A Sonification System for Process Monitoring as Secondary Task

Tobias Hildebrandt
Faculty of Computer Science
University of Vienna, Austria
tobias.hildebrandt@univie.ac.at

Thomas Hermann
Ambient Intelligence Group
CITEC, Bielefeld University, Germany
thermann@techfak.uni-bielefeld.de

Stefanie Rinderle-Ma
Faculty of Computer Science
University of Vienna, Austria
stefanie.rinderle-ma@univie.ac.at

***Abstract*—This paper presents a novel process monitoring system to explore and evaluate the potential of real-time sonifications (i.e. non-verbal auditory representations) for supporting awareness of process states and for detecting and resolving critical process situations. Different from established auditory alarms and warnings, our sonifications convey analogue, i.e. continuous information in form of a process-data-driven soundscape that can be easily blended out in favor of a primary task, yet which is designed to attract the user's attention even before things become critical. We argue that a subsymbolic, implicit and rich display connects better to the human sense to establish auditory categories and develop sensitivity to changes within and in between. In consequence, users may profit from prerational automatic information processing mechanisms so that the cognitive resources remain free for another 'primary task'. Our system allows to generate a prototypical process in real-time, and we present it by two novel sonification approaches. It encompasses software components to log user performance on the monitoring task, and to engage the user in a primary task. We discuss our sonification designs and first experiences with test users, and outlook on studies that are planned to be conducted using our system.**

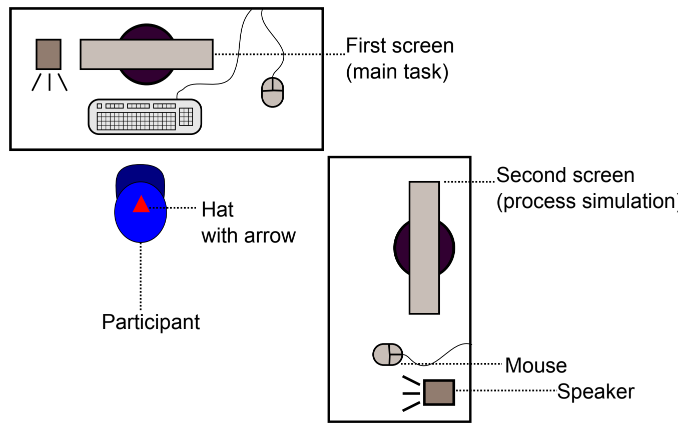

Fig. 1: The operator's workplace for user studies on the utility of sonifications. The monitoring task is presented on the second screen, right of the user's frontal orientation. The angle between the screens is deliberately larger than 90°.

I. Introduction

While complexity and importance of processes for the success of organizations are rising, *real-time monitoring* is becoming more and more important, especially for high-frequency executions of activities in parallel processes. Many companies therefore employ specific applications that offer an overview of running processes by visual means such as charts and graphs. This of course comes with some drawbacks: users cannot keep an eye on these visualizations at all times while performing other tasks. Thus, they either have to dedicate their full attention to process monitoring, or they risk to miss time-critical events or alerts or receive them too late. In situations where auditory alerts are used, these alarms/warnings normally come only after a condition has become critical, which is a rather problem-solving, instead of a problem-prevention style of operation. We argue that these and other drawbacks of state-of-the-art process monitoring techniques can be tackled by supplementing common visualization techniques with methods from the area of *sonification*.

Sonification is the systematic, objective and reproducible transformation of data into (mostly non-speech) sound [1]. Different from scientific visualizations, sonifications just aim at another sensory channel (ear vs. eye). Our human listening skills are different than visual pattern recognition skills; for instance, we have a higher temporal resolution in listening than in seeing, and can better cope with overlaid information in the auditory than in the visual domain. Furthermore our listening is omnidirectional, allowing us to integrate information even if the source is outside our visual focus of attention. This makes sonification very suitable for process monitoring, as it allows operators to focus their visual attention on other tasks while they get informed aurally about their processes' status. Another advantage is that users can habituate to typical soundscapes and remain sensitive to even subtle changes which in turn attract the users' attention. Although some research exists mainly from the area of alerts and alarms, the mechanisms and modes of how more complex sonifications can be used in monitoring settings are largely unknown.

These arguments suggest to support operators in monitoring complex processes by adding data-driven sound. But how? What sounds to use? What sound parameters shall be driven/manipulated by what data conditions? How to design sonifications so that they operate on the periphery of our conscious attention, causing controlled distraction? In this paper we develop a process monitoring system that allows us to study how users perform in situations where the monitoring is a secondary (background) task while the main focus of attention is constantly drawn towards a primary, cognitively demanding task. Our system allows both visual-only and combined audio-visual presentations and thus to compare the performance under these conditions. Furthermore, it allows us to reproduce (and prepare) process situations so that we can assess the utility, acceptance, annoyance, pleasantness, long-term-use compatibility and various other characteristics of competing sonification approaches. Thus, our system offers a versatile test-bed for designing and adapting process monitoring sonifications. Since one area in which real-time monitoring of events and states is especially crucial is manufacturing, our developed system bases on a simulation of a production process as an example domain.

The paper starts with the identification of productive application opportunities for production monitoring in Sec. II, followed by a short review of related work in sonification for process monitoring. This provides requirements and overview which leads to the development of the *SoProMon* system specification presented in Sec. IV. We demonstrate the performance and modes of the system in interaction examples in Sec. V. In the subsequent discussion we focus on evaluation paths for which we regard the SoProMon system as particularly apt.

II. Application opportunities of Sonification for production monitoring

Our auditory perception has several properties that make sonification ideal for process monitoring and offer improvements along the following dimensions:

A. Acceleration of response to critical situations

In order to be able to receive information about critical situations without delay, in state-of-the-art visualization-based process monitoring users have to focus their active attention to their monitoring application at all times. On the other hand, in most real-life situations there are no personnel that are charged to monitor processes in full-time, but instead technicians and supervisors will mainly work on other tasks with process monitoring as a side task. In such cases, auditory alerts will enable users to get immediately informed about critical situations, even without actively paying attention. Furthermore, sound is processed faster than visual signals, allowing us to decrease reaction times, which may be important for time-critical situations that might occur during process execution.

B. Constant awareness of process states with minimal distraction from other tasks

Auditory process monitoring gives operators more physical freedom but it also allows them to focus their visual attention on other tasks while they can at the same time listen to (i.e. aurally monitor) a sonification. As sound can be processed more passively than visuals, auditory process monitoring can enable an unobtrusive perception of process states without distracting users from other tasks. As our auditory perception is able to habituate to regular soundscapes, such a sonification can remain unobtrusive during 'normal' operation, while even small changes in sound over time are able to immediately grab our attention in case of process deviations or undesired process behavior.

C. Anticipating critical situations

Current process monitoring systems are founded on the concept of alerts and alarms, which are conveyed (either visually, or by using simple auditory signals) whenever e.g. a predefined threshold value has been exceeded. In a production scenario, this could for instance be the case when the stock level of a resource has dropped below a critical level, or when a temperature sensor of a machine measures a critical temperature, indicating imminent machine failure. However, this has several drawbacks: on the one hand, if rules that define trigger values of alerts are defined too *conservative*, i.e. requiring strong evidence before issuing positive classifications, potentially critical situations such as machine failures might occur without issuing an alert. On the other hand, if the values are defined too *liberal*, i.e. risking high false positive rates, the resulting flood of (in many cases unnecessary) alerts and alarms might lead to an information overload of the user. In such situations, frustrated users may decide to ignore the alerts and alarms altogether. Furthermore, in many scenarios engineers are not able to define all states and values that might lead to a critical situation beforehand. But even if all possibly critical situations are covered by alerts and alarms, in most cases operators might prefer to be informed even *before* a situation might become critical, thus being able to anticipate, intervene and avoid the problem. A constant awareness of states and values through an auditory ambient information systems might enable such an anticipation of critical situations.

In factories and plants, machine maintenance experts have been listening to the acoustic patterns machines produce for decades. They are often able to evaluate if a machine is about to break down, or a specific part needs to be replaced soon, by listening to the frequencies and patterns of the sounds a machine produces, a technique referred to as vibration analysis [2]. However, by using sonification methods these inherent properties of audio can be leveraged and made accessible to a wider range of people. On the one hand, sonification can decrease the need for experience that is necessary for vibration analysis significantly, as the data can be aggregated and filtered according to the individual users information needs, while at the same time optimizing the resulting audio for our cognitive abilities. On the other hand, users who need to monitor production processes (like e.g. engineers and supervisors) often work in offices that are distanced from the factory floor, e.g. in control rooms. Such rooms can contain systems that produce a high number of auditory alerts that are often conceived as straining. If the mechanisms of how audio can direct attention are better understood, sonification can help in decreasing the need of such obtrusive auditory alerts while at the same time conveying a better awareness of process states in a more pleasant and less straining way.

III. Research in Sonification for Process Monitoring

In the *ARKOLA Simulation*, the production processes of a bottling plant have been conveyed in a multi-modal representation that combines visual and auditory means [3]. The system sonified events during a production process using real-world recordings of such events, such as spills of liquid. The authors concluded that the auditory feedback helped in diagnosing problems in the production process. Another approach that uses an auditory-based system to monitor industrial production environments has been presented in [4]. Apart from the usage of sonification to monitor production processes, auditory-based systems have also been researched in various other application domains such as computer program execution (e.g. [5]), web server and computer network monitoring (such as the *Peep* system [6]) or business process executions (see [7]).

The existing research on sonification for process monitoring gives evidence that sonification increases the performance of operators. Unfortunately, it does not provide detailed guidance on how to effectively sonify what is useful for guiding attention and what aesthetic and interaction design aspects need to be considered. Furthermore, existing research focuses mainly on scenarios where the monitoring of a specific system is the user's main, or even only task. However, especially in small and medium-sized manufacturing enterprises, potential users such as engineers, maintenance personnel and supervisors typically have to dedicate their main attention to tasks other than process monitoring, while they still need timely information on process changes, updates, deviations or alerts/problems. Auditory process monitoring systems that support such use cases need to be designed differently than ones that require permanent active attention, e.g. more unobtrusive, as they need to be listened to for long periods of time. Furthermore, monitoring as a second task needs different means of attention allocation, a fact that has been catered to by several researchers. Especially in peripheral or serendipitous-peripheral monitoring situations, visual means are not well suited, as pointed out by [8, ch. 18, p. 455]. On the other hand, sonifications are claimed to be able to keep users informed about background activities without being disruptive and that in peripheral monitoring, auditory cues are more useful than visuals and less intrusive. But sonification is not only useful as sensory substitution of visual displays, it can also serve as a supplement during direct monitoring tasks, for instance in case that the visual displays are very complex or if many parameters need to be conveyed simultaneously. Thus, in both direct and indirect monitoring, attention allocation is an important aspect that needs to be considered in designing auditory process monitoring.

Ideally, during normal operation auditory monitoring systems should hardly be actively perceived at all. In cases that require the users' attention, such as exceptional or even potentially dangerous situations, the sonification should nonetheless be able to attract the users' full attention. This leads to a trade-off between awareness and disturbance. Generally, the more information a sonification conveys, the greater is the risk of being disturbing. This trade-off has been, among others, researched by Gaver et al. [3]. There is a wide selection of research that investigates how sonifications can guide the users' attention (e.g. [9]). McClimens and Brock [10] investigated the effectiveness of auditory displays to improve dual-task performance. Kilander and Lönnqvist [11] suggest that natural sounds are better accepted as part of the environment, especially if they constitute a *constant murmur* instead of a stream of individual sounds. Caldwell and Viraldo [12] suggest to investigate the usage of complex sonification to convey state-based information in control-room scenarios in order to tackle the problem of information overload and alarm flooding.

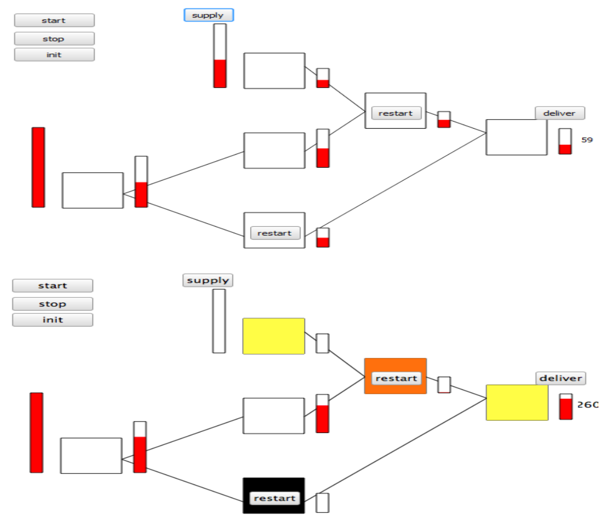

Fig. 2: Visualization of a SoProMon test process: filling levels are depicted in red, machines include buttons to restart, further buttons allow the management of buffers (supply / deliver).

IV. The SoProMon system for evaluating dual-task sonifications

Our system provides a generic testbed for evaluating different information systems for process monitoring with a focus on the before-mentioned challenges of attention allocation in dual-task settings. For that it was first necessary to find suitable and at the same time realistic primary (main) and secondary (monitoring) tasks. In a real-life setting, users interrupt their main tasks in more or less frequent intervals, to observe the status of production. If necessary, these users then perform certain process-related actions such as ordering new raw material or repairing a broken machine.

In order to design the scenario as lifelike as possible, to provide the users with a natural motivation to observe the process sonifications and to grab their immediate attention in severe cases, we decided to base the sonification on a real-time simulation of a simple production process. This production simulation consists of several production steps that partially run in parallel and require input of one or more previous production steps. The simulation has been designed in such a way as to require several actions by the users, in order to measure the performance of auditory monitoring in attention allocation and in interrupting the users during their main task. Furthermore, as a delayed or non-existing response of the users to certain situations can influence the simulated process and consequently the total number of units that are being produced, this provides an elegant way to measure the users' monitoring performance. In order to be able to perform an adequate number of user tests and in order not to exhaust the users excessively, we planned a single evaluation time to be between 15 and 20 minutes. Thus, we designed the simulation such that it provides an adequate number of interruptions within the evaluation time. The simulation requires different types of user actions as depicted in Fig. 2 (top).

supply: One of the machines requires the user to refill the resource input of this machine in more or less regular intervals. In order to simulate a realistic environment, each machine contains a random factor that influences the time between an input resource has been taken and the resulting material resource has been produced.

deliver: One of the machines requires the user to clear the output buffer (i.e. initiate a delivery / transport of goods) so that there is new space for assembled goods to be buffered.

restart: Two machines can encounter conditions of malfunction/maintenance stops, which require active attention of the operator. In our case these situations can be anticipated, as they occur regularly every pre-defined number of steps. However, this number is so large, that it is rather surprising / interrupting for the regular work on the primary task. For resolving, the user has to click on the 'restart' button, yet of course this could also be modified so that the operator needs to go to a different place in the monitoring room to resolve the problem (which would probably be more realistic).

In order to discourage users to just perform unnecessary actions (e.g. prophylactically clicking "supply" every once in a while), we introduced a small waiting period after an action has been performed, in which the respective machine is paused. This is a realistic behavior, as real machines often require a short downtime when they are being refilled, and a longer one if they are being maintained or repaired. While the upper screenshot of Fig. 2 shows the *normal* state of simulation, in which all machines are working, the lower screenshot shows a critical state, in which several machines are out of order, either due to empty resource levels or due to maintenance (i.e. "restart") needs.

A. Main task

We identified the following requirements for the main task: (a) it should be rather complex and cognitive demanding, (b) performance should be measurable, specifically task completion time and correctness of the result, ideally in an automated way, (c) the task should be repeated, rather frequent and of short duration, so that we can observe performance effects on small time scales, finally (d) it should be possible to interrupt the task in favor of attending to the processes to be monitored without consequences (i.e. the main task would then just wait). While requirements (a) and (d) aim to design the evaluation in a way that it resembles the real-life working conditions of supervisors as much as possible, (b) and (c) are needed in order to measure the performance and distraction of the test subjects in an effective way. Typical tasks in real-world scenarios may be, depending on the user group, documenting, processing documents such as emails or postal inquiries, or planning/scheduling. These are rather domain-specific and heterogeneous in type. We argue that for the sake of binding the cognitive resources much simpler tasks are already effective to allocate and bind the users' full attention. For the sake of simplicity we selected the *adding numbers* task, which is a simple mental arithmetic task of summing two numbers (each smaller than 50). The result is to be entered to a text field using the computer keyboard. On hitting the return key the task, the result and the timestamp are logged and the next random numbers are drawn and presented. Figure 1 shows the operator's workplace. The task is displayed in the center of the laptop screen which is otherwise empty to reduce distractions.

B. Sonification Design

With the sonification design we aim at meeting the following requirements:

(a) the sonification should provide an awareness of the ongoing process steps, allowing the listener to recognize when the overall process state has significantly changed

(b) the sonification should represent information on the underlying process in a continuous manner, i.e. dimensions of the sound vary in tight connection with underlying data, so that the user can infer and anticipate states

(c) the sonification should be compatible with verbal interaction, i.e. it should enable to engage in verbal conversations with limited / controlled disturbance

(d) the sonification should be unobtrusive and support participants in letting it perceptually disappear in the periphery of attention, same as we can ignore car engine sounds while driving – yet surprising changes shall draw the attention to the sound

(e) it should be possible to discern the sounds from different machines, allowing users to learn and associate distinct sounds with particular machines

(f) the sound should be compatible with the acoustic environment in which the process monitoring situation is embedded.

The requirements for the sonification design result from combining our analysis of productive application opportunities in Sec. II with our understanding of the possibilities of sonification and our defined goals. To fulfill the requirements with particular designs, we employ selected sonification techniques (e.g. Earcons, Auditory Icons, Parameter Mapping Sonification or Model-based Sonification) and optimize the acoustic features and the mapping in iterative design cycles to better meet the requirements (a)–(f).

We started with a first design, which we call *basic process identification sonification* and progress further to our latest design, the *process data driven soundscapes*, which we explain in detail.

C. Basic Process Identification Sonification

A basic assumption is that each machine execution is an elementary (atomic) step that transforms an input into an output situation. The execution can be conceptualized as an event that takes place at a certain time. It seemed straightforward to represent these elementary events to corresponding acoustic events, leading to the perception of sound streams for each machine of rather repeated sound events. Thus, the sound aggregates have both the characteristics of identifiable events and a continuous stream. The rate at which the machines operate can be perceived by the rate of sound events, allowing an intuitive association: the higher the number of sounds per second the higher the execution speed (same as a Geiger counter where radioactivity corresponds to the number of tick sounds).

To distinguish different machines, it is necessary to create machine-specific sounds that are easily discernible, yet coherent in their structure so that they can be perceived as belonging to the same auditory display. As the human listening system already performs (by means of the tonotopic organization of the cochlea) a frequency analysis, we started with short percussive tones, tuned to a set of pitches so that tones can be easily discerned. Specifically, we use a complex timbre synthesized from a source filter model with 4 resonant filters at the fundamental f_0, second harmonic, $3.4{\cdot}f_0$, and the seventh harmonic. There is no strong argument behind these timbre vector, yet it sounds pleasant and a bit like a glass/metal object. Still, in the selection of pitches lies a huge design space as these tones can either be chosen to be equidistant on a pitch perception scale, resulting in equal musical intervals between tones (e.g. a third, or fifth), or they can be selected deliberately to form a coherent musical chord according to a harmonic system. As first starting point, we chose 8 semi-tones between sound streams. However, as pitch is such a salient feature, it would be underused for barely identifying the machine. Instead we slightly increase the pitch depending on the degree to which the machine output buffer reaches the maximum. Perceptually this compares to the increasing of pitch in sound when filling a bottle with water, indicating in an analogue and intuitively understood manner that ‘something runs full’. Specifically, the pitch range spans 7 semi-tones for the range of semi-full to full output buffer.

A critical point is that a cacophony of parallel playing sounds makes it difficult to attend to the relevant sound streams. We thus decided to map the output buffer filling level to the sound level as well, so that empty output buffers (i.e. no problem) result in quiet sound events. Specifically, the sound level increases by 21 dB as the output level increases from 75% to 100%. To enable listeners to anticipate that input buffers run empty, we furthermore add a noise as the transient phase of the sound, mapping emptiness of input buffers to increasing noise levels. Thus, the more noisy the sounds are, the more critical the input situations become. The increase of noisiness becomes thus more and more discernible as buffers run empty. As the sound events have only a single sound level, we use this variable as indicator of any kind of criticality: individual mappings for the cases ‘filling output buffers’, ‘emptying of input buffers’ and ‘machine failure’ have a mapping to amplitude, so that the more critical the situation, the higher the amplitude becomes. Finally, the maximum of the three amplitudes is used as overall sound event amplitude. In consequence, sound events corresponding to events where no problem is apparent become quiet, and depending on settings almost inaudible. In case of a failure of a machine, the machine sound is repeated at loud volume and low rate. This is immediately understood as an alarm condition, grabbing for immediate attention. Our design has been iteratively refined by the authors while carefully balancing parameters to achieve subjectively acceptable sound streams for a sustained monitoring situation. However, the applied Earcons are neither really well designed, nor meeting the acceptability threshold for extended (e.g. full day) use. We concluded that more *natural* sounds as encountered everyday a thousand times when interacting with real-world objects would be attractive, and yield better compatibility with long uses. This brought us to the following design.

D. Process-data-driven Soundscapes

Since the previous approach appeared to us as promising in terms of information richness and interpretability yet suboptimal in terms of the acceptability of the sonic texture we mostly worked in this design on this issue. Our starting point was to question the way how coherence between acoustic representations of machines is established within the display. Interpreting the process as a ‘virtual world’, a closed ecology which is distinct from our natural acoustic ecology, we see the opportunity to design the virtual acoustic environment based on the same guiding principles that structure natural soundscapes. For instance, in most ecosystems, animal sounds optimized so that animal voices allocate disjunct spectra and thus misclassifications are reduced. One design seed in this direction was to associate different bird motifs with machines, and thus have the running process to yield ’the soundscape of birds’, which is mostly regarded as relaxing and as a nice ambience. However, in the situation of our machine simulation, where events occur at rather constant rate, the sounds are too frequent for complex bird motifs, and using bird chirps soon became exceedingly annoying. Looking at soundscapes in general we see that most everyday sounds in interacting with objects are transient, percussive yet stochastic. For instance, footsteps, closing a door, opening a bottle, etc. So we defined as soundscape the *forest theme*, a soundscape with positive connotations and selected sound events from this area that work well, such as a cracking branch, a bee, a woodpecker, rustle of leaves, water drops and a snippet of a brook sound. The mappings are largely identical to the ones explained in the basic identification sonification above. This design, which emphasizes non-pitched sounds, is to our experience much less obtrusive, and much better capable to fulfill requirements (c), (d), (e) and (f). Since playback rate (and thus implicit pitch) is not as salient as pitch in the other design, (b) is slightly less clear, yet this sacrifice is probably acceptable taken the improvements in the other categories. Of course, these observations are from our subjective impressions and require formal testing in empirical studies, as followed on in Sec V. Based on the concept of *Cognitive Infocommunication Channels* [13], this approach applies both Low-Level Direct Mappings (as all production steps directly result in respective parameterized Auditory Icons) and Structural Mappings, as for instance the mapping of filling buffer levels to pitch is an analogic representation.

To demonstrate the sonifications we provide an example video as supplementary material [1]. The video S1 shows the visualization of the process. We suggest first to watch the video with muted audio so that the ‘silent’ visual-only monitoring can be understood. As discussed in Sec. IV all detailed information is rendered in the visual display, including the emptying of buffers (see time code: 0’50”), the anticipation of machine failures (see time code: 1’50”) and the filling of buffers (see time code: 3’00”). Now, please un-mute the sound and look and experience how sound augments the perception of the process. The sonification is the *process-data-driven soundscape*. In general, the audio volume is being adjusted by

[1] see http://doi.org/10.4119/unibi/2695709

the user before starting the simulation, according the hardware that is being used and the users' hearing capabilities. In the video example, the sound has been recorded with a very high volume, so it is recommended to adjust it to a low volume which is above just audible. See also whether you can recognize which sound belongs to what machine. After 1-2 times watching the video, we recommend to look away from the screen and try to identify from listening alone, when it is time to attend to the process. From informal tests with few subjects we learned that it takes a short time to become familiar with the sonification and how it represents the information as sound.

For comparison, we provide the sonification example S2 [2], which presents the first design *Basic Process Identification Sonification*, using the pitched sounds to identify machines. Again, from a handful of test listeners so far we received the feedback that this is much more obtrusive than the process-data-driven soundscape design. Figure 3 shows a visualization of log data of selected buffer levels as well as the condition of a machine. The data have been recorded during a pre-test experiment. The three different types of user actions ('supply', 'deliver', 'restart') are clearly visible.

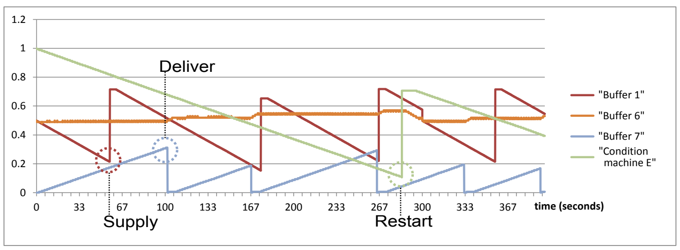

Fig. 3: Visualization of experimental log data

V. Conclusion

We have introduced a novel system to simulate and sonify processes as secondary tasks while offering a system to measure primary task performance. Our system is implemented in the sound programming language SuperCollider3 and thus a solid basis for experiments to assess the utility, acceptability and performance of sonification as complement to visual process monitoring aids, and importantly, as a means to conduct basic research in sonification to compare alternative designs. We have introduced two specific sonification strategies and discussed the resulting sounds, which set the starting point for subsequent studies. These and their results will be presented in subsequent papers. First pre-tests suggest that with the developed system, users are indeed able to infer states and the need to intervene, however, whether it improves performance over visual-only and simple auditory-warning based systems in a significant manner still has to be proven. Specifically, we plan to conduct extensive evaluations that compare three groups: one that uses a visual-only system, another one using our developed multi-modal approach and a third one combining the visual system with simple auditory alerts. Between these three groups, we not only try to find out if and how the sonification influences or distracts people from their main task (in this case adding numbers), but also if the proposed system can increase the users' performance in the monitoring task. The performance in the main task is planned to be evaluated by the number of correctly solved additions, while the performance in the second task, among other factors, will be based on the number of produced units. In terms of sonification technique development, we try to improve the current sound design in a way that makes it more natural, and thus less straining to listen to for a longer period of time, e.g. by using a set of slightly different sound recordings for each of the machines in order to create variety.

[2] see http://doi.org/10.4119/unibi/2695709

Acknowledgment

We thank the German Research Foundation (DFG) and the Center of Excellence 277 Cognitive Interaction Technology (CITEC). This work was partially supported by the Commission of the European Union within the ADVENTURE FP7-ICT project (Grant agreement no. 285220).

References

[1] T. Hermann, "Taxonomy and definitions for sonification and auditory display," in *Proceedings of the 14th International Conference on Auditory Display*, Paris, France, 2008.

[2] J. T. Renwick and P. E. Babson, "Vibration analysis — a proven technique as a predictive maintenance tool," *IEEE Transactions on Industry Applications*, vol. IA-21, no. 2, pp. 324–332, 1985.

[3] W. W. Gaver, R. B. Smith, and T. O'Shea, "Effective sounds in complex systems: the ARKOLA simulation," in *Proc. of the SIGCHI conference on Human factors in computing systems: Reaching through technology (CHI'91)*. ACM, 1991, pp. 85–90.

[4] M. Rauterberg and E. Styger, "Positive effects of sound feedback during the operation of a plant simulator," in *Selected papers from the 4th International Conference on Human-Computer Interaction*, ser. EWHCI '94. London: Springer, 1994, p. 35–44.

[5] D. Sonnenwald, B. Gopinath, G. Haberman, I. Keese, W.M., and J. Myers, "InfoSound: an audio aid to program comprehension," in *Proceedings of the Twenty-Third Annual Hawaii International Conference on System Sciences*, vol. ii, 1990, pp. 541 –546 vol.2.

[6] M. Gilfix and A. L. Couch, "Peep (the network auralizer): Monitoring your network with sound," in *Proceedings of the 14th USENIX conference on System administration*, ser. LISA '00. Berkeley: USENIX Association, 2000, pp. 109–118.

[7] T. Hildebrandt and S. Rinderle-Ma, "Toward a sonification concept for business process monitoring," in *19th International Conference on Auditory Display (ICAD 2013)*, Lodz, 2013, p. 323–330.

[8] P. Vickers, "Sonification for process monitoring," in *The Sonification Handbook*. Logos Berlin, Dec. 2011.

[9] F. J. Seagull, C. D. Wickens, and R. G. Loeb, "When is less more? attention and workload in auditory, visual, and redundant patient-monitoring conditions," *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 45, no. 18, pp. 1395–1399, 2001.

[10] B. McClimens and D. Brock, "Sonification of NRL dual-task data," in *16th International Conference on Auditory Display (ICAD2010)*, Washington, D.C., USA, 2010.

[11] F. Kilander and P. Loennqvist, "A whisper in the woods - an ambient soundscape for peripheral awareness of remote processes," in *Proceedings of the 8th International Conference on Auditory Display (ICAD2002)*, Kyoto, Japan, 2002.

[12] B. S. Caldwell and J. E. Viraldo, "Improving control room state awareness through complex sonification interfaces," in *Human Interface and the Management of Information. Information and Knowledge in Applications and Services*, ser. Lecture Notes in Computer Science. Springer International Publishing, 2014, no. 8522, pp. 317–323.

[13] P. B. Adam Csapo, "A unified terminology for the structure and semantics of CogInfoCom channels," *Acta Polytechnica Hungarica*, vol. 9, pp. 85–105, 2012.