Olli Kietäväinen

Crafting Desolation

Designing Post-Apocalyptic Soundscapes Using Audio Synthesis

Metropolia University of Applied Sciences

Bachelor of Culture and Arts

Film and Television

Thesis

November 23, 2023

Abstract

Author(s):	Olli Kietäväinen
Title:	Crafting Desolation: Designing Post-Apocalyptic Soundscapes Using Audio Synthesis
Number of Pages:	46 pages + 9 appendices
Date:	November 23, 2023
Degree:	Bachelor of Culture and Arts
Degree Programme:	Film and Television
Instructor(s):	Aura Kaarivuo, Senior Lecturer Kai Ansio, Senior Lecturer
Commissioning Party:	Heckmann Audio GmbH

This thesis examines whether audio synthesis is a worthwhile method of designing soundscapes for post-apocalyptic fiction. Its theoretical part covers three subject matters: soundscapes and their elements, the effects of an apocalyptic event on the soundscapes of the world, and audio synthesis. For its practical application, three surreal soundscapes were synthesized with marginally different methods. It was determined that audio synthesis can be successfully used to create immersive and convincing soundscapes given that enough time is allocated to the process.

Keywords: ambiance, ambience, ambient, apocalypse, apocalyptic, art, audio, crafting, creating, design, designing, Fallout, Fallout franchise, Mark Morgan, modular, music, nuclear, nuclear holocaust, nuclear war, post-apocalypse, post-apocalyptic, semi-modular, sound, sound design, sounds, soundscape, soundscapes, synth, synths, synthesis, synthesizer, synthesizers, synthesizing, u-he, Zebra, Zebra 2, Zebra 2.9.3, Zebra Legacy

Tiivistelmä

Tekijä(t):	Olli Kietäväinen
Otsikko:	Lohduttomuuden luomistyö: maailmanlopun jälkeisten äänimaisemien suunnittelu äänisynteesin keinoin
Sivumäärä:	46 sivua + 9 liitettä
Aika:	23.11.2023
Tutkinto:	Medianomi (AMK)
Tutkinto-ohjelma:	Elokuva ja televisio
Ohjaaja(t):	Lehtori Aura Kaarivuo Lehtori Kai Ansio
Tilaaja:	Heckmann Audio GmbH

Tämä opinnäytetyö tutkii, onko äänisynteesi varteenotettava työkalu äänimaisemien suunnitteluun maailmanlopun jälkeistä fiktiota varten. Sen teoriaosa kattaa kolme aihepiiriä: äänimaisemat ja niiden yksittäiset osat, apokalyptisen tapahtuman vaikutukset maailman äänimaisemiin ja äänisynteesi. Opinnäytetyötä sovellettiin käytäntöön syntetisoimalla kolme surrealistista äänimaisemaa hiukan eri menetelmin. Sovelluksen aikana selvitettiin, että äänisynteesiä on mahdollista käyttää immersiivisten ja uskottavien äänimaisemien tekemiseen, kunhan niiden suunnitteluun on varattu tarpeeksi aikaa.

Avainsanat:	ambient-musiikki, Fallout, Fallout-sarja, luominen, luomistyö, maailmanloppu, maailmanlopun jälkeinen, Mark Morgan, modulaarinen, modulaarisuus, modulaarisynteesi, post-apokalyptinen, synteesi, syntetisaattori, syntetisaattorit, syntetisointi, taide, u-he, ydinsota, ydintuho, Zebra, Zebra 2, Zebra 2.9.3, Zebra Legacy, äänet, ääni, äänimaisema, äänimaisemat

---

Contents

1 Introduction

What does the post-apocalypse sound like? Can a soundscape be synthesized for a world that has met its end?

The goal of my thesis was to examine these questions through two key variables. Firstly, I decided to use the semi-modular virtual synthesizer Zebra 2.9.3 (u-he n.d.) by the German synthesizer and effects company u-he to create three distinct soundscapes. Secondly, I aimed to specifically create the soundscapes in the surreal, post-apocalyptic style of Mark Morgan, who composed the original scores for the video games Fallout (Fallout, USA 1997) and Fallout 2 (Fallout 2, USA 1998) created by Black Isle Studios and released by Interplay.

The theoretical side of my thesis consists of three parts. Firstly, some light is shed on the terminology used to describe soundscapes and several of their elements. Secondly, some of the details of what would happen to the world of sounds if the human species were to face an apocalyptic event are examined. Finally, the basics of audio synthesis are glossed over, and the most relevant functions of my main tool, Zebra 2, are described.

The practical application side of my thesis first depicts how I aimed to recreate two of Mark Morgan's Fallout soundscapes as closely as I could only using Zebra 2. Finally, in the penultimate chapter I delve into developing my own take on a post-apocalyptic soundscape using the same techniques, after which a collection of my thoughts about the process are presented in the conclusions.

It is worth noting, that while there are far easier ways of creating soundscapes for post-apocalyptic fiction – ones that can even yield richer and more nuanced outcomes – using a virtual synthesizer, such as Zebra, does have its benefits. The most obvious of these benefits becomes apparent when comparing file sizes. While high quality PCM soundscape files in formats such as WAV or AIFF can often take several hundreds of megabytes of hard disk space each, a soundscape patch made for Zebra fits in mere kilobytes without any intrinsic shortcomings in sound quality as it only contains control data for the synthesizer. This may not always remain a relevant factor, but at the time of writing in 2023, we are yet to obtain the luxuries of infinite hard drives or other forms of storing data, and especially laptop computer users may appreciate the smaller patch files of Zebra.

Even more significant a benefit is the malleability of these synthesizer sounds. A PCM sound file is usually distributed as is, and can only be altered by using post-processing, by manipulating the playback device and/or environment or by editing the audio file itself on a binary level. In my experience, these limitations are often far from ideal. A Zebra sound patch, however, can be modulated in several ways that the designer of the sound can themselves finetune to their liking.

Finally, there is the exciting prospect of challenging oneself to doing things the hard way. While building a physical noise box from a container, metal rods, screws, springs, and a piezo microphone can be fun – and playing one with bows and mallets through delays and reverbs definitely is, I wanted to figure out just what I could achieve if I dedicated some time to doing essentially the same thing with a semi-modular virtual synthesizer.

As is obvious, most soundscapes to which we are exposed daily are not synthesized. Even in audiovisual fiction, so-called room tones or other ambience recordings of actual environments are often used to convey a sense of space in situations where the recording location itself was insufficient. But when it comes to the imaginary and the surreal, audio synthesis allows sound designers to conjure up spaces that do not exist, and where no microphone can be placed.

And since our world has not ended, yet, some kind of detour into sound design has to be taken to invoke its aural apocalypse.

2 What Are Soundscapes?

In this chapter the term ‘soundscape’ and several of the individual elements of soundscapes are defined based on literary references.

2.1 Defining Soundscapes

International Organization for Standardization defines soundscapes as "acoustic environments as perceived or experienced and/or understood by a person or people in context" (ISO 12913-1 2014). They consist of sound sources "generated by nature or human activity" and specifically entail "sound at the receiver from all sound sources as modified by the environment". Furthermore, the "acoustic environment can be actual or simulated, outdoor or indoor, as experienced or in memory."

Composer and author R. Murray Schafer illustrated the idea with somewhat grounded terms in his seminal book The Tuning of the World: "The soundscape is any acoustic field of study. We may speak of a musical composition as a soundscape, or a radio program as a soundscape or an acoustic environment as a soundscape. We can isolate an acoustic environment as a field of study just as we can study the characteristics of a given landscape." (Schafer 1977, 7.)

In other words, depending on the context and the intent of observation, a soundscape can entail virtually anything audible and virtually anything audible can justifiably be considered a soundscape, if it has an observer capable of perceiving it in a space shared by both.

2.2 Dissecting Soundscapes

Every soundscape can be divided into individual elements. And while it is even possible to focus on minute details of single sounds, analysing the frequency contents, individual samples or the modulating harmonic series of lengthy soundscapes may not be the best use of anyone’s time.

Instead, I opt for Schafer’s broader divisions, in which various parts of soundscapes are sorted into *keynote sounds*, *signals* and *soundmarks*.

2.2.1 Keynote Sounds

In music, keynotes are notes or frequencies that classify the key or tonality of a composition. In simple terms a keynote gives a musical composition its proverbial anchor onto which all the other notes in the used scale(s) latch in harmony.

Likewise, a keynote sound is a prominent element of a soundscape that may not be its most defining feature, but without which its other elements would lack a reference point. They are sounds that “do not have to be listened to consciously; they are overheard but cannot be overlooked, for keynote sounds become listening habits in spite of themselves” (Schafer 1977, 9).

For example, if a soundscape consists of two people talking in a room that has an air conditioning unit humming in the background, the hum of that air conditioner would be a clear keynote sound of that soundscape.

2.2.2 Signals

If keynote sounds can be subtle or so omnipresent in nature that they may in fact be difficult to detect without conscious effort, signal sounds are their opposites. They are “foreground sounds and they are listened to consciously” (Schafer 1977, 10). Signals spring forth from the humming air conditioners and traffic noises and demand attention to themselves by being either loud or distinct enough.

To reuse our previous example, signal sounds would be the voices of the two people talking. More specifically yet, the way these people laugh or enunciate certain words could also be considered signal sounds.

Furthermore, every single type of alarm is an example of a signal sound.

2.2.3 Soundmarks

A soundmark is to a soundscape what a landmark is to a landscape. The concept "refers to a community sound which is unique or possesses qualities which make it specially regarded or noticed by the people in that community" (Schafer 1977, 10). And while it is true that every sound is unique to some extent, soundmarks refer to sounds the uniqueness of which is easy to recognize and differentiate.

For example, as Schafer suggests, church bells can be heard tolling all across the Christian world, but in many places, the bells are set to play ostinatos or melodies personal enough for them to become unique, especially in their respective environments. In areas where sounds of particular distinctiveness are prevalent, locals or those familiar with the environment recognize these unique soundmarks and intuitively discern any alterations or abnormalities.

However, soundmarks do not have to be signals. A river flowing free in its environment, or the exact way the hum of traffic sounds in certain parts of the city can also qualify as a soundmark. What soundmarks must be, is distinct.

Finally, I would argue that certain types of environmental echoes or reverberations could similarly be considered as soundmarks in their uniqueness.

An example of such a peculiar effect can be experienced behind a large furniture store building in the district of Porttipuisto in Vantaa, Finland. There, as one walks close to the metal overpass of the facility, the echoes of their footsteps form a type of comb filtering effect. Anyone who has walked that path daily has been exposed to the effect, but whether they have paid attention to it, that is up to their personal level of interest and the way in which they observe sound.

A more famous echo of comparable nature can be heard in front of the Parliamentary building in Helsinki, Finland. The distinct stone stairs and columns of the establishment, designed by J.S. Sirén in 1924 (Eduskunta Riksdagen n.d.), form a series of rapid echoes that are familiar to most people who have attended demonstrations in front of the granite structure.

2.2.4 Noise

There are numerous definitions for the term noise, even in the context of audio, two of which are used for the purposes of this thesis. First, in chapter three, the concept of noise pollution is used. These are the types of sounds that are considered unwanted in a human habitat, and which can cause various health issues with prolonged exposure, including “annoyance, sleep disturbance, and cardiovascular and metabolic issues” (European Environment Agency 2023).

In chapters five through eight, the types of noises to which I refer are so-called coloured noises. They are “sounds that can perfectly be synthesized by filtering white noise” (Hanna & Louis & Desainte-Catherine & Benois-Pineau 2004, 1). White noise, in turn, is “the effect of the complete range of audible sound-wave frequencies heard simultaneously, analogous to white light, which contains all the frequencies of the light spectrum” (Britannica, 2023).

Furthermore, the signal-to-noise ratio is of importance when soundscapes are examined. R. Murray Schafer writes about the phenomenon in terms of Hi-Fi and Lo-Fi soundscapes. The more hi-fi a soundscape is, the lower is the amount of its unwanted elements, and vice versa. “In the ultimate lo-fi soundscape the signal-to-noise ratio is one-to-one and it is no longer possible to know what, if anything, is to be listened to” (Schafer 1977, 71).

When soundscapes are synthesized, noise is often a pivotal part of the design process since many of the sounds of our daily soundscapes have noisy elements to them. Especially the many states of wind and other similarly ubiquitous drones often find an optimal starting point in layers of noise. From this foundation, adjustments can be made to accentuate or diminish their specific aspects, which helps integrating the designed sound into its desired context.

3 The Human World and How to Destroy It

The end of the world could occur in countless of ways, and especially in fiction, they take various shapes and forms. In this chapter the end of the world is examined on a worldwide scale, especially through the genre lens of what are referred to as post-apocalyptic and post-post-apocalyptic fiction.

All fiction reflects the era of its conception, but this is particularly tangible in the case of post-apocalyptic fiction since it forces us “to face urgent socio-political questions such as danger of globalization, effect of neoliberal capitalist hegemony, ecological disasters, fragility of human civilization, and so on” (Moon 2014, 2). Whereas depictions of the post-apocalypse explore life and events after the fall of civilizations, portrayals of post-post-apocalypses envision the possibilities of what we could accomplish if we survive Armageddon and begin “fighting for the imagined utopia rising from the ashes” (Roby 2013, 77).

One of the most intrinsic questions when imagining the kinds of soundscapes that could be perceived after the end of the world, is the nature of the calamity itself through which the apocalypse would occur. The human civilization with all its manifestations cannot endure indefinitely, but it definitely cannot be sustained if there are not enough people to maintain it. If humanity faces any widespread disaster that substantially diminishes our numbers, much of our infrastructures will be lost to entropy with no way of preventing it.

The degree of devastation to our surroundings, and with it the nature of the soundscapes available to any entity still capable of experiencing them, will vary immensely if the alternatives are, for example, a pandemic that wipes out the entire human species without affecting other animals or our infrastructures, and an all-out nuclear war that annihilates most lifeforms on the planet without selectivity and turns all the human constructions it touches into dust and rubble.

Before we can appreciate the variety of sound emerging from beneath the noise floor of the world, however, we must first pay deliberate attention to all the sounds that would vanish along with our species.

3.1 Our Lo-Fi World

Even if it can be argued that the world has already taken several steps into the apocalyptic direction during our lifetime, when it comes to noise, we still have ways to go before it can be said to have ended on a global scale. Our species is still the loudest and the most constant proponent of noise pollution on the planet, and our numbers are rising (United Nations 2022).

This in turn leads to increasing requirements on the fields that produce the most unwanted noise. Our traffic and industries are responsible for volumes of sound matched in nature only by few sources. More on them in chapter 3.3.

Traffic is the one part of the urban soundscape it feels impossible to evade without taking a train to the countryside (Nugent & Blanes & Fons & Sáinz de la Maza 2016). The constant hum and whirr of traffic follows us into our homes, into our places of work and even into the mouldy cellars where we wash our laundry.

Regardless of the reason because of which the population of the world would decrease, a reduction in our numbers would inevitably lead to the sustained noises of traffic dwindling away. They could decline to a point where, instead of the constant keynote sounds of masses of vehicles grinding against the tarmac, the many sounds of motor vehicles would become signals.

Easy to pinpoint from afar with no other traffic to mask them, the sounds of approaching vehicles would announce themselves to anyone listening within a several-kilometre radius. Instead of an everyday non-event, they could even begin to signal danger. After all, whom can one trust, when our various social contracts expire, and no-one is left to enforce our arbitrary laws?

The silencing of industry, due to an insufficient workforce, would lead to the cessation of heavy machinery, factories, and essential facilities. Consequently, many goods deemed integral to our lifestyles would cease production, only to fade into obscurity for future generations.

We have built a world that seems to hum and murmur by itself. But no noise is created without something causing it. Even soundscapes that feel all-encompassing consist of individual elements, and the ones created by humans will not last forever without someone to keep the wheels turning.

3.2 Working After the End of the World

If the remaining humans had to focus on mere survival, at least for a time, the meaning of work itself would change in both the post-apocalypse and the post-post-apocalypse. The surviving peoples would have to re-evaluate their priorities upon facing an apocalyptic event, and without management and shareholders many industries would wither away as people would see their redundancy. For many, work would either readjust to or be replaced by means of survival such as hunting, gathering and small-scale agriculture. In many ways, a simpler time would arise from the ruins of the old.

Hunting generates sparse signal sounds, especially if firearms or hunting dogs are used. But compared to the enormous livestock production facilities where our carnivorous desires are met, hunting is still limited as a soundscape element.

Whereas supermarkets and food processing facilities generate plenty of both keynote and signal sounds, gathering mushrooms, plants, berries, or roots is practically silent in comparison. And while both industrial and small-scale agriculture create keynote and signal sounds alike, the difference in the amplitudes between the two are immense.

If there would be enough people to form towns or communities after the apocalyptic event, some sounds lost after the Industrial Revolution could eventually return to our lives. Borrowing another example from Schafer, the town blacksmith's trade, characterised by the roaring furnaces and the rhythmic clinking of hammers on anvils, might regain exceptional respect after the Armageddon. Simultaneously, horseback riding or carriages could reclaim their status as preferred methods of transportation, as an eventual shortage of gasoline and car batteries would render modern cars obsolete.

Some survivors might consider for example electricity to be a commodity important enough to spend time operating a power plant. The hum of electricity and all the sounds related to its generation would hence continue in the vicinity of these facilities for as long as they would be maintained and operated.

By necessity, should some cataclysm wipe most of us off the face of the Earth, our ways of working would quiet down along with our dwindling numbers. In our stead, what would remain of nature would once again claim the position of the single most prominent composer of the soundscapes of the world.

3.3 Nature After Us

As indicated in chapter 3.1, only few sources in nature can match the immense amounts of noise generated by our traffic and industries. For example, tempests, oceans, fires, types of geysers, earthquakes and most infrequently erupting volcanoes can reach some of the noise levels of our technologies.

But if a Doomsday scenario were to obliterate most of humanity, our downfall would unveil the myriad voices of nature long overshadowed by the inadvertent soundscapes of our creation. In a world devoid of our excellence, even some of the nature’s smaller sounds could stand out as keynotes or signals.

Without the relentless droning of traffic, the many tones of wind would be easier to differentiate. As it would howl through collapsed cities, banging loosened doors and windows against their frames, those who would remain to truly listen to the wind, free from distractions, might notice the nuances of even a gentle breeze. And when it intensifies, how a small gust can gradually start to scream.

Schafer remarks how the sounds of nature are essential to our existence: “The keynote sounds of a landscape are those created by its geography and climate: water, wind, forests, plains, birds, insects, and animals. Many of these sounds may possess archetypal significance; that is, they may have imprinted themselves so deeply on the people hearing them that life without them would be sensed as a distinct impoverishment.” (Schafer 1977, 9–10.)

We tend not to realize how significant many of the simpler things in life are before they are taken from us. If someone from our time were to survive an apocalyptic event, I suspect these realizations would dawn on them in abundance as the impossibility of ever returning to the ways of the old would unfold.

I have always found the pattering of rain to be one of the most soothing sounds in the world. Whether it is tapping the foliage of a summer forest or drumming the tiles of a roof, the gentle falling of rain creates a natural soundscape into which I can let my thoughts drift with ease. It is also one of the sounds the masking of which I lament the most amidst the urban cacophony. One cannot fully enjoy rainfall next to a freeway.

We may not have lost these types of natural sounds altogether, but by masking them we have denied ourselves many of the most pleasing elements in the soundscapes of the world. As Schafer later remarks: “Even the most ordinary sounds will be affectionately remembered after they disappear” (Schafer 1977, 180). If meaning and purpose in soundscapes are lost to noise pollution, what will remain but sensory overload? What joy or message can birdsong deliver if no-one can hear it?

Depending on the type of Armageddon humanity would face, other inhabitants of the planet might still survive the cataclysm. Nature tends to spread everywhere unless it is stopped with intent, and without us our derelict cities would quickly become new homes for all kinds of surviving critters and vegetation.

Any human living in a world devastated by an apocalyptic event would meet countless of challenges in their everyday live, but one thing is certain: they would face those challenges in a world far quieter than the one of today. Whether or not they would be able to appreciate the more nuanced soundscapes of their environments, that would be up to them.

3.4 A World of Decay

As human infrastructures degrade, a plethora of sound sources emerges. Once natural elements, especially water, infiltrate our systems, the inevitable failure of these systems leads to the crumbling of our cities (Weisman 2007, 30).

While buildings already produce sounds by themselves, the subtle creaks and thumps we can nowadays perceive would pale in comparison to the symphonies of noise generated by the decay of unattended structures. With water eroding everything and with no-one to fix leaks or overflow, the foundations of buildings would fall apart. With time, everything from skyscrapers, bridges and overpasses to highways, tunnels and sewers would collapse. (Weisman 2007, 24–46.)

As coolant water supplies would evaporate, nuclear reactors exposed to air would begin to melt or explode, and fires would break in both the reactor facilities and the depots for used radioactive fuel (Weisman 2007, 268–272). By themselves, radioactive materials are silent, but the eruptions they cause are not.

The creaking of bending metal, the rumble of collapsing buildings, the clattering of crumbling concrete and the bursts of water rushing through the cracks of our creations would all be common if sparse elements in the soundscapes of the world for a while as entropy would gnaw its way through the relics of humanity.

It would take a long time, but if we were to go extinct with no chance of restoration, there would eventually be a moment in time during which the last sound of even our indirect making would fade away. At the latest, it would happen as the radiant gas giant of our Solar System would engulf the planet and our remains along with it. Yet, if some humans or other resilient animals were to survive the apocalypse before this cosmic conclusion, the echoes of life would persist.

As long as there is developed life on Earth, at least the kind that still resembles us or our neighbours, it will keep clattering on and using its plurality of voices within the soundscapes of the world. But life is uncertain, and sometimes it takes a form which develops the means to bring about its own end.

3.5 A Nuclear World

The final apocalyptic scenario addressed in this thesis involves a self-destructive nuclear holocaust, devastating both humanity and significant portions of the world with it. The character of Dr. Ian Malcolm famously berates a venture capitalist in the motion picture Jurassic Park with a classic of a line: “Your scientists were so preoccupied with whether or not they could, they didn’t stop to think if they should” (Jurassic Park, USA 1993). Even if, on the surface, he speaks about recreating the dinosaur, it is the hubris of man that is the value at play underneath it.

Next to nothing displays the same level of hubris as the number of nuclear weapons our collective ingenuity has created. Hoarding more of the deadliest explosives we have ever invented just to have more than your supposed competitor, and to end up having enough to burn the entire surface of the planet several times over, is a recipe for, if not disaster, at least for intense fiction.

Unfortunately, atomic warfare has not always remained strictly within the realms of fiction, where the very real horrors of these technologies are often easier to forget for the sake of comfort and psychological distance. But despite the imaginary worlds we enjoy, we must never forget what has transpired in our own.

For on August 6, 1945, a shadow so dark was cast over our collective subconsciousness, that in many ways, we are still recovering from its effects. That was the day when the orders to use the recently acquired nuclear weapon against the city of Hiroshima, Japan, were actualised. The orders were authorised by the president of the United States of America, Harry S. Truman, and they lead to the destruction of the city and many of its inhabitants. (Miller 2020.)

Science author Brian Clegg depicts the events as such: “The bomb exploded. With a temperature at its core of around 60 million degrees Celsius, far hotter than the surface of the Sun, the initial flash vaporized some individuals and turned everyone openly exposed to it for around half a mile around into a fused carbon relic. Immediately after came the shock wave, which combined with that initial flash killed around seventy thousand people – flattening nearly as many buildings in the process” (Clegg 2010, 57–58.) As horrific as that may sound, however, the worst was arguably yet to come for the survivors of the blast.

The nuclear radiation emitted from the explosion resulted in radiation sickness that would “nearly double the initial death toll to well over 100,000 people during the subsequent year. That was out of a population of 350,000. Some have put the eventual total at double this. Although the explosive power of nuclear bombs is terrible, this is the truly terrifying aspect of such weapons: the silent, invisible, deadly action of radioactivity.” (Clegg 2010, 60.)

The images taken from Hiroshima and Nagasaki after the nuclear detonations paint a harrowing picture. The surviving people with severe burns (Appendices 1–2) averted the fate of the ones who were disintegrated (Appendix 3). Only few structures can be seen standing, and most of the landscapes has been rendered to debris (Appendices 4–9). Eventually, both cities were rebuilt, but what if that were not possible? An uneasy quiet would settle into the ashes of nuclear fire.

For a time, only the whispers of the wind would travel the ruins, sighing through scorched rubble and grime. People might know to avoid the radioactive areas, but some animals could wander into ground zero, not knowing the severity of their straying. If they were to call for others, there would be few walls to echo their signals. The eventual collapse of the remaining structures would mark another sparse signal of the soundscape. Eventually, if another person were to move in, theirs would be the first human sounds in the area after an apocalypse.

The nuclear weapon stockpiles of the world, while in slow decline (Kristensen & Korda & Johns & Kohn 2023) are still capable of ending most of human existence and a sizeable portion of life on the surface of the planet if even a fraction of them were to be deployed. The blasts alone would cause an irreversible catastrophe to our civilizations, but combined with the amount of radiation the explosions would release, it would all but render the surface of the Earth uninhabitable for the possible survivors of such a calamity. However, if a contingency plan involved an attempt to outlive the nuclear fallout underground, at least in fiction, humanity might still be able to get back on its feet.

4 The Few Sounds of Fallout

The Fallout franchise has gained a popularity so vast that many of the people reading this may very well be familiar with it already. After the franchise was purchased by Bethesda Game Studios in 2007, Fallout has evolved from a niche computer roleplaying game into a media juggernaut and even a household name recognised even by many people who do not themselves play video games. Fallout 3, Fallout: New Vegas (Kollar 2015) and Fallout 4 (Makuch 2017) have sold over ten million copies each, the massive multiplayer online game Fallout 76 has over ten million players in total (Mäki 2021), and a TV series based on the world of Fallout is slated for release in 2024 (Petski 2023).

However, not as many know of the series' origins, and fewer still have experienced its first two instalments for themselves. Which is why a brief recapitulation of the first Fallout game may be in order.

4.1 The Story of Fallout

Fallout takes place in a world that is destroyed in an all-encompassing nuclear war in 2077. In 2161, ahead of schedule, one of the great underground vaults, where survivors have taken shelter for generations, has to open its doors when their water purification chip suddenly malfunctions. The player character is sent into the irradiated wasteland to find a replacement for the broken part with no information about where to locate one. Alone and at odds with the few surviving people and the many strange creatures that now roam the deserts of Southern California, they are their vault's only hope for survival. (Bethesda n.d.)

4.2 The Distribution Technology of Fallout

To appreciate the soundscapes of the original Fallout game, contemporary audiences may benefit from a review of the technological limitations within which the game and its music were released. Back in 1997, mainstream video games were often much simpler affairs than the ones published nowadays, and even by the standards of the late nineties, Fallout was a comparably small production.

The game was distributed on a single CD-ROM (Nukapedia Fallout Wiki 2023), which were optical discs capable of storing up to 700 megabytes of data. This meant that the game with its executables, assets, animations, few cutscenes and everything else had to fit on the same disc with the game audio, and compromises had to be made. Given these limitations, it would have made little sense to cram soundscape recordings on the disc along with the dedicated compositions.

This led, either directly or indirectly, to the music also having to carry the atmosphere of the game with its presence. Sound effects were used sparingly and were mostly limited to character interactions such as operating of doors and containers and firing of weapons. As far as I can remember, not a single environmental sound effect unrelated to player actions, characters or creatures can be heard in the entire game.

4.3 The Music of Fallout

In hindsight, the limitations of technology turned into an immense advantage, because the ambient music compositions created by Mark Morgan for the first two Fallout games – in my opinion – have stood the test of time better than any of their contemporaries. What Morgan created for the music of Fallout are still among the finest achievements in the history of video game scoring, and I have now carried those tracks with me for almost 25 years.

More conventional ambient music often revolves around various kinds of clearly synthesized pads, smooth drones, few sound effects and simplified, sparse melodies. And while this kind of traditional approach might also have worked for Fallout, Morgan’s style that has been sometimes described as ‘aggressive ambient’ forged a world of sound far more unique.

Morgan’s Fallout compositions can be split into two categories: those that focus on pads and drones, and those driven by percussions.

4.3.1 The Drones

Mark Morgan's pads and drones embody the barren, forlorn aspects of the world of Fallout, and the compositions that employ these sounds form the vast majority of the soundtrack of the game.

The ingenuity of these compositions lies in their gradual transformation, which continues to captivate even after hundreds of hours of listening. They feel like they are constantly in motion and shift through a range of nuances and motifs, no matter how slow that movement is.

Morgan achieves additional sense of progress with sparse use of sound effects both subtle and abrupt. These can be, for example, murmurs from a radio, the clacking of a keyboard, squeals of unknown origin, or ominous church bells.

I would argue that a more static approach could not have achieved an experience as profound and would most likely have bored its listener with its more noticeable repetition even during a single playthrough. These pieces of music are played on repeat until the player leaves a particular area, and this sometimes takes a while. Hence, having intriguing music that can be played on repeat without causing upset is of utmost importance for setting the tone for the entire game.

The soundscapes that are present in these compositions are obviously surreal by intent. Most of the haunted synthesizer pads and drones swaying in various degrees of malevolence are not meant to evoke sensations of actual spaces or natural sound sources. Instead, they excel at evoking emotions that would feel entirely justified when navigating such locations. They radiate feelings of loss aimed at an audience who have not experienced the end of the world and can therefore appreciate what has been reduced into ruins in the world of Fallout.

With all its audiovisual might, the game forces its audience to face the impermanence of our world and acts as a warning sign for everything we could lose. But not without also showing how we might endure.

4.3.2 The Drums

There are four compositions in Fallout where Mark Morgan prominently features percussions. These are A Trader’s Life, Moribund World, Khans of New California, and City of Lost Angels. Set against the bleak drones and pads of the other compositions, these tracks feel like signs of life in comparison. Even if they only form a small part of the whole, they are still a very memorable part of it.

Sparse reverberated percussions such as djembes, metal pipes and timpani can be heard in other Fallout compositions as well, but in those they take the roles of background instruments that do not demand attention to themselves.

5 The Tools and How to Use Them

To understand the soundscape designs that are illustrated in chapters six through eight, it is important to have a basic grasp of how audio is shaped in synthesis.

5.1 Audio Synthesis

Audio synthesis is the process of generating sounds with no acoustic sources (Truax 1999). These technologies offer sound designers, audio engineers and musicians the means of creating sounds that have not existed beforehand, and they have therefore become pivotal to contemporary audio production methodology. Various systems and instruments have been created for this very purpose, and while similar kinds of results can be achieved by remarkably diverse means, some forms of synthesis are better suited for specific applications.

What makes audio synthesis a fine pairing for surreal soundscape design, is the ability to approach a goal of something which does not exist by adding individual elements to a blank slate until auditory satisfaction is reached. Layers upon layers of synthesized sounds can be added to imbue nuance and meaning to the soundscape, enhancing the immersiveness of the listening experience.

5.1.1 Types of Synthesis

The most common type of audio synthesis is called subtractive synthesis. It starts with audio generators called oscillators, which produce harmonically rich audio waveforms such as sawtooth waves, square waves, pulse waves, or white noise, and by then subtracting information from them within either the amplitudinal, the spectral or the temporal domains (Underdog Electronic Music School 2023).

Another popular type is called Frequency Modulation synthesis, or FM synthesis for short, where the sound generators are often called operators instead of oscillators (Huang 2018). In FM synthesis, simpler waveforms such as sine waves or their multiples are often used to modulate the frequencies of other oscillators, which produces new harmonic content into the signal and can help the sound designer to create exceedingly complex sounds. It is worth noting that these are not the kinds of signals of which Schafer often writes.

Additive synthesis is the process of generating sounds both simple and complex by manipulating numerous sine waves tuned to distinct frequencies (Jordan 2017). This follows the principles of how all the sounds of the world can be formed and dismantled. Theoretically, any sound can be created through additive synthesis, but due to how cumbersome it can be to edit multiple sine waves separately, many additive synthesizers utilize methods of manipulating them at once, which makes working on an additive sound somewhat less of a challenge.

Due to its ease of use and flexibility, at the moment, the most popular type of audio synthesis is wavetable synthesis. It functions by arranging and cycling through a user-defined series of single-cycle waveforms called wavetables (Reverb 2018). These wavetables can consist of any samples of sound, which makes it possible to use excerpts of, for example, physical instruments or any other sounds of the world within a synthesizer. This opens up possibilities that would be arduous to reproduce through many other means of synthesis.

More types of synthesis exist, but since these are the ones used in Zebra 2, I will refrain from describing the rest of them. For now.

5.1.2 Oscillators

The two most common oscillator waveforms are the sawtooth and the square. Both are often used because of the richness of their harmonic series, which leads to full sounds that are suitable for subtracting frequencies from them.

A sawtooth, or a saw, as it is often abbreviated, is a waveform the single cycle or oscillation of which starts from its signal minimum and rises to the maximum forming a sort of ramp when visualised with an oscilloscope. A reversed sawtooth is also quite common, and otherwise identical to a regular saw, but instead of ascending, their ramps descend from maximum to minimum. (Truax 1999a.)

A square wave sharply alternates between the maximum and the minimum of its signal, forming, as the name implies, a sort of square waveform when visualised in an oscilloscope. A common parameter for modifying these kinds of waveforms modulates its pulse width, meaning the allocation of time between the minimum and maximum amplitudes of the waveform, though it is worth noting that only a temporally even pulse wave is called a square wave, and the ones with either shorter or longer pulses are called pulse waves or rectangle waves with different widths to them. (Stolet 2009.)

The key difference between the two is that a sawtooth waveform contains both even and odd harmonics whereas a square wave only has odd harmonics.

Oscillators can most often be controlled at least in frequency or pitch. This is determined by the number of oscillations per second, meaning that a frequency or Hertz value of an oscillator is equal to the number of times it repeats its one cycle within a second. For example, an oscillator running at 60 Hz will repeat its cycle 60 times per second, and an oscillator running at 440 Hz does the same 440 times per second. The higher the Hertz value, the higher the pitch.

A common exception to oscillators having a pitch is a noise oscillator. While noise oscillators often have some tonal filtering options, most noise oscillators cannot be controlled in pitch since noise is essentially a signal containing all frequencies.

5.1.3 Envelopes

Envelopes are arguably the most common modulators for shaping the different signal amplitudes of a synthesizer over time. In their most common use, envelopes dictate how the amplitude of a sound behaves, but they can just as well be set to control other parameters of a synthesizer.

Many distinct types of envelopes exist, but perhaps the most straightforward of them is an AD envelope, which only consists of the Attack and Decay stages of an envelope signal.

Attack parameters control the amount of time it takes for a signal to reach its maximum after the signal is initiated. Decay controls how quickly the signal will fall back to its initial level after reaching the end condition of the previous stage. After the parameters are set in an AD envelope, both stages will unfold in sequence by triggering the gate of the envelope.

In sound synthesis, gates are control signals most often generated by the keyboard player pressing a key for a time of their choosing, or by a sequencer signalling the gate to open and close at specific times.

However, if one wishes to play sustained sounds, they require a different envelope. The simplest option is to use an AR envelope, which is otherwise identical to the AD envelope, but instead of proceeding directly onto the Decay, an AR envelope waits for the gate to be closed before entering its Release stage.

It is a functional option for many signals, but having more envelope stages can be of benefit, which is why most synthesizers tend to feature ADSR envelopes. On top of Attack, Decay and Release, an ADSR envelope allows its users to set the level of signal at which the envelope will Sustain it after the Decay stage.

Envelopes can be divided into as many stages as the manufacturer of the synthesizer wants to make available. The most complex envelopes are often referred to as Multi-Stage Envelopes. More on them in chapter 8.3.5.

5.1.4 Filters

Filters control the frequency content of the signals of synthesizers, most notably in audio form. Some attenuate or completely remove selected frequencies while leaving others untouched, others mangle signals in ways more elaborate than for what I have pages left for describing them.

Most filters used in synthesizers work by manipulating the phases of the signals sent into them (Reid 1999). By shifting the phase of the frequencies above or below a selected cutoff frequency, a filter can cause unwanted parts of the signal to cancel out. Depending on the filter, this phase shift may also lead to harmonic distortion of other frequencies due to overlap.

The basic filter types used in synthesizers are Lowpass, Highpass, Bandpass and Bandreject or Notch. Furthermore, Allpass is a common filter type, but it is less frequently used in the context of synthesizers.

Lowpass Filters allow the signal below a chosen cutoff frequency of the filter to pass through it. Sometimes referred to as Highcut Filter, this is the most widespread audio filter type in musical instruments. If a synthesizer sound loses its high frequency content and becomes ‘darker’, it is extremely likely that a Lowpass Filter of one type or another is being applied to it.

Highpass, Bandpass and Notch Filters work in a similar fashion, but all of them allow different frequencies to pass. A Highpass Filter allows the signals above the cutoff point to pass, a Bandpass Filter eliminates frequency content above and below the cutoff point and a Notch Filter removes the cutoff point from the signal, allowing signals both above and below it to pass through the filter.

Finally, Allpass is a filter type notably used in algorithmic reverbs. It only introduces a phase shift to the cutoff frequency of the filter but allows signals to pass through it completely. It sounds like a type of smearing effect around the cutoff frequency, and it can be particularly useful in designing soundscapes if one is about to run out of reverbs or other effects with which to smudge their creations.

5.1.5 Modulation

Modulation is arguably the most interesting part of audio synthesis. Turning one sound into another can be extremely rewarding, and hearing it happen in an instant because of one's own actions is a reward in and of itself.

To modulate a signal within the context of audio synthesis, one needs a system with a modulation source, or a modulator, and a modulation destination. Both can be either audio or control signals, but most commonly modulators are control signals, for example envelopes, and destinations are audio signals.

That being said, many contemporary synthesizers disregard these kinds of limitations and offer more comprehensive modulation feature sets, and hence an instrument with modulators capable of modulating other modulators is nowadays a common sight.

5.1.6 Low Frequency Oscillators

Low Frequency Oscillators, or LFOs, as they are usually abbreviated, are oscillators designed mainly as modulators instead of sound sources. Their cycle frequency is restricted to, as one may surmise, frequencies so low that they often remain inaudible. Instead, they can be used to sway signals in several ways.

For example, if an LFO is used to modulate the amplitude of a synthesizer, the sound can be heard growing stronger and weaker along with the cycles of the LFO. If an LFO modulates a Lowpass Filter, its slow oscillations that close and open the filter alternate between making the sound darker and brighter in turn.

While most LFOs are fixed to a set of basic waveforms, many contemporary synthesizers allow their users to also define the shapes of their LFOs. The simplest of waveforms suffice in most uses, but if one is looking for organic shapes and unpredictability, custom LFO waveform shapes may offer an answer.

5.2 Modular and Semi-Modular Synthesis

When freedom in designing synthesizer sounds is crucial, modular systems remain unmatched. The ability to select individual modules with which to generate and modify signals is powerful to the point of bordering on limitless.

With enough knowledge and a proper set of modules, a sound designer can indeed make almost anything sound like almost anything else, and as technology evolves, the types of sounds unattainable by synthesis diminish in numbers.

In a fully modular synthesizer system, any modules can be chosen, the modules can be set in any order, and any module can be set to modulate any other.

In a semi-modular system, some limitations apply. For example, the order in which the modules are set can be fixed, or not every module is allowed to modulate all the others. In virtual synthesis, semi-modular systems often have these kinds of limitations due to the exponential complexity of full modularity.

5.3 Virtual Synthesizers and DAWs

Virtual synthesizers are software components or applications specialised in generating sound by using the computing power of the device running them. They can be used either as dedicated executables, or as software modules called plug-ins within their host programs designed for audio and music production.

Several distinct types of host programs exist, but in the context of using virtual synthesizer plug-ins, the most common types are called Digital Audio Workstations, which is usually abbreviated to DAW.

DAWs are comprehensive software platforms designed as all-in-one solutions for audio production, composing, recording, editing, mixing, and mastering. Such programs include, but are not limited to, titles such as Bitwig Studio (Bitwig n.d.), Reaper (Reaper n.d.) or Renoise (Renoise n.d.), which are all examples of various kinds of DAWs with substantially different sets of features.

5.4 Virtual Synthesizer Formats

Several plug-in formats are available for different operating systems and uses.

The most common format for virtual synthesizers is the VST (Steinberg n.d.). It stands for Virtual Studio Technology and was created in 1996 by the German company Steinberg, who still license the technology today. The most advanced version of this plug-in format is the VST3, which was first released in 2008.

While VSTs are the most common and arguably the most popular audio plug-ins, alternatives to it also exist on different operating systems.

The most notable of these is the Audio Unit, used on macOS systems (Sweetwater 2022).

On Linux systems, one of the more prominent plug-in types is the LV2, which stands for Linux Audio Developer's Simple Plugin API Version 2 (LV2 n.d.). It is a royalty-free set of standards for audio plugins and host programs, and while it may not compare to VST in popularity because of its operating system limitations, technically it is relatively similar.

Finally, the most recent addition to the assortment of plug-in formats is the cross-platform, open source endeavour CLAP, which is short for CLever Audio Plug-In (Clever Audio Plug-In n.d.). Released by Bitwig Studio and u-he in 2022, CLAP surpasses its main alternative VST3 in both supported features and flexibility of licensing, making it an interesting option for plug-in users and developers alike.

Despite the many formats available, ultimately it is the instrument, its usability and its sound that matter considerably more when designing sounds.

5.5 u-he Zebra 2

Zebra 2 is a semi-modular virtual synthesizer first released October 9th in 2006. Its current version is 2.9.3, which will most likely remain its final release before Zebra 3 is published.

In my introduction, I highlighted some justifications for selecting Zebra 2 for this project, but they were not my only reasons for this choice. Perhaps most importantly, I was already familiar with the workings of Zebra 2. Having used it for over a decade, I knew its capabilities, and understood what kinds of sounds have previously been created with it. I knew how to get the most out of its selection of modules, and I knew how to work around most of the restrictions that its semi-modularity presents.

Zebra 2 is also versatile enough for creating soundscapes. While a fully modular system might be more powerful for sounds with these levels of complexity, none of the limitations of Zebra prevent me from creating anything toward which I strive.

While it makes little sense to go through all the features of Zebra 2 here, some of its functionality which I will be using may still benefit from an explanation.

5.5.1 The Modules

Compared to entirely modular systems, it may not offer the broadest selection of modules, but despite its age and the limits inherent to semi-modular synthesizers, this equid still has some kick in it.

The oscillators of Zebra 2 come in three varieties. The first of them are the wavetable oscillators. Every Zebra 2 patch can hold up to four wavetable oscillators, each of which can use one of four different oscillation modes. Up to sixteen waves can be arranged within a single oscillator, and up to two of twenty-six different spectral effects can be used in each oscillator. Each oscillator also has four unison options and dedicated parameters for amplitude, tuning, phase, panorama, and unison width, all of which can be modulated.

Unlike many contemporary wavetable synthesizers, Zebra 2 cannot directly convert PCM sounds into wavetables, but a workaround exists in the form of external converter applications. Alas, no user guides exist for the converters, so some trial and error may have to be involved in their utilisation.

The second of the three oscillator types of Zebra 2 are the FM oscillators. As with the wavetable oscillators, every Zebra 2 patch can hold up to four FM oscillators. Each of them uses one of four FM modes or a Ring Modulation mode, and each has eight waveforms available. The FM oscillators also feature the modulable parameters of amplitude, tuning, panorama, and stereo width.

The final oscillator type of Zebra 2 are its two noise oscillators, which are also commendable in their features. They have four operation modes, built-in Lowpass and Highpass filters, and modulable parameters for panorama, amplitude, and stereo width. There are not as many types of noises in Zebra 2 as there are in many newer virtual synthesizers, but especially for soundscape design they offer a solid foundation upon which one can build their spaces.

The flexibility of these oscillators already offers plenty of sound design possibilities as such, but where the possibilities are multiplied, is the filter section. The six filter modules and three Cross-Modulation filters are great for shaping sounds in many expected ways, but the Comb Filters of Zebra 2 are special. Comb Filters are named this way because when they are visualised, they form several peaks in the rough shape of a comb. They are filters often used in modelling of physical sound sources such as wooden or metallic textures, and they function by generating feedback to extremely short delay lines (Acoustic Sciences Corporation 2010).

In Zebra 2, the Comb Filters have several elaborate parameters for finetuning the relationships between the peaks of the filters. By modulating these parameters with several sources, sounds can be drastically morphed from one texture to another, which can give life to even the simplest of sounds but is often crucial to achieving a sense of motion in soundscape patches.

Another terrific way of introducing movement into a synthesized soundscape is to modulate the parameters of its waveshaper and saturation modules. In Zebra 2, these functions are handled by the modules Shape, Distortion and Fold. Most of their parameters can be modulated, and especially grimier sounds may benefit from the occasional bursts of distortion and screeching this can introduce.

5.5.2 The Three Lanes

The modules of Zebra 2 are divided into three lanes: Generators, Effects and Modulators. Generators and effects are routed in series, meaning that signals from the Generator lane can pass into the Effects lane. However, the only modules capable of modulating parameters in Zebra 2 are the ones that can be added onto the Modulators lane.

5.5.3 The Modulation Options

Parameter modulation can be achieved in Zebra 2 by several means. Most obviously, the internal modulators of Zebra can modulate most of the parameters of the synthesizer. These modulators include Envelopes, Multi-Stage Envelope Generators, Low Frequency Oscillators, Global Low Frequency Oscillators that are otherwise like the regular LFOs but affect all the voices playing within a patch, Modulator Mixers, and Modulation Mappers.

Parameters can also be modulated by DAW automation, whether it is accomplished by inserting data points or by using external modulators available in DAWs such as Bitwig Studio or Renoise. Furthermore, the parameters of Zebra can be modulated using MIDI data. Most notably, I used Aftertouch or Pressure and the Modulation Wheel in the creation of my soundscape patches.

Finally, Zebra 2 has four XY pads per patch. These are two-dimensional control interfaces that can modulate multiple parameters with varying amounts by moving a handle along the X and Y axes of a pad. These pads are ideal for transforming sounds from one state to another, especially if the user has access to a touch screen. I used all four pads for each of my soundscape patches.

5.5.4 The Delay and the Reverbs

The three Effects lane modules that allow the most spatial manipulation are the Delay, and the two reverberation effects Reverb and NuRev.

First, the Delay is yet another example of Zebra having far surpassed its contemporaries in features and flexibility. Two Delays can be added to the Effects lane, and each can be set to one of four delay modes that differ from each other in the ways in which they set up delay line seriality and their panoramas, filtering, and feedback loops (u-he n.d.a, 108–109). This offers immediate options for the ways in which the delays behave. All modes have modulable parameters for Lowpass and Highpass filters, feedback, cross-feedback from other delay lines, mixing and panorama. But what makes the delay module of Zebra interesting, are the parameters that allow the finetuning and modulation of the four delay times, even ones synchronised to the project tempo settings.

These parameters allow for a sound designer to set their delays in a manner which, especially with longer delay times and in combination with the reverb modules, can help create a sense of immense spaces within a sound. This proved beneficial in the creation of my final soundscape.

There are significant differences to the two reverberation modules of Zebra 2. The newer of them, NuRev, is an algorithmic plate type reverb with basic parameters for controlling its tone and the length of its decay. It sounds smooth without effort and overall is a fitting choice for creating a resemblance of an acoustic environment into a sound patch.

Reverb, the older of the two, tends to produce a more metallic tone compared to NuRev. With its broader array of parameters, Reverb can also be more challenging to operate than its newer counterpart. In particular, careless adjustments to feedback settings on Reverb can easily induce booming resonances in signals. However, given the crucial role of spatial processing in creating depth within one's soundscapes, combining both reverberation modules often proves to be the most favourable option.

6 Soundscape Design, Part I

In this chapter, first of the three soundscapes I created will be analysed. The main goal was to recreate one of Mark Morgan's most iconic soundscapes within a single Zebra 2 patch that could be played by holding only a single note.

6.1 Soundscape Goals

For the purposes of this thesis, I focused on discrete parts of musical compositions as soundscapes, simulated environments, and synthesized surrealism. I employed the ideas of keynotes and signals and applied something akin to soundmarks into my three designs.

Finally, despite the fabricated nature of my sounds, one of my key aspirations was to be able to suggest a sense of space with my soundscape design choices. And while absolute realism remained out of my reach because of the limitations inherent to audio playback, fortunately, evoking feelings does not require realism. It requires but just enough incentive for our imagination to fill in the blanks.

6.2 Placing the Wind

The first soundscape I approached is from Mark Morgan's composition Desert Wind (Banzay27 2012) from the Fallout soundtrack. It is a deceptively simple piece of ambient music that feels like the wind blowing through a bare desert. But amidst that wind, several minute details can be identified.

The natural deserts are one of the few environments of the world that would most likely undergo the least change during a nuclear apocalypse. Since deserts are typically characterised by their scarcity of resources, there would be little reason for any warring nation to directly bombard them. It is not impossible that someone would, perhaps driven by a desire to eliminate an oil field, but it is relatively unlikely. Carried by the wind, an all-encompassing nuclear fallout could still impact the safety of the already hazardous sand plains, although the landscape itself would likely remain unchanged if it were not a target in the bombardment.

It is that idea of radiation carried by the wind that I most associate with this particular soundscape. Judging by the end results it seems likely that Morgan did not try to replicate natural wind in this composition, for real wind does not sound this malignant. The gusts of this soundscape are grounded in reality but ultimately reject it to evoke surreal feelings of dread in the listener.

6.3 Planning the Wind

The windy textures are accompanied by pad sounds of similar nature that occupy the lower and higher mid-ranges of the frequency spectrum of the composition. They all sound synthetic, but without confirmation about Morgan's process, it is difficult to discern the exact methods with which they were created. Most prominent of the lower pads features a quiet fundamental around 147 Hz, a pronounced second partial and a series of upper harmonics in a gradually filtered series. The higher mid-range pads have significant variety to their instrumentation, but several of them are so complex that emulating them would quickly become a problem.

Furthermore, several distant clangs can be heard in the background of the composition. They can be samples, or they could even be synthesized, but regardless of their origin, it was beginning to look like there were too many elements to the composition before I had even started.

6.4 The Design Process

My first action was to set up a simple template for the project. I inserted a ModFX in chorus mode and a NuRev reverb module on the second Effects lane of Zebra 2 and then routed it into a Delay module on the third lane. This way, I could have at least rudimentary spatialisation from the get-go.

6.4.1 The Noise Oscillator

The shaping of the wind started with a Noise module. Since the wind textures heard in the original mostly occupy the higher mid-frequencies, I decided to filter most of the lower and middle frequencies with the built-in Highpass Filter of the Noise Oscillator. I opted for the stereo mode in the hopes of adding some side information into the dry signal.

I added a Comb Filter after the Noise Oscillator and found a fitting texture using the Dissonant mode, notable dampening, and negative feedback. The wind would still require some more filtering and nuance, but this was a promising start.

6.4.2 A Series of Failures

I turned my focus to creating the lower mid-frequency drone. I started by turning off the pitch tracking of Oscillator 1 and matching its tuning with that of the original drone to ensure identical results regardless of the note used to play the patch. Using the sidechain signal of a spectrum analyser, finding the proper volumes for different partials was a straightforward process using the additive features of the oscillators of Zebra 2. Modulating the detune parameter and both spectral effects of Oscillator 1 with a Modulation Mapper, I introduced movement to its tone. But despite the partials of the patch closely mirroring the timbre of the original, it was clear that the drone sounded different.

The first attempts to remedy this involved waveshaping, more noise, and even more oscillators. They all contributed to a sense weight and fullness, but no matter what I tried, I failed to replicate the lower drone exactly. I ended up with a sound that was interesting and full of movement but did not resemble the original.

Even though I now began to grasp the many challenges of recreating the drone, I was not ready to give up. I saved the patch for further use and modifications and proceeded to create a new one. This time, I would start with the drone and recreate the more straightforward wind textures after I had found its timbre.

The combination of a rich sawtooth wave and a Comb Filter seemed appropriate as an initial step and experimenting with the spectral effects of the first oscillator, I managed to limit its frequency content into a narrow band suitable for the imitation drone. However, setting up the Comb Filter proved more difficult than I had anticipated as the proper texture was eluding my attempts. The original drone sounds like it should be possible to recreate using a Comb Filter, but no matter how many combinations I tried, I could not produce one that would yield the results I was seeking within a reasonable time frame.

As daylight faded, several passes of decreasing effort fell through. I had been trying to decipher an equation to which I only had one of several variables, and the method of trial and error had not led me to its solution. Since the lower pads had proved this troublesome, I already shunned the notion of trying to replicate the more elaborate high-mid frequency pads. Having come so close to my goal only to realize how far away from it I still was, I felt disappointed. Regardless, it was time to simplify the soundscape.

6.4.3 A Simpler Wind

Even if Zebra 2 could produce both the pad sounds along with the wind textures within a single patch, using more time on them would be pointless without a new method with which to approach the issue. Still, the wind textures I had created the previous day were not too dissimilar to the ones of the original, so I decided to try again focusing solely on the sounds of the wind.

Newfound fervour accompanied this attempt from the beginning as the howls of the wind started to materialize. I once again started by limiting the frequency bandwidth of a stereo Noise module to the upper mid-frequencies of the spectrum. Most likely, this would be the only sound source needed for this patch.

I proceeded to add a Comb Filter onto the second Generator lane, and to my surprise found a fitting texture in the Blown mode, which I started to modulate with all four Modulation Mappers, all of which were modulated by LFOs running in Random Glide mode. Some light waveshaping and a Sideband module helped to add aggression and animation to the sound, and I finished tweaking its timbre by adding a Cross-Modulation Filter straight after the Noise module on the first lane, and a resonant Bandbass Filter on the second.

While the soundscape was off to a great start, it still felt like it lacked nuance. Routing the filter at the end of the second lane into a wavefolder on the third lane, I managed to give it a disparate texture. I applied a narrow Bandpass filter behind it cutting everything but the very highest of hisses from the signal. I further filtered and saturated the hiss with a Distortion, and finally routed the signal into the remaining Comb Filter, which I set to Dissonant mode. Positive feedback, dampening and tuning the filter up by sixteen semitones resulted in an eerie whistle, which fit the previous elements perfectly as a background element.

6.4.4 Affecting the Wind

Adding a Reverb module at the beginning of the effects chain felt like a natural way of adding a sense of place for the wind, but I quickly realised that I also wanted to be able to mix the dry signal in without it having to go through the entire chain, so I added a slow Compressor as the first module of the chain.

From the compressor, I divided the signal into its dry and reverberated parts. The dry parts would end up going through a Cross-Modulation Filter in Bandpass mode, a Mix module that I added solely for the purposes of panorama, and an equalizer that made four selective cuts, without which the signal had started to sound too full for its purposes.

The reverberated side of the signal ended up being simpler than I expected. After the Reverb, I added an equalizer that I used to separate the reverberated signal from the dry by lifting it around 1000 Hz. The trusty combination of ModFX in chorus mode and NuRev tied the sound together. For stereo imagining, simplified delay lines on the third lane proved most fitting for the source, and anything more elaborate only risked turning the entire soundscape into an aimless concoction of noise. I finished the timbre by adding a Distortion module before the Delay to further separate the echoing signal from its source.

Finally, I spent some time boosting the amplitude of the patch, and added modulation options for all four XY pads, Modulation Wheel and Aftertouch. And with these final adjustments, the first soundscape was finished.

6.5 Afterthoughts

Despite the many drawbacks that I had to face during the design process of the first soundscape, I am pleased with the results. It is often best to move on and try again later if some aspect of a project refuses to work out, and this time it led to me creating a usable sound that shares plenty of features with its influencing design. It may not be as rich in detail as I would have liked, but it is finished.

7 Soundscape Design, Part II

The second soundscape I created was my interpretation of Mark Morgan's equally iconic City of Lost Angels, more specifically the version heard in the prologue of Fallout 2 (The Fallout Wiki Channel 2014). For the most part, it is identical to the composition heard in the original Fallout game, but a handful of new sounds were added to it to emphasize certain moments in the speech read by Ron Pearlman.

7.1 Placing the City

Despite the surreal nature of the composition, several elements that parallel our reality can be identified within its soundscape. First of these is, again, the howling of the wind. Compared to the full textures heard in Desert Wind, the keynote sounds of City of Lost Angels sound hollower and haunted. The entire composition sounds like the lamenting of the billions of lives lost to nuclear annihilation, but especially its many whining pads feel like witnessing a great evil.

The second surreal soundscape element related to our world are the many percussions signals heard in the composition along with the pad sounds. They have a steady, slow rhythm to them, and within the context of the otherwise ghostly soundscape the sparse floor tom, djembe and bongos feel like they represent the remaining people of a city desperate to rebuild itself.

As its name suggests, City of Lost Angels is heard in Fallout when the player character enters a small community surviving in the ruins of Los Angeles. The area is but a husk of its former self, and the slow rhythms of the composition symbolize a way of life that has had to wind down since the apocalypse. Gone is the constant traffic, and every facility in the industrial district has fallen. Not even the hum of electricity remains. Building on this imagery, the entire composition becomes a symbol for the modest rhythms of the new world setting the pace on top of the echoes of the old in this masterpiece of a composition.

7.2 Planning the City

Before I even began, it was clear that I would not be able to create this soundscape within a single Zebra 2 patch. There are far too many separate elements in the composition for me to effectively modulate Zebra between the different required states to represent them all, and there are not enough modules or lanes to create and trigger them all individually, either.

Because I would have to separate the elements into discrete patches, this meant that the DAW would have to share some of the responsibilities of Zebra during the creation of this soundscape. With this approach, I would have more flexibility and available options, but with each sound getting a patch of its own, the design process would also consume much more time.

7.3 The Design Process

As there are many more elements to the soundscape of City of Lost Angels, I will also refrain from describing the design process in as much detail as in the previous chapter and will only delineate the methods I used in a broad manner.

Since the eerie sound of the wind is the keynote sound most central to the composition, it felt like the appropriate choice for the first sound I would design. Once again, I started by filtering and comb filtering a noise oscillator, but quickly added a second one and a sine wave oscillator as this would simplify the process.

After some equalization and chorusing, I used the two reverb modules and one of the two delays to finish the first sound patch of my adaptation.

The second sound I designed was the prominent floor tom heard throughout the composition. First, I created a MIDI pattern that mirrored the original and started to adjust two oscillators that would form the backbone of the sound. The first one created the transient of a mallet hitting a membrane, and the second one was a simple sine wave that gave the drum its fundamental note. A comb filtered noise oscillator added most of the timbre, and some lowpass filtering and compression tied the sound together. For final additions, I used only NuRev and one Delay since anything more was not necessary.

Two additional percussions were created with a similar method. A sine wave would form a fundamental note, upon which a comb filtered noise would lay a fitting texture. After some light saturation, further filtering, and reverberation the sounds were finished. I could not get the ostinatos they play in the original quite right, but at least some of the intent is still conveyed in their patterns.

Next, I turned my attention to the snare-like drum sound of the original that accents the second and the fourth beat of each bar. I do not think it is an actual snare drum, but it does not quite sound like a tom drum either. Whatever it was, my adaptation needed something of similar nature.

Finding the fundamental note was again the first phase of my design process. I recorded a snare drum hit, with the snares turned off, and turned the result into a wavetable that barely resembled a drum and routed it into an FM oscillator with a simple sine wave, which was modulated by an envelope. With the help of a waveshaper, this added a fitting snap to the transient of the drum. A filtered and comb filtered noise oscillator once again formed the mid-frequency content of the drum, which I further accented by adding a second comb filter. Finally, I added another noise oscillator in the hopes of adding even more snap to the higher frequencies of the drum sound. I could not quite replicate the upper frequency tones of the original, but they were close enough. After some equalization, compression, reverberation and delaying, the patch was finished.

It was time to create the eerie hums and shrieks of the original composition. The first of them would prove difficult, for the original is a soft sound of a male choir, and even the most basic human voices can be relatively complex to synthesize. For the body of the sound, I used a noise oscillator that I routed into a resonant Formant Filter. Further filtering, equalization and reverberation resulted in a playable pad sound that did have some human qualities because of its highlighted formants. It was not a close match with the original, but again, it was close enough for the purposes of this project. I was pleasantly surprised to find that the sound could also be played an octave higher to support the wind textures.

The soundscape was starting to come together. For the final element of my adaptation, I decided to recreate the higher mid-frequency shriek that was added to the version heard in the prologue of Fallout 2. I used a saturated sine wave with a touch of noise running through a comb filter as the body of the sound, and then used all four filters in various modes to hone its tone. For the effects, I used two choruses, another filter, two equalizers, two reverbs and a delay to create a sense of otherworldliness. Considering how significant the original sound is, my adaptation succeeded surprisingly well in reaching its spirit.

There were still elements that I had not figured out in the original soundscape of City of Lost Angels, but I was beginning to notice the limitations of my skills. My adaptations were close to the original sounds, but to create exact matches, I would still have needed more time, of which I had little left.

7.4 Afterthoughts

I am unsure about how to feel about this second soundscape. While I appreciate many of the individual elements I created for it as playable sounds in their own right, I cannot help but feel like I would have needed more time to reach the nuanced nature of the original. My adaptation does resemble the original enough for the influences to come across, but comparing the two side by side, it is clear that they are different soundscapes altogether.

Regardless of the results, the method itself was a viable option that allowed for most control in creating an immersive space within a soundscape. However, synthesizing individual elements may end up taking more time than working with recordings, and the endless finetuning of details is a pitfall upon which one may stumble. But if enough time is available, the method is worth considering.

8 Soundscape Design, Part III

Finally, it was time to design a post-apocalyptic soundscape of my own.

8.1 Placing the Soundscape

Since this soundscape was an original design, there was plenty of creative freedom to figuring out what kind of an environment could have produced it. Upfront, I only knew that I wanted distant clanking of metal in it, but after some contemplation, a location built around that soundmark started to suggest itself.

The place I imagined was once a small suburbia bustling with life. A gentle summer breeze would have swished through the trees growing by the avenues, birdsong would have filled the air, and only infrequent cars would have purred as they pulled to the driveways. Laughter and juvenile squabbles of children playing outside would have carried in through the windows, and the serene rhythm of family life would only have been disrupted each weekday morning by mechanical alarm clocks, their chimes having signalled the start of another workday. On Sundays, the bells of the local church would have tolled, and people from across the area would have gathered in service.

Walking on cracked tarmac, only the faint idea of a neighbourhood can be perceived. The stalks of dead trees stick out from the ashes of the houses and the people who once lived in them. The cars have not moved in decades, and even the wind seems to have halted in this graveyard. The only signals are imaginary, as the distant church bells still cry their detuned alarms to a community that would not react in time.

8.2 Planning an Original Soundscape

With this final soundscape, I wanted to take Zebra 2 and my skills to their limits and create a soundscape that would feel as full as possible, all within a single patch played only by holding a single note. It would be tricky, but it could be done.

8.3 The Design Process

I began designing the final patch with only the concept of some kind of imaginary bells as a soundmark. Furthermore, I wanted a slowly transforming drone as its pivotal keynote sound. I wanted the timing of the patch to be as random as possible, and I wanted it to feel oppressing yet constantly captivating. Everything else I would figure out as the obstacles would present themselves.

8.3.1 Creating a New Wavetable

With my first oscillator, I wanted to approach an organic drone texture, so I picked up my small djembe and started brushing it with the back of my hand close to my Røde NT1-A microphone recording into Bitwig Studio. The idea was to create a short pattern that would have just enough nuance and movement yet little in terms of dynamics. I soon found a fitting phrase and began processing it.

While I knew how little detail of the already minute sound would translate into the Zebra wavetable, it still needed to be much louder. With a simple amplitude boost and saturation, the loop found a style that suited it, after which I cut it into sixteen parts for wavetable transformation. When exported to Zebra, the wavetable sounded nothing like the loop, but at least it had enough nuance for a drone.

8.3.2 Finding the Keynote Sound

After setting up a basic reverb to delay signal chain and doing two prompt tone tests, I found something interesting on the third attempt. Playing with the spectral effects of the oscillator, I created a drone sound that reminded me of Mark Morgan’s Fallout composition Vault of the Future without feeling like its copy.

As with the first soundscape, the allocation of available effects would also be the main issue I would face creating this patch. While full of potential, the drone needed filtering because of its sizzling top end and its booming low end, but I would also need filters for other purposes, so I was facing inevitable compromises. However, since this drone was to become the defining keynote sound of the soundscape, I decided to use two of the four regular filters and one of the two cross-modulation filters of the Generator lane to forge its tone.

The regular filters I set to four-band Allpass and resonant Bandpass respectively and gave both some movement with two LFOs and two Modulation Mappers. After some brief saturating, it was time to move onto the soundscapes signal sounds, namely the metallic clanging.

8.3.3 The Metal Signals

I wanted these clangs to be as random as possible within reason, so using the Crackle mode of the Noise module was a natural choice. The Crackle mode is great if one wants to be able to control the sparsity and tone of the noise grains, and if they then patch these crackles onto a Comb Filter set to create a metallic resonance, they will become randomly generated percussions the tone of which can be moulded by setting or modulating just a few parameters.

Contemplating between bell-like tones and a percussion sound not unlike a djembe, I ended up choosing the former because it better fit the ominous drone. The bells did not need much finetuning, but I ended up using the second cross-modulation filter to further saturate and resonate them.

The soundscape was already coming to its own, but it did not feel finished just yet. I had three wavetable oscillators, one Noise module, one Comb Filter and two regular Filter modules left, and the bottom end of the soundscape felt like it could still use some bolstering.

8.3.4 The Low Frequencies

I spent some time creating and tweaking a djembe-like percussion much larger than its real-world origin, but after two hours of sculpting, I instead chose to turn it into a lower drone sound to better fit the mood of the soundscape. This was the right decision, since immediately after I turned the percussions into a low-end drone, I started to feel like the sound was beginning to near its completion. However, it still needed more movement.

8.3.5 Resetting The Multi-Stage Envelope Generators

The Multi-Stage Envelope Generators that I had set to create the percussion rhythms were easy to convert into swaying drones by removing or lowering some of the envelope peaks and lowering their angles. At first, they were all running in sequences of equal lengths, but I ended up using the third MSEG to modulate some parameters for extra movement twice for every longer MSEG sequence.

For this purpose, I opted to use MSEGs instead of LFOs or Modulation Mappers because with MSEGs I could both have the low-end drone move with nuance but also design recurring moments of drama into it. This felt like the proper way of introducing structure into the otherwise fairly randomised soundscape.

8.3.6 The Finishing Touches

After the envelopes were set, it was time to move onto the Effects lane. The first order of business was to finetune the tones of the reverbs to my liking. I had hastily set both up when I was working on the first oscillator, but some more work had to be done for them to better fit the soundscape. Mostly it was a matter of finding the right frequency response, but I also ended up switching the place of the Rev module, further reverberating the main drone before NuRev.

The bells were processed with the final Delay module and some equalization, after which I finalised the sound by allocating the XY pads, setting up slight Aftertouch modulation and making the Modwheel work as a transition tool that dulls the sound for fading purposes. I tinkered with several nuances, but one problem I still had to solve was the overall amplitude of the patch. Despite having boosted it on several occasions, the volume level had dropped much lower than I would have liked.

The final volume boost would have to happen right before the final reverb and delay, but since I was out of compressor modules, I turned to the waveshaper module Shaper. By only boosting its output volume, I was able to lift the sound up to a comfortable soundscape level of -30 Loudness Units relative to Full Scale without much distortion to the sound. And with that, the sound was finished.

8.4 Afterthoughts

Overall, I am extremely pleased with the results of this final soundscape. It turned out almost exactly like I wanted, and it is simply a pleasure to listen to. Full of nuance, intrigue, and ill will, it paints a picture so vivid that if I did not know better, I could very well believe that it belonged to the original Fallout soundtrack.

The greatest challenge I faced while designing this patch was the one I had issued to myself. Even if the technical limitations of Zebra 2 may have gotten in the way on some occasions, this time they were not insurmountable. The soundscape would still benefit from proper mixing and mastering, but as an illustration of what Zebra 2 is capable, I feel like the patch succeeds beautifully.

As was demonstrated in chapter six, the method of creating a complete soundscape inside a single patch file often requires the sound designer to compromise, but if the goal is to create complete-sounding patches for a sound bank, because of the simplicity of having everything within a single patch file, this is arguably a better method than expecting the user to combine different patches.

A one-note soundscape patch such as this is a perfect example of a sound that would also likely benefit from being created in a comprehensive modular environment, but this proves beyond all doubt that, within reason, immersive soundscapes can also be created using a synthesizer with a fixed set of modules.

9 Conclusions

Soundscapes both captivating and convincing can indeed be created using audio synthesis. It may not be the most efficient way of designing them, and the method may not fit all purposes, especially if utmost realism is required, but where a pinch of surrealism is allowed, synthesizing a soundscape can be rewarding and provide results that bear their creator's unique signature in their echoes.

Furthermore, there is considerable potential in layering synthesized soundscape elements alongside recorded ambience tracks to accentuate specific story beats, emotions, or settings without using more pronounced music. These elements can be subtle and natural-sounding to the point of almost being unnoticeable yet emphasize any recorded elements in countless of ways and contexts, but as they are synthesized, overt signal sounds are a valid option as well.

Since synthesizer sounds can be modified and tweaked to fit any purpose, they offer far more total control of their individual elements in ways that cannot be matched by regular PCM files. A synthesized soundscape that is specifically created to fit a project is hence only limited by the skills of its designer, by the tools of their choosing and by the time they can spend on creating it.

Synthesizing a soundscape full of subtlety and detail can take several hours, but especially in a virtual environment, once the patch has been created, recalling it takes no more time than inserting a sample onto the project timeline, and with only a few tweaks it can sound different enough to justify its use in another project.

Choosing the right synthesizer for a project is an important part of the process, and one to which I have no answers apart from gradually gaining experience by learning to use synthesizers in general, and by then trying out many different instruments until personal favourites are found. For this project, Zebra 2 was a natural choice because I was already familiar with it. However, I am certain that the process could have been replicated with many other synthesizers with modular or semi-modular architectures. The comb filters of Zebra are unique and powerful, but they are not necessary for synthetic soundscape design. But would I have created the exact same sounds had I used another synthesizer? Not likely.

The idea of unlimited modules can sound appetizing, and many have built modular synthesizers the size of walls in search of this freedom, but choosing a tool with limitations can inspire creativity in a way that may not seem apparent if one has access to an option where no such limitations exist. While working on this project, there were several times during which I found myself double and triple checking the popup menus of Zebra in the hopes of finding one more filter or another reverb module I had neglected to notice. No more modules were left, but the foolish hope often remained, and the bitter realization of having to compromise in my designs stung every time. However, had I been able to simply add more modules whenever I ran into an obstacle, I also would not have gotten to solve these particular problems.

Because I had to face these challenges, I got to surprise myself with several moments of success. Because of these limitations, when I managed to create a particular element with limited modules, or especially when I succeeded in creating my own post-apocalyptic soundscape that can be played just by holding one note, it felt like I had solved a problem with my skills and perseverance instead of merely throwing additional tools at it.

Experience is gained when one is faced with having to overcome various kinds of obstacles enough times to know how to approach and ascend them and even those adjacent to them. Being able to surpass an obstacle by always adding more is not the most fruitful of solutions, because if a situation then presents itself where the available tools are limited, more problems may surface as a result. Suddenly one would need to develop a new set of skills to which they may not have been prepared, or to which they may not have enough time. A sound designer or other artist with enough experience, and hence, with a diverse enough skill set, will also succeed when ‘more’ is not an option.

To some degree, all art entails imitation. We tend to most appreciate artists whose personal styles shine in their work on top of sheer skilfulness, but without influence from others, we would have trouble creating art in the context of our culture. Imitation can lead to growth, for by imitating others we can learn new techniques, and even new skills altogether. And as we grow more comfortable with our tools and technique, our own individual touches will start to show.

When I started making music, one of my great wishes was to be able to create the kind of music Mark Morgan had composed for Fallout. That soundtrack enamoured me with its immense sense of place and countless of sound design choices that still tickle my imagination. Back then, I had no idea about how to even approach creating something like it. Now, over twenty-four years later, I am much more confident in my capacity to produce similar music. But what is of more interest to me, is developing my original ideas, and cultivating a voice of my own. Because at least for me, that is how the most meaningful works of art are created.

Fallout did not start as the giant it is today. At first, it was a handful of people working after hours with only pizza as their reward (GameSpot 2012). It was Tim Cain’s passion project that first resonated with his co-workers and then an audience large enough to warrant a sequel. The rest is history because of that passion. What began as a dream and willingness to work for it became a world of desolation that has brought joy for millions.

In a way, my willingness to still spend time analysing the soundscapes of its music all these years later is proof of even small passion projects being capable of profoundly moving others and changing their lives for the better. As a personal truth, it proves that creating art, in one form or another, is a worthwhile cause for as long as something that resembles humanity exists.

Our world has not ended, yet, and while we are still here, by crafting art with meaning we can reach out and join our voices into its infinite soundscapes.

References

ISO 12913-1 2014. Acoustics — Soundscape. https://www.iso.org/obp/ui/#iso:std:iso:12913:-1:ed-1:v1:en (referenced October 11, 2023).

Clegg, Brian 2010. Armageddon Science: The Science of Mass Destruction. New York: St. Martin's Griffin.

Hanna, Pierre & Louis, Nicolas & Desainte-Catherine, Myriam & Benois-Pineau, Jenny 2004. Audio Features for Noisy Sound Segmentation. Talence: Université de Bordeaux. https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=f41f9f17729c15c0927ba4b11c921b9eb874d7d3 (referenced November 10, 2023).

Bitwig n.d. Bitwig | Home | Bitwig. Webpage. Bitwig. https://www.bitwig.com/ (referenced November 13, 2023).

CLAP n.d. Clever Audio Plug-In. Webpage. CLAP. https://cleveraudio.org/ (referenced November 13, 2023).

Roby, David Allen 2013. Crust Punk: Apocalyptic Rhetoric and Dystopian Performatives. Texas: Texas A&M University. https://oaktrust.library.tamu.edu/bitstream/handle/1969.1/151027/ROBY-THESIS-2013.pdf (referenced November 12, 2023).

Bethesda n.d. Fallout. Webpage. Bethesda. https://fallout.bethesda.net/en/games/fallout (referenced November 12, 2023).

Nukapedia Fallout Wiki 2023. Fallout. Webpage September 17, 2023. https://fallout.fandom.com/wiki/Fallout (referenced November 13, 2023).

The Fallout Wiki Channel 2014. Fallout 2 Intro Narration. Online video January 8, 2014. YouTube. https://youtu.be/A4qcb_6RVX4 (referenced November 13, 2023).

Kollar, Philip 2015. Fallout 4 could be a bigger hit than Skyrim. Webpage November 10, 2015. Polygon. Archived at web.archive.org. https://web.archive.org/web/20151113201845/http://www.polygon.com/2015/11/10/9673936/elder-scrolls-bigger-than-fallout-sales-data-report (referenced November 10, 2023).

Makuch, Eddie 2017. Fallout 4 Surpasses Skyrim To Become Bethesda's Most Successful Game Ever. Webpage February 8, 2017. GameSpot. https://www.gamespot.com/articles/fallout-4-surpasses-skyrim-to-become-bethesdas-mos/1100-6447621/ (referenced November 10, 2023).

GameSpot 2012. Fallout Classic Revisited. Online video March 9, 2012. YouTube. 02:55. https://www.youtube.com/watch?v=Xa5IzHhAdi4&t=175s (referenced November 10, 2023).

Petski, Denise 2023. 'Fallout' TV Series From 'Westworld' Creators Based On Games Gets Prime Video Premiere Date. Webpage October 23, 2023. Deadline. https://deadline.com/2023/10/fallout-tv-series-premiere-date-westworld-creators-video-game-trailer-1235580537/ (referenced November 10, 2023).

Huang, Andrew 2018. FM SYNTHS in under 4 minutes. Online video February 22, 2018. YouTube. https://youtu.be/vvBI3YUBUyY (referenced November 12, 2023).

Acoustic Sciences Corporation 2010. Acoustic Basics. Webpage. Acoustic Sciences Corporation. Archived at web.archive.org. https://web.archive.org/web/20100507124237/http://www.asc-hifi.com/acoustic_basics.htm#2 (referenced November 13, 2023).

Jordan, Ben 2017. How Additive Synthesis Works. Online video October 14, 2017. YouTube. https://youtu.be/wGkdb6YlLgg (referenced November 12, 2023).

Jurassic Park, 0:36:07–0:36:11. USA 1993. Writers Crichton, Michael & Koepp, David. Sound Designers Rydstorm, Gary & Scarabosio, Christopher. Composer Williams, John. Director Spielberg, Steven. Amblin Entertainment, Universal Pictures: Kennedy, Kathleen & Molen, Gerald R. & Ryan, Lata & Wilson, Colin. Blu-ray. 127 min.

LV2 n.d. LV2. Webpage. LV2. https://lv2plug.in/ (referenced November 13, 2023).

Banzay27 2012. Mark Morgan – Vault Archives [Fallout music]. Online video November 1, 2012. YouTube. 01:13:42. https://youtu.be/_cLRP-QQxn4&t=4422 (referenced November 13, 2023).

European Environment Agency 2023. Noise. Webpage October 4, 2023. European Environment Agency. https://www.eea.europa.eu/en/topics/in-depth/noise (referenced November 11, 2023).

Reid, Gordon 1999. Of Filters & Phase Relationships. Webpage August 1999. Sound On Sound. https://www.soundonsound.com/techniques/filters-phase-relationships (referenced November 12, 2023).

Steinberg n.d. Our Technologies. Webpage. Steinberg. https://www.steinberg.net/technology/ (referenced November 13, 2023).

Eduskunta Riksdagen n.d. Parliament's buildings and art. Webpage. Eduskunta Riksdagen. https://www.eduskunta.fi/EN/naineduskuntatoimii/eduskunnan_talot_ja_taide/Pages/default.aspx (referenced November 13, 2023).

Moon, Hyong-jun 2014. The post-apocalyptic turn: A study of contemporary apocalyptic and post-apocalyptic narrative. Milwaukee: The University of Wisconsin. https://www.proquest.com/openview/794f4b9acd5910d2b081361357e2d37a/ (referenced November 12, 2023).

Nugent, Colin & Blanes, Núria & Fons, Jaume & Sáinz de la Maza, Miquel 2016. Quiet areas in Europe. Luxembourg: Publications Office of the European Union. https://www.eea.europa.eu/publications/quiet-areas-in-europe (referenced November 10, 2023).

Reaper n.d. REAPER | Audio Production Without Limits. Webpage. Reaper. https://www.reaper.fm/ (referenced November 13, 2023).

Miller, Jake A 2020. VJ-Day 75th Anniversary. Webpage August 11, 2020. National Archives. https://unwritten-record.blogs.archives.gov/2020/08/11/vj-day-75th-anniversary/ (referenced November 23, 2023).

Renoise n.d. What Is Renoise? Webpage. Renoise. https://www.renoise.com/products/renoise (referenced November 13, 2023).

Truax, Barry 1999. Sound Synthesis. Webpage. Cambridge Street Publishing. https://www.sfu.ca/sonic-studio-webdav/handbook/Sound_Synthesis.html (referenced November 13, 2023).

Truax, Barry 1999a. Sawtooth Wave. Webpage. Cambridge Street Publishing. https://www.sfu.ca/sonic-studio-webdav/handbook/Sawtooth_Wave.html (referenced November 13, 2023).

Kristensen, Hans & Korda, Matt & Johns, Eliana & Kohn, Kate 2023. Status Of World Nuclear Forces. Webpage. Federation of American Scientists. https://fas.org/initiative/status-world-nuclear-forces/ (referenced November 11, 2023).

Underdog Electronic Music School 2023. Synthesis for beginners. Online video February 22, 2023. YouTube. https://youtu.be/eu0zpa7OiYA (referenced November 12, 2023).

Stolet, Jeffrey 2009. Square and Rectangle Waves. Webpage. University of Oregon. https://pages.uoregon.edu/emi/14.php (referenced November 13, 2023).

Mäki, Jonas 2021. Todd Howard: Fallout 76 is one of our most played games with 11 million players. Webpage November 11, 2021. Gamereactor. https://www.gamereactor.eu/todd-howard-fallout-76-is-one-of-our-most-played-games-with-11-million-players/ (referenced November 10, 2023).

Schafer, R. Murray 1977. The Tuning of The World. New York: Knopf.

Reverb 2018. Wavetable Synthesis Explained | Reverb. Online video December 12, 2018. YouTube. https://youtu.be/_w1h4vknfow (referenced November 12, 2023).

Sweetwater 2022. Which Plug-in Format Do I Need For My DAW? Webpage November 17, 2022. Sweetwater. https://www.sweetwater.com/sweetcare/articles/which-plug-format-need-for-my-daw/ (referenced November 13, 2023).

The Editors of Encyclopaedia Britannica 2023. White Noise. Webpage September 28, 2023. Britannica. https://www.britannica.com/art/white-noise-music (referenced November 9, 2023).

United Nations 2022. World Population Prospects 2022. Webpage. United Nations Department of Economic and Social Affairs Population Division. https://population.un.org/wpp/Graphs/Probabilistic/POP/TOT/900 (referenced November 12, 2023).

Weisman, Alan 2007. The World Without Us. New York: Picador.

u-he n.d.a. Zebra 2 User Guide. Operating documentation. u-he. https://uhe-dl.b-cdn.net/manuals/plugins/zebra2/Zebra2-user-guide.pdf (referenced November 23, 2023).

u-he n.d. Zebra Legacy. Webpage. u-he. https://u-he.com/products/zebra-legacy/ (referenced October 11, 2023).

Analysed Media

Fallout. USA 1997. Tim Cain. CD-ROM. Interplay 1997.

Fallout 2. USA 1998. Feargus Urquhart. CD-ROM. Interplay 1998.

Appendices

Appendix 1. A mother and her child, survivors of Nagasaki. (Yamahata, Shogo. https://www.independent.co.uk/news/world/asia/nagasaki-one-day-after-the-atomic-bombing-seen-in-newlydiscovered-pictures-9472178.html.)

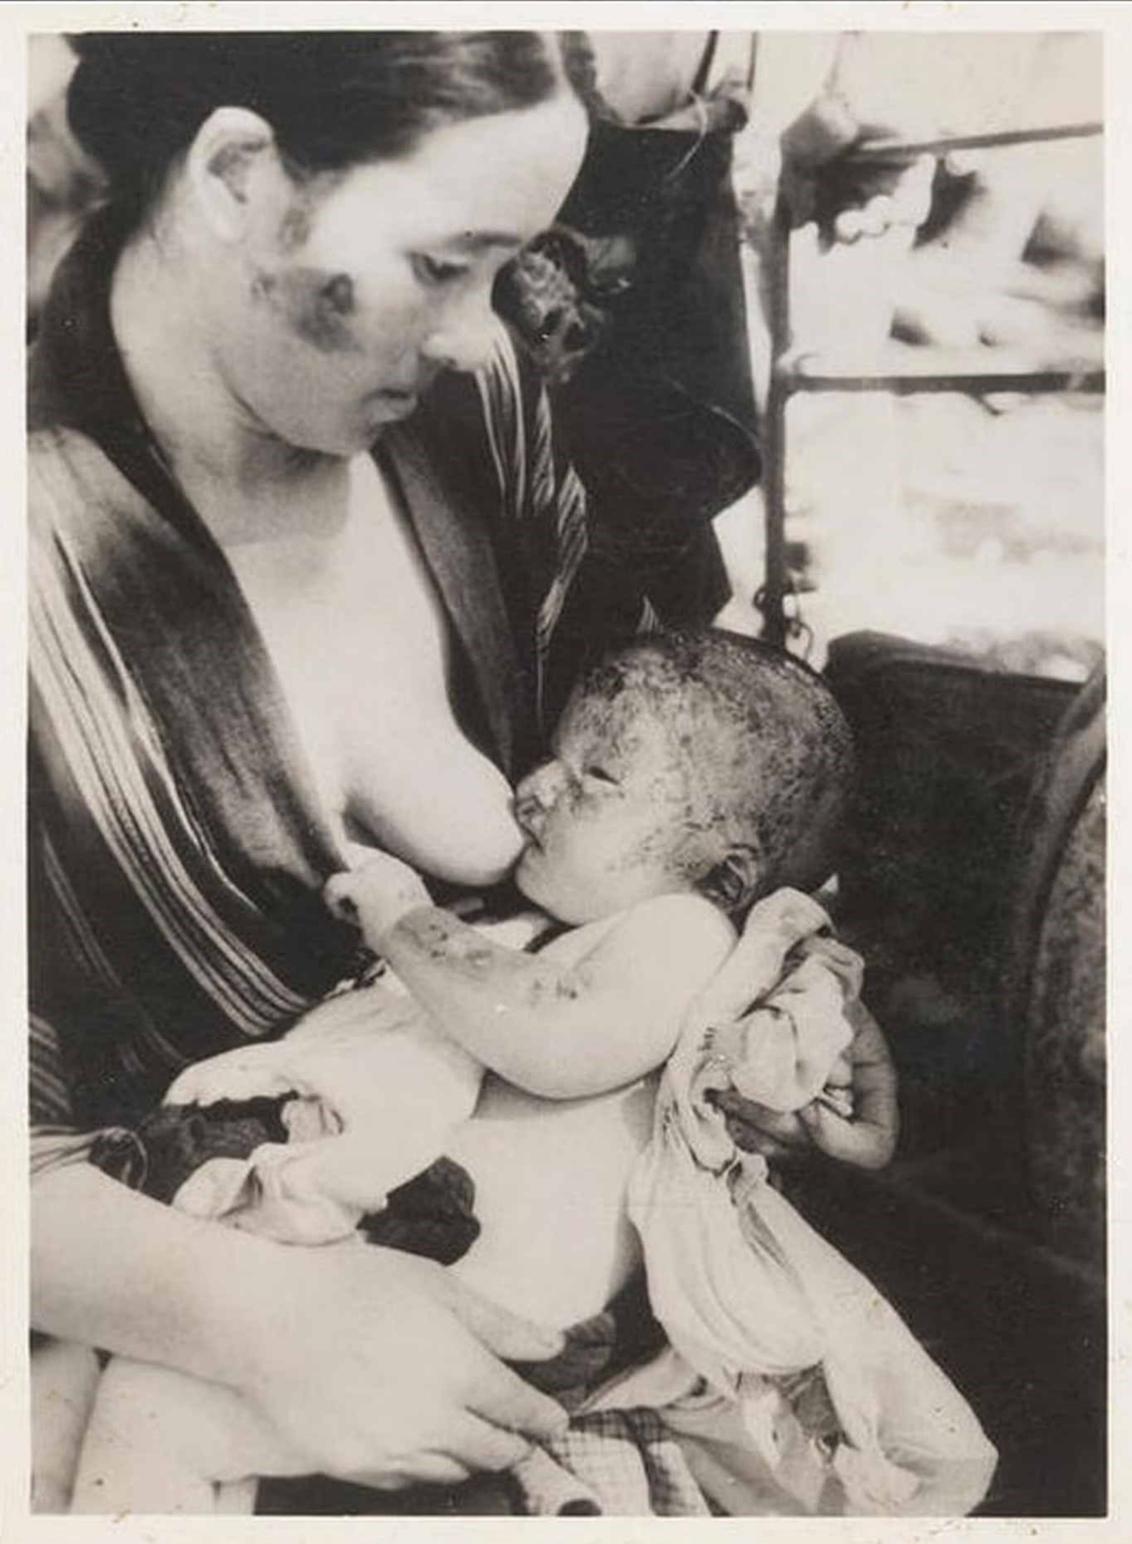

Appendix 2. A surviving mother nursing her child after the Nagasaki bombing. (Yamahata, Shogo. https://www.independent.co.uk/news/world/asia/nagasaki-one-day-after-the-atomic-bombing-seen-in-newlydiscovered-pictures-9472178.html.)

Appendix 3. A 'shadow' of a person who shielded a part of the steps as they were disintegrated in the destruction of Hiroshima. (Universal History Archive/UIG via Getty Images. https://allthatsinteresting.com/hiroshima-aftermath-pictures#6.)

Appendix 4. Buildings reinforced to withstand earthquakes were the only ones still standing after detonation. (Imperial War Museums via Getty Images. https://allthatsinteresting.com/hiroshima-aftermath-pictures#10.)

Appendix 5. A man looking over the ruins in Hiroshima. (Getty Images. https://allthatsinteresting.com/hiroshima-aftermath-pictures#11.)

Appendix 6. A picture of Hiroshima taken from air in 1946. (Library of Congress. https://allthatsinteresting.com/hiroshima-aftermath-pictures#23.)

Appendix 7. An aerial photograph of Nagasaki after the bombing. (RWU Archives and Special Collections/Digital Commons. https://www.washingtonpost.com/news/worldviews/wp/2015/08/09/what-nagasaki-looked-like-before-and-after-the-bomb/.)

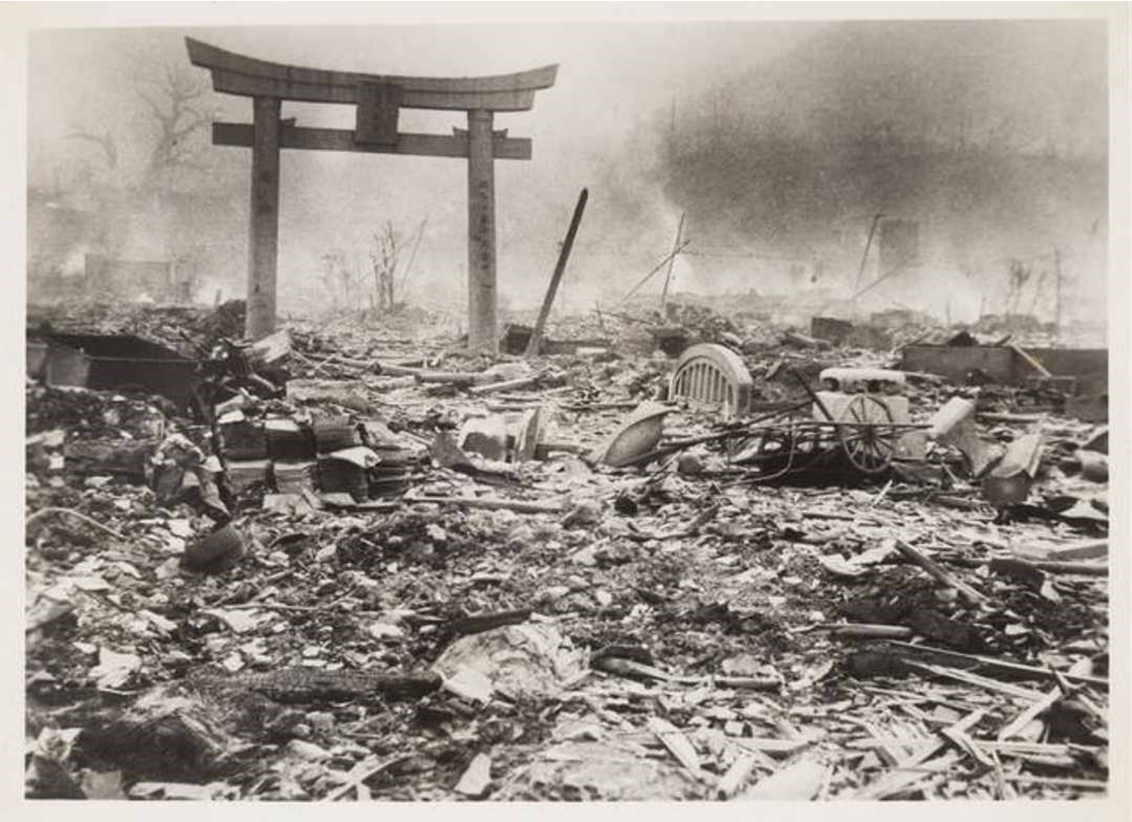

Appendix 8. An archway standing in the ruins of Nagasaki. (Yamahata, Shogo. https://www.independent.co.uk/news/world/asia/nagasaki-one-day-after-the-atomic-bombing-seen-in-newlydiscovered-pictures-9472178.html.)

Appendix 9. Surviving people traversing the ruins of Nagasaki. (Yamahata, Shogo. https://www.independent.co.uk/news/world/asia/nagasaki-one-day-after-the-atomic-bombing-seen-in-newlydiscovered-pictures-9472178.html.)