Designing Through The Stack: The Case for a Participatory Digital Security By Design

Ian Slesinger
Lizzie Coles-Kemp
Niki Panteli
Ian.Slesinger@rhul.ac.uk
Lizzie.Coles-Kemp@rhul.ac.uk
Niki.Panteli@rhul.ac.uk
Royal Holloway University of London
England

René Rydhof Hansen
rrh@cs.aau.dk
Aalborg University
Denmark

ABSTRACT

Whilst participatory practice is increasingly adopted in end user studies, there has been far less use of a participatory approach when designing lower down the software stack. As a result, end users are often presented with security controls over which they have no control but for which they retain the responsibility. Conversely, hardware and software engineers struggle to innovate new security control designs that are resilient to new and emerging threats. In a study utilising ethnographic research and stakeholder interviews, we show that there is a siloing of communities of practice between hardware security engineers, software engineers and coders, manufacturers in the technology supply chain and end users. Our findings indicate that this siloing and a lack of participatory practice impedes the development of a more cohesive digital security design that integrates security through the stack from the hardware layer upwards to the OS and application layers. These barriers make difficult the negotiation between what is possible lower down the stack with what is needed and wanted higher up the stack. Our findings suggest that a more holistic and comprehensive participatory design approach is required to negotiate a digital security by design paradigm that more evenly distributes power over and responsibility for security controls throughout the stack. Working with the HCI literature on co-production in design, this paper will suggest that a pathway for breaking through this impasse is to utilise objects in the design process of the hardware secure instruction set architecture as a feedback mechanism to incorporate other sets of designers and users in the design process to create a more workable stack.

CCS CONCEPTS

• **Security and privacy** → *Social aspects of security and privacy*; **Security in hardware**; • **Human-centered computing** → *User studies.*

NSPW '22, October 24–27, 2022, North Conway, NH, USA

ACM ISBN 978-1-4503-9866-4/22/10…$15.00
https://doi.org/10.1145/3584318.3584322

KEYWORDS

participatory design, hardware security engineering, software security engineering, digital security by design

ACM Reference Format:
Ian Slesinger, Lizzie Coles-Kemp, Niki Panteli, and René Rydhof Hansen. 2022. Designing Through The Stack: The Case for a Participatory Digital Security By Design. In *New Security Paradigms Workshop (NSPW '22), October 24–27, 2022, North Conway, NH, USA.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3584318.3584322

1 INTRODUCTION

The UK government-sponsored Digital Security by Design (DSbD) programme sets out to establish hardware security as a disruptive solution to address memory-related security vulnerabilities that will break the cycle of what the programme's architects call a "market failure" in cybersecurity innovation [45]. DSbD is a UK government programme designed to further develop a technology called CHERI (Capability Hardware Enhanced RISC Instructions) and support its journey to market and adoption. CHERI originally emerged from a joint research project of SRI International and the University of Cambridge. It is important to note that CHERI is not only a UK initiative but since 2010 was originally funded by DARPA CRASH, MRC, and SSITH programs, as well as other DARPA research and transition funding. Since 2019 the UK's research and innovation fund (UKRI) has has supported the development of Arm's experimental CHERI-enabled Morello processor, SoC, and board.

CHERI takes a novel approach to systems security that re-examines and re-designs fundamental concepts and building blocks of hardware and software systems with a strong focus on security and enabling development of secure systems. The perceived market failure is a result of the low risk appetite of technology manufacturers to take on the cost of innovating hardware solutions without guaranteed market penetration or adoption. As part of our work interrogating the disruptive nature of the DSbD proposition from a socio-technical perspective on cybersecurity, we present a study that asks what roles participatory design might play in breaking the impasse. In this paper we focus on two primary research questions. Firstly, what are the current barriers and challenges to participatory practice within the DSbD ecosystem? Secondly, how might participatory design and co-design approaches be used to cohere hardware security with the wider concept of computer security?

Building upon research on participatory design from the Human Computer Interaction (HCI) literature, this paper will suggest a paradigm shift in how security is designed through the software stack from the hardware level at the base of the stack, through the application layer and onto the human-computer interaction extensions that sit on top of the traditional 7-layer OSI reference model [6]. The applicability of traditional user-centred and participatory design techniques for developer-centred security design has been questioned [46] in the context of security development. In response to this critique, we place emphasis on meaning-making within the participatory design process and explore how collaborative meaning-making might work for the different stakeholders in the development and adoption of CHERI, rather than emphasise how such stakeholders might adopt participatory and user-centred techniques. Such a change will necessitate a more systemic approach to how the stack is designed that puts multi-stakeholder participation at the heart of the design process. Such an approach will help shift security practice away from the current 'patch and pray' approach to managing security vulnerabilities, as well as anticipate potential unforeseen vulnerabilities in future security solutions, such as CHERI, that seek to address existing vulnerabilities foundationally at the hardware level at the base of the stack.

This paper begins by setting out how participatory design incorporates democratic values in the innovation process, and how the role of participatory design in the human-computer interaction literature indicates its value for security design. We then present our study and its findings. Finally, we discuss the implications of our findings for security design at the base of the stack and set out the ways in which participatory design might contribute to a more integrative digital security by design approach.

2 BACKGROUND LITERATURE

Participatory design is an approach to designing digital technology and technologically-enabled processes that seeks to actively include the intended users of the digital products in the design process. The core principle of participatory design is that users become equal partners in the design process, working alongside with the ICT designers [41]. Participatory design has its roots in Scandinavian sociotechnical design culture [41]. In 1993, Communications of the ACM published a special issue on participatory design and the breadth of the topics covered in this issue reflected the range of participatory design activities, including co-design [42]. In this literature review section, we first set out the meaning of the concept and its adaptation from sociological Science and Technology Studies (STS) to HCI, and then to computer security more specifically. In the following subsections we examine how this principle has shaped the software design paradigm.

2.1 Design and The Participatory Perspective

Design scholar Pelle Ehn explains that "participatory design started from the simple standpoint that those affected by a design should have a say in the design process" [25]. He goes on to note that this process is not necessarily a means to achieve a passive consensus, but rather a way to explore and negotiate the "controversies and conflicts around an emerging design object." Design scholar DiSalvo considers participatory design to be a means of "critical making... through which the political qualities of an issue are materialised by participatory means" [21, p.96]. A key aspect of this is to draw out the practical knowledge possessed by members of a community of practice to "inform the design of products they might use, in order to make them more useful and usable" and to include them as stakeholders "in the design and use of products that might alter their practice" [21, p.97]. However, participatory design is not a purely cooperative enterprise in which perspectives and approaches to design naturally cohere [34]. Mogensen and Trigg [40, p.55] point out how confrontations between how designers contextualise artifacts and how users and practitioners conceptualise their own work practices and experience are formative in shaping "new understandings of current practice as well as possible futures."

Initially the participatory design literature tended to focus on innovation within an organisation rather than across organisations [2]. Grudin [31] and Carlile [15] both argue that the division of responsibilities within a large organisation is an impediment to innovation and the adoption of participatory design techniques. Likewise, Sanderson (1996) evaluates the use of participatory design in small scale computing innovation, emphasising that there is a scalar dimension to the implementation of participatory design that distinguishes it from more traditional software development processes. More recently, researchers have adopted participatory design principles to support the growth and sustainability of online communities (e.g. [14, 30]).

2.2 Design Through a Political Lens

Participatory design makes explicit the politics of technology design in a process where participants are challenged as to what values are being reflected through the design process [14]. According to Empsak [26, p.111] participatory design is an inherently political technique to enable "robust democracy" by moving "decision making possibilities" and "control over technologies" away from "elites" to allow for "the population to be able to exercise meaningful control over the technologies that determine our lives." However, the distribution of responsibility is not only structural but also political. Responsibility distribution is also fundamental to the power redistribution dimension of participatory design. The principle of all participants having the responsibility of looking after each other helps to ensure that power is fairly distributed across the design process [54]. The question of power distribution extends beyond the communities that are participating but also places questions of social justice and lifelong responsibility for the digital technologies that are produced at the centre of the participatory design process [7]. From the beginning it was argued that it was important to recognise the political power of participatory design so that it is not confused with a narrower user-centred design agenda. For example, Van den Besselaar cautions that the emancipatory and democratic potential of participatory design is being supplanted by a "narrower" agenda of "improving systems for users" in service of naturalising technology as a "neutral" fact [53].

Structurally, redistribution of responsibilities within the design process is a hallmark of participatory design [31]. For example, Kensing et al. [35] highlight that in participatory design it is the responsibility of IT professionals to critically reflect on the framing

of the software to be designed. To reflect this principle, in the participatory design framework that Kensing et al. developed, the IT professionals were assigned the responsibility to choose the participatory approach and those managing the process were assigned the responsibility to ensure that users had sufficient time to participate in the process. Beirne and Ramsay [9] who examined the complex structures of this approach, showed evidence of where users were selectively incorporated into formal design processes depending upon their position and status, rather than their knowledge per se. As such, according to these authors, established power structures may constrain effective participatory design in practice, promoting instead a conservative and managerialist orientation.

2.3 The Growth of Participatory Design

Initially there were debates as to whether participatory design was possible to implement beyond the Scandinavian regulatory and cultural environment [41]. Nevertheless, this approach with its emancipatory characteristics captured the attention of researchers elsewhere especially in North America with the work of Greenbaum, J. and Kyng, 'Design at Work' [29] being a notable example of the attention that participatory design had received internationally. Furthermore, with the rise of the study of practice in HCI [37] and movements such as digital civics [44, 57] that study both the design and use of technology within its social, economic and political settings, participatory design has become a core component of a growing number of HCI studies [56]. Computer science is a broad church of different disciplinary positions and research methods are a key means of differentiating between those positions [51]. The methods used to encourage participation in technology studies are distinctive in their use of creative engagement techniques [24, 56] which are often used as a means to blend traditional social science approaches with design and computer science methods.

2.4 Participatory Design and Computer Security

Participatory design in computer security accompanied the move towards usable security and people-centred security design. Saltzer and Schroeder [48] produced one of the first papers that sets out the importance of ensuring that security technology is usable by its intended user base. Usable security scholars such as Adams and Sasse [1] argued that as a result of this focus on mathematical approaches, technical security practices such as application-specific passwords are often poorly adapted to the needs and behaviours of users. They suggest that an alternative paradigm for designing technological security is to adopt a user-centred design approach that better engages with users to co-produce a design that better meets their needs in real-world environments and applications. Whilst this paper is less clear in setting out what a user-oriented design process ought to look like, the relevance of participatory design to user-centred approaches to security is clear. Similarly, Ashenden and Ollis [4, p.35] suggest that attending to the "lived experience" and "social practices" of software developers can be a means to better integrate computer security in software development. Such moves within the computer security literature reflect a desire to both engage with and understand the needs of those designing and using security technologies.

One way to better understand the lived experience is to use a participatory approach both to research and to technology design. There has been a growth in participatory studies to identify the security and privacy needs of vulnerable communities. For example, Balcerzak et al. [5] provide the example of security specific participatory design in the creation of an app for older users in Poland as a way of better developing HCI to meet the needs of a specific vulnerable user group. Slupska et al. [50] develop such an approach further by conceptualising security design as a set of practices rather than properties in a system. On this basis, participatory research becomes beneficial for developing a user-centred approach to threat modelling. Their approach critiques the role of "expert" knowledge in cybersecurity as monolithic and ignorant of what counts as security for marginalised and underserved groups of users. Participatory design is also used to re-imagine universal security functionality such as authentication [27].

As part of this growing interest in participatory research and design in computer security, NSPW is an interdisciplinary computer security venue that has promoted a number of participatory design paradigms in its history. Wang [58] published an NSPW paper on inclusive security and privacy that championed the use of value sensitive design and participatory approaches to achieving this. In 2014 NSPW included a paper by HCI scholars that set out the use of participatory design featuring creative techniques to elicit data that was otherwise hard to reach on issues relating to information sharing and protection [23]. NSPW has also supported experimental work that takes participatory design beyond its more traditional use of exploring end user requirements and interface design, into fundamental questions of computer security architecture. For example, in 2021 NSPW included a paper that considers how participatory action might work in practice in order to collectively set access control rules [28]. NSPW has also been an influential forum to challenge the limitations of participatory and user-centred design in the context of privacy and security [46].

Collaboration with the HCI community has resulted in a strong creative engagement practice within computer security studies and the rise of creative engagement methods that integrate social and political perspectives on security with the form and structure of security technologies. Brandt et al. [13] set out standard participatory design techniques that can be deployed for the telling and re-telling of security issues. Collard and Briggs [19] reviewed the participatory engagement projects in the UK's socio-technical security programme, TIPS. They identified a range of participatory design toolkits that could be used to explore social aspects of computer security. Collard and Briggs' study revealed that participatory design toolkits were being used, adapted and extended to explore numerous aspects of socio-technical security. In particular, the participatory design approach uses creative engagement techniques that enable participants to reflect on and articulate their views, concerns and hopes related to information sharing and control [24]. Participatory design techniques also offer approaches that enable researchers to work with hard-to-reach communities such as older technology users [23]. Such techniques also enable multi-generational consultation on topics of information sharing and protection. For example Bowyer et al. [12] use a participatory toolkit to elicit the views of families on the sharing of personal data to deliver statutory and civic services. Participatory design

is particularly useful when trying to develop security and privacy techniques in resource-constrained environments [55]. Participatory design also enables the recognition dimensions of community security practice not typically included in computer security designs. For example, Chouhan et al [16] used community design to develop a community oversight function for privacy and security decisions. Collard and Briggs [19] also point to the fact that the use of participatory design is not only in the research domain but also in the practice domain [43].

2.5 The Significance of Boundary Objects

There are a number of challenges to deploying participatory design approaches in a community of technological security practitioners that Zurko and Simon [62, p.28] describe as prioritising "mathematical rigour" over "usability" in a way that ignores the needs of users. One way to overcome some of these challenges is to examine how the acts of coding and digital making might become sites of participation and collaboration across different communities of technologists.

Understanding the process of meaning-making in participatory design requires attending to the pivotal role in the design process of artifacts. This topic is best articulated in the Science and Technology Studies (STS) literature. Particularly useful is a consideration of what S.L. Star [52, pp.251–252] calls boundary objects. Star describes boundary objects as "both plastic enough to adapt to local needs and constraints of the several parties employing them, yet robust enough to maintain a common identity across sites." Harper et al. [33] conceptualise the computer file as a boundary object that provides a "grammar of action... for both users and engineers" through which they can "cohere their plans, their doings, their goals, around files and the other things they want to do in the age of the cloud." Lee [38, pp. 314, 333] suggests a more nuanced distinction be made between boundary objects as a basis for standardisation that is negotiated between stakeholders, and "boundary negotiating artifacts" which allow for "incipient, non-routine and novel collaborations" that can "create shared understanding about specific design problems." In specific relation to computer security, Chung et al. [18] align the notion of data as a boundary negotiating artifact with the need to address the privacy and security concerns of users in patient-healthcare provider data sharing interactions.

Putting boundary negotiating artifacts at the centre of a participatory design project can allow an account to emerge for the diverse ways in which different stakeholders (e.g. computer scientists, engineers, business people, civil servants, cybersecurity analysts, end users) attach different meanings to it. This allows designers to better understand how an object works, how to use it, and what its affordances and possible applications are.

In the following section we set out the context for the study by providing an overview of the DSbD programme and the CHERI and Morello technologies within that programme.

3 THE CONTEXT FOR THE STUDY

The context for our study is the Digital Security by Design (DSbD) programme, a UK government programme designed to support the further development of the CHERI technology and support its journey to market and adoption. Originally developed through international programmes such as the DARPA CRASH programme, CHERI re-examines and re-designs fundamental concepts and building blocks of hardware and software systems with a strong focus on security and enabling development of secure systems. The CHERI design practices and approach highlight how the boundaries between hardware and software engineering is blurring. In essence, CHERI is a model or an architecture for adding architectural capabilities to an instruction set architecture (ISA), allowing fine-grained control over memory access. This is achieved by extending the ISA with capability aware instructions enforcing that reads and writes from and to memory must be authorised by a corresponding capability. In particular this prevents attackers from "hijacking" and abusing raw pointers and enables software to define memory regions for data and code that are isolated from each other, thus implementing compartmentalisation, that further prevents a successful attack from propagating and/or escalating to other more sensitive memory regions. By targeting memory safety, the CHERI project is directly addressing several of the largest classes of security vulnerabilities commonly found in software, including buffer overflows and use-after-free bugs, that are notoriously difficult to detect and avoid.

The first version of the CHERI ISA was completed in 2010, and since then the ISA has undergone a series of updates [60, pp.29–30]. To ensure the correctness of the design, the CHERI extension has been formally modelled and verified using mechanised theorem proving. The same formal models also form the basis for synthesised simulators and automated tests. The ISA extension is sufficiently generic that it is relatively straightforward to port it to most common ISAs including ARM, x86, and RISC-V. It is furthermore designed as a hybrid extension that is able to mix traditional memory models with the capability based memory model allowing for full backwards compatibility and a gradual switch to the new model. In practice this is essential, since the hardware security can not stand alone: it requires significant investment and resources in developing new (or enhancing old) development software tools and systems, not least compilers and operating systems, to take advantage of the capability model.

3.1 DSbD Ecosystem

Figure 1 shows a map of the DSbD ecosystem produced through an ethnographic exercise carried out by the research team to map the links connecting the current DSbD programme stakeholders, and their relation to other sets of stakeholders they must bring on-board for CHERI to be successfully adopted. The diagram is initially divided into three grey boxes that characterise the generalised roles of stakeholders, and within these the colour-coded rectangular boxes represent categories or groupings of organisations and the triangles represent specific organisations related to each of these categories. The technologists are the loci of technical expertise both invested in bringing CHERI to fruition, and working on the technical engineering of CHERI to make it operable. However, the ten SME projects occupy a somewhat liminal space within the technologists grouping as not all of them are strongly linked to the key CHERI insiders, as indicated by the unidirectional arrows between. Furthermore, there is a geographical link between Arm—which is headquartered in Cambridge—and the university,

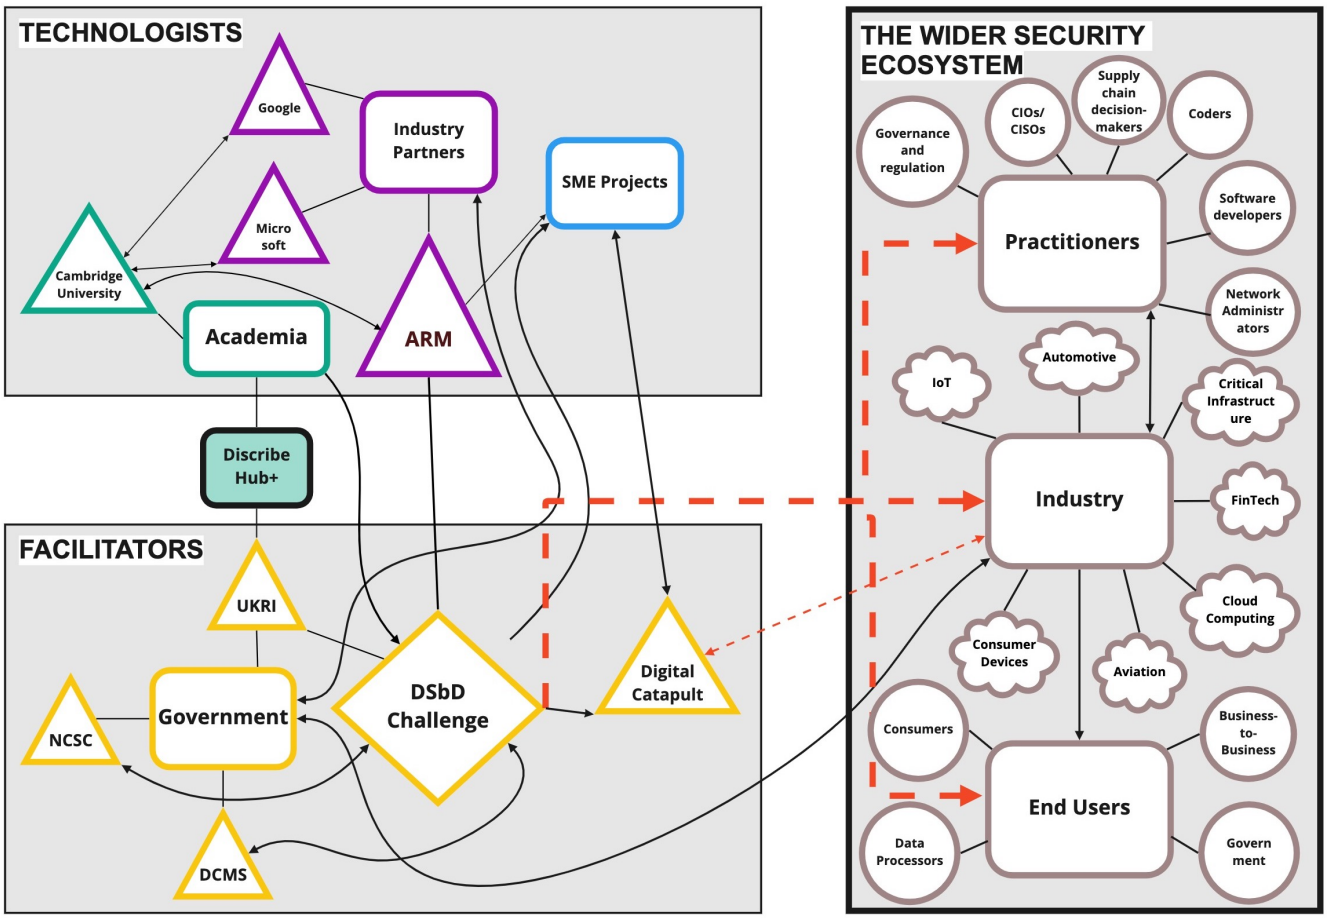

Figure 1: DSbD Ecosystem. (Rectangular boxes represent organisational groupings with corresponding colour-codes for each grouping; triangles represent specific organisations within those groupings; black lines connected to boxes indicate belonging to a grouping; black lines with arrows represent existing links between groupings and organisations [note directionality of arrows]; dashed red lines represent underdeveloped links [note directionality and thickness]; circles represent specific roles; clouds represent industry sectors.)

which facilitates knowledge exchange and the flow of personnel between the two organisations. The Facilitators are organisations with mostly non-technical or socio-technical missions who are engaged in the socio-technical, political and economic work of furthering CHERI adoption. It is perhaps notable that this grouping currently almost entirely consists of UK government departments and UK government-sponsored third sector organisations. Another significant feature is the central role of the DSbD challenge and the restrictions on the way in which the challenge might be interpreted. The 'wider security ecosystem' category contains other potential stakeholders that are outside the CHERI programme, and whom DSbD is attempting to bring on-board. This separation is denoted in the diagram by the bold border around the 'wider security ecosystem' which emphasises the strong barrier between the CHERI insiders and outsiders. Likewise, the dashed red lines indicate the key pathways of communication that are not yet developed, and whilst the government-sponsored Digital Catapult innovation incubator has some relationships with industry, the unidirectional flow of communication between the main DSbD Challenge and industry is, at the time of writing, the most significant barrier for the adoption of CHERI, which is illustrated using bold red lines.

The DSbD also contains social science research activities. The Discribe Hub+ (which the present research is part of) is a UK Economics and Social Research Council-funded research project within the wider DSbD programme that applies an inter-disciplinary social scientific approach to investigate questions around adoption, regulation, social attitudes and futures relating to the CHERI technology.

3.2 Digital Security by Design and Morello

In 2019, the Digital Security by Design programme was launched by UKRI in cooperation with DCMS and the NCSC with funding from the Industrial Challenge Research Fund (ICRF). The DSbD programme seeks to develop a security ecosystem in which CHERI

can operate. The stated purpose of DSbD is to break the cycle of a perceived "market failure" blocking the wholesale adoption of hardware security. This market failure is defined in two ways. One is the 'chicken-and-egg' cycle whereby technology enterprises don't want to take the financial risk of developing new security hardware without the software ecosystem to support it. The other is the unwillingness of the market to pay a negative externality for digital security where the return on investment and benefits are unclear or intangible. To date one of the most tangible outputs of DSbD is Arm's Morello programme which has developed an experimental hardware demonstration board implementing CHERI within the ARM ISA. This experimental prototype is being trialled by SME and larger-scale enterprises to identify uses, relevant features and to develop some of the software interface necessary to create a viable CHERI stack.

In the following study we examine the case for participatory design in the context of CHERI and of DSbD more generally, we identify some of the implications and impacts to the DSbD ecosystem of not using a participatory design approach, and we pinpoint where the desire for a more participatory approach surfaces in the DSbD context.

4 STUDY DESIGN

We conducted two packages of interviews. The first set of baseline interviews sought to gain an overview of the CHERI design process and the goals of the CHERI programme from key individuals within the DSbD programme and the wider hardware security community. The second package of interviews with SME stakeholders who received DSbD funding was designed to understand how the design process played out in practice. An anonymised list of stakeholders we interviewed across both interview packages is shown in (Table 1).

Because the baseline interviewees are individuals either closely affiliated with DSbD programme or the wider industry with a potential stake in the CHERI technology, there was greater sensitivity around their participation. To reduce the level of sensitivity, we made the decision to operate on a notes-only basis and to use these interviews to gather a background understanding of DSbD. However, we chose to take the more traditional approach of recording and transcribing interviews with the SME research because participants were less directly involved in the core DSbD programme, and we believed recording interviews would allow us to obtain direct quotes and a more nuanced analysis.

Baseline Interviews: To gain a baseline perspective on how key stakeholders understand CHERI and the related DSbD concepts, the research team initially conducted eight semi-structured interviews with individuals in senior professional positions closely linked to the DSbD programme in industry, government and academia. This involved engaging with participants in a guided discussion that addressed key themes and talking points rather than covering a fixed order of questions. This approach encouraged the interviewer to ask probing questions, and for participants to expand upon and elaborate their ideas. Specific themes and questions included asking participants to define CHERI, DSbD and Morello, then asking them to map out how these key concepts related to one another. We also asked participants to define the key stakeholders in DSbD from their perspective and evaluate the extent to which their own understanding of CHERI has changed over the course of their involvement with DSbD. As a final provocation, we asked participants to speculate on what the likely first application of CHERI would be, and what actors would be sufficiently motivated to bring it to market.

As explained at the start of this section, notes were generated in the first round of interviews. In order to ensure greater accuracy a second researcher took notes as well. These notes were then collated and sent back to the participant for additional input and approval. This collaboration provided verification to offset any potential inaccuracies or missed information in our notes and allowed us to create a "version of record" that we could use as a basis for analysis and evidence to support our findings and discussion.

SME Interviews: A second package of interviews was then conducted with five SME organisations who were awarded six months funding for projects that tested the CHERI stack in relation to their business needs through a virtual platform environment (VPE) based around the FreeBSD platform. These interviews focused on the SMEs' experiences with the projects including successes, challenges, their approach to CHERI in addition to their understanding of the socio-technical landscape surrounding CHERI. The interviews were recorded and professionally transcribed. In addition to the same questions asked in the baseline interviews described above, we asked these participants to reflect on their experiences in conducting their respective DSbD project, including successes and challenges faced, as well as whether their expectations of the technology were met. We also asked a sequence of questions on what the participants anticipated their regulatory responsibilities would be if they were to implement CHERI technology on a commercial basis, and how they would define the failure of a CHERI-enabled chip and what the consequences of such a failure might look like.

A full list of questions is provided in Appendix A and B. Both sets of interviews were then qualitatively analysed to identify key and recurrent themes across the data. Furthermore, ethical approval for the research was obtained in line with the primary institution's research ethics policy. In addition, the paper was also reviewed by the Discribe Hub+ advisory board prior to completion.

5 FINDINGS

In this section we present the findings from our interviews. Woven through the findings that we present in this section, a clear desire is expressed for a more participatory approach to the development of CHERI. There is an understanding that DSbD is potentially a transformational programme and there is motivation to participate in that transformation from across the ecosystem. However, our evidence suggests that at present potential users are being excluded from having a meaningful stake in the design process and a say in what a security ecosystem with CHERI ought to look like in practice, both within and across market sectors and use cases. We argue this exclusion is due to the organisation and practices of the DSbD programme.

Our findings are presented in six themes: Broken Market where interviewees highlighted the difficulties that the economic model posed to bringing in changes to security at the hardware level; the

Baseline interviews	Participant 1	Academic with significant role in DSbD Programme
	Participant 2	Industry expert affiliated with DSbD
	Participant 3	Academic with significant role in DSbD Programme
	Participant 4	Industry expert not affiliated with DSbD
	Participants 5 & 6	Civil servants affiliated with DSbD (from same department, interviewed together)
	Participant 7	Civil servant affiliated with DSbD
	Participant 8	Industry expert affiliated with DSbD
	Participant 9	Industry expert affiliated with DSbD
SME interviews	Participant 1	CEO of Aerospace and Defence Software SME
	Participant 2	Co-founder of Healthcare Technology SME
	Participant 3	CEO of Privacy Technology SME
	Participant 4	Co-founder of Enterprise Software Services SME
	Participant 5	Co-founder of Software Security Consultancy

Table 1: List of participants

framing of DSbD as an enabler of Disruptive Technology innovation, the metaphor of CHERI technology as a Magic Chip; Barriers to Engaging at Scale affecting DSbD; Siloed Communities within the DSbD ecosystem; and finally Desire for a More Participatory Approach to engagement in the design process within DSbD and at its periphery. These themes speak directly to the lack of a participatory design process in the DSbD programme, the challenges to adopting a participatory approach, and the impact of not following a participatory approach on the programme's objectives.

5.1 Broken Market

The economic model underpinning technology development shapes the nature of participatory engagement. This is partly because resources have to be made available to undertake participatory activity. It is also because the economic model is a major factor in distributing responsibility for innovation, and in the case of DSbD the responsibility for security innovation. Our findings reflect a debate about responsibility and resource allocation through the metaphor of the broken market.

A recurring narrative within the DSbD community is that the programme's main objective is to fix the "market failure" or "broken market" around innovation in hardware security. However, what the perceived breakage or failure is and the reasons for failure were multiple. The first explanation was that there is a cyclical dependency whereby the lack of new hardware development leads to a lack of software development that would support the new hardware, and that new hardware is not being developed as there is an absence of software innovation to make the hardware operable. The second type of market failure is that there is an "energy-investment" problem in industry in which companies are unwilling to take on the costs and heavy financial risks of hardware innovation without a near-guaranteed and significant return-on-investment. The third failure relates to the mismatch between the "four to five-year" length of the research and development cycle for new silicon, and the one to two year cycle for software development, as well as for consumer devices. The fourth problem identified by participants' is more specific to hardware security. This claims that security is perceived as a negative externality by businesses, meaning it is an additional cost to be incurred to hedge risk, rather than a positive driver of growth or greater return. This means that companies are less willing to pay higher costs than minimally necessary for security.

5.2 Disruptive Technology

Our findings reflect a dissensus over the purpose and uses of the CHERI technology within the DSbD ecosystem, which partly arise as a result of a lack of participatory engagement across the ecosystem. Disruption was seen as a means of calling out this dissensus and changing the direction of hardware security. A frequent narrative trope in our discussions around DSbD with participants and in the public-facing discourse being projected by the DSbD programme is that CHERI is a disruptive technology, and similarly that DSbD is a catalyst for disruption of the "broken" market cycle. According to one industry participant involved in the work of DSbD, the CHERI chip is disruptive because it overcomes industry inertia against adopting hardware security solutions, which is partly due to their hard to quantify benefit that fails to attract buy-in. Another DSbD insider from academia stated that "if you don't disrupt things you aren't going to fix the problem."

Our findings indicate that a disagreement exists between DSbD stakeholders about where and how to force change in the market. Stakeholders tended to see the likely pathway for implementing CHERI differently based on their role and sector and also therefore where they see the potential for market disruption. Within this general lack of consensus two sets of visions for a probable adoption pathway emerged. One vision—which was most popular with those invested from a sector agnostic high-level perspective—sees CHERI adoption as being driven by economic forces or regulatory mechanisms.

Economic forces and regulatory mechanisms both redistribute responsibilities across the DSbD ecosystem but do so in different ways. The economic position suggests that the multi-billion dollar scale of chip development and manufacture necessitates that the technology will either appear in an area with a wide-scale use such as mobile devices or a high-value area that requires security, such as the automotive sector. The regulation argument suggests that regulators will mandate requirements that will enable CHERI to be utilised by industries as a means for evidencing compliance. It should be noted that the way regulation works is not prescriptive, but rather provides a set of standards that must be adhered to

without necessarily spelling out how those standards must be met. One government stakeholder went as far as to suggest that they would deliberately push to encourage regulation as a lever for adoption if they thought it would be the best way of achieving adoption of CHERI in the market.

The other pathway to adoption sees the benefits of applying CHERI in terms of an interlocutor's specific sector or agenda. For example, an SME with a product offering in digital healthcare setting sees data privacy as the catalyst for adoption. Likewise, another participant with an interest in IoT thought that a safety critical market such as autonomous vehicles would be where CHERI first appears.

5.3 Magic Chip

Another narrative centred around DSbD's security proposition is the notion of DSbD as a "magic chip" that would provide "a panacea" to memory-targeting attacks. This narrative provoked significant dissensus amongst the DSbD insiders. One participant explained the objective of DSbD as "is to get the CHERI magic into the actual chips in the market... so my smart phone will contain the magic chip which will protect us from cyber attacks." However, this description of CHERI attracted censure from several quarters within the DSbD ecosystem. According to one participant from academia, calling the product of DSbD a "magic chip" undermines the scope of the project, which they believe is to "change the architecture for all future chips" in "everything digital." Some industry figures and SME demonstrators were sceptical about the notion that CHERI would entirely eliminate unsafe C/C++ coding practices merely by implementation in a chip. Several SME and industry participants expressed concern that training in secure coding, ensuring best practice and CHERI-specific coding knowledge would need to develop alongside the implementation of CHERI technology. One industry participant put forth that "CHERI is necessary but not sufficient," in that it must be understood in relation to other components of the security ecosystem.

5.4 Barriers to Engaging At Scale

A lack of participatory engagement and the lack of a sense of a broad DSbD developer community were cited in interviews as barriers to engaging at scale in the DSbD ecosystem. The interviews revealed several barriers to adoption that became evident in both Morello and the broader DSbD programme during our research. One SME demonstrator felt that the main DSbD stakeholders are "preaching to the already converted [...]. What you have to do is not just bring people into the tent, you also have to tell people outside the tent what's happening." Some stakeholders inside the DSbD "tent" were also acutely aware of this issue, with one stakeholder expressing "our main worry is that [all of the] enthusiastic people are within the ecosystem, but the programme doesn't really have any broader awareness in the market." One SME participant questioned how "the scope of DSbD" was being delimited, and whether its main objective was to "seed the software ecosystem... things like language run times and operating systems and libraries, just those building blocks of software and whether the programme goes beyond that" or whether it needed to include "broader industry stuff as within the scope of DSbD."

Discussion on the nature and purpose of the DSbD programme frequently focussed on the market economy of hardware innovation within the technology industry. According to one SME participant DSbD's "biggest challenge is a business challenge" rather than "the technical challenge" because they are trying to create a "fundamental shift in technology, how you architect systems... and at the end of the day it's not the best technology that wins, it's what changes the heart and minds of the market that wins." However, the same participant explained that they were "curious" about what CHERI could do and how it worked but for them a key question was "does that curiosity turn into something that means we can increase our bottom line?" Several participants also emphasised concerns over how DSbD must be made viable in relation to the large economy of scale required to make silicon manufacturing profitable. One industry insider posited that "the biggest risk is that this will not lead to pull from industry... [if] the outputs of all of these things are not actually convincing to the people who would ultimately need to say "yeah, we'll buy a million of those, or more likely 100 million of those."

At the time of writing this paper CHERI is only a prototype and collaboration. The importance of collaboration to engage at scale was present in many of the responses, knowledge exchange and knowledge co-production is necessary for DSbD to be truly transformative. One SME participant pointed out that considerable interaction was needed with the Morello project team in order to co-produce joint knowledge of how the Morello boards work. The participant went on to point out that there was much that was still unknown about how the Morello boards could be used.

5.5 Siloed Communities

Our interviews identified siloed stakeholder communities both as a direct impact of not using a participatory design approach and at the same time one of the main constraints to the adoption participatory approach to DSbD technology. Our respondents articulated a broad range of different meanings to CHERI and Morello and this reflects the breadth of the DSbD ecosystem and the complexity of the user community. Unsurprisingly, hardware technologists gave a highly technically complex explanation of what CHERI is, whereas SME participants gave a more practical definition that emphasised its ability to offer memory protection and compartmentalisation. Likewise, several SME participants defined CHERI in terms of its applicability to their business objectives. However, this did not necessarily translate to how certain technologists viewed CHERI as a more generalised mathematical and engineering problem. Such differences are not frictionless. One participant from the hardware engineer community found the question of application to be irrelevant, indicating evidence of schisms between groups in the DSbD ecosystem.

Technical knowledge is a key divider between groups within the DSbD ecosystem. A key reason identified by our participants for why CHERI has struggled to gain traction is its highly technical and complex nature. One SME demonstrator explained that "some of our clients are not super technical themselves, so they're much more interested in a broad brush of how do you solve my problem than the fine detail, and I think getting from that abstract conversation to the specifics of CHERI and why it's different from

everything else is probably not going to be productive." Several participants believed that the technical nature of CHERI required navigating between domains of both technical and non-technical expertise, as well as between domains of technical expertise. One interviewee from a government agency involved in DSbD explained "technical expertise is important here because CHERI is in the middle of a supply chain and the consumers of this product are highly technical. This requires that technologists need to understand how chips are bought by companies... Technical concepts are hard to translate to non-engineering people, and [in the past] the things that cleared those boundaries were non-technical." There was an awareness that non-technical aspects relate to how a technological innovation materialises socio-technical concepts, as was elaborated by an industry executive involved in DSbD, who acknowledged that "CHERI won't necessarily mean trust will automatically happen, but is a component in establishing trust such as in the healthcare sector or data anonymisation. Security is multi-layered through the stack and at the upper layers of this stack there are societal and social elements at work such as organisational ones."

The silos between technical communities within the DSbD ambit became very apparent in the course of our research. One of the most significant silos is the divide between hardware security specialists and software developers and coders, which manifested in an often antagonistic discourse where the coding communities were often represented as the source of security vulnerabilities in the stack. This message is re-enforced in the CHERI proposition as outlined in the DSbD grey literature promises to mitigate "the risk of programming mistakes" in C and C++ [22].

Furthermore, many of our participants who were aware that DSbD and Morello needed to circulate the CHERI proposition from the specific niche of hardware security to software practitioners, manufacturers and other technical users. The same participant quoted above made clear that "CHERI is not enough on its own – people need to know how to code to make software run on CHERI, but the question is where in the stack do you need to enable CHERI coding? Low-level talking to the chip, at the hypervisor level or at the app level?" Another industry participant working on Morello addressed the practical implications of adoption for software engineers and coders: "CHERI protects legacy software, but requires software to be recompiled. This is not a simple push-button exercise, but will have impacts, but these haven't been quantified in detail."

The importance of collaboration and shared responsibility across the silos was acknowledged by one participant who explained the requirements for adoption in terms of the supply chain: "I think the architect has to design a thing that, you know, really does have the kind of properties that they need, but then the manufacturer has to translate that into actual silicon, and there are many ways to get that wrong too... [then] the toolmakers, so people who write compilers, they have to correctly target it, operating systems have to use it correctly, there's a huge list." It was also acknowledged that this collaboration had to exist beyond the research. One SME participant felt that part of the issue was "a lack of clarity" about whether there would "be future support in terms of translating some of that research to be more integrated with the product as well" and were concerned that they "didn't get the feeling that was the route that the project was going down."

A participatory approach might be desirable to encourage engagement with CHERI and facilitate work to enable its marketisation. However, a countervailing imperative at work is that the siloing of expertise is a deliberately designed feature of the CHERI proposition due to the fact that CHERI has been designed as a solution by an elite group of hardware experts in order to negate issues that occur in software development. This means that an internal contradiction exists within DSbD over who designs what and where expertise lies in the stack. This tension between collaboration and the insular form of expertise intrinsic to hardware security design will need to be worked through and negotiated in DSbD, and this will profoundly influence the CHERI adoption story.

5.6 Desire For a More Participatory Approach

The dependency of the SME demonstrators on two specific organisations at the top of the DSbD hierarchy for the necessary software infrastructure and support necessary to make their projects workable highlights a distinct lack of organic user engagement or participation in the design process, even within a technically specialist user community that would likely be amenable to such an approach. This was picked up on by one of the SME participants who made clear "we would love to continue our engagement with DSbD... but also [we need] collaboration and introduction to the wider ecosystem and network, because there's a lot you can learn from others and there is a lot of support that is out there, that could make our journey a lot quicker."

The above quotation highlights a desire for a more inclusive design process from the technical community who would be the users of CHERI technology. DSbD instantiates a clear desire to widen engagement and participation in the CHERI innovation process, which is evident in the underlying rationale for the programme. One academic stakeholder explained that "DSbD is an opportunity to gather feedback on what needs tweaking, as it is expensive and hard to fix hardware once it is shipped." According to another industry participant "DSbD aims to explore what the software stack would look like and how to make CHERI portable to other ISAs, so that ARM, RISC-V or Intel can use the same set of standards." The Morello programme in particular has sought to create openings for participation to shift the CHERI design and adoption cycle from, as one industry participant involved with Morello put it, a "Build it and they will come mindset" to iterative design vis-à-vis demonstrators. This potential for co-design was in fact materially instantiated within the technical design of the programme. According to another industry participant, the Morello boards were designed to include an array of potential features to encourage open-ended experimentation by demonstrators, which might or might not be included in a commercial version of a CHERI-enabled Arm chip depending on what users find to be useful.

6 DISCUSSION

The intended purpose of DSbD is to facilitate a wider engagement between stakeholders in an innovation project at a supra-organisational level. However, so far it has met with mixed success in facilitating this agenda. This is partly due to a lack of development of considered strategies for the meaningful inclusion of diverse perspectives. This emanates perhaps from underlying issues

of top-down and siloed thinking being disseminated by dominant figures in the DSbD network that is often unaware of the different security ontologies of the stakeholders outside the hardware security milieu. A recurring theme throughout our findings is the siloed nature of the DSbD ecosystem as well as a desire for a more integrated approach to developing, or even co-producing, digital security by design. A conclusion from our findings might therefore be that all designers and users within the security ecosystem including hardware engineers, coders, manufacturers, government facilitators, regulators, policy makers, systems architects, supply chains, businesses across scales from SMEs to "big tech" and consumers need to develop a more holistic understanding of the needs of other users within the DSbD ecosystem. One pathway to encourage engagement across this broad range of stakeholders is to adopt a pluralist approach to defining and actualising computer security that encourages a dialogic consideration of the multiple security understandings and different security outcomes for different stakeholder groups. This is a challenge because concepts of security and responsibility are understood differently by different groups of stakeholders. The challenge is to find places through the stack where interactions between stakeholder groups might take place. In the following discussion we look at where such interactions might take place and the barriers, challenges and opportunities for such interactions. This work falls within the remit of the Discribe Hub+ to proactively investigate and evaluate questions surrounding motivations and barriers to adoption to facilitate a solution to the Market Failure problem.

6.1 Barriers to Participatory Design At The Bottom of The Stack

The evidence presented in the previous section indicates how the structure and organisation of DSbD created barriers to a participatory design of a hardware-enabled secure by design stack. These barriers need to be understood in relation to the economic scale at which DSbD strives to effect change, and the elite networks and practices within the DSbD ecosystem. Whilst software engineering and related communities are culturally complex and diverse [8] and the motivations for including security functionality are equally complex [39], our interviews indicate that those working at the base of the stack with technologies such as CHERI form a more insular community. In the case of DSbD this is caused by both the circulation of shared goals within a closed network of like-minded organisations and individuals, as demonstrated in the network diagram in Figure 1, and the practices used to actualise those goals. These practices include business events at elite spaces, talks by prominent figures within DSbD that speak at an undefined audience, rather than engaging with them, and formal networking sessions where business figures are meant to build connections with each other. However, these practices and sites often exclude meaningful engagement between technical communities and cross-fertilisation between them. They also exclude potential stakeholders who are unfamiliar or uncomfortable with these prescribed ways of collaborating.

The limitations of DSbD's attempts at engaging with a participatory design process highlights wider questions about the relationship between participatory design and economic scale. Van den Besselaar et al. (1998) [53] suggest that participatory design can work in large-scale organisational contexts and can be economically scalable as long as there is "technological and organizational flexibility" within a design programme. Whilst our research does not dismiss the possibility of participatory design working at scale out-of-hand, it does illustrate that within DSbD, there is a lack of malleability and ability to work 'outside of the box' to encourage greater participation in the design process. We posit several reasons for this. Firstly, DSbD is not merely a large-scale organisational project, but an intra-organisational project working across industry sectors and scales, as well as extensive public, private and third sector establishments. This results in differences in epistemologies, working practices, jargons and acronyms, levels and specialisms of technical expertise and areas of interest. This requires the development of deliberate strategies and adapting ways of working to bridge these gaps or promote meaningful dialogue between them. Secondly, the CHERI proposition definitionally entails creating a permanent and fixed solution etched in silicon. This intentional rigidity has an affordance that gravitates towards an inflexible design process, which is exacerbated by the sheer complexity of the technology. Pieczul et al. (2017) [46] argue that security design cannot "anticipate all uses in advance" and that one must "design to support design after design." However, CHERI's security proposition makes it difficult, if not impossible, to continue the design dialogue after the hardware is physically created, and indeed the CHERI designers have made this a deliberate choice.

6.2 The Disruption Discourse and Participatory Design

This conceptual use of disruption is perhaps meant to align DSbD with the business idea of disruptive innovation developed by Clayton Christensen and Joseph Bower [11], which has gained wide traction in the culture of the Silicon Valley technology industry. However, DSbD's use of disruptive innovation—as well as much of its use in the Silicon Valley milieu and elsewhere—differs significantly from the orthodox version posited by Christensen et al. [17] and [36]. The original concept describes "a process whereby a smaller company with fewer resources is able to successfully challenge established 'incumbent businesses' by offering a better version of a product and/or lower cost in order to gain significant 'mainstream' market share and thus disrupt the existing business paradigm in that sector" [17]. However, DSbD is in essence promoting an almost polar opposite approach, whereby major incumbent industry players including Arm, Microsoft, Google are collaborating with the UK government to facilitate the introduction of a new technology and pattern of innovation in the digital technology market.

Two possible ways to interpret the emphasis on putative disruption from the DSbD community might be as follows. One is that framing DSbD as disruptive innovation is a disingenuous marketing exercise intended to frame DSbD as a 'cool' and insurgent practice, despite it being a decidedly top-down approach, in order to boost the project's cachet and credibility. Another possible interpretation is that the disruption discourse is intended to signal the project's intent to create change in how security is embedded in hardware manufacturing. In the latter analysis, the disruption discourse can

be understood as an attempt to bring various stakeholders on-board and make them feel included in a cutting-edge transformational process.

However, rather than widening inclusion, the disruption discourse has potentially exacerbated the silos between the technical communities who must coalesce as stakeholders for CHERI to be adopted. The disruption discourse in its present form casts the practices of coders and software engineers as a problem that must be solved by hardware security, rather than enrolling the software community in developing the code libraries and software infrastructure required to make the CHERI technology viable in hardware form. Underlying this failure of inclusion is the ontological nature of what disruption does in how it is applied in the DSbD discourse. Disruption is in essence about breaking the existing order to supplant it with something new, and ostensibly better, rather than securing the ontological stability security technology is meant to uphold. However, it can also be argued that disruption in the DSbD context can be mobilised as a catalyst for an agonistic or polemical form of participatory design. Whether such an agonistic approach is necessary or viable would be an area for further research into the hardware security adoption cycle, possibly drawing evidence from past cases such as the adoption of Trusted Platform Modules (TPMs).

6.3 Metaphors and Their Roles in DSbD

Analogies and metaphors play a significant role in creating meanings and understandings of security technologies between domains and fields of practise, and this is acutely noticeable in the case of DSbD. As Wolff [61, p. 4] points out, metaphors are valuable as a means of "mapping between source domains, or concepts that we are familiar with and new, unfamiliar ideas." Slupska [49, p. 1] draws particular attention to the "generative" power of metaphors to "structure our understanding... by imposing mental models of both the problem and possible solutions" when used to "describe and understand new technologies" in cyber security policy and governance contexts. Betz and Stevens [10, p. 149] see this role of metaphors as a cause for concern as reliance on familiar themes and discourses can result in the "'dulling' of critical faculties" and make policy makers vulnerable to manipulation. This concern relates to a position first articulated by Halesz and Moran [32] that analogic reasoning, in which they include metaphors, is imprecise and oversimplifies the complexity of computing systems, models and concepts in a way that inhibits meaningful understanding of how they work. However, Demjaha et. al. [20] interrogate this perspective by using an HCI methodology to evaluate users' understandings of metaphors for a specific security technology. In doing so they draw on a distinction between 'structural' and 'functional' mental models to differentiate between metaphors that seek to convey the technical system of how a security technology works (structural) versus metaphors that work selectively to convey the functional purpose of said technology without attempting to convey specific knowledge of how it works. They found that it was challenging to convey meaningful understandings through metaphors at all, although functional metaphors caused less confusion for users than structural metaphors, and thus are a better basis for developing metaphors for technological security. Our research emphasises a somewhat different perspective on the role of metaphors in communicating technological security concepts, which should be understood in the context of the programme of innovation in which the research was conducted. Our findings suggest that metaphors do active work in creating understandings between domains, rather than passively circulating understandings between them. This productive meaning making occurs through the engaged processes of stakeholders comparing and evaluating the accuracy of metaphors, contesting their validity and negotiating upon determining the best metaphor to convey a desired meaning, such as in the discussions around the 'Magic Chip' metaphor.

6.4 Tearing Down the Silos: Morello as a boundary negotiating artefact

We suggest that a potential way forward in enabling greater inclusion in the design process of a usable CHERI component in hardware and expansion of the network of stakeholders enrolled in the DSbD endeavour is to make more effective use of Morello as a boundary negotiating artefact. This builds on Pieczul et al.'s call [46] for an "artifacts-first approach" to addressing complexity in security design, and uses DSbD and Morello as a case study to develop what such a methodological approach might look like and what work it can do. As evidenced in our findings, SME stakeholders perceived the Morello programme as having the potential to encourage participation in the design project, and to act as a basis for facilitating the growth of a user community to co-engineer implementations of the hardware-software interface required to make CHERI usable in the stack. In this, the Morello board can do work as a boundary negotiating artefact that can be circulated between specific technical and user communities to realise a shared purpose, even when the understanding of what the object is and how it works is understood in different ways by each stakeholder.

Using Morello as a medium to coalesce the diverse interests involved in DSbD has the potential to innovate change in how the information stack is conceptualised and designed, and how security is considered and incorporated within the design of the stack. More specifically, this object-centred approach invites a more holistic consideration of how the stack should work as a whole rather than as layers of discrete components laid on top of each other. This includes consideration in design of the interaction between layers and how information flows through the stack. Furthermore, CHERI can facilitate the design of different stacks for different use cases and users, for example a stack that emphasises real-time control of creative content in television virtual production environments. Therefore, designers of CHERI-enabled stacks must ask and evaluate what does the stack look like when CHERI is included and how does this alter the security principles and rituals embedded in the stack?

Using Morello as a boundary negotiating artefact must also be interrogated in relation to the heterogeneous design practices in technical communities, businesses and academic research, as well as the economic scale at which DSbD is pitched, in contrast to the majority of participatory design and co-design processes. Put simply, is it possible to scale up the "hack-a-thon" design approach and rebrand it in large business-friendly terms, or does scaling up

design to fit the economic necessities of hardware security innovation require entirely new sets of design philosophies and practices? Further investigation and analysis of DSbD as the programme progresses can hopefully be an empirical resource for evaluating this question, building on existing interventions from the body of participatory design and HCI literatures as set out in our background literature section, and specifying what participatory design might or ought to look like in the particular exigencies of the hardware security field.

6.5 Responsibility Versus Liability

Concerns of how responsibility is framed and delegated are implicit in questions of who and what designs, and how design methodologies are executed. By taking ownership in the design process a given stakeholder also takes responsibility for how the outcomes of the process will play out and what social principles are made material in the design of that technology.

At present there are several potential pathways for how the CHERI proposition delegates responsibility that are being negotiated amongst the CHERI community. On the whole it places current responsibility for failure, i.e. exploits of memory vulnerabilities, solely onto coders, software developers, system administrators and end users for ignoring memory-safe coding practices or not understanding specific security policies. This encourages a process of circulating responsibility whereby blame for failure is passed up the stack eventually placing the final burden of liability on end users for not adequately taking ownership of their own security [3]. The "magic chip" claim proposes shifting the delegation of responsibility for data security onto the technology itself as a safer option, thus absolving stakeholders in the security ecosystem of responsibility for how data is handled and secured in the CHERI stack. This is inimical to co-production as it disincentivises stakeholders from taking responsibility for establishing what good security looks like, and ownership for the consequences if the technology were to fail.

An alternative pathway being advocated by some of the key players in CHERI, and by some of the SME stakeholders at the periphery of the DSbD community, is the "necessary but not sufficient" proposition. This suggests that security in terms of who controls data must be delegated across the security ecosystem and must include designers and users. This entails treating CHERI as a redundancy measure that backs up greater knowledge of how memory safety works at the silicon level and safe C/C++ coding practices. Alongside this, CHERI technology must be manufactured and sourced in a way that is transparent and alongside other measures that can mitigate supply chain vulnerabilities. Whilst this is a more sophisticated and nuanced adoption route than the "magic chip" hyperbole, it is much harder to articulate convincingly to different groups of stakeholders, particularly those who are non-technical.

6.6 Extending The Participatory Design Paradigm

Our findings reflect that within DSbD there is an appetite for a more participatory approach from stakeholders engaged with security concerns located in the middle and upper layers of the stack, to which key players in DSbD, whose focus is on the hardware layer of the stack, are oblivious. Furthermore, our interviews also reflect the view that for technologies such as CHERI to be commercially viable it requires synergistic thinking and collaboration between existing and potential stakeholders across multiple domains of technical, business and policy expertise. For technologies to be secure, security must be incorporated through the stack and not just at a particular layer, and security design must be attuned to this fact. Embedding of security principles in the hardware level of the stack must therefore articulate technical security principles in a way that formalises the socio-technical security needs and concerns of a plurality of users in order to offer greater resilience to a broad range of threats. Incorporating a participatory design process in how hardware security is designed and brought into relation with the other layers of the stack offers a pathway to actualise such a pluralistic conceptualisation of security.

Placing CHERI within a participatory design paradigm is not straightforward because CHERI exceeds the intra-organisational scale and focus of earlier participatory design approaches. Likewise, the subaltern critique and framing of security in more recent participatory design and computer security work (e.g. [50]) might seem somewhat incongruous with CHERI's security proposition of embedding more secure access control at the base of the stack. And yet, the embedding of more secure access control at the base of the stack might free end users from some of the security responsibilities that have been pushed upon them [47], creating capacity to more easily co-produce secure ways of living and working. For such an integrative approach to security to exist, participatory approaches to connecting both ends of the stack need to be found.

Our findings place a spotlight on the importance of the Morello board as a potentially powerful site of participation. As CHERI is a theoretical model, it is more conceptually useful to evaluate the meaning-making around it through its instantiation in the Morello board [59] as a 'boundary artifact' through which pluralistic understandings of security are negotiated, articulated and co-produced. Therefore participatory design in the CHERI context could be "up-scaled" to the expert level and incorporate multiple forms of technical expertise in a way that would differ drastically from participatory design as a counter-hegemonic strategy. Instead, such an up-scaling enables a more equitable re-distribution of power and responsibility between technologists, the technology industry and end users and results in a stronger, more resilient, security proposition.

The participatory studies found in HCI often do not equate to a design environment such as the one required for CHERI. The power and social justice debates found in the canon of participatory studies in HCI literature do not map to the CHERI debate. Here it is actually non-hardware security technical experts who are being marginalised e.g coders, network architects. Therefore the notion of who ought to be considered a participatory user needs to be adjusted or redefined. Given the nature of hardware security in general and CHERI in particular, participatory design in this context is bound to be less concerned with what the user wants and more with what is possible in hardware. This is an important distinction that is reflected in the engineering goals, the methods used and the style of communication. Consequently, successful co-design and co-development must be rooted in a participatory approach based on meaningful joint activities, e.g., joint coding sessions possibly using simulated hardware prototypes, and targeted exploration of ways to

make well-crafted documentation that is relevant, useful, and easily accessible for developers. This includes developing alternatives to traditional documentation, e.g. ways to co-produce community resources such as Stack Overflow.

The above insights offer a start point for rethinking participatory design at the base of the stack. Looking up through the stack, it becomes clear that the scale and magnitude of complexity in negotiating the diverse positions and knowledges of stakeholders across organisations, fields including hardware and software engineering, business strategy and academic computer science, industries and sectors from the automotive industry to consumer devices to chip manufacturers to government regulators, will require a qualitatively different approach to participatory design that must be evinced further.

When considering the insights offered in this paper, it is important to keep in mind that CHERI will not be a stand-alone technology, but will be integrated as a component within other technologies. Since it will be located upstream in the supply chain, its potentially disruptive effects would likely create an iterative cycle of second- and third-order effects cascading down the supply chain. Therefore, rather than treating each layer of the CHERI-enabled stack as discrete, it will be necessary to account for and facilitate interactivity between the languages and operations up and down the stack. Locating sites for interactivity and how to facilitate the necessary exchange between domains of expertise is the specific purpose of this research and the Discribe Hub+ overall.

6.7 Limitations

Methodological limitations to this research include the small sample size of research participants. This was primarily due to the limited size of the DSbD network at the start of our research, as well as challenges in obtaining replies to our participation requests from several of the *de minimis* project SMEs that formed the sample group for our SME research. Based on our more developed understanding of the security policy landscape, we intend as the next phase of our research to conduct interviews with both digital security-facing regulators, as well as company executives in Chief Information Officer(CIO)/Chief Information Security Officer (CISO) or similar roles. Another source of limitations arose from the temporal position of our research at an early stage in the DSbD programme in which the innovation process was in an immature state. This meant we lacked some clarity on who we should recruit for interviews outside of the core DSbD programme stakeholders as it was unclear who would be relevant decisionmakers for the adoption and commercialisation of CHERI in the wider security ecosystem. This is due to both the disruptive intention of the DSbD programme to fundamentally shift the locus of memory security to CHERI-based hardware platforms, as well as the fact that power relations in the innovation process—as with any social process—will transform over time. Because the present research is looking at the early potential socio-technical configurations and pathways for adoption, we recommend that a longitudinal study be conducted towards the end of the innovation process. This can provide a comparative basis for analysis of how the actualisation or failure of CHERI occurred, and how meanings and understandings of the technology and its context evolved over the course of the DSbD programme.

7 CONCLUSION

In this paper we have examined potential reasons for a lack of participatory design lower down the software stack, and presented a study that illustrates the potential implications of an absence of participatory design in hardware security innovation. Building on this, we have proposed how participatory design might be re-conceptualised to overcome some of the challenges of carrying out an ambitious project of transformation in technological security. Contemporary hardware and software engineering practices, the emergence of virtualised development environments and the shift towards a society that is digital by design make it increasingly important that the approach to security design is integrative. By connecting the design of security at the base of the stack to the design of security in the layers above we are better able to produce a form of digital security that is resilient to a wider range of attacks and to create capacity for a greater focus on co-producing secure ways of living and working. We propose a reconfiguration of participatory design at an economic and supra-organisational scale that could contribute to the realisation of this positive outcome.

ACKNOWLEDGMENTS

We would like to thank our participants for the time and effort they spent engaging with us. Contributions from Slesinger, Coles-Kemp, and Panteli were funded by ESRC, grant number: ES/V003666/1.

REFERENCES

[1] Anne Adams and Martina Angela Sasse. 1999. Users are not the enemy. *Commun. ACM* 42, 12 (1999), 40–46.
[2] Jonathan P Allen. 1992. Enabling Participatory Design in a Hierarchical, Tightly Integrated Setting. In *PDC*. 73–79.
[3] Ross Anderson and Tyler Moore. 2009. Information security: where computer science, economics and psychology meet. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 367, 1898 (2009), 2717–2727.
[4] Debi Ashenden and Gail Ollis. 2020. Putting the sec in devsecops: Using social practice theory to improve secure software development. In *New Security Paradigms Workshop 2020*. 34–44.
[5] Bartłomiej Balcerzak, Wiesław Kopeć, Radosław Nielek, Sebastian Kruk, Kamil Warpechowski, Mateusz Wasik, and Marek Węgrzyn. 2017. Press F1 for help: participatory design for dealing with on-line and real life security of older adults. In *2017 12th International Scientific and Technical Conference on Computer Sciences and Information Technologies (CSIT)*, Vol. 1. 240–243.
[6] Ben Bauer and Andrew S Patrick. 2004. A human factors extension to the seven-layer OSI reference model. *Retrieved January* 6 (2004), 2004.
[7] Christoph Becker, Ann Light, Chris Frauenberger, Dawn Walker, Victoria Palacin, Syed Ishtiaque Ahmed, Rachel Charlotte Smith, Pedro Reynolds Cuéllar, and David Nemer. 2020. Computing professionals for social responsibility: The past, present and future values of participatory design. In *Proceedings of the 16th Participatory Design Conference 2020-Participation (s) Otherwise-Volume 2*. 181–184.
[8] Sarah Beecham, Nathan Baddoo, Tracy Hall, Hugh Robinson, and Helen Sharp. 2008. Motivation in Software Engineering: A systematic literature review. *Information and software technology* 50, 9-10 (2008), 860–878.
[9] Martin Beirne and Harvie Ramsay. 1992. *Information technology and workplace democracy*. CUP Archive.
[10] David J Betz and Tim Stevens. 2013. Analogical reasoning and cyber security. *Security Dialogue* 44, 2 (2013), 147–164.
[11] Joseph L Bower and Clayton M Christensen. 1996. Disruptive technologies: Catching the wave. *The Journal of Product Innovation Management* 1, 13 (1996), 75–76.
[12] Alex Bowyer, Kyle Montague, Stuart Wheater, Ruth McGovern, Raghu Lingam, and Madeline Balaam. 2018. Understanding the family perspective on the storage,

sharing and handling of family civic data. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
[13] Eva Brandt, Thomas Binder, and Elizabeth BN Sanders. 2012. Ways to engage telling, making and enacting. *Routledge international handbook of participatory design. Routledge, New York* (2012), 145–181.
[14] Oliver K Burmeister. 2009. What users of virtual social networks value about social interaction: The case of GreyPath. In *Virtual social networks*. Springer, 114–133.
[15] Paul R Carlile. 2002. A pragmatic view of knowledge and boundaries: Boundary objects in new product development. *Organization science* 13, 4 (2002), 442–455.
[16] Chhaya Chouhan, Christy M LaPerriere, Zaina Aljallad, Jess Kropczynski, Heather Lipford, and Pamela J Wisniewski. 2019. Co-designing for community oversight: Helping people make privacy and security decisions together. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–31.
[17] Clayton M. Christensen, Michael E. Raynor, and Rory McDonald. 2015. What is Disruptive Innovation? *Harvard Business Review* (Dec. 2015). Retrieved May 25, 2022 from https://hbr.org/2015/12/what-is-disruptive-innovation
[18] Chia-Fang Chung, Kristin Dew, Allison Cole, Jasmine Zia, James Fogarty, Julie A Kientz, and Sean A Munson. 2016. Boundary negotiating artifacts in personal informatics: patient-provider collaboration with patient-generated data. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*. 770–786.
[19] Helen Collard and Jo Briggs. 2020. Creative Toolkits for TIPS. In *European Symposium on Research in Computer Security*. Springer, 39–55.
[20] Albese Demjaha, Jonathan M Spring, Ingolf Becker, Simon Parkin, and M Angela Sasse. 2018. Metaphors considered harmful? An exploratory study of the effectiveness of functional metaphors for end-to-end encryption. In *Proc. USEC*, Vol. 2018. Internet Society.
[21] Carl DiSalvo. 2014. Critical making as materializing the politics of design. *The Information Society* 30, 2 (2014), 96–105.
[22] DSbD. [n. d.]. *How it works*. Retrieved May 25, 2022 from https://www.dsbd.tech/how-it-works/
[23] Paul Dunphy, Andrew Monk, John Vines, Mark Blythe, and Patrick Olivier. 2014. Designing for spontaneous and secure delegation in digital payments. *Interacting with Computers* 26, 5 (2014), 417–432.
[24] Paul Dunphy, John Vines, Lizzie Coles-Kemp, Rachel Clarke, Vasilis Vlachokyriakos, Peter Wright, John McCarthy, and Patrick Olivier. 2014. Understanding the experience-centeredness of privacy and security technologies. In *Proceedings of the 2014 New Security Paradigms Workshop*. 83–94.
[25] Pelle Ehn. 2008. Participation in design things. In *Participatory Design Conference (PDC), Bloomington, Indiana, USA (2008)*. 92–101.
[26] Frank Emspak. 1996. Participatory design: examples and institutional Needs. In *PDC*. 111–115.
[27] Matthias Fassl, Lea Theresa Gröber, and Katharina Krombholz. 2021. Exploring User-Centered Security Design for Usable Authentication Ceremonies. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
[28] Kevin Gallagher, Santiago Torres-Arias, Nasir Memon, and Jessica Feldman. 2021. COLBAC: Shifting Cybersecurity from Hierarchical to Horizontal Designs. In *New Security Paradigms Workshop*. 13–27.
[29] Joan Greenbaum and Morten Kyng. 2020. *Design at work: Cooperative design of computer systems*. CRC Press.
[30] Catrin Griffiths, Niki Panteli, Deirdre Brunton, Ben Marder, and Heidi Williamson. 2015. Designing and evaluating the acceptability of Realshare: an online support community for teenagers and young adults with cancer. *Journal of health psychology* 20, 12 (2015), 1589–1601.
[31] Jonathan Grudin. 1990. Obstacles to participatory design in large product development organizations. In *PDC*. 14–21.
[32] Frank Halasz and Thomas P. Moran. 1982. Analogy Considered Harmful. In *Proceedings of the 1982 Conference on Human Factors in Computing Systems* (Gaithersburg, Maryland, USA) *(CHI '82)*. Association for Computing Machinery, New York, NY, USA, 383–386.
[33] Richard Harper, Siân Lindley, Eno Thereska, Richard Banks, Philip Gosset, Gavin Smyth, William Odom, and Eryn Whitworth. 2013. What is a File?. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 1125–1136.
[34] Per-Anders Hillgren, Anna Seravalli, and Mette Agger Eriksen. 2016. Counter-hegemonic practices; dynamic interplay between agonism, commoning and strategic design. *Strategic Design Research Journal*, 9 (2): 89-99 May-August 2016 (2016).
[35] Finn Kensing, Jesper Simonsen, and Keld Bodker. 1998. MUST: A method for participatory design. *Human-computer interaction* 13, 2 (1998), 167–198.
[36] Andrew A King and Baljir Baatartogtokh. 2015. How useful is the theory of disruptive innovation? *MIT Sloan management review* 57, 1 (2015), 77.
[37] Kari Kuutti. 2013. 'Practice turn' and CSCW identity. *ECSCW 2013 Adjunct Proceedings* (2013), 39–44.
[38] Charlotte P Lee. 2007. Boundary negotiating artifacts: Unbinding the routine of boundary objects and embracing chaos in collaborative work. *Computer Supported Cooperative Work (CSCW)* 16, 3 (2007), 307–339.
[39] Tamara Lopez, Helen Sharp, Thein Tun, Arosha Bandara, Mark Levine, and Bashar Nuseibeh. 2019. Talking about security with professional developers. In *2019 IEEE/ACM Joint 7th International Workshop on Conducting Empirical Studies in Industry (CESI) and 6th International Workshop on Software Engineering Research and Industrial Practice (SER&IP)*. IEEE, 34–40.
[40] Preben Holst Mogensen and Randall H Trigg. 1992. Using artifacts as triggers for participatory analysis. *DAIMI Report Series* 413 (1992).
[41] Michael J Muller and Allison Druin. 2012. Participatory design: the third space in human–computer interaction. In *The Human–Computer Interaction Handbook*. CRC Press, 1125–1153.
[42] Michael J Muller and Sarah Kuhn. 1993. Participatory design. *Commun. ACM* 36, 6 (1993), 24–28.
[43] NCSC. 2019. *You shape security*. Retrieved May 25, 2022 from https://www.ncsc.gov.uk/collection/you-shape-security
[44] Patrick Olivier and Peter Wright. 2015. Digital civics: Taking a local turn. *Interactions* 22, 4 (2015), 61–63.
[45] Georgios Papadakis. [n. d.]. *Unleash of the release – Enabling the Digital Security by Design (DSbD) ecosystem*. Retrieved May 25, 2022 from https://www.ncsc.gov.uk/collection/you-shape-security
[46] Olgierd Pieczul, Simon Foley, and Mary Ellen Zurko. 2017. Developer-centered Security and the Symmetry of Ignorance. In *Proceedings of the 2017 New Security Paradigms Workshop*. 46–56.
[47] Karen Renaud, Stephen Flowerday, Merrill Warkentin, Paul Cockshott, and Craig Orgeron. 2018. Is the responsibilization of the cyber security risk reasonable and judicious? *Computers & Security* 78 (2018), 198–211.
[48] Jerome H Saltzer and Michael D Schroeder. 1975. The protection of information in computer systems. *Proc. IEEE* 63, 9 (1975), 1278–1308.
[49] Julia Slupska. 2021. War, health and ecosystem: generative metaphors in cybersecurity governance. *Philosophy & Technology* 34, 3 (2021), 463–482.
[50] Julia Slupska, Scarlet Dawson Dawson Duckworth, Linda Ma, and Gina Neff. 2021. Participatory threat modelling: Exploring paths to reconfigure cybersecurity. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–6.
[51] Matt Spencer. 2019. The difference a method makes: methods as epistemic objects in computational science. *Distinktion: Journal of Social Theory* 20, 3 (2019), 313–327.
[52] Susan Leigh Star. 1989. The structure of ill-structured solutions: Boundary objects and heterogeneous distributed problem solving. In *Distributed artificial intelligence*. Elsevier, 37–54.
[53] Peter van den Besselaar. 1998. Democracy & technological change: Limits to steering. In *Proceedings of the Participatory Design Conference, Chatfield, R., Kuhn, S. & Muller, M.(Eds.), Seattle USA*.
[54] Maja Van der Velden, C Mörtberg, J Van den Hoven, PE Vermaas, and I Van de Poel. 2014. Participatory design and design for values. *Development* 11, 3 (2014), 215–236.
[55] Aditya Vashistha, Richard Anderson, and Shrirang Mare. 2018. Examining security and privacy research in developing regions. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*. 1–14.
[56] John Vines, Rachel Clarke, Peter Wright, John McCarthy, and Patrick Olivier. 2013. Configuring participation: on how we involve people in design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 429–438.
[57] Vasillis Vlachokyriakos, Clara Crivellaro, Christopher A Le Dantec, Eric Gordon, Pete Wright, and Patrick Olivier. 2016. Digital civics: Citizen empowerment with and through technology. In *Proceedings of the 2016 CHI conference extended abstracts on human factors in computing systems*. 1096–1099.
[58] Yang Wang. 2017. The third wave? Inclusive privacy and security. In *Proceedings of the 2017 New Security Paradigms Workshop*. 122–130.
[59] Robert NM Watson. 2019. *UKRI Digital Security by Design: A £190M research programme around Arm's Morello – an experimental ARMv8-A CPU, SoC, and board with CHERI support*. Retrieved May 25, 2022 from https://www.lightbluetouchpaper.org/2019/10/18/ukri-digital-security-by-design-a-190m-research-programme-around-arms-morello-an-experimental-armv8-a-cpu-soc-and-board-with-cheri-support/
[60] Robert N. M. Watson, Peter G. Neumann, Jonathan Woodruff, Michael Roe, Hesham Almatary, Jonathan Anderson, John Baldwin, Graeme Barnes, David Chisnall, Jessica Clarke, Brooks Davis, Lee Eisen, Nathaniel Wesley Filardo, Richard Grisenthwaite, Alexandre Joannou, Ben Laurie, A. Theodore Markettos, Simon W. Moore, Steven J. Murdoch, Kyndylan Nienhuis, Robert Norton, Alexander Richardson, Peter Rugg, Peter Sewell, Stacey Son, and Hongyan Xia. 2020. *Capability Hardware Enhanced RISC Instructions: CHERI Instruction-Set Architecture (Version 8)*. Technical Report 951. University of Cambridge, Computer Laboratory. Retrieved May 25, 2022 from https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-951.pdf
[61] Josephine Wolff. 2014. Cybersecurity as metaphor: policy and defense implications of computer security metaphors. In *2014 TPRC Conference Paper*.
[62] Mary Ellen Zurko, Rich Simon, and Tom Sanfilippo. 1999. A user-centered, modular authorization service built on an RBAC foundation. In *Proceedings of the 1999 IEEE Symposium on Security and Privacy (Cat. No. 99CB36344)*. 57–71.

A RESEARCH TOPIC GUIDE: SME PROJECT BASELINE INTERVIEWS

After introductions, carrying out informed consent protocol, and the run through of the study information sheet, the following questions were asked:

Contextual information

- Ask the participant to explain their SMEs work/product
  - Sector/market niche
  - History of organisation
  - Number of employees
- Ask individual to describe their role in the SME organisation, and their position in relation to the project
- Ask to describe what their DSbD project was, and which people from the organisation were involved in the project/what were their roles
  - How did they get involved in the DSbD challenge, and what commitments did they need to make

Participant's baseline understandings

- Ask for their understanding of what is DSbD
- What is the purpose of DSbD
- How does Morello relate to DSbD
- Has their understanding of DSbD changed from before they started the project to after, and how
- Probe any interesting or unexpected insights from their explanation
- Seek clarification about relevant terminology and concepts
- Does participant think Morello will be useful in the long run, and why/how
- Who are the internal and external stakeholders engaged with the project, and who will be the direct users of the technology being developed through Morello
- What is the likely first application of CHERI technology? Ask who would be sufficiently motivated to bring DSbD hardware to implementation and what would be their motivation.

Participant's experience with the CHERI SME project

- Identify participant's reasons/motivations for getting involved with the DSbD challenge.
- Looking back at the end of the project, what were the benefits of participation; Did they face any challenges or drawbacks
- What did participant expect DSbD technology to do
- Were their expectations of the technology met
- Does the participant hope or expect to continue to be involved with the next stage of the programme, if so in what ways, and what resources would this require

5. Understandings of threats, risks and regulatory implications

- What responsibilities, if any does participant believe their organisation has in implementing and using DSbD technologies
- Were any factors of responsibility, failure or risk considered in designing the participant's project
- What would a failure of the DSbD chip look like, and what would happen in such a scenario
  - Who would be responsible in such a scenario. Probe further as to when and why.
- Does the participant foresee any likely vectors of attack if CHERI-enabled technology were widely adopted.
- What is the worst consequence that could result if a CHERI-enabled chip were hacked, and who would be likely worst affected by it
- How would you like DSbD to be regulated? As an organisation would you have involvement in the regulation of this technology

B RESEARCH TOPIC GUIDE: BASELINE INTERVIEWS WITH KEY STAKEHOLDERS

After introductions, carrying out informed consent protocol, and the run through of the study information sheet, the following questions were asked:

Contextual information

- Ask how participant is connected to the Morello project
- Discuss how/why they became involved with Morello and/or CHERI

Participant's definition of Morello/CHERI/DSBD

- Discuss what is Morello from participant's viewpoint
  - What is Morello
  - What is its purpose
  - Probe any interesting or unexpected insights from participant's explanation
  - Seek clarification about relevant terminology and concepts
- Ask questions on how do Morello, CHERI and DSbD relate to one another and how are they distinct from another
- How does participant think Morello is useful (can be general and/or specific to their professional domain)
- Who are the external stakeholders engaged with the project, and who will be the direct users of the technology being developed through Morello
- What is the likely first application of CHERI technology? Who would be sufficiently motivated to bring DSbD hardware to implementation and what would be their motivation.
- Ask if their understanding of what Morello and/or CHERI is, or what it/they can do changed since becoming involved in the project