Locus Diffuse: An Agent-Based Sonic Ecosystem for Collaborative Musical Play

Rory Hoy and Doug Van Nort

DisPerSion Lab rorydavidhoy@gmail.com, vannort@yorku.ca

Abstract. Locus Diffuse is a networked multi-user instrument populated by a simulated slime mold and four human players. Mimicking the biological behavior of slime mold and establishing a virtual living network between player nodes, the system sonifies interaction along these connections. Participants use a browser based interface to play the multi-user instrument, and access an accompanying stream for audio and visual output of the system. Player responses from various play sessions are explored and reported in relation to sonic ecosystems as a product of sound sources intersected with agent behavior, defining interaction through personal connection to agents, an aural vs visual understanding of the system, and various frames of focus employed by participants in regard to human/machine and inter-human collaboration.

Keywords: agent-based musical systems, multi-user instruments, natural computing, slime mold

1 Introduction

Musical play has acted as a vessel for a communal engagement, identity, exploration, and expression throughout history [6]. While the style of play may vary from recital of composed works to free improvisation (and every permutation in between/beyond), a common thread is that emergent group playing dynamics are revealed through the complex interactions between each player [2]. This aspect of musical collaboration is a social ritual in which participants are afforded a medium of aural communication beyond the verbal. Players can be represented as nodes within a network of participants that expresses interpersonal playing decisions, and the resulting sonic landscape can be seen as an emergent form of this established network. Viewed in this way, collective action results in a cumulative sound field that is the product of each node's (player's) input. An interactive instrument/environment, named *Locus Diffuse* was developed to investigate and facilitate these emergent participatory network structures within collaborative musical play for four players. This is mediated by an instrument in which users can "play" a space through interaction both with its population of simulated agents and with each other. Situated at the crossroads of sonic ecosystem design, agent-based musical systems, multi-user instruments, and networked performance, *Locus Diffuse* draws on a network of practices to produce a system that is used to interrogate the outcome of their resulting collaborative human/machine interplay. The system was initially planned for a full scale room implementation within the DisPerSion Lab at York University, however due to social distancing restrictions caused by the global COVID-19 pandemic, the project was required to pivot to a distributed virtual performance space. Players and spectators access a live audio/visual stream as a collective hub for generated activity, while controlling their input within an additional browser window or mobile device. During this time of relative isolation, the project's aesthetic themes of connection and collaboration were heightened through this additional networking component, facilitating the communal play of all participants.

The behavior of the system's population of agents is modelled on networking structures found within the biological form of slime mold. Harnessing natural processes of emergent form and community, these organisms have been demonstrated to have repeatable emergent behaviors of aversion and attraction to environmental stimuli. Most notably their structure takes the form of thin physical networks between food sources, and through implementing approximations of this behavior, *Locus Diffuse* generates flowing and reactive networks of autonomous agents moving between player positions. We argue that these organisms are well suited as a metaphorical frame that mirrors the collaborative generative network-like structure found within musical performance, and that mapping various interaction responses can result in compelling ecosystemic behavior.

2 Related Works & Literature Review

2.1 Harnessing Biology - Artistic & Computational Implementations

Natural Computing studies the application of natural phenomena within ecological systems and biological structure to a multitude of computational tasks [18]. These implementations can come in the form of mimicry, approximation, and inspiration from structures found within natural systems. Slime mold, specifically Physarum polycephalum, exhibits extraordinary behavior for an organism which contains no explicit sensory organs, capable of tactile, chemical, and photoreceptive sensing. The body consists of a single cell, but can produce many flexible space-searching tubules and can change their thickness to allow for a greater flow of cytoplasm in order to move. The body attempts to move in a direction towards food/positive stimulus or away from negative stimulus [5]. The slime mold is able to then retract, reinforcing a minimal path between all available food sources within even complex spatial layouts such as mazes [14]. Computational models of slime mold have resulted in creating logical gates, solving resource heavy computation, and achieving primitive memory [1]. Artistic applications of slime mold have been advancing in tandem with computational implementations. Miranda et al. [12] constructed a sound synthesis project which allowed for recordings of voltage at various locations through the electrical activity of a slime mold network across a series of food nodes. This data was then used within a granular synthesis engine to generate sonic events.

2.2 Sonic & Performance Ecosystems

Sonic ecosystems refer to interactive systems defined by the generation of a reactive audio environment in which self observing behavior and participant input result in audible dynamic feedback [4]. Such systems explore the relationships and outcomes established between human, machine, and ambient environment. A central question in the context of ecosystemic design is the role of the human participation within an established work, and what constitutes “interaction”. Some systems generate a sonic environment purely mediated by an established machine/ambience relationship, while others find room for human interaction to extend these interactions. Di Scipio [4] describes this ability of system self observation as “a shift from creating wanted sounds via interactive means, towards creating wanted interactions having audible traces”, and claims that it is through these traces that compelling sonification can occur.

The original, in-person formulation of *Locus Diffuse* was initially planned to play off of the self-observing vocal & ambient feedback found within the design of the *dispersion.eLabOrate* project [9], a system exploring collaborative sounding within a Deep Listening-inspired sonic meditation [15] context. Within *Locus Diffuse*, self-observation occurs at the agent level. Each agent is only aware of its own state (vs a sense of other’s or environmental current states) and acts according to its sensory input from the environment. Environmental changes and subsequent sonification are a result of the interplay between players and the system’s agents.

2.3 Multi-User Instruments & Networked Music

Intended to promote close relationships between multiple players and resulting play techniques, multi-user instruments allow many participants to perform through a singular instrument. Designing for a multi-user instrument context requires explicit consideration of the intricacies and collaborative experiential content which the instrument/system needs to convey. Jordà [11] outlines key aspects of multi-user instruments that facilitate shared collective control within a musical system. These properties include number of users, user roles, player interdependencies/hierarchies, and the flexibility of each of these components.

Creation of mutual-influence via networked sound data has been explored by pioneering groups such as the League of Automated Music Composers and The Hub [8]. More recently, these networks have also been explored within the realm of telematics, employing the internet as a medium for musical collaboration [16]. Weinberg [21] presents the concept of an Interconnected Musical Network (IMN), live performance collectives in which player interdependencies result in dynamic social relationships and reactive playing. Weinberg states a successful musician network would promote “interpersonal connections by encouraging participants to respond and react to these evolving musical behaviors in a social manner of mutual influence and response”, positioning the performance of group-based music as a social ritual. Additionally, exploring a biological metaphor of the established network, Weinberg [21] states: “Such a process-driven environment, which responds to input from individuals in a reciprocal loop, can be likened to a musical ‘ecosystem.’ In this metaphor, the network serves as a habitat that supports its inhabitants (players) through a topology of interconnections and mutual responses which can, when successful, lead to new breeds of musical life forms...”. This parallels the key ecosystemic theme of *Locus Diffuse* and points back towards the culmination and amalgam of these disparate practices as viable in fostering a connected musical collaborative space.

3 Artistic Intention & System Overview

Locus Diffuse introduces a simulated being that reacts to the movement of players, permeating the environment as a traversable medium. Sonification of the system is achieved when interacting with this mediating entity as well as through participant movement in virtual space, and can therefore only exist/function through the symbiotic relationship of players to it and to each other. Control is not centralized to one participant, nor surrendered to the simulated organism. This control facilitates the musical composition of space, sculpting a form which the simulated organism populates spatially and aurally. This emergent structure and reactive behavior can be paralleled within the participants of the social ritual of "musicking" [19], in which each player has a sensory experience of the whole while also contributing to it. Participating within the shared audio space means enacting this social ritual of musical play, thus the roles and capabilities of players along with the function of environmental agents were established such as to rely on all players. These inter-human and human-agent relationships are critical to explore the resulting network structure. The simulation is contained within Max, employing JS for directing agent positions and control data for grain sonification, JWeb in Max is used for visual feedback in an HTML page, and audio synthesis control patches additionally developed within Max.

3.1 Simulated Agents

Agent behaviour is modeled after the biological structures of Physarum polycephalum, but does not represent an exact scientific model of the organism. Player positions are represented as purple radial gradients within the simulation. Player positions act as food deposits for the simulated agents, and movement results in variations of the environmental structure sensed by the collective simulated slime mold. The simulation is informed by the research of Vogel et al. [20] and inspired by Jones [10], who outlines the mechanics of Physarum polycephalum.

An initial population of 500 agents spawn in the centre of the simulated environment and are given a random starting vector. Each agent is equipped with two sensors positioned at an angular offset of 45 degrees left and right, and a set distance ahead of the agent. The simulated world is quite large (1000^2 pixels) in relation to the size of the cellular bodies (2 pixels), necessitating sensors that have a far reach (default 350 pixels), allowing them to "smell" food sources and trails from a reliable distance. As mentioned in Jones [10], this large distance would normally be considered remote sensing separate from the body of an agent, however this distance also acts as the "overlapping actin-myosin mesh of the plasmodium gel system", allowing the cells to understand their position relative to each other and to nutrient sources. Optimization of the agent network is achieved through a decaying chemoattractant trail deposited and sensed by each agent. Trails are deposited when an agent senses food or another trail, resulting in deposits towards food. As trails diminish over time, an established network is strengthened when searching agents return from an unsuccessful search, or travel along the stream, continually depositing additional trails. Agent sensors check for light values representing chemoattractant strength, average the data collected, and then determine the direction to face. Agents remember the last strongest "smell" they've sampled and choose what to do based on the current reading, always orienting towards the highest value. Agents are in search of energy to keep moving and find more food. Each agent mimics the cytoplasmic streaming behavior of a slime mold, and represents a theoretical main concentration node of this cytoplasm. Energy is a value held by each agent and player attractant node, which maps to qualities of each granular sonification, movement, and rotation speed. Losing energy will cause them to slow or enter a hibernation-like state when approaching zero. Agents which gain energy again can be "revived" from this hibernation state if passed over by a player. Simulated agents actively gain energy while upon a player, while passively losing energy during movement/wandering between nodes. Players regain energy by being in close proximity to others. Agents keep individual energy values as opposed to distributing energy, allowing for unique sonifications based on the amount of energy one contains.

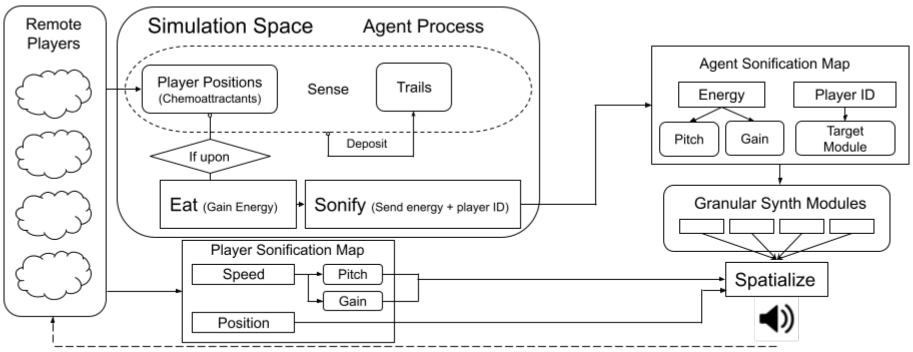

Fig. 1. Human and machine interaction with resulting data flow to sonification

3.2 Sonification

Unique source material was used to ensure an identifiable timbre for each player. Play sessions were done in two waves and audio sources were edited between waves for both refinement of sonification aesthetics, and to gauge changes in play due to these varied timbres. Wave 1 sources were textural in nature, using viscous drips, synth drones, running water, and a filtered conversation as audio material. Wave 2 sources were chosen to result in crisp sonification - timbrally in line with clicking, dripping, droning, and swarming noises. Sonification of a given audio grain was triggered by an instance of an agent "eating" at a particular player location, when an agent takes energy from the player's representational chemoattractant. The sounding potential of a grain triggering is randomized to a 1 in 500 chance upon an agent eating to avoid continuous audio output from a single agent, while also mimicking variance in time needed to break down and process energy from food sources (i.e. a second artistic liberty taken with the model). Messages are sent from the logic JS running in a JWeb, routed to one of four granular synthesis engines, corresponding to a different player. These include the energy value of the agent, and the player ID acting as their source of energy. The granular synthesis patches contain a Petra buffercloud object [13], allowing for accurate single-grain firing (5-50ms long). Energy values are mapped to a pitch multiplier of the source material and gain level. As energy values range from 0 - 100, values are scaled to an appropriate pitch range multiplier between 0.5(±0.2) and 1.7(±0.2), and gain ranging from -30dB to 0dB. Granular synthesis output is then spatialized to the corresponding player position. Player movement is sonified by high frequency sine tones. Unique frequencies are assigned per player, then modulated based on movement speed, with slower movements being modulated down (with higher gain), and faster movements modulated up (with lower gain), which may produce a beating depending on relational position/speed of multiple players. These tones are spatialized in a virtual binaural space using IRCAM's Spat [3].

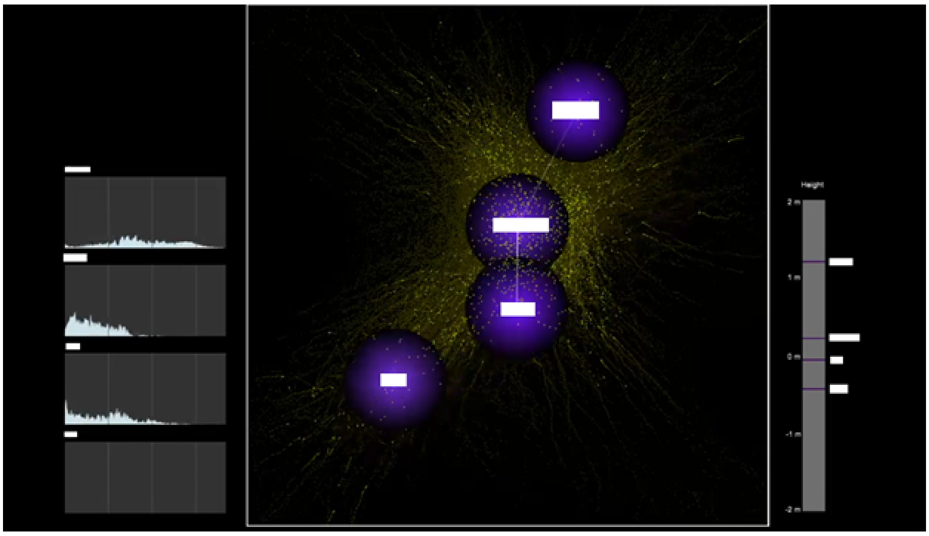

Fig. 2. Stream view of the simulated environment depicting State 4 (participant names censored)

3.3 Networked Interaction & Visualization

Accommodating different devices and network connections was essential in order for public accessibility to players and audience. Control of player movement occurs in the browser through a provided URL, and can be accessed with a browser or touch enabled device. The interface contains a panel for each player consisting of identical controls, including a centre square for position input, and a right-hand slider for vertical movement. The left-hand boxes show spectroscope representations of current sonic activity for each participant. Visual output of the system (Fig. 2) was hosted on a public live stream. This situates the stream as a centralized audio and visual hub for the experience of the instrument, and resulting sonic ecosystem.

4 Survey

Play sessions were held in two waves as open calls on set dates, and public exhibitions following the weekly electro-acoustic improvisation series *DisPerSion Relation X*. The first wave of play sessions focused on a single behavior state of the agents, while the second wave presented players with four varied states. Sessions lasted roughly 25 minutes (with some lasting up to 60 minutes). Following play, participants were asked to complete an anonymous web form. As Wave 2 was centered around four different behavior states, the response form was updated to include questions on state comparisons. Questions focused on perception of the system from two perspectives: relations with other players and the resulting sonification. The questions for Wave 1 (W1) were presented as follows:

1. What was your sense of playing in this virtual environment?
2. What was your sense of connection to the others in the virtual space? (Other players or agents)
3. How did you perceive your own "voice" while playing? (Location, timbre, relation to environment and others)
4. How would you describe your ability (or lack of) to perform expressive musical action?

Wave 2 (W2) introduced states which altered agent trail decay, sensory distance, "death" threshold, birth odds, and agent energy decay. Players were not primed on the behavior of each of these states. The transition between each state was announced to prompt the players that they will be interacting with new behavior. States progressed sequentially through 1-4, but could be revisited following the session. The experienced states were:

- **S1 - Solitary**: Fast trail decay, low sensory distance, default death threshold, low birth odds, and default agent energy decay
- **S2 - Needy**: Slow trail decay, low sensory distance, lower death threshold, high birth odds, and very fast agent energy decay
- **S3 - Lively**: Fast trail decay, high sensory distance, default death threshold, high birth odds, and slow agent energy decay
- **S4 - Starving**: Slow trail decay, high sensory distance, lower death threshold, low birth odds, and very fast agent energy decay

The names provided before the description of each state were given by the first author through personal interpretation of their behavior and were not told to players. Questions from W1 were all asked again, including "For each state:" before a given question. One additional question was asked:

- How would you describe the behaviours of each state? (changes in response, characteristics, etc)

Answer lengths were not prompted to be short or long, allowing players to provide as much detail as they wished. 10 player responses were recorded for both W1 and W2, and a thematic analysis was conducted on this data. Most players had little or no prior experience with participatory musical systems. A small amount had extensive prior experience with improvisational musical play.

5 Responses & Analysis

Participant responses outline a range of interpretations for *Locus Diffuse*, as various natural metaphors were attributed to the audio and visuals. One participant noted "I definitely had the sensation of being immersed in a medium – fluid. The dynamics of the particles, of course, were responsible for evoking this sensation, but so were the sounds and the way that they transformed". While players were primed that the agent behavior was emulative of slime mold, their natural metaphors for the agent behavior tended towards more commonly encountered phenomena of the natural world. Natural processes such as swarms of bugs, flowing rivers, and immersion within a fluid substance were noted as a reaction to both the aural and visual content of the simulation. While similar sources were used as granular input across both waves, the sense of a natural process was far from 1:1 with source audio, and rather was in reaction to both timbre and agent behaviour. This points to the *perception of an emergent sonic ecosystem that is influenced both by variations of the sound source and by agent behavior.*

Interaction with agents, guided by personal connection/narrative, was also a key feature of participant responses. One player noted, "There was a certain appeal to doing things like building 'bridges' between myself and other users, and seeing the cells speed up and slow down made it feel like we were almost taking care of the cells in a way". This was exemplified within states S2 (Needy) and S4 (Starving) of Wave 2, where accelerated agent energy decay and earlier death resulted in huddles of player positions protecting a core population of agents. Players attributed direct and/or implied characteristics towards agent and environmental behavior throughout each of the states. Players would alter the target of these characteristics, displaying that these changes were felt on either an agent or environmental level. Environmental-related characteristics tended to be a product of the visual aspects of the system, noting "busyness" and "growth" of agents within S2 & S4 when trail decay was reduced. For agent behavior characteristics, S1 (Solitary) was perceived as "independent", resulting in localized areas of attraction with distant agents acting indifferent to the presence of energy. One response attributed 'interest' as a quality the agents possessed, stating that "agents seem to be highly invested in the actions of players when they are sharing energy, but seem to actively avoid players who are not working together to share energy". One player noted that these states "rewarded stillness", where one's interaction felt more impactful to the sonification by waiting and allowing the agents to move towards and through them.

Audio and visual cohesion of the system was found to be necessary for players to internalize a complete understanding of the resulting agent behaviors. An interesting trend is shown in some player responses to seemingly *lean towards a visual characterization of the system state vs an aural one*. This can be seen in responses comparing S1 and S3 (Lively). Reports on each states sonic activity were contrasting, noting S1 as reserved and stable, but S3 as busy or chaotic. Although perceived as sonically contrasting, most participants noted S1 and S3 being similar due to visual qualities of the simulation, mainly trail decay rate. Similar reports occurred between S2 and S4 due to their low trail decay rate. This visual bias may also be a product of the relatively low familiarity of participants with musical systems/play experience.

Varied experiences of connectedness were reported: a lack of connection, connection mainly with the simulated agents, and connection to the meditative qualities of *communal movement.* Reports of a lack of connection were attributed to a desire to have dialogue with fellow players (voice/text) in order to coordinate, or again due to a focus upon the visuals, noting “Being able to see where other players were going and patterns they were following made the connection much strong across all states”. The second focus is discussed above (“interaction with agents”). The third focus was on immersion in the system and the sonification of player movements, which divided into two groups. The first noted a distinct gravitation to the mediating agent system and its behaviours, with one participant stating “I could feel each of there positions in a unique way. It was as if they were taking up space in a room”. The second focused on inter-human sounding while immersed in the mediating virtual space, with one participant stating, “I found myself being more consciously aware of the other players’ positions/motions, and adjusted my own motions in relation to theirs”.

6 Conclusions & Future Work

Blending aspects of sonic ecosystems, agent-based musical systems, multi-user instruments, and networked performance to establish a communal musical play context, *Locus Diffuse* depicts these disparate fields of study as complimentary in their nature to establish compelling emergent behavior through various levels of interaction, sounding, group structure, and process. Employing natural computing for the mimicry of biological systems allows for flexible and dynamic collaborative musical agents by speeding up natural processes to allow them to be used in real-time musical computation tasks. The provided system overview allows for detailed understanding of agent mechanics and human/machine interplay resulting in sonification. Play sessions with *Locus Diffuse* resulted in four key observations from participant responses:

- The perception of a sonic ecosystem was tied to variation in sound sources intersected with agent behavior.
- Narrative-based personal connection between players and agents mediated interaction characteristics.
- There was a bias towards a visual understanding of the system vs an aural one.
- The “locus” of experiences of connection were more varied, ranging from a lack of systemic connection, focus on inter-human collaboration, to human-agent collaborative sounding.

Each of these outcomes is a product of the interaction between system behavior, player action, and aural & visual aesthetic decisions, constituting various networks at play between the project’s amalgam of practices, communal musical goals, and telematic structure.

In a post social distancing time, an in-person room scale version of the system will be created in order to explore the translation of the current network-based musical instrument design back into the originally intended space. Translating to this physical space, perceptions related to embodied movement as a control source can be explored within the established agent-based sonic ecosystem. Further research into potential narrative outcomes of inter-human and human-agent-based collaboration may yield interesting results within sonic ecosystems. Sessions aimed at varied levels of musical experience may reveal interesting trends related to aural vs. visual attention, and attention to inter-human or human-agent interactions. Additional telematic sessions with reduced visual feedback may also shed light on how much a purely sound-based system can express these ecosystemic interactions.

References

1. Adamatzky, A.: Physarum Machines: computers from slime mould. World Scientific, Singapore. (2010)
2. Borgo, D.: Sync or Swarm: Musical Improvisation and the Complex Dynamics of Group Creativity. LNCS, vol. 4060, pp. 1–24. (2006)
3. Carpentier, T., Noisternig, M., Warusfel, O.: Twenty Years of Ircam Spat: Looking Back, Looking Forward. ICMC. pp. 270–277. Denton. (2015)
4. Di Scipio, A.: 'Sound is the interface': from interactive to ecosystemic signal processing. Organised Sound, vol. 8(3), pp. 269-–277, United Kingdom. (2003)
5. Durham, A. C., Ridgway, E. B.: Control of chemotaxis in Physarum polycephalum. The Journal of Cell Biology, vol. 69(1), pp. 218-–223. (1976)
6. Frith, S.: Music and Identity. Questions of Cultural Identity, pp. 108–128. (1996)
7. Gale, E., Matthews, O., De Lacy Costello, B. Adamatzky, A.: Beyond Markov Chains, Towards Adaptive Memristor Network-based Music Generation. International Journal of Unconventional Computing. vol. 10. (2013)
8. Gresham-Lancaster, S.: The aesthetics and history of the hub: The effects of changing technology on network computer music. LMJ, vol. 8(1), pp. 39–44. (1998)
9. Hoy, R., Van Nort, D.: Augmentation of Sonic Meditation Practices: Resonance, Feedback and Interaction through an Ecosystemic Approach. In: Kronland-Martinet R., Ystad S., Aramaki M. (eds) Perception, Representations, Image, Sound, Music. CMMR 2019. LNCS, vol 12631. pp. 591–599, Springer, Cham. (2021)
10. Jones, J.: Characteristics of Pattern Formation and Evolution in Approximations of Physarum Transport Networks. Artificial Life, vol. 16(2), pp. 127–153. (2010)
11. Jordà, S.: Multi-user Instruments: Models, Examples and Promises. NIME, Vancouver. pp. 23–26. (2005)
12. Miranda, E. R., Adamatzky, A., Jones, J.: Sounds synthesis with slime mould of Physarum Polycephalum. Journal of Bionic Engineering, vol. 8(2), pp. 107-–113. (2011)
13. Müller, M. W.: petra for max. `http://circuitmusiclabs.com/projects/petra-for-max/` (2016)
14. Nakagaki, T., Yamada, H., Toth, A.: Intelligence: Maze-Solving by an Amoeboid Organism. Nature, vol. 407, p. 470. (2000)
15. Oliveros, P.: Deep Listening: A Composer's Sound Practice. iUniverse, Lincoln. (2005)
16. Oliveros, P., Weaver, S., Dresser, M., Pitcher, J., Braasch, J., Chafe, C.: Telematic Music: Six Perspectives. Leonardo Music Journal, p. 19. (2009)
17. Reid, C. R., Latty, T., Dussutour, A., Beekman, M.: Slime mold uses an externalized spatial "memory" to navigate in complex environments. Proceedings of the National Academy of Sciences, vol. 109(43), pp. 17490-–17494. (2012)
18. Rozenberg, G.: Handbook of Natural Computing. Springer, Berlin (2012)
19. Small, C.: Musicking: The Meanings of Performing and Listening. Wesleyan University Press, Middletown. (1998)
20. Vogel, D., Gautrais, J., Perna, A., Sumpter, D., Deneubourg, J., Dussutour, A.: Transition from Isotropic to Digitated Growth Modulates Network Formation in Physarum polycephalum. Journal of Physics D: Applied Physics, vol. 50, iss. 1. (2016)
21. Weinberg, G.: Interconnected Musical Networks - Bringing Expression and Thoughtfulness to Collaborative Group Playing. MIT, Cambridge. (2003)