An Anthropological Approach to Studying CSIRTs

Sathya Chandran Sundaramurthy, John McHugh, Xinming Ou
S. Raj Rajagopalan, and Michael Wesch

Abstract

Cultural and Social Anthropology is well known for its research method — participant observation. We have adopted this methodology in conducting ethnographic fieldwork in a CSIRT. We explain the notion of "tacit knowledge" and its power in understanding CSIRT operations. We also talk about the unique advantages participant observation brings to this research, compared with other approaches such as pure observation and interviews. We then give an account from "ground zero" of how we started the ethnographic study — how we built relationship with the CSIRT, how the "research" started and what obstacles were met, how we learned from the obstacles and produced breakthroughs in moving the research effort forward, and what insights we have gained so far. We close by re-emphasizing that researchers, by putting their boots on the ground, can make significant progress in understanding CSIRTs challenges and helping address them.

Keywords

C	Computer Systems Organization
C.2	Communication/Networking and Information Technology
C.2.0	General
C.2.0.f	Network-level security and protection
K	Computing Milieux
K.6	Management of Computing and Information Systems
K.6.m	Miscellaneous
K.6.m.b	Security

Sathya Sundaramurthy eyed his screen with a paradoxically detached intensity as he scanned through a small bit of the over 70GB of log data accumulated that day. As an analyst for the university Security Operations Center (SOC), his job was to find a simple data match that could move their latest investigation forward. This was no glamorous or exciting game of catching the hacker. Instead, it was rather simple, yet intense, tedium. He had to be alert and move fast to close the ticket quickly but this particular part of the job - scanning data and numbers - did not exactly engage his capacities for complex thought and analysis. He calculated that he was in the midst of a series of ongoing five-minute cycles. Receive an alert, scan the logs (three minutes), look up an address (one minute), find the user information (another minute), repeat. Despondently, he calculated that this would be the 47th cycle this week. Such is the life of an analyst.

Sundaramurthy is no analyst. He is part of a cybersecurity research team that has spent years trying to understand the workings of SOCs using more traditional methods in order to build tools they hoped would help the analysts. They never seemed to get the data, knowledge, and insight they

needed and it was difficult to convince the analysts to try out their tools in operations. Now they had added an anthropologist to the team to advise them on the use of ethnographic methods and participant observation. In his first month at the SOC, Sundaramurthy was starting to understand why the old methods had not worked. Much of what he learned in that first month was not easy to talk about. It was tacit, embodied knowledge. Much of it would not seem interesting or relevant enough to mention in an interview and the environment, battling a flurry of unseen adversaries in a basement room, did not exactly foster trust. He thought "Why should I believe that the interviewer is not trying to obtain some key information that would allow him to hack the system? Who has time for an interview anyway? More tickets are flowing in as we speak. I have to close tickets." By now Sundaramurthy understood that analysts have little time for interviews with researchers, are unlikely to trust them, and even if they do make the time and the leap of faith to trust, they might not know what to say. Much of what they know is not easily put into words because it is unconscious, inappropriate to discuss, apparently unimportant, or irrelevant.

Forward a year. Sundaramurthy has spent nearly 500 hours working as an analyst. The challenges of previous research methodologies did not simply go away. It took many hours and a great deal of effort to move past the barriers - to earn the trust of the other analysts and upper management, to reflect on the knowledge that became embodied through the fieldwork process and make it more explicit. This allowed him to analyze the broader context of policies, social relations, and bureaucratic requirements that continually change and reshape the context in which the work takes place.

1 Anthropology Today

Normally we imagine cultural anthropologists working at the far edges of mainsteam civilization, studying peoples still isolated from the transformations of technology in order to understand human behavioral characterisitics. Over the past several decades, more and more cultural anthropologists have been using their research methods to explore familiar territories, often with surprisingly practical and productive results. In the late 1990s, Leinbach and Sears brought the methods of anthropology to the design of Recreation Vehicles. Spending six months on the road living in a giant RV staying in RV campsites, they learned that the design of the RV was mismatched to the needs and desires of everyday life within the RV culture; *e.g.*, most RV campers never use the in-vehicle shower (they prefer the high-pressure campground showers) using the shower space as an extra closet. RV designers imagined what RV'ers would need – but never lived with them. Leinbach and Sears built a prototype embodying their findings that was so successful the manufacturer dropped all other models to meet the demand. [9]

There are many other success stories. Genevieve Bell, a cultural anthropologist at Intel explored the social and cultural aspects of ubiquitous computing with Paul Dourish [6], significantly shaping ubiquitous computing research methodologies. Go-Gurt, the handheld yogurt that children can enjoy like a popsicle while appeasing moms' interest in nutrition, was based on an anthropological study by Susan Squires [12]. It is widely, and wrongly, believed that the ubiquitous big green button on Xerox machines was invented by an anthropologist. In 1979 anthropologist Lucy Suchman filmed user interactions with a printer Xerox had just been put on the market; widely disparaged by users as too complicated. Suchman's footage became famous for the meme-like image, "When User Hits Machine," that shows PARC researchers Allen Newell and Ron Kaplan trying to use the machine. The machine already had the big green button on it but when users pushed it the machine would

not do what they wanted it to. From this, Suchman made the point that no machine or tool is self-explanatory, "We need time to make unfamiliar devices our familiars."

As we make a tool our familiar we often change its use and function from that originally intended. In this way, tool users become, in part, their authors as well. While traditional research methods such as interviews and surveys might elicit somebody's explicit understandings, needs, and desires, only participant observation can help us understand the embedding context, underlying assumptions, and tacit practices that shape how tools might actually be used, adopted, and adapted in different work contexts.

As the ethnographic methods developed and honed by anthropologists in the most remote areas of the world have been adapted and used in more practical applications, "ethnography" has come to designate a wide variety of research practices. By the mid 1990s, cultural anthropologist Michael Agar expressed a growing concern among his fellow anthropologists that there was now a "dangerous delusion" that a few hours with a focus group could pose as "ethnography" in some circles [1]. Nonetheless, even mild forms of ethnography were having profound effects: allowing people to see and understand new perspectives and gain new insights. Instead of drawing a line in the sand and denying short-term fieldwork and focus groups the status of "ethnography," Agar drew a broad spectrum of "ethnographic" possibilities from focus groups and workplace observations on one end to long-term intensive participant observation where researchers join the community of their research subjects for months or years, on the other. While many domains of computer science now use ethnographic techniques to great effect, there are few examples of intensive long-term participant observation, our methodology for studying Computer Security Incident Response Teams, or CSIRTs.

2 Bringing Anthropology to the Study of CSIRTs

The traditional approach taken by the academic cybersecurity researchers is to identify problems and areas for improvement by studying the research literature, then developing tools and methodologies to address those problems. This process often results in "solutions," that are not usable in the real world. We believe the problem results from a discrepancy between what the security practitioners actually need and the researchers' perception of that need. As a result, research results rarely find their way into practical use.

Researchers in similar domains have observed comparable problems and adopted ethnography as a key component of their research. The ongoing challenge is how to apply ethnographic methods and close the gaps between researcher and participant, designer and user, and researcher and designer. This continues to enrich the methodological toolkits available to Human-Computer Interaction (HCI), Participatory Design (PD), and Computer Supported Cooperative Work (CSCW) communities, among others. Conceptual and theoretical toolkits continue to grow as well. As Crabtree and Rodden have pointed out [3], new concepts emerging from ethnographic work alert and sensitize us to issues and insights that otherwise might remain unnoticed. While many researchers using ethnography in computer science have moved toward short-termed techniques, tightening their research scope and using multiple methods to match increasingly short production cycles, we find some unique challenges in studying CSIRTS which require us to return to a more traditional form of anthropological fieldwork — long-term participant observation.

2.1 The Biggest Challenge: Gaining Trust

SOC culture cannot be easily understood from the outside. The main challenge facing a researcher undertaking ethnographic fieldwork in a SOC is earning trust. There are a number of reasons why SOC employees are unlikely to trust a researcher. Due to the sensitive environment, sharing information with outsiders is discouraged or prohibited. A SOC is a frenzied and stressful workplace with an intensive workload. Employees are evaluated by how many tickets / incidents they close per day. Analysts do not have the time or desire to share information / knowledge among themselves, much less with outsiders. The "tribal culture" of SOCs dictates that newcomers learn the routine on their own. The greatest obstacle facing an ethnographer doing research in a SOC is lack of trust from the analysts. Trust is not part of the culture; why should they trust the researchers? The ethnographer's presence raises other concerns: absorbing analysts' time without clear benefits, compromise of sensitive data, steep learning curves, short-term presence, and inability to fit the tribal culture.

These were exactly the difficulties experienced by our fieldworkers when they started work. The authors' past experiences designing cybersecurity tools and techniques for SOC analysts showed similar problems. Over the years, we have heard similar anecdotal stories from many sources. There is a significant gap between researcher and practitioner cultures.

2.1.1 Bridging the Chasm

The authors have been involved with both the academic and operational security communities for several decades. Our observations of and participation in both communities has led us to a number of conclusions that shed light on the problems that the current work tries to address. With few exceptions (*e.g.*, Lawrence Berkeley Laboratory), academic researchers in cybersecurity have little contact with practitioners. They work under different conditions, attend different conferences, and have developed a mutual feeling of distance from one another. Academics are accustomed to creating, contesting, and sharing their knowledge and expertise, while practitioners have long worked on their own and created their own expertise without sharing detailed expositions of that knowledge. This is because (i) the culture of cybersecurity practice is closed. One is not encouraged to talk about attacks and how they were found. (ii) there is a mismatch between the kind of knowledge possessed by practitioners and the typical content of research papers. Practitioners almost never get papers published in academic conferences, (iii) At popular "pseudo-practitioner" forums (DefCon, Black Hat, *etc.*) the focus is more on the "wow" factor or "coolness" and less on scientific merit. (iv) Many real practitioner forums are held in classified settings.

Due to the tacit nature of the practitioners' knowledge, fostering conversations between the two camps does not necessarily transfer knowledge. As Michael Polanyi noted [11], "We can know more than we can tell." Cybersecurity practitioners often work from "hunches" or "intuitions." They know what to do, where to look, and how to investigate a case, but often cannot state this knowledge explicitly. SOC jobs such as incident response and forensic analysis have become so sophisticated and expertise-driven that it is nearly impossible to understand the processes without doing the job. Few researchers do "fieldwork" to understand the problem first-hand before attempting to design models and algorithms that purportedly solve the problem.

Researchers perceive SOCs researchers as shrouded in secrecy and find it hard to get research data or to deploy and evaluate research prototypes. SOC analysts are guarded when approached by researchers, who are often perceived to be more interested in publications than solving real prob-

lems. Most academic research gets little credibility with and has little impact on the practitioner community. Studies conducted by researchers tend to be plagued by lack of scientific rigor (*e.g.*, the 2010 NSPW and 2011 ACSAC LASER panels, LASER conferences) and are largely irrelevant or divorced from reality (see Felix Lindner's keynote, "On Hackers and Academia" at EC2ND 2010). The result is infrequent technology transfer from the academic research to security practitioners.

Long-term Participant Observation is the Key The dynamics of research change completely when researchers are willing to join the analysts' community in a long-term participant observation effort. The researchers learn the job themselves, earn trust from the analysts, and gain insights into the SOC's operations. They reflect on their observations to identify real needs and can then design tools readily accepted by the SOC community, because the researcher is now part of it and the tool is embedded directly in its day-to-day operations. This is not unlike the "old-day" anthropological work where anthropologists join a remote indigenous community to understand its culture (as Co-author Wesch did for 18 months, studying the effects of writing on an indigenous culture in a Papua New Guinea). This type of study often takes years. It sometimes takes months for the anthropologist just to gain trust of the local community. Gaining trust is incremental and usually takes a breakthrough after an enduring painful "being pushed to the side" period. We observed this in our SOC fieldwork. Short-term ethnographic work (days or weeks) used in other CS domains simply does not work in a SOC.

Our work shares many of the foundational assumptions of Participatory Design (PD). As Czyzewski, *et al.* noted at the dawn of the PD movement [5]: (1) Our research "turns the traditional designer-user relationship on its head, viewing the users as the experts – the ones with the most knowledge about what they do and what they need." (2) We view the tools we develop "not in isolation, but rather in the context of a workplace; as processes rather than products." (3) Our work "rejects the assumption that the goal of computerization is to automate the skills of human workers, instead seeing it as an attempt to give workers better tools for doing their jobs." While we generally agree with these principles, we also believe that long-term anthropological fieldwork, focused on earning trust, building rapport, and ultimately facilitating long-term collaboration with a relatively small group of analysts is better suited to working with SOC analysts.

2.2 Even the Real Experts cannot Tell How

SOC analysts handle a variety of events related to cyber attacks. They use numerous tools and follow set procedures in handling incidents. The authors have talked with SOC analysts at many organizations. All are uniformly unhappy with current solutions for forensics and incident response. Commercial solutions like Security Information and Event Management (SIEM) systems do not address operational needs, probably because vendors and researchers do not understand how analysts think and work.

SOC analysts often perform sophisticated investigations where the process required to connect the dots is unclear even to analysts. Incident response is not just a technical problem; it involves people with a variety of skills interacting in a closed culture, using specific workflows for each type of incident. Current solutions are not informed by these workflows and are only partially helpful.

The analyst's job is highly dynamic and requires dealing with threats that are constantly evolving. Doing the job is more art than science. *Ad-hoc*, on the job, training for new analysts is the norm. Rookies are intentionally left to find answers by themselves, even though the more experienced analysts could provide hints. The common SOC wisdom is that one must learn the trade

through pain as a necessary "initiation" process. This private learning process contributes to the persistence of hidden or tacit knowledge.

This phenomenon has long been studied in the social sciences, especially anthropology. The term "tacit knowledge" [11] means knowledge that cannot easily be put into words. The tasks performed in a CSIRT job are sophisticated but there is no manual or textbook to explain them. Even an experienced analyst may find it hard to explain exactly how he discovers connections in an investigation. The fact that new analysts get little help in training is not surprising. The profession is so nascent that the how-tos have not been fully realized even by the people who have the knowledge.

Again, Long-term Participant Observation is the Key Cultural and social anthropology could be described as the study of tacit knowledge [7]. Anthropologists do intensive long-term fieldwork in an attempt to reveal and make explicit the "native point of view," which is not just "what natives say" (explicit knowledge) but more importantly, the underlying concepts, presuppositions, and know-how that make up tacit knowledge. While the native point of view is never fully attainable [8], the foundational anthropological method of "participant observation" allows anthropologists to explore the subjects' perspectives and practices by actually taking part in their daily lives and activities (participation) while also standing back from them to gain new perspectives (observation) [2]. If one wants to gain access to these tacit forms of knowledge, one must become embedded in the community of practice. It will take substantial amount of time for the participant, who is also the researcher, to reflect upon the observations and make what is tacit explicit. This has been thoroughly confirmed by our ongoing experience during 15 months' fieldwork.

3 Ethnographic Fieldwork at a University SOC

Author Sundaramurthy has been conducting ethnographic fieldwork at a university SOC for 15 months. As a member of the operations team, he performs many of the analysts' tasks, experiencing their pains, frustrations, and occasional triumphs.

3.1 Fieldwork Setup

The SOC where Sundaramurthy is conducting his fieldwork consists of the Chief Information Security Officer (CISO) and four analysts. Each analyst has specific responsibilities such as incident response, firewall and network management, Payment Card Industry compliance, and Antivirus maintenance. Sundaramurthy and two other PhD students, were initially introduced as "student helpers" who were paid by their adviser (Ou) as graduate research assistants there to learn operational security by doing the work themselves. The CISO was very supportive of the effort.

It took six months for the research team and the SOC to find the best model for managing this relationship. For the first few months the SOC analysts did not understand Sundaramurthy's role. He was asked to perform mundane tasks like maintaining out of date Anti-virus servers and developing hardware requirements for upgrading them. These were day-to-day tasks that the analysts did not have time for, but were not tasks that would allow Sundaramurthy to gain insights into how incident investigations are done. The fieldworkers' work schedules also caused problems. Graduate students do not usually work fixed schedules but the SOC requires dependability and accountability, especially when incidents need to be handled. Ou and the CISO worked to find

mutually agreeable solutions. The CISO released Sundaramurthy from tasks not related to the study goals. The student fieldworkers agreed to work fixed schedules in the SOC and notify the SOC when they could not be present.

We realized the best way to do the study was to do the SOC tasks, be they ticket handling or documentation, in parallel with the research. Later we learned that building tools to help analysts improve their job efficiency was the best way for us to access and use the tacit knowledge embodied in the SOC. This arrangement benefited both the SOC and the researchers. For the first few months, Sundaramurthy spent 15 hours a week in the SOC (currently 6). In 15 months he has worked over 460 hours in the SOC. Another three PhD students worked in the SOC at various times; 2 working 160 hours each over four months; the other, still there, has worked 190 hours over eight months.

3.2 Earning Trust is the Breakthrough

Sundaramurthy was immediately introduced to the tedium of the job's more frustratingly time-consuming and repetitive low-level tasks. The SOC receives alerts on malicious network traffic from a number of trusted sources as well as from their own intrusion detection system (IDS). The alerts contain the IP address (usually that of the NATing firewall) of the infected host and the external IP address with which it was communicating. The real internal IP address has to be extracted from the firewall logs; the MAC address identifying the infected host from DHCP logs. Finding the log entry for a given event and looking up the associated information to resolve the ticket takes about 5 minutes. This repeats and repeats. Before long, Sundaramurthy was fully absorbed. He wasn't creating tools. He was a tool.

It is surprising to the researchers that such a simple event correlation problem does not find an easy solution in any of the off-the-shelf SIEM products of which we are aware. We know from the literature and anecdotal stories that this type of problem is common in other SOCs (commercial, government, and educational).

There were times when Sundaramurthy felt that all his SOC time was spent on carrying out repetitive operational tasks. He felt frustrated because he did not feel he was gaining any insight at all. Many times he tried to talk to the chief incident response/forensics analyst, hoping to learn more advanced investigation skills but was told just to handle the tickets, the highest priority. The chief analyst spent most of his time handling incidents, some too sensitive to give to Sundaramurthy. Sundaramurthy was tied up with the low-level repetitive tasks and he felt that he was not doing anything useful for his research.

The typical SOC mentality focuses on getting incidents processed quickly. It does not have the time for contemplating a long-term vision of improved efficiency. Sundaramurthy was consumed by this mentality. He lost his perspective by becoming part of the SOC! He thought that this is the way things were done and there was nothing he could do to change that. It wasn't until Sundaramurthy discussed his work with his adviser and the research team that he realized that things did not have to be done this way.

After suggestions from co-author Rajagopalan and a senior industrial security analyst collaborating on the project, Sundaramurthy decided to find ways to speed up the ticket handling by building a database of connections and an IP address to MAC address mapping. Noting that most active alerts are a week or less old, he decided to build a caching database retaining seven days of mapping information. Sundaramurthy tried using MySQL which was not able to index the inputs in real time. He then chose MongoDB which stores data as JSON type (schema-less) objects and

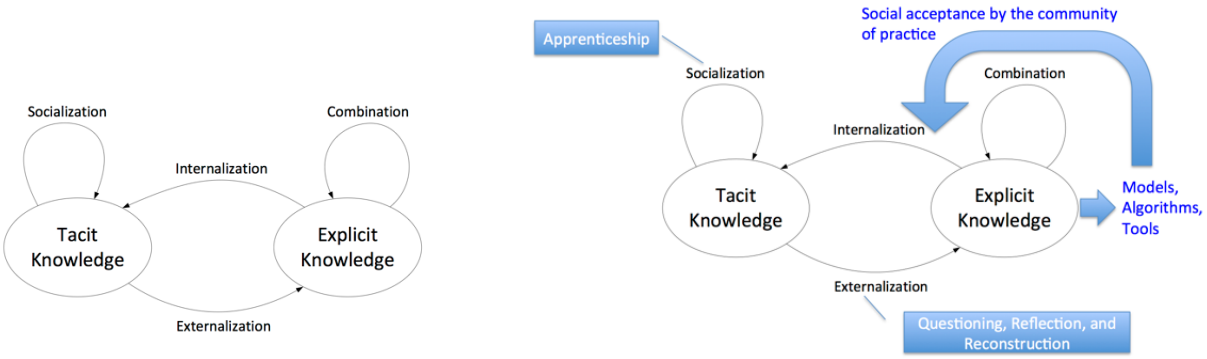

(a) SECI Diagram

(b) Tool-oriented SECI

Figure 1: Knowledge Conversion Models

has a sufficiently high ingest capability.

After the database was operational, Sundaramurthy asked the incident response analyst to use it. The analyst was extremely happy with the performance improvement which reduced ticket handling time from from five minutes to two seconds. The most important result was that the analyst became enthusiastic, willing to talk to Sundaramurthy about possible tool extensions and providing data to expand the database. The two did a long brainstorming session on the whiteboard and arrived at a "Threat Intelligence Framework" that added information sources and relationships among them to the database, allowing a variety of incidents to be handled.

This is the kind thing that Sundaramurthy had wanted to do from the beginning but was not able to do until he understood the workflow and demonstrated his worth. Doing the mundane and repetitive tasks did not appear support his real objective, designing better tools for SOC analysts but were key in extracting the tacit knowledge necessary to build the tool. Once the analyst saw the simple tool Sundaramurthy had built, he completely changed his attitude. Sundaramurthy found his first "informant." The key to this event is to have the subject's trust, which was gained by providing something useful for the analyst. Creating the tool required an understanding of the analysts' process which could only be obtained by performing the task. This first success is a great example showing that being "in the field" allows the fieldworker to see an opportunity to move from "peripheral participation" to "full participation." Now that he had trust and acceptance, Sundaramurthy might also be able to see the world and the tasks more and more through the analyst's eyes and gain a much deeper understanding of the SOC's operations. In other words, we opened the door to access "tacit knowledge" and make it explicit by embodying it in a tool.

3.3 Enhanced Participant Observation Facilitated by Tool Building

Over the past 20 years, researchers in organizational studies have expanded upon and refined a model for how tacit knowledge can be accessed and ultimately transformed into explicit knowledge. Most prominent and relevant for our discussion is Ikujiro Nonaka's SECI model of "knowledge conversion" (Figure 1(a)) which defines four modes of knowledge creation: Socialization (sharing tacit knowledge through apprenticeship and mentoring), Externalization (conversion from tacit to explicit through questioning, reflection, and reconstruction), Combination (accumulation of explicit knowledge), and Internalization (converting new forms of explicit knowledge into tacit knowledge).

The model has been refined over the years. Today it is understood that tacit and explicit knowledge exist on a continuum [10], and that movement along this continuum from tacit to explicit is best achieved through reflection, ongoing practice, and above all, subjecting oneself to a diverse range of alternative experiences that can help disorient oneself just enough to see the tacit dimension — to stop taking for granted the taken-for-granted. In this regard, practice is essential. It is not enough to "join the community," as the elements of knowledge are not necessarily shared across all members of a community. One must be fully engaged in the day-to-day practice.

Based on this strategy and the lessons we learned in earning the trust from the analysts, the fieldwork method we adopt for studying the SOC works as shown in Figure 1(b): (1) researchers become apprentices of CSIRTs analysts (Socialization); (2) researchers reflect, question, and reconstruct what they do (Externalization); (3) researchers design models and algorithms and build tools to help the analysts' job (Combination); (4) researchers take the tool into the work and use the tool as a vehicle to open up more discussions with the analysts and identify more tacit knowledge (Internalization).

As an example of this process at work, after Sundaramurthy developed the Threat Intelligence Framework and released a tool for the analysts to use, the chief incident response analyst wanted to enhance it to handle other incident types. These adaptations were unexpected. One enhancement helps find stolen laptops. Here the perpetrator is usually a student. If the CSIRT knows the MAC address of the stolen laptop, the perpetrator can be apprehended using Access Point (AP) information if he uses the University's authenticated wireless service. If the perpetrator uses any campus service that requires authentication, even through the unauthenticated guest wireless service, we can use its logs together with the other information collected from the authenticated services in the Threat Intelligence Framework to identify him.

Analysts also applied the framework to phishing scam detection. Whenever a phishing email is identified the CSIRT responds from a honeypot university email address using a fake user ID. The CSIRT then watches for a login from that user ID in the future. The IP address associated with that activity is noted and matched against other logins made using the same address but different user IDs since the attacker usually harvests numerous University accounts and tries them in quick succession. Using the framework, a watch can be placed for logins from the honeypot account and other accounts that are possibly compromised can be automatically identified reducing the time the analysts spend in responding to phishing scams.

The success of these efforts created a demand for similar automation by other analysts. Two of successful adaptations illustrate the strength of the approach. Executable attachments to emails are intercepted. The manual process of extracting these, evaluating them and submitting malicious samples to the AV vendor was completely automated. An automated process was developed to detect machines running an out of support OS (Windows XP). This process brought together information from browser "UserAgent" strings obtained from deep packet inspection with address and platform information allowing automatic tracking of presence of such machines and their removal from the network.

The framework automated the handling of multiple types of incidents and enabled the SOC to turn them over to the University's lower-level Network Operations Center.

Tool co-creation: Users vs. Creators In this ongoing collaboration, the true "author" of the tool becomes blurred. The researcher develops a tool which is taken up by analysts and used in ways the researcher might never have imagined. This virtuous cycle produces findings and tools

at the same time. This research methodology differs from both traditional cybersecurity research and anthropological research. Instead of building the algorithms and tools first, the researchers base their model on concrete ethnographic fieldwork data, which yields algorithms and tools that demonstrably help the analysts. The CSIRTs community no longer resists adopting the research prototype since the tool builder is seen as "one of their own." Most importantly, the tool provides an opportunity for analysts to brainstorm with the researchers on additional problems that the tool could be enhanced to address, opening up more venues for sharing the tacit knowledge. In a few instances, we observed that the analysts had difficulty explaining how the tool should be enhanced, illustrating the process of converting the "tacit" knowledge in their mind into explicit forms. We observe that this knowledge conversion process seems to be most effective in this tool-oriented ethnographic fieldwork. From the cybersecurity researchers' perspective, we no longer view the practitioners' role as helping us evaluate our research prototype or providing data. Rather, we view them as *the* experts who possess knowledge that will inform tool building. In some sense, the analysts are co-creating the tools with the researchers. The tool building process reveals tacit knowledge and makes it explicit, but the tools it builds are the key to the fieldworker's acceptance. In this iterative process, we identify and document the key findings concerning the CSIRT analysts' job: how they do the job, and how to make it better.

3.4 Observations from the Fieldwork to Date

The fieldwork has produced observations in two areas. The first relates to the fieldwork, itself. There, we believe that we have a new paradigm in which a participant observer who is also a researcher in the operational area can produce significant findings in both academic and operational areas. The second helps address the longstanding problem of developing tools that are both useful and likely to get used.

3.4.1 On the Research Methodology

Past researchers have realized the tacit nature of knowledge in IT security operations and adopted short-term participant observation [13]. We believe long-term participant observation, on the order of years, is needed to gain insights that reveal deep problems in SOC operations today. Our four fieldworkers have conducted nearly 1000 hours of fieldwork in one SOC over a 15 month period. It took us three months to just earn trust from the analysts and start discussion through tool co-creation. We plan to continue the fieldwork for at least two more years and expand to include more SOCs as the opportunities arise.

Sundaramurthy's realization that he could build a tool to help speed up incident response illustrates a paradox of fieldwork. On the one hand, the farther you are from the community you want to study the more daring your ideas can be, but without being part of the community, your ideas may lack relevance or be unimplementable. Fieldworkers have to be members of the community they are studying and remind themselves often that they are observers as well. Subjective findings are inevitable. It is important for researchers to practice reflexivity — stepping out of the subject role to reflect upon and question what one does and how things are perceived. Anthropologists exist "betwixt-and-between" the world of the researcher and subject. For the rich tacit knowledge existing in environments like CSIRTs, this approach is necessary since the subjects themselves cannot identify and articulate the critical relevant information. Observation *and participation* in the target environment by researchers in a long-term effort is critical to understanding the problem.

Our approach is also markedly different from the classical design ethnography process, where there is a distinct difference between researchers (anthropologists), designers (tool builders), and users (participants). In our work, the three roles are combined into one and our fieldworkers are doing ethnographic fieldwork while designing new tools to be used by themselves and other SOC analysts. This unique mode of ethnography is determined by the nature of the SOC environment — one would not be able to simply observe the analysts using tools built by a third party and draw the same deep insights. There is a tight collaboration between the researcher and the research subjects and the fieldworker's role is the perpetual trinity of researcher, designer, and user all in one.

To move from "peripheral participation" to "full participation," fieldworkers need to be fully accepted into the community. The most straightforward way for fieldworkers to gain full acceptance is by designing tools that can help the analysts do their jobs. Tacit knowledge about process and organizational structure can only be teased out if researchers have deep discussions with the SOC analysts within the concrete contexts. The tools Sundaramurthy built started, catalyzed, and fostered this type of discussions. We have gained much deeper insights than we would have if we had not proposed building tools to change the SOC's workflow. Our fieldwork not only helped researchers understand the SOC's workflow and processes, but also helped the SOC improve them.

3.4.2 On Tools and Technologies' role in the SOC

Our own experience, like that of other anthropologists engaging in similar practical applications of participant observation, has been one of a continuous flow of subtle and sometimes not so subtle insights that continually reshape our understanding. Workplaces are complex social environments, made even more complex by the use of complex systems and tools. We do not just use tools. As the tool is used it changes our routines and habits. It changes the way we think about and address problems. It might change who we collaborate with and how we collaborate with them. It might even be the catalyst for a complete restructuring of an organizational chart or workflow. This is perhaps best illustrated by Sundaramurthy's tool which changed the workflow of the SOC's incident response, enabled the SOC's analysts to automate the simple repetitive tasks, and formed a standard operating procedure to be handled by lower-level less skilled analysts. As John Culkin (invoking the insight of his colleague, Marshall McLuhan) noted, "We shape our tools and thereafter our tools shape us." [4] We recognize that the relationship between humans and their tools is always going to be a complex one.

4 Closing Thoughts

This work started in 2012 at a meeting of McHugh, Ou, and Rajagopalan. We had two questions: (1) how can we make SOCs and in particular CSIRTs more effective (by any reasonable metric) and (2) how can cybersecurity researchers play a significant role in the improvement of SOCs and CSIRTs. We have managed to make some progress on both fronts. The article describes our approach and findings, summarized below:

- **The use of anthropology is effective**

  With the addition of a professional anthropologist (Wesch) to our team our methodology has solidified. Under his guidance, our fieldworker, Sundaramurthy, has made significant

progress. We have uncovered new understandings of the reality of CSIRT operations that we have not encountered either in print or apocrypha. New findings will continue to emerge as the research progresses.

- **The importance of participation**

  Starting with the naive view that observation is sufficient for understanding CSIRTs, we changed to a robust participation model. The fieldworker actively participates in the environment. Our experience has shown the effectiveness of the engagement and yielded immediate results in terms of rapport with the CSIRT staff.

- **The importance of tool co-creation**

  An obvious indicator of our effectiveness is the level of acceptance of our embedded fieldworkers and the tools created jointly by the research team and the CSIRT staff. Having CSIRT staff co-own the tools enriches the tools and enables its acceptance in the CSIRT environment.

- **The criticality of tacit knowledge**

  Our experience has shown very clearly that the problems of CSIRT operations are not merely technological. They are exacerbated by fast turnaround and high volume. None of our tools were technically sophisticated. Vendor tools present in the environment were ignored by the SOC team due to a mismatch between the CSIRT personnel's internal model of the process that they undertake and the model implicit in the vendor tools' designs. No tool can be effective until that internal model is made explicit. Every case where our tool was accepted in the CSIRT environment is an instance of tacit knowledge converted into explicit form.

- **The importance of reflection and the open-ended nature of the learning process**

  We believe that most security researchers (and computer scientists in general) lack introspective and reflexive skills. As a result, they get caught up in their solutions without taking the time to consider the import of the work or its place in a spectrum of real-world problems. We note that computer science started as a discipline devoted to developing tools to solve computational problems from other areas with many of the first generation of CS researchers having their roots in mathematics, physics, engineering, business, or other areas with hard computational problems. As the field came into its own, many academic computer scientists drifted away from real-world problems and concentrated on more tractable abstractions or simplified cases. In security, the use of noise-free test data is an example. By joining forces with anthropology, we have learned the importance of understanding the world of the SOC analyst, and have been able to gain their trust through participant observation. By reflecting on our observations, we were able to bridge the gap between building an arbitrary "computer sciency" tool and a tool that supports the tacit needs of the analyst's workflow. This is part of a learning process for both the researcher / tool builder and the analyst. There are no definite endpoints to such a learning process. We must be continuously aware of how our presence, and the presence of the tools we build, might shape the research environment itself. The entire process is inherently reflexive and demands ongoing commitment to a careful and critical analysis of our own biases and assumptions. As we learn about the CSIRT, the external world keeps changing and the CSIRT has to continuously adapt. New knowledge that gets incorporated into the CSIRT is very likely to be tacit because of the experiential nature of the problem. It must be converted to explicit form as we go forward.

4.1 Future Work

We are extending and expanding this effort to include additional SOCs. We would like to find more partners to work with, so that our study can be more representative. We need our collaborators to dedicate some human resources to doing fieldwork. Collaborating organizations will benefit from a third-party perspective of operational effectiveness, intra-team interactions, *etc.*, in the context of cybersecurity operations. They may also benefit from tools that the fieldworkers build or help build for the organization. At the end of the project, we expect to write a training manual for organizations employing cyber security operations personnel. The manual should be useful to commercial, academic, and government SOCs.

References

[1] Michael H. Agar. *The Professional Stranger: An Informal Introduction to Ethnography*. Emerald Group Publishing, 2 sub edition, 1996.

[2] H. Russell Bernard. *Research Methods in Anthropology: Qualitative and Quantitative Approaches*. AltaMira Press, 5th edition, 2011.

[3] Andy Crabtree and Tom Rodden. Ethnography and design. In *International Workshop on 'Interpretive' Approaches to Information Systems and Computing Research*, 2002.

[4] John Culkin. A schoolman's guide to Marshall McLuhan. *Saturday Review*, March 18:51–53, 71–72, 1967.

[5] Paul Czyzewski, Jeff Johnson, and Eric Roberts. Introduction: Purpose of PDC 90. In *PDC 90 Conference on Participatory Design*, pages ii–iii. CPSR, 1990.

[6] Paul Dourish and Genevieve Bell. *Divining a digital future: mess and mythology in ubiquitous computing*. MIT Press, 2011.

[7] Julia Elyachar. Before (and after) neoliberalism: Tacit knowledge, secrets of the trade, and the public sector in Egypt. *Cultural Anthropology*, 27(1):76–96, 2012.

[8] Clifford Geertz. From the native's point of view: On the nature of anthropological understanding. *Bulletin of the American Academy of Arts and Sciences*, 28(1):26–45, 1974.

[9] Charles Leinbach. Managing for breakthroughs: A view from industrial design. In Susan Squires and Bryan Byrne, editors, *Creating Breakthrough Ideas: The Collaboration of Anthropologists and Designers in the Product Development Process*, pages 3–16. Greenwood Publishing, 2002.

[10] I Nonaka and G von Krogh. Tacit knowledge and knowledge conversion: Controversy and advancement in organizational knowledge creation theory. *Organization Science*, 20(3):635–652, 2009.

[11] Michael Polanyi. *The Tacit Dimension*. DoubleDay & Company, 1966.

[12] Susan Squires and Bryan Byrne. *Creating Breakthrough Ideas: The Collaboration of Anthropologists and Designers in the Product Development Industry*. Greenwood Publishing, 2002.

[13] Rodrigo Werlinger, Kirstie Hawkey, and Konstantin Beznosov. Security practitioners in context: Their activities and interactions. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems*, 2008.

Author Biographies

Sathya Chandran Sundaramurthy is a PhD candidate in Computer Science at Kansas State University in Manhattan, KS. His current research interest is studying security operation centers (SOCs) using anthropological methods. In the past, he has also worked on applying probabilistic and artificial intelligence techniques for automating intrusion analysis. He has a bachelor's degree in Computer Science and Engineering from Anna University, India. Contact him at sathya@ksu.edu.

John McHugh is the Senior Principal and Chief Analyst at RedJack, LLC. and an Adjunct Professor of Computer Science at the University of North Carolina. He previously held the Canada Research Chair in Security and Privacy at Dalhousie University in Halifax, NS, was a senior member of the technical staff at CERT CC, part of the Software Engineering Institute at CMU, held a Tectronix Professorship and served as chair of the CS department at Portland State University, has taught at UNC and at Duke and worked for RTI in North Carolina. His research interests include network data analysis and operational security. He has a PhD in Computer Science from The University of Texas, a MS in Computer Science from the university of Maryland, and a BS in Physics from Duke. Dr. McHugh is a senior life member of the IEEE. Contact him at mchugh@cs.unc.edu.

Dr. Xinming "Simon" Ou is associate professor of Computer Science at Kansas State University. He received his PhD from Princeton University in 2005. Before joining Kansas State University in 2006, he was a post-doctoral research associate at Purdue University's Center for Education and Research in Information Assurance and Security (CERIAS), and a research associate at Idaho National Laboratory (INL). Dr. Ou's research interests focus on designing better technologies to facilitate cyber defense. He is a recipient of 2010 NSF Faculty Early Career Development (CAREER) Award, and three-time winner of HP Labs Innovation Research Program (IRP) award. Contact him at xou@ksu.edu.

Raj Rajagopalan is a Senior Principal Scientist with Honeywell Automation and Control Systems (ACS) Research Labs where he leads a team of researchers who work on providing appropriate cybersecurity tools for Honeywells portfolio of control systems. He was previously a Research Scientist at HP Labs Cybersecurity Research in Princeton, New Jersey. His research interests include safety and security for operational control systems. He has a PhD in Computer Science from Boston University. Contact him at siva.rajagopalan@honeywell.com.

Michael Wesch is an associate professor of cultural anthropology at Kansas State University. He received his PhD from the University of Virginia in 2006 after two years of field research on the effects of writing on a remote indigenous culture in Papua New Guinea. Since then his research has focused on the effects of media and technology on global society. Contact him at mwesch@ksu.edu.

Acknowledgment

This research is supported by the National Science Foundation under Grant No. 1314925. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.