

Kent Academic Repository

Full text document (pdf)

Citation for published version

Axon, Louise and Creese, Sadie and Goldsmith, Michael and Nurse, Jason R. C. (2016) Reflecting on the Use of Sonification for Network Monitoring. In: Tenth International Conference on Emerging Security Information, Systems and Technologies, July 24 - 28, 2016, Nice, France.

DOI

Link to record in KAR

<http://kar.kent.ac.uk/67489/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Reflecting on the Use of Sonification for Network Monitoring

Louise Axon, Sadie Creese, Michael Goldsmith, Jason R. C. Nurse

Department of Computer Science, University of Oxford,
Parks Road, Oxford, UK

Email: {louise.axon, sadie.creese, michael.goldsmith, jason.nurse}@cs.ox.ac.uk

Abstract—In Security Operations Centres (SOCs), computer networks are generally monitored using a combination of anomaly detection techniques, Intrusion Detection Systems (IDS) and data presented in visual and text-based forms. In the last two decades significant progress has been made in developing novel sonification systems to further support network monitoring tasks. A range of systems has been proposed in which sonified network data is presented for incorporation into the network monitoring process. Unfortunately, many of these have not been sufficiently validated and there is a lack of uptake in SOCs. In this paper, we describe and reflect critically on the shortcomings of traditional network-monitoring methods and identify the key role that sonification, if implemented correctly, could play in improving current monitoring capabilities. The core contribution of this position paper is in the outline of a research agenda for sonification for network monitoring, based on a review of prior research. In particular, we identify requirements for an aesthetic approach that is suitable for continuous real-time network monitoring; formalisation of an approach to designing sonifications in this space; and refinement and validation through comprehensive user testing.

Keywords—*Sonification; Network Security; Anomaly Detection; Network Monitoring; Research Agenda.*

I. INTRODUCTION

The monitoring capabilities of the Security Operations Centres (SOCs) within and on behalf of organisations are vital to enterprise cybersecurity. SOCs are run by security analysts who monitor and aim to maintain network and systems security. In the face of a constantly evolving set of threats and attack vectors, and changing business operations, there is a requirement for effective monitoring tools in SOCs.

One of the key challenges that SOCs face in monitoring large networks is the huge volume of data that can be present on the network. This is both the data created by the day-to-day operations of the enterprise, and data created by security tools. For real-time monitoring, tools that present this data in a form that can be processed quickly are essential. Intrusion Detection Systems (IDS) and visualisations are general examples of classes of tools that are widely used to convey information pertaining to network security in a form that can be easily understood by analysts. The anomaly detection techniques that usually underlie such tools have certain limitations, and can produce false positive and negative results [1] [2]. Detecting attacks, and recognising which risks must be prioritised over other attacks and malign activities is difficult, and the degree of inaccuracy of detection systems can make it even more so.

Over the last two decades, the incorporation of sonification of network data into the monitoring activity of SOCs has been considered. Sonification is the presentation of data in a sonic (generally non-speech) form. Some of this prior work has provided sound evidence of the role sonification could

play in improving SOC monitoring capabilities. It has already been shown, for example, that using sonification techniques enables users to detect false positives from IDS more quickly [3]. Based on the state of the art, there are, however, clear requirements for further research and testing to validate the usefulness of sonification for efficient network monitoring, and to develop appropriate and effective sonifications to enhance network monitoring capabilities.

In this paper, we review the major developments in sonification and multimodal systems for network monitoring over the last two decades. In particular, we consider approaches to design and user testing, since we have identified these as two areas in which further research is needed. A key contribution of this paper is a consolidation of existing work, and an analysis of the approaches taken thus far to sonifying network monitoring tasks. We also derive and outline a research agenda for advancing the field; specifically, we aim to highlight directions in which work is needed in order to validate and improve sonification techniques for network monitoring tasks in SOCs. We identify a requirement for comprehensive assessment of the extent to which, and ways in which, sonification techniques can be useful for network monitoring tasks in SOCs through extensive, in-context user-specific testing. We also identify a requirement for the development of aesthetic sonifications appropriate for use in continuous network monitoring tasks, and a requirement for a formalised approach to designing sonifications for network monitoring.

This paper is structured in four further sections to achieve the research aims set out above. In Section II, we present traditional approaches to network monitoring and detail their shortcomings. In Section III, we review prior work in using sonification for network monitoring, and highlight outstanding challenges in the field. Section IV presents a research agenda for sonification for real-time network monitoring. In Section V, we give our conclusions and future work.

II. TRADITIONAL APPROACHES TO NETWORK MONITORING

Network monitoring is generally conducted by security analysts, who observe activity on the network – usually using a variety of tools – in order to detect security breaches. According to the UK government’s information security breaches survey for companies across the UK in 2015, 90% of large organisations reported that they had suffered a security breach, the median number of security breaches for a large organisation was 14, and the average cost to a large organisation for its worst security breach of the year was £1.46m–£3.14m [4]. In the face of such frequent and potentially costly breaches, network monitoring and attack detection capabilities are of extremely high importance.

A variety of tools are used in network monitoring: IDS, Intrusion Prevention Systems (IPS), visualisations, textual presentations, and firewalls are some of the tools with which analysts conduct their monitoring tasks. The subject of our research paper is primarily detection, rather than prevention capabilities. We therefore focus on two key approaches to the detection phase – IDS and visualisation – and on the anomaly detection techniques that often underlie these.

Anomaly detection techniques describe methods for the detection of changes in systems that may be of interest from a monitoring perspective. In anomaly detection, the state of the network is monitored and compared with a specified “normal” baseline. Anomalous activity is that which exceeds an acceptable threshold difference from this baseline. Anomaly detection often informs the output of IDS and visualisations. There are several reports reflecting on the state of the art in anomaly detection techniques: [1] [5] [6]. In general, we can divide anomaly detection methods into three categories [1] [7]: detection methods based on statistics, in which values are compared against a defined acceptable range for deviation [8] [9]; detection methods based on Knowledge Systems, in which the current activity of the system is compared against a rule-based “normal” activity [10]; and detection methods based on Machine Learning, automated methods in which systems learn about activities and detect whether these are anomalous through supervised or unsupervised learning [5] [11].

Network monitoring is largely based on alerts given by IDS. Many IDS have been based on Denning’s model [12]. In general, there are two types of IDS. Statistical anomaly-based IDS monitor network traffic, and compare it against an established baseline (based on bandwidth, protocols, ports, devices, and connections that are “normal”). Signature-based IDS compare packets monitored on the network against a database of signatures/attributes from known malicious threats [1]. Leading SOCs typically craft their own signatures, defined by analysts in the form of rules. Recent advances automate the collection and analysis of data from a range of sources such as logs and IDS alerts using novel Machine Learning and Data Mining approaches.

Much of the presentation of network monitoring data is conveyed through visualisation systems. There are a number of recent surveys of the state of the art in visualisation techniques for security monitoring. Conti *et al.* in [13] and Zhang *et al.* in [14] present reviews as of 2007 and 2012 respectively, reporting research into improving graphical layout and user interaction techniques [15] [16]. Visualisations generally work by mapping network data parameters to visual parameters, such that analysts can observe the changes in the visualisation presented and from this deduce changes in, and information about, the network. The design of effective visualisation involves identifying mappings that represent the data in an intuitive way that can be understood by security and network analysts, in SOCs for example, without inducing cognitive overload, and can convey as effectively as possible any information pertaining to the security of the computer network.

There are certain drawbacks to current approaches to the monitoring and analysis of security data. Anomaly detection techniques can be unreliable or inaccurate, and may produce false positives and false negatives [1] [2]. A shortcoming of existing visualisation-based network monitoring systems is the requirement that operators dedicate their full attention to the

display in order to ensure that no information is missed – for real-time monitoring especially – which can restrict their ability to perform other tasks. Furthermore, the number of visual dimensions and properties onto which data can be mapped is limited [17].

Based on these shortcomings in existing monitoring techniques, we identify ways in which monitoring capability in SOCs might be improved. While many promising advances have been made recently in novel data analytics approaches in particular, we highlighted that network monitoring systems do not always produce reliable outputs. It is, therefore, important that the human operator has situational awareness and an understanding of the network state, in order that he can interpret the alerts given by the detection systems used, and accurately decide their validity. Such awareness could also enable analysts to detect patterns, recognise anomalous activity and prioritise risks differently to their systems. Techniques that provide analysts with a continuous awareness of the state of the network require further investigation. Research is also needed into novel methods for improving the presentation of network data, the main technique for which is currently visualisation. In particular, it is important to design representations of large volumes of network data that are as easy as possible for analysts to use, understand and act on.

III. NETWORK MONITORING USING SONIFICATION

We believe that sonification could address the requirements with which we conclude Section II in a number of ways. Presenting network data as a continuous sonification may improve analysts’ awareness of the network state, and furthermore may enable the analyst to detect patterns in the data, acting as a human anomaly detector of sorts. These are both areas for investigation and are detailed as research questions in Section IV. Sonification could also offer a solution to the shortcomings of visualisation techniques for network monitoring, as another human interface alongside the visualisation, using a different sense. Firstly, sound can be presented for peripheral listening – a secondary task – and, if designed appropriately, engage the listener’s attention as required, allowing operators to perform other tasks in the meantime; secondly, using sound offers another set of dimensions in addition to visual dimensions to which data can be mapped.

In this section, we present background on sonification generally. Following this, we review prior work in the application of sonification to network security monitoring, and in multimodal systems (combining visuals and sound) for network monitoring.

A. Sonification: A Background

Sonification is used in numerous fields, such as financial markets, medicine (Electroencephalography (EEG) monitoring [18], image analysis [19]) and astronomy. User testing has validated that the presentation of sonified data can improve certain capabilities in a number of applications: improved accuracy in monitoring the movement of volatile market indices by financial traders [20], and improved capabilities for exploratory analysis of EEG data [21], for example.

A variety of techniques and guidelines have been developed for the design and implementation of sonification [22] [23] [24] [25]. Throughout sonification literature there are three main approaches recognised: earcons/event-based sonification (discrete sounds representing a defined event), parameter mapping

sonification (in which changes in some data dimensions are represented by changes in acoustic dimensions), and model-based sonification (in which the user interacts with a model and receives some acoustic response derived from the data).

The current state of the art in sonification for network and server monitoring is summarised in Hildebrandt [17], in which systems for sonification of computer security data are identified, in various stages of maturity. It is concluded that there is a lack of formal user and usability testing, even in those systems that are already fully developed [26] [27] [28]. Our work differs from Hildebrandt in two key ways. Firstly, while Hildebrandt’s survey gives an overview of the design approaches taken in some existing sonification systems, our survey provides much greater detail on the sonification design of existing systems in terms of sonification techniques, sound mapping types, the network data and attack types represented and the network monitoring scope. Secondly, we make recommendations in our research agenda for advancing sonification system design for the network monitoring context through aesthetics and formalisation, as well as defining the research questions to be answered through user testing.




Figure 1. A summary of the existing relationship between traditional monitoring techniques and their potential relationship with sonification systems in SOCs.

Figure 1 shows the existing relationship between raw data, anomaly detection techniques, IDS and visualisations in SOC monitoring, and the position we envisage sonification might take in this setup.

B. Applications of Sonification to Network Monitoring

PEEP, a “network auralizer” for monitoring networks with sound, is presented in [26]. PEEP is designed to enable system administrators to detect network anomalies – both in security and general performance – by comparing sounds with the sound of the “normally functioning” network. The focus of PEEP is on the use of “natural” sounds – birdsong, for example – in sonifying network events. Recordings are mapped to network conditions (excessive traffic and email spam, for instance), and are played back to reflect these conditions. Abnormal events are presented through a change in the “natural” sounds. PEEP represents both network events (when an event occurs it is represented by a single natural sound) and network state (state is represented through sounds played continuously, which change when there is a change in some aspects of the state, such as average network load). There is, however, no validation of the performance of PEEP and its usefulness for monitoring networks.

The Stetho network sonification system is given in [29]. Stetho sonifies network events by reading the output of the Linux tcpdump command, checking for matches using regular

expressions, and generating corresponding Musical Instrument Digital Interface (MIDI) events, with the aim that the system creates sounds that are “comfortable as music”. The aim of Stetho is to convey the status of network traffic, without a specific focus on anomaly detection. The research includes an evaluation experiment in which the Stetho system is used – users’ ability to interpret the traffic load from the sounds generated by Stetho is examined. The experiment shows that this monitoring information can be recognised by users from the sounds created by Stetho; however, only four users are involved in the evaluation experiment.

Network Monitoring with Sound (NeMoS) is a network sonification system in which the user assigns network events, and the system then associates these events with MIDI tracks [30]. The system is designed to allow monitoring of different parts of a potentially large network system at once, with a single musical flow representing the whole state of the part of the system the system manager is interested in. The focus is not on network security but on monitoring network performance in general; printer status and system load, for example, can be represented through two different sound channels.

More recently, Ballora *et al.* look to create a soundscape representation of network state which aids anomaly detection by assigning sounds to signal certain types and levels of network activity such as unusual port requests [31] (“soundscape” definition given by Schafer [32]). The concept is a system capable of combining multiple network parameters through data fusion to create this soundscape. The fusion approach is based on the JDL Data Fusion Process Model [33], with characteristics of the data assigned to multiple parameters of the sound. The authors aim, firstly, to map anomalous events to sound and, secondly, to represent the IP (Internet Protocol) space as a soundscape in which patterns can emerge for experienced listeners. No validation is carried out as to the usefulness of the system in network anomaly detection tasks.

Vickers *et al.* use a soundscape approach to sonify meta properties of network traffic data [34]. The aim of the system is to alert the system administrator of abnormal network behaviour with regard to both performance and security; it is suggested, for example, that a distributed denial-of-service (DDoS) attack might be recognisable by the system’s representation of an increase in certain types of traffic. There is, however, no evaluation of users’ ability to recognise such information using the system. Vickers *et al.* then extend that work to further explore the potential for using sonification for network situational awareness [35]. For this context, i.e., continuous monitoring for network situational awareness – it is argued that solutions based on soundscape have an advantage over other sonification designs in this context, and that there is a need for sonifications that are not annoying or fatiguing and that complement the user’s existing sonic environment.

A soundscape approach is also adopted in the InteNtion system [27] for network sonification. Here, network traffic analysis output is converted to MIDI and sent to synthesisers for dynamic mixing; the output is a soundscape composed by the network activity generally rather than the detection of suspicious activity specifically. It is argued that the system could be used to help administrators detect attacks; however this is not validated through user testing. DeButts is a student project available online in which network data is sonified with the aim of aiding security analysts to detect anomalous

incidents in network access logs [36].

García-Ruiz *et al.* investigate the application of sonification as a teaching and learning tool for network intrusion detection [37] [38]. This work includes an exploratory piece in which information is gathered regarding the subjects' preferred auditory representations of attacks. Sonification prototypes are given for the mapping of log-registered attacks into sound. The first uses animal sounds – auditory icons – for five different types of attack (“guess”, “rcp”, “rsh”, “rlogin”, “port-scan”); the second uses piano notes at five different frequencies as earcons to represent the five types of attack. Informal testing was carried out for these two prototypes, and suggested that the earcons were more easily identifiable, while the subjects could recall the attack types more easily using the auditory icons. While this is a useful start to comparing approaches to sonification design for network data, the mappings tested are limited, and further research is required into mappings involving other sound and data types.

Systems have been proposed to sonify the output of existing IDS, and to act as additions to the function of these systems. The CyberSeer [39] system uses sound to aid the presentation of network security information with the aim of improving network monitoring capability. Sound is used as an additional variable to data visualisation techniques to produce an audio-visual display that conveys information about network traffic log data and IDS events. The requirement for user testing to establish the most effective audio mappings is recognised, but no testing is carried out.

Gopinath's thesis uses JListen to sonify a range of events in Snort Network Intrusion Detection System to signal malicious attacks [3] (Snort is a widely used open-source network intrusion detection system for UNIX derivatives and Windows). The aim is to explore the usefulness of sonification in improving the *accuracy* of IDS; usability studies indicate that sonification may increase user awareness in intrusion detection. Experiments are carried out to test three hypotheses on the usability and efficacy of sonifying Snort. The findings are: musical knowledge has no significant effect on the ability of subjects to use the system to find intrusions; sonification decreases the time taken to detect false positives; immediate monitoring of hosts is possible with a sonified system. As noted in Hildebrandt [17], however, the comparison is somewhat biased since the control group without auditory support had to conduct the tasks by reading log files, without access to the visualisation-based tools to which those tested with auditory support had access.

Multimodal systems, that combine visualisation and sonification for network monitoring, have also been explored. Varner and Knight present such a system in [40]. Visualisation is used to convey the status of network nodes; sonification then conveys additional details on network nodes selected by the user. This multimodal approach is useful because it combines advantages of the two modalities – the spatial nature of visualisation, and the temporal nature of sonification – to produce an effective and usable system. García-Ruiz *et al.* describe the benefits and pitfalls of using multimodal human-computer interfaces for analysing intrusion detection in [41]. A sonification method is proposed for network intrusion detection systems (NIDS) as part of a multimodal interface, to enable analysts to cope with the large amounts of information contained in network logs.

Qi *et al.* present another multimodal system for detecting intrusions and attacks on networks in [42]; distinctive sounds are generated for a set of attack scenarios consisting of denial-of-service (DoS) and port scanning. The same approach is adopted by Brown *et al.* [43]: the bit-rates and packet-rates of a delay queue are sonified in a system for intrusion detection. The sounds generated by the system, which maps bit rate and packet rate to sound, are tested (not tested on users, but listened to by the authors) for DoS and port scanning attack scenarios. It is concluded that the sounds generated could enable humans to recognise and distinguish between the two types of attack. However, user testing is needed to validate this conclusion and investigate the extent to which this approach is effective.

NetSon [28] is a system for real time sonification and visualisation of network traffic, with a focus on large-scale organisations. In this work, there are no user studies, but the system is being used at Fraunhofer IIS, a research institution, who provide a live web stream of their installation [44]. Microsoft have a multimodal system, *Specimen Box*, for real-time retrospective detection and analysis of botnet activity. It has not yet been presented in a scientific publication, but description and videos of the functioning system are presented online [45]. The system has not been subject to formal evaluation, but is used in operations at the Microsoft Cybercrime Centre.

Mancuso *et al.* conduct user testing to assess the usefulness of sonification of network data for military cyber operations [46]. The aim of the testing, in which participants were tasked with detecting cyber attacks using either a visual display only, or both visual and sonified displays, was to assess the extent to which sonification can improve the performance and manage the workload of, and decrease the stress felt by, operators conducting cyber monitoring operations on military networks. The testing results show that the sonifications did not affect the performance, workload or stress. However, only one method of sonifying the data was tested, in which each possible source and destination IP address was represented by a different instrument and note, and the loudness increased if a threshold packet size was exceeded. The results do not, therefore, show that sonification does not affect performance, stress and workload in this context, but demonstrate only that this particular method of sonifying the data is ineffective.

C. Outstanding Challenges

In Table I, we summarise the sonification systems previously developed (solutions for which full systems or prototypes have been developed) for network monitoring; from this we have identified three key areas in which research is lacking: user testing, sonification aesthetics, and formalisation of an approach to designing sonification systems for network monitoring.

In general, a weakness in the articles in which user studies are conducted is the small number of users involved. Table I shows that little user testing has been carried out, and of that which has, none has specifically targeted security analysts, and none has been conducted in a SOC environment. Table I shows also that there has been little (and no comprehensive) evaluation of the usefulness of existing sonification systems for network security monitoring tasks. Extensive user testing is required in this field to validate the usefulness of the approach and of proposed systems, and to refine the sonification design.

The systems listed vary in the data they represent. Some map raw network data to sound, some map the output of IDS

TABLE I. REVIEW OF APPROACHES TO AND USER TESTING IN EXISTING SONIFICATION SYSTEMS FOR NETWORK MONITORING, ORDERED BY YEAR.

Author	Year	Sonification approach description	User testing	Number of participants	Nature of participants	Network data type mapped	Sound type	Sonification technique	Monitoring scope	Evaluates usefulness for security monitoring?	Multimodal
Gilfix [26]	2000	"Natural" sounds mapped to network conditions	N			Raw data	Natural (wildlife and nature) sounds	Parameter-mapping; soundscape	Anomaly detection: conditions such as high traffic load and email spam are mapped to sound	N	N
Varner [40]	2002	Multimodal system: visualisation conveys status of network nodes; sonification conveys additional details on network nodes selected by the user	N			Not specified	Not specified	Not specified	Network attack detection	N	Y
Kimoto [29]	2002	Maps parameters of sound to raw network data	Y	4	Subjects familiar with network administration	Raw data	Musical	Parameter-mapping	General network activity	N	Y
Malandrino [30]	2003	Associates MIDI tracks to user-defined network events	N			Raw data	Musical	Event-based	Network performance	N	N
Gopinath [3]	2004	Instrument and pitch mapped to IDS alert intrusion type	Y	20	Computer Science students and staff	IDS alerts	Real-world (man-made/natural) and musical	Parameter-mapping	Anomaly detection: used alongside IDS logs to improve detection capability	Y	N
Papadopoulos [39]	2004	Combines network events rendered as spatial audio with 3D stereoscopic visuals to form a multimodal representation of network information. Sounds are created in response to changes in data patterns using Gaussian Mixture Modelling	N			A variety	Real-world and musical	Parameter-mapping	Anomaly detection: network data presented for pattern recognition, and IDS output sonified	N	Y
Qi [42]	2007	Maps traffic pattern (based on a classification-based mitigation system) to audio; byte rate and packet rate mapped to frequency and intensity of audio respectively	N			Classification-based attack mitigation system	Musical	Parameter-mapping	Network attack detection (DoS, port scanning)	N	N
El Seoud [38]	2008	Auditory icons (non-instrumental) and earcons (instrumental) mapped to attack type	Y	29	Telematics Engineering students	Logged attacks	Real-world and musical	Event-based	Network attack detection	N	N
Brown [43]	2009	Maps raw network traffic to sound to convey information on network status	N			Raw data	Musical	Parameter-mapping	Network anomaly detection (increase in traffic; HTTP error messages; number of TCP handshakes)	N	N
Ballora [31]	2011	Parameter mapping-based soundscape for overall IP space; obvious sound signals for certain types of activity levels	N			Raw data	Musical	Parameter-mapping; soundscape	Anomaly detection: anomalous incidents sonified, and network state presented to enable human pattern recognition	N	N
Giot [27]	2012	MIDI messages mapped to data outputted by SharpPCap library network traffic analysis; MIDI messages mixed to produce a soundscape	N			Raw data	Musical	Event-based; soundscape	General network activity and attack detection	N	N
deButts [36]	2014	Maps distinct notification tones to anomalous network events; visualises network traffic activity (multimodal)	N			Raw (access logs)	Single tones	Event-based	Anomaly detection: anomalous incidents mapped to sounds	N	Y
Vickers [34]	2014	Parameters of each sound generator (voice) mapped to the log return values for the network's self-organised criticality	N			Network's self-organised criticality	Natural	Parameter-mapping	Network performance and attack detection	N	N
Worrall [28]	2015	Multimodal system for large-scale network data. Maps data parameters and events to sound; parameter mapping sonification approach using melodic pitch structures to reduce fatigue.	N			Raw data	Musical	Parameter-mapping	General network activity	N	Y
Mancuso [46]	2015	Multimodal system for representing data on military networks, in which each source and destination IP is mapped to an instrument and pitch, and the loudness is increased when a packet size threshold is exceeded.	Y	30	Local population and air force base personnel	Raw data	Musical	Parameter-mapping	Network anomaly detection (packet rate threshold, source and destination IPs sonified)	Y	Y

systems, while some aim to map attacks to sounds; however, there is no comparison of the efficacy of these approaches, or of the usefulness of sonic representations of different attack types. The sonification design approaches (event-based, parameter mapping, and soundscape-based) also vary, as do the sound types (natural sounds, sounds that are musically informed) but there is as yet no comprehensive investigation into, or comparison of, the usefulness of these methods. Based on this, we propose that comparative research into the sonification

aesthetics most appropriate to the network monitoring context is crucial, in order to inform sonification design. We further identify a requirement for the development of a formalised approach to designing sonifications in this field, to underpin developments and enable comparison. Next, we outline our detailed research agenda to address these issues.

IV. RESEARCH AGENDA

We present our research agenda in three parts: comprehensive user studies, improved aesthetics, and formalisation.

A. Comprehensive User Studies

Section III indicates that of the proposals made for sonification systems for network monitoring, very few have conducted any user testing, and none have conducted such testing to the extent required for an appropriate understanding of the use of such systems and their suitability for actual deployment in security monitoring situations. As such, we outline a requirement for significantly more in-context user testing of sonifications for network monitoring tasks, carried out with security analysts in SOCs, to inform the design and investigate the advantages and disadvantages of the approach. It is important that sonification systems are tested in the SOC environment, in order to investigate how well they incorporate with the particular characteristics of SOCs – a variety of systems running simultaneously, collaborative working practice, high levels of attention required from workers.

We will conduct user testing to investigate the hypothesis that sonification can improve the network monitoring capabilities of security analysts. This hypothesis is proposed in light of prior work in other fields in which it is proven that certain capabilities can be improved by the presentation of sonified data, as outlined in Section III, and of the limited experimental evidence that shows that sonification can be useful for tasks involving network data specifically [3] [29].

For the validation of sonification as a solution to improving network monitoring capabilities, there are certain key research questions that need to be answered through user testing.

- 1) **To what extent, and in what ways, can the use of sonification improve the monitoring capabilities of security analysts in a SOC environment?** User testing is required to establish, firstly, the extent to which the presentation of sonified network data can improve the analyst's understanding and awareness of the network, as mentioned in Section II. Secondly, the extent to which this awareness can improve the ability of analysts to interact with, and decide the accuracy of the output of, their existing monitoring systems. Further investigation is also needed to establish whether the presentation of network data through sonification can enable analysts to "hear" patterns and anomalies in the data, and in this way detect anomalies more accurately than systems in any cases. Given the strong human capability for pattern recognition in audio representations [47] [48], and for contextualising information, it is plausible that a system that presents patterns in network data may enable the analyst to detect anomalies with greater accuracy than the traditional rule-based systems.
- 2) **Are there certain types of attack and threat that sonify more effectively than others, and what implications does this have for the design of sonification systems for network monitoring?** It is important that user testing is carried out to establish the cases in which sonified network data is most useful for network monitoring tasks. For example, it may be the case that certain types of attacks are better-represented through sonification than others, and that some attacks sound anomalous in a way that is particularly easy for analysts to use while others do not sonify well. Findings on this subject should inform sonification system design by distinguishing the attacks and threats in relation to which sonification performs best, and the areas in which the technique therefore has the

potential to be most effective.

Answers to these questions will provide a greater understanding of the role sonification can play in improving monitoring capabilities in SOCs, the limits of the approach, and the extent to which it can be reliable as a monitoring technique. In conducting this testing, we expect to draw from existing research on conducting user studies in general, and in a security context [49] [50].

B. Improved Aesthetics

While there has been some work in aesthetic sonification, as shown in Section III, it has not been heavily applied in the context of network monitoring. Prior work indicates that sonification aesthetic impacts on its effectiveness. In an experiment comparing sonifications of guidance systems, for example, it was shown that sonification strategies based on pitch and tempo enabled higher precision than strategies based on loudness and brightness [51]. It was also shown in [52] that particular sonification designs resulted in better participant performance in identifying features of Surface Electromyography data for a range of different tasks involved. The aesthetics of the design are an important factor in producing sonifications that are suitable for continuous presentation in this context. In particular, the sounds should be unobtrusive and should achieve a balance in which they are unobtrusive to the performance of other tasks while drawing sufficient attention when necessary to be suitable for SOC monitoring. While there are other techniques that may be useful, we propose an approach to this design that draws on techniques and theories of musical composition. We can draw on work in aesthetic sonification by Vickers [47], and on work in musification, i.e., the design of sonifications that are musical. Some key questions to be answered regarding sonification aesthetics for network security monitoring are given below.

- 1) **Can the development of aesthetic sonifications based on techniques of musical composition alleviate the fatiguing nature of sonifications previously reported, and, secondly, are such sonifications more appropriate for continuous network monitoring tasks?** A drawback to sonification for network monitoring is the fatigue or annoyance that listeners can experience as a result of long-term exposure to sonification [35] [47]. The question of how this can be prevented if sonifications are to be developed that are appropriate for continuous network monitoring, and the suitability of a sonification approach based on techniques of musical composition for the network monitoring context requires investigation.
- 2) **To what extent does musical experience affect the ability of security analysts to use musically-informed sonification systems in network monitoring tasks?** The effect of users' musical experience on their ability to understand and make use of the sonification systems design will require investigation. Here, musical experience refers to the level of prior theoretical and aural musical training attained by the user. For this SOC monitoring context, analysts' use of the systems should not be impaired by a lack of musical experience.

Besides aesthetics, aspects of human perception must influence the design: the prior associations sounds may hold for users and the way in which this affects interpretation; the effect of musical experience on perception. It is important that the

design takes into account factors in perception such as cross-field interference (in which different parameters of sound – pitch and tempo, for example – interact in a way that affects perception of either) and does not induce cognitive overload for the user.

C. Formalisation

To address the requirements above we need an underpinning framework which enables us to architect and experiment with sonifications in a flexible way, utilising heterogeneous models alongside each other in order to compare performance. No such framework currently exists, and we therefore identify that there is a need for a formal grammar for designing sonification for network monitoring. We propose the mathematical definition of a grammar for the representation of network data through parameter-mapping sonification that is derived from the results of user testing in a network monitoring context, techniques of musical composition, and the science of auditory perception.

A formal grammar for designing sonifications for use in the network monitoring context could tailor aspects of sonification design such as cross-field interference to produce sonifications that are appropriate for network monitoring tasks. A simple example is a simultaneous change in two network parameters: a statistically significant increase in traffic load, and messages received from an IP address that is known to be malicious (these two changes would generally be found by the statistical anomaly-based IDS and signature-based IDS, respectively, described in Section II). This could be the result of a DoS attack, and the sonification system should therefore attract the attention of the analyst. Cross-field interference could be leveraged in this case (with a mapping to higher pitch and increased tempo – two sound parameters which interact such that each appears more increased than it really is – for the two data parameters respectively) to ensure that the attack is highlighted by the sonification.

In order for the representation to be unfatiguing, we propose that a rule-based approach to aesthetic sound generation may be appropriate. In particular, a defined formal grammar for representing network data as sound could be applied to a variety of genres of music to generate a set of different-sounding sonifications of the same network data. We hypothesise that with this approach, users could be allowed to move between a set of musical genres at choice, all of which would sonify the network data according to the same grammar, and this could reduce the fatigue caused by the sounds. Below, we give the key questions to be investigated on formalising sonification systems for network monitoring.

- 1) **To what extent can a defined formal grammar produce sonification systems that improve network monitoring capabilities in SOCs?** We believe that a combination of factors, including human sound perception, sonification aesthetics and intuitiveness of mappings, will affect and, if addressed correctly, improve the usability of the sonifications designed. A formal grammar can, if designed correctly, combine the solutions to these requirements in a thorough and considered manner and we therefore hypothesise that such an approach may produce highly usable and useful sonification systems for network monitoring. The extent to which this is the case requires investigation through user testing of the systems produced according to the grammar.

- 2) **Is a rule-based approach to generating a set of different-sounding sonifications of the network data, which enables users to change between musical preferences, appropriate for a network security monitoring context?** This question is in two main parts. Firstly, does the presentation of a choice of different-sounding sonifications alleviate the fatigue induced by continuous sonifications, as reported in prior work? Secondly, does such a presentation affect the usefulness of the sonification for network monitoring tasks in any way? For example, it is viable that the presentation of a number of different-sounding sonifications may cause more confusion for users than the presentation of a single sonification, and that this may detract from the usability of the systems designed. If this is found to be the case, then methods for mitigating this effect, or other approaches to designing unfatiguing continuous sonification systems, will need to be investigated.

V. CONCLUSION AND FUTURE WORK

We conclude that there is a growing requirement for the validation of using sonification in SOCs as a means of improving certain monitoring capabilities. The current state of the art provides evidence of the potential of sonification in advancing network security monitoring capabilities. Systems proposed and in use have been shown to be as effective as, or more effective than, other network monitoring techniques insofar as a limited amount of testing has been performed [17].

We recognise a requirement for development of the field in three main directions. Firstly, the performance of extensive user testing to validate claims about and show the extent of the suitability of sonification for network monitoring. Secondly, the need for aesthetics of the sonification design for continuous network monitoring, and the development of methods for generating sonifications that are unfatiguing, unobtrusive, and yet convey the data to an appropriate extent. Lastly, the formalisation of a design approach – a formal grammar that is defined by the requirements: representation of data, perception of sound, and aesthetics.

As future work, we intend to research the potential for sonification to match, or improve on, the performance of current monitoring systems in the SOC context. We will also define a formal grammar for the design of sonification for network monitoring, based on the results of user testing of mappings; in particular, this grammar will enable the design of rule-based aesthetic sonification for this context by drawing on music theory.

REFERENCES

- [1] P. Garcia-Teodoro, J. Diaz-Verdejo, G. Maciá-Fernández, and E. Vázquez, “Anomaly-based network intrusion detection: Techniques, systems and challenges,” *computers & security*, vol. 28, no. 1, 2009, pp. 18–28.
- [2] S. Axelsson, “The base-rate fallacy and its implications for the difficulty of intrusion detection,” in *Proceedings of the 6th ACM Conference on Computer and Communications Security*. ACM, 1999, pp. 1–7.
- [3] M. Gopinath, “Auralization of intrusion detection system using Jlisten,” *Development*, vol. 22, 2004, p. 3.
- [4] PWC, “2015 Information Security Breaches Survey,” 2015, PWC in association with Infosecurity Europe.
- [5] A. Lazarevic, L. Ertöz, V. Kumar, A. Ozgur, and J. Srivastava, “A comparative study of anomaly detection schemes in network intrusion detection,” in *Proceedings of SIAM International Conference on Data Mining*, 2003, pp. 25–36.

- [6] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Computing Surveys (CSUR)*, vol. 41, no. 3, 2009, p. 15.
- [7] V. Kumar, J. Srivastava, and A. Lazarevic, *Managing cyber threats: issues, approaches, and challenges*. Springer Science & Business Media, 2006, vol. 5.
- [8] D. E. Denning and P. G. Neumann, "Requirements and model for ides—a real-time intrusion detection expert system," Document A005, SRI International, vol. 333, 1985.
- [9] N. Ye, S. Emran, Q. Chen, and S. Vilbert, "Multivariate statistical analysis of audit trails for host-based intrusion detection," *IEEE Transactions on Computers*, vol. 51, no. 7, 2002, pp. 810–820.
- [10] D. Anderson, T. F. Lunt, H. Javitz, A. Tamaru, and A. Valdes, *Detecting unusual program behavior using the statistical component of the Next-generation Intrusion Detection Expert System (NIDES)*. SRI International, Computer Science Laboratory, 1995.
- [11] C. Tsai, Y. Hsu, C. Lin, and W. Lin, "Intrusion detection by machine learning: A review," *Expert Systems with Applications*, vol. 36, no. 10, 2009, pp. 11 994–12 000.
- [12] D. Denning, "An intrusion-detection model," *IEEE Transactions on Software Engineering*, no. 2, 1987, pp. 222–232.
- [13] G. Conti, *Security data visualization: graphical techniques for network analysis*. No Starch Press, 2007.
- [14] Y. Zhang, Y. Xiao, M. Chen, J. Zhang, and H. Deng, "A survey of security visualization for computer network logs," *Security and Communication Networks*, vol. 5, no. 4, 2012, pp. 404–421.
- [15] R. F. Erbacher, K. L. Walker, and D. A. Frincke, "Intrusion and misuse detection in large-scale systems," *Computer Graphics and Applications*, IEEE, vol. 22, no. 1, 2002, pp. 38–47.
- [16] B. Shneiderman, "Dynamic queries for visual information seeking," *IEEE Software*, vol. 11, no. 6, 1994, pp. 70–77.
- [17] S. Rinderle-Ma and T. Hildebrandt, "Server sounds and network noises," in *Cognitive Infocommunications (CogInfoCom)*, 2015 6th IEEE International Conference on. IEEE, 2015, pp. 45–50.
- [18] Z. Halim, R. Baig, and S. Bashir, "Sonification: a novel approach towards data mining," in *Proceedings of the International Conference on Emerging Technologies*, 2006. IEEE, 2006, pp. 548–553.
- [19] T. Hinterberger and G. Baier, "Parametric orchestral sonification of EEG in real time," *IEEE MultiMedia*, no. 2, 2005, pp. 70–79.
- [20] P. Janata and E. Childs, "Marketbuzz: Sonification of real-time financial data," in *Proceedings of the International Conference on Auditory Display*, 2004.
- [21] T. Hermann, "Sonification for Exploratory Data Analysis," Ph.D. dissertation, 2002, Bielefeld University.
- [22] G. Kramer, *Auditory display: Sonification, audification, and auditory interfaces*. Perseus Publishing, 1993.
- [23] A. de Campo, "Toward a data sonification design space map," in *Proceedings of the International Conference on Auditory Display*, 2007, pp. 342–347.
- [24] S. Barrass and C. Frauenberger, "A communal map of design in auditory display," in *Proceedings of the International Conference on Auditory Display*, 2009, pp. 1–10.
- [25] S. Barrass et al., "Auditory information design," Made available in DSpace on 2011-01-04T02: 37: 33Z (GMT), 1997.
- [26] M. Gilfix and A. Couch, "Peep (the network auralizer): Monitoring your network with sound," in *Proceedings of the Large Installation System Administration Conference*, 2000, pp. 109–117.
- [27] R. Giot and Y. Courbe, "Intention–interactive network sonification," in *Proceedings of the International Conference on Auditory Display*, Georgia Institute of Technology, 2012, pp. 235–236.
- [28] D. Worrall, "Realtime sonification and visualisation of network meta-data," in *Proceedings of the International Conference on Auditory Display*, 2015, pp. 337–339.
- [29] M. Kimoto and H. Ohno, "Design and implementation of stetho—network sonification system," in *Proceedings of the International Computer Music Conference*, 2002, pp. 273–279.
- [30] D. Malandrino, D. Mea, A. Negro, G. Palmieri, and V. Scarano, "Nemos: Network monitoring with sound," in *Proceedings of the International Conference on Auditory Display*, 2003, pp. 251–254.
- [31] M. Ballora, N. Giacobe, and D. Hall, "Songs of cyberspace: an update on sonifications of network traffic to support situational awareness," in *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics, 2011, pp. 80 640P–80 640P.
- [32] R. Schafer, *The soundscape: Our sonic environment and the tuning of the world*. Inner Traditions/Bear & Co, 1993.
- [33] O. Kessler et al., "[functional description of the data fusion process]," 1991, Office of Naval Technology Naval Air Development Center, Warminster PA.
- [34] P. Vickers, C. Laing, and T. Fairfax, "Sonification of a network's self-organized criticality," arXiv preprint arXiv:1407.4705, 2014.
- [35] P. Vickers, C. Laing, M. Debashi, and T. Fairfax, "Sonification aesthetics and listening for network situational awareness," in *Proceedings of the Conference on Sonification of Health and Environmental Data*, 2014.
- [36] B. deButts, "Network access log visualization & sonification," Master's thesis, Tufts University, Medford, MA, US, 2014.
- [37] M. Garcia-Ruiz, M. Vargas Martin, B. Kapralos, J. Tashiro, and R. Acosta-Diaz, "Best practices for applying sonification to support teaching and learning of network intrusion detection," in *Proceedings of the World Conference on Educational Multimedia, Hypermedia and Telecommunications*, 2010, pp. 752–757.
- [38] S. El Seoud, M. Garcia-Ruiz, A. Edwards, R. Aquino-Santos, and M. Martin, "Auditory display as a tool for teaching network intrusion detection," *International Journal of Emerging Technologies in Learning (IJET)*, vol. 3, no. 2, 2008, pp. 59–62.
- [39] C. Papadopoulos, C. Kyriakakis, A. Sawchuk, and X. He, "Cyberseer: 3d audio-visual immersion for network security and management," in *Proceedings of the ACM workshop on Visualization and data mining for computer security*. ACM, 2004, pp. 90–98.
- [40] P. Varner and J. Knight, "Monitoring and visualization of emergent behavior in large scale intrusion tolerant distributed systems," Technical report, Pennsylvania State University, 2002.
- [41] M. García-Ruiz, M. Martin, and M. Green, "Towards a multimodal human-computer interface to analyze intrusion detection in computer networks," in *Proceedings of the First Human-Computer Interaction Workshop (MexIHC)*, Puebla, Mexico, 2006.
- [42] L. Qi, M. Martin, B. Kapralos, M. Green, and M. García-Ruiz, "Toward sound-assisted intrusion detection systems," in *On the Move to Meaningful Internet Systems 2007: CoopIS, DOA, ODBASE, GADA, and IS*. Springer, 2007, pp. 1634–1645.
- [43] A. Brown, M. Martin, B. Kapralos, M. Green, and M. Garcia-Ruiz, "Poster: Towards music-assisted intrusion detection," 2009, poster presented at IEEE Workshop on Statistical Signal Processing.
- [44] "Fraunhofer IIS Netson," 2016, URL: <http://www.iis.fraunhofer.de/en/muv/2015/netson.html> [accessed: 21/03/2016].
- [45] "Specimen Box, The Office for Creative Research," 2014, URL: <http://ocr.nyc/user-focused-tools/2014/06/01/specimen-box/> [accessed: 21/03/2016].
- [46] V. F. Mancuso et al., "Augmenting cyber defender performance and workload through sonified displays," *Procedia Manufacturing*, vol. 3, 2015, pp. 5214–5221.
- [47] T. Hermann, A. Hunt, and J. Neuhoff, *The sonification handbook*. Logos Verlag Berlin, GE, 2011.
- [48] E. Yeung, "Pattern recognition by audio representation of multivariate analytical data," *Analytical Chemistry*, vol. 52, no. 7, 1980, pp. 1120–1123.
- [49] J. Rubin and D. Chisnell, *Handbook of usability testing: how to plan, design and conduct effective tests*. John Wiley & Sons, 2008.
- [50] J. R. C. Nurse, S. Creese, M. Goldsmith, and K. Lamberts, "Guidelines for usable cybersecurity: Past and present," in *Proceedings of the Third International Workshop on Cyberspace Safety and Security (CSS)*. IEEE, 2011, pp. 21–26.
- [51] G. Parsehian, C. Gondre, M. Aramaki, S. Ystad, and R. K. Martinet, "Comparison and evaluation of sonification strategies for guidance tasks," *IEEE Transactions on Multimedia*, vol. 18, no. 4, 2016, pp. 674–686.
- [52] S. C. Peres, D. Verona, T. Nisar, and P. Ritchey, "Towards a systematic approach to real-time sonification design for surface electromyography," *Displays*, 2016.