Artificial Intelligence Exploration (AIE) Opportunity
DARPA-PA-19-03-02
Artificial Intelligence Mitigations of Emergent Execution (AIMEE)

I. Opportunity Description

The Defense Advanced Research Projects Agency (DARPA) is issuing an Artificial Intelligence Exploration (AIE) Opportunity, inviting submissions of innovative basic research concepts in the technical domain of machine learning classifier diversity.

This AIE Opportunity is issued under the Program Announcement for AIE, DARPA-PA-19-03. All proposals in response to the Artificial Intelligence Mitigations for Emergent Execution opportunity, as described herein, will be submitted to DARPA-PA-19-03-02. If selected, it will result in an award of an Other Transaction (OT) for prototype project. The total award value for the combined Phase 1 base and Phase 2 option is limited to $1,000,000. This total award value includes Government funding and performer cost share, if required.

A. Objective and Technical Scope

Modern computing systems demonstrate strong propensity for unintended, emergent computations and the related unintended, emergent programming models (colloquially known as "weird machines"[1] that enable or amplify cyber-attacks. Computing mechanisms built for a particular purpose and with particular intended models of execution in mind prove to be capable of executing unintended computing tasks outside of their original specification and their designers and programmers' mental models.

Today, we start examining systems for signs of emergent behavior—with methods such as fuzz-testing—only after they are fully built. Despite the empirical prevalence of emergent execution phenomena, theory describing emergent execution is scarce.

However, recent research strongly suggests that a system's exploitability models and propensity for emergent execution arise—and can also therefore be mitigated—already at the design stage, when the system's programming abstractions and intended behaviors at a particular layer are translated into the more granular states and logic of the next computing substrate layer down in the computing stack. As programming abstractions must be translated to ancillary layers of more granular states and logic to implement and optimize these abstractions, each layer creates latent opportunities for emergent execution and unintended programmability, which are inherent in a

[1] *E.g., Weird machines, exploitability, and provable unexploitability*, Thomas Dullien, IEEE Transactions on Emerging Topics in Computing, December 2017
Framing Signals—A Return to Portable Shellcode, Erik Bosman, Herbert Bos, IEEE Symposium on Security and Privacy, 2014
Exploitation as Code Reuse: On the Need of Formalization, Sergey Bratus, Anna Shubina, Information Technology, vol. 59, no. 2, p. 93, 2017
Exploit Programming: from Buffer Overflows to Weird Machines and Theory of Computation, Sergey Bratus, Michael E. Locasto, Meredith L. Patterson, Len Sassaman, Anna Shubina, USENIX ;login:, 2011

few key design decisions of the translation, and will manifest themselves regardless of a wide variety of other implementation details and choices.[2]

The Artificial Intelligence Mitigation of Emergent Execution (AIMEE) will address the problem of anticipating, at a system's design stage, the models of emergent execution inherent in its design, and thus mitigate its propensities for exploitability before they lead to actual vulnerabilities in complete deployed systems.

The AIMEE AI exploration will explore whether a combination of recent advances in AI techniques such as autoencoders, evolutionary programming, deep representation learning, neural sketch learning, etc., can be used to detect, describe, and model the primitives of emergent execution directly in design-level prototypes of performance optimizations and programming abstractions rather than in complete-system implementations.

AIMEE proposals will propose case studies of how AI methods could be applied to design prototype-level representations of several common computing layers known to manifest emergent execution behaviors and programming models ("weird machines"), to enable effective anticipation of these behaviors and models. Strong proposals will discuss ways to generalize these case studies and methods for use with the exemplar designs such as layered APIs, ABIs, or CPU microarchitectures.

The work proposed must include the development of:

(a) A theory and formal model of emergent execution and emergent programmability.

(b) Case studies of emergent execution and programming models that apply the theory in (a) to design-level prototypes of computing abstractions and demonstrate the use of AI methods and representations to anticipate emergent behaviors.

(c) An outline for a general algorithmic approach to constructing representations of computing system designs suitable for the application of AI methods to aid system designers and programmers in mitigating emergent execution at system design and prototyping time.

Proposers should specify (1) the problem domain(s) (e.g., layered APIs, ABIs, CPU microarchitectures or other abstraction layers of computing stacks), and (2) the classes of emergent execution phenomena that their work will address. Covering all possible types of emergent execution is likely beyond the scope of this effort, so proposers should pick representative classes such that success against them would be strong evidence that the approach will generalize.

B. Structure

Proposals submitted to DARPA-PA-19-03-02 in response to this AIE Opportunity must be UNCLASSIFIED and have an 8-page limit. Proposals must address two independent and sequential project phases: a Phase 1 Feasibility Study (base) and a Phase 2 Proof of Concept (option). The periods of performance for these phases are 9 months for the Phase 1 (base) effort and 9 months for the Phase 2 (option) effort. Combined, Phase 1 and Phase 2 efforts for this AIE

[2] *E.g., Spectre is here to stay: An analysis of side-channels and speculative execution,* Ross Mcilroy, Jaroslav Sevcik, Tobias Tebbi, Ben L. Titzer, Toon Verwaest, https://arxiv.org/abs/1902.05178, 2019

Opportunity should not exceed 18 months. The Phase 1 (base) award value is limited to $500,000. The Phase 2 (option) award value is limited to $500,000. The total award value for the award is limited to $1,000,000. This total award value includes government funding and performer cost share, if required.

C. Schedule/Milestones

Phase 1 fixed milestones for this program must include the following:

Month 1: An outline of the formal approach to modeling emergent execution. The outline must include the specific classes of formal methods and representations to be used, and demonstrate how these representations and methods model empirically known phenomena of emergent computation.

Month 3: A plan of case studies consistent with the formal approach and representations provided in the outline. The plan must specify several empirically known and practically significant examples of computing designs in which unintended programming models arose and leveraged emergent execution phenomena. The plan should show how the formal methods and representations will facilitate the study of these examples.

Month 6: Initial results from the case studies, showing application of AI methodologies to anticipate emergent execution phenomena, and refinement of the formal approach.

Month 9: Reports of successful case studies with application of AI methods, with methodological summaries, algorithms, datasets, and finalized formalisms consolidated in a report. For unsuccessful case studies, a report of the encountered obstacles and summaries of the approaches attempted.

Phase 2 fixed milestones for this program must include the following:

Month 12: Generalization of the accomplished case study examples and approaches, unifying the case study methodologies to a single formal methodology applicable at the design stage of systems.

Month 15: Expansion of empirical base to common patterns of system design, demonstrated by reports of successfully applying the methodology to a new set of case studies in different system domains.

Month 18: Produce final report demonstrating effectiveness of AI methods for anticipating emergent execution at system design time.

All proposals must include the following meetings and travel in the proposed schedule and costs:

- To foster collaboration between teams and disseminate program developments, there will be two meetings in Phase 1 (kick-off and at month 9), and two meetings in Phase 2. In both phases of the effort, there will be one meeting in the Washington, D.C. area and one at a location to be determined. For budgeting purposes, assume the second meeting will be San Diego, CA. These meetings should be attended not just by principal investigator(s), but also by any key students or developers who have been involved in the project.
- Regular teleconference meetings will be scheduled with the government team for

progress reporting as well as problem identification and mitigation. Proposers should also anticipate at least one site visit per phase by the DARPA program manager during which they will have the opportunity to demonstrate progress towards agreed-upon milestones.

D. Deliverables

Performers will be expected to provide at a minimum all reports and data required by the individual Phase 1 and Phase 2 milestones. These may include registered reports, experimental protocols, publications, intermediate, and final versions of software libraries, code, and APIs, including documentation and user manuals, and/or a comprehensive assemblage of design documents, models, modeling data and results, model validation data, and algorithm-generated lexicons or grammar rules.

II. Award Information

Selected proposals that are successfully negotiated will result in award of an OT for Prototype project. See Section 3 of the AIE Program Announcement (DARPA-PA-19-03) for information on awards that may result from proposals submitted in response to this notice.

Proposers must review the model OT for Prototype agreement provided as an attachment to the AIE Program Announcement (DARPA-PA-19-03) prior to submitting a proposal. DARPA has provided the model OT in order to expedite the negotiation and award process, and ensure DARPA achieves the goal of AIE, which is to enable DARPA to initiate a new investment in less than 90 days from idea inception. The model OT is representative of the terms and conditions that DARPA intends to award for all AIE awards. The task description document, schedule of milestones and payments, and data rights assertions requested under Volume 1, Volume 2, and Volume 3 of the AIE Opportunity will be included as attachments to the OT agreement upon negotiation and award.

Proposers may suggest edits to the model OT for consideration by DARPA and provide a copy of the model OT with track changes as part of their proposal package. Suggested edits may not be accepted by DARPA. The Government reserves the right to remove a proposal from award consideration should the parties fail to reach agreement on OT award terms and conditions. If edits to the model OT are not provided as part of the proposal package, DARPA assumes that the proposer has reviewed and accepted the award terms and conditions to which they may have to adhere and the sample OT agreement provided as an attachment, indicating agreement (in principle) with the listed terms and conditions applicable to the specific award instrument.

In order to ensure that DARPA achieves the AIE goal of award within 90 days from the posting date (October 9, 2019) of this announcement, DARPA reserves the right to cease negotiations when an award is not executed by both parties (DARPA and the selected organization) on or before January 6, 2020.

III. Eligibility

See Section 4 of the AIE Program Announcement (DARPA-PA-19-03) for information on who may be eligible to respond to this notice.

IV. AIE Opportunity Responses

Responses to this AIE Opportunity must be submitted as full proposals to DARPA-PA-19-03 as

described therein. All proposals must be unclassified.

A. Proposal Content and Format

All proposals submitted in response to this notice must comply with the content and format instructions in Section 5 of DARPA-PA-19-03. All proposals must use the templates provided as Attachments to the PA and the "Schedule of Milestones and Payments" Excel Attachment provided with this AIE Opportunity and follow the instructions therein.

Information not explicitly requested in DARPA-PA-19-03, its attachments, or this notice may not be evaluated.

B. Proposal Submission Instructions

DARPA/I2O will employ an electronic upload submission system (https://baa.darpa.mil/) for proposals responding to this AIE Opportunity topic (AIMEE – DARPA-PA-19-03-02)

First-time users of the DARPA BAA Submission website must complete a two-step account creation process at https://baa.darpa.mil/. The first step consists of registering for an Extranet account by going to the above URL and selecting the "Account Request" link on the right side of the page, using the Chrome browser. Upon completion of the online form, proposers will receive two separate emails; one will contain a user name and the second will provide a temporary password. Once both emails have been received, proposers must go back to the submission website and log in using that user name and password. After accessing the Extranet, proposers must create a user account for the DARPA BAA Submission Website by selecting the "Register Your Organization" link at the top of the page. The DARPA BAA Submission Website will display a list of solicitations open for submissions. Once a proposer's user account is created, they may view instructions on uploading their proposal.

Proposers who already have an account on the DARPA BAA Submission website may simply log in at https://baa.darpa.mil/, select this solicitation from the list of open DARPA solicitations and proceed with their proposal submission. Note: Proposers who have created a DARPA BAA Submission website account to submit to another DARPA Technical Office's solicitations do not need to create a new account to submit to this solicitation.

However, the proposer should verify that the account is still active and access can be achieved prior to the day that proposal submissions are due.

All submissions submitted electronically through DARPA's BAA website must be uploaded as zip files (.zip or .zipx extension). The final zip file should contain only the files requested herein and must not exceed 50 MB in size. Only one zip file will be accepted per submission.

Note: Submissions not uploaded as zip files will be rejected by DARPA.

Proposers submitting a proposal via the DARPA Submission website MUST (1) click the "Finalize" button in order for the submission to upload; AND (2) do so with sufficient time for the upload to complete prior to the deadline. Failure to do so will result in a late submission, which will NOT be accepted or reviewed.

Technical support for the DARPA BAA Submission website may be reached at Action@darpa.mil and is typically available during regular business hours (9:00 AM – 5:00 PM ET, Monday-Friday).

Questions regarding submission contents, format, deadlines, etc. should be emailed to

AIMEE@darpa.mil.

Since proposers may encounter heavy traffic on the web server, it is highly recommended that proposers not wait until the day proposals are due to request an account and/or upload the submission. Proposals should not be submitted via Email. Any full proposals submitted by Email will not be accepted or evaluated.

C. Proposal Due Date and Time

Proposals in response to this notice are due no later than Friday, **November 8, 2019, 12:00 NOON (ET)**. Full proposal packages as described in Section 5 of DARPA-PA-19-03 must be submitted per the instructions outlined therein *and received by DARPA* no later than the above time and date. Proposals received after this time and date may not be reviewed.

Proposers are warned that the proposal deadline outlined herein is in Eastern Standard Time and will be strictly enforced. When planning a response to this notice, proposers should take into account that some parts of the submission process may take from one (1) business day to one (1) month to complete (e.g., registering for a Data Universal Numbering System (DUNS) number or Tax Identification Number (TIN).

DARPA will acknowledge receipt of complete submissions via email and assign identifying numbers that should be used in all further correspondence regarding those submissions. If no confirmation is received within two (2) business days, please contact AIMEE@darpa.mil to verify receipt.

V. Proposal Evaluations and Selection

Proposals will be evaluated and selected in accordance with Section 6 of DARPA-PA-19-03. Proposers will be notified of the results of this process as described in Section 7.1 of DARPA-PA-19-03.

VI. Administrative and National Policy Requirements

Section 7.2 of the AIE Program Announcement (DARPA-PA-19-03) provides information on Administrative and National Policy Requirements that may be applicable for proposal submission as well as performance under an award.

VII. Point of Contact Information

Sergey Bratus, Program Manager, DARPA/I2O, AIMEE@darpa.mil

VIII. Frequently Asked Questions (FAQs)

All technical, contractual, and administrative questions regarding this notice must be emailed to AIMEE@darpa.mil. Emails sent directly to the program manager or any other address may result in delayed or no response.

All questions must be in English and must include name, email address, and the telephone number of a point of contact. DARPA will attempt to answer questions publically in a timely manner; however, questions submitted within 7 days of the proposal due date listed herein may not be answered.

DARPA will post an FAQ list under the AIE Opportunity on the DARPA/DSO Opportunities page at: http://www.darpa.mil/work-with-us/opportunities?tFilter=&oFilter=2&sort=date. The list will be updated on an ongoing basis until one week prior to the proposal due date. In addition

to the FAQ specific to this notice (DARPA-PA-19-03-02), proposers should also review the Program Announcement for AIE General FAQ list on the DARPA/DSO Opportunities page under the Program Announcement for AIE (DARPA-PA-19-03).