Agile Applied Research for Cybersecurity: Creating Authoritative, Actionable Knowledge When Speed Matters

Rick Linger
Cyber and Information Security Research Group
Oak Ridge National Laboratory
Oak Ridge, TN 37831
lingerr@ornl.gov

Luanne Burns Goldrich
Johns Hopkins University
Applied Physics Lab
Laurel, MD 20723
luanne.burns@jhuapl.edu

Matt Bishop
Dept. of Computer Science
University of California, Davis
Davis, CA 95616
mabishop@ucdavis.edu

Melissa Dark
Computer and Information Technology Department
Purdue University
West Lafayette, IN 47907
dark@purdue.edu

Abstract

Securing information systems from attack and compromise is a problem of massive scope and global scale. Traditional, long-term research provides a deep understanding of the foundations for protecting systems, networks, and infrastructures. But sponsors often need applied research that will create results for immediate application to unforeseen cybersecurity events. The Agile Research process is a new approach to provide this type of rapid, authoritative, applied research. It is designed to be fast, transparent, and iterative, with each iteration producing results that can be applied quickly. The idea is to engage subject-matter experts fast enough to make a difference. Agile Research requires new levels of collaboration and performance, plus adaptive organizational structures that support this new way of working. In addition to its application in Government, Agile Research is being employed in academic settings, and is influencing how research requirements and researchers are identified and matched, and research traineeship.

1. The Need for Agile Research

Traditional, long-term research often involves extensive requirements definitions, comprehensive proposals, competitive awards, distributed organizational structures, complex funding protocols, and long-term performance that can extend for years or even decades [1]. These processes are embedded in the national research and development infrastructure, as embodied, for example, in Defense Advanced Research Projects Agency (DARPA), National Science Foundation (NSF), National Institutes of Health (NIH), Department of Energy (DOE), and other Government organizations. When the scope and scale of research requirements are large, as, for example, with autonomous vehicles or alternative energy, these traditional processes and their management and review procedures are essential to maintaining control across collaborating organizations and reducing risks of overruns and nonperformance. As such, they serve a vital role in conducting large-scale, long-term research projects to achieve national goals.

These broad national research goals will always be with us. But events occur in cybersecurity areas that require fast and decisive responses in order to protect national well-being and even survival. These responses would benefit from rapid and authoritative analysis by the best minds and organizations. The traditional research infrastructure is ill-suited for this level of fast engagement and immediate application, leaving a pressing need for institutional innovations in the research infrastructure.

An Agile Research process is being developed and implemented to address the need for fast and effective exploratory applied research in situations where speed is an overarching requirement. When attempts have been made to apply traditional methods in these situations, the research results, no matter how comprehensive and valuable, are often too late to be of use in the current cybersecurity event, and wind up as shelfware.

Wells and Smyth [2] presented an agile approach to developing research methodology, arguing that qualitative research is emergent (which it is), so instead of thinking about the research methods a priori, the re-

URI: http://hdl.handle.net/10125/41883
ISBN: 978-0-9981331-0-2

search methods should really be thought of iteratively and agilely. While agile applied cybersecurity research can include adaptive research methodologies, this effort is more broadly focused on institutional innovation. This effort is relatively new, and at this point in time, the main contribution of this work is to merge two different styles of research: the agile, exploratory method that focuses on applied research with the academic, broader method that focuses on foundational research. We argue that the two are complementary, and this synergy can lead to advances both in foundational research and applied research in a way that, taken separately, the two cannot achieve. The first key idea is that the agile research develops a specific target, leading into questions that drive foundational research that in turn lead to advances in applied research. The second key idea is that training students in such a way that they understand the bidirectionality of the flow between applied and foundational research is critical for developing research infrastructure to support the advancement of knowledge and applications of that knowledge in cybersecurity.

We begin with a discussion of the nature and role of institutional innovation in technical innovation. Next we present our innovative approach, namely the framework for agile research formalizing that process, and the principles that underlie it. From these principles, we derive a waypoint that exposes fundamental questions not suitable for applied research, but eminently suitable for deep, long-term foundational work. We then discuss incorporating these methods into academia and other research-oriented institutions.

2. Technical and Institutional Innovation

Research and development produce technical change. This technical change is carried out in institutions such as research universities, national laboratories, industrial research laboratories, and experiment stations. The work within and among these institutions is shaped by the physical, social, economic, and cultural environments within and around them. While institutions are places such as a research university or a national lab, institutions are also (perhaps more so) the social roles played by these places and the social rules that specify how these places will interact with each other. Institutional innovations then are changes to the roles that are played by these places, and changes to the rules that shape how these places interact with each other. An historic example of an institutional innovation in research and development is the U.S. Bayh-Dole Act of 1980, which fundamentally changed the nation's system of technology transfer by enabling universities to retain title to inventions and take the lead in patenting and licensing groundbreaking discoveries.

While institutions (roles, rules, and places) need to be stable for extended time periods in order for progress to occur, at times institutions need to change in order for technical innovation to continue. Past institutional contributions to technical progress can and do get thwarted. Growing disequilibria eventually creates sufficient demand for institutional changes. We contend that there is sufficient demand in cybersecurity research and development, making now the time for a dramatic shift to an agile-based approach.

3. Agile Research Principles

Agile Research is a method for conducting exploratory applied research. It is organized around *sponsors*, who pose research questions to be answered, and *researchers*, who conduct the research and produce results. Sponsors and researchers may be in the same or different organizations, and may be organized in any number of ways provided the following principles are satisfied.

- *Principle of Predefined Infrastructure*. Resources and logistics must be defined and allocated before research needs emerge to permit immediate deployment for fast engagement when needed. Agreements between sponsors and researchers regarding organizational roles, research capabilities, contracting, funding, and intellectual property must be in place and ready to be instantiated in unforeseen circumstances with no delays. This "load-and-go" approach permits fast reaction using pre-positioned resources to meet unpredictable research needs unburdened by logistical constraints. It also provides sponsors and researchers with flexibility to redeploy additional resources at each iteration of the work should that prove necessary.
- *Principle of Incremental Research*. Agile research is structured into iterative, short-term, accumulating increments that each produces actionable results. Increments focus on understanding the problem, and progress to solution strategies, and then to incremental solutions. Understanding how to organize applied research into a series of accumulating, referentially transparent increments requires careful planning that should be revisited frequently as the work progresses. Early increments must provide a framework for inserting and composing later increments such that results accumulate with little or no revision of prior work.
- *Principle of Incremental Management.* The incremental research process provides built-in, short-

term checkpoints for sponsors to understand researcher progress, and to direct subsequent work based on incremental findings. Agile Research projects can be quickly refocused based on changes in both fast-paced problem environments and on intermediate shortfalls and windfalls in the research with minimal loss of work and time. Visibility, transparency, and clear and forthright communication between researchers and sponsors are essential for informed management decision making, and researchers must be prepared to change direction as necessary to achieve desired outcomes.

- *Principle of Transferability*. Agile Research projects may be carried out by one group of researchers, but ready transfer of results from one group to another must be possible if necessary. As research increments are completed and changes in direction are made, mechanisms for quickly repositioning the research and resources to a new team must be in place. This includes knowing where the research expertise exists for the next increment, as well as providing supporting documentation and consultation that permits a new team to pick up the work seamlessly and rapidly.

4. The Agile Research Process

Agile Research projects proceed through up to four stages, each culminating in researchers delivering results, either through briefings, white papers, tools, or a combination of these. At the completion of each stage, the sponsor decides whether and how to proceed. This process is summarized in Figure 1.

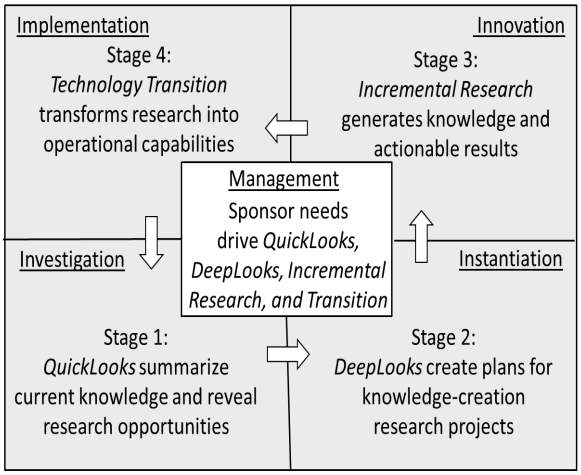

Figure 1. The Agile Research process

1. The *QuickLook Stage* generally takes days or weeks. It answers the question of what is known now about the problem. During this stage, the research team clarifies the research needs with the sponsor, develops appropriate hypotheses, explores the existing knowledge base, identifies subject-matter experts (SMEs), and provides recommendations to form a foundation for the research effort. This stage is deliberately made flexible to accommodate urgent or even emergency needs. In extreme situations, this stage could be accomplished by teleconference or email with subject-matter experts.
2. The *DeepLook Stage* generally takes weeks. Based on results from the QuickLook stage, it answers the question of what the applied research can be expected to accomplish and how should it be done. It defines the research goals and plans in terms of iterative, accumulating increments that produce useful results for sponsors. During this step, the hypotheses proposed in the QuickLook Stage are refined, and the approach to testing them determined.
3. The *Incremental Research Stage* consists of multiple incremental steps, generally performed in weeks or months per increment. Every iteration adds to an evolving solution to the problem. This step-wise approach permits sponsors to modify incremental research goals and apply results based on the intermediate findings as the work progresses. The results may lead to a change or refinement of hypotheses or the methods of testing them or both.
4. Finally, if a project requires technology transfer, the *Technology Transfer Stage*, generally performed in months, provides specifications, prototypes, and support to guide technology implementation and operational use of intermediate and final research results.

Agile Research is flexible. A project might require only a QuickLook to determine the state of knowledge for a particular problem. Or, a project could continue to a DeepLook to understand what the research could accomplish were it continued to the next stage, and how the research in that stage should be structured. The sponsor could then initiate the incremental research. This research portfolio management process is depicted in Figure 2. A set of anticipatory QuickLooks can be created and periodically updated when necessary as the environment evolves. Most importantly, QuickLooks can be initiated in response to unforeseen cybersecurity events. QuickLooks of current importance can be selected for DeepLooks to be prepared should further research become necessary.

Agile Research embodies properties that support fast response and flexibility while maintaining intellectual rigor. Among these properties are:

- *Speed*: Agile Research is structured to avoid unproductive time-sinks that delay the applied research and the application of the findings. This approach emphasizes quick reaction to sponsor needs. Given their level of knowledge and experience, subject-matter experts can quickly provide findings and recommendations at the QuickLook level, and can work effectively with sponsors to develop research approaches and plans at the DeepLook level. Research increments are planned to expect results to be generated within defined intervals to ensure progress toward a timely solution.
- *Quality*: Moving quickly does not mean sacrificing rigor. Researchers are interested in seeing their work make a difference, and understand that in applied research dealing with immediate problems, slow research can be overtaken by events and become irrelevant to those problems. Agile Research is not slowed down by logistics. It permits researchers to focus solely on the problem at hand, and builds in peer review through the daily give-and-take of multi-disciplinary, multi-organizational teams. The incremental approach itself can improve applicability and rigor through opportunistic adaptation to unforeseen results.

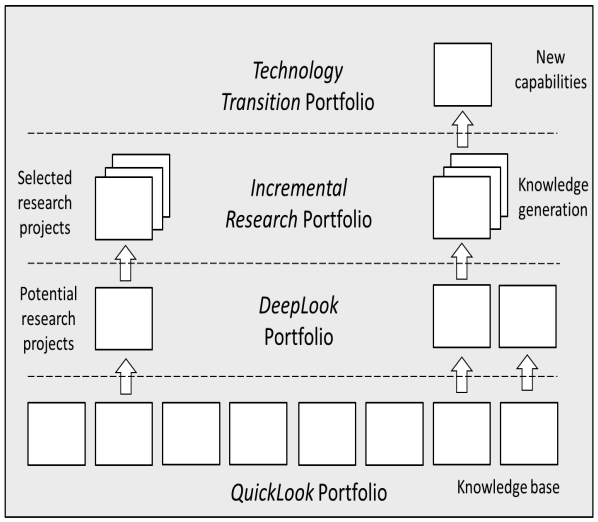

Figure 2. Agile Research portfolio management

- *Visibility*: Agile Research is a working partnership between sponsors and researchers. The incremental process provides transparency for sponsors through briefings, white papers, tools, or other research output, followed by decisions to proceed or not. This approach offers opportunities to reconfigure the remaining hypotheses, work, and resources for maximum effect.
- *Effectiveness*: Agile research is designed to produce incremental and actionable results that accumulate into a complete solution. Organization of a program of applied research into accumulating increments requires rigorous, yet adaptive planning that can help reveal the internal structure of a problem and the building blocks required to produce actionable results. The planning itself can become part of the solution process by avoiding false starts and wasted effort.
- *Impact*: Applied research is useless if it has no real impact on the problem. Agile Research keeps the problem statement at the forefront of all activities. The incremental process is geared to providing a series of partial solutions that reduce the remaining unsolved parts of the problem at each step.
- *Opportunism*: Research work is inherently unpredictable. Reducing that unpredictability prescribes an effective management process based on disciplined incremental development. Each increment may produce unforeseen results to which the process must adapt. Shortfalls are valuable because they show what will not work, and help guide the research into alternate strategies. Windfalls are valuable because they confirm existing approaches. In either case, the management process is to opportunistically adapt to intermediate findings to achieve best results.

5. An Agile Research Example

The Institute for Information Infrastructure Protection (I3P), formerly led by Dartmouth College and now led by George Washington University and SRI International, is an organization of 26 leading universities, national laboratories, and Federal Funded Research and Development Centers (FFRDCs) dedicated to advancing national cybersecurity capabilities. The I3P has employed Agile Research in a demonstration performed for a government agency. A QuickLook study was carried out by I3P subject-matter experts to investigate Data Tagging research issues with respect to the following (abbreviated) requirements provided by the sponsor. The applied research objectives flowed directly from these requirements:

- Examine existing information control data tagging for attribute-based access control (ABAC), with access controlled by policy-based attributes and data tags employed by an enterprise-scale system that processes substantial volumes of data.

- Identify technologies that can be adapted, combined, or extended for data tagging needs.
- Conduct research on how to use data tagging to incorporate definition, evolution, auditing, and management of access and retention policies and their implementation.
- Identify additional relevant research objectives.

The following research findings are drawn from the executive summary of the QuickLook study, substantially abbreviated to fit within the space constraints of this paper. Each recommendation is stated as a direct, actionable task, and each was fully elaborated in the report to provide guidance on how to carry it out. These findings, provided by subject-matter experts, created a framework the customer on how to proceed with deeper research for the Data Tagging project. The recommendations were organized into three areas: Way Forward, Solution Space, and Requirements Analysis.

5.1. Data Tagging Way Forward: Findings and Recommendations

- Define a path forward in light of the complexity of the problem.
  - Recommendation: Organize the complexity of the problem through structured, divide-and-conquer refinement of goals and requirements.
  - Recommendation: Explore the existing data tagging solution space for cost-effective application to the problem.
- Conduct incremental research and development.
  - Recommendation: Develop a hierarchical goal set to address agency needs.
  - Recommendation: Conduct research into tag representation and management as a rigorous foundation for information sharing.
  - Recommendation: Develop a proof of concept system to explore and evaluate potential solutions.

5.2. Data Tagging Solution Space: Findings and Recommendations

- There are promising existing commercial solutions.
  - Recommendation: Run a public challenge for data tagging to elicit potential solutions.
  - Recommendation: Conduct data tagging product evaluations.
- The agency is beginning to pilot solutions for enterprise data tagging in several areas.
  - Recommendation: Study data tagging design patterns of (agency name elided).
- Other Government organizations are beginning to tackle enterprise data tagging.
  - Recommendation: Evaluate design patterns used in (agency name elided).
  - Recommendation: Investigate an earlier (agency name elided) information discovery and assured access study.

5.3. Data Tagging Requirements Analysis: Findings and Recommendations

- The problem domain is too complex to tackle with traditional requirements specification.
  - Recommendation: Conduct a structured engineering assessment to define incremental development and deployment stages.
- An information architecture is needed for data tags.
  - Recommendation: Develop a data tagging Concept of Operations (CONOPS).
  - Recommendation: Conduct an organizational inventory of attribute data.
  - Recommendation: Assess taxonomies and ontologies for representing tags.
  - Recommendation: Conduct a tradeoff study of tagging-data-at-rest vs. tagging-data-on-the-fly.
- Tagging technologies and mechanisms must be secured.
  - Recommendation: Develop definitions of potential threats and vulnerabilities.
  - Recommendation: Develop security reference architectures for data tagging.
  - Recommendation: Assess efficacy of Identity-Based Internet Protocol (IBIP) to secure the data tagging network.

This first step clearly indicates several paths through the DeepLook Step. It also suggests several more foundational research questions.

6. Organizing for Agile Research

Because logistical structures must be predefined to enable rapid response to research needs, the Agile Research process can provide a turn-key capability to sponsors. For example, a Research Provider Organization (RPO) could establish advance relationships with groups that could conduct this type of applied research, and serve as a single point of contact for sponsors in forming multidisciplinary teams for particular needs.

Figure 3 depicts an example RPO serving sponsors through team solicitation and formation based on predefined relationships with key sources for performing

applied research in UARCs/FFRDCs, academic institutions, and national laboratories.

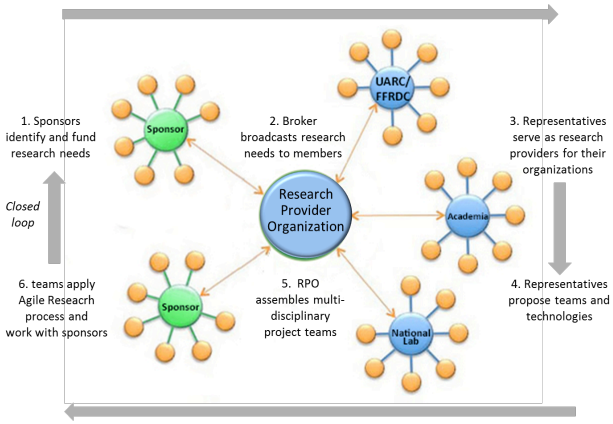

Figure 3. An example Agile Research structure

In this hub and spoke model, the RPO maintains knowledge of subject-matter expertise through representatives in its member organizations, and can move quickly to address research needs. Sponsors are freed from the need to establish these relationships and maintain this information, and can simply "pick up the phone" to initiate research requests through the RPO.

7. Agile Research in Academia

In addition to its application in government and industry, the Agile Research framework is well-suited for teaching students the applied research process, and preparing graduates capable of working in fast-paced, mission-oriented environments that require research competencies. Agile Research captures many traditional elements of long-term research, such as locating and understanding primary research literature; formulating a research problem; designing a research study; analyzing data; presenting data coherently and effectively; interpreting data; and preparing research materials for publication or presentation. Agile Research permits teaching these skills on a time scale compatible with academic semester and term structures. The work done in such a program also may well lead to questions that require fundamental research, thereby demonstrating to the student the value of that long-term research.

The first step in all research is to understand the problem being studied. Perhaps it is one with immediate real-world applications, such as whether a particular clustering method will enable an intrusion detection system to correlate alarms quickly. Perhaps it is more foundational, such as whether P = NP. The researcher must have a clear understanding of the problem, and the parameters within which it is to be analyzed.

The next step is to see what others have done to solve the problem, or problems related to it. This typically involves searching literature; it may also involve contacting experts working in the field to see whether they have extended their reported results. Then students analyze this earlier work to determine its applicability. If the work is not applicable, they say why it is not; if it is, the researchers distinguish their approach, or decide how to advance the previous work to make a further contribution.

The third step is to plan the research. Often the plan is incremental in the sense that it has specific sub-goals at which intermediate results can be presented, and the researchers can determine whether the research should end, a change of direction is necessary, or the research should continue. In academia, these sub-goals usually result in academic papers; in industry, they result in changes to existing products, new products, or new directions to pursue.

When the research is complete, or a sub-goal results in a prototype product or system, the new technology is transferred to the sponsor. In some cases, the sponsor will have an outside organization (such as a commercial firm) develop the prototype into a robust, functional tool with a usable interface. In other cases, the sponsor will be a commercial organization that will transfer the prototype to a production unit that will then productize the system or tool.

Compare these four steps to the four stages of the Agile Research process. They are essentially identical, with the primary differences being the time involved and the communication between the sponsor and the researchers. In the Agile Research process, the time frame for the first two steps is greatly compressed. Throughout all the steps, communication between the sponsor and the researchers is much tighter for the Agile Research process than for the traditional research process. That way, the sponsor can give immediate feedback to the work as it progresses to ensure that it is useful, and can understand both the work being done and the results produced. The sponsor can, if needed, retarget the work as it progresses.

The current predominant model for teaching research competencies in the United States includes research methods and content matter classes, which serve as the foundation for the research experience that is manifest as a thesis or dissertation. The thesis or dissertation is a sole endeavor by the student designed to confirm the student's ability to conduct research and advance the state of understanding in the field. The experience is highly individual, and while necessary, it is no longer sufficient to develop the types of research competencies required of graduates.

The evolving nature of research problems requires researchers be more ready to participate in multi-

disciplinary, multi-institutional, collaborative, cooperative, persistent, and distributed research teams. Graduate education largely focuses on development of the researcher through learning opportunities that construct research as an individual, episodic, in-discipline effort. While graduate students interact with others in their research group, their faculty mentors and advisors, and (usually rarely) the research sponsors, the ultimate goal of a graduate student researcher is individual: to complete their graduate project, thesis, or dissertation and graduate. Thus, the implications for research traineeship are that students need educational experiences that teach them how to work in an interdisciplinary environment, work in research teams with researchers from diverse types of organizations, leverage existing data and tools, work with increasing data volumes and varieties, develop and leverage research networks, and manage research projects.

The Agile Research framework can be used to provide these types of research traineeship experiences. Currently, the INSuRE (Information Security Research Education) [3] project is applying the Agile Research framework to research traineeship.

The INSuRE project is developing a partnership among ten successful and mature Centers of Academic Excellence in Information Assurance Research (CAE-R) and the National Security Agency (NSA), several national labs, two state agencies, and one military base in order to design, develop, and test the research network. INSuRE is a self-organizing, cooperative, multi-disciplinary, multi-institutional, and multi-level collaborative research project that can include both unclassified and classified research problems in cybersecurity.

The project permits students to work on real-world problems, as well as to be mentored by practitioners, rather than focusing solely on faculty-led research. Students benefit from the guidance of multiple, inter-disciplinary research faculty from several institutions. The student-led research may provide solutions for pressing national problems. While INSuRE is still in a pilot phase, students teams have produced solutions that are being transitioned to practice in sponsor organizations.

To facilitate scientific discovery, learning, and collaboration, an open source software platform called HUBzero® is employed. HUBzero includes a content management system designed to support scientific activities. Users on a hub can write blog entries and participate in discussion groups. They can work together on projects, publish datasets and computational tools with Digital Object Identifiers (DOIs), and make these publications available for others to use as live, interactive digital resources. Simulation and modelling tools published on a hub can be accessed with the click of a button. They run on cloud computing resources, campus clusters, and other national high-performance computing (HPC) facilities.

INSuRE, using the Agile Research framework, is innovating dimensions of research traineeship. The problems provided by sponsors are multidisciplinary in nature. This requires students to evaluate the nature of the problems and the types of disciplinary knowledge required to solve them, and to form teams that bring the requisite knowledge to bear. Students are expected to identify (and learn how to recognize) needed skills and expertise outside their background, whether these are to be learned by the student or brought in through search/choice of collaborators. Within INSuRE, the research problems are worked on across multiple institutions, either concurrently or sequentially. The project repository keeps incremental reports from previous Quick and Deep Looks. Teams are expected to know where other expertise resides within the network, and to leverage that expertise in their approach. Students in INSuRE are required to work in teams with peers with different levels of skills, knowledge, expertise, and research experience. The teams are fluid in that they are not just close peers (students in a single institution's class) but also include other students brought in on an ad hoc basis, other professors, and subject-matter experts within sponsor organizations. Teams are expected to instrument results and reports specifically so that they can be picked up in operational contexts directly, and be passed on to another team for further work. Thus, the INSuRE project exposes students to constructing research as a continual effort embodied in smaller tasks than a thesis or dissertation. It also exposes students to the evolution of research problems by rapidly and iteratively involving them in increments, reports, refocused problem setting, and re-engaged research work.

As an example, one project in a recently-completed INSuRE class involved looking at data leakage from mobile devices. Current approaches involve static and dynamic analysis, used in combination. A recent paper [4] described how to combine the two, and presented impressive results. The students examined the work described in the paper, and especially the limitations. They noticed that, under certain conditions (specifically, when variables used have unknown values), only one branch of a conditional would be taken (this is "approximation mode" and reduces the number of paths to be analyzed). A limitation of the model prevents the dynamic analysis from determining these values at run time, again under specific conditions. The students changed the analysis approach to use symbolic rather than specific values, and as a result improved the coverage and were able to detect previously unknown leaks. The applied research question, tackled by the Agile Research method, was to deter-

mine how to detect data leakage. The foundational question that arose from this research is to determine how much the use of symbolic execution reduces the number of false positives and negatives compared to not using it, and how the use of symbolic execution affects testing performance.

The limited time frame and the need for guidance suggest that the Agile Research process will be pedagogically more effective than the traditional approach. Following completion of the class, the students can of course continue the research in a more traditional framework. It may well turn out that experience with Agile Research in an academic setting will transfer directly into work performance with organizations requiring fast and authoritative results to deal effectively with unforeseen cybersecurity events.

8. Matching Sponsors with Researchers

In recognition of the value of research in non-traditional settings where speed is a factor, organizations are increasingly adopting procedures to systematize and streamline the process of matching sponsor needs with researcher capabilities.

A powerful approach to this matching step, itself a precursor to an effective Agile Research project, is embodied in a system named REQcollect (Requirements Collection Repository). REQcollect was developed by Johns Hopkins University Applied Physics Laboratory (JHU/APL) to enhance the R&D mission of Federal departments and agencies in terms of organizing and managing research work.

The system is used to gather and store research requirements and research project information, facilitate correlations between requirements and projects, and assist in launching transitions. It is the central repository for storing information about elicited requirements, discovered projects and technologies, matches between requirements and technologies, integration activities, and lessons learned. REQcollect uses an automated Apache Lucene [5] matching algorithm to complete a Google-like full-text search over project descriptions and requirement keywords to suggest prioritized lists of matches between requirements and projects.

After matches are made, users may select elements and characteristics for technology transition. This centralization and standardization of project and requirement data provides automated, suggested matches and discovery [6]. Before development of REQcollect, matching a research requirement to a research project was a time-consuming, manual process. The algorithm used by REQcollect streamlines the approach and supports increased objectivity in the selection process by eliminating human bias.

The web-based interface of REQcollect allows for easy insertion and editing of requirements and projects. Reporting utilities provide an easy interface for extracting requirements and generating reports sorted on fields of the user's choosing: for example, requirements by priority, by organization, by category or by multiple fields. Requirements can be deprecated and reports can include or exclude deprecated requirements. Figure 3 shows the REQcollect home page.

REQcollect systematizes research project management, reduces risk by facilitating productive matches between sponsor requirements and performer capabilities, and enables fast team formation and entry into performance mode. The government agency for which the Data Tagging QuickLook was produced employs REQcollect as part of a sophisticated embedded research and development process targeted to achieving efficient and effective results from funded research.

9. Future Work

Agile Research is a new paradigm that provides a basis for pursuing applied research, and seeding fundamental research, by demonstrating the relevance of that research to sponsors' needs, and by giving sponsors an idea of what they can gain from that research. This is done by producing deliverables early in the process, thereby enabling sponsors to focus more tightly on funding research that will meet their needs.

Figure 3. REQcollect web-based frontend

Critical to the success of an Agile Research program is matching sponsor requirements and needs with groups that can carry out the research in the required time frame. Currently, there is no systematic way to do

the matching. Many sponsors have databases of researcher capabilities, but these are often imprecise and frequently out of date. How to create matching tools so they can be used effectively and how to ensure capabilities stay up to date, are complex questions. One obvious problem is that the language used to describe the requirements must be compatible with the language used to describe the capabilities. How to do this semantic comparison is an interesting question in the theory of natural language processing.

As the paradigm is new, there is no set of best practices or guidelines for conducting it. In many cases, structures supporting research will require adaptation to accommodate Agile Research; indeed, for some organizations, adopting this model may be counterproductive or simply not possible. An interesting question is how to determine when Agile Research rather than, or in addition to, traditional long-term research, is the right choice. A set of best practices and guidelines would help to determine this, as one could then match these with the organization considering the new paradigm. This could convince organizations to adopt a business model that supports this type of research plan. Fortunately, the cost of entry is low; QuickLooks are fast and inexpensive, and permit organizations to gain experience with the process.

An interesting research question is the notion of incremental results that a sponsor will find immediately useful. This concept of "incremental deliverables," where each deliverable builds on its predecessors, is a key technical aspect of Agile Research. This is in some sense similar to a requires-provides model of attack [7]. In that model, an attacker must have certain capabilities to take a step towards compromising a system; once that step is taken, she gains additional capabilities that enable the attack to advance further. Here, the "attackers" are the researchers and the "steps" are the incremental results.

All this raises a very interesting question: how can sponsors and researchers develop intermediate goals so that incremental results are useful, will enable the sponsor to provide further guidance to the research group, and (especially in an academic setting) provide insight into the foundational research necessary to provide deeper understanding of the problem and, possibly, long-term solutions. For example, perhaps the researchers find that the first incremental goal they agreed upon cannot be met given the context of the problem. The sponsor and the research group can then work to define another useful goal that is attainable. This cycle of problem refinement will help focus the research, and help establish limits on what can be done so that sponsor's expectations become more realistic. Structuring research goals so that useful intermediate objectives can be met is a difficult, yet needed, research problem in itself.

Finally, the Agile Research method has been used only in limited circumstances. How does it work in the general research environment? In academia, the IN-SuRE program may help answer some of these questions because, as noted above, the work being done by the students essentially follows the Agile Research process. In any event, Agile Research exhibits properties that are critical to research involvement in the fast paced and unpredictable world of cybersecurity.

Acknowledgements:: This manuscript has been authored by UT-Battelle, LLC, under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. This submission was written by the author(s) acting in their own independent capacity and not on behalf of UT-Battelle, LLC, or its affiliates or successors.

This material is based upon work supported by the National Science Foundation under Grant Number DGE-1303211 to the University of California at Davis and Grant Number DGE-1303048 to Purdue University. Matt Bishop and Melissa Dark thank NSF for this support. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

9. References

[1] Committee on Innovations in Computing and Communications, *Funding a Revolution: Government Support for Computing Research*, National Academies Press, Washington DC (1999). ISBN: 978-0-309-06278-7

[2] H. Wella and H. Smyth, "An Agile Approach to the Real Experience of Developing Research Methodology and Methods," *Designs, Methods, and Practices for Research of Project Management*, Beverly Paisan, *ed.*, Gower Publishing, Surrey, UK (2016). ISBN 978-1-4094-4880-8.

[3] M. Dark, M. Bishop, R. Linger, and L. Goldrich, "Realism in Teaching Cybersecurity Research: The Agile Research Process," *Information Security Education Across the Curriculum: Proceedings of the 9th World Information Secury Education Conference* pp. 3–14 (May 2015). DOI: 10.1007/978-3-319-18500-2_1.

[4] M. Xia, L. Gong, Y. Lyu, and Z. Qi, "Effective Real-Time Android Application Auditing," *Proceedings of the*

2015 IEEE Symposium on Security and Privacy pp. 899–914 (May 2015). DOI: 10.1109/SP.2015.60.

[5] L. Branscomb and P. Auerswald, "Between Invention and Innovation an Analysis of Funding for Early-Stage Technology Development," National Institute for Standards and Technology, NIST GCR 02–841 (Nov. 2002). URL: http://www.atp.nist.gov/eao/gcr02-841/contents.htm.

[6] L. Goldrich, S. Hamer, M. McNeil, T. Longstaff, R. Gatlin, and E. Bello-Ogunu, "REQcollect: Requirements Collection, Project Matching and Technology Transition," *Proceedings of the 47th Hawaii International Conference on System Sciences* pp. 4887–4894 (Jan. 2014). DOI: 10.1109/HICSS.2014.599.

[7] S. Templeton and K. Levitt, "A Requires/Provides Model for Computer Attacks," *Proceedings of the 2000 New Security Paradigms Workshop* pp. 31–38 (2000). DOI: 10.1145/366173.366187.