

Kent Academic Repository

Full text document (pdf)

Citation for published version

Axon, Louise and Alahmadi, Bushra and Nurse, Jason R.C. and Goldsmith, Michael and Creese, Sadie (2020) Data Presentation in Security Operations Centres: Exploring the Potential for Sonification to Enhance Existing Practice. *Journal of Cybersecurity*, 6 (1). ISSN 2057-2085.

DOI

<https://doi.org/10.1093/cybsec%2Ftyaa004>

Link to record in KAR

<https://kar.kent.ac.uk/79587/>

Document Version

Publisher pdf

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>

Research paper

Data presentation in security operations centres: exploring the potential for sonification to enhance existing practice

Louise Axon ^{1,*} Bushra A. AlAhmadi,¹ Jason R. C. Nurse,²
Michael Goldsmith¹ and Sadie Creese¹

¹Department of Computer Science, University of Oxford, Oxford OX1 3QD, UK; and ²School of Computing, University of Kent, Canterbury, UK

*Correspondence address: Department of Computer Science, University of Oxford, Oxford OX1 3QD, UK. Tel: 01865 610805; E-mail: louise.axon@cs.ox.ac.uk

Received 15 March 2019; revised 9 October 2019; accepted 24 January 2020

Abstract

Security practitioners working in Security Operations Centres (SOCs) are responsible for detecting and mitigating malicious computer network activity. This work requires both automated tools that detect and prevent attacks, and data presentation tools that can present pertinent network security monitoring information to practitioners in an efficient and comprehensible manner. In recent years, advances have been made in the development of visual approaches to data presentation, with some uptake of advanced security visualization tools in SOCs. Sonification in which data are represented as sound, is said to have potential as an approach that could work alongside existing visual data presentation approaches to address some of the unique challenges faced by SOCs. For example, sonification has been shown to enable peripheral monitoring of processes, which could aid practitioners multitasking in busy SOCs. The perspectives of security practitioners on incorporating sonification into their actual working environments have not yet been examined, however. The aim of this article, therefore, is to address this gap by exploring attitudes to using sonification in SOCs and by identifying the data presentation approaches currently used. We report on the results of a study consisting of an online survey ($N=20$) and interviews ($N=21$) with security practitioners working in a range of different SOCs. Our contributions are (i) a refined appreciation of the contexts in which sonification could aid in SOC working practice, (ii) an understanding of the areas in which sonification may not be beneficial or may even be problematic, (iii) an analysis of the critical requirements for the design of sonification systems and their integration into the SOC setting and (iv) evidence of the visual data presentation techniques currently used and identification of how sonification might work alongside and address challenges to using them. Our findings clarify insights into the potential benefits and challenges of introducing sonification to support work in this vital security monitoring environment. Participants saw potential value in using sonification systems to aid in anomaly detection tasks in SOCs (such as retrospective hunting), as well as in situations in which peripheral monitoring is desirable: while multitasking with multiple work tasks, or while outside of the SOC.

Key words: cybersecurity; sonification; security operations centres; usable security

Introduction

The threats to the cybersecurity of today's organizations are numerous, vastly varied and constantly evolving. Security Operations Centres (SOCs) run within and on behalf of organizations and are responsible for the security of networks and critical infrastructure. In SOCs, security practitioners work, often under high pressure [1], interacting with a range of security tools to detect and prevent malicious activity. There is a requirement for monitoring tools for use in SOCs that are effective and meet the needs of security practitioners. A key role of these tools is presenting pertinent security information to practitioners in a way that is comprehensible. In recent years, advances have been made in the development of visual methods of presenting security data. The incorporation of sonification in which data are represented as sound, into SOCs has also been considered.

Sonification is defined as 'the use of non-speech audio to convey information' [2]. The outputs of sonification systems are often referred to as 'sonified displays' or 'auditory displays'. A body of research exists in the use of sonification for monitoring processes, exploring data and alerting [3]. Based on existing research, the properties afforded by sonification align with some known requirements of SOCs. Articles exploring the sonification of network security data indicate its promise as a technique for attack detection [4–7], improved methods for which are critical to SOCs. Furthermore, sonification is an effective medium for peripheral monitoring of information as a non-primary task [8]. This could be useful to busy practitioners in bustling SOCs. On the other hand, there are concerns about the fatigue and distraction that could be caused by sonification, which raise questions about its true utility in these dynamic environments.

Despite these potential benefits, there has to date been no research exploring practitioners' perspectives on the contexts in which sonification could integrate into SOC workflow. It is therefore unclear how these practitioners regard the incorporation of sonification in SOCs. Understanding the needs of users, however, is crucial to incorporating new technologies into their working environment [9]. To address this gap, we consulted with practitioners working in SOCs, to further understand the current practice in presenting information to practitioners in SOCs, and to explore their perspectives on incorporating sonification into this unique setting.

We envisage that sonification might work alongside and enhance existing data presentation approaches. To support exploration of this, we initially aimed to gather evidence on the existing use of visual data presentation approaches in SOCs. By visual data presentation approaches, we refer to any methods by which data are presented to be observed visually by security practitioners: security visualizations and text-based presentations of data are common examples. We then aimed to identify and refine contexts of use for sonification in SOCs, analyse integration and design requirements, as an initial stage in the user-centred design process [10].

This article reports on the results of a study involving an online survey and semi-structured interviews with security practitioners working in SOCs. Both the online survey and interview involved two different sets of questions on (i) the current use of visual data presentation approaches in SOCs and (ii) the potential for using sonification in SOCs. The results obtained through these questions are reported in this article. In addressing (i), we aimed to understand which visual approaches are currently used in SOCs to present information about low-level network data and about security-tool output to security practitioners.

To address (ii), we began by designing tentative use cases for sonification in SOCs, using information gathered from existing

literature and the responses of security practitioners in the online survey. We then discussed sonification in the interviews, beginning by presenting participants with a network-packet sonification prototype we developed, in order to familiarize them with the concept of sonification. The proposed tentative use cases were then explored, and participants' views on integration and design discussed. We thus refined contexts of use, discarding use cases that were not considered to have promise and analysed user needs with regard to integration and design [11]. In our analysis, we consider the implications of the evidence gathered about visual data presentation approaches currently used for the use of sonification and for the possibility of interaction between new sound-based approaches and existing visual approaches to data presentation.

This article extends a previously published paper [12] and makes the following contributions to the usable security, human-computer Interaction (HCI) and cybersecurity domains:

- Presents evidence of the current use of visual data presentation approaches in SOCs, which can inform researchers focusing on tool development in this area.
- Identifies and refines the contexts in which sonification systems could improve working practice in SOCs.
- Establishes an empirical understanding of the challenges of integrating sonification into the SOC setting.
- Extracts design requirements for sonification tools that would be effective and usable for SOC practitioners.
- Identifies directions for research into the potential for sonification to solve challenges in using, and to work alongside, existing visual data presentation approaches.

Our findings can inform sonification interface development and future studies into the use of sonification in SOCs. The rest of this article is structured as follows. In section 'Background and related work', we present relevant background and related work on SOCs and sonification. We describe the methodology followed in this study in section 'Methodology'. In section 'Online survey and interview participants', we present the demographics of participants in the online survey and interviews. The data gathered during the online survey and interviews on the use of visual data presentation approaches in SOCs are presented in section 'Existing visual data presentation approaches'. In section 'Development of tentative use cases for sonification in SOCs', we present the results relating to sonification from the online survey and our analysis of them to produce initial use cases for sonification in SOCs; the interviews in which we explored these use cases are then reported in sections 'Interview results: perspectives on the utility of sonification use cases' and 'Interview results: perspectives on the integration and design of sonification'. In section 'discussion and implications', we reflect on the results presented on sonification and discuss their implications for the use of sonification in SOCs. We also consider the results on sonification in the context of our findings on visual data presentation to identify how sonification might work alongside existing practice. In section 'Conclusion and future work', we conclude this article and describe directions for future work.

Background and related work

We begin with an overview of the work of security practitioners in SOCs. We then review HCI studies on how SOCs work, and applications of sonification to network security tasks.

SOCs and existing data presentation approaches

The objective of an SOC is primarily to mitigate cybersecurity threats towards the organizations for which they are responsible [13]. Internal SOC's are responsible for the organizations they are placed within, while multitenanted SOC's monitor network security on behalf of multiple client organizations. Figure 1 [14] is an example of an SOC, with security data presented to security practitioners on computer monitors. Practitioners are frequently required to work long shifts, including night shifts, looking at multiple screens for extended time periods [15]. The resulting pressure and demanding nature of SOC work have been highlighted in HCI research [1, 15].

Security practitioners interact with automated security tools, such as signature- or anomaly-based intrusion detection systems (IDSs), which produce security events. These data are often collated in integrated security incident and event management (SIEM) solutions [15]. The role of security analysts can include preliminary detection, triage of events and responding to customer tickets. Security engineers are also responsible for maintaining infrastructure and creating detection rules, for example [15]. By 'security practitioner', we denote a person who works in an SOC (an analyst, engineer, or manager).

It is important that humans are presented with network security information, such that they are best able to detect anomalies that do not fit automated detection profiles and to triage machine-based detection inaccuracies [16]. The provision of effective techniques for presenting security data to humans is important for SOC's, and an area of continuing academic research [17]. Security visualizations and text-based interfaces present automated system output, as well as unparsed network packets, which can enable security practitioners to recognize anomalous activity [18]. While we have seen continued advances in the visual presentation of security data in recent years [17], there is a need to further refine presentation approaches that are usable and enable users to comprehend complex data [19].

HCI studies in SOC's

A number of HCI articles have focused on examining the work of security practitioners in SOC's, and the challenges faced. This has included interview-based research [20–22] and ethnographic fieldwork [15, 23, 24]. Below, we reflect on some of the most pertinent to our research.

Sundaramurthy *et al.* conducted anthropological fieldwork in SOC's spanning 4 years. Students trained in anthropological methods were embedded in three different SOC's as security analysts [1, 13, 14, 25]. Activity Theory was used to model SOC operations, and the successes and failures encountered in integrating new technologies into SOC's studied. The implications of the findings for improving SOC operations were described, including the need for useful new tools to be dynamic and constantly resolve emerging conflicts [25]. Factors contributing to security analyst burnout, rates of which are consistently high, were modelled as a cycle linking factors concerning skills, empowerment, creativity and growth [1].

Werlinger *et al.* used interviews and participatory observation to identify the interactions of security practitioners [21, 24, 26]. They found that the existing tools used were not sufficient to support complex security tasks, with the high number of false positives produced by IDSs highlighted [21]. In extended research, Werlinger *et al.* used semi-structured interviews to understand security incident response practices [22]. Findings included a tendency for complication of incident diagnosis by usability issues with security tools and by a need for practitioners to rely on their own knowledge.

D'Amico *et al.* investigated the workflow, decision processes and tool use of security practitioners in SOC's using cognitive task analysis [23]. Cognitive challenges including the massive amounts of network data were identified. D'Amico *et al.* also explored the perspectives of security practitioners on the use of security visualizations in their work [16]. Findings indicated that visualizations could support data analysis.

While HCI studies have identified approaches to improving SOC operations, approaches using sonification have not been explored.



Figure 1: A SOC (US Department of Defense photo).

The use of sonification has been examined only insofar as its utility in network security tasks has been assessed, in studies not specific to SOCs (reported in section ‘Sonification for network security monitoring’). Incorporating sonification into SOCs has not, to our knowledge, been explored from an HCI perspective.

Sonification for network security monitoring

Prior work has applied sonification in security monitoring tasks. Axon *et al.* surveyed existing articles [4], highlighting sonification systems designed for network attack detection [5, 27–32]. The utility of sonification for SOCs is proposed, based on the challenges SOCs face, and evidence of the potential benefits of sonification [4]. Sonification can enable humans either to identify a general change in status, without knowing exactly what changed, or to actually understand the meaning of the information represented.

Researchers have reported the ability to hear attacks using a range of mappings from network traffic features to parameters of sound [5, 32]. Qi *et al.* mapped network traffic parameters to sound and stated that a range of attack scenarios was distinguishable [32]. Ballora *et al.* sonified network traffic with a view to aiding anomaly detection and reported the ability to hear patterns associated with port-scanning and distributed denial-of-service (DDoS) attacks [5]. Gilfix *et al.* detected unusual network conditions such as excessive traffic using a mapping from network traffic to natural sounds [28].

User studies have been carried out with sonification systems for network security monitoring tasks. Gopinath sonified a range of security events in Snort IDS [33]. Results indicated that sonification may increase user awareness in intrusion detection [33]. Studies by Debashi and Vickers showed that humans could detect certain network attack conditions (including denial-of-service attacks and botnet activity) more accurately using a sonification system than an IDS [6] and than three leading anomaly detection systems that use machine learning [7].

Kaczmarek *et al.* found that non-expert participants’ failure rates in carrying out security-critical tasks were lower when auditory cues were played [34]. Less complex stimuli improved performance, while more complex stimuli worsened it [35]. These results are consistent with the Brain Arousal model: moderate noise can improve cognitive performance, while excessive or insufficient noise is detrimental [36]. The findings support the potential for the improvement of network security monitoring task performance through audio cues designed with appropriate levels of complexity.

While the potential utility of sonification for conveying network security information is evidenced in prior work, and the integration of sonification into SOCs has been proposed, users’ perspectives on this technology have not been explored. This is the research gap that our article seeks to address.

Methodology

We describe the set-up of the online survey and interviews, our recruitment of participants and ethical approval in section ‘Online survey and interview set-up, recruitment and ethical approval’. We describe how we used the survey and interviews to collect information on the current use of visual data presentation approaches in SOCs in section ‘Exploring current approaches to data presentation in SOCs’. We then describe how we used a different set of questions asked during the same online survey and interviews to explore the potential for using sonification in SOCs in section ‘Exploring the potential for sonification’.

Online survey and interview set-up, recruitment and ethical approval

Recruitment

We recruited a convenience sample of 20 participants for the online survey and 21 participants for the interview. Participants were security practitioners who worked in SOCs with which we had previously established relationships and were recruited through spoken or email contact with those responsible for the SOC. We targeted organizations that ran internal or multitenanted SOCs; in section ‘Online survey and interview participants’, we describe the demographics of participants in more detail. There was likely some overlap between survey and interview participants since the same SOCs were involved in each. The extent of this overlap is unknown since survey responses were anonymized.

Survey and interview set-up

The survey was hosted online and emailed to recruited participants. Face-to-face semi-structured interviews took place at the organizations at which participants worked, in rooms exterior to the SOC. The exception was two participants who were interviewed through a live video chat due to travel constraints. Interviews were audio-recorded and lasted approximately 30 minutes. We chose to conduct semi-structured interviews, with the aim of extending discussion based on the flow of conversation.

Reliability

To ensure face validity [37] of the online survey and interview questions, both were discussed with, and incorporated feedback from, a field expert (a researcher in HCI), and three subject matter experts (who worked, or had previously worked, in SOCs). Both the survey and interview questions were also answered by subject matter experts in a pilot study.

Ethical approval

Ethical approval for this study was granted by the Central University Research Ethics Committee, University of Oxford (reference: R48822/RE001). We ensured the ethical handling of collected data through an informed consent process for participants, storage of data in password-protected files viewed only by the researchers and anonymization of published results.

Exploring current approaches to data presentation in SOCs

As part of the online survey, we asked two questions that focused on the existing data presentation approaches used in SOCs: the use of security visualizations and text-based representations of data. The first of these questions asked participants to state whether their SOC work involved the use of either security visualizations or text-based data presentations (we provided examples of text-based data presentations—Wireshark, tcpdump and Nmap—to create a common understanding). The second question asked participants to rate the importance of security visualizations to their work using a five-point Likert-type scale from ‘very unimportant’ to ‘very important’.

These survey data were supplemented by further exploration of existing data presentation techniques during the interviews: participants were asked to describe the ways in which network security information is presented to them in their SOC work, and also to express their views on the benefits of and challenges to using these approaches. The survey and interview results on current approaches to data presentation in SOCs are presented in section ‘Existing visual data presentation approaches’, and their implications for the use of

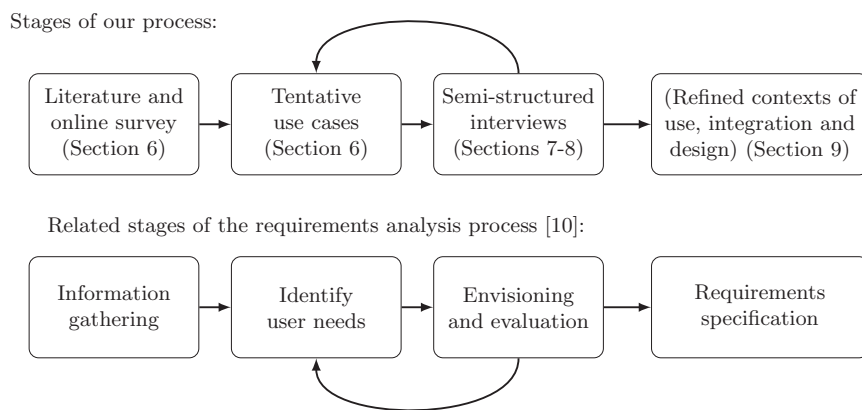


Figure 2: Study methodology: Requirements Analysis Process.

sonification in SOCs are discussed in section ‘Implications of findings on existing visual data presentation approaches’.

Exploring the potential for sonification

Research approach

Having identified existing approaches to data presentation used in SOCs, we aimed to explore the potential for sonification to be used. In particular, we aimed to identify requirements in sonification design and integration, and contexts of use for sonification in SOCs, as part of the user-centred design process [10]. By contexts of use, we refer to the conditions under which sonification could be used in SOC work [38]. The stages of the process we used to explore the potential for sonification and design and integration requirements are shown in Fig. 2, in relation to the requirements analysis process. This requirement analysis approach is widely used in prior literature and described by Maguire *et al.* [10].

As illustrated in Fig. 2, we drew on existing literature and the results of an online survey to design tentative use cases: descriptions of conditions in which sonification might be used in SOCs, the development of which is presented in section ‘Development of tentative use cases for sonification in SOCs’. We refined those use cases that participants felt had some utility in the interviews, to produce contexts of use.

By exploring the use cases in interviews, we identified the potential for integrating sonification into SOCs and challenges. Questions remain to be answered, however, before a full requirements specification (the final stage in Fig. 2) can be produced. The refined contexts of use, and integration and design requirements that we contribute in this article are initial work that can form the basis of a requirement specification. In section ‘Discussion and implications’, we highlight the areas that remain to be addressed experimentally and through further interaction with users in the construction of a full requirement specification.

Developing tentative use cases for sonification

We drew on existing literature in developing ideas for tentative use cases for sonification in SOCs [5, 8, 14, 16, 17, 32, 39–44]. From this, we identified areas requiring validation and constructed questions on these aspects. For example, one area in need of validation was the extent to which security practitioners were required to use multiple screens in SOCs, and for this we developed questions

Table 1: Sonification prototype mappings

Packet property	Musical property
IP/port commonness	Consonance of pitch
Source/destination IP/port	Octave of pitch
Packet size	Amplitude
Direction of traffic	Pan of sound

pertaining to the need to watch multiple monitors or dashboards. These questions were asked in the online survey: participants were asked to indicate their level of agreement with six assertions.

Level of agreement with assertions was indicated using a Likert-type scale [45] with response categories: from ‘Strongly disagree’ (=1) to ‘Strongly agree’ (=5). We selected the Likert-type scale as an efficient method of collecting participants’ attitudes [46]. Based on these responses, we designed five tentative use cases, which are presented in section ‘Development of tentative use cases for sonification in SOCs’.

Semi-Structured interviews to explore the use of sonification

First, participants were introduced to the sonification prototype, a system that maps properties of network packets to music. The system reads a packet capture in (mock) real-time and generates sound events based on sampled packets.¹ Table 1, which we provide to enable replication of our research approach, describes the mappings used from properties of packets to musical properties. The prototype design is not the focus of this article, so we do not detail the implementation. Further details on our technical approach can be found in Axon *et al.* [4].

The prototype was pre-recorded running on a synthetically generated dataset containing port scan, DDoS and data exfiltration attacks. Our aim here was to familiarize participants with the concept of sonification; this was particularly important given that the technique is relatively little known, and not operational in SOCs. Early prototyping is key to user-centred design, to convey to users an understanding of the system, elicit ideas for discussion and enable users to play a role in the iterative design process [11]. This is also crucial for creating security interfaces that are effective, yet usable [47].

The researchers described the system and mappings from data to sound. Participants then listened to an audio recording of the prototype using headphones. Next, the interview took place, guided by the questions presented below.

¹ The sonification prototype sound clip is available at <https://soundcloud.com/user-71482294/socs-interview-network-sonification> (10 February 2020, date last accessed).

[1-5.] We are considering the use of sonification for [tentative Use Cases 1-5] in SOCs. What is your view on the potential of sonification in this use case? This can include this particular prototype, and also the concept of sonification as a whole for SOCs.

Before these questions were asked, participants were given the five tentative use cases on paper. Participants then answered each question, and discussion ensued with the researchers, expanding on topics brought up by the participant such as other use cases and challenges in integration. We encouraged both criticisms and positive responses. Throughout the interview, we highlighted that the participant could consider different sonification designs to the prototype presented. We ensured that this was clear since the aim of the interview was to discuss the potential for the concept of sonification in SOCs in general.

Participants were then asked to rate the potential utility of each of the five tentative use cases presented, using a Likert-type scale: 'Please rate the potential utility of sonification in this use case, from 1: not at all useful, to 5: very useful'. This rating stage was placed at the end of the discussion of each use case to allow participants to formulate their views.

Data analysis

Given the discrepancy in the community as to how to treat Likert scale data [48-50], we calculated the mode and median to analyse both the responses to the assertions in the online survey and the ratings given to each use case in the interviews. We considered that a mode or median rating higher than 3 constituted overall agreement with an assertion since 3 was the middle value. We also calculated comparison of non-neutral scores (CNNS) in which we took the ratio of scores less than (1, 2) and greater than (4, 5) the neutral value (3). The three measures support the same conclusions, considered alongside the analysis of the interview data.

We analysed the interviews using template analysis [51]. This technique is useful for qualitative data analysis in which the researcher has some understanding of the concepts to be identified. We first developed a priori themes to be identified in the data: use-case utility; integration questions; and design requirements. We manually transcribed our interview recordings, producing transcripts for each discussion, and spent time becoming familiarized with the data. We then coded the interview transcript data set initially, attaching relevant parts of the transcriptions to the a priori themes. Relevant sections of data that did not fit into these themes were assigned new codes.

We thus produced an initial template of codes, which we then developed through iterative application to the data set, modifying the template as appropriate to the data. Through this refinement, we produced a final template and data set coded according to it. We then interpreted the data and wrote up the findings within the themes of the template. During the interpretation and write-up process, we engaged in frequent reflections to avoid bias and the influence of personal beliefs.

Online survey and interview participants

The online survey was completed by 20 participants working in SOCs between January and April 2017: 2 SOC managers; 14 security analysts, 5 of whom were 'senior' security analysts; and 4 security engineers (2 senior).

Table 2: Interview participant (P) demographics

Position	Internal SOC	Multitenanted SOC
Manager	3 (P1/P2/P17)	1 (P6)
Senior analyst	0	3 (P7/P15/P16)
Analyst	7 (P3/P4/P13/P18-P21)	3 (P10-P12)
Engineer	2 (P5/P14)	0
Analyst and engineer	0	2 (P8/P9)

We interviewed 21 participants between May and June 2017. Participants were security practitioners working in seven different SOCs. From 3 different internal SOCs, responsible for the security of a single organization, 12 participants were interviewed. We also interviewed nine participants from four different multitenanted SOCs, who provided managed services for client organizations. Of the participants, 4 were SOC managers; 10 were security analysts (3 senior); 2 were both security analyst and engineer; 2 were security engineers. Table 2 shows the job role and organization type of each participant.

Existing visual data presentation approaches

In this section, we present the evidence gathered during the online survey and interviews on the visual approaches taken currently to present both low-level network data and security-tool output in SOCs. These details provide an insight into the approaches currently taken to data presentation, which can inform future research on this topic. The evidence gathered also provides a basis for understanding the potential for sonification in SOCs, and allows us to consider the possible interactions between sonification and other data presentation approaches. We reflect on our findings in section 'Implications of findings on existing visual data presentation approaches', exploring the possible interactions between sonification and existing data presentation approaches.

Survey results

The results on use of data presentation approaches gathered in the online survey are shown in Figs 3 and 4. Figure 3 shows that a larger proportion of participants used text-based data presentations in their SOC work (79%) than used security visualizations (37%). Figure 4 presents participants' views on the importance of security visualizations and shows that on a five-point scale from 'very unimportant' to 'very important', security practitioners most frequently rated security visualizations as 'important' to their work.

Interview results

We now describe the evidence gathered during the interviews on the data presentation approaches currently used in SOCs. This section is structured in two parts. In the first, we focus on approaches to presenting low-level network data. This refers to data that have been subject to no, or minimal, parsing by security tools; packet captures and logs of activity on machines meet this description, for example. In the second part, we describe techniques to present the data outputted by security tools, such as IDS alerts.

Presentation of low-level network data

Text-based approaches to presenting low-level network data were described by participants. Packet captures are reviewed by security practitioners, presented as text using tools such as Wireshark.²

² A network packet capture viewer and protocol analyser are available at <https://www.wireshark.org/> (10 February 2020, date last accessed).




Figure 3: Use of security visualizations and text-based data presentation.




Figure 4: Participants' views on the importance of security visualizations to their work.

It was noted that reviewing packet captures manually can be difficult due to the large quantities of data they contain:

We will have all of the packet captures that we get from a customer's network coming into [the PCAP viewing tool] ... you can manually go and review it, but with the quantities of data and stuff in there, it's going to be very difficult unless you have got something to point you in the right direction. (P7)

Participants discussed the use of visualizations such as time series in SOCs to represent low-level network data. The utility of such approaches for enabling practitioners to detect unusual patterns, or 'spikes' was described, particularly 'when you're talking about long term trends' (P13).

The types of changes described as being quickly noticeable using plots included 'a sudden spike on a day there isn't normally a spike, a sudden increase on an interface throughput, a sudden increase on an egress point' (P6). When there is a requirement to focus on specific parts of the low-level traffic, such as traffic between particular IP addresses, filtered visualizations can aid in detecting changes: 'visualizations are useful for spotting spikes. ... graphs and counts of logs where you've filtered it down to something specific' (P15).

Visualizations can be beneficial in enabling security practitioners to convey information to people with less technical security expertise—this might include their customers or management, for example. A participant explained the utility of plots displaying the traffic crossing a network boundary in conveying data loss to a client:

One of the things our senior customer is really interested in is how much data is crossing his boundary outbound every day ... so we show a graph that says, this much [data] went on this day, and this much went on this day and by the end of the month you had lost four terabytes of data. (P1)

Techniques for visually displaying the locations of traffic flow and activities on the network were described, enabling an understanding of 'what the network really looks like' (P1) in an 'easy format to take in' (P7). One participant described the inclusion in the SIEM tool they used of: 'a graphical representation of where the bulk of the traffic flows flowed through the various zones of the firewalls and the network' (P2). As well as security monitoring, such visualizations of the network traffic flow can be useful for monitoring network operations, for 'monitoring throughput of networks, so how

much traffic is going through a device, is the threshold reached' (P7).

These approaches to visualization that display the layout of the network can help to represent a variety of information relevant to network security monitoring. An example is the spread of malware through a network's hosts:

If you see something on a host it will show you a nice diagram of the other host it has been seen on. ... if a bit of malware spread, are there the same changes on all of the other hosts that it has been seen on? (P16)

Incorporating the physical structure of the organization's building into the representation can be beneficial for monitoring for hardware activity such as Universal Serial Bus insertions: 'I saw a brilliant tool ... you can zoom into rooms, and there's a picture of the room and you can go "right, that room, that port, has just plugged something in"' (P1).

After the detection of a change, either through alerting by a security tool such as an IDS or through a practitioner spotting an anomalous pattern, low-level data representations can be used to investigate it. One participant described the process that might ensue following the detection of anomalous file additions to a server:

They [the analyst] will then follow that to its logical conclusion, so who interacted with the server in that spike period, what segment of the network were they on, what privileges did they have, why might they have suddenly started putting files onto that server, what files have suddenly turned up on that server, is there anything unusual. (P6)

It was noted that at this point the practitioner might discover that the anomaly was caused by a benign activity, such as a 'quarterly report that all departments in the company will put into this share' (P6). This highlights that a benefit of using low-level data representations is that they can provide humans with enough information to understand the security data based on their knowledge of the operational context, which can aid in decision-making. The sheer volume of low-level network data was cited as a challenge to creating useful visualizations since they are often:

Too noisy: a lot of network based stuff again is volumetric ... it is very hard to actually pick up on a big map what is normal and what is not. So it is more trying to automate and give an alert from what is flagging up rather than just having it on a map. (P16)

Presentation of security-tool output

There was variation in the way alerts are presented across SOCs, with some using text only, and others supplementing text with visual or colour-coded representations. A participant who used text-based representations of alerts described the types of information that would usually be presented about the alerts:

The alarm will have the security event data inside it, so, and it takes particular factors out of it such as the IP, the time, the date, the value and so on, and it puts that in different fields, it's better for the analyst to read. (P11)

The use of packet capture viewing tools to inspect the packets relating to an alert more closely was described.

Some participants were familiar with the use of colour coding of alerts presented as text, with focus on colour associations, such as red being a 'bad' (i.e. severe) alert. This includes traffic-light colouring: 'red, amber, green alerts' (P13). Higher resolution colouring is also used in which a higher number of alert severity levels are represented: 'we've got a kind of risk-based priority which will give you a number and a colour associated with it ranging from zero with a clear colour, through to red and being 99' (P15). A practitioner described the use of coloured lights to signal alerts and noted its attention-grabbing value: 'I love the light ... any other indicators for me are really useful, because I don't focus on any one thing' (P16).

An advantage of using colour coding is that it draws attention to events that are severe and, therefore, reduces the possibility of attackers using distraction tactics (such as using one activity to distract practitioners while the real attack is carried out separately): '... flashing red light screams "quickly, move, do this". So I'm not sure how somebody would actually distract us' (P1). The difficulty of tuning the SOC's detection systems to detect anomalies accurately enough for colour-based solutions to be of real value was explained—although we note that this is a difficulty that applies to SOC practice in general and not only to colour-based data presentation solutions.

Examples of the types of alert visualization used in SOCs were also given:

Dashboards with pie charts and things on and what's your most common alert at this particular time ... when you see a deviation from what you're normally seeing, that's when you know you need to investigate a bit more. (P13)

Visual approaches are used to give an overview of the SOC's activity, including the presentation of ongoing and resolved alerts, for example. This type of information can be displayed on large shared screens in the SOC, and the view was expressed that this type of information is useful for conveying activity to management, but perhaps more challenging for analysts to use: 'great for management so they can see how many alerts are there, but try reading that from a few rows back' (P16).

Development of tentative use cases for sonification in SOCs

We summarize our development of ideas using existing literature on SOC working practice and sonification, indicating potential uses for sonification in SOCs. We present the outstanding questions (OQs) that we identified and addressed to support the evolution of these ideas and their formulation into assertions to be asked in the online survey. Finally, we present the five tentative use cases derived.

Developing ideas using existing literature

Anomaly detection approaches for security monitoring are widely researched, including visualization-based techniques to enable detection of abnormal activity by humans [17, 42]. A wide array of experimental results evidence the utility of sonification for detecting anomalous patterns in data in fields, including Medicine and Astrophysics, for example [39, 41, 43, 44]. Furthermore, prior work has supported the use of sonification for hearing network attacks [5, 32]. We therefore posit that it is important to explore the potential for sonification to enable humans working in SOCs to detect anomalies in the network traffic, and seek to address the following question:

OQ1. Do security practitioners feel capable of detecting anomalies directly from the network traffic?

Security practitioners may be required to carry out other tasks while monitoring the network, for example, managing email inboxes [15]. Prior literature indicates the utility of sonification as a solution to enabling monitoring as a non-primary task. Hildebrandt *et al.* showed that using sonification to monitor a process as a secondary task while performing a different primary task had no significant effect on performance in either task [8]. The use of sonification for peripheral monitoring may extend to cases in which security practitioners wish to continue to monitor whilst outside of the SOC. We consider that this may be true particularly for practitioners alone on shift while taking breaks, for example. To support the evolution of this idea, we seek to address the following question:

OQ2. To what extent are security practitioners required to multi-task while monitoring in SOCs?

The information required for monitoring in SOCs is often distributed across multiple monitors used by security practitioners [16], including large screens at the front of the SOC. Security practitioners may therefore be required to focus their visual attention in multiple directions, yet it has been shown that visual perceptual clutter leads to increased errors in judgement [40]. Furthermore, security practitioners, depending on their role, can be required to monitor screens for extended time periods, focusing on visual representations of the data and monitoring alerts from SIEM solutions, for example [15], which may lead to visual fatigue. Presenting sonified data could reduce the emphasis on visual monitoring. This could mean either reducing the number of directions in which visual focus is required or providing an alternative monitoring method for visually fatigued practitioners. We seek to address the following question in developing this idea:

OQ3. To what extent are security practitioners required to visually monitor information presented on multiple screens?

Exploring ideas using the online survey

The six assertions developed to assess the OQs, and participants' responses to them in the online survey, are presented in Table 3. Five of the assertions obtained mode and median ratings greater than 3, which we consider agreement, as explained in section 'Methodology'. The exception is 'Assertion 2', which indicates that while practitioners feel capable of detecting anomalies, they are less confident that their existing monitoring set-ups enable this, and this is supported by the CNNS. This result supports experimentation with new methods of enabling this capability. The survey results can therefore be seen to affirm the three OQs.

Table 3: Online survey results: number of responses of each value to each assertion [Resp, ordered from ‘Strongly disagree’ (=1) to ‘Strongly agree’ (=5)]: mode, median (Med), and CNNS: Disagree (1–2):Agree (4–5) (CNNS: D:A).

Assertion	Resp					Mode	Med	CNNS
	1	2	3	4	5			
Anomaly detection by humans (pertains to OQ1)								
Assertion 1: human analysts monitoring the network are capable of detecting network anomalies missed by automated systems	0	1	5	10	4	4	4	1:14
Assertion 2: the monitoring set-up I use enables me to detect network anomalies that are missed by automated systems	0	4	11	4	1	3	3	4:5
Assertion 3: I sometimes rely on my experience and intuition to detect network anomalies rather than monitoring system alerts	0	2	7	7	4	3.5	4	2:11
Multitasking/non-primary task monitoring (pertains to OQ2)								
Assertion 4: I am required to monitor the network, while carrying out other tasks simultaneously (e.g. responding to emails)	0	2	2	13	3	4	4	2:16
Monitoring across multiple screens (pertains to OQ3)								
Assertion 5: In monitoring, I am required to watch multiple monitors depicting different data at one time	0	1	2	12	5	4	4	1:17
Assertion 6: I am required to watch multiple dashboards on the same monitor depicting data at one time	0	3	4	9	4	4	4	3:13

Tentative use cases

Based on the survey results presented in Table 3, and the prior literature, we derived the following five tentative use cases to carry forward to the interviews.

1. Detecting anomalies in the network traffic: presenting high-resolution sonifications of the network traffic, to enable humans to hear network anomalies.
2. Monitoring as a non-primary task: sonifying network security data to be monitored as a secondary task, enabling the user to carry out a separate primary task simultaneously.
3. Monitoring data presented across multiple screens: sonifying parts of information that are currently presented across multiple screens, reducing the directions for focus of visual attention by users.
4. Alleviating fatigue from monitoring screens: enabling users to monitor with reduced strain on visual attention, by providing the option to use sonification.
5. Enabling monitoring whilst outside of the SOC: enabling users to continue monitoring work (e.g. using wireless earpieces) whilst outside of the SOC.

Use Case 1 was supported by the assertions of survey participants that detecting anomalies directly from the traffic was a capability of practitioners. The requirement to monitor across multiple screens motivated the development of Use Case 3. Use Case 4 was supported by the requirement for extended periods of visual monitoring reported in prior literature [15]. The requirement affirmed by the survey to multitask while monitoring the network justified the development of Use Cases 2 and 5. We considered that multitasking might occur while carrying out other work inside the SOC, or while carrying out activities when away from the SOC, but still on duty.

Interview results: perspectives on the utility of sonification use cases

We present the interview results on each tentative use case, as well as new use cases proposed by participants. Participants’ views on the potential utility and challenges of each use case are analysed, with the aim of refining promising contexts for the use of

sonification in SOCs. In section ‘Discussion and implications’, we reflect critically on the requirements for these contexts of use.

Use Case 1: detecting anomalies in the network traffic

Overall, participants felt that sonification had potential in this use case. A number of participants felt that humans were capable of detecting anomalies when presented with network data. The belief that it was mostly humans who detect network anomalies was expressed (P15), and it was suggested that humans have the capacity to recognize more subtle anomalies than machines: ‘there’s still a lot of human analysis, and a machine can only determine the really obvious ones’ (P8). Security visualizations were frequently specified as a class of tools that enabled participants to detect anomalies, showing anomalous spikes in traffic volume, for example.

Possible benefits of sonification over existing anomaly detection approaches were explored. The potential of sonification for detecting anomalies not apparent from visualizations was described because ‘the thing with a graph is – it’s not how much you can see, it’s how much you can present’ (P6). The trustworthiness of the information conveyed by the sonification was highlighted as an advantage over automated approaches, which can produce false-positives: ‘it can’t ever lie because it’s just going on what it’s seeing, it’s not saying it’s malicious it’s just saying that that’s what I am seeing’ (P10).

The ways in which anomalies might be detected using sonification were discussed; in particular, the potential to learn some baseline sound of the network, and from this basis detect anomalies. This included hearing deviations to greater traffic throughput. The potential to ‘get used to the sound’ such that deviations were apparent was also highlighted: ‘when say a DoS attack or some other form of attack would take place, I’m sure it would stand out because you would get used to hearing a certain type of tune or hum from day-to-day activity’ (P15).

In general, participants felt that sonification had promise in this use case. Assessment of the key points highlighted is required—the ability to hear deviations from a ‘baseline’ sound, and the comparison of a sonification-based approach to anomaly detection with automated and visualization-based approaches. This comparison is important given that, while many participants believed humans

could detect anomalies in the data, some felt that anomaly detection currently was predominantly machine-driven. Another participant noted that the real solution maybe somewhere in between, i.e. that anomaly detection capabilities differ between individuals (P16).

Use Case 2: Monitoring as a non-primary task

The general consensus was that sonification could prove valuable in this use case. Participants stated that they were required to multitask in their role, and reported a range of tasks during which they were required to multitask whilst monitoring. These included researching new threats, composing reports, sending emails or investigating cyber incidents:

One issue we have is that when we see something of interest, and we are researching ... you're no longer monitoring. So, at points in time where you're not monitoring, if there was an audible cue that "oh actually, there is something happening right now, maybe my attention should be back there". (P13)

The current requirement to use a visual means in multiple tasks was highlighted as a challenge: 'If we're investigating something else ... I've only got three screens, and I've only got one pair of eyes' (P10). Participants described the potential value of sonification for monitoring without focused visual attention: 'you could just be monitoring or listening to that background rather than having to keep looking up' (P8). This extended to the use of sonification for monitoring alerts generated by automated systems, removing the need to keep 'viewing the alarm view while I'm doing other things' (P7).

The discussion of both sonified network traffic, and of auditory alerts, brings into question the types of information most appropriate for sonification in this multitasking application. The information content of sonified network packets, compared with auditory alerts, was highlighted as advantageous by one participant: 'the music can tell me, something else has happened ... and not just as an - alert, alert' (P8).

In summary, perspectives on this use case were positive, subject to some design and capability questions. A key question was the type of information to be sonified—both network packets, and alert data, were discussed as advantageous. Participants voiced concerns about the possible effect of monitoring using sonification on their primary task, and vice versa. While Hildebrandt *et al.* showed that these effects were not significant in a different context [8], assessment with SOC-specific tasks, which are often time-pressured, and require high levels of attention, is required. The nature of SOC tasks could affect the performance of users multitasking using sonification.

Use Case 3: monitoring data presented across multiple screens

The potential for sonification in this use case divided opinion. First, the extent to which practitioners felt that they were required to monitor across multiple screens differed between SOCs. Some (8/21) stated that multiple screens (between 2 and 7) were used to show live alerts and incidents (often displayed in an SIEM tool), email, chat feeds, ticketing systems, or were used to do research, for example. In other SOCs, all monitoring information was presented within a single pane of glass (6/21).

Some challenges in the use of multiple screens were reported. Information could be missed because of its distribution across multiple screens. Missed information on monitors at the front of the SOC was reported, if practitioners were engaged by other screens: 'something on this [front] screen could be red, but if they're already

doing a priority 1, they're not going to look over there seeing the other priority 1' (P6).

For these participants, monitoring across multiple screens was a challenge sonification could help alleviate. Both sonification of alerts and of network traffic were mentioned:

There are analysts sitting down there, and you have a massive dashboard, so they are still required to be looking at that at all times, and looking at their own screen. Sound will help in minimising that, just looking, as it avoids constant attention. (P17)

For some participants, however, the use of multiple screens did not pose a challenge, and it was considered convenient to have dedicated screens for executing commands, for example. These participants stated sonification would not be useful in this application and did not wish to reduce the number of screens: 'I will still use 7 [screens], even if I have all the sound in the world' (P12). One participant reported that reducing the number of screens would cause inconvenience: 'If I don't have enough screens, I've got to constantly minimise, maximise, and copy this and go here and it can be very difficult' (P7).

On the whole, participants were divided as to whether sonification had the potential to be useful in this use case. The type of information that might usefully be represented by the sonification was unclear, and a number of participants did not desire any fewer screens. While it is clear that the spread of screen locations can cause information to be missed, it is likely that other technologies would be more effective solutions than sonification, meeting the needs of a greater proportion of security practitioners. Some participants suggested that the combination of this information into a single pane of glass would be a solution preferable to sonification in this instance.

Use Case 4: alleviating fatigue from monitoring screens

In general, sonification was not perceived to have potential as a solution here. Some participants (6/21) stated that they were sometimes visually fatigued by their monitoring work in the SOC, yet others stated that they were not visually fatigued as they were accustomed to looking at screens. It was suggested that the extent to which fatigue was felt differed between individuals and types of role: 'nowadays I am doing stuff all the time, but there was a period when I was just staring at, I think it was, 3 different monitors at once' (P9).

Methods used for mitigating fatigue currently included encouraging workers to take regular breaks. Another approach adopted was automating as much as possible. Participants questioned the practicality of using sonification as an alternative for visually fatigued practitioners. If the sonification played only when practitioners were fatigued, their ability to interpret information from it might be limited:

I can see it as an alternative to visualization for when you get to a point when your eyes are tired ... the thing is if you only switch it on when you get to that point, then I think you won't really understand what normal would be, so you would still need it on in the background to some extent. (P15)

A number of participants felt that sonification would not be useful for them in this application. Visual fatigue was already prevented through other approaches (automation and regular breaks), such that participants were not (or were unaware that they were) fatigued by visual monitoring work, stating that they would continue to look at screens even with sonification. The utility of sonification in this use case was questionable, and the ways in which it might work in practice unclear.

Use Case 5: monitoring whilst outside of the SOC

Participants generally felt sonification had strong potential in this use case: ‘if you were just going out and you pop a pair of headphones on or whatever and you can hear, something is going on, I can jump back in’ (P10). Specific times that could necessitate monitoring whilst outside of the SOC included during fire alarms, and while making drinks, on break, or going to the shop. This was particularly true for participants who were required to work one-man shifts: ‘today it’s only me here, and I did have to leave to the shop earlier’ (P11).

It was noted that using sonification in this application could be particularly useful for practitioners alone on shift: ‘the first ever job I did in a SOC I was the only person in the room. You could definitely say that would help with that one’ (P9). Smaller SOCs in which one-man shifts occurred, as well as companies running their own SOCs, were mentioned as situations in which this capability might be especially helpful: ‘the guy running his own SOC, the SOC won’t be his only task, he might be plumbing computers in the main office, and want to come back in if something big happens’ (P6).

It was reported that there were existing approaches to monitoring whilst outside of the SOC. This included emails sent to cell-phones, and a sonic alarm used when on break: ‘when I leave, I unmute it, so that I can go and put my feet up, and then if there’s an alarm I would come’ (P11). The potential value of a more informative sonic approach (than the simple alarm currently used) was discussed by this participant:

If we had a melody like yours representing that, and I knew what the melody was playing and what it was, then maybe I wouldn’t have to come and look at it [on-screen alerts], because I would be like ok it’s something normal for this time ... with the current beep, we don’t know until we actually log in. (P11)

The placement of monitoring screens in the break room was another approach currently used to indicate to practitioners that they were required in the SOC. A number of participants discussed being waved at through the window by other SOC workers, to attract their attention while on break. This was particularly true for analysts with higher skill levels, required for specific events. Participants felt that sonification could be useful as an approach to informing practitioners on break that they are required in the SOC, played through speakers in break areas (e.g. the kitchen), or through an ear-phone worn while on break: ‘they wouldn’t need to rush back, keep checking, they could just go about their business and know “right, when I hear that sound, I need to take whatever action”’ (P7).

The desire to use sonification for monitoring outside of the SOC might differ. For example, one SOC manager was of the opinion that monitoring should not be continued whilst on break, as it would defeat the purpose of the break. In general, however, this use case was considered a promising solution to actual challenges faced by security practitioners.

Use case ratings

Table 4 presents participants’ ratings for each use case.

Use Cases 1, 2, 3 and 5 obtained mode and median ratings greater than 3, which we consider indicates overall agreement with potential utility (see section ‘Methodology’). The CNNS shows that these four use cases were rated above neutral by most participants. Based on these results, we selected Use Cases 1, 2 and 5 to form the basis of our refined contexts of use, presented in section ‘Discussion and implications’. Although Use Case 3 also scored ratings greater than 3, we chose to omit it from the contexts of use, based on the qualitative interview analysis, from which we concluded that other

Table 4: Use case rating: number of each rating value given to each use case by participants [Resp, ordered from ‘Not at all useful’ (=1) to ‘Very useful’ (=5)]: mode, median (Med), and CNNS: Not useful (1–2):Useful (4–5) (CNNS: N:U).

Use Case	Resp					Mode	Med	CNNS
	1	2	3	4	5			
1. Anomaly detection	1	1	7	7	5	3.5	4	2:12
2. Multitasking	1	4	4	4	8	5	4	5:12
3. Multiple screens	3	2	2	8	6	4	4	5:14
4. Visual fatigue	3	6	4	3	5	2	3	9:8
5. Outside SOC activities	0	1	1	3	16	5	5	1:19

solutions to the challenge of multiple screens may be more appropriate for SOCs.

Other use cases suggested by participants

Aside from the use cases we presented, other uses were suggested by participants, falling under the following themes.

Occasional use

It was suggested that the sonification could be used to occasionally check the sound of the network: ‘I might listen to it once an hour, and go “... it doesn’t sound the same at 1 o’clock today as it did at 1 o’clock the last three days”’ (P6). Sonification could be played for the duration of particular events, which could be useful for conveying the length of events, since: ‘sometimes looking at data you might not fully understand when it started and when it ended’ (P11). Similarly, sonification could be played in the background particularly at times when high-severity incidents were being dealt with, to act as an indicator for SOC workers when a new incident may require their attention (P8).

Hunting for anomalies

One participant suggested the use of sonification as a threat-hunting tool, for analysts required to search data for anomalies retrospectively: ‘if I put that on for five minutes, and it sounds anomalous, then I know there’s five minutes’ worth of packets worth looking at. Otherwise, I might spend an hour just looking at some packets with nothing particularly interesting in’ (P6).

The potential to listen to the sonification at increased speed (fast-forwarding), both for conducting audio reviews of data retrospectively (‘if you’ve got an alarm or a period that you’re interested in’) and for real-time monitoring, was discussed:

If you had a baseline amount of traffic, you could go “I’ll listen to a minute of that, now I’m going to listen to a minute of what has just gone through the sensors”, maybe accelerated, you will then start straight away going “that doesn’t sound right”. (P6)

Improving SOC workflow

It was suggested that a continuous soundtrack could improve SOC workflow by making practitioners aware more efficiently of events that are relevant to them, without the need for others to escalate to them:

At the minute, it relies on the first person who sees those events to recognise it’s bad, to then escalate ... if you heard lots of anomalies, the people who it would be eventually escalated to would instantly know that, and could maybe start on it earlier. (P8)

A manager suggested that sonification could take over some of their alert-handling workload, by verbally presenting the queue of alerts and their severity ratings: ‘an audio prompt would give me more time: . . . if it’s not shouting numbers out, I don’t need to look at the queue’ (P6).

Interview results: perspectives on the integration and design of sonification

We present key themes identified relating to the integration of sonification into the SOC environment and to sonification design. We consolidate these results in section ‘Discussion and implications’ and highlight challenges and implications for system development.

Headphones or speakers?

A number of participants discussed whether the sonification would be best played through speakers or personal headphones. Some participants highlighted potential problems with playing the sonification through speakers—for example, that if the sonification was made the soundtrack to the SOC, practitioners who were not monitoring would still have to listen to it ‘when they’re trying to concentrate on doing something else’ (P2).

Some participants, however, felt that headphones were not always a desirable solution, as wearing headphones could isolate practitioners or hamper collaboration. Alternative solutions were suggested, including the use of a single earpiece rather than headphones, suggested by two different participants, to enable practitioners to continue to collaborate.

Existing SOC workflow and soundscape

Some participants focused on integrating sonification into necessary SOC workflow in an unobtrusive way, noting that it should not prevent ‘people being able to talk about what’s going on’ (P10). A need to standardize responses to sounds heard was suggested: ‘everything we do is based around a procedure, so I’m not sure how you would . . . get everyone to conform to, “when you hear this, you do this”’ (P7).

Participants described existing SOC soundscapes. In some SOCs, there was currently a soundtrack, such as radio for the whole room. In others, there was no deliberate noise, with practitioners listening to music at times through headphones: ‘we don’t have any audio . . . Occasionally people use headphones to listen to music, and on the odd occasion we will put music on’ (P15). If integrated into SOCs sonification must work appropriately with this range of existing soundscapes.

Complexity of networks

Participants discussed the difficulty of finding unusual behaviour in networks: ‘the more complex your network is, the more difficulty you have working out what is unusual’ (P6). A suggested approach to dealing with large amounts of network traffic was filtering sound by particular IP addresses or assets.

One participant highlighted the issue of network complexity in the multitenanted SOC they worked in: ‘I think if it was for an internal SOC for a specific company that probably would work better. Here because we’re a managed services provider, I think there would be too many things going on’ (P10). Further research into differences in required design solutions for different SOC types is needed, such as filtering sound to focus on single networks for multitenanted SOCs.

Sonification of alerts

Sonification of alerts was mentioned by a number of participants (6/21), as an approach to communicating critical events, or alerting on particular systems: ‘using this would benefit us, if only the DDoS mitigation stuff that we use, or a subset of alarms or devices alerts us to anomalies via sound’ (P7). It was suggested that sonified alerts could be layered with the sonified low-level network traffic: ‘you could tell the system to play music not just based upon the packet captures but based upon outputs of other things, signatures, outputs of x, y, z. Then you can build up two layers of that, so you could listen to the underlying traffic as part of an incident’ (P8).

Mitigating fatigue

A number of participants (8/21) stated that they felt they would be fatigued by continuous exposure to the sonification. The potential for occasional use of the sonification was discussed in the context of listening fatigue: ‘I guess you could use it as and when, but I think if you put that on somebody’s head for a day, I think you would struggle with that’ (P6).

The potential for the sonification to be unobtrusive unless required was highlighted: ‘music you can switch off to, but equally the anomalies in there, your brain is going to pick up on them and go that’s changed, that’s different’ (P8). Designing sonifications that are unobtrusive in this way is a potential approach to mitigating fatiguing effects.

What next? Approaches to investigating anomalies heard

The role of sonification as a tool for enabling anomaly detection was discussed. Sonification could be used as the initial indicator that something was wrong, resulting in some follow-up by the analyst:

It takes care of the first bit for you. So you’re going to after, go and investigate it yourself, and you’re going to have to ask a question – why has there been an increase, or why this anomaly has occurred. (P7)

Participants raised the question of how to make hearing deviations in the sonification actionable (P8, P13). It was suggested that without enabling investigation of anomalies heard, the sonification would be less useful:

Saying you heard something weird is great, but unless you can quantify that, in an actual investigation, then you know something is bad, but you don’t know what that bad thing is. (P8)

The suggestions of participants on ways of making sonification more actionable fell into two categories. The first involved making the sonification itself informative in particular ways; the second involves using existing data presentation methods to enable further exploration of information notified by the sound.

An example of an approach to making the sonification itself more informative was suggested by participants working in multitenanted SOCs. Design suggestions for monitoring for multiple customers were made, including ways of linking events from the same client, while listening across all clients. One suggestion was a voice-based approach to this linking across multiple customers in which a voice would speak the customer number simultaneously with the network sonification playing, for example (P8). Another idea used musical methods of conveying information about which customer was affected:

Different sound sets for different customers. So if I hear the DDoS sound and the malware sound for customer x and they’re

at pitch y , then I can go “ok yeah I recognise that. Hold on those are both from the same customer”. (P6)

Assessing the quantity and granularity of information extraction possible by humans listening to sonification, and the learning curve required to achieve it, will be key to understanding the potential of such approaches.

Outputting to visualizations was another approach suggested for the linking of anomalies heard to information content (P8, P10). Visual representation of the music itself, such that recognition of times at which events were heard was possible, was suggested (P10). Addressing this question will involve assessing the amount of information that can be extracted from the sonification, and the implications this has for the way in which further information should be conveyed.

Discussion and implications

We reflect on the results presented, with a view to summarizing our four main contributions, listed in section ‘Introduction’. We refine contexts of use, then consider the implications for sonification design and integration. This can guide interface designers in developing sonification systems for SOCs.

Refined contexts of use

Based on the interview results relating to use-case utility, we refine contexts of use, identifying the potential actors, key usage scenarios and relevant SOC workflow factors [10].

Monitoring whilst outside of the SOC

There are times when it is ideal for security practitioners to be able to continue their monitoring work whilst outside of the SOC. This is particularly true for smaller SOCs in which workers undertake one-man night shifts. Such workers, who might leave the SOC for a short time (e.g. to make a drink) or for a longer time (e.g. to go to the shop), saw potential value in the use of sonification to enable their monitoring work to continue. In larger SOCs, listening to sonification in break areas could improve SOC workflow for more experienced practitioners, who might currently be called (e.g. waved at physically) back into the SOC by others when their expertise is needed.

This capability could be enabled through wireless earpieces worn by workers when venturing outside the SOC or by speakers playing sonification in break areas. As well as the network packet sonification approach, of which the prototype presented was an example, sonified alert streams were highlighted as information that could be monitored at times outside the SOC. Sonification designs that enable both packet and alert representation, individually or in combination, would therefore be appropriate. Monitoring accuracy and attention during out-of-SOC activities should be compared with inside-SOC capabilities, to support the development of this use.

Detecting anomalies in network traffic

Situations in which sonification of network traffic has potential value as an anomaly detection approach include long-term, continuous listening to the sonification for real-time detection of deviations. To support this use, anomaly detection capabilities using sonification should be compared with those using security visualizations and automated tools. Prior sonification work indicates that malicious network activity can be detected using sonification [5, 32], but does not make this comparison.

Short-term anomaly detection uses include occasional checking of the sonification—for example, once per hour—to compare with previous times. Another promising short-term use is a retrospective analysis. Practitioners tasked with hunting retrospectively through data for anomalies suggested sonifications of the data could enable location of interesting packets requiring closer inspection. Research is needed into approaches to enable users to link anomalous sounds heard to the relevant data (in a text or visual form). For such tasks, listening to sonification played at increased speeds could enable users to sift through data from extended time periods more quickly.

Multitasking whilst monitoring as a non-primary task

Sonification is potentially useful for aiding security practitioners in carrying out monitoring tasks while conducting other primary tasks. It is important to assess this capability experimentally; in particular, the effect of primary tasks on secondary task sonification monitoring, and vice versa. Such work can draw on the aforementioned work of Hildebrandt, which showed that such monitoring had no significant effect on either primary or secondary task [8]. However, context-specific assessment is important, using primary tasks relevant to SOCs: sending emails, writing threat intelligence reports and investigating incidents were some tasks described.

Subsequently to the publication of the original paper that this article extends [12] and prior to writing this extended version, we carried out two studies in which we experimented with some of the use cases above. First, we found that participants could use the same prototype network traffic sonification system as had been presented during the interviews reported in this article to detect and identify four different types of network attack accurately and efficiently, including attacks occurring in combination [52]. Secondly, we experimented practically with the use of sonification of both network packets and IDS alerts for monitoring as a primary and as a non-primary task [53]. Our results showed that a number of aspects of the monitoring performance of security practitioners were improved when they used sonification alongside a Security Information and Event Management (SIEM) tool compared to when they used a SIEM alone, in an experimental setting. As we highlight in section ‘Conclusion and future work’, there is nevertheless a need for further experimentation to establish the utility of sonification in these use cases and explore the integration of the approach into the SOC environment.

The need for flexibility in interface design

Some key differences in opinion were highlighted, with implications for the sonification design. Participants differed in their opinion on whether the sonification should use headphones, speakers or single earpieces, and whether continuous or occasional use would be most appropriate. It is clear that different approaches may suit different users and scenarios. It is therefore appropriate for sonification designs to be flexible, depending on the use case and user preference. Playing the audio through all mediums discussed should all be viable (e.g. spatialization of different sounds through different ears is unsuitable for single earpiece listening), and the sonification approach should support both continuous and occasional use.

The analysis highlights a difference in requirements between multitenant and internal SOCs. A participant working in a multitenant SOC described the potential difficulty of using sonification in that environment, with large amounts of data for many customers, compared with an internal SOC. Further research into differences in the required design solutions across different SOC types is necessary. A solution for multitenant SOC environments might be

the provision of tool features to filter sound by the single SOCs to be monitored.

The prototype design presented in this study initiates the participatory design process [11]. This should be iterative, and as such future design of sonification systems for this application can draw on the design requirements we identified. Consulting users in the development process are especially important given that the technology is not operational in SOCs.

Challenges in integrating sonification into SOCs

Some challenges in integrating sonification into SOCs emerged from the interview responses. Appropriate integration of sonification with the existing SOC soundscapes reported is key if the technology is to be unobtrusive to users. In SOCs where the soundscape is silent, headphones or a sonification design that is unobtrusive could be used. Equally, the existing soundscape may affect sonification listening: the sounds produced may be drowned out in noisy SOCs.

It was highlighted that sonification should not distract users in a way detrimental to SOC activity. Sonification systems should be designed with appropriate sound complexity for particular tasks since the complexity of auditory stimuli has been shown to affect cognitive performance [36]. Reducing cognitive load is a key consideration for creating usable security interfaces [47]. Less complex sound is needed for non-primary tasks since less complex background auditory stimuli have been shown to improve the performance of security-critical tasks [35]. Mapping highly complex network data to low-complexity sounds will pose a challenge.

The copious amounts of complex data present on networks exacerbate the challenge of designing sonification systems suitable for the SOC environment since it makes finding a baseline of ‘normal’ behaviour difficult. Concerns were voiced in interviews that sonification systems representing such data could become cacophonous, and tuning systems to some network baseline would take time. The need to train users to use these systems, and understand the sounds of the networks monitored such that abnormalities could be identified, was also discussed as a potential challenge. The time required for adequate training of users, and for tuning of systems to networks, is a key factor affecting the utility of the approach.

Listening to sonification for extended time periods may be fatiguing. Fatigue caused by previous sonification designs has been reported [3] and was highlighted as a potential pitfall by a large number of participants. In integrating sonification into SOCs, therefore, it is important to consider mitigating fatiguing effects. Kramer argued that developing aesthetic sonifications can reduce listener fatigue, and prior work in such aesthetic sonification can be drawn on [54]. Another approach to mitigating fatigue, to be assessed experimentally, is to enable personalization of the sounds listened to.

Implications of findings on existing visual data presentation approaches

We reflect on our findings reported in section ‘Existing visual data presentation approaches’ on the approaches currently used in SOCs to present low-level network data and security tool output. We consider the way in which new sonification tools might solve the described challenges to using existing visual data presentation approaches. We also make recommendations for research into designing sonification systems that can interact effectively with visual data presentation approaches when working alongside these approaches. While we have already considered the integration of sonification with data presentation approaches to some extent in exploring the integration of sonification into SOC working practice

more broadly, the aim of this section is to provide a more specific identification of the types of data that might be suited to presentation through visual and sonic media, and how the two might interact, drawing on the more in-depth results on existing visual data presentation approaches reported in section ‘Existing visual data presentation approaches’.

Based on the described challenges to using existing visual data presentation approaches, we posit that there are a number of challenges that sonification might address. It was noted that there is a limitation to how much information can be conveyed visually, and that sonification might usefully be deployed to display extra information that could not be displayed through only visual means. In exploring whether visual and sonic data presentation methods can interact effectively in this way, it will be important to understand where the limit to the amount of information that can be conveyed sonically lies, as well as the limit to the amount of information that can be conveyed visually.

As described in the interviews, it can be challenging to manually review low-level presentations such as packet captures (during hunting tasks, for example) without guidance to point analysts ‘in the right direction’. Sonification could be an approach to providing such guidance during manual reviews through the sonic highlighting of anomalous packets and flows, and this strengthens the case for exploring the utility of sonification in hunting tasks (Use Case 1). The need for techniques that reliably draw the attention of security practitioners when required (for high-severity alerts, for example) was described. Currently, some SOCs use approaches such as colours and lights to address this need, and this is a clear area where sonification (particularly of alerts) may have the potential to attract attention while providing some information.

Identifying the potential for sonification to reduce challenges to using current visual data presentation approaches in the ways described above will require a greater understanding of the merits of using visual versus sonic approaches to present different types of information (e.g. traffic flows between particular IPs, severity of alerts) and to enable practitioners to conduct different types of task (e.g. hunting, monitoring of alerting systems). Such understanding should inform the distribution of content and tasks across these media. Participants in the interviews described the use of visual representations of the network and its zones to represent traffic flow and malware propagation, for example. Intuitively, we might expect that this type of spatial information is more effectively portrayed visually than sonically.

This question of information distribution across media, as well as the question of the limits to the amount of information representable through each media, could be explored through comparative user studies. The display of visual and sonic information needs to be tested in conjunction: it is likely that the amount of visual information users can comprehend will be lower when the user needs to simultaneously comprehend visual information, for example. It is important to note when designing such studies that participants may be more accustomed to using text-based data presentation than security visualizations, as indicated by Fig. 3.

It was noted that the visualizations of low-level network data are particularly effective when used to observe long-term trends. Sound, however, is temporal and if a sound that is part of a continuous sonification is not heard when it plays, it cannot be observed later (as a visual plot that remains on a screen can be). This has implications for the way in which sonification of low-level data can be used in comparison with visualization: in particular, if sonification is to be used to observe long-term trends over a time period, then it must be listened to continuously during that time period. Therefore, the

short-term use cases developed by participants (such as listening to a small stretch of the sonification once per hour) would not be suited to uses in which it is necessary to observe long-term trends. This should be taken into account when conducting further research into the use of sonification for network security monitoring: a possible solution may be to play long stretches of sonification compressed (sped up) over a short-time span.

Study limitations

Owing to the nature of the semi-structured interviews we conducted, there was variation in the level of detail in which different participants discussed each question. Furthermore, this article can report only those contexts of use, challenges in integration and design requirements highlighted in this study. It is possible that others would emerge in conversation with other participants. Consolidation of these findings through further studies would ensure coverage of all requirements.

The presentation and discussion of the technology with practitioners could have caused acquiescence bias in which participants agreed with statements by default. To mitigate this, we encouraged discussion around criticisms as well as positive responses in the interviews and explored challenges raised by participants pertaining to environmental factors, such as the noisiness of the SOC, complexity of networks and the distractions that could be caused by sonification.

Conclusion and future work

Working alongside existing visual data presentation approaches, sonification has promise as an approach to improving security practitioners' working practices in SOCs, based on SOC workflow and challenges, and evidence of the benefits sonification can offer. Using an online survey and semi-structured interview responses from practitioners, we explored perspectives on the incorporation of sonification into SOCs and identified key elements of current SOC working practice and data presentation. Our results show that security practitioners see high potential for the use of sonification in a range of use cases; in particular, for peripheral monitoring—while multitasking with other work tasks, or whilst outside of the SOC. Participants also saw value in using sonification for anomaly detection, in an approach similar to the existing visualization techniques used in SOCs.

We identified challenges in integration, and requirements for design, which should be addressed in future research. In order to be appropriate for a range of different SOC types, SOC soundscapes and practitioners' job roles, sonification tools should be flexible in design. More specifically, sonification should be playable through a range of mediums and suitable for a range of different types and lengths of use. Sonification of alerts was a key area highlighted for further design investigation, as well as approaches to mitigating listener fatigue.

As future work, we intend to address the design and integration questions highlighted in this study and to explore the possible interactions identified between sonification and existing data presentation methods. We also intend to validate experimentally the capability of SOC practitioners to use sonification in our refined contexts of use, compared with other SOC tools. We have noted the results of the experimentation we carried out subsequently to the research reported in this article, which suggest the utility of sonification when applied in some of the use cases identified, in an experimental setting [52, 53]. Experimentation with sonification in real SOC settings, and in realistically complex networks, will be key

to assessing the utility of sonification for SOCs and the effect of sonification on the SOC, and vice versa.

Funding

This work was supported by the Engineering and Physical Sciences Research Council (grant number EP/K035606/1).

Conflict of interest statement. None declared.

References

- [1]. Sundaramurthy SC, Bardas AG, Case J. *et al.* A human capital model for mitigating security analyst burnout. In: *Symposium on Usable Privacy and Security (SOUPS)*, Ottawa, Canada: USENIX, 2015, pp. 347–59.
- [2]. Kramer G, Walker, B, Bonebright, T. *et al.* The sonification report: status of the field and research agenda. Report prepared for the national science foundation by members of the international community for auditory display. In: *Proceedings of the International Conference on Auditory Display, Santa Fe, New Mexico*, 1999.
- [3]. Hermann T, Hunt A, Neuhoff J. *The Sonification Handbook*. Berlin: Logos Verlag, 2011.
- [4]. Axon L, Nurse JRC, Goldsmith M. *et al.* A formalised approach to designing sonification systems for network-security monitoring. *Int J Adv Secur* 2017;10:26–47.
- [5]. Ballora M, Giacobbe N, Hall D. '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. 80640P.
- [6]. Debashi M, Vickers P. Sonification of network traffic flow for monitoring and situational awareness. *PLoS One* 2018;13:e0195948.
- [7]. Debashi M, Vickers P. Sonification of network traffic for detecting and learning about botnet behaviour. In: *IEEE Access*, vol. 6, 2018, pp. 33826–39.
- [8]. Hildebrandt T, Hermann T, Rinderle-Ma S. Continuous sonification enhances adequacy of interactions in peripheral process monitoring. *Int J Hum Comput Stud* 2016;95:54–65.
- [9]. Bevan N. 'International standards for hci and usability'. *Int J Hum Comput Stud* 2001;55:533–52.
- [10]. Maguire M, Bevan N. User requirements analysis. In: *IFIP World Computer Congress, TC 13*. Boston, MA: Springer, 2002, pp. 133–48.
- [11]. Gulliksen J, Göransson B, Boivie I. *et al.* Key principles for user-centred systems design. *Behav Inform Technol* 2003;22:397–409.
- [12]. Axon LM, Alahmadi B, Nurse JR. *et al.* Sonification in security operations centres: what do security practitioners think?. In: *Workshop on Usable Security (USEC, NDSS)*, San Diego, California, 2018.
- [13]. Sundaramurthy SC, Wesch M, Ou X. *et al.* Humans are dynamic—our tools should be too. *IEEE Internet Comput* 2017;21:40–6.
- [14]. Navy Petty Officer 1st Class Shane Wallenda. *Colorado Springs, CO - U.S. Northern Command Joint Operations Center*, 2013, U.S. Northern Command, [https://www.northcom.mil/Images/igphoto/200002015/10 February 2020, date last accessed](https://www.northcom.mil/Images/igphoto/200002015/10%20February%202020,%20date%20last%20accessed%20).
- [15]. Sundaramurthy SC, Case J, Truong T. *et al.* A tale of three security operation centers. In: *Proceedings of the 2014 ACM Workshop on Security Information Workers*. Scottsdale, Arizona: ACM, 2014, pp. 43–50.
- [16]. D'Amico A, Buchanan L, Kirkpatrick D. *et al.* Cyber operator perspectives on security visualization. In: *Advances in Human Factors in Cybersecurity*. Florida: Springer, 2016, pp. 69–81.
- [17]. Zhang Y, Xiao Y, Chen M. *et al.* A survey of security visualization for computer network logs. *Secur Commun Netw* 2012;5:404–21.
- [18]. Botta D, Werlinger R, Gagné A. *et al.* Towards understanding it security professionals and their tools. In: *Proceedings of the 3rd Symposium on Usable Privacy and Security*. Pittsburgh, Pennsylvania: ACM, 2007, pp. 100–11.
- [19]. Hall P, Heath C, Coles-Kemp L. Critical visualization: a case for rethinking how we visualize risk and security. *J Cybersec* 2015;1:93–108.

- [20]. D'Amico A, Whitley K. 'The real work of computer network defense analysts. In: *VizSEC 2007*. Springer, 2008, pp. 19–37.
- [21]. Werlinger R, Hawkey K, Beznosov K. An integrated view of human, organizational, and technological challenges of it security management. *IMCS 2009*;17:4–19.
- [22]. Werlinger R, Muldner K, Hawkey K. *et al.* Preparation, detection, and analysis: the diagnostic work of it security incident response. *Information Management & Computer Security 2010*;18:26–42.
- [23]. D'Amico A, Whitley K, Tesone D. *et al.* Achieving cyber defense situational awareness: a cognitive task analysis of information assurance analysts. In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 49, No. 3, Los Angeles, CA, SAGE Publications, 2005, pp. 229–33.
- [24]. Werlinger R, Hawkey K, Beznosov K. Security practitioners in context: their activities and interactions. In: *CHI'08 Extended Abstracts on Human Factors in Computing Systems*. Florence, Italy: ACM, 2008, pp. 3789–94.
- [25]. Sundaramurthy SC, McHugh J, Ou X. *et al.* 'Turning contradictions into innovations or: how we learned to stop whining and improve security operations. In: *Symposium on Usable Privacy and Security (SOUPS)*, 2016.
- [26]. Werlinger R, Hawkey K, Botta D. *et al.* Security practitioners in context: their activities and interactions with other stakeholders within organizations. *Int J Hum Comput Stud 2009*;67:584–606.
- [27]. Brown A, Martin M, Kapralos B. *et al.* Poster: towards music-assisted intrusion detection. In: *poster presented at IEEE Workshop on Statistical Signal Processing*, 2009.
- [28]. Gilfix M, Couch A. Peep (the network auralizer): monitoring your network with sound. In: *Proceedings of the Large Installation System Administration Conference*. San Diego, California: USENIX, 2000, pp. 109–17.
- [29]. Giot R, Courbe Y. Intention–interactive network sonification. In: *Proceedings of the International Conference on Auditory Display*. Atlanta, Georgia: Georgia Institute of Technology, 2012, pp. 235–6.
- [30]. Mancuso VF, Greenlee ET, Funke G. *et al.* Augmenting cyber defender performance and workload through sonified displays. *Procedia Manuf 2015*;3:5214–221.
- [31]. Papadopoulos C, Kyriakakis C, Sawchuk A. *et al.* Cyberseer: 3d audio-visual immersion for network security and management. In: *Proceedings of the ACM Workshop on Visualization and Data Mining for Computer Security*. Washington DC: ACM, 2004, pp. 90–8.
- [32]. Qi L, Martin M, Kapralos B. *et al.* Toward sound-assisted intrusion detection systems. In: *On the Move to Meaningful Internet Systems 2007: CoopIS, DOA, ODBASE, GADA, and Is*. Vilamoura, Portugal: Springer, 2007, pp. 1634–45.
- [33]. Gopinath M. Auralization of intrusion detection system using Jlisten. *Development 2004*;22:3.
- [34]. Kaczmarek T, Kobsa A, Sy R. *et al.* An unattended study of users performing security critical tasks under adversarial noise. In: *Proceedings of the NDSS Workshop on Useable Security, San Diego, CA*, Internet Society, 2015, p. 14.
- [35]. Berg B, Kaczmarek T, Kobsa A. *et al.* An exploration of the effects of sensory stimuli on the completion of security tasks. *IEEE Priv Secur 2017*;15:6.
- [36]. Söderlund G. Positive effects of noise on cognitive performance: explaining the moderate brain arousal model. In: *Proceedings of the 9th Congress of the International Commission on the Biological Effects of Noise*. Mashantucket, Connecticut: *Leibniz Gemeinschaft*, 2008, pp. 378–86.
- [37]. Nevo B. Face validity revisited. *J Educ Meas 1985*;22:287–93.
- [38]. Maguire M. Context of use within usability activities. *Int J Hum Comput Stud 2001*;55:453–83.
- [39]. Baier G, Hermann T, Stephani U. Event-based sonification of EEG rhythms in real time. *Clin Neurophysiol 2007*;118:1377–86.
- [40]. Baldassi S, Megna N, Burr DC. Visual clutter causes high-magnitude errors. *PLoS Biol 2006*;4:e56.
- [41]. Ballora M, Cole RJ, Kruesi H. *et al.* Use of sonification in the detection of anomalous events. In: *Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications. Proceedings of SPIE - The International Society for Optical Engineering*. International Society for Optics and Photonics, 2012, vol. 8407, pp. 84070S.
- [42]. Etoty RE, Erbacher RF. A survey of visualization tools assessed for anomaly-based intrusion detection analysis. DTIC Document, Technical report. Adelphi: Army Research Laboratory, 2014.
- [43]. Merced D, Wanda L. Sound for the exploration of space physics data. Ph.D. Dissertation, University of Glasgow, 2013.
- [44]. van Ee R, van Boxtel JJ, Parker AL. *et al.* 'Multisensory congruency as a mechanism for attentional control over perceptual selection. *J Neurosci 2009*;29:11641–9.
- [45]. Likert R. A technique for the measurement of attitudes.' *Arch Psychol 1932*;22:55.
- [46]. Kaptein MC, Nass C, Markopoulos P. Powerful and consistent analysis of Likert-type rating scales. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Atlanta, Georgia: ACM, 2010, pp. 2391–4.
- [47]. Nurse JRC, Creese S, Goldsmith M. *et al.* Guidelines for usable cybersecurity: past and present. In: *2011 Third International Workshop on Cyberspace Safety and Security (CSS)*. Milan, Italy: IEEE, 2011, pp. 21–6.
- [48]. Jamieson S. Likert scales: how to (ab) use them. *Med Educ 2004*;38:1217–8.
- [49]. Norman G. Likert scales, levels of measurement and the 'laws' of statistics. *Adv Health Sci Educ 2010*;15:625–32.
- [50]. Robertson J. Likert-type scales, statistical methods, and effect sizes. *Commun ACM 2012*;55:6–7.
- [51]. King N. Template analysis. In: Symon G, Cassell C (eds), *Qualitative Methods and Analysis in Organisational Research: A Practical Guide*. London: Sage Publications Ltd, 1998, pp. 118–34.
- [52]. Axon L, Happa J, Goldsmith M. *et al.* Hearing attacks in network data: an effectiveness study. *Comput Secur 2019*;83:367–88.
- [53]. Axon L, Happa J, Janse van Rensburg A. *et al.* Sonification to support the monitoring tasks of security operations centres. In: *IEEE Transactions on Dependable and Secure Computing*, 2019.
- [54]. Kramer G. *Auditory Display: Sonification, Audification, and Auditory Interfaces*. Boston: Addison-Wesley Longman Publishing Co., Inc, 2000.