3174045. URL: https://dl.acm.org/doi/10.1145/3173574.3174045 (visited on 10/18/2021).

[16] Vishal Sharma, Neha Kumar, and Bonnie Nardi. “Post-growth Human–Computer Interaction”. en. In: *ACM Transactions on Computer-Human Interaction* 31.1 (Feb. 2024), pp. 1–37. ISSN: 1073-0516, 1557-7325. DOI: 10.1145/3624981. URL: https://dl.acm.org/doi/10.1145/3624981 (visited on 12/18/2023).

[17] Federico Savini. “Post-Growth, Degrowth, the Doughnut, and Circular Economy: A Short Guide for Policymakers”. en. In: *Journal of City Climate Policy and Economy* 2.2 (Jan. 2024), pp. 113–123. ISSN: 2816-7414, 2816-7414. DOI: 10.3138/jccpe-2023-0004. URL: https://jccpe.utpjournals.press/doi/10.3138/jccpe-2023-0004 (visited on 02/08/2024).

[18] Jason Hickel. *Less is more: how degrowth will save the world*. eng. UK USA Canada Ireland Australia India New Zealand South Africa: Penguin Books, 2022. ISBN: 978-1-78609-121-5.

[19] Timothée Parrique. *Ralentir ou périr: l’économie de la décroissance*. Paris XIXe: Éditions du Seuil, 2022. ISBN: 978-2-02-150809-3.

[20] Bill Tomlinson, M. Six Silberman, Donald Patterson, et al. “Collapse informatics: augmenting the sustainability & ICT4D discourse in HCI”. en. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Austin Texas USA: ACM, May 2012, pp. 655–664. ISBN: 978-1-4503-1015-4. DOI: 10.1145/2207676.2207770. URL: https://dl.acm.org/doi/10.1145/2207676.2207770 (visited on 03/17/2026).

[21] Jay Chen. “A strategy for limits-aware computing”. en. In: *Proceedings of the Second Workshop on Computing within Limits*. Irvine California: ACM, June 2016, pp. 1–6. ISBN: 978-1-4503-4260-5. DOI: 10.1145/2926676.2926692. URL: https://dl.acm.org/doi/10.1145/2926676.2926692 (visited on 11/16/2021).

[22] Esther Jang, Matthew Johnson, Edward Burnell, et al. “Unplanned Obsolescence: Hardware and Software After Collapse”. en. In: *Proceedings of the 2017 Workshop on Computing Within Limits*. Santa Barbara California USA: ACM, June 2017, pp. 93–101. ISBN: 978-1-4503-4950-5. DOI: 10.1145/3080556.3080566. URL: https://dl.acm.org/doi/10.1145/3080556.3080566 (visited on 11/16/2021).

[23] Roel Roscam Abbing. “‘This is a solar-powered website, which means it sometimes goes offline’: a design inquiry into degrowth and ICT”. en. In: *LIMITS Workshop on Computing within Limits* (June 2021). DOI: 10.21428/bf6fb269.e78d19f6. URL: https://limits.pubpub.org/pub/lecuxefc (visited on 10/18/2021).

[24] Daniel Pargman and Björn Wallsten. “Resource Scarcity and Socially Just Internet Access over Time and Space”. en. In: *Proceedings of the 2017 Workshop on Computing Within Limits*. Santa Barbara California USA: ACM, June 2017, pp. 29–36. ISBN: 978-1-4503-4950-5. DOI: 10.1145/3080556.3084083. URL: https://dl.acm.org/doi/10.1145/3080556.3084083 (visited on 11/16/2021).

[25] Bonnie Nardi, Bill Tomlinson, Donald J. Patterson, et al. “Computing within limits”. en. In: *Communications of the ACM* 61.10 (Sept. 2018), pp. 86–93. ISSN: 0001-0782, 1557-7317. DOI: 10.1145/3183582. URL: https://dl.acm.org/doi/10.1145/3183582 (visited on 11/16/2021).

[26] Chris Preist, Daniel Schien, and Eli Blevis. “Understanding and Mitigating the Effects of Device and Cloud Service Design Decisions on the Environmental Footprint of Digital Infrastructure”. en. In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. San Jose California USA: ACM, May 2016, pp. 1324–1337. ISBN: 978-1-4503-3362-7. DOI: 10.1145/2858036.2858378. URL: https://dl.acm.org/doi/10.1145/2858036.2858378 (visited on 10/05/2021).

[27] Anaëlle Beignon, Aurélien Tabard, and Nolwenn Maudet. “How Interface Design Choices Lead to Indirect Environmental Impacts Through Use Intensification”. en. In: *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. Barcelona Spain: ACM, Apr. 2026, pp. 1–18. ISBN: 979-8-4007-2278-3. DOI: 10.1145/3772318.3790305. URL: https://dl.acm.org/doi/10.1145/3772318.3790305 (visited on 04/30/2026).

[28] Christian Remy and Elaine M. Huang. “Addressing the Obsolescence of End-User Devices: Approaches from the Field of Sustainable HCI:” in: Stockholm, Sweden, 2014. DOI: 10.2991/ict4s-14.2014.12. URL: https://www.atlantis-press.com/article/13430 (visited on 03/17/2026).

[29] Paul Dourish. “HCI and environmental sustainability: the politics of design and the design of politics”. en. In: *Proceedings of the 8th ACM Conference on Designing Interactive Systems*. Aarhus Denmark: ACM, Aug. 2010, pp. 1–10. ISBN: 978-1-4503-0103-9. DOI: 10.1145/1858171.1858173. URL: https://dl.acm.org/doi/10.1145/1858171.1858173 (visited on 11/04/2025).

[30] Yolande Strengers. *Smart energy in everyday life: Are you designing for Resource Man? | ACM Interactions*. 2014. URL: https://interactions.acm.org/archive/view/july-august-2014/smart-energy-in-everyday-life-are-you-designing-for-resource-man (visited on 09/24/2021).

[31] Yann Seznec and Sandra Pauletto. “The Singing Shower: A Melody-Sensitive Interface For Physical Interaction and Efficient Energy Consumption”. en. In: St Etienne, France, 2022, p. 8. DOI: 10.5281/zenodo.6798253.

[32] Hronn Brynjarsdottir, Maria Håkansson, James Pierce, et al. “Sustainably unpersuaded: how persuasion narrows our vision of sustainability”. en. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Austin Texas USA: ACM, May 2012, pp. 947–956. ISBN: 978-1-4503-1015-4. DOI: 10.1145/2207676.2208539. URL: https://dl.acm.org/doi/10.1145/2207676.2208539 (visited on 10/18/2021).

[33] James Pierce and Eric Paulos. “Beyond energy monitors: interaction, energy, and emerging energy systems”. en. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Austin Texas USA: ACM, May 2012, pp. 665–674. ISBN: 978-1-4503-1015-4. DOI: 10.1145/2207676.2207771. URL: https://dl.acm.org/doi/10.1145/2207676.2207771 (visited on 04/20/2022).

[34] Donella H. Meadows, Jørgen Randers, and Dennis L. Meadows. *The limits to growth: the 30-year update*. White River Junction, Vt: Chelsea Green Publishing Company, 2004. ISBN: 978-1-931498-51-7 978-1-931498-58-6.

[35] Magnus Ramage and Karen Shipp. *Systems Thinkers*. eng. SpringerLink Bücher. London: Springer London, 2009. ISBN: 978-1-84882-525-3. DOI: 10.1007/978-1-84882-525-3.

[36] Elina Eriksson and Daniel Pargman. "Exploring Inner Transition: Expanding Computing forSustainability". In: *Ninth Computing within Limits 2023*. Virtual: LIMITS, June 2023. DOI: 10.21428/bf6fb269.22dea4ad. URL: https://limits.pubpub.org/pub/ov23zjwm (visited on 06/27/2023).

[37] Minna Laurell Thorslund, Oliver Bates, Elina Eriksson, et al. "Meta-crisis computing and you: Finding agency through the Two Loops model of change". en. In: *Proceedings of the sixth decennial Aarhus conference: Computing X Crisis*. Aarhus N Denmark: ACM, Aug. 2025, pp. 127–139. ISBN: 979-8-4007-2003-1. DOI: 10.1145/3744169.3744178. URL: https://dl.acm.org/doi/10.1145/3744169.3744178 (visited on 05/06/2026).

[38] Kyle Devine. *Decomposed: the political ecology of music*. Cambridge, Massachusetts: The MIT Press, 2019. ISBN: 978-0-262-53778-0.

[39] Tara Rodgers. *Pink noises: women on electronic music and sound*. Durham [NC]: Duke University Press, 2010. ISBN: 978-0-8223-4661-6 978-0-8223-4673-9.

[40] Mathis Wackernagel, Niels B. Schulz, Diana Deumling, et al. "Tracking the ecological overshoot of the human economy". en. In: *Proceedings of the National Academy of Sciences* 99.14 (July 2002), pp. 9266–9271. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.142033699. URL: https://pnas.org/doi/full/10.1073/pnas.142033699 (visited on 01/11/2023).

[41] Jem Bendell and Rupert J. Read, eds. *Deep adaptation: navigating the realities of climate chaos*. Cambridge, UK ; Medford, MA: Polity Press, 2021. ISBN: 978-1-5095-4685-5.

[42] *Is it too late to prevent climate change? - NASA Science*. en-US. Section: Climate Change. Sept. 2014. URL: https://science.nasa.gov/climate-change/faq/is-it-too-late-to-prevent-climate-change/ (visited on 05/06/2026).

[43] Michelle Bastian. "Inventing Nature: Re-writing Time and Agency in a More-than-Human World". en. In: *Australian Humanities Review* 47 (Dec. 2009). ISSN: 18358063, 13258338. DOI: 10.22459/AHR.47.2009.10. URL: http://press-files.anu.edu.au/downloads/press/p40881/pdf/09.pdf (visited on 01/04/2023).

[44] Patricia Ciobanu and Oskar Juhlin. "Me, the Hill and My Browser – Investigating the Role of Time in Posthuman Interaction". en. In: *Nordic Human-Computer Interaction Conference*. Aarhus Denmark: ACM, Oct. 2022, pp. 1–12. ISBN: 978-1-4503-9699-8. DOI: 10.1145/3546155.3546651. URL: https://dl.acm.org/doi/10.1145/3546155.3546651 (visited on 01/04/2023).

[45] Daniel Fallman. "Design-oriented Human—Computer Interaction". en. In: *NEW HORIZONS* (2003).

[46] B. P. Abbott, R. Abbott, T. D. Abbott, et al. "Observation of Gravitational Waves from a Binary Black Hole Merger". In: *Phys. Rev. Lett.* 116 (6 Feb.

2016), p. 061102. DOI: 10.1103/PhysRevLett.116.061102. URL: https://link.aps.org/doi/10.1103/PhysRevLett.116.061102.

[47] Alexander Refsum Jensenius. *Sound Actions: Conceptualizing Musical Instruments*. en. The MIT Press. Cambridge: The MIT Press, 2022. ISBN: 978-0-262-54463-4 978-0-262-37220-6.

[48] Davis Baird. *Thing Knowledge: A Philosophy of Scientific Instruments*. eng. Berkeley, CA: University of California Press, 2004. ISBN: 978-0-520-23249-5 978-0-520-92820-6. DOI: 10.1525/9780520928206.

[49] Thor Magnusson. "Designing Constraints: Composing and Performing with Digital Musical Systems". en. In: *Computer Music Journal* 34.4 (Dec. 2010), pp. 62–73. ISSN: 0148-9267, 1531-5169. DOI: 10.1162/COMJ_a_00026. URL: https://direct.mit.edu/comj/article/34/4/62-73/94484 (visited on 03/06/2024).

[50] Thor Magnusson. *Sonic writing: technologies of material, symbolic and signal inscriptions*. eng. New York, NY: Bloomsbury Academic, 2019. ISBN: 978-1-5013-1386-8 978-1-5013-1388-2 978-1-5013-1387-5.

[51] Jakob Nielsen and Rolf Molich. "Heuristic evaluation of user interfaces". en. In: *Proceedings of the SIGCHI conference on Human factors in computing systems Empowering people - CHI '90*. Seattle, Washington, United States: ACM Press, 1990, pp. 249–256. ISBN: 978-0-201-50932-8. DOI: 10.1145/97243.97281. URL: http://portal.acm.org/citation.cfm?doid=97243.97281 (visited on 01/20/2025).

[52] Ben Shneiderman, Catherine Plaisant, Maxine Cohen, et al. *Designing the User Interface: strategies for effective human-computer interaction*. eng. Sixth edition. Always Learning. Boston, Mass. [u.a]: Pearson Education (US), 2017. ISBN: 978-0-13-438038-4.

[53] Cathal Gurrin, Gareth J. F. Jones, Hyowon Lee, et al. "Mobile access to personal digital photograph archives". In: *Proceedings of the 7th international conference on Human computer interaction with mobile devices & services*. MobileHCI05: 7th International Conference on Human-Computer Interaction with Mobile Devices and Services. Salzburg Austria: ACM, Sept. 19, 2005, pp. 311–314. DOI: 10.1145/1085777.1085842. URL: https://dl.acm.org/doi/10.1145/1085777.1085842 (visited on 01/14/2025).

[54] Joseph 'Jofish' Kaye, Janet Vertesi, Shari Avery, et al. "To have and to hold: exploring the personal archive". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI06: CHI 2006 Conference on Human Factors in Computing Systems. Montréal Québec Canada: ACM, Apr. 22, 2006, pp. 275–284. DOI: 10.1145/1124772.1124814. URL: https://dl.acm.org/doi/10.1145/1124772.1124814 (visited on 01/14/2025).

[55] Jay Cunningham, Gabrielle Benabdallah, Daniela Rosner, et al. "On the Grounds of Solutionism: Ontologies of Blackness and HCI". en. In: *ACM Transactions on Computer-Human Interaction* 30.2 (Apr. 2023), pp. 1–17. ISSN: 1073-0516, 1557-7325. DOI: 10.1145/3557890. URL: https://dl.acm.org/doi/10.1145/3557890 (visited on 03/24/2026).

[56] Lars Hallnäs and Johan Redström. "Slow Technology – Designing for Reflection". en. In: *Personal and Ubiquitous Computing* 5.3 (Aug. 2001),

pp. 201–212. ISSN: 1617-4909. DOI: 10.1007/PL00000019. URL: http://link.springer.com/10.1007/PL00000019 (visited on 01/11/2023).

[57] William Odom, Mark Selby, Abigail Sellen, et al. “Photobox: on the design of a slow technology”. en. In: *Proceedings of the Designing Interactive Systems Conference on - DIS ’12*. Newcastle Upon Tyne, United Kingdom: ACM Press, 2012, p. 665. ISBN: 978-1-4503-1210-3. DOI: 10.1145/2317956.2318055. URL: http://dl.acm.org/citation.cfm?doid=2317956.2318055 (visited on 06/21/2022).

[58] William Odom, Samuel Barnett, Jordan White, et al. “Becoming Watchful on the Trail and at Home: Understanding Experiential Outcomes of Capra in Long-Term Use”. en. In: *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. Barcelona Spain: ACM, Apr. 2026, pp. 1–29. ISBN: 979-8-4007-2278-3. DOI: 10.1145/3772318.3790407. URL: https://dl.acm.org/doi/10.1145/3772318.3790407 (visited on 04/16/2026).

[59] Amy Yo Sue Chen, William Odom, Ce Zhong, et al. “Chronoscope: Designing Temporally Diverse Interactions with Personal Digital Photo Collections”. en. In: *Proceedings of the 2019 on Designing Interactive Systems Conference*. San Diego CA USA: ACM, June 2019, pp. 799–812. ISBN: 978-1-4503-5850-7. DOI: 10.1145/3322276.3322301. URL: https://dl.acm.org/doi/10.1145/3322276.3322301 (visited on 11/03/2025).

[60] Jordan White, William Odom, Nico Brand, et al. “Memory Tracer & Memory Compass: Investigating Personal Location Histories as a Design Material for Everyday Reminiscence”. en. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Hamburg Germany: ACM, Apr. 2023, pp. 1–19. ISBN: 978-1-4503-9421-5. DOI: 10.1145/3544548.3581426. URL: https://dl.acm.org/doi/10.1145/3544548.3581426 (visited on 04/26/2023).

[61] Audrey Desjardins, Timea Tihanyi, Freesoul El Shabazz-Thompson, et al. “The Inner Ear: Capturing and Physicalizing Home Vibrations”. en. In: *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. Pittsburgh PA USA: ACM, July 2023, pp. 594–607. ISBN: 978-1-4503-9893-0. DOI: 10.1145/3563657.3596070. URL: https://dl.acm.org/doi/10.1145/3563657.3596070 (visited on 05/17/2024).

[62] William Odom and Tijs Duel. “On the Design of OLO Radio: Investigating Metadata as a Design Material”. en. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Montreal QC Canada: ACM, Apr. 2018, pp. 1–9. ISBN: 978-1-4503-5620-6. DOI: 10.1145/3173574.3173678. URL: https://dl.acm.org/doi/10.1145/3173574.3173678 (visited on 04/25/2024).

[63] William Odom, Ron Wakkary, Jeroen Hol, et al. “Investigating Slowness as a Frame to Design Longer-Term Experiences with Personal Data: A Field Study of Olly”. en. In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Glasgow Scotland Uk: ACM, May 2019, pp. 1–16. ISBN: 978-1-4503-5970-2. DOI: 10.1145/3290605.3300264. URL: https://dl.acm.org/doi/10.1145/3290605.3300264 (visited on 10/24/2025).

[64] William Odom, Richard Banks, and Dave Kirk. “Reciprocity, deep storage, and letting go: opportunities for designing interactions with inherited digital materials”. en. In: *Interactions* 17.5 (Sept. 2010), pp. 31–34. ISSN: 1072-5520, 1558-3449. DOI: 10.1145/1836216.1836224. URL: https://dl.acm.org/doi/10.1145/1836216.1836224 (visited on 04/25/2024).

[65] Jed R. Brubaker and Gillian R. Hayes. “”We will never forget you [online]”: an empirical investigation of post-mortem myspace comments”. en. In: *Proceedings of the ACM 2011 conference on Computer supported cooperative work*. Hangzhou China: ACM, Mar. 2011, pp. 123–132. ISBN: 978-1-4503-0556-3. DOI: 10.1145/1958824.1958843. URL: https://dl.acm.org/doi/10.1145/1958824.1958843 (visited on 04/25/2024).

[66] Verena Fuchsberger, Martin Murer, and Manfred Tscheligi. “Time and Design: Seven Sensitivities”. en. In: *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction*. Stanford California USA: ACM, Jan. 2015, pp. 581–586. ISBN: 978-1-4503-3305-4. DOI: 10.1145/2677199.2687911. URL: https://dl.acm.org/doi/10.1145/2677199.2687911 (visited on 04/25/2024).

[67] James Pierce. “In Tension with Progression: Grasping the Frictional Tendencies of Speculative, Critical, and other Alternative Designs”. en. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Yokohama Japan: ACM, May 2021, pp. 1–19. ISBN: 978-1-4503-8096-6. DOI: 10.1145/3411764.3445406. URL: https://dl.acm.org/doi/10.1145/3411764.3445406 (visited on 09/04/2023).

[68] Carl DiSalvo. *Adversarial design*. eng. Design thinking, design theory. Cambridge, Massachusetts: The MIT Press, 2015. ISBN: 978-0-262-52822-1.

[69] Anthony Dunne. *Hertzian tales: electronic products, aesthetic experience, and critical design*. eng. 1. MIT Press paperback ed. Cambridge, Mass. London: MIT Press, 2008. ISBN: 978-0-262-54199-2.

[70] Anthony Dunne and Fiona Raby. *Speculative everything: design, fiction, and social dreaming*. en. Cambridge, Massachusetts ; London: The MIT Press, 2013. ISBN: 978-0-262-01984-2.

[71] Ron Wakkary, Doenja Oogjes, Henry W. J. Lin, et al. “Philosophers Living with the Tilting Bowl”. en. In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Montreal QC Canada: ACM, Apr. 2018, pp. 1–12. ISBN: 978-1-4503-5620-6. DOI: 10.1145/3173574.3173668. URL: https://dl.acm.org/doi/10.1145/3173574.3173668 (visited on 10/27/2023).

[72] Ron Wakkary, William Odom, Sabrina Hauser, et al. “Material Speculation: Actual Artifacts for Critical Inquiry”. en. In: *Aarhus Series on Human Centered Computing* 1.1 (Oct. 2015), p. 12. ISSN: 2445-7221. DOI: 10.7146/aahcc.v1i1.21299. URL: https://tidsskrift.dk/ashcc/article/view/21299 (visited on 07/24/2023).

[73] Doenja Oogjes and Audrey Desjardins. “A temporal vocabulary of Design Events for Research through Design”. en. In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*. Honolulu HI USA: ACM, May 2024, pp. 1–12. ISBN: 9798400703300. DOI: 10.1145/3613904.3642560. URL: https://dl.acm.org/doi/10.1145/3613904.3642560 (visited on 05/17/2024).

[74] Johan Redström. *Making Design Theory*. en. The MIT Press, 2017. ISBN: 978-0-262-34184-4. DOI: 10.7551/mitpress/11160.001.0001. URL: https://direct.mit.edu/books/book/2277/making-design-theory (visited on 02/02/2024).

[75] William Gaver. "What should we expect from research through design?" en. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Austin Texas USA: ACM, May 2012, pp. 937–946. ISBN: 978-1-4503-1015-4. DOI: 10.1145/2207676.2208538. URL: https://dl.acm.org/doi/10.1145/2207676.2208538 (visited on 05/16/2022).

[76] Jeffrey Bardzell, Shaowen Bardzell, Peter Dalsgaard, et al. "Documenting the Research Through Design Process". en. In: *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. Brisbane QLD Australia: ACM, June 2016, pp. 96–107. ISBN: 978-1-4503-4031-1. DOI: 10.1145/2901790.2901859. URL: https://dl.acm.org/doi/10.1145/2901790.2901859 (visited on 03/18/2026).

[77] Bill Gaver and John Bowers. "Annotated portfolios". en. In: *Interactions* 19.4 (July 2012), pp. 40–49. ISSN: 1072-5520, 1558-3449. DOI: 10.1145/2212877.2212889. URL: https://dl.acm.org/doi/10.1145/2212877.2212889 (visited on 03/26/2023).

[78] Jonas Löwgren. "Annotated portfolios and other forms of intermediate-level knowledge". en. In: *Interactions* 20.1 (Jan. 2013), pp. 30–34. ISSN: 1072-5520, 1558-3449. DOI: 10.1145/2405716.2405725. URL: https://dl.acm.org/doi/10.1145/2405716.2405725 (visited on 03/26/2023).

[79] John Bowers. "The logic of annotated portfolios: communicating the value of 'research through design'". en. In: (2012).

[80] Martin Stojanov, Daniel Pargman, Mike Hazas, et al. "How do we arrive at constraints? Articulating limits for computing". In: *Ninth Computing within Limits 2023*. Virtual: LIMITS, June 2023. DOI: 10.21428/bf6fb269.a317d18f. URL: https://limits.pubpub.org/pub/oojdthry (visited on 06/27/2023).

[81] Fatemeh Bakhshoudeh and Rob Comber. "Designing with the Solar Internet: Towards Constraint-Based Design for Sustainable Consumption". en. In: *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Yokohama Japan: ACM, Apr. 2025, pp. 1–18. ISBN: 979-8-4007-1394-1. DOI: 10.1145/3706598.3713101. URL: https://dl.acm.org/doi/10.1145/3706598.3713101 (visited on 03/18/2026).

[82] Matthew Chalmers and Ian MacColl. "Seamful and seamless design in ubiquitous computing". In: *Workshop at the crossroads: The interaction of HCI and systems issues in UbiComp*. Vol. 8. 2003, p. 10.

[83] Matthew Chalmers. "Seamful Design and Ubicomp Infrastructure". en. In: ().

[84] Curtis Roads. *Microsound*. eng. Cambridge (Mass.) London: the MIT press, 2004. ISBN: 978-0-262-68154-4.

[85] Yann Seznec. "I am the Ball". en. In: *Proceedings of the 20th International Audio Mostly Conference*. Coimbra Portugal: ACM, June 2025, pp. 216–222. ISBN: 979-8-4007-2065-9. DOI: 10.1145/3771594.3771616. URL: https://dl.acm.org/doi/10.1145/3771594.3771616 (visited on 12/13/2025).

[86] Knut H. Sørensen and Merete Lie. *Making technology our own? domesticating technology into everyday life*. eng. Oslo: Scandinavian University Press, 1996. ISBN: 978-82-00-22758-8.

[87] Patrick Feaster. *Pictures of sound: one thousand years of educed audio : 980-1980*. eng. OCLC: 1013281501. Atlanta, GA: Dust-to-Digital, 2013. ISBN: 978-1-938922-23-7.

[88] Yann Seznec. "The Book of Knowledge of Impractical Musical Devices". In: *Sound and Music Computing 2025 (SMC 2025), Graz, Austria, July 10-12, 2025*. 2025, pp. 290–297.

[89] Yann Seznec and Nadia Campo Woytuk. "The Period Instrument". In: *Proceedings of the International Conference on New Interfaces for Musical Expression*. Ed. by S M Astrid Bin and Courtney N. Reed. Utrecht, Netherlands, Sept. 2024, pp. 91–99. DOI: 10.5281/zenodo.13904792. URL: http://nime.org/proceedings/2024/nime2024_14.pdf.

[90] Lindsay Balfour. "Traumatic Ruins and the Archeology of Sound: William Basinski's The Disintegration Loops". In: *Journal of Sonic Studies* 16 (2018). URL: https://pureportal.coventry.ac.uk/en/publications/traumatic-ruins-and-the-archeology-of-sound-william-basinskis-the.

[91] David Layzer. "The Arrow of Time". In: *Scientific American* 233.6 (1975), pp. 56–69. ISSN: 00368733, 19467087. URL: http://www.jstor.org/stable/24949962 (visited on 03/26/2026).

[92] Michel Chion. *Audio-vision: sound on screen*. eng. New York: Columbia University Press, 1994. ISBN: 978-0-231-07898-6 978-0-231-07899-3.

[93] Giacomo D'Alisa, Federico Demaria, and Giorgos Kallis, eds. *Degrowth: a vocabulary for a new era*. New York ; London: Routledge, Taylor & Francis Group, 2015. ISBN: 978-1-138-00076-6 978-1-138-00077-3.

[94] Matthew Chalmers and Areti Galani. "Seamful interweaving: heterogeneity in the theory and design of interactive systems". en. In: *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques*. Cambridge MA USA: ACM, Aug. 2004, pp. 243–252. ISBN: 978-1-58113-787-3. DOI: 10.1145/1013115.1013149. URL: https://dl.acm.org/doi/10.1145/1013115.1013149 (visited on 11/12/2025).

[95] Tom Everrett. "Writing sound with a human ear: reconstructing Bell and Blake's 1874 ear phonautograph". en. In: *Science Museum Group Journal* 12.12 (Nov. 2021). ISSN: 2054-5770. DOI: 10.15180/191206. URL: http://journal.sciencemuseum.ac.uk/browse/issue-12/writing-sound-with-a-human-ear/ (visited on 01/05/2022).

[96] Don Ihde. *Technology and the lifeworld: from garden to earth*. eng. Nachdr. The Indiana series in the philosophy of technology 560. Bloomington, Ind.: Indiana University Press, 1996. ISBN: 978-0-253-20560-5 978-0-253-32900-4.

[97] Donella Meadows. *Leverage points-places to intervene in a system*. Compiled by Dr. Chapman, 2015.

Included papers

Paper A

2022 International Conference on ICT for Sustainability (ICT4S)

Music Within Limits

Yann Seznec
Division of Media Technology and Interaction Design
KTH Royal Institute of Technology
Stockholm, Sweden
yannse@kth.se

2022 International Conference on ICT for Sustainability (ICT4S) | 978-1-6654-8286-8/22/$31.00 ©2022 IEEE | DOI: 10.1109/ICT4S55073.2022.00028

***Abstract*—The creation of music relies heavily on, and is tightly intertwined with, mass production and industrialisation, making it virtually inseparable from many of the fundamental issues that make up the climate crisis today. Music (and the arts more generally) has the potential to play a role in influencing behaviour and effecting structural change by showcasing the issues at play and raising awareness around the magnitude of the crisis. However, digital music making contains an embedded political ecology and problematic history which must be assessed and acknowledged in order to effectively leverage this power. Therefore this paper will situate digital music within a cultural history of industrialisation, with all of the ecology contained therein. Building upon this, a set of approaches will be laid out that draw from established ICT4S concepts of limits theory and heuristic design, combined with artistic practice, proposing potential ways forward that will encourage digital music making to move towards a sustainable future and position musical tools to challenge societal norms.**

***Index Terms*—sustainability, music, political ecology, design, limits, music technology**

Political Ecology of Music Tools

In 1874 Alexander Graham Bell and Clarence J Blake removed the middle ear from a human cadaver and attached it to an apparatus most likely cobbled together from available scientific equipment [11], with the aim of visually inscribing sound waves onto glass plates. They were successful in this endeavour, with the resulting drawings made by their "ear phonautograph" representing an incremental advancement in audio technology, forming a part of the movement towards the electrification of sound more generally:

> ...the human ear became, in the ear phonautograph, a legible point of consolidation in the emergence of a more general 'tympanic principle' that enabled the development of a number of historically significant sound technologies, the telephone being only one. [11, p. 24]

If we accept the importance of the ear phonautograph in terms of the technological development of audio and musical tools, we must also accept the social and cultural issues that it represents. The creation of this device was in no way neutral - it required a human death and a human body, and fitting with the general practice of western scientific research in the 19th century the provenance of this body is vague at best:

> These anonymous ears connect the distanced brutality that is fundamental to the development of modern sound technologies to the more viscerally violent practices of lynching, body snatching, public dissection, and capital punishment. And these are all violent acts that are committed on nameless, socially marginalized, objectified, and fragmented bodies. [23, p. 109]

This gruesome episode in the development of modern sonic technology illustrates how the history and culture of music follows the same patterns of industrialisation and exploitation seen in other technological sectors. Much as with other 19th century technological innovations, the development of sound recording was built on unjust societal structures and justified through problematic framings of a neutral scientific method and notions of progress. Similarly, the revolution of musical tools beginning in the mid-19th century and continuing to today closely parallels and is tightly intertwined with the modern industrial and digital revolutions [41]. The piano can be seen as a prime example of this - in 1853 the Chickering piano company was able to build the largest industrial building in the United States [20] due to technical innovations and manufacturing efficiencies associated with the industrial revolution [36]. This revolution would probably have happened without music, but modern music most certainly would not exist without it.

This relationship can be seen throughout the modern history of music - colonial exploitation allowed for the cheap harvesting of shellac to produce records in the 1920s [9], and the design of magnetic audio tape in the mid-20th century stemmed mainly from partnerships with the military sector [35]. Tara Rodgers describes how the entire approach to audio "reproduction" reflects a deeply gendered approach to music technology, with women representing and embodying roles that tie women to "age-old notions of passivity, receptivity, and materiality" [35]. Vágnerová and others [44, 4] have also analysed gender dynamics with regards to design and fabrication of technology, with deeply ingrained ideas of "nimble fingers, small hands, proclivity for working on the small scale, tolerance for tedium, [and] obedience", reinforcing subordinate roles for women in the sector. This can be seen materially in the fabrication of modern electric guitars, vinyl records, and drum machines being done primarily by women [35], while credit for both designing and using these items

This work is supported by the Sound for Energy Project (see: https://soundforenergy.net) which is funded by Swedish Energy Agency (Project No. 51645-1)

978-1-6654-8286-8/22/$31.00 ©2022 IEEE
DOI 10.1109/ICT4S55073.2022.00028

Authorized licensed use limited to: KTH Royal Institute of Technology. Downloaded on December 01,2022 at 11:47:27 UTC from IEEE Xplore. Restrictions apply.

is given almost exclusively to men. This reflects a broader pattern of social injustice within ICT, such as the exploitation of Navajo women in the design and fabrication of circuit board design and continuing through to the manufacture of modern integrated circuits [27].

A contemporary example of the social shaping of music can be found in arguably the most important, or at least most recognizable audio effect in popular music in the past 20 years: auto-tuning. The technology stems directly from the oil drilling sector, relying on feature extraction methodology first developed for reflection seismology (used to identify suitable locations for deepwater drilling) before being applied to audio [26]. It is no accident that a development history intertwined with non-renewable energy and climate change is also culturally and creatively problematic. A tool designed to iron out supposed musical imperfections is by definition ignoring countless musical traditions that take different approaches to pitch, tuning, performance, or harmony. It is both fitting and unsurprising that a tool inextricably linked with rampant capitalist expansion can also be seen to be creatively ethnocentric and colonialist. [31, 32, 8, 39]

Musicological discussions of music technology most commonly focus on musical form and output, often at the expense of social and material concerns, as detailed by Kerman [16]. This paper will instead reckon with the consequences of the musical industrial complex, analyzing and critiquing it through the lens of ICT and sustainability.

As has been discussed in a number of papers [29], identifying the precise environmental toll of any ICT sector is complex at best, and the music industry is no different. This is compounded by the difficulty of defining the "music industry" in general, as this term can be used to refer to performing music, recording studios, the manufacturing of compact discs, or the maintenance of streaming services, among many other things. These music industries are of course interrelated but will also have distinct goals and even competing priorities [45].

Therefore in order to precisely situate it within the ICT4S discourse, this paper will focus primarily on the design, manufacture, and distribution of electronic musical instruments (synthesizers, drum machines, effects, etc), and the use of computers to create and manipulate music. This forms an easily identified group with shared characteristics, and overlaps most clearly with the ICT sector. The (admittedly imperfect) phrase "digital music" will be used to refer to the act of using these electronic musical instruments and computer systems[1].

History and contemporary state of environmental sustainability in digital music

The music software industry grew alongside the computer software industry, starting in the 1980s with computer sequencing systems and eventually moving towards fully fledged "Digital Audio Workstations" (DAW) which at one point seemed to be replacing so-called 'outboard' hardware instruments. Indeed, for a time the music industry appeared to be a rare example of actual dematerialisation [22] in ICT, as sales of hardware synthesisers, sequencers, drum machines, and effects units plummeted in the early 2000s due to increased quality of computer-based DAW systems that could recreate and in some cases outperform the very systems that they were emulating, often at a fraction of the cost.

Whether this shift represents true dematerialisation is debatable - whilst during this time it definitely became possible (and in many ways fashionable) to compose, perform, and record with considerably less hardware [42], this was only possible because of the reliance on desktop computers, which themselves are of course complex and resource-intensive devices. However these computers would have performed a number of other functions as well, and it is hard to deny that the production of music circa 1998-2005 relied on significantly less material technology as a whole than previous eras [39], thanks to an intersection of efficient software design, computing power reaching a threshold of acceptability for audio, and a fashion for digital production encouraged by the creative possibilities of new tools.

In any case, whatever dematerialisation did happen was short lived - an array of factors over the past decade have led to a vigorous return to hardware sales [15], as well as a move towards software that requires more intensive local and "cloud" computing. The rebirth of the hardware market is particularly notable, and in many ways parallels the surprising return of vinyl as a music listening format [9]. Within ICT this pattern appears to be unique to the music industry (or perhaps sets an ominous precedent).

It is difficult to source reliable data on the music hardware market, but the issues in terms of sustainability are easy to extrapolate - virtually all synthesisers, drum machines, samplers, and effects units use similar integrated circuit systems as any other computer or smartphone, and nearly all instruments of these categories will make extensive use of plastic for their casing. In many cases the instruments require batteries, either to operate at all or as a backup mechanism for clocks or memory systems. The vast majority of electronic instruments are manufactured and assembled in China or elsewhere in Asia, and with the majority of the customers located in North America and Europe there are obvious ecological concerns in terms of shipping. Kyle Devine argues that due to a similar set of circumstances, music listening (vinyl/CDs/streaming) is primarily a petrochemical industry [9], and this framing can be equally applied here. Digital music making relies on oil at every step from production through to sale - it is heavily dependent on plastic, rare earth materials, oil, and (at best) problematic supply chain systems. This is virtually the same set of issues that can be seen throughout ICT.

It is worth noting that the concept of sustainability does not seem to have reached deeply into the music creation industry. Unlike most other ICT sectors, the companies involved do not feel the need to publish even symbolic statements about their efforts to move towards more sustainable production or

[1]This will include "analog synthesizers", which are of course electronic but use analog circuitry. This distinction is significant within the music technology sector but does not translate well to ICT vocabulary.

Authorized licensed use limited to: KTH Royal Institute of Technology. Downloaded on December 01,2022 at 11:47:27 UTC from IEEE Xplore. Restrictions apply.

designs. Of the major music technology manufacturers (Korg, Yamaha, Behringer/MusicTribe, inMusic) only Korg mentions sustainability or environmental awareness in any public way[2], and even then only as a part of a small experimental design studio in Berlin. Yamaha has published a detailed web page regarding their use of lumber in pianos [40], but do not mention any of the plastics or rare earth materials involved in their wide variety of electronic instruments. When MusicTribe (parent company of a number of brands including the extremely popular low-cost Behringer line of instruments and effects) opened a new 300,000 square meter "city" factory in Zhangshan City in 2018 with integrated supermarket, bakery, laundry service, hospital, and living quarters for 3,000 employees, the press release briefly mentions "green" without giving any indication of what this could mean [33]. To date no data has been released to support any so-called "green" activity.

A final consideration regarding the environmental impact of digital music is how it forms a part of (and relies upon) the wider ICT ecosystem of personal computers and software infrastructure. Put simply, modern digital music at virtually every level is dependant upon computers, whether to produce an album, document a modular synth improvisation, or compose a score. Digital music has, by taking advantage of personal computers, become far more accessible. Yet the ease of access - an ability to create perfect drum loops with a few taps, or record vocals with a quality that would have been unimaginable 20 years ago [3], is "encumbered" by the massive infrastructure required to enable this, as described by Thrift within the context of non-representational theory [43]. Software that runs flawlessly on a computer or tablet is dependent on decades of hardware and software design and manufacturing, as well as a complex set of global production and maintenance requirements and the environmental impact contained therein. The consequences are invisible and perhaps obscure, but as a society we have accepted this encumbering in exchange for smartphones and cloud storage and cheap microphones. The tradeoff for a frictionless digital music making experience is burning coal to power server farms [7], children mining for cobalt in the Congo [38], and the severed ear of a nameless convict in 1874 [11].

MAKING MUSIC WITHIN LIMITS

With this broad understanding of the digital music sector, we can attempt to construct a methodology for modern music making that accounts for the issues faced, beginning by applying the classic *in* vs. *by* (or *through*) concepts [34] to digital music and sustainability.

Sustainability *in* digital music can be understood as an approach that would encourage or legislate for greater environmental awareness and hopefully more efficient and ethical practices in the design and manufacture of musical instruments. Ultimately, within this context the responsibility lies primarily with the manufacturers who produce the products, which are overwhelmingly created using unsustainable methods as described above. The potential 'solutions' that can be applied within this framework would be for manufacturers to replace as much plastic as possible with more sustainable materials, implement effective recycling of rare earth metals and components, and integrate efficient manufacturing and shipping systems. This type of strategy has been widely adopted, at least symbolically, by virtually every major computer and smartphone manufacturer [13]. However, as mentioned above to date no major music manufacturer has publicly committed to these types of reforms, or even acknowledged the issue, with the possible exception of Swedish company Teenage Engineering who partnered with the Research Institutes of Sweden (RISE) to develop a speaker design using sustainable materials in 2018 [1] (this design does not appear to have been used for the final version that is currently for sale). We are left in the unenviable position of hoping that these companies will self-regulate, or consumer pressure will result in reforms. Both of these appear deeply unlikely - there is little evidence to show that corporate entities will ever willingly undergo sustainability reforms, and there is no appetite within the music technology market to hold companies responsible for their actions. For instance, Behringer appears to have suffered no ill effects after publicly ridiculing a journalist with anti-Semitic abuse [37], and Moog are still among the most highly respected manufacturers despite serious allegations of sexual harassment within the company [6].

On a smaller level, the increased accessibility of electronic parts and more rapid knowledge transfer due to communities on the internet has led to a striking rise in the number of small (and micro) musical instrument designers, most notably working on modular and "boutique" synthesisers, samplers, and effects. There is some potential to consider that these offer an opportunity to effect direct change with regards to sustainable practice in music technology, though the overall effect is likely to be small.

What, then, would sustainability *by* digital music look (or sound) like? It is most tempting to consider how music can be used to promote environmental awareness or sustainability issues more generally - music, like all art forms [17], can be used effectively as both a form of protest (and, indeed, to reinforce hegemony or status quo). Certainly a number of artists have composed music in response to the climate crisis [28], and this is likely to continue and may indeed form a crucial part of a fundamental cultural shift towards more sustainable practice.

However any music that is made using existing systems will be materially linked to the musical industrial complex, as well as continuing the problematic social lineage that resulted in their design. Meaningful and powerful digital music should and will certainly continue to be made, and will likely offer a great deal to the conversation, but systemic change must also be attempted - the "master's tools [or indeed, his master's voice] will never dismantle the master's house" [21]. This type of approach can acknowledge that digital music is a technology embedded with decades of social injustice, which has

[2]https://korg.berlin

Authorized licensed use limited to: KTH Royal Institute of Technology. Downloaded on December 01,2022 at 11:47:27 UTC from IEEE Xplore. Restrictions apply.

resulted in a domain firmly rooted within capitalist structures that exploit resources without regard for sustainability.

Therefore this paper proposes adopting a design framework that imposes limits on the design of music making technology. This is composed of the four main critiques raised in the paper - the general disregard within the industry for sustainability; the basic material issues with building musical devices; the problematic historical and cultural context of music technology; and the resulting potential musical output. Each of these form the basis of a design limitation.

Whilst it is unreasonable to consider that these limitations would be applied to every new musical instrument, using them as heuristic guidelines could function as a way of developing musical interfaces that reflect William Gaver's concept of "threshold devices" or "ludic action research" - something that functions partially as an intervention, encouraging or even forcing the user to reckon with a specific context and set of issues. As Michael and Gaver ask - "What ways are there that might promote this poetical engagement with the complexity, ambiguity, and so on of technonatures?" [25, p. 2]

The four heuristics described below therefore overtly and consciously limit the design of a digital musical interface [5]. They form a framework through which a new instrument can be created that explicitly engages with the issues of sustainability and associated social justice concerns laid out in this paper. The result would ideally be a tool that can be used to create meaningful and powerful sound, as well as function as a provocation towards a music making sector which is perhaps more accustomed to a cornucopian design approach [30].

Sustainable power

Whilst the definition of "sustainable" is of course up for a great deal of debate, it is nevertheless useful to begin by explicitly acknowledging that electronic instruments by definition must consume energy. The actual electricity used in a musical interface is relatively small, especially on systems without screens and amplifiers, being comparable to a single LED lightbulb (and presumably used far less). However there is considerable symbolic weight to the idea of overtly engaging with this energy consumption, and by extension the generally invisible use of resources in modern life [19]. Therefore a design that requires sustainable power as a fundamental principle can in some ways have an outsized impact, particularly within a sector that has, as shown above, so far neglected to take any stance (however symbolic) on sustainability issues.

The most obvious approach to this could result in an electronic instrument that runs only on solar, wind, or hydropower. This could potentially follow the calculations laid out by projects such as the Solar-Powered Website [2] in order to find the most efficient method for powering whilst taking into account a life cycle assessment of the materials required. However as musical interfaces will almost never require the consistent uptime demanded by web servers, the battery requirements can be considerably less or possible nil - several experimental musical devices have been created which bypass batteries altogether, relying on direct solar energy voltage, in the case of Electronic Etymology, for example [46].

Decoupling the visual and the auditory

Modern computer interfaces are primarily a screen-based visual medium. The limitations of this have been much analysed and discussed, particularly within the sonification and "auditory display" communities [12, 18, 10], but it is an undeniable reality that the vast majority of modern music is produced with the heavy use of a screen [24]. It is not unusual for recording studios to prioritise the computer screen in their setups, with a monitor in the centre of the speaker setup. Music software has developed a distinct visual language over the past two decades, with blocks of waveforms, timelines, grids, and other design elements being shared across virtually all software packages [3]. Creating music now requires the act of looking perhaps as much as the act of listening.

This is a compromise, stemming from the intertwining of music with computers [42]. The interface tradeoff for the enormously increased power and accessibility offered by running music software on personal computers is that the software must necessarily conform to the design conventions dictated by the computer. If, therefore, for sustainability reasons we want to decouple the creation of music from the personal computer ecosystem we can must remove the primary interface - the screen.

Of the heuristics described here this is certainly the least revolutionary, as a large number of digital instruments are already produced and sold with no screen (and the modular synthesis community often points to this as a fundamental appeal).

However a new digital musical instrument musical instrument that does not run on a visual-interface-based computer can prioritise listening, rather than looking, and would also both symbolically and tangibly distance itself from personal computers. It would require fewer parts to manufacture and be generally less encumbered by both the ICT ecosystem and the interaction design considerations inherent in any screen-focused interface.

Break the myth of neutrality

In many ways, modern electronic musical tools are designed to give the impression of neutrality [3]. They are supposedly blank slates, upon which a player is meant to put in their own skills, craft, and personal touch. As a whole they are nearly always sold on the premise that they are empty vessels that provide a user with the ability to be creative. The automatic drum pattern generator for Logic Pro, for example, is said to allow you to "choose from dozens of drummers across many musical genres, and direct their performances using simple controls. Edit the drum pattern in real time, including volume, complexity, and swing. Drummer can even follow along to a specified track and adjust its playing accordingly — much like a live drummer would."[3]. The tool is positioned as something

[3]https://www.apple.com/logic-pro/

Authorized licensed use limited to: KTH Royal Institute of Technology. Downloaded on December 01,2022 at 11:47:27 UTC from IEEE Xplore. Restrictions apply.

that is empowering the user, whilst providing a limited palette that will always achieve a clean and apparently correct sound.

This supposed neutrality can be seen as resulting directly from the problematic political ecology described in the first half of this paper. From a creative and cultural standpoint, the rise of accessible musical technology has demonstrably led to less variety in musical output, raising uncomfortable questions around cultural dominance and bias. Kyle McDonald's work has shown, for example, how tempos and keys have stabilized over recent years, with the majority of released music now using a 120bpm tempo and C Major key, which happens to be the default setting for most DAWs [4]. These design choices will all have been influenced by deeply entrenched biases on behalf of the designers who create these systems - companies which are overwhelmingly white, male, young, and located in wealthy urban centers in western countries. A default 120bpm setting reflects this, being almost exactly the average tempo setting for techno and house music that is popular amongst that socio-cultural group (and significantly higher than the average for genres such as Dub or Hip-Hop which skew towards minority groups).

Therefore, any instrument designed within this framework must strive to avoid "preset" decisions made on behalf of the user. The user must be empowered to decide for themselves - this will almost certainly generate less clean results at the outset, but this should be seen as a positive outcome.

Meangingful musical output

Musical instruments must make and/or manipulate sound. This sonic content, and the way that it is generated and controlled, is what forms the identity of an instrument. Thus, it is important that the sound itself also be subject to a level of constraint through a heuristic that contributes to the overall aim of creating music sustainably. There is quite a wide degree of scope in this heuristic, and a great deal of subjective choice. It is deliberately provocative - the aim is to force awareness and discussion of the content, positioning the issue as inescapable.

Within this context, the sonic content that is contained within the instrument must be derived from data or recordings that can be directly linked to core sustainability principles. For example, many instruments allow a player to add reverberation - in the case of this instrument, the reverberation included could be recorded at an abandoned coal mine, or in a petroleum tanker. The oscillations for a synthesiser, rather than being supposedly neutral sine waves, could be derived from fluctuations in ocean temperature, whilst sawtooth waves could be replaced with data representing rising per capita emissions. A drum machine could be designed entirely from sounds of now-extinct animals, or field recordings from scrap heaps of abandoned computers.

This method of foregrounding the source content has been used previously by artists such as Matthew Herbert, whose album "The End of Silence" was built entirely from the sound of a bomb being dropped in Libya, for example, resulting in an uncomfortable, nuanced, and challenging work that forces the listener to reckon with their own relationship to injustice [14]. In some ways this limitation represents most completely the potential of sustainability *through* music, as it interweaves the resulting musical content with the underlying issues.

[4]https://gist.github.com/kylemcdonald/1fab024c9878106b486deb1a26bc2079

Conclusion

Digital music is inextricably linked with ICT and a wider history of industrialisation. It can, and should, therefore be analyzed in terms of the technology it uses, as with any other sector within ICT. The theories and discussions within the ICT4S community thus offer a powerful framework that can be applied to music, and doing so raises uncomfortable questions around the political ecology and sustainability issues inherent in modern musical systems. This paper attempts to broadly acknowledge these issues as well as outline a methodology for designing new musical instruments, drawing on heuristic design techniques that embrace limits. This is intentionally provocative, particularly for a music sector which may be more willing to discuss musical output or new instruments than the problematic history of their tools.

Sound is energy - to make music is to literally generate and manipulate energy waves, however minuscule and short-lived. Music must be *made*, both performatively and physically, but energy can not be created, only converted. Therefore the act of making music is an act of industry, with all its associated environmental and social outcomes. In effect, this paper has proposed that a form of life cycle assessment can be performed on musical energy, and by extension it has questioned how the resulting environmental and social energy input can be justified by the musical output.

References

[1] *A loudspeaker for treehuggers with golden ears - the first speaker ever built using holocellulose*. RISE. URL: https://www.ri.se/en/press/a-loudspeaker-for-treehuggers-with-golden-ears-the-first-speaker-ever-built-using (visited on 01/28/2022).

[2] Roel Roscam Abbing. "'This is a solar-powered website, which means it sometimes goes offline': a design inquiry into degrowth and ICT". In: *LIMITS Workshop on Computing within Limits* (June 14, 2021). DOI: 10.21428/bf6fb269.e78d19f6. URL: https://limits.pubpub.org/pub/lecuxefc (visited on 11/16/2021).

[3] Adam Patrick Bell. "Can We Afford These Affordances? GarageBand and the Double-Edged Sword of the Digital Audio Workstation". In: (), p. 23.

[4] Georgina Born and Kyle Devine. "Music Technology, Gender, and Class: Digitization, Educational and Social Change in Britain". In: *Twentieth-Century Music* 12.2 (2015), pp. 135–172. DOI: 10.1017/S1478572215000018.

Authorized licensed use limited to: KTH Royal Institute of Technology. Downloaded on December 01,2022 at 11:47:27 UTC from IEEE Xplore. Restrictions apply.

[5] Jay Chen. “A strategy for limits-aware computing”. In: *Proceedings of the Second Workshop on Computing within Limits*. LIMITS ’16: Workshop on Computing within Limits. Irvine California: ACM, June 8, 2016, pp. 1–6. ISBN: 978-1-4503-4260-5. DOI: 10.1145/2926676.2926692. URL: https://dl.acm.org/doi/10.1145/2926676.2926692 (visited on 11/16/2021).

[6] Veronica Coit. *The misogyny at Moog*. URL: https://ashevilleblade.com/?p=4042 (visited on 01/20/2022).

[7] Gary Cook, Elizabeth Jardim, and C Craighill. “Clicking Clean Virginia The Dirty Energy Powering Data Center Alley”. In: *Greenpeace Available from: https://www. greenpeace. org/usa/reports/click-clean-virginia/[Accessed 21 Mar 2020]* (2019).

[8] Caroline Creton and Julien Bellanger. “Pas de pétrole, mais de l’Auto-Tune!” In: *Multitudes* 1 (2021), pp. 84–90.

[9] Kyle Devine. *Decomposed: the political ecology of music*. Cambridge, Massachusetts: The MIT Press, 2019. 316 pp. ISBN: 978-0-262-53778-0.

[10] A C Eldridge. “Issues in Auditory Display.” In: (), p. 27.

[11] Tom Everrett. “Writing sound with a human ear: reconstructing Bell and Blake’s 1874 ear phonautograph”. In: *Science Museum Group Journal* 12.12 (Nov. 3, 2021). ISSN: 2054-5770. DOI: 10.15180/191206. URL: http://journal.sciencemuseum.ac.uk/browse/issue-12/writing-sound-with-a-human-ear/ (visited on 01/05/2022).

[12] Sam Ferguson. “Sonifying Every Day: Activating Everyday Interactions for Ambient Sonification Systems”. en. In: *International Conference on Auditory Display (ICAD 2013)* (2013), pp. 77–84.

[13] *Guide to Greener Electronics 2017*. Greenpeace USA. URL: https://www.greenpeace.org/usa/reports/greener-electronics-2017/ (visited on 01/28/2022).

[14] Matthew Herbert. “Listen while you work: negotiating power and meaning in post-concrete music.” PhD Thesis. Canterbury Christ Church University, 2020. URL: https://repository.canterbury.ac.uk/item/8x378/listen-while-you-work-negotiating-power-and-meaning-in-post-concrete-music.

[15] Roeder Jonathan. *The Market for Musical Gear Is Amping Up Like Never Before*. Bloomberg.com. URL: https://www.bloomberg.com/news/articles/2021-03-17/musicians-are-hoarding-musical-gear-like-never-before (visited on 01/20/2022).

[16] Joseph Kerman. *Musicology*. London: Fontana Press : Collins, 1985. 255 pp. ISBN: 978-0-00-197170-7 978-0-00-686002-0.

[17] David Kirby. “The Future is Now: Diegetic Prototypes and the Role of Popular Films in Generating Real-world Technological Development”. In: *Social Studies of Science* 40.1 (Feb. 2010), pp. 41–70. ISSN: 0306-3127, 1460-3659. DOI: 10.1177/0306312709338325. URL: http://journals.sagepub.com/doi/10.1177/0306312709338325 (visited on 01/28/2022).

[18] Dan Lockton, Flora Bowden, and Claire Matthews. “Powerchord: Exploring Ambient Audio Feedback on Energy Use”. In: *Living Labs*. Ed. by David V. Keyson, Olivia Guerra-Santin, and Dan Lockton. Cham: Springer International Publishing, 2017, pp. 297–308. ISBN: 978-3-319-33526-1 978-3-319-33527-8. DOI: 10.1007/978-3-319-33527-8_23. URL: http://link.springer.com/10.1007/978-3-319-33527-8_23 (visited on 09/24/2021).

[19] Dan Lockton et al. “Powerchord: Towards Ambient Appliance-Level Electricity Use Feedback through Real-Time Sonification”. en. In: *Ubiquitous Computing and Ambient Intelligence. Personalisation and User Adapted Services*. Ed. by Ramón Hervás et al. Vol. 8867. Series Title: Lecture Notes in Computer Science. Cham: Springer International Publishing, 2014, pp. 48–51. ISBN: 978-3-319-13101-6 978-3-319-13102-3. DOI: 10.1007/978-3-319-13102-3_10. URL: http://link.springer.com/10.1007/978-3-319-13102-3_10 (visited on 10/11/2021).

[20] Arthur Loesser. *Men, women, and pianos: a social history*. New York: Dover Publications, 1990. 654 pp. ISBN: 978-0-486-26543-8.

[21] Audre Lorde. *Sister outsider: Essays and speeches*. Crossing Press, 2012.

[22] Sylvia Lorek. “Dematerialization”. In: *Degrowth*. Routledge, 2014, pp. 111–113.

[23] Ming-Yuen S. Ma. *There is no soundtrack: rethinking art, media, and the audio-visual contract*. Rethinking art’s histories. Manchester: Manchester University Press, 2020. 242 pp. ISBN: 978-1-5261-4212-2.

[24] Ian Macchiusi. ““Knowing is Seeing”: The Digital Audio Workstation and the Visualization of Sound”. PhD thesis. Toronto: York University, 2018.

[25] Mike Michael and William Gaver. “Home Beyond Home: Dwelling With Threshold Devices”. In: *Space and Culture* 12.3 (Aug. 2009), pp. 359–370. ISSN: 1206-3312, 1552-8308. DOI: 10.1177/1206331209337076. URL: http://journals.sagepub.com/doi/10.1177/1206331209337076 (visited on 11/10/2021).

[26] Greg Milner. *Perfecting sound forever: an aural history of recorded music*. 1st paperback ed. New York: Faber and Faber, 2010. 416 pp. ISBN: 978-0-86547-938-8.

[27] Lisa Nakamura. “Indigenous Circuits: Navajo Women and the Racialization of Early Electronic Manufacture”. In: *American Quarterly* 66.4 (2014), pp. 919–941. ISSN: 1080-6490. DOI: 10.1353/aq.2014.0070. URL: http://muse.jhu.edu/content/crossref/journals/american_quarterly/v066/66.4.nakamura.html (visited on 01/17/2022).

[28] Alexis Petridis. “The greatest songs about the climate crisis – ranked!” In: *The Guardian* (Oct. 28, 2021). ISSN: 0261-3077. URL: https://www.theguardian.com/music/2021/oct/28/the-greatest-songs-about-the-climate-crisis-ranked (visited on 01/28/2022).

Authorized licensed use limited to: KTH Royal Institute of Technology. Downloaded on December 01,2022 at 11:47:27 UTC from IEEE Xplore. Restrictions apply.

[29] Johanna Pohl, Lorenz M. Hilty, and Matthias Finkbeiner. “How LCA contributes to the environmental assessment of higher order effects of ICT application: A review of different approaches”. In: *Journal of Cleaner Production* 219 (May 2019), pp. 698–712. ISSN: 09596526. DOI: 10.1016/j.jclepro.2019.02.018. URL: https://linkinghub.elsevier.com/retrieve/pii/S095965261930397X (visited on 10/05/2021).

[30] Chris Preist, Daniel Schien, and Eli Blevis. “Understanding and Mitigating the Effects of Device and Cloud Service Design Decisions on the Environmental Footprint of Digital Infrastructure”. en. In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. San Jose California USA: ACM, May 2016, pp. 1324–1337. ISBN: 978-1-4503-3362-7. DOI: 10.1145/2858036.2858378. URL: https://dl.acm.org/doi/10.1145/2858036.2858378 (visited on 10/05/2021).

[31] Catherine Provenzano. *Auto-Tune, Labor, and the Pop-Music Voice*. Vol. 1. Oxford University Press, Oct. 18, 2018. DOI: 10.1093/oso/9780199985227.003.0008. URL: https://oxford.universitypressscholarship.com/view/10.1093/oso/9780199985227.001.0001/oso-9780199985227-chapter-8 (visited on 01/28/2022).

[32] Catherine Provenzano. “Making Voices: The Gendering of Pitch Correction and The Auto-Tune Effect in Contemporary Pop Music”. In: *Journal of Popular Music Studies* 31.2 (2019), pp. 63–84.

[33] Ben Rogerson published. *Behringer’s mindblowing new factory is actually a city*. MusicRadar. Mar. 27, 2018. URL: https://www.musicradar.com/news/behringers-mindblowing-new-factory-is-also-a-city (visited on 01/20/2022).

[34] Christian Remy and Elaine M Huang. “Addressing the Obsolescence of End-User Devices: Approaches from the Field of Sustainable HCI”. In: (), p. 11.

[35] Tara Rodgers. *Pink noises: women on electronic music and sound*. Durham [NC]: Duke University Press, 2010. 322 pp. ISBN: 978-0-8223-4661-6 978-0-8223-4673-9.

[36] Craig H. Roell. *The piano in America, 1890-1940*. 2. print. Chapel Hill: Univ. of North Carolina Pr, 1991. 396 pp. ISBN: 978-0-8078-4322-2 978-0-8078-1802-2.

[37] Janus Rose. *Behringer Created This Fake Product To Attack A Journalist*. URL: https://www.vice.com/en/article/wxe7qx/a-major-synth-company-created-this-fake-product-to-attack-a-journalist (visited on 01/20/2022).

[38] Benjamin K. Sovacool. “The precarious political economy of cobalt: Balancing prosperity, poverty, and brutality in artisanal and industrial mining in the Democratic Republic of the Congo”. en. In: *The Extractive Industries and Society* 6.3 (July 2019), pp. 915–939. ISSN: 2214790X. DOI: 10.1016/j.exis.2019.05.018. URL: https://linkinghub.elsevier.com/retrieve/pii/S2214790X1930084X (visited on 04/04/2022).

[39] Robert Strachan. *Sonic technologies: Popular music, digital culture and the creative process*. Bloomsbury Publishing USA, 2017.

[40] *Sustainable Resource Use - Sustainability - Yamaha Corporation*. URL: https://www.yamaha.com/en/csr/environment/sustainable_resource_use/ (visited on 01/28/2022).

[41] Paul Théberge. *Any sound you can imagine: making music/consuming technology*. Music/culture. Hanover, NH: Wesleyan University Press : University Press of New England, 1997. 293 pp. ISBN: 9780819553072 9780819563095.

[42] Paul Théberge. “Digitalization”. In: *The Routledge Reader on the Sociology of Music*. Routledge, 2015, pp. 345–354.

[43] N. J. Thrift. *Non-representational theory: space, politics, affect*. International Library of Sociology. Milton Park, Abingdon, Oxon ; New York, NY: Routledge, 2008. 325 pp. ISBN: 9780415393201 9780415393218.

[44] Lucie Vágnerová. “‘Nimble Fingers’ in Electronic Music: Rethinking sound through neo-colonial labour”. In: *Organised Sound* 22.2 (Aug. 2017), pp. 250–258. ISSN: 1355-7718, 1469-8153. DOI: 10.1017/S1355771817000152. URL: https://www.cambridge.org/core/product/identifier/S1355771817000152/type/journal_article (visited on 01/07/2022).

[45] John Williamson and Martin Cloonan. “Rethinking the music industry”. en. In: *Popular Music* 26.2 (May 2007), pp. 305–322. ISSN: 0261-1430, 1474-0095. DOI: 10.1017/S0261143007001262. URL: https://www.cambridge.org/core/product/identifier/S0261143007001262/type/journal_article (visited on 03/28/2022).

[46] Lucas Yasunaga. *Electronic Entomology*. URL: http://electronicentomology.com/ (visited on 01/20/2022).

Authorized licensed use limited to: KTH Royal Institute of Technology. Downloaded on December 01,2022 at 11:47:27 UTC from IEEE Xplore. Restrictions apply.

Paper B

The Memory Cloud: Personal media libraries as affordance and constraint

Yann Seznec
yannse@kth.se
KTH Royal Institute of Technology
Stockholm, Sweden



Figure 1: Playing a Memory

Abstract

The Memory Cloud is a musical instrument that uses a player's own library of personal recordings as sonic material. This paper presents the design of the instrument, situating it within sustainability HCI studies and constraints-based design, before describing the instrument being used by two musicians in a professional context. Over 2000 sounds from the musician's personal cloud library, dating back over 10 years, were placed in the instrument as the only sonic material available for exploring. I argue that a radically small scale and personal approach could be one strategy for addressing the issues of longevity in NIME, and I suggest that using personal media libraries presents a potential affordance and constraint for musical instrument design.

Keywords

instrument, nime, cloud, sustainability, performance, design

1 Introduction

My cloud library is full. 34,638 photogaphs and 1,308 videos are taking up my 2TB allocation. As a designer working with sound it is not difficult to see the videos in particular as a form of sonic content - it is raw, recorded mainly on my phone and full of wind noise, clipping, muffled by a poorly positioned hands. And yet it is unmistakeably mine. No one else even has access to this library, built up almost accidentally over many years, representing virtually all of the major events in my life. Chances are, you have one of these libraries too (though you may manage yours better than mine). What would it mean to *use* these sounds? What kind of musical relationship can we build through encounters with this media?

This paper presents an instrument that takes a user's entire cloud library and repurposes it for musical use: The Memory Cloud. Pressing a button calls up a memory from the past and allows the player to break it apart or reconfigure it in any number of ways. In doing so the player builds a relationship with their own sonic history, unfolding new affordances for musical expression.

2 Background

2.1 Cornucopian design

Our overly saturated cloud media libraries, and indeed the entire existence of cloud storage, can be situated within what has been called the *cornucopian paradigm* of computing [16]. This is an approach to designing computer technologies with a fundamental assumption that all options and data should be made instantly available to the user at all times. Cloud libraries provide the potential for users to store all of their media, and the priority from a design standpoint is for the user to be able to seamlessly access

This work is licensed under a Creative Commons Attribution 4.0 International License.
NIME '25, June 24–27, 2025, Canberra, Australia
© 2025 Copyright held by the owner/author(s).

it anytime and anywhere. This is clearly unsustainable from a number of perspectives [9], with a growing body of sustainable HCI (SHCI) research calling for post-growth [20], degrowth [3], computing within limits [2, 15, 22] and deeper inner transitions [4] when considering the design of computing technologies.

2.2 Affordances and constraints

The work described here engages both with constraints-based design from a sustainability standpoint as well as in terms of the creative opportunities it presents. As discussed by Magnusson [11], among others, constraints and affordances are "two sides of the same coin", at least for musical instruments (pg 71). Design constraints, such as the number of strings on a guitar or the number of polyphonic voices on a synthesizer, are things that a musician discovers and engages with through exploration and play, potentially creating a relationship. That embodied, situated, and evolving dynamic can be seen as the affordances of an instrument. As Holzer argues through Barad's notion of agential cuts [8], the affordances of an instrument are thus comprised of the specific set of possibilities that are built between an instrument and player.

2.3 Longevity and NIME

The instrument described in this paper is by definition very personal, and thus it is unlikely to be developed in any large-scale way. There will be no "future developments" section in this paper where I suggest that a commercial version could be released; indeed, the paper will argue rather that the radically small-scale nature of this work is precisely what generates the musical affordances in the instrument. Nevertheless there are aspects of this work that link it to some of the recent discussions around longevity in NIME [5, 12]. In particular this work engages with concepts of *ownership*, which has been suggested as a potential mechanism for encouraging long term use of an instrument [14, 23]. Additionally, by foregrounding the conceptual positioning of the instrument within sustainable HCI we align with recent calls for more "outward-looking" work [13] to encourage wider discussions of the socio-political ramifications of designing musical technologies.

2.4 Personal sonic content

The specific constraint that is being using in the Memory Cloud is the sound library. Once the instrument has been set up for the player, there is no way of adding new sounds or recording samples into the machine[1]. Instead, the only sounds available to be loaded into the granular engine come from the player's own cloud library. Admittedly this can be quite large (and in that sense pushes the boundaries of the term "constraint"), but nevertheless virtually guarantees that the sonic output will be associated with historical and emotional connection to the sonic source. Using what are essentially field recordings as musical content creates obvious parallels with the musique concrète practices of Schaefer [18] and Pade [1], among many others. Through the use of personal sounds in particular, the work described here aligns strongly with Matthew Herbert's concept of the "audible subject" [7] in terms of how the context of the source sound is crucial to the relationships that are built between listener, performer, and instrument (moving away somewhat from a strict acousmatic approach).

[1]As discussed in the "Design" section, this is not a technical so much as a design limitation, as the engine is certainly capable of adding or recording new sounds, and anyone with a screwdriver and Raspberry Pi experience could load new sounds into the instrument. However these features are hidden from the user and thus for the purposes of this work can be considered impossible.

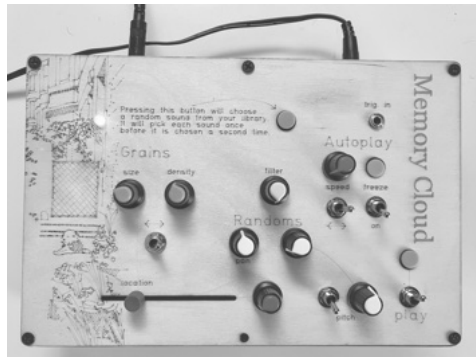


Figure 2: The Memory Cloud

3 The Design

The Memory Cloud is an instrument based on a granular synthesis engine made in Pure Data. This engine has been developed over several years and used for a number of different projects[2]. The precise nature of the granular code is not important for the purposes of this paper, and the specifics of granular synthesis are sufficiently well understood by the NIME community and do not need to be outlined here. For details on how the instrument itself functions and sounds please see the instructions document that is included in the code repository.

The instrument itself is comprised of the granular engine running on a Raspberry Pi, with a Teensy microcontroller taking analog control input from the interface elements and sending them over Serial to the Pure Data patch using the *comport* object[3]. The interface is made up of 8 knobs, one slider, four switches, and three buttons. The majority of these controls are mapped to granular manipulation parameters such as grain size, grain location, grain direction, pitch and location randomness, etc.

The sounds contained in the instrument are perhaps the most important, or at least the most distinctive aspect of the Memory Cloud. Although technically any folder of sounds could be used, the instrument is intended specifically to contain the sounds from all of the videos of the user's personal media library. In initial tests I used approximately 1,200 audio files extracted from videos in my own iCloud photo library, dating back 12 years. The version that was made for a musician (which will form the bulk of the analysis of this paper) included 2,632 sound files from roughly 10 years worth of cloud storage.

The sound library works as follows: a folder of sounds is exported from all of the videos in a user's cloud library and loaded into the local memory of the Raspberry Pi. When switched on for the first time the Pure Data patch in the device will "shuffle" these sound files and select one to be loaded into the granular synthesis engine for manipulation. A dedicated button on the instrument allows the user to change this sound file - when pressed, the next sound in the shuffled order will be selected and loaded. There

[2]The code is fully open source and available for download and use at https://github.com/yannseznec/ys.granular, along with a detailed explanation video.
[3]https://git.iem.at/pd/comport

is no way of going back to the previous sound: Each time the players loads one sound they are also leaving the current sound behind. In order to play with a sound file a second time, the user must first load every other sound in the folder.

Preparing, exporting, and transferring these files is something of an undertaking. The most efficient method I devised involved downloading all of the videos from an iCloud library to a local drive, using a command line batch audio processing system (ffmpeg) to extract the audio, and finally batch converting it with SoX to the suitable audio format for the granular engine. The resulting sound files could then be loaded over SSH or USB to the Raspberry Pi's local storage.

This process ensures that each Memory Cloud is absolutely unique. While the engine which creates the sonic output, and the interface for controlling it, can be seen as a relatively straightforward and even fairly standard granular system, the content that forms the basis of this sonic output is tied inextricably to the owner of the device.

The initial design of the Memory Cloud took the form of a first-person Research through Design project, in which I developed and created the system for my own use using my own cloud library. The design stems from thinking around sustainability issues with regard to cloud storage and cornucopian design paradigms, as described previously, and the resulting instrument acts as something to think with [17], with knowledge emerging [6] from the object itself and its use. Working together with another musician extended this exploration one step further, giving me a glimpse of another perspective on the instrument and allowing me to play the role of mediator, as the process for setting up the instrument required a certain amount of specialist guidance. In this way the device offered a way to expand the inquiry of the design space beyond my own designing. It and the unfolding use became a carrier of the research questioning in which I was engaged around radically small-scale values embedded into the instrument.

4 Using the Instrument

The Memory Cloud was used for several months by musicians Sam Beste and Yoshika Colwell, using Sam's personal sound library. During the initial familiarization period the instrument was in Sam's studio and was used by visitors:

> *Sam:* So it's just been sort of set up in my studio, and just as people have been passing through they've been using it, which has been really nice, and that's the way that I've most enjoyed interacting with it, in a way, is with me playing an instrument and somebody else actually playing [The Memory Cloud] so that there's an interaction.

After some time it became clear that the instrument could work well for a concert at a large gallery in London, particularly as Yoshika built a relationship with it:

> *Sam:* I made this collaborative EP with this singer, Yoshika, which kind of came together through us improvising together in a way, and then building songs out of those improvisations. And so when we were offered this opportunity to do this thing at the Gallery, you know, we only have four songs, and they wanted us to play 40 minutes. And at first I wasn't really sure how we would to do that. But then, actually, when [Yoshika] came into the studio, the idea was like, well, why don't we play the four songs but in between the songs, we sort of just improvise? And having the Memory Cloud, it was actually quite a transformative experience for her, because she's often been in environments where she isn't given control of an instrument, or she's...been stuck in certain kind of ways of perceiving herself and others perceiving her. And actually having this box, she was completely amazing at performing with it as well...
>
> *Yoshika:* I played it a few times leading up to that gig and then I played it at the gig and it was amazing. I absolutely love it, it's one of my favorite things I've ever used. Because I'm not someone who plays with samplers [or other music technology]. I often feel slightly alienated by them...With the Memory Cloud I immediately took to it and I found it very very accessible, which was so exciting for me.

The inherently personal nature of the device was one of the key elements that activated the instrument for Yoshika in particular:

> *Yoshika:* The beauty of it is the personalized element, so I'm sitting there playing with Sam and I can hear "oh that's your kids" and it's just so beautiful and moving. All of those sounds immediately have a weight to them, even though I'm not in his family. There's something very humanizing about that, I really like that.
>
> *Sam:* I think part of her joy in using it was knowing that it's completely singular and that it's completely related to my own life, and that the things that she's relating to and is playing are so personal and it's like this sort of process of discovery...That's what makes it magical.

The impermanence of the sound selection seems to have been an important part of what made the instrument work well in this context, with Sam eventually beginning to see future potential for building work around it:

> *Sam:* But she also just really loved the concept of it, and loved how the kind of impermanence of it, you know that you can't get attached to anything, because you'll never get there again. And that really worked for the nature of performance, and I think that kind of impermanent nature of something in your setup just invites you to kind of be present in a way, and not be sort of tied to things that are going to occur every time and you know are going to be there. It kind of protects you from stagnation in some way, which I really love. And I'm not entirely sure how I'm going to live without it, to be honest, because I sort of feel like I keep imagining ways that it's going to be involved in different aspects [of my work].

As the designer of the object I was surprised to hear from Sam that he had so readily given this instrument, the sonic contents of which were so personalized and particular to his own life, over to others to play. However he explained that this was a result of his desire to respond musically (mainly through the piano) to processed sound recordings:

> *Sam:* It has fit so well into a world that I was already exploring, but didn't have a means of doing it in a kind of this generative way, live. And so it

> was always going to be somebody else playing it in a way, because I'm always interested in the interaction between the piano and sound...So actually, having somebody else control it sort of just makes a lot of sense.

5 Reflections

The use of the Memory Cloud suggests three specific themes that can be discussed here: the specific potential of radical small scale design, how the design of the instrument can be seen as a constraint that encourages creative output, and how the use personal sounds acts as a musical affordance.

5.1 Radically microscale design, politics and longevity

The Memory Cloud is an instrument that demands care. From the outset, my own role of designer involves a significant amount of practical work and guidance, with the eventual goal of collecting a user's own personal sounds and creating a customized version of the instrument. On the one hand this is an illustration of the inefficiency and biases of the technological systems that we are entangled with. Each step of the process is frictional and demonstrates how the systems on which we rely for managing our personal media are simply not designed for these purposes. Clicking and dragging 2000 video files from a photo management application immediately creates graphical and functional slowdowns of the operating system even on a relatively new computer. Exporting that many files is simply not expected behaviour - the iCloud browser system, for example, only allows a user to download 1000 videos at a time (and does not function reliably).

Indeed, the process of creating Sam's Memory Cloud from his cloud library was challenging. I needed to guide him through the process of exporting his entire video library from iCloud, installing and using command line batch processing tools for the first time, and sending me the resulting sounds to load onto the instrument. Working through these steps with Sam, whilst living in two different countries, took several days (and many text messages), and gave him some understanding of both the scale and power of his own media:

> *text message from Sam:* in trying to find a way through i accidentally opened all 4.5k videos[4] with quicktime omg! it almost collapsed my system! haha! But was really emotional to see all the videos popping up one by one...quite amazing that we have this archive of our lives

This process was essentially reactivating a relationship between Sam and his own media library, a relationship in which I as the designer was acting as something of a mediator or guide, carefully helping him collect and prepare his personal memories for use in his own instrument.

While the technical system that underpins the Memory Cloud is a relatively straightforward granular synthesis engine, the use of the personal media library as the only sonic content available in the instrument is a constraint that serves to strengthen its emotional power. The Memory Cloud takes a radically anti-innovation approach that explicitly uses an existing technical engine and instead is about building a connection between a specific user and their own personal memories. Put simply, each Memory Cloud is unabashedly a singular instrument for a single person. My own Memory Cloud could not possibly have the same power for someone else - and that is precisely the point.

Thus the core elements that come together to form what I call "radically microscale design" are:

- A user willing to engage with their own media library and confront the biases embedded in our technologies
- A thoughtful and caring relationship between designer and user
- An overt avoidance of technological innovation

The way this concept dovetails with sustainability and computing issues as detailed in the background also points towards a form of media reuse, which could form the basis of more work in the NIME community. Where systems such as sample libraries, physical modeling, and latent spaces all push for creating *new* sounds, what about the thousands of sounds that we are all carrying around in our pockets? Perhaps these sounds are worth considering as material for creative expression - particularly as many of them will contain far more emotional weight than anything generated by machine learning or other methods. Therein lies the political stance in the work, which is both embedded into the design of the instrument and is revealed through setup and use.

The way this was manifested in the Memory Cloud, through the use of the vast personal library spanning a number of years, prompted reflections from Yoshika around the nature of these sounds and how they differ from musical material that is often used in musical composition and performance:

> *Yoshika:* It's such a scrapbook of somebody's life, so intimate. That ambient noise, it's not constructed, which does resonate with the philosophy of the [instrument]. It's like...all of the background stuff is all of the important stuff, really. The stuff of everyday life, the noise, the laughing, the rustling, it's not performative in any way. And I think that is so mirrored in the instrument, because it feels like it's a tool to stimulate your thoughts about those sounds.

It has been suggested that a preoccupation with technical innovation and an avoidance of political stances has potentially contributed to the problems with NIME longevity [13]. I would like to position the Memory Cloud within these discussions that have taken place within the NIME community, suggesting that the creative process for this instrument stemmed from a specific political stance based on anti-solutionism, sustainable HCI, care, and radically small scale design. Whether or not this results in an instrument with longevity remains to be seen, but I present this work as a way of attempting to actively engage with the discussions.

5.2 Constraints of personal sonic material

Having briefly touched upon the results of a constraints-led design process, it is worth considering exactly how these are applied in the design of the instrument and how the resulting interaction manifests.

It is immediately noticeable that the instrument itself has no screen. Whilst some discussions around musical instrument design has suggested that moving away from screens might be a small step towards sustainability [19], in practice there are also cost and technical reasons for this approach. The instrument runs on a Raspberry Pi, which could certainly have a screen attached

[4] the final number was smaller due to file compatibility, doubles, and naming issues, but was deemed sufficient in any case

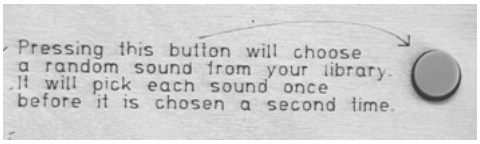


Figure 3: Choose a new sound, leave the old sound behind

through HDMI, but integrating this into the user interface would likely have created more problems than it solved. Nevertheless an instrument that requires a sound file to be loaded in order to start making sound (and contains potentially thousands of files to load) will certainly require some method of choosing a sound file to load.

This constraint is what led to the button for loading the sound into the granular engine. This button is prominently featured in the middle of the device, accompanied by an explanation (Figure 3). As mentioned earlier, the entire sound library is shuffled, so the order of loading the sounds is predetermined but unknown to the player.[5] The next sound file may be from a video recorded on the same day as the previous, or could be years apart. It could be a short clip of a rehearsal or it could be the sound of a dog running on a beach. The only thing that is certain is that pressing the button will leave the current sound behind, and you will not hear it again until you have heard every other sound in the library - with over 2000 sounds in Sam's library it is unlikely that he will be hearing any given sound again any time soon.

There is a peculiar balance to this, where there is a combination of unknown (which sound will come next) and known (all of the sounds are from a single and personal source). It is likely that the player will recognize any given sound, though the order in which they are presented will be surprising and potentially jarring. This mirrors the dual nature of the design constraint - on the one hand this is an instrument filled with a staggering amount of sonic content. In Sam's case the total size of the library came to 46 gigabytes, or roughly 80 hours of audio. This could, admittedly, be seen as adopting the very cornucopian design paradigm that I was trying to avoid, in the sense of providing the user not just with quite a lot of sounds, but with *every single sound they have ever recorded.* However, by only providing a single button as the interface for selecting a sound (and by overtly avoiding a screen to enable selection or give information on upcoming files), this glut of sound is transformed instead into a constraint where each button press must be carefully considered. Is it time for a new sound? Is it worth sticking with this one a bit longer? When will I hear it again?

For Yoshika in particular this constraint was a liberating experience from her very first time playing the Memory Cloud:

> *Yoshika:* I was playing one sample and [Sam] was like "the fun thing about it is you go forward and you can never get back [to the previous sound]"...So you can never get attached to it, which I absolutely love. And immediately you make your peace with that, just psychologicaly for me, you can't get stuck. I think that's because you immediately accept that you can't be attached to a sound...I think in terms of not being attached, and how generative that is creatively, I think that's amazing and very freeing because you feel like you are improvising with it, it feels like a two way thing with the machine...you're not attached to the outcome being exactly the same. You want to capture the feeling and the essence rather than keep something in a very constricted space.

5.3 Personal media as affordance

Sam's relation with his Memory Cloud was intensely personal, describing it as "singular and...completely related to my own life". This is perhaps not surprising as it contained only his own sounds from his own media library, though it also seems likely that the process of exporting the videos and preparing the sounds contributed also to his connection to the finished instrument. Playing the instrument did not exist in a vacuum - Sam began building a relationship with the library of sounds well before he received the Memory Box in the post. By the time he received it he was already aware of the personal nature of the sounds that it would create, as well as having a handle on the sheer quantity of files and the associated technical challenges of using them.

These interwoven relationships between myself (the designer), Sam and Yoshika (the performers), the instrument, and the personal media content can be seen as reflecting the complex and entangled nature of musicking with technology [10, 21]. The Memory Cloud creates affordances not primarily through interface or sound design, but instead through the use of personal media. It leverages sonic content which by its very existence will already be filled with meaning, particularly for the one individual who created that content. The affordances that make the Memory Cloud function are built upon and with the user's past - potentially years of memories that are re-configured and presented back to the creator, generating encounters with their own history and giving the opportunity to dissect a singular moment from their own experience.

6 Conclusion

The Memory Cloud is an instrument that uses personal media to generate new affordances for musical interaction. It leverages our vast libraries of videos towards a radically small scale and personal experience, embracing constraints as a design strategy and a political stance.

7 Ethical Standards

This research took place at KTH Royal Institute of Technology in the context of my doctoral studies. The research is exempt from formal ethical approval under the interpretation of Swedish law[6]. Data collection does not directly relate to sensitive personal data. The musicians involved in testing the instrument (Sam Beste and Yoshika Colwell) were informed from the beginning of the research context of the work and Sam fully consented to using his own sound recordings in the instrument, aware that this formed an integral part of the design of the instrument and the resulting research. Both Sam and Yoshika consented to our conversations being recorded and transcribed for the paper. They were not compensated financially for their participation. The research also necessarily involved the use of my own personal media library in the development of the instrument. While my supervisory team

[5]When the device is switched on it will automatically reload the same sound that was in the buffer when the instrument was switched off. The shuffled order is retained.

[6]https://www.riksdagen.se/sv/dokument-och-lagar/dokument/svensk-forfattningssamling/lag-2003460-om-etikprovning-av-forskning-som_sfs-2003-460/

was fully supportive of this work, the use of my own sounds was entirely of my own volition and was not the result of any coercion from supervisors or anyone else in a position of power.

Acknowledgments

This work is supported by the Energimyndigheten (Swedish Energy Agency) project *Interaktionsdesign för ett Soldrivet Internet* (Interaction Design for a Solar Powered Internet), 51885-1.

I would like to thank my supervisor Rob Comber for advice and support throughout this project.

Special thanks also to Sam and Yoshika for their open minded but critical embrace of the concepts, constraints, and outcomes of using the instrument.

References

[1] Andrea Bak. 2009. *Else Marie Pade: en biografi.* Gyldendal, København. OCLC: 667231490.

[2] Jay Chen. 2016. A strategy for limits-aware computing. In *Proceedings of the Second Workshop on Computing within Limits*. ACM, Irvine California, 1–6. https://doi.org/10.1145/2926676.2926692

[3] Giacomo D'Alisa, Federico Demaria, and Giorgos Kallis (Eds.). 2015. *Degrowth: a vocabulary for a new era.* Routledge, Taylor & Francis Group, New York ; London.

[4] Elina Eriksson and Daniel Pargman. 2023. Exploring Inner Transition: Expanding Computing forSustainability. In *Ninth Computing within Limits 2023*. LIMITS, Virtual. https://doi.org/10.21428/bf6fb269.22dea4ad

[5] Alessandro Fiordelmondo, Giada Zuccolo, Sergio Canazza, and Raul Masu. 2024. Longevity in NIME research: a case study using time-based media art preservation models. (Oct. 2024). https://doi.org/10.5281/ZENODO.13904858 Publisher: Zenodo.

[6] William Gaver, Peter Gall Krogh, Andy Boucher, and David Chatting. 2022. Emergence as a Feature of Practice-based Design Research. In *Designing Interactive Systems Conference*. ACM, Virtual Event Australia, 517–526. https://doi.org/10.1145/3532106.3533524

[7] Matthew Herbert. 2020. *Listen while you work: negotiating power and meaning in post-concrete music.* PhD Thesis. Canterbury Christ Church University. https://repository.canterbury.ac.uk/item/8x378/listen-while-you-work-negotiating-power-and-meaning-in-post-concrete-music

[8] Derek Holzer. 2024. A Piano for Visuals: Affordances of Scan Processing Instruments. *Leonardo* 57, 6 (Dec. 2024), 613–621. https://doi.org/10.1162/leon_a_02584

[9] Esther Jang, Matthew Johnson, Edward Burnell, and Kurtis Heimerl. 2017. Unplanned Obsolescence: Hardware and Software After Collapse. In *Proceedings of the 2017 Workshop on Computing Within Limits*. ACM, Santa Barbara California USA, 93–101. https://doi.org/10.1145/3080556.3080566

[10] Alexander Refsum Jensenius. 2022. *Sound Actions: Conceptualizing Musical Instruments.* The MIT Press, Cambridge.

[11] Thor Magnusson. 2010. Designing Constraints: Composing and Performing with Digital Musical Systems. *Computer Music Journal* 34, 4 (Dec. 2010), 62–73. https://doi.org/10.1162/COMJ_a_00026

[12] Raul Masu, Fabio Morreale, and Alexander Refsum Jensenius. 2023. The O in NIME: Reflecting on the Importance of Reusing and Repurposing Old Musical Instruments. (May 2023). https://doi.org/10.5281/ZENODO.11189120 Publisher: Zenodo.

[13] Fabio Morreale, S. M. Astrid Bin, Andrew McPherson, Paul Stapleton, and Marcelo Wanderley. 2020. A NIME Of The Times: Developing an Outward-Looking Political Agenda For This Community. (June 2020). https://doi.org/10.5281/ZENODO.4813294 Publisher: Zenodo.

[14] Fabio Morreale and Andrew McPherson. 2017. Design For Longevity: Ongoing Use Of Instruments From Nime 2010-14. (June 2017). https://doi.org/10.5281/ZENODO.1176218 Publisher: Zenodo.

[15] 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

[16] 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 *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 1324–1337. https://doi.org/10.1145/2858036.2858378

[17] Johan Redström. 2017. *Making Design Theory.* The MIT Press. https://doi.org/10.7551/mitpress/11160.001.0001

[18] Pierre Schaeffer. 1998. *À la recherche d'une musique concrète* ([nachdr. der ausg.] paris, seuil, 1952 ed.). Seuil, Paris.

[19] Yann Seznec. 2022. Music Within Limits. In *2022 International Conference on ICT for Sustainability (ICT4S)*. IEEE, Plovdiv, Bulgaria, 172–178. https://doi.org/10.1109/ICT4S55073.2022.00028

[20] Vishal Sharma, Neha Kumar, and Bonnie Nardi. 2024. Post-growth Human–Computer Interaction. *ACM Transactions on Computer-Human Interaction* 31, 1 (Feb. 2024), 1–37. https://doi.org/10.1145/3624981

[21] Christopher Small. 1999. Musicking — the meanings of performing and listening. A lecture. *Music Education Research* 1, 1 (March 1999), 9–22. https://doi.org/10.1080/1461380990010102

[22] Martin Stojanov, Daniel Pargman, Mike Hazas, Rob Comber, and Jorge Luis Zapico. 2023. How do we arrive at constraints? Articulating limits for computing. In *Ninth Computing within Limits 2023*. LIMITS, Virtual. https://doi.org/10.21428/bf6fb269.a317d18f

[23] Isaac Wallis, Todd Ingalls, Ellen Campana, and Catherine Vuong. 2013. Amateur Musicians, Long-Term Engagement, and HCI. In *Music and Human-Computer Interaction*, Simon Holland, Katie Wilkie, Paul Mulholland, and Allan Seago (Eds.). Springer London, London, 49–66. https://doi.org/10.1007/978-1-4471-2990-5_3 Series Title: Springer Series on Cultural Computing.

Paper C

Three Ways to Destroy a Sound

Yann Seznec
yannse@kth.se
KTH Royal Institute of Technology
Stockholm, Sweden



ABSTRACT

This paper presents a set of three different designed objects which each take a different approach to the same concept - the active destruction of sound recordings. I situate this work within sustainability studies, political ecologies of technology, and constraints-based design. Following this I describe the design and development of the devices themselves, introduce the final built objects, and attempt to analyze what they might represent. The paper takes the form of an annotated portfolio, in which the objects are shown, analyzed, and critically evaluated.

CCS CONCEPTS

• **Human-centered computing → Sound-based input / output**; HCI theory, concepts and models; *Interaction design theory, concepts and paradigms.*

KEYWORDS

sound, memory, technology, destruction, deleting, sustainability, time, alternative design

ACM Reference Format:
Yann Seznec. 2024. Three Ways to Destroy a Sound. In *Adjunct Proceedings of the 2024 Nordic Conference on Human-Computer Interaction (NordiCHI Adjunct 2024), October 13–16, 2024, Uppsala, Sweden.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3677045.3685496

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
NordiCHI Adjunct 2024, October 13–16, 2024, Uppsala, Sweden
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0965-4/24/10
https://doi.org/10.1145/3677045.3685496

1 INTRODUCTION

As a 41 year old who was first introduced to digital technologies as a child, I have already lived through at least four generations of lost media. Floppy disks (let us generously group the two different floppy disks together), zip drives, DAT tapes, and writable CDRs are all now, to various degrees, effectively lost media formats that are at best difficult to access and at worst degraded entirely. It is not difficult to imagine the millions of photographs, sound recordings, videos, and documents around the world that are lost despite being "backed up" onto these supposedly permanent media. We were promised Pure Perfect Sound Forever [10].

This is not an unrecognized problem - an entire community of "media archaeologists" work, such as in the videogame community [1], towards finding methods to get obsolete software and hardware running again so the original experiences can be preserved. And yet despite these efforts, and the recognition of past mistakes, we are currently living through another cycle of promise and loss. As we enter the twilight years of the spinning hard drives, awaiting the day when our Time Machine or Drobo backups will slowly degrade or suddenly fail, our files and memories are now being outsourced to a mysterious "cloud". This time, we don't need to worry, the cloud is limitless and invisible and all-powerful.

If we again take the example of videogames, this scenario does not bode well. The move towards trading control for convenience, accepting software updates and walled garden app stores in exchange for stability and security, has meant the disappearance of an untold number of games that did not meet upgrade requirements or simply pay the fee to be listed on app stores. I, personally, now have virtually no way to play any of the iOS videogames that I developed and released in the early 2010s. These too have been lost (though perhaps in that particular case it might be for the best).

It is effectively guaranteed that the cloud services that are now generally accepted as a “permanent" storage solution will one day fail us. It is difficult to guess precisely how - perhaps general system failure linked to climate change, combined with changing company policies or surprise bankruptcies or software updates or viruses [18] - the list is endless but the outcome is clear.

The illusion of data permanence sits neatly within the current paradigms of technology design, in which limitless options and storage and speed are all presented as the ideal. This so-called “cornucopian" [27] approach to technology mirrors the general glorification of economic growth [9, 16, 31], and is equally unsustainable. We simply can not depend on endless server space - while it may seem unimaginable now, resources are never limitless. On a micro scale this is already noticeable in the experience of constant warnings: “you are out of cloud storage space, upgrade your plan to keep your files". We expand to fill whatever storage space we have (helped by rapidly growing media file sizes due to higher quality capturing devices), far outpacing whatever storage we have available to us.

I have made three devices that explore these issues through the medium of digital sound. As a group I refer to these as the Sound Destruction Machines. The three devices can be summarized thus:

- **The Memory Boxes**: A set of wooden boxes which contain an audio file. Pressing a button simultaneously plays and permanently destroys the sound
- **The Kopiator**: A device which records sound and immediately plays it back and destroys it on loop until it disappears
- **The Sound Destroyer**: A interface that provides fine grained control for how a sound can be recorded, looped, and destroyed

This paper is about these devices, how they work and what they might mean. I will first provide some background on the design, sound, and sustainability concepts that have informed the projects. Following this I will introduce each of the devices, describing them and offering some critical analysis of what they do and represent.

2 BACKGROUND

2.1 Alternative Designs

The creation of the devices described here is quite clearly the result of a Research Through Design process [3, 14, 34], in which the manifestation of the concepts (in this case both the code and the hardware) is prioritized and is seen to contain knowledge in of itself. In particular I draw on Pierce’s definition of *alternative designs* [26] which “prefigure possible, though improbable productions" and resist further development. The friction that I encountered in the development of these designs, and which continues in their realized states, is precisely the point and I intend to avoid the caveats and qualifications that so often accompany this type of work. Rather this paper is perhaps best seen as a form of annotated portfolio [5, 13], in the sense that it describes a set of designs and attempts to tease out a set of themes and “family resemblances to be reasoned about". A paper that describes a set of physical, tangible objects will necessarily be somewhat limited, and is a poor subsitute for holding and listening to an instrument as it destroys a cherished memory or destructively moulds a sound in an unexpected and delightful way. This paper is thus accompanied by a video file showing all three of the instruments as they destroy sounds in their own unique ways, as well as a demonstration of the underlying code running an demonstration on a desktop computer environment. The video and this paper together function as the annotated portfolio, and at NordiCHI will be presented alongside the devices themselves for people to interact with.

2.2 Political Ecology of Sound

Modern sound technology relies on the paradigm of reproducibility - the ability to record, revisit, and endlessly reproduce sound “non-destructively". Tara Rodgers and others have analyzed how the language used for these processes is not accidental [29], but rather reflects how recording technology reflects and perpetuates hegemonic norms; the ideal sound “reproduction" medium is generally considered to embody feminine-coded traits of silence, passivity, and invisibility. Meanwhile, an audio source is often described as “punchy", “clear", or “strong" - and discussions of audio fidelity are generally positioned as the struggle for these masculine traits to be perceived despite the weaknesses of the reproductive surface. These analysis reveals a political ecology of recorded sound, which can be traced back to the use of disembodied human ears gathered from nameless cadavers to build the first sound recording devices [11], continuing through to gendered approaches to design and manufacturing[33]and recently with the application of deep-sea oil drilling technologies to auto-tuning [28, 30].

It is no wonder, then, that there is an embedded expectation for sonic technology to be endless, seamless, and neutral - the aforementioned cornucopian attitude to design and technology that is well integrated into all computing technology.

From a sound design perspective the cornucopian design mentality can be seen not only as harmful from a social justice and material sustainability standpoint but also philosophically problematic. Sound as a medium is uniquely unsuited to cornucopian design principles, due to its fundamentally time-based nature. As Chion and others have explored, sound is completely reliant on (and therefore particularly illustrative of) time [8], and as such listening back to a sound requires inhabiting the same amount of time as it took to record it. Although a library of photographs or digital videos can be perceived at a glance, a hard drive full of sound recordings could take hours to listen to, if not significantly more. Recording sound is effectively an attempt to force a wedge into the continuation of time, on the problematic assumption that we will one day be able to simply make the time to listen back. Thus if sound is inherently linked to our perception of time, why then does modern technology and interaction design insist on breaking this paradigm, relying instead on a cornucopian paradigm that robs sound of the exact thing that makes it powerful?

To work with sound is to confront the issues faced by the world as a whole - a problematic relationship between the past and the present, as well as our attempts to control an uncertain future. Sound is thus uniquely positioned as a medium of experience and expression with regards to sustainability, in the sense of physically embodying the effects of time. Listening to a sound is the act of physically experiencing the consequence of a release of energy in the past, just as the seemingly intractable issues of sustainability

stem primarily from our current experience of actions taken in the past.

Perhaps, then, we must reject the core assumption of recorded sound - that it exists in order to allow us to hear a sound again. What if we could only ever hear a recording once, or indeed if the very act of listening to a recording was the same act that destroyed it? If a recorded sound is no longer reproducible, perhaps it goes some way towards challenging the norms of listening and experiencing sound, as well as the greater notions of political ecologies, limitless growth, and endless storage that are so embedded into technology design and media consumption.

2.3 Limits

Within the Sustainable HCI (SHCI) community, cornucopian design has been heavily critiqued from a sustainability perspective [7], with a number of designers and researchers proposing a "limits" or constraints-based approach [17, 19, 21, 25, 32] to move towards a more socially and technologically sustainable future. This project takes this concept and considers how it overlaps with our use of media and storage. As with most of the issues in sustainability studies, the facts are clear (it is already too late to stay below 1.5°, and infinite data storage is impossible) even if the calls to action are confusing at best (recycle your straws and buy more cloud storage).

I would like to position the work in this paper alongside other HCI approaches to the overwhelming mass of media and the heavy weight of history that goes along with it. The three devices described here are referencing projects like the OLO Radio [23], which both encourages (and to a certain extent problematizes) a listener's relationship with their past, and the tradition of autodestructive music and art typified by Marclay's *Record Without a Cover* or even Lucier's *I'm Sitting In A Room* [15, 20] among others. There are also echoes of the research projects looking into the vast array of digital content being left behind when people die [6, 22], and by extension how designers can consider time [12]. Where those projects consider the digital artifacts people want to keep forever, SHCI research points out that our digital footprints necessarily need to shrink (alongside our more recognizably material footprints).

This project thus accepts that we must move towards deletion. As designers we need to consider how we can provide the tools to delete and explore the experience of destroying media. The Sound Destruction Machines represent an approach that embraces the limits in order to focus on intentionality and meaning. They are here to help.

3 HOW TO DESTROY A SOUND

3.1 The sound destruction engine

The three devices in this set all use the same underlying system for destroying the sound. As such in some ways they represent how the presentation and interaction are what actually create the experience (more so than the technical system), but nevertheless let us describe how it works.

In order to position the work within contemporary media consumption and technology design the project focuses exclusively on digital sound files. The act of sonic destruction is accomplished through a wavetable playback system that overwrites the soundfile. This allows for the destruction to be permanent, constant, to a certain degree controllable, and arguably aesthetically pleasing (or at least understandable).

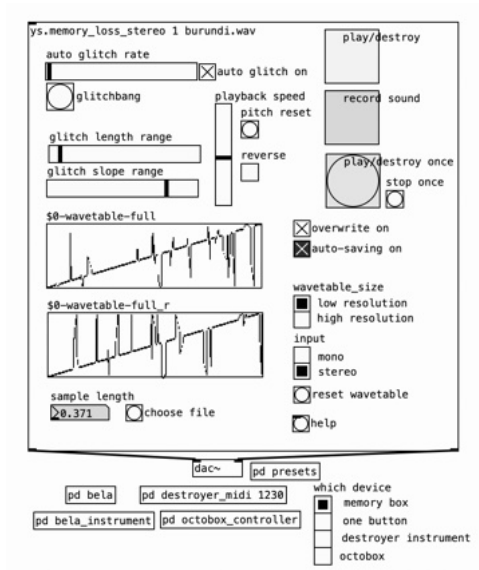


Figure 1: The sound destruction engine, destroying a sound

A sound file exists on a digital storage medium, most likely a flash disk. For the purposes of this engine it is comprised of 44100 data points per second of sound, stored in an array, each point representing a position over time of the reproduction medium that will eventually be driven to make the sound audible.

In order to drive the playback of this sound file, a second array is created that represents time. This is essentially a graph that begins as a line running from 0 to 1 over the same length of time as the source sound. Applying this graph to the audio file results in a linear playback of the sound - this is known as a wavetable.

Adjusting the wavetable to be something other than a diagonal line introduces variability into the playback of the sound file. If a section of the line has a negative slope then the corresponding section of sound will play backwards. Sharp slopes will play the sound file more quickly, and a horizontal line will stop playback entirely. All of the resulting sound is still being generated from the original source media data - introducing a number of "errors" to the wavetable can be compared to cutting up and re-assembling a photograph.

The destruction engine can introduce these errors, or "glitches", into the wavetable automatically using a set of semi-randomized parameters. Two controls are available in the code to control the size and characteristics of a glitch, dictating the range of random numbers that will be used to generate the width (time) and height (playback rate) of a glitch somewhere in the wavetable. These controls

are set and then the resulting glitches are generated and applied on command. The glitch command can be triggered individually or be scheduled to occur at a given rate (relative to the length of the original source file). The rate of playback can also be controlled, including whether it is played forward or in reverse.

In order for the destruction to be permanent, when the audio is played back the output is routed to overwrite the original file with a short delay (to allow for the playback and overwriting to happen simultaneously). In the engine this feature is switchable, so that the glitched playback can occur with or without overwriting the audio file. This method for making the destruction of the sound permanent can only occur while the sound is being played back, since it is the sonic output that is being stored into the array. The listener therefore hears the destruction as it happens.

The bulk of this paper will address what happens when this system is given form through interaction. Each of the three devices described below reveal different aspects of the destruction engine to the user through interface elements, and other aspects to obscure or control automatically. Naturally, each of these decisions changes how the resulting interaction takes shape and what type of device is generated.

3.2 The Memory Box

The first device that uses the destruction engine is The Memory Box, a device designed with the aim of generating meaningful sonic destruction. As such it requires a meaningful sound recording in order to function. Related to this, the sound contained within the Memory Box needs to be the only remaining copy of that sound recording. Only with these two rules satisfied can the device provide meaningful destruction of the sound.

I have built three of these devices. They are hexagonal wooden boxes that plug into a mains power outlet and contain a single visible button, which glows white when ready. Some programming work is required to load a sound into the electronics inside. When ready, pressing and holding the button activates playback, which is heard through a built-in speaker that vibrates the box. The sound plays for as long as the button is held, and a glitch is introduced on each loop. The overwrite function is activated in the code, so each time the sound is heard and a glitch is introduced the change is permanent. It is difficult to estimate how long a sound will last in the Memory Box, but after roughly 100 loops the original sound has been completely lost and all that remains is slowly fading noise.

What constitutes a meaningful sound? If I open my cloud photo library I can instantly listen to the sounds that accompany 1,215 video files recorded over the past ten years. Presumably I could randomly and blindly choose any one of those and delete it, and I would not ever notice it was gone. However it is equally likely that if I opened one at random to actually watch (something approaching the Photobox mechanic [24]) I would become unreasonably attached to that video and the feeling of deleting it permanently would land somewhere between disappointment and rage. The personal attachments I have to these files are, then, perhaps recreated whenever I perceive them - however this is something that I very rarely do. I simply do not have the time or mental capacity to browse 1,215 files.



Figure 2: Three Memory Boxes

How, then, would I go about selecting a sound that would be destroyed and whose destruction would carry meaning and weight? Even more challenging, how can I work with someone else to help them select a sound for themselves? And if we do manage to find a sound that they are willing to consider permanently destroying, how do they accept this?

Ensuring the uniqueness of the sound recording in the Memory Box is something of a challenge, due to the nature of modern file storage and sharing. Deleting a video file from an iCloud photo library, for example, simply moves it to a password-protected “trash" which is is only “emptied" after 30 days unless you do so manually. Even if this trash is deleted there is no guarantee that it does not continue to exist on a server, and any number of permutations exist wherein the file will not be removed from some devices (if files sync locally and are cached on a backup, for example). Messaging a file will result in multiple copies stored on various servers; cloud storage systems bias strongly against permanent deletion; the list goes on.

Due to these factors, the loading of the sound into the Memory Box devices is an action that can only be performed by me, the designer. It requires a set of relatively specialist knowledge regarding how the code needs to be set up and how the sound is loaded onto the machine. I therefore become a mediator of the media. I take a sound, prepare it for the device, load it on, and ensure that it is deleted everywhere else. Only then does it begin to function.

The outcome of these design choices is a relatively passive device that contains a sound that can only be heard if it is also simultaneously destroyed. The device can sit quietly and never be used, preserving the sound and becoming something of a visual representation of the sound file that exists only on the memory chip inside. Where photos can be printed and hung on the wall, this device offers the potential for a sonic memory to form a part of the ambient decoration of a space. However the soft glowing light and gentle form also comes as something of a threat - pressing the button, and listening to the sound that is generated, is a permanently destructive act.

Giving The Memory Box to others to use would require a careful, and very personal, approach. One scenario for giving these devices



Figure 3: The Kopiator



Figure 4: The Sound Destroyer

to others could involve working very closely with two volunteers, each of us having one of the devices. We would each explore our media libraries and identify a recording to be placed in the Box. If each person feels comfortable sharing, we could listen to each other's choices and discuss why it is meaningful. I would then "put" the recordings in the boxes and we would work together to ensure that every possible copy of the recording has been deleted elsewhere to the best of our knowledge. We would then say goodbye to the sounds, and each of us would take the devices home to live with them for some time.

3.3 The Kopiator

The second implementation of the destruction engine is designed for use in an environment with many people passing through, perhaps a hallway or lounge or similar liminal space. A prominent microphone protrudes from an otherwise nondescript metal enclosure that would not feel out of place in a conference room or lecture hall. Below the microphone is a single glowing button. The microphone will record any sound that is made while the button is held down, and the system will play the resulting recording on loop the moment the button is released. The sound will continue to play on loop, with the auto-glitch system in the software introducing errors on each playback, which are permanently written to disk. The glitch settings are high enough that the sound is destroyed and completely disappears within a couple of minutes. Pressing the button quickly at any time will stop the playback.

Where the Memory Boxes function as a representation of the past, the Kopiator is primarily a reflection of the immediate present. Any sound that is recorded represents the now, and the resulting playback is a sonic manifestation of the time period the button was held down. Any sound that occurred during that time is then played back and repeated until it disappears. Thus the device blurs time - a sound which would normally have already become a part of the past is instead prolonged. It still fades away, but it takes a bit more time and the listener is much more aware of the mechanism for how it is destroyed.

It is also, effectively, a toy. It lends itself to funny interactions, especially as people first learn to use it. With no instructions on the device, it can take a few button presses before people understand how it works (if, indeed, they figure it out at all). This often takes the form of someone saying "what does this do?" as they press the button, resulting in a fragment of the sentence being immediately played on loop with compounding errors "this do? this do? thiis do? tehisd do? ttechissdo? tetechischooo?"

3.4 The Sound Destroyer

Considering this work as a set of devices that put form to the same underlying code, it is tempting to imagine a sort of gradient between two extremes of control. Perhaps on one side is a physical manifestation of the code which is entirely automated with no user control whatsoever. Arguably this would be so extreme as to cease to function, since by removing any human interaction there would be no method for generating or selecting a sound. One step into the gradient brings us to something like the Memory Box, where a single button applies a number of predetermined actions. This creates a highly personalized experience which requires no prior training or knowledge but also offers no control (apart from the choice of whether or not to push the button) and relies on the mediation of an expert.

On the other extreme, then, we can make every control available. The resulting object begins to function as something of an instrument, in both the musical and scientific sense. The Sound Destroyer has four knobs, seven buttons, and an LED which indicates whenever a sound has been recorded or a glitch implemented. A footswitch input allows for hands-free recording and looping and audio is routed using input and output jacks.

Simply the existence of these jacks immediately moves this into a different category of interaction - they require (or assume) a level of existing infrastructure and by extension knowledge. In order for the instrument to function as a sonic object the user requires a microphone, possibly a preamp, and some kind of amplification or listening system. There is no volume control on the inputs or outputs, this has been effectively outsourced to the user along with an assumption that they have the knowledge for how an input signal can be generated and controlled and an output signal managed and routed.

The knobs and buttons are labeled and mapped to each of the available parameters in the code. While a number of these (record,

play, forward/reverse, etc) are digital signals that have only two possible states, the potentiometer knobs are analog inputs that can send a range of values. These values are then scaled to an appropriate range for the associated parameter. The total control that the instrument suggests is thus something of an illusion, since this scaling must be decided at some stage - the autoglitch speed parameter, for example, ranges from 1-10x per loop. Why not 50? 100? 1000? Why can the playback pitch knob go from 0-4?

These decisions, made by myself based on my own opinions and aesthetic desires, can be seen as analogous to tuning an instrument. A flute would vibrate a column of air in exactly the same way with a wide range of placement and number of holes, but their precise location has been chosen for aesthetic purposes, and after many generations the specific design has crystallized and standardized. The Sound Destroyer instrument has nowhere near the maturity of even a relatively new instrument, and if it were to develop further than perhaps the "tuning" of early version like this would be considered frightfully old fashioned.

This is an instrument that is designed to be played. It is unwieldy and somewhat mindbending to keep track of what is happening at any given point, particularly given the lack of screen. The visual feedback is extremely limited, forcing the player to focus on the sound. The existence of all of the controls can also lead to problematic, or at least unintended results. The pitch knob, for example, changes the playback speed of the current sound buffer. However if the "permanence" switch is activated, the overwriting of the buffer will still occur at the original playback rate (due to limitations in the audio library used for writing signals). It is virtually impossible to anticipate the outcome of this mismatch, as it depends on a complex relationship between the relative timings.

Nevertheless, with repeated use it becomes possible to reach a level of understanding to the point where a player can improvise a performance from an audio source. The instrument acts as something in between a looper and an effect, and the act of destruction can be as permanent or as temporary as desired. If the permanence switch has not been activated, the reset button can bring the original sound back from apparent death. Or, with some care, a destroyed sound can be "reset" to a semi-destroyed state and re-destroyed in a different fashion. The aesthetics of the destruction can be adjusted in real time; a high value "height" combined with a short "length" will generate noticeable but short glitches, while the opposite gives the potential for large sections of the sound to be momentarily slowed to a crawl before continuing at the original rate.

I say "a player", with an apparently impartial tone, but the reality is I am the only one to have played it. It remains to be seen whether this design, and the aesthetic choices embedded in the destruction engine, would be equally expressive and satisfying to someone else. I suspect that the instrument would function best with a player that has had the system explained to them, though perhaps my own knowledge is in fact holding back my ability to perform with it in an expansive way.

4 FINAL THOUGHTS

Sound is a time-inefficient medium. It is much faster for you to read this sentence than it would be for me to read it out loud. Certainly the same is true of this method for destroying a sound - it would be far quicker to simply delete the file with the press of a button rather than listen to it slowly decay. However this project has never been about efficiency, but is rather built on sustainability concepts of degrowth, anti-cornucopian design, social justice, and care. As such the destruction engine and the three manifestations of the code into objects suggest three different approaches to deletion and destruction that apply these concepts to different degrees whilst eschewing solutionist tendencies of optimization.

From a purely material standpoint, it is also worth acknowledging that these processes of "destruction" do not include any material change. Although the scales are unimaginably small, the same amount of physical space is being taken up by the audio buffers, whether or not the content within them has been distorted beyond all recognition. If a 3 second audio file is in the Memory Destroyers, an audio buffer of 132300 points is holding that sound, and each glitch simply shifts where those points are without shrinking the overall size. While this work has never claimed to provide a solution to the constantly growing cloud and the illusion of infinite resources, from a design perspective it is important to keep this in mind and consider what this represents. If we compare the destruction engine to a more analog process, the equivalent could be to pass a magnet over a section of audio tape. This makes some sense from a sustainability point of view, as the tape (or, indeed, the digital disk space) is thus made available for re-use, whereas any other method of destruction - perhaps cutting and destroying - would necessarily involve some form of waste (slash and burn, one could say). Where projects such at ChronoTape [4] and The Disintegration Loops [2] have taken advantage of the materiality of tape in particular, the Destruction Machines focus on how this can be approached with digital formats.

Indeed, what can actually be destroyed? This work separates out the need for material degrowth and destruction from the experiential act of deleting and letting go. On an emotional level it seems callous and insensitive to discuss destruction in such a casual way. When describing this work to others it is not unusual for people to react viscerally when faced the prospect of losing treasured media artifacts for the sake of some conceptual interaction design research project. A technology that kills a sound each time it is heard punishes the very act of listening and the instinct to preserve memories. However in doing so, this approach also positions a recorded sound as something that *can* be killed. It effectively breaks the concept of a recording as something that can be revisited at any point. Rather, a sound becomes in some way alive, and something that must be listened to with care.

ACKNOWLEDGMENTS

This work is supported by the Energimyndigheten (Swedish Energy Agency) project *Interaktionsdesign för ett Soldrivet Internet* (Interaction Design for a Solar Powered Internet), 51885-1. I would also like to thank Rob Comber for support and guidance in writing the paper, and Linus Remahl and Dee Tobin for help with the design and construction of the devices.

REFERENCES

[1] John Aycock. 2016. *Retrogame Archeology.* Springer International Publishing, Cham. https://doi.org/10.1007/978-3-319-30004-7

[2] Lindsay Balfour. 2018. Traumatic Ruins and the Archeology of Sound: William Basinski's The Disintegration Loops. *Journal of Sonic Studies* 16 (2018). https://pureportal.coventry.ac.uk/en/publications/traumatic-ruins-and-the-archeology-of-sound-william-basinskis-the

[3] Jeffrey Bardzell, Shaowen Bardzell, and Lone Koefoed Hansen. 2015. Immodest Proposals: Research Through Design and Knowledge. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, Seoul Republic of Korea, 2093–2102. https://doi.org/10.1145/2702123.2702400

[4] Peter Bennett, Mike Fraser, and Madeline Balaam. 2012. ChronoTape: tangible timelines for family history. In *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction*. ACM, Kingston Ontario Canada, 49–56. https://doi.org/10.1145/2148131.2148144

[5] John Bowers. 2012. The logic of annotated portfolios: communicating the value of 'research through design'. (2012).

[6] Jed R. Brubaker and Gillian R. Hayes. 2011. "We will never forget you [online]": an empirical investigation of post-mortem myspace comments. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*. ACM, Hangzhou China, 123–132. https://doi.org/10.1145/1958824.1958843

[7] Jay Chen. 2016. A strategy for limits-aware computing. In *Proceedings of the Second Workshop on Computing within Limits*. ACM, Irvine California, 1–6. https://doi.org/10.1145/2926676.2926692

[8] Michel Chion. 1994. *Audio-vision: sound on screen*. Columbia University Press, New York.

[9] Giacomo D'Alisa, Federico Demaria, and Giorgos Kallis (Eds.). 2015. *Degrowth: a vocabulary for a new era*. Routledge, Taylor & Francis Group, New York ; London.

[10] Kieran Downes. 2010. "Perfect Sound Forever": Innovation, Aesthetics, and the Re-making of Compact Disc Playback. *Technology and Culture* 51, 2 (April 2010), 305–331. https://doi.org/10.1353/tech.0.0442

[11] Tom Everrett. 2021. Writing sound with a human ear: reconstructing Bell and Blake's 1874 ear phonautograph. *Science Museum Group Journal* 12, 12 (Nov. 2021). https://doi.org/10.15180/191206

[12] Verena Fuchsberger, Martin Murer, and Manfred Tscheligi. 2015. Time and Design: Seven Sensitivities. In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, Stanford California USA, 581–586. https://doi.org/10.1145/2677199.2687911

[13] Bill Gaver and John Bowers. 2012. Annotated portfolios. *Interactions* 19, 4 (July 2012), 40–49. https://doi.org/10.1145/2212877.2212889

[14] William Gaver. 2012. What should we expect from research through design?. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Austin Texas USA, 937–946. https://doi.org/10.1145/2207676.2208538

[15] Stina Hasse. 2012. I am Sitting in a Room. *Body, Space & Technology* 11, 0 (Jan. 2012). https://doi.org/10.16995/bst.71

[16] Jason Hickel. 2022. *Less is more: how degrowth will save the world*. Penguin Books, UK USA Canada Ireland Australia India New Zealand South Africa.

[17] Lorenz M. Hilty and Bernard Aebischer. 2015. ICT for Sustainability: An Emerging Research Field. In *ICT Innovations for Sustainability*, Lorenz M. Hilty and Bernard Aebischer (Eds.). Vol. 310. Springer International Publishing, Cham, 3–36. https://doi.org/10.1007/978-3-319-09228-7_1 Series Title: Advances in Intelligent Systems and Computing.

[18] Esther Jang, Matthew Johnson, Edward Burnell, and Kurtis Heimerl. 2017. Unplanned Obsolescence: Hardware and Software After Collapse. In *Proceedings of the 2017 Workshop on Computing Within Limits*. ACM, Santa Barbara California USA, 93–101. https://doi.org/10.1145/3080556.3080566

[19] Bran Knowles, Oliver Bates, and Maria Håkansson. 2018. This Changes Sustainable HCI. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–12. https://doi.org/10.1145/3173574.3174045

[20] Christian Marclay. 1985. Record without a Cover.

[21] 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

[22] William Odom, Richard Banks, and Dave Kirk. 2010. Reciprocity, deep storage, and letting go: opportunities for designing interactions with inherited digital materials. *Interactions* 17, 5 (Sept. 2010), 31–34. https://doi.org/10.1145/1836216.1836224

[23] William Odom and Tijs Duel. 2018. On the Design of OLO Radio: Investigating Metadata as a Design Material. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–9. https://doi.org/10.1145/3173574.3173678

[24] William Odom, Mark Selby, Abigail Sellen, David Kirk, Richard Banks, and Tim Regan. 2012. Photobox: on the design of a slow technology. In *Proceedings of the Designing Interactive Systems Conference on - DIS '12*. ACM Press, Newcastle Upon Tyne, United Kingdom, 665. https://doi.org/10.1145/2317956.2318055

[25] Daniel Pargman and Barath Raghavan. 2014. Rethinking sustainability in computing: from buzzword to non-negotiable limits. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational*. ACM, Helsinki Finland, 638–647. https://doi.org/10.1145/2639189.2639228

[26] James Pierce. 2021. In Tension with Progression: Grasping the Frictional Tendencies of Speculative, Critical, and other Alternative Designs. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–19. https://doi.org/10.1145/3411764.3445406

[27] 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 *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 1324–1337. https://doi.org/10.1145/2858036.2858378

[28] Catherine Provenzano. 2018. *Auto-Tune, Labor, and the Pop-Music Voice*. Vol. 1. Oxford University Press. https://doi.org/10.1093/oso/9780199985227.003.0008

[29] Tara Rodgers. 2010. *Pink noises: women on electronic music and sound*. Duke University Press, Durham [NC].

[30] Yann Seznec. 2022. Music Within Limits. In *2022 International Conference on ICT for Sustainability (ICT4S)*. IEEE, Plovdiv, Bulgaria, 172–178. https://doi.org/10.1109/ICT4S55073.2022.00028

[31] Vishal Sharma, Neha Kumar, and Bonnie Nardi. 2024. Post-growth Human–Computer Interaction. *ACM Transactions on Computer-Human Interaction* 31, 1 (Feb. 2024), 1–37. https://doi.org/10.1145/3624981

[32] Martin Stojanov, Daniel Pargman, Mike Hazas, Rob Comber, and Jorge Luis Zapico. 2023. How do we arrive at constraints? Articulating limits for computing. In *Ninth Computing within Limits 2023*. LIMITS, Virtual. https://doi.org/10.21428/bf6fb269.a317d18f

[33] Lucie Vágnerová. 2017. 'Nimble Fingers' in Electronic Music: Rethinking sound through neo-colonial labour. *Organised Sound* 22, 2 (Aug. 2017), 250–258. https://doi.org/10.1017/S1355771817000152

[34] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 493–502. https://doi.org/10.1145/1240624.1240704

Paper D

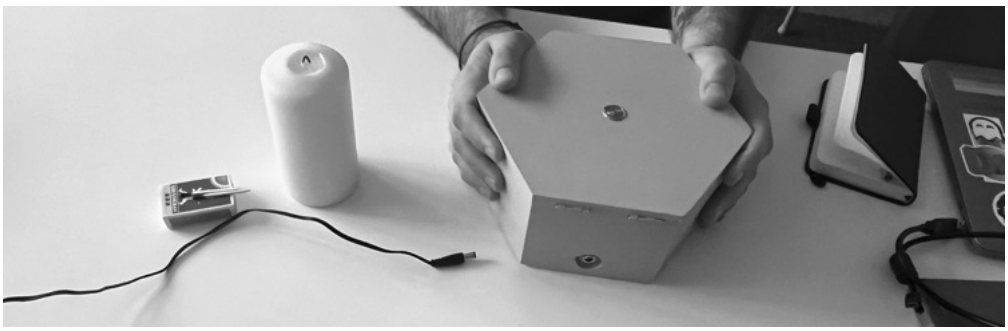
Play/Destroy: A portfolio of sound destruction devices

Yann Seznec
Department of Media Technology and Interaction Design
KTH Royal Institute of Technology
Stockholm, Sweden
yannse@kth.se

Andreas Lindegren
Department of Media Technology and Interaction Design
KTH Royal Institute of Technology
Stockholm, Sweden
alindeg@kth.se

Frida Lindegren
No affiliation
Stockholm, Sweden
lindegren.frida@gmail.com

Rob Comber
Department of Media Technology and Interaction Design
KTH Royal Institute of Technology
Stockholm, Sweden
rcomber@kth.se



Abstract

Digital media operates on a curious boundary between storage and loss. While each new storage format promises a permanent solution to our exponentially expanding media libraries, they inevitably fail or otherwise become unusable. This paper reflects on a long-term design process that attempts to bring a different paradigm to the experience of personal digital media: destruction. We present an annotated portfolio of a set of sound listening devices, critically unpacking the particular temporal, perceptual, and experiential qualities that emerge when designing for the loss of personal media. These annotations show how destruction comes to matter in designing against the traditional bias towards growth and accumulation in HCI.

CCS Concepts

• **Human-centered computing → Interaction design theory, concepts and paradigms**.



This work is licensed under a Creative Commons Attribution 4.0 International License.
CHI '26, Barcelona, Spain
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2278-3/26/04
https://doi.org/10.1145/3772318.3791701

Keywords

Research through Design, sound, impermanence, sustainability, destruction, seamful design

ACM Reference Format:
Yann Seznec, Andreas Lindegren, Frida Lindegren, and Rob Comber. 2026. Play/Destroy: A portfolio of sound destruction devices. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain*. ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3772318.3791701

1 Introduction

Media storage has a problematic relationship with permanence [5]. Floppy disks, zip drives, DAT tapes, and writable CDRs are all now, to various degrees, effectively lost or disappearing media formats that are at best difficult to access and at worst degraded entirely - despite each of these being portrayed when launched as a permanent solution to the problems of digital storage. Sustainable HCI (SHCI) [4] research implies that it is effectively guaranteed that the cloud services that are now generally considered "permanent" storage solutions will one day fail us too [65]. It is difficult to guess precisely how — future failures could come as a result of climate change, changing company policies, compromised security systems, surprise bankruptcies, software updates, ransomware or viruses [31]. The digital data systems on which we now rely are increasingly built on the impossible promise of endless growth

[38, 50, 59], with exponentially expanding storage sizes supported by correspondingly increasing energy and resource use [16]. Nevertheless they are just as prone to failure and disappearance as previous generations of storage.

And yet the bias towards permanence is an approach that is so ingrained in the design process as to be unquestionably implemented in digital tools and accepted wherever possible [5]. When our media libraries inevitably outgrow our current allocation it is far easier to purchase a storage upgrade than to delete items. After all, deleting our own personal media is a fraught experience. The combined forces of high quality media capture, seamless storage, and integrated software experience design (among many other factors) result in a user taking potentially dozens of photographs or videos of what are arguably the same "moment", either with the intention of revisiting them to choose the best one or perhaps without even the pretense of that curation at all. When faced with the inevitable reminder that we have filled our remaining device or cloud storage we experience stress through loss aversion [32], as we desperately trawl through our libraries, ruthlessly deleting four or five of the twenty pictures of the rainbow we saw at the beach, perhaps with a sense of injustice that this was not "supposed" to happen [40]. In a somewhat ironic twist, the standard deletion process itself is not even as final as it may seem, as it rarely actually removes the data anyway [19], putting users into the odd situation of being regularly forced to make a decision that does not have the impact that it appears.

To engage with the interwoven issues of digital permanence and loss, this paper presents an annotated portfolio of devices for listening to and simultaneously destroying a sound. These devices both encapsulate the sound and are the site of their end: using the device will also make the recording slowly disappear forever. First, we cover media storage and sustainability in HCI, showing how they overlap in this design space. Then we present a lineage of designed devices, before presenting annotations centered around the themes of destroying time, impermanence in media, tensions of storing/remembering, and decay. Finally, the portfolio is discussed in relation to sustainable computing and HCI, proposing alternative interactions stemming from a design space that is generated by embracing the realities of digital impermanence and destruction. The work offers three contributions: A portfolio that re-integrates seamfulness into the experience of digital sound, a strategy for embracing impermanence in design for personal media, and a call to integrate deliberate destruction into HCI.

2 Background

2.1 Seamlessness and the Cornucopia

For a number of years the limitations of media storage were an inescapable reality embedded into the design and use of technology. The very existence of the removable storage formats mentioned in the introduction was due to these limitations: digital files needed to be backed up to physical media, text messages had strict size and memory restrictions [63], standard film rolls could hold 24 or 36 exposures, and the storage of sound required a decision to choose between fidelity and length [37].

These limitations could be considered as a type of seam [11–13]: a moment when the technology underpinning an interface becomes overtly noticeable. Seams can enable a deeper understanding of the technology involved, and offer certain powerful moments of interaction - choosing exactly when to pause a cassette tape, or deciding which text messages are least important. Nevertheless, the bias towards growth in technology has encouraged a form of seamlessness, in which the limitations of a system are made as invisible as possible. Sustainable HCI literature has described this bias as the *cornucopian* paradigm[50], in which seamless and endless storage and speed are considered the ideal towards which all computing should strive. This approach to technology mirrors the general glorification of economic growth as a measure of progress and is equally unsustainable.

Despite efforts to portray our media ecology as limitless, a number of seams remain in digital media, such as the occasional stark reminder of our need to delete some photos from our somehow constantly full cloud library, or the difficulty of finding a specific cherished video in a vast sea of files. The challenges of dealing with storage and presentation of media are not unknown to HCI, as we can see research into growing photo archives dating back at least 20 years [24], which led towards understandings of the limitations of "find and retrieve" as a mechanic in personal archiving practices [33]. Partially in response to this, a number of design research projects have since addressed how to explore and give meaning to large media archives, most notably within the "slow technology" movement [25]. This includes Will Odom's Olly [43] and the OLO Radio[41], both of which empower users to explore their music listening past. Other recent instances which apply this philosophy include The Chronoscope [14], that uses metadata to allow the user to browse through photographs on specific temporal cross sections, and Photobox [42], which occasionally prints random photographs from a user's past. All of these designs push our interaction models for the temporality of our media archives, from the momentary access to an old photograph, to the granular navigation of a listening past. They weave a backwards projection from 'now' to 'then' to enrich the otherwise invisible media of our past. These projects all engage significantly with large scale media libraries, considering how to allow users to rebuild meaning with them and recreate connections through surprising encounters with their own past. Overall they do not necessarily question the existence of these libraries, or reflect on the hoarding of media as a result of the embedded assumptions in the design of the technologies.

A separate body of work in HCI has considered different specific situations where media needs to be managed through deletion. These explorations have included grappling with post-mortem deletion of social media accounts [18], the need for "explainable" or "understandable" deletion [52, 57], and helping people delete media as a form of "letting go" as part of a breakup or bereavement process [55, 56]. Other works such as Delete by Haiku [60] attempted to enliven the act of deleting by encouraging users to delete text messages by playfully turning them into haiku. This work all recognizes the need for interaction design research to engage with deletion and loss, though the deletion itself is still enacted within standard HCI paradigms.

Hardware destruction has also been explored as an interaction, such as with the Obscura Camera: a counterfunctional camera that must be smashed to pieces in order to access the photos [49].

Rossmy et al. described several projects exploring "irreversible interactions" [54], which looked at the experience of irreversibility (a glass breaks and can not be unbroken), pointing out how it clashes with the computing heuristic of reversibility (endless undo).

When taken as a whole, this body of work points towards a general bias against deletion in technology systems, which is manifested through obscure or misleading interaction design. The work presented in this paper seeks to extend the design work associated with deletion. This is partially a response to the problematic growth of personal media libraries, and the embedded assumptions about it in design, from both a sustainability perspective as well as an experiential one.

In adopting deletion at the core of our design, we use a constraints-based approach [1, 30, 34, 38, 45, 61] to addressing the fundamental SHCI concepts of cornucopian paradigms and collapse computing through design. A constraints-based approach is an iterative process of design away from growth, asking what is needed, can be removed or discarded [1, 46]. Pierce [46] argues for multiple means of negation in and through design —such as displacement and erasure— with the possibility to close off forms of over-consumption and unsustainability in design. Bakhshoudeh and Comber [1] suggest introducing friction in interaction, while simultaneously and iteratively stripping away the unnecessary clutter of our digital interactions. Applying this approach to deletion in particular led to the questions examined in this paper — how can we examine the concept of deletion in a way that explores the experience and points towards strategies and methods that can help us embrace the inherent impermanence of digital media?

2.2 Destruction in contemporary art and music

The designs that form the basis for our work focus on sound, leveraging the particular relationship between sound and time as described by Chion and others [15, 67]. Within sound art the concept of auto-destructiveness has been explored notably by William Basinski with The Disintegration Loops [2], by Christian Marclay with his Record Without a Cover [36], and perhaps most famously by Alvin Lucier with I'm Sitting In a Room [26]. All of these works looked at the aesthetics of destruction, embracing the imperfections that are gradually amplified by time, revealing the underlying characteristics of the materials and spaces involved in creating and listening to the original sound.

Similarly, other non-sound related artwork has used destruction as a creative tool. Yoko Ono's Cut Piece used scissors to reveal embedded gender violence [9], Banksy's Love is in the Bin destroyed a painting to question the value of visual art[1], and the KLF allegedly burnt a million pounds to underscore their complex relationship with capitalism [29]. In each of these cases (of which many more can be found) the destruction functioned as a method for exposing existing social tensions, much as Basinski's and Lucier's work exposed the characteristics of their sonic materials.

In summary, this project looks at the design space of personal media libraries, viewing them as a particular application of a seamless cornucopian computing design paradigm. This paradigm suggests suggests that HCI is fundamentally biased against deletion and erasure. We build on previous work and draw on HCI research as well as contemporary art and music with the aim of exploring an explicitly anti-cornucopian strategy for engaging with these issues. In this paper we offer a set of artifacts resulting from a long-term design process that prioritized and privileged the seams of destruction, deletion, and permanent loss.

3 Methodology

The method for examining these issues focused on building a set of objects to play and think with. This work has thus taken the form of a Research through Design (RtD) project, in which the process of designing, and the resulting outcomes, form the basis from which insight and knowledge can be drawn. This is specifically in line with Gaver's description of RtD [22], which emphasises that research can emerge from the design work [6, 23], and is well established in HCI [27].

We would like to also draw on the work of Desjardins and Oogjes, who argue that linear temporal narratives of Research through Design projects can be unhelpful at best [44]. Realistically the beginnings of this work can be traced as far back as eight years, with the Book of Knowledge of Impractical Musical Devices project[2], and involved a number of what Desjardins and Oogjes might refer to as "pauses" and "transitions" —not the least the first author's shift from independent artist/performer to full-time academic researcher. Much of what is portrayed here as "background" or "context" was in fact revealed throughout the design process, and conversely many of the concepts that were used as background to generate the design space from which the project emerged were later discarded. The work therefore is a product of a long and irregular process, spanning many years and different contexts. The outcomes, both in terms of design and knowledge, reflect this temporal irregularity and are enriched by it.

Related to this, throughout the paper, we will alternate between the first person singular and plural, I and we, to reflect the role of the first author as designer (I) and the other authors as co-discussants of the design work (We).

More concretely, the method for presenting the work will be an annotated portfolio, as outlined by Gaver, Löwgren, and Bowers, among others [6, 20, 35]. Löwgren recognizes the annotated portfolio as a form of intermediate level knowledge that can act as a bridge between design practice and academic research [35] and Bowers in particular argues that while the insights they generate may remain relatively local to the particular portfolio, they can nonetheless perform a similar function to design theory in clarifying knowledge [6]. Although in many cases annotated portfolios are used to analyze a set of different designs (for example [7]), in other cases a portfolio can be comprised of an iterative lineage of work or the development of a single prototype [62, 66]. The designs here all share a number of characteristics but changed over the course of a number of years. Presenting an annotated portfolio is also a deliberate choice to help distinguish the work from other design or science traditions which might come with embedded statistical, solutionist, or data-driven expectations. Rather the designed exemplars outlined here can be best contextualized by Pierce's definition of *alternative designs* [47] which "prefigure possible, though improbable productions" and

[1] https://www.sothebys.com/en/articles/latest-banksy-artwork-love-is-in-the-bin-created-live-at-auction

[2] https://www.impracticaldevices.com

resist further development. Examining the tensions in the development of any productions, however improbable, in response to sustainable HCI concepts leads also to our use of Pierce's concept of undesign [46] to help us manage the apparent conflict between proposing new designs that attempt to address entrenched issues of previous technologies.

The paper will thus continue by presenting a portfolio of work. This will provide as much detail as possible in terms of the technical design, interaction possibilities, and the process that led to their conception. It is worth noting at this stage that while Bowers in particular has used annotated portfolios to present sound-based work [8], this remains the exception rather than the rule. A visual portfolio (and indeed the PDF required to satisfy the requirements of academic publishing) will always be a challenging method to showcase audio-focused designs, particularly when the actual material of the sound is being discussed (and, indeed, destroyed). Therefore the "portfolio" as presented here will focus on text-based descriptions of the work, accompanied by a range of images in the document showing the software, sonic output, and interface designs as much as possible. Links to further video and audio documentation and a downloadable version of the software are available as supplementary material.

Following this portfolio, we offer four annotations — Destroying Time, Impermanence, Storing/Remembering, and Decay. These annotations bring together themes that emerge from the portfolio, and will be followed by a discussion which looks at how the annotations can be related to broader questions of the use of destruction as a form of media interaction and as a method of exploring alternative paradigms in the service of sustainable computing.

4 Designs

The designs described here all operate on the principle that hearing and destroying a sound are intrinsically linked — each time the sound is heard, it is also slightly but permanently destroyed. Although technically this could be applied to any sound, for the purposes of this portfolio the designs here are all intended to be used as personal media storage devices, containing (at least conceptually) the only remaining copy of a sound recording.

This underlying principle — linking hearing with destruction — stemmed originally from a discussion with an art conservator who works primarily with oil paintings. She characterized her work as a balance between the enjoyment of art and its eventual inevitable decay. This is because the visible light that enables a viewer to see an artwork is the very same light that is effectively destroying the painting. A painting could theoretically last many thousands of years if kept in a pitch black cave, as the colors would never be robbed of their pigments by visible light, however that would come at the somewhat obvious tradeoff of never being seen (which is arguably the whole point of the painting). This led to a deceptively simple question originally posed by text message from the first author's wife[3]: could this be applied to sound? **What if listening to a digital sound recording was the same act that made it disappear?**

4.1 The Sound Destroyer Engine

The designs described here all run a version of the same code. The hardware designs described in the following sections all contain this code and it is thus tempting to consider the code simply as an integral part of each of these devices. However as the code was developed first and given that it is capable of being used on its own (on a computer, rather than exclusively on an embedded device), it is worth presenting it as a separately designed element.

The Sound Destroyer Engine (fig. 1) is a set of code written in Pure Data[4] that can be configured in a number of different ways via user interface, but is principally a system for controlling how a digital sound file can be played back with varying degrees of permanent (or semi-permanent) destruction. This is accomplished using wavetable playback, where a sound is loaded into a buffer, the timing of which which is then controlled by a separate buffer (the wavetable). When the wavetable takes the form of an unbroken diagonal line this will generate a linear playback of the sound. The slope of the line dictates the speed of playback, which is perceived as pitch.

The destruction is accomplished by inserting semi-randomized glitches into the wavetable, resulting in corresponding timing glitches to the sound playback. This playback is then (when desired) overwritten into the audio file buffer. The parameters which control the glitches can be controlled to various degrees. Using a visual metaphor, this method can be compared to taking a photograph, cutting sections out (representing a "glitch") and rearranging them into new patterns.

As mentioned previously, this engine can be run on any standard desktop computer. With the correct settings a sound file on the computer's hard drive can be permanently modified. This process takes some time, in that even with the parameters at their maximum values the sound must still be looped and heard in order to overwrite each glitch to the source file. The destruction is only enacted when it is audible.

A consequence of the overwriting mechanism used in the engine is that the destroyed sound, when looped, will often fade to silence. This was unintentional and is likely due to phase differences in the overwritten audio, but the effect is aesthetically pleasing and thus was retained. When taken to this conclusion the digital file representing the sound recording will still exist (and, in fact, continue to be roughly the same filesize) but will be replaced by silence. Due to the method of overwriting and glitching it is impossible to recreate the original recording from the modified sound file.

A link to download and run this code is available in the supplementary material, as well as video files showing the sound destruction process.

4.2 The first physical version

The sound destruction engine was written with the primary intention of running on an embedded device in a piece of hardware. When the engine was therefore sufficiently completed to achieve this, a period of physical prototyping was undertaken. The first

[3] a text message from 2015: "*A random potentially cheesy idea for an artwork - you create something in a digital format (it could be a digital piece of music or a game or a video) and release it to the world, but it should have a built in ability programmed into it somehow to slowly decay the more it is used/listened to etc like a physical analogue piece of media. Maybe it's been done before or maybe it's cheesy but I quite like it as an idea*"

[4] https://puredata.info

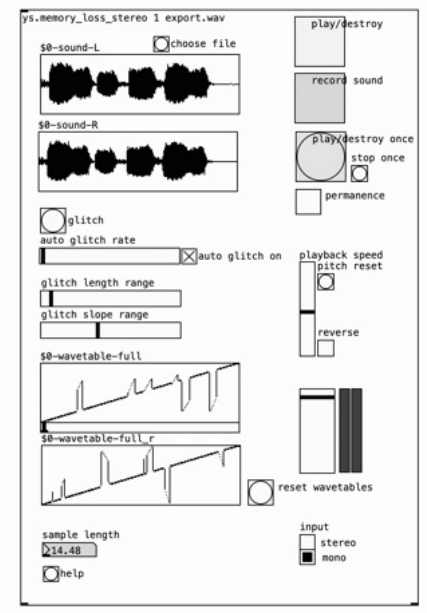


Figure 1: Screenshot of the Engine. The currently loaded sound waveform is visible at the top. The wavetables for playing each sound channel are below, and have been glitched approximately seven times.

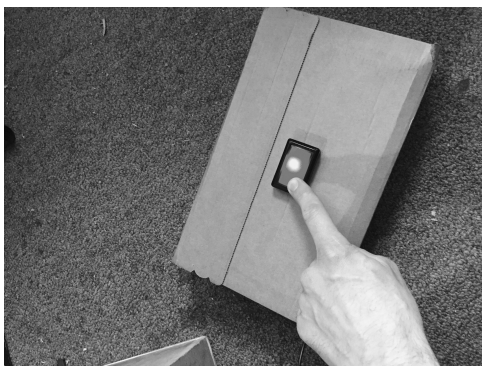


Figure 2: The cardboard destroyer



Figure 3: Screenshot from the video of the second prototype, housed in a cigar box. A link to the video is provided in the supplementary material.

version, using a Bela embedded computer running the Pure Data code, was built into a cardboard box with illuminated button (fig. 2). This was significant in retrospect mainly in terms of containing nearly all of the interface and design elements that were present in every future version of the device. These were:

- An enclosure containing an embedded computer
- A non-latching button
- A light acting as visual feedback
- A surface transducer speaker

Each of these elements were chosen in response to the initial design concept of linking listening and destruction.

4.2.1 An enclosure. Whilst the code functioned satisfyingly from the early stages, even generating surprisingly emotional outcomes in early non-permanent tests[5], in order for the destruction to be permanent it is more or less required for the device to be standalone, with the associated interface restrictions that go along with that. A (relatively) non-accessible system with a limited user interface element makes it far more likely for the destruction to be permanent, whereas a desktop computer offers many more options for copying or making backups. All of the versions that used an enclosure also used a surface transducer to generate the sound output (as opposed to a more standard speaker cone or headphone jack output). This was based primarily on a desire to limit the ability to re-record any output as a way of circumventing the intended permanence of the changes to the sound. A secondary outcome is that the enclosure, by vibrating to transmit the sound, effectively *becomes* the sound, reinforcing the concept of the original sound being "inside" the machine.

4.2.2 The button. Restricting the user interface, through the medium of a dedicated enclosure, necessitates the introduction of some form of interaction for playing back the sound. This could be accomplished by any number of inputs (switch, latching button, motion sensor, etc) but I settled on a momentary button. This was primarily for ease of use, as a single interaction could play, destroy, and stop the sound, because the act of pressing guarantees that the button will also be released. A momentary button press input could also be

[5]See "Destruction engine video demonstration" video in the supplementary material

fairly flexible, being used either as a press-and-hold or potentially as a press-once trigger.

4.2.3 The light. A single light was able to perform several functions from the outset, beginning with the most practical — indicating that the device was switched on and correctly running the code (not always a given in the early stages of an prototype embedded system). Eventually a brief flash was used to indicate that the playback had looped, which carried the second but perhaps more important implication that any glitches that had been introduced had also been overwritten onto the original file. The first cardboard prototype pictured here used an illuminated button for this light, while some subsequent versions had a separate light.

4.3 Second version

A second version, nearly identical to the first cardboard prototype, was built into a cigar box. This was primarily to improve the performance of the surface mount transducer, which was not effective on cardboard. This second version is notable because for the first time I named the device "Everything You Love Will One Day Be Taken From You" and, also for the first time, made a video showing it destroying a sound of my toddler laughing at my dog[6].

Attempting to create documentation of the new design made clear some of the challenges inherent in the work as a whole, in the sense that any recording would run counter to the intended permanence of the destruction. Making a video of the device in action would necessarily capture a version of the original sound before it had been glitched and permanently overwritten. In the case of the video that I ended up making, the sound of my toddler giggling at my dog is now preserved permanently in the video file (and on the servers of YouTube). This in effect devalues, or even completely negates, the destruction. This tension was left unresolved (indeed, the video remains online and is a useful way of explaining the work).

4.4 Everything You Love

The work on this device had thus far taken place within a larger project, The Book of Knowledge of Impractical Musical Devices [58], which involved developing a set of experimental sound devices that each explored a different design limitation. The final version of Everything You Love Will One Day Be Taken From You thus took the form of an instrument within this family, with a custom built wooden enclosure complete with design by artist Tommy Perman[7] (fig. 4). The main functional difference between this and the previous instances was the addition of a hidden audio input and button, allowing the user to record a sound directly into the device and thus avoid the danger of any copy of the sound remaining on an external recording machine or computer. This input and recording button was inside the box and required some knowledge to understand and access, effectively making it only usable by me.

I brought this version of the instrument back to my home shortly after finishing it, and set it up in my living room to take pictures. My toddler son showed some interest and asked if he could press the button (without remembering the word), saying "Can I press the beep?". I managed to record those words just before his pudgy little finger pressed down on the button labeled Play/Destroy (fig. 5), which immediately began looping and permanently glitching the recording we had just made.

[6]See supplementary material to view this video
[7]https://surfacepressure.net



Figure 4: The final version of Everything You Love

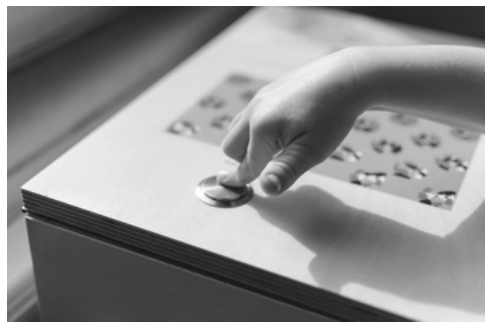


Figure 5: Playing and Destroying

4.5 A pause and a transition

Though not strictly a design instance, at this stage I would like to invoke Desjardins and Oogjes' concept of a "pause" and a "transition" [44] into the story of the design development. Shortly after completing this last version, just weeks after recording and beginning to destroy the sound in the final version of the instrument, I moved to a different country and literally left this work behind. Two years after that I moved to yet another country and eventually managed to get the Everything You Love instrument shipped to me at my new home. Seeing it for the first time in around four years, I was aware



Figure 6: The three Memory Boxes

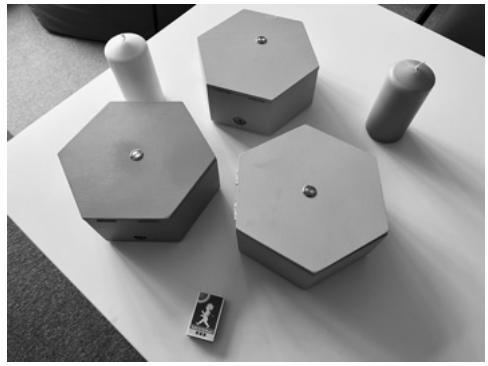


Figure 7: Getting ready for destruction

that in all likelihood the sound that my toddler son had recorded and begun to destroy was still "inside" the instrument, and I did not feel emotionally able to power it up and test it. I preferred to remember the moment as it was, and live with the knowledge that it still existed to a certain extent within the machine. This continues today - as I write this paper the device sits above my desk and the sound contained therein remains unplayed and undestroyed.

My transition from artist-teacher to academic research provided an opportunity to engage with sustainability studies, particularly with regards to sustainable computing and HCI. As my research focused on how sustainability could be considered in relation to media (and sound in particular), I once again reached towards the idea of destruction. Influenced by concepts of collapse computing [65], degrowth [28], hospicing [17], composting [64], and other frameworks for thinking about ways to counter the incessant and irrational desire for growth that is embedded in technology design, I revisited the Everything You Love project from a different angle.

4.6 The Memory Boxes

The Memory Boxes are a set of three hexagonal wooden enclosures with a single visible illuminated button. Their core functions are the same as the last Everything You Love device: the light indicates when a sound is ready to be heard and destroyed. Pressing and holding the button plays back the sound and simultaneously introduces permanent glitches which will eventually make the sound disappear. The sound is played back through a surface transducer, making the box itself vibrate. The Memory Box is distinguished from earlier iterations of the concept primarily through intention. This device was designed to be used not just be myself as the creator, but by anyone willing to engage. To enact this intention it was much simpler to build, had a more generic visual aesthetic, and the engine was modified slightly to enable a soundfile to be loaded into the box. This was still a process requiring some knowledge and in practice could only be performed by me.

These boxes were intended to be used to house a *meaningful* recording from the user's past, and for the sound contained therein to be the only copy remaining in the world. These two characteristics (meaningfulness and uniqueness) were conceptually linked. In order for the destruction to feel consequential it was crucial for the sound to be both important to the listener and not exist elsewhere. Being conceptually linked it is perhaps no surprise that these two issues were equally challenging to address in the design. How one defines any piece of media as "meaningful" is deeply unclear, and guaranteeing that it is the only remaining copy is virtually impossible due to the complicated and opaque nature of cloud backups and hard drive deletions. The design intention therefore necessitated the creation of a set of steps guiding the process of selecting and using the machines, functioning as both a technical and emotional mechanism for activating the device. These were designed on the assumption that several people would do this together. This full process was performed by the first two authors of this paper.

(1) Select a video from your personal media library.
(2) Show the video to the other participant(s), discussing why it is personally meaningful.
(3) Light a candle, representing the video
(4) Export the video from the media library
(5) Extract the audio portion from the digital video file
(6) Delete all copies of the video permanently from the disk and the cloud, leaving only the audio
(7) Connect the computer to the Box and load the audio file into the Box via USB
(8) Permanently delete the file from the computer, so the audio in the Box is the only remaining copy of the sound
(9) Blow out the candle

5 Annotations

5.1 Destroying Time

The method of destruction that has been developed for this work is inherently temporal. This is perhaps no surprise, since sound is completely reliant on (and therefore particularly illustrative of) time [15, 67], and as such listening back to a sound requires inhabiting the same amount of time as it took to record it. Although a library of photographs or even digital videos can be perceived at a glance, a hard drive full of sound recordings could take hours to listen to,

if not significantly more. It is, by its nature, a slow technology [42]. A recorded sound is effectively an attempt to force a wedge into the continuation of time, on the problematic assumption that we will one day be able to simply 'make' the time to listen back.

The designs in this portfolio weave this characteristic into the experience in several ways. For the physical devices, the sound itself is only heard while the button is held down, and thus requires active pressure from a listener in order to be heard. This forces the person pressing the button to physically embody the amount of time that passed in the original sound recording, resulting in a keen awareness of the amount of time that passed during that now-disappearing moment.

Crucial to this method of erasure, the destruction enacted by the Sound Destroyer Engine is also temporal — the glitches that are creating the destruction are not being applied directly to the sound but rather to the playback system which dictates the rate at which the audio file is heard (the wavetable). The digital buffer containing the audio file exists continually on the embedded device as a list of numbers that represent the position of the reproducing membrane that could move in order to generate sound. This data does not change and is not corrupted or destroyed if the engine is never engaged. Pressing play will insert glitches into the wavetable, and it is only when those glitches are heard and experienced that they are overwritten onto the sound. Destruction in this case requires time (listening to the sound as it is destroyed) and demands a deliberate choice (press play).

The standard HCI paradigm for making a piece of digital media disappear is deletion — usually through a two or three step process involving sending it to the "trash" of an operating system, which is then emptied after confirmation. All modern operating systems have failsafe mechanisms in place to guard this process; confirmation dialogues where the user is asked if they are certain the file should be deleted, or some form of deletion purgatory where a "deleted" file is kept for a further 30 days or similar. Indeed, in the process for placing a sound into the Memory Boxes the user will be confronted numerous times with confirmation dialog messages: "Are you sure you want to permanently erase the items in the Bin? You can't undo this action."

Despite the several step process required to delete, the standard method could be considered as atemporal, as the time spent on the deletion process is not related in any meaningful way to the time that could be spent experiencing the media. Furthermore, the time taken to 'delete' the file is not associated with the time taken to destroy it. In most cases what is actually being erased instantaneously is a reference to the data rather than the data itself, which would require a form of overwriting to fully erase. The actual erasure happens later, invisibly, as the operating system treats the previously used space as somewhere to write new data.

As this is a sound-based project it is also worth acknowledging that in most cases each step in the process of digital deletion in HCI is accompanied by an "earcon" [3] such as the sound of crumpling paper to represent a file being moved to the trash [21]. This use of earcons as a representation of the concept of deletion serves as an illustration of how the designs described here are the product of a different paradigm entirely. Our use of sound is not a method for delivering a message or confirmation of an action, but rather the sound engine itself *is* the deletion, a deletion that must be heard, experienced, and embodied in order to function.

In the designed experiences presented here, the slow destruction of a piece of media forces (or allows) the user to confront the fact of deletion. You can hear as the audio disappears, piece by piece, replaced eventually by the slowly fading hiss of white noise. This could be considered as an audible version of secure deletion, since the data is being actively overwritten. The destruction within this context is therefore a form of temporal erasure, in contrast to the immediate, seamless, but often illusory and consequence-free experience of digital deletion that is embedded into technology design. A temporal approach is particularly well suited to audio, and can be seen as amplifying the innate material connection between sound and time.

5.2 Impermanence

This presents us with a potential axis for considering our interaction with digital media: destruction vs loss. The generally accepted design paradigm of computing has biased strongly towards loss, with users accepting a certain degree of inevitability about data becoming eventually corrupted or inaccessible. Peeling CDRs and drawers full of SCSI drives are simply the reality of modern digital lives. The earlier devices described in this portfolio reflect this, most noticeably with the name "Everything You Love Will One Day Be Taken From You", which takes agency away from the listener or "owner" of the sound. You can press the button if you want to speed up the process, but it will be taken from you one day no matter what, one way or another. The design reflects the sense of loss that is embedded in the design and the use of media.

The other end of this axis points towards destruction. The Sound Destroyer engine that forms the core of these designs is an enactment of destruction as a deliberate and active process: the user can choose precisely how the sound will be destroyed (and when). The Memory Box design, having dropped the (admittedly morose and impractical) name shifted the focus to something broader and more accessible, moving the experience towards empowerment. It's up to you to put the sound in, it's up to you to delete the file, it's up to you to press the button. Either way you've already come to terms with the disappearance.

5.3 Storing/Remembering

The designs described here illustrate an uneasy relationship between storage and memory. In many cases digital storage can be seen as the opposite of human memory — it is a system that allows for our memory to be outsourced, especially with modern cloud storage that affords the potential of storing far more sounds (and videos and images) than we could ever hope to have the time to experience again. Their existence on the cloud is meant to be somehow relieve us of the fear of their disappearance, and perhaps by extension our own disappearance. The Memory Boxes in particular reveal this tension through the selection of a single sound from a single video in a vast sea of cloud media. This somewhat artificial elevation of an otherwise ordinary recorded moment provides an opportunity to engage with its instability and inevitable loss. It takes that sound and moves it away from our storage and towards

our memory (not entirely unlike the computing process of taking files and loading them into the RAM for display).

The physical manifestation of this in most of these designs included the "placing" of a sound in an enclosure, which then would become the source of the sound by vibrating in response to the surface transducer. This was enacted in two different ways, either by recording directly into the devices (as with the All the Things You Love version) or by loading a selected sound from a library and deleting the original (as with the Memory Boxes). While the first of these was an effective method for ensuring that the sound in the machine was the only recorded copy of that moment, it was also impractical and became almost paralyzing. Meanwhile the "final" version of All the Things You Love, a device which took over a year to design and build, only generated sound for perhaps a couple of minutes at most after the sound was recorded directly into the device. The feeling of finality, which underscored the fragile and fleeting nature of parenthood, meant that the Play/Destroy button has not been pressed since the day the sound contained within was recorded.

The Memory Box approach to filling the box with sound involved a process of selection or curation. Choosing a piece of media to destroy involved finding a video that was meaningful enough to be poignant, but not so valuable as to be painful. Ensuring that no other copies remained needed careful thought and consideration (had I messaged the video to my parents? Was it backed up to a secondary cloud service?). These requirements, which resulted in the pseudo-ritual that we developed to guide the selection of the sounds, reflected the complexity of our relationship with our own media as well as the entangled nature of the technology itself. The steps we followed were designed to ensure a careful but thorough experience, creating a space for discussing why we had selected the particular videos, how it might feel to lose them, and how we can be certain that we were fully enacting the design intention of permanent destruction.

Discussing this work with others it is not unusual for people to find the idea of deliberately and permanently destroying a cherished memory to be upsetting or offensive, and there is a certainly a level at which a piece of media becomes too important to lose. The first video of a newborn baby? The last recording of a cherished relative? An apparently unremarkable image from a walk on the beach which represents something entirely different for only one person? The process we developed for choosing a sound together went some way towards addressing this concept, though we stop short of presenting this as a solution. Rather, the takeaway is simply that the design required a selection process, which underscored the complicated relationship we have with storing and remembering.

5.4 Decay

Experientially, destruction is already built into sound: by the time you hear *any* sound the event that generated the sonic energy has already passed. In our everyday life this is mostly imperceptible, with sound illustrating the passing of time on a micro scale when listening to something nearby, or an unimaginably macro scale when listening to galaxies collapsing [10]. It is perhaps most jarring somewhere in the middle, such as hearing a sound from some distance — the delay between seeing movement across a field and hearing the clanging of a hammer making you suddenly aware that you are situated in time and space. The terminology used in the audio domain reflects this destructive characteristic; sound *decays*. In acoustics terms the decay refers to the amount of time it takes for the audible remnants of a sound to disappear as the agitations in the listening medium (usually air) bounce around a space[8]. Sound is by its very nature a self-destroying medium. The energy that becomes audible is impossible to sustain and must surrender to entropy.

Digital audio, however, prioritizes reproducibility and the illusion of permanence: the launch of the compact disc in 1985 was accompanied by the slogan "Pure Perfect Sound, Forever"[9]. Audio technology promises not only a faithful reproduction of sound (a concept with curious gender implications [53]) but also consistent repetition. Each time you press "play" you will hear the same sound, it won't have changed, your audio and your memories remain intact. As with all computing, there is an embedded expectation for sonic technology to be endless, seamless, and neutral, even if sound itself is a medium that is impermanent and self-destructive by nature. Just as the HCI heuristic of reversibility clashes with our lived experience [54], recorded sound is meant to represent an opportunity to revisit the past.

This seamlessness is challenged by the approach presented here, particularly with All the Things You Love, which functions as both a recording and playback machine, albeit one which overtly promises that the the moment you record will never again be perfectly reproduced. This re-aligns our expectations of sound recording more generally, reintroducing the seam of decay into our experience of listening to a recording. This seam was especially noticeable in the challenges that arose with the creation and documentation of the second version in the cigar box. Documenting the prototype at all necessitated the use of recording tools that were dependent on the paradigm of endless reproduction, which worked in opposition to the ideals of the work.

6 Discussion

Following the annotations above, this section discusses what this work means for HCI in terms of offering solutions (or a lack thereof), how these devices can be seen as part of a larger artistic context, how the temporal aspects compares to earlier work, and how this body of work centers destruction. These discussions points each suggest a seamful [13] approach to addressing the notions of cornucopia, time, and our relationship to the impermanent nature of media.

6.1 This is not a solution

The background section of this paper opened with a discussion of the inherent instability of storage media. This reality is formed of two interwoven elements: the guaranteed failure of all digital recording mediums [5], and the embrace of a *new* storage medium which promises to solve the problem — this time forever! This cycle forms an integral part of the computing technology ecosystem.

The cornucopian design principles [50] that contribute to this pattern lead to unsustainable promises of endless growth, while

[8]Interestingly, "decay" is also a common term in audio technology but with the slightly different usage, most often as the "D" in a standard "ADSR" envelope, referring to the amount of time it takes for the envelope to reach a static "sustain" level.
[9]https://www.youtube.com/watch?v=GaPlDhR-i_I

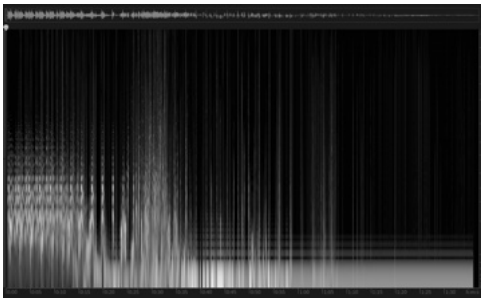


Figure 8: A sound disappears: A spectrograph showing a sound played on loop through the sound destroying engine. The image represents approximately one and a half minutes of the sound looping. The glitches from the engine initially add more spectral activity before eventually the sound fades to noise. A link to video of this spectrograph recording is included as supplementary material.

denying the reality of the instability and impermanence of computing. The design bias against deletion inevitably leads to situations where users are forced to suddenly delete files under pressure or lose them entirely due to unforeseen circumstances.

The sound destroyers in this portfolio do not offer a realistic or practical alternative to this. The act of deleting one file, over the course of a longer period of time, cannot diminish the media hoards, let alone outpace the capture of new materials to add it to. The design work underpinning the sound destroyers instead is an inquiry into how we might peel back the assumptions of permanence and reveal the depth of meaning in making (and destroying) of our media archives. As annotated above, the devices offer an embrace of deletion through a seamful and embodied time-based approach. In doing so they can illustrate what is missing in the standard design paradigms and prefigure alternative (if improbable, impractical, or "functionally oppositional" as suggested by Pierce [47, 48]) methods for acknowledging and working with the qualities of loss already built into our media technology.

6.2 Artistic context

The overtly non-solutionist attitude of this work suggests a relationship with the large body of destruction-based sound, music, and art. To a certain extent the sound destruction engine is an application of Basinki or Lucier's ideas to the digital domain, relying on iterative repetition and overwriting to reveal the characteristics of the medium. Although the Disintegration Loops or I'm Sitting in a Room could rely on the nature of tape and resonance to generate the decay, digital sound is a medium that by nature can be copied and repeated virtually endlessly. Thus the sound destruction engine actively inserts glitches in order to provide the initial impulse towards decay, which is then reinforced through repetition and overwriting (as can be seen in fig. 8).

The insertion of the glitches, of course, requires a user (in HCI parlance) or a player (in musical parlance). Rather than being a piece of static media that can be started and experienced without further interaction, these devices build on the model of an instrument that is capable of continuous (albeit ephemeral) use. The listener *plays* the instrument by pressing and holding the button for as long as they want to hear it, at the cost of the steady disappearance of the embedded sound.

The finality that is implied in many of the instances in the portfolio also invites comparisons to the contemporary art examples of destruction. While the Disintegration Loops is examining the sonic textures that can be generated from repetition and decay, the finality is limited — when you reach the end of the recording, you can listen to it again from the beginning and remind yourself what it sounded like at the start.

In contrast, Yoko Ono's clothing could not be un-removed at the end of a performance of Cut Piece, and the KLF could not un-burn their cash. The tension I encountered in the documentation of the project reflected these two different approaches. Recording a video of the sound being destroyed in the second version aligned the work more closely to sound art, while the later version could not be documented in the same way without compromising the core concept, and thus became more aligned with destructive practice in contemporary art.

A primary distinction with regards to the work in this portfolio is how it emphasizes a embodied form of destruction that is intertwined within the act of listening. Pressing and holding the button to both play and destroy the sound contained in the enclosure, whilst hearing the decay in real time, is a particularly personal and physical act.

6.3 Neither slow nor fast

The background and annotations above describe both the impermanence of digital storage as well as the inevitable decay of acoustic sound. These two very different perceptual experiences also operate on two different timescales — while a sound decays nearly instantaneously, digital files (despite their instability) could last for years or decades.

The designed experience of destruction offers a third, intermediate timescale. The artifacts presented in this portfolio can be seen as simultaneously a method of speeding up an eventual inevitable loss of media, allowing a listener to hear a digital sound disappear in a matter of minutes rather than years, and a method for extending the decay of sound, allowing a listener to stretch out an otherwise inevitable loss. A potential parallel can be made to sonification practice which often involves a significant manipulation of the time domain (transposing the duration of radio waves from the sun by a factor of 42,000 [51], playing two years of Covid data in 2 minutes [39], etc).

This intermediate timescale is further expanded by the physical interaction involved in all of the designs discussed here, in which the user is embodying the playback and the destruction. Reflecting Chion's discussions of listening as an act that illustrates time [15], holding down the button to play back a sound (and enact the destruction) is a method of relating to the specific window of time in which the original sound occurred. Your finger must be held on

the button for precisely the time that it took to record the original sound: you are physically enacting the listening and the destruction in "real time" in two senses of the phrase - the destruction happens while you are pressing, and the time spent listening is a simulacrum of the original moment.

6.4 Making destruction matter

There is an undeniable contradiction built into the conceptual premise of this work. If the design process here stemmed from notions of countering the bias towards endless growth in computing, what can be gained from offering a new set of objects built from digital technologies?

The designs in this portfolio are all manifesting destruction. They were built in order to destroy. However the destruction that they can inflict is in many ways potential, rather than enacted. The media is not destroyed (or indeed, changed at all) until a button is pressed by an active listener/destroyer, and as in the case of the final All the Things You Love instrument, the sound can remain untouched, unheard, and unharmed for an undetermined but likely extended period of time; going back to the original motivation for the project, the recording of my then-two-year-old son in that particular artifact is like a painting stored in a pitch-black cave, retaining its color until it is perceived.

Nevertheless these designs are all primed for destruction. But how exactly is that destruction taking place? While a flashing light indicates that the destruction has been overwritten to disk, the actual process is an activation of the real-time playback of a wavetable controlling the time domain of the digital sound file. The light is a representation of the process, analogous to the familiar crumpling paper earcon, while the actual destruction is being enacted by the material of sound itself. There is no conversion, no metaphor. The process *is* the sound, and the sound *is* the process: they become a single technical and perceptual event. Where the Sound Destroyer Engine makes this visible in the user interface of the code, the later embedded devices physically vibrate using the integrated surface transducer to emphasize this.

Furthermore, considering how the wavetable overwriting sound destruction engine is comparable to a photograph being cut into pieces and reassembled, the destroyed sound in the engine is not taking up materially less space. The buffer is precisely the same size throughout the entire destruction process, just as the same amount of paper remains from a cut-up photograph as existed in the undisturbed picture. What is being destroyed is rather the meaning of the sound, the particular combination of data points that (when played back very fast and amplified) once resulted in a recognizable baby's laugh or dog's bark. While standard approaches to deletion rely on the illusion of freeing up material space, this method of destruction focuses instead of the experience of loss, reflecting the meaning contained in the media.

This can be contextualized through Pierce's discussion of *undesigning* [46] or negation of technology by design. While we accept that this portfolio is responding to a surfeit of media with yet more technology, the work is not merely meant to propose an alternative method for shrinking our media footprint (indeed, that was neither an intention or an outcome of the project). Rather, the development of a method for combining the perception of destruction and the destruction itself represents an intentional negation of both technology and interaction with media. This negation is a reflection of the alternative design paradigm that we propose, prefiguring forms that an anti-growth computing ecology could take.

7 Conclusion

As I write this section an email notification arrives in my inbox, telling me that my cloud storage is full. Our digital lives are currently full of these moments where the responsibility of limiting our resource use is suddenly shifted to us personally. Our technology encourages us to create and share media without limits, but then punishes us when we do so. What content will I delete to free up space? How will it feel? The tools at my disposal do not enable this in any way that reflects the meaning of the content I must delete. And neither does the act of deleting them reflect or illustrate the actual act of deletion.

This work explored destruction as a design space, focusing on digital sound as a form of personal media. The annotations we developed analyzed aspects of time, impermanence, storage, memory, and decay as properties that emerged from this exploration. The annotations and discussion are generated from a seamful approach to the design of digital media, which engages with the limitations and impermanence inherent in the technology. This embrace then opens up a new set of possibilities for interaction with media archives. Our work has particularly explored the deliberate destruction of sound as one of these possibilities, suggesting the potential for more methods that can expand this space.

With this in mind, the portfolio itself is the first contribution of this paper, offering a set of designs that have iteratively examined how destruction can be implemented in digital sound. This in turn revealed the inherent impermanence in digital media, a quality that HCI is both fundamentally biased against and tends to attempt to mitigate through cornucopian and seamless approaches to design. The work presented here points towards an embrace of impermanence as a design strategy. Finally, as an overtly anti-solutionist approach, the work is necessarily partial and provisional, and we call on HCI to further examine how deliberate destruction can enable deeper meaning making with digital media. Where many previous alternative designs for media technology have proposed methods for revisiting forgotten or overlooked pieces of personal media, a destruction-led paradigm of design can embrace the inevitability of digital decay.

Acknowledgments

This project was funded by the Swedish Energy Agency (*Energimyndigheten*), award number P2020-90326.

References

[1] Fatemeh Bakhshoudeh and Rob Comber. 2025. Designing with the Solar Internet: Towards Constraint-Based Design for Sustainable Consumption. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–18. doi:10.1145/3706598.3713101

[2] Lindsay Balfour. 2018. Traumatic Ruins and the Archeology of Sound: William Basinski's The Disintegration Loops. *Journal of Sonic Studies* 16 (2018). https://pureportal.coventry.ac.uk/en/publications/traumatic-ruins-and-the-archeology-of-sound-william-basinskis-the

[3] Meera Blattner, Denise Sumikawa, and Robert Greenberg. 1989. Earcons and Icons: Their Structure and Common Design Principles. *Human-Computer Interaction* 4, 1 (March 1989), 11–44. doi:10.1207/s15327051hci0401_1

[4] Eli Blevis. 2007. Sustainable interaction design: invention & disposal, renewal & reuse. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 503–512. doi:10.1145/1240624.1240705

[5] Raymond Blum and Betsy Beyer. 2019. Achieving digital permanence. *Commun. ACM* 62, 5 (April 2019), 36–42. doi:10.1145/3303872

[6] John Bowers. 2012. The logic of annotated portfolios: communicating the value of 'research through design'. (2012).

[7] John Bowers and Owen Green. 2018. All the noises: Hijacking listening machines for performative research. In *New Interfaces For Musical Expression*. Virginia Tech, 114–119.

[8] John M Bowers, John Richards, Tim Shaw, Robin Foster, and Akihiro Kubota. 2023. Raw Data, Rough Mix: Towards an Integrated Practice of Making, Performance and Pedagogy. (May 2023). doi:10.5281/ZENODO.11189230 Publisher: Zenodo.

[9] Julia Bryan-Wilson. 2003. Remembering Yoko Ono's "Cut Piece". *Oxford Art Journal* 26, 1 (2003), 99–123. http://www.jstor.org/stable/3600448

[10] Davide Castelvecchi and Alexandra Witze. 2016. Einstein's gravitational waves found at last. *Nature* (Feb. 2016), nature.2016.19361. doi:10.1038/nature.2016.19361

[11] Matthew Chalmers. 2003. Seamful design and ubicomp infrastructure. (2003), 577–584.

[12] Matthew Chalmers and Areti Galani. 2004. Seamful interweaving: heterogeneity in the theory and design of interactive systems. In *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques*. ACM, Cambridge MA USA, 243–252. doi:10.1145/1013115.1013149

[13] Matthew Chalmers and Ian MacColl. 2003. Seamful and seamless design in ubiquitous computing. In *Workshop at the crossroads: The interaction of HCI and systems issues in UbiComp*, Vol. 8. 10.

[14] Amy Yo Sue Chen, William Odom, Carman Neustaedter, Ce Zhong, and Henry Lin. 2023. Exploring Memory-Oriented Interactions with Digital Photos In and Across Time: A Field Study of Chronoscope. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg Germany, 2023-04-19). ACM, 1–20. doi:10.1145/3544548.3581012

[15] Michel Chion. 1994. *Audio-vision: sound on screen*. Columbia University Press, New York.

[16] Miyuru Dayarathna, Yonggang Wen, and Rui Fan. 2016. Data Center Energy Consumption Modeling: A Survey. *IEEE Communications Surveys & Tutorials* 18, 1 (2016), 732–794. doi:10.1109/COMST.2015.2481183

[17] Elina Eriksson and Daniel Pargman. 2023. Exploring Inner Transition: Expanding Computing forSustainability. In *Ninth Computing within Limits 2023*. LIMITS, Virtual. doi:10.21428/bf6fb269.22dea4ad

[18] Katie Gach. 2019. A Case for Reimagining the UX of Post-Mortem Account Deletion on Social Media. (2019). doi:10.13140/RG.2.2.29754.75205 Publisher: Unpublished.

[19] Simson Garfinkel. 2024. Complete Delete: In Practice, Clicking 'Delete' Rarely Deletes. Should it? *MIT Case Studies in Social and Ethical Responsibilities of Computing* (Sept. 2024). doi:10.21428/2c646de5.26e95ae4

[20] Bill Gaver and John Bowers. 2012. Annotated portfolios. *Interactions* 19, 4 (July 2012), 40–49. doi:10.1145/2212877.2212889

[21] William Gaver. 1986. Auditory Icons: Using Sound in Computer Interfaces. *Human-Computer Interaction* 2, 2 (June 1986), 167–177. doi:10.1207/s15327051hci0202_3

[22] William Gaver. 2012. What should we expect from research through design?. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Austin Texas USA, 937–946. doi:10.1145/2207676.2208538

[23] William Gaver, Peter Gall Krogh, Andy Boucher, and David Chatting. 2022. Emergence as a Feature of Practice-based Design Research. In *Designing Interactive Systems Conference*. ACM, Virtual Event Australia, 517–526. doi:10.1145/3532106.3533524

[24] Cathal Gurrin, Gareth J. F. Jones, Hyowon Lee, Neil O'Hare, Alan F. Smeaton, and Noel Murphy. 2005. Mobile access to personal digital photograph archives. In *Proceedings of the 7th international conference on Human computer interaction with mobile devices & services* (Salzburg Austria, 2005-09-19). ACM, 311–314. doi:10.1145/1085777.1085842

[25] Lars Hallnäs and Johan Redström. 2001. Slow Technology – Designing for Reflection. 5, 3 (2001), 201–212. doi:10.1007/PL00000019

[26] Stina Hasse. 2012. I am Sitting in a Room. *Body, Space & Technology* 11, 0 (Jan. 2012). doi:10.16995/bst.71

[27] Sjoerd Hendriks, Mafalda Gamboa, and Mohammad Obaid. 2024. The Undertable: A Design Remake of the Mediated Body. In *Designing Interactive Systems Conference* (IT University of Copenhagen Denmark, 2024-07). ACM, 2591–2610. doi:10.1145/3643834.3660698

[28] Jason Hickel. 2022. *Less is more: how degrowth will save the world*. Penguin Books, UK USA Canada Ireland Australia India New Zealand South Africa.

[29] John Higgs. 2013. *The KLF: chaos, magic and the band who burned a million pounds* (paperback edition ed.). Phoenix, London. OCLC: 853507033.

[30] Lorenz M. Hilty and Bernard Aebischer. 2015. ICT for Sustainability: An Emerging Research Field. In *ICT Innovations for Sustainability*, Lorenz M. Hilty and Bernard Aebischer (Eds.). Vol. 310. Springer International Publishing, Cham, 3–36. doi:10.1007/978-3-319-09228-7_1 Series Title: Advances in Intelligent Systems and Computing.

[31] Esther Jang, Matthew Johnson, Edward Burnell, and Kurtis Heimerl. 2017. Unplanned Obsolescence: Hardware and Software After Collapse. In *Proceedings of the 2017 Workshop on Computing Within Limits*. ACM, Santa Barbara California USA, 93–101. doi:10.1145/3080556.3080566

[32] Daniel Kahneman and Amos Tversky. 1979. Prospect Theory: An Analysis of Decision under Risk. 47, 2 (1979), 263. doi:10.2307/1914185

[33] Joseph 'Jofish' Kaye, Janet Vertesi, Shari Avery, Allan Dafoe, Shay David, Lisa Onaga, Ivan Rosero, and Trevor Pinch. 2006. To have and to hold: exploring the personal archive. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Montréal Québec Canada, 2006-04-22). ACM, 275–284. doi:10.1145/1124772.1124814

[34] Bran Knowles, Oliver Bates, and Maria Håkansson. 2018. This Changes Sustainable HCI. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–12. doi:10.1145/3173574.3174045

[35] Jonas Löwgren. 2013. Annotated portfolios and other forms of intermediate-level knowledge. *Interactions* 20, 1 (Jan. 2013), 30–34. doi:10.1145/2405716.2405725

[36] Christian Marclay. 1985. Record without a Cover.

[37] Greg Milner. 2010. *Perfecting Sound Forever - An Aural History of Recorded Music*. fsgbooks, New York.

[38] 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. doi:10.1145/3183582

[39] Tri-Long Nguyen and Mads Kamper-Jørgensen. 2023. The sound of waves: tracking the pandemic. *The Lancet Infectious Diseases* 23, 1 (Jan. 2023), 43. doi:10.1016/S1473-3099(22)00791-5

[40] Nathan Novemsky and Daniel Kahneman. 2005. The Boundaries of Loss Aversion. 42, 2 (2005), 119–128. doi:10.1509/jmkr.42.2.119.62292

[41] William Odom and Tijs Duel. 2018. On the Design of OLO Radio: Investigating Metadata as a Design Material. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC Canada, 2018-04-19). ACM, 1–9. doi:10.1145/3173574.3173678

[42] William Odom, Mark Selby, Abigail Sellen, David Kirk, Richard Banks, and Tim Regan. 2012. Photobox: on the design of a slow technology. In *Proceedings of the Designing Interactive Systems Conference* (Newcastle Upon Tyne United Kingdom, 2012-06-11). ACM, 665–668. doi:10.1145/2317956.2318055

[43] William Odom, Ron Wakkary, Jeroen Hol, Bram Naus, Pepijn Verburg, Tal Amram, and Amy Yo Sue Chen. 2019. Investigating Slowness as a Frame to Design Longer-Term Experiences with Personal Data: A Field Study of Olly. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, Glasgow Scotland Uk, 1–16. doi:10.1145/3290605.3300264

[44] Doenja Oogjes and Audrey Desjardins. 2024. A temporal vocabulary of Design Events for Research through Design. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–12. doi:10.1145/3613904.3642560

[45] Daniel Pargman and Barath Raghavan. 2014. Rethinking sustainability in computing: from buzzword to non-negotiable limits. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational*. ACM, Helsinki Finland, 638–647. doi:10.1145/2639189.2639228

[46] James Pierce. 2012. Undesigning technology: considering the negation of design by design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Austin Texas USA, 957–966. doi:10.1145/2207676.2208540

[47] James Pierce. 2021. In Tension with Progression: Grasping the Frictional Tendencies of Speculative, Critical, and other Alternative Designs. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–19. doi:10.1145/3411764.3445406

[48] James Pierce and Eric Paulos. 2014. Counterfunctional things: exploring possibilities in designing digital limitations. In *Proceedings of the 2014 conference on Designing interactive systems* (Vancouver BC Canada, 2014-06-21). ACM. doi:10.1145/2598510.2598522

[49] James Pierce and Eric Paulos. 2015. Making Multiple Uses of the Obscura 1C Digital Camera: Reflecting on the Design, Production, Packaging and Distribution of a Counterfunctional Device. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, Seoul Republic of Korea, 2103–2112. doi:10.1145/2702123.2702405

[50] 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 *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 1324–1337. doi:10.1145/2858036.2858378

[51] Michael Quinton, Iain McGregor, and David Benyon. 2016. Sonifying the Solar System. In *Proceedings of the 22nd International Conference on Auditory Display - ICAD 2016*. The International Community for Auditory Display, Canberra, Australia, 28–35. doi:10.21785/icad2016.003

[52] Kopo Marvin Ramokapane and Awais Rashid. 2023. ExD: Explainable Deletion. In *New Security Paradigms Workshop*. ACM, Segovia Spain, 34–47. doi:10.1145/3633500.3633503

[53] Tara Rodgers. 2010. *Pink noises: women on electronic music and sound.* Duke University Press, Durham [NC].

[54] Beat Rossmy, Nađa Terzimehić, Tanja Döring, Daniel Buschek, and Alexander Wiethoff. 2023. Point of no Undo: Irreversible Interactions as a Design Strategy. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–18. doi:10.1145/3544548.3581433

[55] Corina Sas and Steve Whittaker. 2013. Design for forgetting: disposing of digital possessions after a breakup. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Paris France, 1823–1832. doi:10.1145/2470654.2466241

[56] Corina Sas, Steve Whittaker, and John Zimmerman. 2016. Design for Rituals of Letting Go: An Embodiment Perspective on Disposal Practices Informed by Grief Therapy. *ACM Transactions on Computer-Human Interaction* 23, 4 (Sept. 2016), 1–37. doi:10.1145/2926714

[57] Brennan Schaffner, Neha A. Lingareddy, and Marshini Chetty. 2022. Understanding Account Deletion and Relevant Dark Patterns on Social Media. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (Nov. 2022), 1–43. doi:10.1145/3555142

[58] Yann Seznec. 2025. The Book of Knowledge of Impractical Musical Devices. (July 2025). doi:10.5281/ZENODO.15837664 Publisher: Institute of Electronic Music and Acoustics, University of Music and Performing Arts Graz.

[59] Vishal Sharma, Neha Kumar, and Bonnie Nardi. 2024. Post-growth Human–Computer Interaction. *ACM Transactions on Computer-Human Interaction* 31, 1 (Feb. 2024), 1–37. doi:10.1145/3624981

[60] Vygandas 'Vegas Simbelis, Pedro Ferreira, Elsa Vaara, Jarmo Laaksolahti, and Kristina Höök. 2016. Repurposing Bits and Pieces of the Digital. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 840–851. doi:10.1145/2858036.2858297

[61] Martin Stojanov, Daniel Pargman, Mike Hazas, Rob Comber, and Jorge Luis Zapico. 2023. How do we arrive at constraints? Articulating limits for computing. In *Ninth Computing within Limits 2023*. LIMITS, Virtual. doi:10.21428/bf6fb269.a317d18f

[62] Anna Ståhl, Madeline Balaam, Marianela Ciolfi Felice, and Irene Kaklopoulou. 2022. An Annotated Soma Design Process of the Pelvic Chair. In *Designing Interactive Systems Conference*. ACM, Virtual Event Australia, 1921–1933. doi:10.1145/3532106.3533469

[63] Alex S. Taylor and Jane Vincent. 2005. An SMS History. In *Mobile World: Past, Present and Future*, Lynne Hamill, Amparo Lasen, and Dan Diaper (Eds.). Springer London, London, 75–91. doi:10.1007/1-84628-204-7

[64] Minna Laurell Thorslund, Oliver Bates, Elina Eriksson, Daniel Pargman, Aksel Biørn-Hansen, Fatemeh Bakhshoudeh, and Arjun Rajendran Menon. 2025. Metacrisis computing and you: Finding agency through the Two Loops model of change. (2025).

[65] Bill Tomlinson, Eli Blevis, Bonnie Nardi, Donald Patterson, M Six Silberman, and panyue@cn.ibm.com Pan. 2013. Collapse informatics and practice. *ACM Transactions on Computer-Human Interaction* 20 (Sept. 2013), 1–26. doi:10.1145/2509404.2493431

[66] Laia Turmo Vidal, Yinchu Li, Martin Stojanov, Karin B Johansson, Beatrice Tylstedt, and Lina Eklund. 2023. Towards Advancing Body Maps as Research Tool for Interaction Design. In *Proceedings of the Seventeenth International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, Warsaw Poland, 1–14. doi:10.1145/3569009.3573838

[67] MinYoung Yoo, William Odom, Arne Berger, Samuel Barnett, Sadhbh Kenny, Priscilla Lo, Samein Shamsher, Gillian Russell, and Lauren Knight. 2024. Remembering through Sound: Co-creating Sound-based Mementos together with People with Blindness. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–19. doi:10.1145/3613904.3641940

Paper E

All of Your Sounds

YANN SEZNEC, KTH, Sweden

All of Your Sounds is a designed sound object that generates a virtually infinite soundscape from a user's personal cloud video library. Thousands of sounds from the past are interwoven, creating an ever-changing ambient portrait of a single person's life through sound. The design and realization of this device represents a conversation with a material formed from the relationship between sound, personal media, and the listener. Designing in response to this relationship leads to three key understandings: the use of sound as an alternative method of engaging with personal media, an awareness of the media ecologies that we form a part of, and the benefits of an overtly anti-solutionist approach to designing personal media devices.

CCS Concepts: • **Human-centered computing → Interaction design theory, concepts and paradigms**.

Additional Key Words and Phrases: research through design, sound, media

1 Introduction

I have 1,293 videos in my cloud library. The earliest video dates to January 2012, and the most recent was taken this morning. If I close my eyes and scroll randomly around the library and pick a video (not a particularly simple maneuver) I can stumble across a random moment from my past, maybe something prosaic like testing a cassette player in my studio or perhaps something more startling, like a three second video taken with my wife, sitting on a park bench in Edinburgh in 2016, holding our infant son. He seems vaguely confused, I giggle and kiss his forehead, my wife gently tucks her hair back, and the video suddenly ends. Before writing this introduction I had never watched this video, or at least I had no memory of it, though I certainly remember the day itself. We had spent the afternoon in Holyrood Park taking pictures to introduce the newest member of the family to the world. I must have given my phone to my sister and the camera app presumably switched to video mode by accident, explaining the brevity and stark ending. It's a video that frankly shouldn't exist, but for some reason it has survived in my library for nine years. How many other videos like that do I have?

A video library with over a thousand items is overwhelming. It is virtually guaranteed that I will never be able to watch them all, and their very existence raises challenging questions about how we deal with permanence, storage, media, and memory. While this paper will not propose a solution, it will begin with a particular quality of these videos that is generally overlooked – their sound. Reframing these videos as personal sound libraries enables the design exploration of a media device for listening, rather than viewing, taking advantage of the particular qualities of both personal media libraries and sound more generally. The process of designing within these boundaries represents an oblique approach to unfolding deeper questions around our relationship to personal media, storage, and technology.

This paper presents a device to examine these issues: All of Your Sounds (AoYS). The design will be introduced and then analyzed in terms of how it engages with sound, personal media, and the listener. I will argue that the combination of qualities across these three result in a design material leading to several key understandings: the use of sound as an alternative method of engaging with personal media, an awareness of the media ecologies that we form a part of, and the benefits of an overtly anti-solutionist approach to designing personal media devices.

Author's Contact Information: Yann Seznec, KTH, Stockholm, Sweden, yannse@kth.se.

2 Background, Theory, Method

The design described in this paper can be linked to slow design as developed by Odom [16], Hallnäs and Redström [8], and others. This paradigm places emphasis on a deliberate and intentional slowing down of the experience of interacting with a technology to encourage reflection and understanding. The objects that apply this philosophy, such as the Drift Table [6] and the Tilting Bowl [27] generate a slowly unfolding experience that encourages questions and conversations rather than attempting to provide answers. A number of slow design projects have particularly engaged with large media libraries. Photobox [15], Olo Radio [14], Olly [17] and others have looked to engage with the meanings and identities embedded into personal media such as photographs and music listening history. Those designs have encouraged thoughtful and contemplative reimaginings of the source material, offering different relationships to both the technology and the original media.

This anti-solutionist attitude is mirrored in the methodology for my work, which took a Research through Design [20] approach focusing on the act of making as a method for generating knowledge. In particular I draw on Pierce's definition of alternative designs [19] which "prefigure possible, though improbable productions" and resist further development, and Gaver's definition of emergence as a core principle of practice-led design knowledge [5]. The All of Your Sounds device is not intended to be a commercial product, or even a prototype for a future production, but is rather a functioning object that came into existence precisely because it stemmed from a design process that assumed that it would not need to be developed into a "finished" thing, similar to Odom's concept of a "research product" [18]. This methodology extends to the analysis, which focuses on my own design and experience of using the device. To present the work I will borrow certain qualities from annotated portfolios [2]: the design will be described and the annotations will be presented and discussed as "reflections".

An additional element of theory that this work engages with is Schön's positioning of design as a conversation with material [24]. This overlaps significantly with Gaver, Bowers, and Pierce's positions, of course, but I would in particular like to propose a sound-focused version of this concept, which is notably biased towards visual design in Schön's original paper. Although sound itself is occasionally discussed as a design material [4], I will instead propose a particular material that is built through the design of AoYS, with which I have had a design conversation.

While this work can be broadly situated within the area of personal media, the artefact analyzed here is a sound making and listening machine. As such it draws significantly from the areas of field recording, generative composition, and sound studies. The device I have built relies on what are essentially field recordings as generative content, creating parallels with the musique concrète practices of composers such as Schaefer [22] and Pade [1], among many others. The use of personal sounds in particular aligns strongly with Matthew Herbert's concept of the "audible subject" (Herbert, 2020), considering how the context of the source sound is crucial to the relationships that are built between listener, performer, and instrument. The way we listen to these sounds can be analyzed in a number of ways, including using Chion's "listening modes" [3].

There is considerable overlap between these ideas and the idea of "soundscape", an admittedly a problematic and contested term which is often used to describe creative sound or music made from field recordings such as the work of Hildegard Westerkamp [28], Luc Ferrari [21], and others. "Soundscape" is also commonly used in a broad sense meaning something akin to "ambient sound" as well as having a very particular definition famously introduced by Schafer [23], which focuses on a set of sounds as perceived by humans. It can also be used to describe the acoustic environment more generally [13], including sounds that may not be readily perceived by the naked ear, or alternatively as a "process

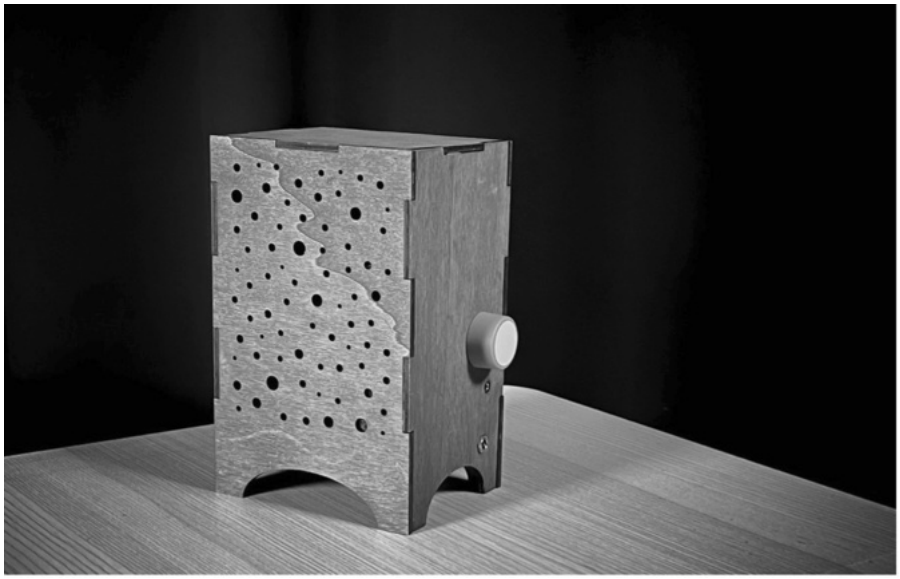


Fig. 1. The All of Your Sounds device, latest iteration

of sonic placemaking" [26]. It is important to acknowledge the significant objections to the term, which point out its problematic colonial connotations [10] as well as how it can ignore the embodied and situated experience of sound [9].

The following work thus takes a Research through Design approach to exploring our personal sounds, and how they can create a material that enables a deeper understanding of our media ecology.

3 The Design

All of Your Sounds (AoYS) is a personal listening device (figure 1). It contains the sounds of every video in a user's cloud library - in my case my AoYS machine contains around 1200 sounds from video files dating back over 12 years. The device takes this personal library and layers up to four sounds at a time, selecting excerpts from each file and fading them in and out in a non-repeating and constantly shifting generative soundscape. A volume knob on the enclosure enables a listener to "tune in" to this generative composition, and indeed to tune out when desired.

3.1 Making the Sound Library

The first step in the design process for AoYS is the creation of a library of personal sounds. This can be a complicated undertaking, as it is not unusual to have several thousand video files in a personal cloud media library. Most commercial cloud library management systems do not offer a method for easily accessing sound from the video files. The most straightforward approach to do so involves using a dedicated media library app (such as Photos for macOS), selecting all the videos in the library and dragging them to a location on the local drive. However, it is likely that most of the video

files will not be on the local drive to begin with and must first be downloaded from the cloud server. When multiplied 1000 times or more the result is a massive slowdown of the operating system. Considering the sheer quantity of media, often in the order of gigabytes, modern computers will catch up with this fairly quickly, perhaps 30 seconds or so, but the user attempting to complete this task must have the awareness and patience to wait for the software to indicate that it is ready to copy the files before releasing the mouse.

With the video files copied, the user must then run a batch processing system for stripping the audio from each video, normalizing the volumes to account for differences across files, and converting it to a suitable format for use in the audio engine. Finally, for use in the AoYS device these resulting files need to be transferred to the Raspberry Pi running the engine via USB or SSH. Though these are not particularly complicated processes, neither are they familiar to average computer users.

3.2 Technical Description

The device runs on a Raspberry Pi, with a sound engine built in Pure Data. A custom built PCB is attached to the Raspberry Pi with a power socket and a small speaker driven by an audio amplifier and volume knob. The sound library is located on the local storage of the Pi.

The Pure Data code is comprised of four instances of a single patch, each running on a separate core of the CPU for performance reasons. These four instances are identical but the random number generators in each are assigned unique seeds to ensure that the various settings are applied differently. In practice these function as four “voices” in the traditional synthesizer design sense.

Each voice can play one sound at a time. The sound library is shuffled, so that (in my case) the order of the 1200 sounds is randomized and each sound only appears once in the list. A counter is set up in each voice which works its way through the list to load the next sound. This ensures that every sound in the library will be played once before a sound is played a second time. In practice this means that when a specific audio file is heard, it is also guaranteed to *not* be heard again for some time.

Each time an audio file is loaded, two random numbers are generated: offset and length. The length number represents how long the sound will be played for - the range of possibilities of this number is dictated by the length of the original file, and is calculated as a percentage. Similarly, the offset number represents the starting point of playback. Thus, a 30 second audio file could be played back for 6 seconds starting 10 seconds into the recording, or perhaps you will only hear the final moments of a 200 second audio file.

Each voice also has an automated volume control. This is set using the absolute value of a slowly oscillating sine wave, fading the sound playback in and out. The precise frequency of this sound wave is set by a random number generator.

The end result is an endlessly generative sound composition, with periods of quiet, almost ambient soundscapes interrupted by snippets of voices and often recognizable moments from a listener’s past. When the four volume oscillators overlap the effect can be sometimes overwhelming, with four continuously changing voices playing back at once, whereas at other times when only one or two sounds are audible it can feel more akin to glancing at a photograph.

3.3 Physical Construction

The AoYS physical design has gone through several iterations, from desktop-computer based code to cardboard prototype (figure 3) to two different wooden enclosures (figure 4 and figure 1). The most recent version is built from laser cut plywood with one side populated by various sized holes for sound output.

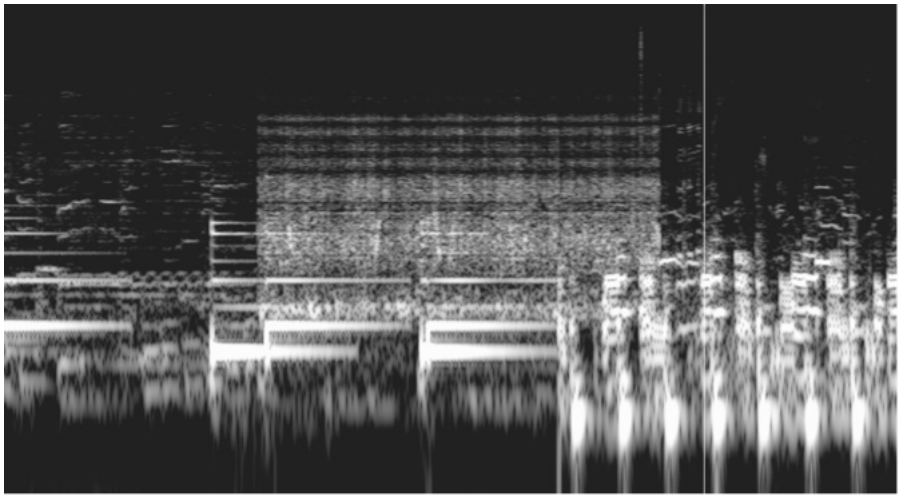


Fig. 2. : 8 Spectrogram showing 8 seconds of sound generated by AoYS. At least four different layered sounds are visible.

While the technical aspects of the build could be considered finished, in the sense that the audio engine, embedded computer system, volume control, and sound output all deliver on the intended user experience, the physical build is still considered provisional and as such will not be analyzed in any depth in this paper. The current designed object serves mainly as a way of running the sound engine in conjunction with the personal library away from a desktop computer, allowing it to play the role of a potentially domesticated object [12].

3.4 Listening to All of Your Sounds

I have built several iterations of the AoYS devices, mainly for use with my own personal library. In addition I have prepared customized versions for several others in an informal way, using them as a method for reflecting on the process of creating the sound libraries and the particularly personal qualities they inhabit. In order to illustrate the experience of the AoYS machine, I invite the reader to watch the documentation video and listen to the sound recording of the sound engine, presented as a video showing a 7.5 minute audio file playing with accompanying sonogram (figure 5). Links to this media are available at the end of the paper (see Appendix).

Additionally, I offer this set of notes from a session of listening to my own AoYS as it generated this soundscape from 12 years worth of home personal recordings. These notes represent my reactions to listening to the machine live, and correspond to the included spectrogram recording.

> Brief kora playing, a concert somewhere? And a child singing, must be my daughter, very young. Chomp chomp! Kicking a football to my dog, very very quiet but noticeable to me.
> Skating, wind noise, guitar playing, who is that?! Whooaaaa, Neenawneenaw
> I remember sending that message to a friend. Sound check on a stage somewhere, so much echo



Fig. 3. Cardboard prototype, with early project name referencing the push notifications from the Mac Photos app

Terrible trumpet playing, why did I record this.
Experimenting with some strange instruments, they sound pretty nice actually
I heard it through the grapevine. That trumpet way in the background, I guess I made several videos
The squigglebox! A thank you video, such excitement I'm sure you can climb that tree, if you just give it another go, my wife giving some solid advice there So many sounds happening all at once now, hard to keep up Happy birthday to…who?
Little kid voices, he didn't know how to talk that much then, he was so nervous on video
Mixing the languages. I remember that broken intercom A ceilidh dance, in Scotland or in America? Not sure Synthesizer experiments and a kid who can't find something…Oui!
Strong scottish accented kid, must be a workshop or something I ran? Giggling tiny little baby Kazoo + piano experiments. Who is playing? Aww I miss that dog so much. So many happy birthday videos
Another overwhelming period, so many sounds at once, dog barking and piano and introducing a train and guinea fowl and more ice skating
A very rainy day, counting in Swedish, a silly game I made where kids shout "Noodles"



Fig. 4. Enter Caption

The kids have such different voices then.
Panting dog, recordings from Ethiopia, an awkward presentation video for a thing I made
Wind noise, another ceilidh (or the same one?)
Playing football with a little kid somewhere, sometime I obviously say “bravo” a lot, who knew. and then Dance music?! When was that? A wedding?
Glitchy electronics, then sounds like a gig. Quiet moment, little blasts of some broken electronics
Testing out a drum machine, recordings from Ethiopia, suddenly so many recordings of children. Playing the Ghana Post Office song on piano. Tu dis merci à Granny? A vinyl record playing somewhere.

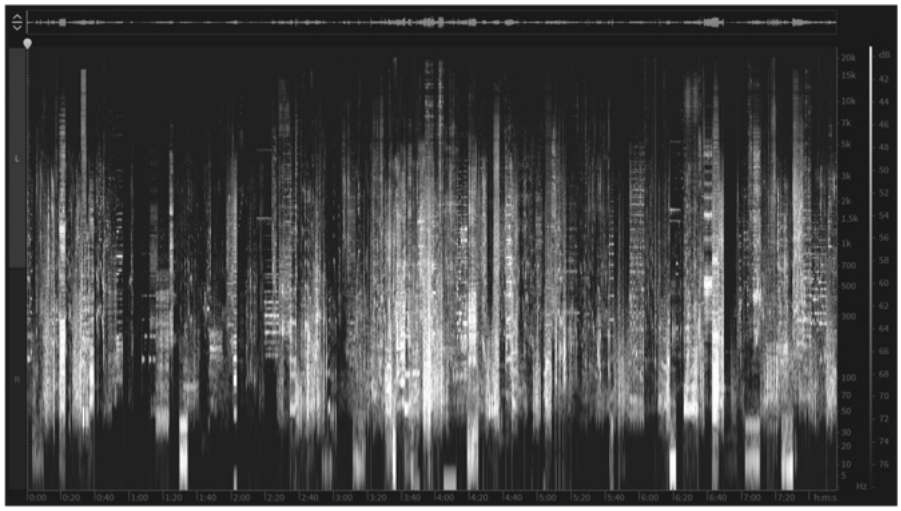


Fig. 5. A spectrogram of approximately 7.5 minutes of sound generated by AoYS

4 Reflections

4.1 The Sound

While the use of sound recording as a personal media capture tool is generally limited, video recording is more popular than any previous period in history. However, watching a video requires a specific form of visual attention – it is very difficult to watch several at once, and a moving picture can often act as a distraction [25]. Listening to the sound of a video offers a different experience entirely, often resulting in a deeper awareness of the moment that was recorded. It is also very possible to listen to several sounds at once, distinguishing them from each other whilst appreciating how they overlap in surprising and unexpected ways.

The ease and casualness of modern smartphone video generates a particular set of qualities to the resulting sounds as well. My library was built up entirely accidentally, or at least not deliberately. The videos are generally low fidelity, full of wind and extraneous noise. Each video was recorded individually as an attempt to capture a single moment, without necessarily any larger strategy in mind for building up a lifetime of content. And yet after twelve years I have quite unexpectedly developed a library of media that represents most of the major events in my life. The library was almost entirely recorded on a series of similar types of (increasingly powerful) mobile phones. This uniformity also means that the library as a whole eschews the pretense of realism that is so pervasive in the field recording sector and associated with problematic representations of nature and purity [10].

In the opening paragraph of the "Design" section of this paper, I described the audio output of All of Your Sounds as a "non-repeating and constantly shifting generative soundscape". This term is used here primarily due to the simple fact that there does not seem to exist a satisfying term in English to describe a form of generative sound that is both built from what could be described as field recordings (albeit not necessarily made with that intention or framework)

and can be approached with multiple listening modes. Certainly, when listening to my own AoYS device I am regularly hearing specific sounds that I can recognize (the voice of a child, whose age I can guess by the way they are speaking, the jingle of a dog collar). This absolutely aligns with Chion's semantic listening [3]: I am listening and identifying sonic moments that have a specific meaning to me and which generate a specific set of emotional reactions. Having these moments presented, and then snatched away from me is both powerful and immensely frustrating. However at other times several sounds will be layered, diminishing my ability to distinguish individual moments. The resulting palette will become more akin to an ambient soundtrack, moving the experience towards reduced listening, in which I am struck by the timbre and abstract sonic qualities of my life and the particular qualities of the sound of handheld mobile phones outside in casual familial situations. A parallel can be made in those moments with musique concrète or creative field recording collage techniques, making AoYS into a procedural composition tool – one that takes advantage of modern digital techniques that allow for endlessly shifting semi-randomized audio generation. In the absence of a specific term for this approach, I have thus reached for *generative soundscape* as a potential description with the intention of borrowing certain aspects of all of these different uses of the term.

The design of AoYS required a set of decisions around the sound engine which dictated the resulting sonic palette. Most important were the minimum and maximum bounds of the potential playback length of a sound, the rate of volume fading of each voice, and the number of possible simultaneous voices. Changing any of these would result in a drastically different listening experience – indeed, in early tests a version was developed that could play 50 sounds simultaneously, with a knob for selecting the precise amount of "memories" to play at once. This was both emotionally and technically overwhelming, and eventually the engine was dialed down to a maximum of 4 to allow me to just about parse the sounds even when all voices are active. My own choices here represent the compositional decisions, and they are in the service of generating a particular sonic aesthetic.

4.2 The Library

The sound library that forms the basis of the generative soundscape is made from an extremely personal batch of files. It is very likely that anyone who makes a sound library for AoYS is literally the only person with access to those files (apart from perhaps a husband/wife/partner or close family member), and consequently the library contains meaning primarily for the owner. While I am happy to see some photos or videos of a friend's small child or a dog, I am unlikely to lose hours scrolling through years worth of photos of that child or dog unless they are my own child or dog. I am not only the owner of those photos and videos, I am also the creator of the content: I was there when the moments occurred and I can either remember them (and feel proud, poignant, whistful, happy, sad, etc) or I have forgotten them (generating feelings of confusion, wonder, self-loathing, surprise, etc).

From the perspective of a designer working specifically with sound, it is worth acknowledging the quiet revolution of recent years that has resulted in many millions of people walking around with vast personal media libraries in their pockets. On the one hand this presents a unique opportunity to explore how personal sounds can be explored and leveraged in new ways, both creatively and in terms of more passive media explorations. However, through the design process of AoYS it became clear that the media ecologies that form the background of these libraries are also responsible for enormous friction and barriers. The process of exporting the videos, for example, is tedious and impractical at virtually every stage. The library management systems are not necessarily actively malicious in their design, but in practice it is certainly difficult to export hundreds, if not thousands, of video files. This is due to a number of factors, primarily how the videos are often by default available only on a cloud server. Clicking and dragging multiple files at once will trigger a download of each of them, creating slowdown – when this is even possible. The iCloud web interface,

for example, allows for a maximum of 1000 download items at a time, presumably a maximum value set by designers who could not imagine that anyone would want to download such a vast repository of media at once. The bias against using the library in a creative way is built into the design of the technology.

The sheer size of these media libraries is the motivation behind a number of design research projects that look at improving search or browsing [7, 11], and many of the popular default photos tools now offer automatic systems for categorizing and organizing personal media with this in mind. Other projects such as Odom's slow design work aim to provide alternative non-solutionist approaches to this issue of scale. AoYS suggests another approach, one in which the existence of the media libraries themselves are also questioned. Presenting the listener with not only some moments from their past but *every single sound they've ever recorded and stored* is accompanied by an implicit question: why do we have these sounds? What was I saving them for? How do I listen to them now?

The boundaries of this sound library should also be considered. The sounds in my own AoYS device date back 12 years, coinciding with the beginning of my cloud storage account. Having ready access only to those sounds, and very few recorded before that, lends artificial weight to that period of my life. I have lived (significantly) longer than 12 years, and yet my AoYS only presents me with a slice of that time. Similarly, the recordings in the device are overwhelmingly positive – dozens of happy birthday videos and recordings of children doing adorable things. As far as I know I have not a single sound recording taken in the middle of a sleepless night with a teething baby or the fourth trip outside with an antsy puppy. These somewhat arbitrary boundaries, temporal and emotional, can be seen as forming a part of the media ecologies that we inhabit which guide and inform how we record and then re-experience our personal memories. While it feels like a library of this size should capture my entire life, AoYS suggests that it is actually a very limited representation of my past.

4.3 The Listener

Listening to the sounds and recognizing (or not) the contents of the library is the activation of the media. In the case of the physical build of AoYS, the listener turns the volume knob as the first step in a moment of active listening. As otherwise unrelated moments are weaved together – a recording of a concert last week overlayed with a giggling baby from ten years ago – the listener forges connections between them and is also overwhelmed by them. Listening to the generative soundscape is an act that folds together the media ecology and the sound. The listener becomes both aware of the vast quantities of data and the characteristics that join it all together, from the recurring characters (human and non-human) to the aesthetic of the phone as a recording medium.

The compositional decisions that went into the generation of the soundscape encourage a particular form of listening, which is both selective and ambient. One the one hand, singular moments are identified, such as in the list presented in Section 3.4. Listening becomes an act of continuous questioning – "what is that sound?". This can be shocking or funny or puzzling, but only really to me as I relate to my own media and memories. Just as I would not expect someone else to have the same emotional response to looking through an album of my family photos, the stream of personal sounds generated by AoYS are not meaningful to a different listener. The listener must form the bridge between the sonic output and the media ecology in order to generate meaning.

Applying this concept to a design results in an object that can treat sound as almost a form of personal decoration. In much the same way that we would not display personal photos of cherished love ones or significant life events in someone else's home, I would not necessarily give my own AoYS to someone other than a close family member. Rather it functions as the sound equivalent of a digital photo frame, enabling me to experience the videos I have recorded over

the course of many years. This is fundamentally different to the act of scrolling through my library and finding a video to watch, a difference stemming from the focus on sound.

4.4 The design as a conversation with the material

The three interwoven elements described here – the sound, the media library, and the listener – come together to form a new material. The design of All of Your Sounds is a conversation with this material in the Schön sense: an exploration of what can unfold when this material is used and explored. The definition and recognition of the material and what it is capable of is perhaps the most crucial part of the process. All of Your Sounds is, after all, only one particular design instance, one possible conversation. Others are possible. As this design space is expanded so to can our understanding of our digital media ecology and the potential of personal sound as a medium. I would like to encourage more design work with the sound of personal media libraries, and as such I offer the following takeaways from AoYS:

- **Sound as a gateway**: Personal media capturing and storage is strongly biased towards visual mediums of photography and video. Overtly shifting focus to sound can offer new avenues for development of designs for personal media.
- **Awareness of our media ecology**: The underlying premise of AoYS was to use every sound in a user's library. The end result makes the listener keenly aware of the frictions and limitations embedded into our media ecology, the scale of time, and the particular sonic aesthetic of their memories.
- **Anti-solutionism**: This design is overtly avoiding a solutionist approach to personal media, suggesting that traditional methods of encouraging browsing or search can be accompanied by techniques drawing from musical composition and slow design to offer a different type of emotional relationship with digital memories.

5 Conclusion

This paper has presented All of Your Sounds, a designed sound object that generates a virtually infinite soundscape from a user's personal cloud media. I have presented the design and analyzed how the artefact represents a conversation with a particular material, formed from the relationship between sound, personal media, and the listener. Unfolding this relationship enables new ways of understanding our relationship to the digital media ecology and how we listen to the vast libraries of sound we create and store.

Acknowledgments

Spectrogram video of All of Your Sounds generated audio file: https://bearnavy.nohost.me/nextcloud/s/YFjQGtjYoNx9TTi
All of Your Sounds device demonstration: https://bearnavy.nohost.me/nextcloud/s/TDYJbS5CJHQARDL

References

[1] Andrea Bak. 2009. *Else Marie Pade: en biografi.* Gyldendal, København. OCLC: 667231490.
[2] John Bowers. 2012. The logic of annotated portfolios: communicating the value of 'research through design'. (2012).
[3] Michel Chion. 2012. The three listening modes. *The sound studies reader* 1 (2012).
[4] Lucy Fife Donaldson. 2017. You have to feel a sound for it to be effective. In *The Routledge companion to screen music and sound.* Routledge, 85–95.
[5] William Gaver, Peter Gall Krogh, Andy Boucher, and David Chatting. 2022. Emergence as a Feature of Practice-based Design Research. In *Designing Interactive Systems Conference.* ACM, Virtual Event Australia, 517–526. doi:10.1145/3532106.3533524
[6] William W. Gaver, John Bowers, Andrew Boucher, Hans Gellerson, Sarah Pennington, Albrecht Schmidt, Anthony Steed, Nicholas Villars, and Brendan Walker. 2004. The drift table: designing for ludic engagement. In *Extended abstracts of the 2004 conference on Human factors and computing systems - CHI '04.* ACM Press, Vienna, Austria, 885. doi:10.1145/985921.985947

[7] Cathal Gurrin, Gareth J. F. Jones, Hyowon Lee, Neil O'Hare, Alan F. Smeaton, and Noel Murphy. 2005. Mobile access to personal digital photograph archives. In *Proceedings of the 7th international conference on Human computer interaction with mobile devices & services*. ACM, Salzburg Austria, 311–314. doi:10.1145/1085777.1085842

[8] Lars Hallnäs and Johan Redström. 2001. Slow Technology – Designing for Reflection. *Personal and Ubiquitous Computing* 5, 3 (Aug. 2001), 201–212. doi:10.1007/PL00000019

[9] Timothy Ingold. 2007. Against soundscape. In *Autumn leaves: sound and the environment in artistic practice*, Angus Carlyle (Ed.). Association Double-Entendre in association with CRISAP, Paris, France, 10–13.

[10] A. M. Kanngieser. 2023. Sonic colonialities: Listening, dispossession, and the (re)making of Anglo-European nature. *Transactions of the Institute of British Geographers* 48, 4 (Dec. 2023), 690–702. doi:10.1111/tran.12602

[11] Joseph 'Jofish' Kaye, Janet Vertesi, Shari Avery, Allan Dafoe, Shay David, Lisa Onaga, Ivan Rosero, and Trevor Pinch. 2006. To have and to hold: exploring the personal archive. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Montréal Québec Canada, 275–284. doi:10.1145/1124772.1124814

[12] Mike Michael and William Gaver. 2009. Home Beyond Home: Dwelling With Threshold Devices. *Space and Culture* 12, 3 (Aug. 2009), 359–370. doi:10.1177/1206331209337076

[13] Paola Moscoso, Mika Peck, and Alice Eldridge. 2018. *Systematic literature review on the association between soundscape and ecological/human wellbeing*. preprint. PeerJ Preprints. doi:10.7287/peerj.preprints.6570v2

[14] William Odom and Tijs Duel. 2018. On the Design of OLO Radio: Investigating Metadata as a Design Material. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–9. doi:10.1145/3173574.3173678

[15] William Odom, Mark Selby, Abigail Sellen, David Kirk, Richard Banks, and Tim Regan. 2012. Photobox: on the design of a slow technology. In *Proceedings of the Designing Interactive Systems Conference on - DIS '12*. ACM Press, Newcastle Upon Tyne, United Kingdom, 665. doi:10.1145/2317956.2318055

[16] William Odom, Erik Stolterman, and Amy Yo Sue Chen. 2022. Extending a Theory of Slow Technology for Design through Artifact Analysis. *Human–Computer Interaction* 37, 2 (March 2022), 150–179. doi:10.1080/07370024.2021.1913416

[17] William Odom, Ron Wakkary, Jeroen Hol, Bram Naus, Pepijn Verburg, Tal Amram, and Amy Yo Sue Chen. 2019. Investigating Slowness as a Frame to Design Longer-Term Experiences with Personal Data: A Field Study of Olly. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, Glasgow Scotland Uk, 1–16. doi:10.1145/3290605.3300264

[18] William Odom, Ron Wakkary, Youn-kyung Lim, Audrey Desjardins, Bart Hengeveld, and Richard Banks. 2016. From Research Prototype to Research Product. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 2549–2561. doi:10.1145/2858036.2858447

[19] James Pierce. 2021. In Tension with Progression: Grasping the Frictional Tendencies of Speculative, Critical, and other Alternative Designs. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–19. doi:10.1145/3411764.3445406

[20] Johan Redström. 2017. *Making Design Theory*. The MIT Press. doi:10.7551/mitpress/11160.001.0001

[21] Brigitte Robindore and Luc Ferrari. 1998. Luc Ferrari: Interview with an Intimate Iconoclast. *Computer Music Journal* 22, 3 (1998), 8. doi:10.2307/3681154

[22] Pierre Schaeffer. 1998. *À la recherche d'une musique concrète* ([nachdr. der ausg.] paris, seuil, 1952 ed.). Seuil, Paris.

[23] R. Murray Schafer. 1977. *The tuning of the world*. Knopf, New York.

[24] D.A. Schön. 1992. Designing as reflective conversation with the materials of a design situation. *Knowledge-Based Systems* 5, 1 (March 1992), 3–14. doi:10.1016/0950-7051(92)90020-G

[25] Daniel J. Simons. 2000. Attentional capture and inattentional blindness. *Trends in Cognitive Sciences* 4, 4 (April 2000), 147–155. doi:10.1016/S1364-6613(00)01455-8

[26] Anne-Sofie Udsen and Kim Halskov. 2022. Soundscape design for historical buildings as a sonic place-making process. doi:10.21606/drs.2022.178

[27] Ron Wakkary, Doenja Oogjes, Henry W. J. Lin, and Sabrina Hauser. 2018. Philosophers Living with the Tilting Bowl. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, Montreal QC Canada, 1–12. doi:10.1145/3173574.3173668

[28] Hildegard Westerkamp. 2002. Linking soundscape composition and acoustic ecology. *Organised Sound* 7, 1 (April 2002), 51–56. doi:10.1017/S1355771802001085