The Sound of Decentralization – Sonifying Computational Intelligence in Sharing Economies

Marinos Koutsomichalis[1], Evangelos Pournaras[2]

[1] Department of Computer Science
Norwegian University of Science and Technology
Trondheim, Norway
E-mail: me@marinoskoutsomichalis.com

[2] Professorship of Computational Social Science
ETH Zurich
Zurich, Switzerland
E-mail: epournaras@ethz.ch

Abstract

Pervasive technologies in socio-technical domains such as smart cities and smart grids question the values required for designing sustainable and participatory digital societies. Privacy-preservation, scalability, fairness, autonomy, and social-welfare are vital for democratic sharing economies and usually require computing systems designed to operate in a decentralized fashion. This paper examines sonification as the means for the general public to conceive decentralized systems that are too complex or non-intuitive for the mainstream thinking and general perception in society. We sonify two complex datasets that are generated by a prototyped decentralized system of computational intelligence operating with real-world data. The applied sonification methodologies are largely ad-hoc and address a series of concerns that are of both artistic and scientific merit. We create informative, effective and aesthetically meaningful soundworks as the means to probe and speculate complex, even unknown or unidentified, content. In this particular case, the sonification represents the constitutional narrative of two complex application scenarios of decentralized systems towards their equilibria.

Keywords

Data, data aesthetics, sonification, decentralization, optimization, Smart Grid, Smart City, sharing economy.

Introduction

Following the 'datalogical turn' of the last few decades [Kitchin, 2014; Chandler, 2015; Clough et al., 2015] several types of data aesthetics have been well laid out, including 'diagram aesthetics' [Heinrich, 2016], visualization/materialization practices [Bjørnsten, 2016], paradigmatic approaches to music composition [Koutsomichalis, 2016a] and database aesthetics in general [Vesna, 2007]. Far from being a shift specific to visual culture, the above turn has dramatically affected music listening practices [Koutsomichalis, 2016b]. Sonification and audification have been well standardized in both artistic and techno-scientific milieux as plausible ways to aesthetisize data [Dombois and Eckel, 2011; Walker and Nees, 2011]. Approaches of a purely technical/scientific scope have been accounted for in the proceedings of the various International Conferences for Auditory Display[1] that occur annually, with a few exceptions, since 1994. There are numerous diverse, even disparate, trends to sonification and while there are numerous examples of a purely scientific or artistic merit, real life practice often transcends classification and it is not rare at all that scientific and artistic mileux are simultaneously relevant to a project [Prudence, 2014; Arsenault, 2002; Wilson, 2002]. *Desoundralization* is a characteristic example of such a case, since our original intention is to both compose 'intriguing', in purely subjective terms, sound-art, as well as to experiment with new means/tactics that help us better scrutinize, perceive, understand and eventually present to an audience large-scale decentralized systems applied in contemporary social-technical ecosystems of power systems and urban environments.

Desoundralization concerns the sonification of rather complex data that exhibit decentralized computational intelligence, i.e. localized data that are an actual result of collective decision-making and peer-to-peer interactions between autonomous agents. In this context, and at least as far as our particular approach is concerned, sonifying such systems proves to be a largely ad-hoc process that brings forth a series of questions that are simultaneously of artistic and scientific merit: how can we sonify systems of decentralized compu-

[1] Available at https://smartech.gatech.edu/handle/1853/49750 (last accessed: November 2016)

tational intelligence that are too complex or non-intuitive to be approached with mainstream thinking and ordinary logic, so that they become 'meaningful' for both specialists and the general public? what structural qualities of the system under scrutiny may be revealed in this way? how can we probe data-sets and their embedded infrastructure when the latter is largely unknown or too complex/large to be known? is it possible to demonstrate alternative decentralized design patterns for computational intelligence in data-intensive systems by means of sonification? how can we aestheticize or foreground those attributes that are relevant to intelligent decentralized systems, e.g. robustness, scalability, privacy-by-design, fault tolerance, fairness and social welfare?

Desoundralization is an ongoing, collaborative and highly-interdisciplinary endeavor to address the above questions. Our approach is pragmatic and largely empirical. We zero in on how we can embody the above questions as well as our particular responses into the actual sonic outcomes so that they are experienced in a phenomenological fashion. In this paper we account for the first outcomes of the project, namely for the first two *studies in decentralized computational intelligence*, and we elaborate on our methodological traits. In particular, we discuss the ways in which we sonify two distinct datasets produced by I-EPOS, a fully decentralized combinatorial optimization system [Pilgerstorfer and Pournaras, 2017]. I-EPOS is applied in two domains of participatory sharing economies: (i) *energy management* for a more sustainable energy usage in a smart grid and (ii) *bicycle sharing* for load-balancing the bicycle stations of a smart city. The sonification process uses as input the two distinguished datasets produced by the same algorithm for each of the two application domains.

Decentralized Computational Intelligence

Several data-intensive techno-socio-economic systems are challenged by collective decision-making outcomes of the actors they consist of. For example, energy consumers choose the level of energy consumption and the moments when they consume energy from the power grid. Their collective choices may result in power peaks that can cause catastrophic blackouts or high overall energy prices [Pournaras et al., 2014a]. Similarly, sharing economies emerging in the context of smart cities such as bicycle or other vehicle sharing [Midgley, 2011], require coordinated decision-making so that the availability of bicycles in stations meets a varied demand. In this way, utility companies and operators do not need to perform expensive manual relocation of bicycles among the stations. In such large-scale, distributed in nature, decision-making problems the aggregate result of local autonomous choices result in system-wide outcomes that influence the overall reliability and sustainability of the environment in which citizens reside and act. Given that citizens' choices are made among different feasible options, or options with a varied cost for each citizen, the computational problem is combinatorial: the number of possible global outcomes is exponential as k^n, where k is the number of options of each citizen and n is the number of citizens participating. This computational complexity can even challenge Big Data infrastructures running MapReduce computational models [Lin et al., 2016].

Combinatorial optimization in techno-socio-economic domains such as the aforementioned one is challenging and not straightforward to apply. Such domains entail complications that go beyond technical ones. Big data infrastructures often require the collection of massive personal data for parallel processing in large, energy-intensive, and expensive data centers. This approach raises several issues of trust, privacy-intrusion, surveillance, discriminatory actions and undermining of autonomy [Hajian et al., 2015; Helbing and Pournaras, 2015]. An alternative approach is to perform a fully decentralized combinatorial optimization using collective intelligence deployed over crowdsourced Internet of Things devices ran by citizens. I-EPOS, the *Iterative Economic Planning and Optimized Selections* [Pilgerstorfer and Pournaras, 2017] is an example of such an alternative decentralized system. I-EPOS builds upon the earlier EPOS optimization mechanism [Pournaras et al., 2014a,b] and adds on a fully decentralized back propagation learning capability. Software agents of I-EPOS run in citizens' devices structured in self-organized [Pournaras et al., 2014c] tree topologies over which they perform a bottom-up and top-down networked exchange of information to locally perform an informed choice for a resource allocation plan to execute. This optimum plan may represent the schedule of energy consumption of a household appliance that minimizes power peaks or maximizes the stations with available bicycles to pick up. Selection is made using a fitness function that receives as input the plans of the agents and the aggregate agent selections made in the bottom part of the tree topology. Figure 1 visualizes the tree topology and the plan selections made at the first iteration of an I-EPOS execution.

For *Desoundralization* we zeroed in on the following local and global output data of I-EPOS:

Selected plans-local The globally optimum plans locally selected by the I-EPOS agents at every iteration.

Standard deviation-global This is the global evaluation criterion at every iteration and it is used as a local minimization criterion in the fitness functions of the agents as well.

Aggregate plans-global The aggregate, computed by summation, of all selected plans at each iteration.

Incentive signal-global The computed cost signal at every iteration used in the fitness function of the optimization process.

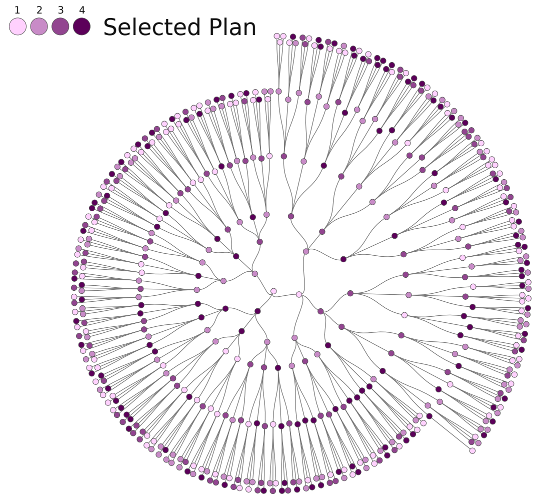

Figure 1. EPOS agents performing collective decision making over a self-organized tree topology for the optimization of the aggregate energy demand. Agent locally generate four energy demand plans and cooperatively make a selection.

Output data are generated using real-world data from the PNW Smart Grid Demonstration Project[2]—concerning the energy consumption for the period 23.07.2014, 01:00-12:00, from 493 households with 4 generated[3] plans per agent to choose from—as well as from the Hubway Data Visualization Challenge[4]—concerning trips from 1000 extracted unique users recorded for the Hubway bicycle sharing system in Paris and showing the available bicycle stations at a two-hour morning time slot (08:00-10:00).

The fitness function of the agents is expressed by a gradient descent that minimizes the standard deviation of the plans collectively chosen over the tree topology using I-EPOS. In the smart grid domain, a minimum standard deviation represents a stable energy demand with minimal changes in the power supply. This decreases generation costs and increases system reliability [Pournaras, 2013]. In bicycle sharing, a minimum standard deviation of incoming/outgoing bicycles among the stations keeps the utilization of all stations balanced and minimizes the operational costs of technical staff moving bicycles between stations.

Agents are self-organized [Pournaras et al., 2014c] in a binary tree topology and can be sorted using several criteria, for instance, how many number of plans they have, or mathematical properties of the plans such as their standard deviation. For the purpose of this work, the following three agent strategies are evaluated:

STRATEGY A Agents with a high total number of plans and high standard deviation in the plans are placed on top. Agents use all their plans throughout the learning process and perform decision-making with a fitness function expressed with standard gradient descent [Burges et al., 2005].

STRATEGY B Agents with a lower number of plans and higher standard deviation are placed on top. Their fitness function is expressed in standard gradient descent [Burges et al., 2005] by considering two plans and adding one more every 10 iterations.

STRATEGY C Agents are randomly placed in the tree topology. They learn with their fitness function expressed in the adam gradiend descent [Kingma and Ba, 2014].

Figure 2 illustrates the learning curves of the three strategies in smart grids and bicycle sharing. The strategies may have comparable performance, nevertheless they traverse different optimization trajectories and their learning curves diverse. Accordingly, the three different strategies become an important empirical means to evaluate the fitness of the sonification and, up to a certain extent, to quantify the complex inter-relationships between abstract aesthetic qualia and measurable statistical information.

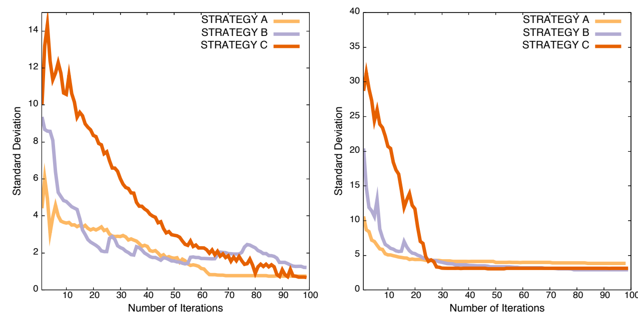

(a) Bicycle (b) Energy

Figure 2. Learning curves of three agent strategies represented by the decrease in standard deviation over 100 algorithm iterations. These curves are a core part of the sonification that turns their measurable difference into an auditable result, an aesthetic perceptual experience of decentralization.

Sonification as Material Speculation

Sonification systems are typically stratified and comprise a series of submodules that retrieve, clean, filter, process and parse data [Walker and Nees, 2011; Koutsomichalis, 2013].

[2]Available upon request at `http://www.pnwsmartgrid.org/participants.asp` (last accessed: November 2015)

[3]Plans are generated by shuffling temporal values of the original energy consumption. Such permutation may represent a change in a user activity, for instance, exchanging the time of taking a shower with the time of cooking.

[4]Available at `http://hubwaydatachallenge.org` (last accessed: November 2016)

Coupling the output of such processes to sound synthesis algorithms is, then, addressed by mapping schemas. We have implemented the sound synthesis algorithm as well as the necessary mapping and scaling schemata using Supercollider[5], a state-of-the-art programming language capable of both real and non-real time sound synthesis. Sonification systems implemented in SuperCollider have had some relevance to our research [Koutsomichalis, 2013; Hermann et al., 2007; Grosshauser and Hermann, 2009; de Campo et al., 2011]. They, nevertheless, do not reflect on the decentralization process, i.e. the actual data locality and peer-to-peer interactions, and therefore, they mostly concern cases where the sonification process ends up being a manifestation of a central meta-authority with full information: a creative sound composition process that does not need to draw parallels or convey concepts of network protocols designed for a decentralized data management. On the other hand, there is an abundance of resources regarding the sonification of decentralized systems, in particular evolutionary systems and flocking simulations [Huepe et al., 2014; Miranda, 2004; Bisig et al., 2007; Blackwell and Young, 2004]. Most of this work, however, concerns simulation models and synthetic data, in contrast to *Desoundralization* that exclusively deals with real-world socio-technical data and the actual prototyped implementation of a decentralized system for computational intelligence applied to sharing economies.

Accordingly, our approach has been largely driven by pragmatic and ad-hoc experimentation. Pragmatic, in that we try out different tactics and then empirically assess the results against their potential aesthetic and scientific merits. Ad-hoc, in that we do not rely on any specific stratagem nor do we necessarily intend to extrapolate generic algorithms to be applied elsewhere. In this sense, and up to a certain extent, there is an analogy between our approach and well-standardized improvised composition practices—these typically concern unique and situated pieces of music meant to be performed in some particular context (space, time, performers,etc). *Desoundralization* also suggest a similar contextual and material specificity that is relevant to particular decentralized systems and the data they generate. We do not suggest that such an approach is unique or particularly innovative. We, nevertheless, emphasize that since we attempt to compose sonifications that are informative, aesthetically intriguing and effective as means to probe complex, maybe even unknown or unidentified content, this element of material specificity becomes of paramount importance—to the extent that our overall practice could be thought of as a way to speculate about the data under scrutiny. Thinking about decentralized data in material terms entails immensely complex interrelationships between agent-dependent/locally-produced information and the system-wide outcomes of agents' coordinated actions. As far as the two datasets under scrutiny herein are concerned, consider the sheer complexity of the optimum plans locally selected by the agents at every iteration (selected plans-local) as well as the ways in which they might relate with one another and with the various globally quantified measures in the scope of each different optimization strategy. We argue that our sonifications do succeed in articulating this sheer complexity, at least up to a certain extent.

In that vein, our approach comprises testing and evaluating different mapping schemata and audio synthesis strategies. In the course of our experimentation with different sonification methods and mappings, complex aspects of the decentralized design and emergent intra/inter-dependencies between localized data become eventually apparent, yet in a haphazard fashion. Still, once some particular aspect or feature that we understand as exemplificatory of some function within the system becomes implicit, it is possible to further ameliorate and fine-tune the sonification so that the former becomes explicit and even foregrounded. It is undoubtedly questionable to what extent we create ourselves those features that we are supposed to probe, but such an uncertainty is integral to the scientific method in many contexts and has been one of the main arguments of those criticizing it [Feyerabend, 1993; Latour and Woolgar, 2013]. In the context of *Desoundralization*, emergent uncertainties of this kind are rather desirable since they allow alternative insights to the systems under scrutiny and, more importantly, since they may also bear aesthetic merit.

Under these premises we discuss a few of our discarded sonification attempts and explain how they pave the way to the ones we finally adopt, so that the pragmatic, ad-hoc and largely speculative nature of our approach is accounted for.

For the bicycle-sharing data, we initially experiment with a few disparate and complex audio generators producing sustained noisy textures in parallel and then we use data to control the various synthesis parameters. However, the resulting audio offers little, if any, insight to the dataset: the transition from one iteration to the other is unnoticed and no particular change occurs as the data progress to the eventual convergence and, in general, nothing in what we would listen delineates the complex inter-dependencies enacted the coordinated action of hundreds of localized agents. Given the particular nature of the locally produced data—i.e. mostly zeros with the occasional appearance of ±1 or ±2—it immediately becomes evident that in order to represent peer-to-peer relationships between the various autonomous agents and to potentially reveal properties still unknown to us, we should focus on aestheticizing the intrinsic 'rhythmicity' of those numbers scattered across the 1000 nodes. Accordingly, we follow a completely different approach, this time revolving around percussive sounds arranged in complex configurations of varying density with respect to the

[5] Available at `https://supercollider.github.io/` (last accessed: November 2016).

localized data. These series of experiments make explicit, albeit fail to properly account for, a series of simultaneous inter-dependent evolving ‘rhythms’ between the localized data and their system-wide quantified effects as the entire system approaches convergence. This observation has been eye-opening in how we should proceed further. Following, we scan the data in parallel from iteration to iteration, experimenting with different granular synthesis generators and trying to figure out which synthesis parameters have to be controlled by what kind of data so that the both the microscopic intra-dependencies of the localized data and their coordinated system-wide effects are properly delineated. In this vein, we experiment with different generators (based on both pure oscillators and the appropriation of real-world recordings), with more sophisticated audio-spatialization and more sophisticated meta-mappings employing formulated ratios and basic statistical analysis on the raw input data.

To begin our orientating within the energy-management dataset, we simply use the localized data (floating point numbers) to control the frequency of simultaneously sounding oscillators—one per agent. The result is rather discouraging: semi-periodic oscillations between the very same timbres—a strong indication that the variance and volatility of the localized data is rather limited. We then seek ways to both emphasize this rhythm and to better probe the subtle modulations that sustain it. The fact that the two datasets have more or less the same structural qualities (hundreds of local agents, hundreds of iterations, eventual convergence of some sort, etc.) is a strong indication that the same mapping strategy employed in the case of the bicycle-sharing dataset is applicable herein, too. However, the energy-management case calls for a completely different approach as far as audio synthesis is concerned: simple, percussive sounds can describe neither how points in a floating-number continuum relate with one another, nor how they are distributed across the hundreds of autonomous nodes that constitute this system. By experimenting with various other kinds of audio generators a peculiarity of the localized data becomes apparent: despite occupying a practically continuous domain and despite being arbitrarily volatile, most of the time they oscillate across fixed attractors. To arrive at sonifications that would be both exemplificatory of this behavior and compositionally intriguing we have to come up with more complex and nonlinear generators that would accentuate subtle changes and occasionally behave in an unpredictable fashion. In this way we could manage to represent all stable, bi-stable and divergent states of the system with more interesting sonic textures.

We draw upon the experience of all our experiments and discarded strategies in order to first identify what exactly we aim at exploring in the data under scrutiny and, then, to properly account for it. Our approach has been largely iterative and empirically-driven and in this sense, it shares some grounds with design science paradigmatic methodology, where the goal is to create a satisfactory artifact via an iterative process [Baskerville et al., 2016; Martens et al., 2012]. The sonification algorithm in its final form embodies the insights gained from previous attempts. It also embodies our own subjective decisions on what emergent features should be foregrounded and in what ways, so that the sonification eventually becomes sound-art worth listening to for its own sake and, more importantly, a process of speculation about the system under scrutiny. While speculation of that kind might be unacceptable in other scientific mileux, it can become an invaluable tool in establishing new means for the general public to conceive complex decentralized systems that are too complex and non-intuitive for the mainstream thinking and general perception in society.

Sound Synthesis and Mapping

Herein, we illustrate the sonification process of the two output data of I-EPOS on energy demand and bicycle sharing. The sonification process does not require input about the tree topology and uses entirely the local and aggregated data of the agents. For each dataset, the approaches are complementary in spirit. For the bicycle sharing data we are inspired by glisson synthesis techniques [Roads, 2004, pp. 121-5]. The sonification of energy management comprises a series of rather sophisticated generators that are active throughout the entire duration of the piece. In both cases the idea is to simultaneously represent the current state of all agents as they ‘perform’ the various algorithmic iterations, occasionally using complex mappings that also involve system-wide statistical measures, so that the emergent complexity of the original data manifests in sonic terms. In this way, as the sonification advances, one may listen to how the overall performance of the system changes at each iteration to arrive at a certain performance outcome—a certain ‘telos’ suggesting the eventual convergence.

Algorithm 1 demonstrates the minimal glisson generator used for the sonification of the bicycle sharing data. The recording of a bicycle bell is used as a source sound to convey the application domain. The generator varies the rate by which a fragment of a pre-recorded sound sample is reproduced so that a glisson is generated. The algorithm scans the data and spawns thousands of glissons with their initial and final playback rate, the duration of their glissando, their amplitude and their positioning in the stereo field parameterized accordingly. Each glisson is also fed through a single-channel parametric equalizer and a low-pass filter, so that their spectral characteristics are also forged algorithmically. 1000 glisson synthesis processes—one for each agent in the original dataset—are initiated simultaneously. The arguments for each glisson are then calculated as follows:

Initial playback speed: Calculated by the factor of the selected plans divided by the factor of the aggregate plans divided by the maximum of all aggregate plans.

Final playback speed: Calculated by the initial playback speed and the difference of the topical standard deviation minus the standard deviation at the final iteration.

Frequency of the equalizer: Calculated by the ratio of the incentive signal and the maximum incentive signal.

Gain or attenuation factor for the equalizer: Calculated as the initial playback speed.

Glisson duration: Calculated with respect to the ratio of the aggregate plans and the incentive signal.

Amplitude: It is a random value.

Algorithm 1 Glisson generator

```
SynthDef(\gen, { arg rate_start = 1, rate_end = 1, pan =
      0, dur=0.5, amp = 0.1, freq = 440, db=0, buf=0;
  var signal = Mix.new(PlayBuf.ar(2,buf,Line.kr(
        BufRateScale.kr(buf)*rate_start,BufRateScale.kr(
        buf)*rate_end,dur),loop:1));
  signal = MidEQ.ar(signal,freq,0.4,db);
  signal = LPF.ar(HPF.ar(signal,40),10000);
  signal = signal * EnvGen.kr(Env([0,1,1,0],[dur/4,dur/2,
        dur/4],[-3,0,3]),doneAction:2);
  Out.ar(0, Pan2.ar(signal,pan) * amp)})
```

Note that in both sonifications the sonic image gradually progresses from monophony to a wide sonic image (stereo or multichannel). The overall speed is calculated as a fraction of the difference between the standard deviation of the last iteration and that of the current, so that the whole system becomes progressively slower (or faster, if desired).

Algorithm 2 demonstrates the audio generator used in the case of energy management. This generator is arguably more complex and comprises the esoteric Astrocade generator[6], a reverb unit, a chaotic noise generator (Crackle), a single channel parametric equalizer and sophisticated multi-mappings hard-coded into the body of the synthesizer. Note that $reg1$, $reg4$, $reg5$, $reg7$ are higher order inputs, which are mapped into several audio synthesis parameters the state of which often depends on more than one of the former ones. Note also that the parameters of the Astrocade are meant to emulate 8-bit programming registers, different bits of which often control different parameters, so that the generator cannot be controlled in a linear fashion. Chaotic noise, non-linear ranges in conjunction with higher-order inputs and multi-mappings built-in the synthesis engine account for generators that would dramatically 'magnify' subtle changes in the input data and that would occasionally destabilize and behave in a chaotic fashion. As discussed in the previous section, such a behavior is desired. The logic of the rest of the algorithm is very similar to that of the bicycle-sharing data: a generator is instantiated for each of the 493 agents at play and subsequently controlled by the input data. The parameters are set with respect to the following schema:

reg1: Calculated by the ratio of the selected plans over the aggregate plans over the maximum of the latter.

reg4: Controlled by the selected plans.

freq: Calculated by the ratio of the incentive signal and the aggregate plans.

reg5: Calculated with respect to the selected plans multiplied with the frequency, as defined above.

reg7: Calculated by the ratio of the aggregate plans and the incentive signal.

Gain or attenuation factor for the equalizer: Calculated as $reg1$.

Amplitude: Calculated by the value of selected plans.

Algorithm 2 Generator used in the sonification of the household energy consumption data.

```
SynthDef(\gen, { arg reg1=0, reg4=0, reg5=0, reg7=0, pan
      = 0, amp = 0.1, freq = 500, db = 0;
  var signal = Astrocade.ar(0, reg1, 0, 0, reg4, reg5,
        15, reg7);
  signal = Select.ar( ((reg4+reg7+reg5) / 385).ceil, [
        signal, FreeVerb.ar(signal * Crackle.ar( 0.9 + ((
        reg1 + reg4 ) / 400), mul: (reg1 + reg7)/500), ((
        reg7 + reg5) / 500), 0.5, 0.7)]);
  signal = signal * LFNoise2.ar(freq / (reg1 + 5));
  signal = MidEQ.ar(signal,freq, 0.2 ,db);
  Out.ar(0, Pan2.ar(signal, Lag.kr(pan,lagTime)) * amp *
        0.2);
})
```

Modulating the various ranges, changing the denominators in the ratios, accelerating or decelerating tempi or the magnitude of the various changes (e.g. by multiplying with a progressively larger/smaller factor) all result[7] in accentuating and foregrounding different system properties. Improvisation can be the means to further probe decentralized data.

Discussion and Conclusion

The outcome of this work has been 6 audio pieces[8] with duration of about 20 minutes each. These represent the soni-

[6]According to the help file of SuperCollider, "a custom 'IO' sound chip driver by Aaron Giles and Frank Palazzolo", which may not be "working as it should, but it's still somewhat fun sounding"

[7]The illustration of the sound synthesis eschews everything relevant to scaling and interpolation schemata for reasons of simplicity. The latter are, nonetheless, rather significant in defining how the outcome eventually sounds and are often defined by exponential distribution curves and even in a reverse order, so that higher input values correspond to lower output ones. Again, employing different interpolation/scaling schemata allow us to pinpoint different properties of the systems under scrutiny.

[8]The audio pieces are available at `http://evangelospournaras.com/shared/Desoundralization-mp3s.zip` (last access: February 2017).

fication of the three different strategies per dataset. The different strategies sound as different narratives leading to more or less the same 'telos'—i.e. the state of convergence for each system. In both cases, the range and the characteristics of change to various globally relevant parameters, such as panning or speed, are fine tuned so that the 'telos' does not come too fast or too sudden and so that the most important micro-modulations at play are accentuated throughout the entire sonification. Several sonic parameters are controlled by means of ratios or other mathematical formulations involving both local and global data as well as various statistical maxima/minima. In this way we can pinpoint those particular abstract properties that are identified as the most important or explicatory during the experimentation stage that has preceded. We can sustain the complex nexi of micro-modulations that seem to support and accentuate the aforementioned narrative so that we also arrive at aesthetically interesting compositions for sound-art enthusiasts.

Comparing the different strategies with one another is revealing. In the case of the bicycle sharing data, STRATEGY A and in particular STRATEGY C are more evocative of the sound origin, that is a bicycle bell: towards the end of the sonification the various abstract glissandi gradually transform to definite bell sounds. STRATEGY B, which is the preferred one in terms of aesthetics, is abstract throughout and is rather characterized by minimal subtle change; the listener does not encounter the dramatic glissandi present in the other two strategies as the placement of the agents in the tree by the self-organization process result in plan selections with a more moderate acoustic impact. The case of the energy management data is very different, both because of the nature of the data and the synthesis generator in use. A more involved and non-linearly controlled sound generator is used as the basis of the entire sonification to account for the unstable, unpredictable and often haphazard nature of electricity. This generator occasionally 'crackles' or 'explodes' and often results in noisy textures while also sounding in the correct pitches. All strategies here seem to eventually self-modulate, repeating more or less the same sonic states over and over. All in all, the different strategies sound (dramatically) different, in this way exemplifying the different processes that are in play, their effects on the scattered autonomous agents, and our compositional decisions regarding which kinds timbres/textures should be associated with each system.

Desoundralization zeroes in on the output of I-EPOS and intends to probe the complex locally generated data as well as the way they inter-depend with system-wide properties and macroscopic statistic measures. The proposed sonifications allow us to gain invaluable insight in the internals of I-EPOS. Traditional visualizations such as the ones depicted in Figure 2 are helpful in delineating the behavior of the aggregate results and in quantifying statistical measures and overall system performance. Yet, it would be immensely difficult, if at all possible, to aestheticize complex qualities with traditional visualizations such as the intra-dependencies forged between the locally generated data as the iterations progress or the relationship between the latter and overall macroscopic qualities of the system. Our proposed sonifications do enable us to delineate the complexities of decentralized systems as manifested both microscopically and macroscopically, to speculate on particular aspects of them and to present our findings in a straightforward phenomenological fashion that may also hold artistic merit. In our approach we take into account all localized agents and the data they generate in a bottom-up fashion, in this way exposing the granularity of the system, while at the same time we keep comparing them with system-wide quantified data directly encoding the results in sound. While it may be impossible for non-specialists to fully understand the deeper implications of such systems, it is still straightforward for the general public to appreciate the various processes at play in their proper granularity and, more importantly, to immediately perceive how overall convergence of the system translates to microscopic modulations in the locally-generated data and vice-versa.

References

Arsenault, L. M. (2002). Iannis xenakis's achorripsis: The matrix game. *Computer Music Journal*, 26(1):58–72.

Baskerville, R., Kaul, M., Pries-Heje, J., Storey, V. C., and Kristiansen, E. (2016). Bounded creativity in design science research.

Bisig, D., Neukom, M., and Flury, J. (2007). Interactive swarm orchestra. In *Proceedings of the Generative Art Conference. Milano, Italy*.

Bjørnsten, T. (2016). Big data between audio-visual displays, artefacts, and aesthetic experience. *MedieKultur*, 59:50–72.

Blackwell, T. and Young, M. (2004). Swarm granulator. In *Workshops on Applications of Evolutionary Computation*, pages 399–408. Springer.

Burges, C., Shaked, T., Renshaw, E., Lazier, A., Deeds, M., Hamilton, N., and Hullender, G. (2005). Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, pages 89–96. ACM.

Chandler, D. (2015). A world without causation: Big data and the coming of age of posthumanism. *Millennium*, 43(3):833–851.

Clough, P., Gregory, K., Haber, B., and Scannell, R. (2015). The datalogical turn. *Non-Representational Methodologies: ReEnvisioning Research*, page 146.

de Campo, A., Rohrhuber, J., Bovermann, T., and Frauenberger, C. (2011). Sonification and auditory display in supercollider. *The SuperCollider Book*.

Dombois, F. and Eckel, G. (2011). Audification. In Hermann, T., Hunt, A., and G., N. J., editors, *The sonification handbook*, pages 301–324. Logos Verlag, Berlin, DE.

Feyerabend, P. (1993). *Against method*. Verso, New York, NY.

Grosshauser, T. and Hermann, T. (2009). The sonified music stand–an interactive sonification system for musicians. In *Proceedings of the 6th sound and music computing conference*, pages 233–238. Casa da Musica, Porto, Portugal.

Hajian, S., Domingo-Ferrer, J., Monreale, A., Pedreschi, D., and Giannotti, F. (2015). Discrimination-and privacy-aware patterns. *Data Mining and Knowledge Discovery*, 29(6):1733–1782.

Heinrich, F. (2016). (big) data, diagram aesthetics and the question of beauty. *MedieKultur*, 59:73–94.

Helbing, D. and Pournaras, E. (2015). Society: Build digital democracy. *Nature*, 527:33–34.

Hermann, T., Bovermann, T., Riedenklau, E., and Ritter, H. (2007). Tangible computing for interactive sonification of multivariate data. In *Proceedings of the 2nd Interactive Sonification Workshop*, volume 3, page 171.

Huepe, C., Colasso, M., and Cádiz, R. F. (2014). Generating music from flocking dynamics. In *Controls and Art*, pages 155–179. Springer.

Kingma, D. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kitchin, R. (2014). Big data, new epistemologies and paradigm shifts. *Big Data & Society*, 1(1):1–12.

Koutsomichalis, M. (2013). *Mapping and Visualization with SuperCollider*. Packt Publishing, Mumbai, IN.

Koutsomichalis, M. (2016a). Catalogue æsthetics: Database in and as music. In Kostagiolas, P., Martzoukou, K., and Lavranos, C., editors, *Trends in Music Information Seeking, Behavior, and Retrieval for Creativity*. IGI-Global., Hershey, PA.

Koutsomichalis, M. (2016b). From music to big music: Listening in the age of big data. *Leonardo Music Journal*, 26.

Latour, B. and Woolgar, S. (2013). *Laboratory life: The construction of scientific facts*. Princeton University Press.

Lin, M.-Y., Yang, C.-W., and Hsueh, S.-C. (2016). Efficient computation of group skyline queries on mapreduce. *GSTF Journal on Computing (JoC)*, 5(1):69.

Martens, W., McKinnon-Bassett, M., and Cabrera, D. A. (2012). Perceptual evaluation of stochastic-event-based percussive timbres for use in statistical sonification. In *Audio Engineering Society Convention 132*. Audio Engineering Society.

Midgley, P. (2011). Bicycle-sharing schemes: enhancing sustainable mobility in urban areas. *United Nations, Department of Economic and Social Affairs*, pages 1–12.

Miranda, E. R. (2004). At the crossroads of evolutionary computation and music: Self-programming synthesizers, swarm orchestras and the origins of melody. *Evolutionary Computation*, 12(2):137–158.

Pilgerstorfer, P. and Pournaras, E. (2017). Self-adaptive learning in decentralized combinatorial optimization–a design paradigm for sharing economies. In *Proceedings of the 12th International Symposium on Software Engineering for Adaptive and Self-managing Systems (SEAMS 2017)*. IEEE.

Pournaras, E. (2013). *Multi-level reconfigurable self-organization in overlay services*. PhD thesis, TU Delft, Delft University of Technology.

Pournaras, E., Vasirani, M., Kooij, R. E., and Aberer, K. (2014a). Decentralized planning of energy demand for the management of robustness and discomfort. *IEEE Transactions on Industrial Informatics*, 10(4):2280–2289.

Pournaras, E., Vasirani, M., Kooij, R. E., and Aberer, K. (2014b). Measuring and controlling unfairness in decentralized planning of energy demand. In *Energy Conference (ENERGYCON), 2014 IEEE International*, pages 1255–1262. IEEE.

Pournaras, E., Warnier, M., and Brazier, F. M. (2014c). Adaptive self-organization in distributed tree topologies. *International Journal of Distributed Systems and Technologies (IJDST)*, 5(3):24–57.

Prudence, P. (2014). Data transmutations: Making sound sense of big data. *Neural*, 48:19–21.

Roads, C. (2004). *Microsound*. MIT press.

Vesna, V., editor (2007). *Database aesthetics: Art in the age of information overflow*. University of Minnesota Press.

Walker, B. N. and Nees, M. A. (2011). Theory of sonification. In Hermann, T., Hunt, A., and G., N. J., editors, *The sonification handbook*, pages 9–40. Logos Verlag, Berlin.

Wilson, S. (2002). *Information arts: intersections of art, science, and technology*. MIT press.