Efficient and Resilient Edge Intelligence for the Internet of Battlefield Things

Maggie Wigness, Tien Pham, Stephen Russell
U.S. DEVCOM Army Research Laboratory
UNITED STATES
{maggie.b.wigness, tien.pham1, stephen.russell15}.civ@mail.mil

Tarek Abdelzaher
University of Illinois
UNITED STATES
zaher@illinois.edu

ABSTRACT

A doctrine of convergence of military capabilities from multiple domains to enhance efficacy heralds a new age in defense, marked by the ability to withstand a higher operation scale and tempo, enabled by increased levels of automation and coordination in the battlefield. Reaping the potential benefits of these technological advances, however, is predicated on finding successful solutions to a myriad of challenges in order to enable more efficient and scalable operation of intelligent, heterogeneous, interacting assets in contested environments. Said differently, increased automation and coordination of defense capabilities call for a smarter "battlefield operating system" – a system that manages complex automated tasks at time-scales that preclude human engagement, while empowering the warfighters with adequate control. We call this operating system, the Internet of Battlefield Things (IoBT). In this article, we focus on challenges in upholding three principles of superiority (in modern conflict) upon which the IoBT is based. Namely, (i) time is a weapon; winners are those who minimize the latency between their sensors and shooters, (ii) IoBT is a fighting network; all functions must withstand an active, determined and technologically sophisticated adversary, and (iii) machine intelligence belongs at the point of need; a new breed of AI solutions is needed that can be projected rapidly to the point of need, where they can survive the austere environment of field operations, as opposed to restricting AI to solutions that run at data centers of higher echelons.

Solutions to the above challenges, developed in the Internet of Battlefield Things Collaborative Research Alliance (an alliance of research institutions from the government and academia, funded by the U.S. Army Combat Capabilities Development Command, known as DEVCOM, Army Research Laboratory (ARL)) are discussed that (i) map the capability envelop (i.e., help understand the fundamental feasibility limits of envisioned IoBT capabilities) (ii) optimize for performance (i.e., improve the IoBT cost/value trade-offs by offering intelligent capabilities at a lower cost), and (iii) ensure resilience (i.e., improve the capacity of developed IoBT capabilities to withstand a broad spectrum of threats in challenging battlefield environments). We focus specifically on functions that involve machine automation and threats that jeopardize the AI itself. While defense science has a long history of investigating solutions for protecting physical assets, once automation enters the loop and is relied upon as a superior alternative to manual operation that automation or Artificial Intelligence (AI) demands the same kinds of emphasis on protection, as it is critical to operational advantage. A key challenge addressed in the Internet of Battlefield Things is therefore to protect the efficiency, efficacy, and integrity of the IoBT itself.

1.0 INTRODUCTION

The Collaborative Research Alliance (CRA) for Internet of Battlefield Things Research on Evolving Intelligent Goal-driven Networks (IoBT REIGN) aims to develop and transition basic science that significantly advances performant and resilient real-time *computational networked services for the battlefield* (henceforth called

battlefield IoT services) that can be used to fight cohesively in any environment, including conditions where the electromagnetic spectrum is denied or degraded. The research focuses on the underlying science necessary for building more resilient command and control (C2) systems that are *robust*, *intelligent*, *efficient*, and *scalable*. The IoBT concept demands supporting C2 systems that connect sensors and effects and that carry out military decision processes in scenarios involving all-domain operations, as captured by the multi-domain effect loop, depicted in Figure 1. The figure was originally introduced in SPIE 2020 [1] and outlines a breakdown of the observation to effect cycle into seven stages. An overarching goal of the research is to provide the intelligence and autonomy that can accelerate the execution of the MDO effect loop. In this context, battlefield IoT services must meet challenges of distribution, heterogeneity, size/scale (in multiple dimensions), high op-tempo, and operation in harsh environments in the presence of adversarial disruption/denial.

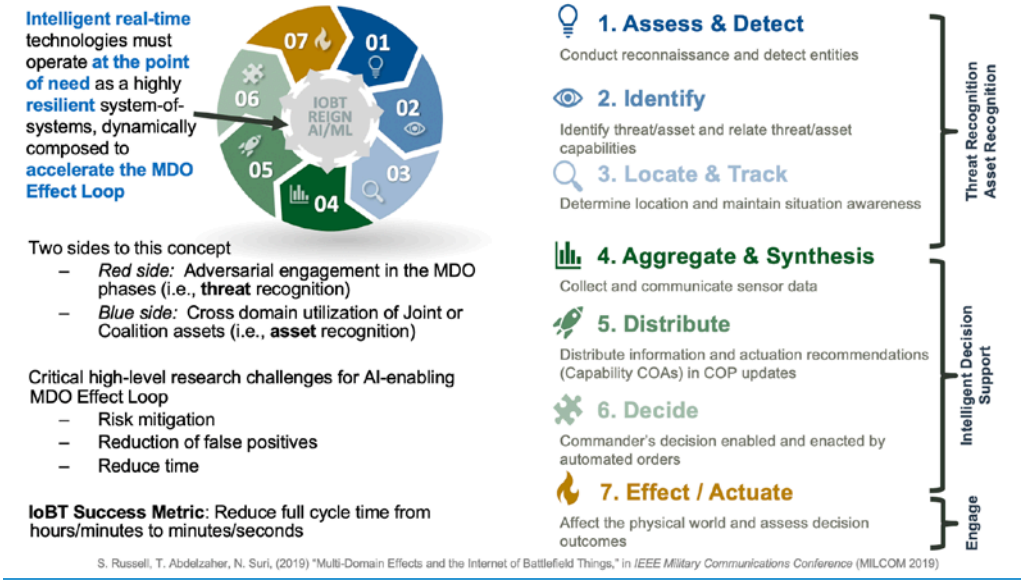

Figure 1: The multi-domain operations (MDO) effect loop.

To understand the need for IoBT research, it is helpful to make an analogy between today's state of the art in battlefield automation and the early days of computing. At the dawn of the computing era, new machines with enhanced processing capabilities enabled significant computational advantages, being faster and more precise than their human counterparts. The multiplicity of applications soon necessitated the development of operating systems to address common challenges such as resource-efficiency, execution robustness, scalability, and responsiveness that the applications required. In a modern battlefield where mission success depends in large part on performance of computational artifacts that optimally operate collaboratively, a new operating-system-like construct is in order. The construct's goal is to ensure that the execution of sensor-to-effect decision loops, which involve multiple intelligent devices and systems, meet the challenges arising from spatial distribution, accelerated mission-tempos, transient resources, uneven environmental conditioning and the potential presence of adversarial activity.

The research contributions discussed in this paper can each be roughly aligned to the seven stages of the MDO effect loop. Given that the MDO effect loop supports C2 decision-making and that a human commander may often be making the decision to engage, i.e., executing the "decide" stage of the loop, most of the algorithms and contributions address the last and the first five stages of the loop and automate the last stage. However, it should

be noted that the decision stage could be fully automated, no humans in the loop, and that the delivered effect may be kinetic or non-kinetic. The IoBT represents the common operating environment where individual technologies exist, are composed to execute a process, and are employed with a military purpose – to deliver an effect. Thus, it is both the aggregated composite technologies and the process in which they are utilized in that form the IoBT common operating environment. Discussion of these contributions is broken down by their emphasis on three critical elements of modern battlefield superiority, time is a weapon, the IoBT is a fighting network, and machine intelligence is necessary at the point of need. It is the aggregation of these elements that enhance decision dominance by accelerating the execution of the full cycle of the MDO effect loop.

2.0 KEY CHALLENGES: FROM SENSORS TO EFFECTS

"[There] are going to be fundamental technologies that will change how we fight in the future." said General John Murray, the Commanding General of the U.S. Army Futures Command in a recent comment on exercises of Project Convergence, "It's artificial intelligence. It's machine learning. It's the network that will support all that. It's autonomy, it's robotics, and it's really the underlying data architectures and how we manage data."

General Murray's quote is an example of changes in military thinking that suggests that battlefields of the future will rely increasingly on *computationally augmented distributed systems*, wherein physical military assets will work hand-in-hand with computational, networking, and data management counterparts to attain the desired objectives. It is in this context that IoBT provides the underlying operating "system" or environment. However, if the IoBT is an operating system/environment, what do the applications that function on or in it look like?

The applications of IoBT, elaborated in [1], are the cyber-physical decision loops that ultimately enable the instantiation of mission effects. While one might perceive the MDO effect loop [1] as a typical sensor-to-shooter activity, the MDO effect loop model generalizes the decision process for delivering any cross-domain effect, kinetic or not [2]. It is helpful to break down these loops into the stages shown in Figure 1. In a typical loop, actuation is informed by observations of opportunities, threats, or other causes for action. It may require subsequent assessment to determine whether the chosen action (i.e., effect) succeeded. Elements involved in the loop may be heterogeneous, and may include assets from the ground, air, space, sea, and cyber domains. While it is common to consider kinetic effects as the actuation engaged, the effects may include a broad menu of non-kinetic alternatives as well. For example, in an intelligent battlefield, where an adversary's assets (e.g., unmanned combat vehicles, tactical operations centers, military units, etc.) need to communicate to accomplish their joint mission, creating interference in the communication medium may temporarily disable decision-making capabilities across these assets and offer a window of uncontested access to their vulnerabilities. In general, examples of effects include (i) stimulating the adversary to action (e.g., in order to reveal a capability, vulnerability, location, or other information), (ii) revealing the adversary, (iii) striking the adversary, or (iv) assessing outcomes of an interaction with the adversary.

It is noteworthy that the MDO effect loop, from observation to effect, may be non-linear and recursive to any individual stage. In other words, a stage can loop back to previous stages to gain additional inputs. Further, given decisions internal to each stage, a nested loop can be executed as necessary for the purpose of stage completion. Elements of the loop may operate at different temporal and spatial scale. Multiple loops may be instantiated in different domains that interact via the aggregation and distribution components. For example, detection components in different domains might operate at the fastest time-granularity, locally. When changes of interest (or of concern) are suspected, additional assets may be brought in to assist with identification, localization, and tracking. While these stages may operate on local objectives, an important capability is to aggregate the information at multiple spatial scales and across domains in order to identify higher-level patterns,

such as coordinated movements of troops across large areas. In turn, any conclusions from such larger-scale analysis may need to be shared with the individual loops and decision makers. Decisions can then be made in accordance with mission objectives and commander intent. Effects are ultimately employed in accordance with executed decisions. Assessment follows to understand the impact of employed effects. An explicit goal is to enable autonomous execution of this loop to provide a much higher scale and faster tempo than manual operation alone might support; all while also executing in a manner that is more resilient to adversarial threats. It is these performance and resilience requirements that motivate IoBT advances.

3.0 TECHNICAL ACCOMPLISHMENTS

A key distinguishing property in executing the above loop in future battlefields lies in the large operation scale and high operation tempo that potentially exceed human capacity to keep up with all dynamics. Future battlefields will exhibit high degrees of lethality, making it important to operate reliably despite a broad variety of threats. The inability of human operators to keep up with the large-scale rapidly evolving loop and adversarial actions (e.g., inability to manually analyze all incoming sensor data for detection and identification purposes, manually aggregate such data across multiple sources to confirm gathered intelligence, manually cross-cue sensors, manually perform adaptation to adversarial action, and manually perform weapon-to-threat assignment) means that machine intelligence must step in, delegating the human to a more supervisory and/or decision role. IoBT fundamental research prepares machine intelligence for this responsibility through advances in fundamental science. Below, we elaborate on the three principles guiding IoBT advancements as mentioned earlier. Namely: (i) time is a weapon, (ii) fighting network, and (iii) intelligence at the point of need.

3.1 Time is a Weapon

It is well known timing can provide advantages and disadvantages. To this point, on a contemporary battlefield time is a weapon. The IoBT must reduce the time elapsed between sensing and effect. Considering the full scope of the MDO effect decision loop, sensing supplies information. Key changes in sensory measurements may thus require timely detection and response. A fundamental question in this context is the following: how quickly can one detect change from weak indicators? In principle, if sensors were perfect, and if the detected phenomena were easily distinguishable and deterministic, then a change would be detected as soon as sensor values changed. This is not the case in the complexity of modern battlefields, where a variety of factors contribute to measurement errors, where the detected features are often stochastic and inconclusive. Fundamental feasibility limits on early detection (i.e., *earliest change detection* results) have recently been developed for classes of dynamic anomalies [3], moving objectives [4], growing anomalies [5], and heterogeneous objectives [31], and extended to the challenging case of distributed detection problems [6] [7].

An especially challenging case arises when an intelligent system can continuously learn. The ability to learn allows the system to adapt to state or mission changes, which is assumed to happen often in a highly dynamic battlefield environment. Thus, rapid adaptation is critical to ensure correct and safe operation to defeat the adversary. Take for example a learning-enabled system, such as a radar-equipped intelligent unmanned aircraft system (UAS), that can develop a model of its environment which it uses for inference tasks. Specifically, the UAS might learn to differentiate targets from decoys. When and how should the system conclude that reality has drifted sufficiently such that the true underlying model has now changed (beyond an acceptable threshold) from the previously learned model? For example, it might be that the adversary has deployed a new type of decoy, rendering the UAS's previously learned model of decoys inaccurate. This challenge can be formulated as a *model change detection* problem [8], and the speed at which these changes can be identified and mitigated is critical to mission success. New adaptive sequential machine learning algorithms [9] [10] have recently been

developed, in the IoBT context, to offer analytical foundations for detecting (and mitigating) model change. Another approach has been developed to ensure model accuracy evaluates model prediction errors to interpret changes in sensors and the environment. Small errors indicate typical environment evolution and larger prediction errors indicate unexpected change, each resulting in their respective model adaptation [38].

The scale of data within an IoBT may be massive in part because of the vast number of sensors within the network, but also because of the increased speeds at which sensors are now able to take readings or measurements. If the machine intelligence that is part of the IoBT is not able to process data at the same rate as an adversary then decision dominance is lost. To begin to address this need, a new neural network framework was designed to prioritize inference execution and reduce latency for mission-critical stimuli [36]. Multimodal information was used to define scheduling of perception-based inference at a subframe level, e.g., parts of images instead of the entire frame, where more critical regions of interest are processed first. Intelligent image resizing based on mission-criticality and resource batch processing has also been designed to improve detection speed [37].

IoBT is by definition a distributed operational environment. Thus, federated learning and inference in the presence of a dynamic real-time environment are key challenges. To reduce data transfer latency, detection and inference tasks must be executed closer to the sensor platform, implying that execution requirements of such tasks must be significantly reduced. A distributed inference approach has been developed to partition neural network inference models across edge and cloud resources by evaluating the dynamic network bandwidth and edge resource capabilities [35]. The combination of progressive model layer slicing and network partitioning enables the approach to maintain low latency and energy use while maintaining high performance accuracy. For the case where the platform must consult a remote server to complete the full inference, a new brand of compression is introduced, where data over the communication link are compressed in a manner guided by the nature of the inference algorithm to be executed [17][32]. Taking the algorithm into account, significant improvements are attained in the compression factor. Recent work also addressed latency optimization in multimodal sensing, where speculative inference is carried out with only a subset of the inputs when some of the fused sensing modalities are inaccessible or delayed [34]. Speculative execution was shown to significantly improve end-to-end data fusion and inference latency with only minimum effects on quality.

In an IoBT environment, aggregation and distribution tasks are essential, as individual sensors and platforms must cue each other to accomplish collaborative tasks, federated learning systems must share information to collectively learn from individual distributed observations, and updates must be sent to headquarters. Significant work was done over the past decades on network flow optimization, but more aggressive measures are needed to alleviate the communication burden imposed by modern sensors, learning systems, and inference algorithms on the contested scarce communication resources of a battlefield. To attain the next leap in network optimization efficacy, IoBT research takes a new stance. Namely, it rethinks resource allocation protocols to consider the *purpose of communication*. Knowing the purpose opens up additional optimization opportunities that significantly improve the attainable trade-off between resources used and results achieved [19]. The work starts with asking some fundamental questions. For example, what are optimal communication patterns for distributed learning [20]? What are the fundamental storage needs of learning algorithms? These needs bound the degree of attainable compression [21][22]. Can federated learning algorithms employ task-dependent compression schemes to dramatically reduce their communication needs [23]? The ideas behind exploiting purpose from communication have so far resulted in 3-4 orders of magnitude improvement in communication efficiency in distributed learning from complex data samples [24]. These efficiencies, in turn, significantly reduce end-to-end latency of systems that rely on intelligent detection and automation.

3.1 IoBT is a Fighting Network

With the advent of machine learning algorithms, the concepts of decoys and camouflage have been extended to exploit properties of machine learning in a manner that produces misclassifications and misprediction [11]. Specifically, adversarial input examples design minimal changes in appearance (that do not significantly alter sensory input) to cause misclassifications in inference algorithms. A key challenge in IoBT systems is therefore to detect such adversarial input manipulation. A recently proposed metric, called *attribution-based confidence* [12], for assessing the likelihood of adversarial manipulation, has shown significant improvements in adversarial input detection [13], leading to better classification.

Also, resilience of identification, localization, and tracking with respect to adversarial inputs is a key problem. For example, a spoofed sensor may advertise an incorrect target location measurement, leading to a failure in localization. It is important for distributed sensors to jointly decide if a subset of them were compromised. This was shown to be an NP-hard problem [18]. Polynomial solutions were developed for the case, where a "sufficient" number of sensors are present (by a metric of sufficiency defined in the original paper) [18]. The discovery informs sensor deployment strategies, where multiplicity is used as a means to combat complexity. Algorithms for ruggedizing learning and optimization against adversarial attacks have also been explored [25] [26], as well as algorithms for secure deep learning inference [33].

AI-based decision support agents must analyze the current state and recommend actions for commander's decision. A resilient formalism is needed to inform courses of action recommendations in the presence of adversarial action. A significant contribution of IoBT lies in developing such a formalism, called *non-cooperative inverse reinforcement learning* [27]. In traditional reinforcement learning [28], agents learn to maximize an objective function by executing well-designed actions to "prod" their environment and observing collected rewards from their actions. Conversely, inverse reinforcement learning [29] refers to the ability of an agent to infer the objective function of another by observing their behavior. This is often compared to learning by apprenticeship [30]. Decision-making in IoBT shares with inverse reinforcement learning the need to understand the objectives of another (namely, the adversary) in order to optimize one's own actions. However, the classical inverse reinforcement learning formalism is not designed for situations where an adversary might purposely obfuscate their intent by executing actions that lead an observer to an incorrect conclusion. To address such situations and ensure resilience with respect to adversarial deceit, the formalism of non-cooperative inverse reinforcement learning was recently proposed [27]. The work offers foundations for developing optimal policies in interacting with strategic adversaries who obfuscate their objectives and solutions that increase the ability to tolerate such obfuscation.

3.1 Intelligence at the Point of Need

A key objective of IoBT research is to enable automation that allows scalable, high-tempo execution of stages of an effect loop that connects sensing and effects when human cognitive capacity fails to keep up. An important question is: how can the freedom (or burden) of initiative be properly apportioned among humans and machines to most effectively attain mission goals, and most judiciously manage risk? How much delegation of authority in decision-making is appropriate to different levels of automation in the effect loop? Importantly, how to offer intelligence where it is needed, when it is needed, without the necessity to rely on expensive resources (including humans) that may not be readily available (or may be hard to contact) in a given dynamic situation. Elements of the effects loop should locally support and enhance team performance and human decision making in unpredictable, fast-paced environments.

The IoBT must ameliorate the inherent trade-off between resource cost of modern identification, localization,

and tracking algorithms on one hand and their quality of results on another. How can one attain both resource economy and quality at the same time? This is an important question for IoBT systems, where data transfer latency might be prohibitively high, yet individual sensors and processors in the field, at the point of need, have size, weight, and power (SWaP) requirements that make it challenging to run complex tasks locally. Recent work attains a significant reduction in inference model size [14][15], as well as model execution latency [16], allowing it to execute on a resource-constrained platform with nearly no sacrifice in quality.

As more processing is moved to the edge, the need for uncertainty awareness in the autonomy becomes critical for decision support. That is, the trade-off between latency, energy and performance continues to exist at the edge and must be encoded in the information produced by automated stages of the effect loop. Although uncertainty quantification research has made general advances in the field, the ability to provide this uncertainty estimation at the edge still remains a challenge. Recent work has developed sparsity and distillation-based methods for compressing expensive Monte Carlo posterior uncertainty computations into resource constrained neural network models [39]. This edge-based uncertainty has implications on real-time multimodal data fusion and commander decision making.

Automation of the effect loop can also include automatic generation of different courses of action for a decision maker, and perhaps action recommendations with associated pros and cons. A more aggressive category of delegation is when IoBTs are endowed with sufficient autonomy to automate their own decision-making and triggering (of effects) in line with commander intent. Game theoretic techniques for goal decomposition can be used in this scenario to assign local objectives to subordinate components while retaining assurances on global end results. As discussed in preceding sections, recent IoBT advances allow limited delegation of authority, while significantly reducing cost and improving outcomes on resulting behavior. This offers new options for bringing intelligence to the point of need and conducting operations, when the scale and tempo of the situation exceeds human response bandwidth.

4.0 INTEGRATED FIELD EVALUATION

The work performed in this IoBT research provides foundational theories and algorithms that help advance performant and resilient real-time services for the battlefield. A key need towards successful defense modernization is to accelerate the pipeline from research lab innovations to warfighter tools and enablers. We seek to address this need through extensive field evaluation of an IoBT that consists of the integrated components previously mentioned that address real time processing, adversarial resilience, and intelligence at the point of need.

One fundamental research question behind evaluation is, how can empirical basic research be executed at IoBT-scale, in a distributed fashion, with real-time access, and implicitly target a variety of cyber-physical effects? Further, how can this be achieved in a military research environment? The concept of a Distributed Virtual Proving Ground (DVPG) was developed to advance IoBT research and enable rapid integrated empirical experimentation. The DVPG provides a widely distributed environment where thousands or more heterogenous intelligent devices and autonomy transparently interoperate within a common operating environment dynamically composed to achieve mission effect-driven decisions. Concretely, the DVPG is a federation of experimental testbeds with the capabilities necessary to allow simultaneous, virtualized experimentation, i.e., we are not limited to be at a physical location or physically with systems/assets, but rather interaction is defined by interfaces, standards and processes.

The virtualized experimentation provided by the DVPG has already been demonstrated, showing integrated

research capabilities, remote connectivity, and continuous integration and testing between NATO partners. In this demonstration, sensor feeds from an U.S. Army facility in New Mexico were ingested by NCIA's VISTA environment integrating with their hyperledger-based trust algorithms and weapons detection object detector algorithms. Furthermore, the DVPG was used in a CWIX 2021 integration exercise, facilitating remote capabilities and enabling broader participation given pandemic-related travel restrictions. The DVPG also incorporates scenario playback, including the NATO developed scenario called ANGLOVA [40] that includes up to 283 nodes over several hours of movement. With ANGLOVA, the DVPG was able to stress communication links between nodes and functions to create realistic tactical network effects. Furthermore, playback of the ANGLOVA data allowed new basic research algorithms to be inserted in the experimental process where the effect of these innovations could be observed, assessed, and also recorded creating a new modified dataset for future research.

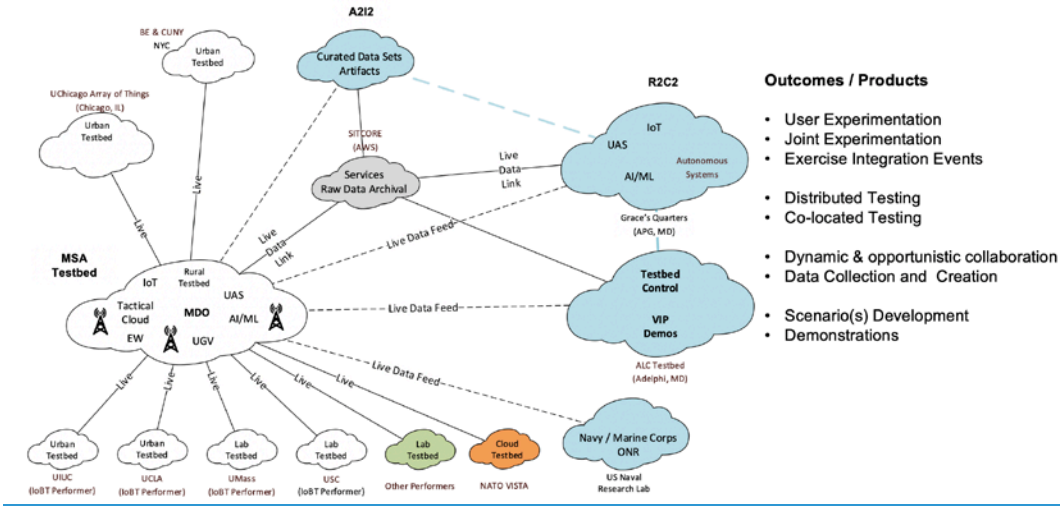

Figure 2: Notional architecture of the Distributed Virtual Proving Ground (DVPG).

Shown notionally in Figure 2, the DVPG architecture consists of a number of connected nodes, where each node represents a different experimental testbed that consists of sensing and/or processing capabilities. There are 3 capstone nodes that provide extended capabilities. The ARL Robotics Research Collaboration Campus (R2C2) node hosts a wide range of robotics and autonomy-centric experimental capabilities. The Army Artificial Intelligence Innovation Institute (A2I2) hosts capabilities that address the U.S. Army scientific challenges associated with artificial intelligence by providing an environment for accelerating the fundamental understanding, development, validation, and transition of AI for the Army. In addition to machine learning development capabilities, the A2I2 also provides a large repository for datasets and data storage. One hub of the DVPG is ARL's Multi-purpose Sensing Area (MSA) node which consists of 51 towers that stand 30ft high with 15-20 heterogenous sensors mounted to each tower. The MSA covers a 14km x 40km area that depicts natural characteristics for contested austere environments and has both edge and command post system processing capabilities. For the experimental evaluation of the IoBT research, we also leverage the connection between the MSA and IoBT CRA university partner nodes that define various lab testbeds with diverse processing clusters.

Field evaluation scenarios are designed to incorporate sensors, processors and platforms that are similar to what might be found in the battlefield to accomplish tasks necessary to deliver multi-domain operation effect loop processing for C2. In these scenarios, we leverage the DVPG that links sensors and processors from multiple

locations that are geographically dispersed. This enables the evaluation phase to address challenges of scale by connecting existing infrastructure to emulate, for example, a large urban city with vast numbers of smart devices that could be used for opportunistic sensing during a mission. We make use of unattended ground sensors and mobile autonomous agents (both ground and aerial) to capture data for tasks related to target detection and localization. The heterogeneity of the platforms provides opportunities to evaluate the battlefield services across varying degrees of processing power, highlighting the true innovations of edge-based processing in the cases where no GPUs are available. Information gathered during the stages of the effect loop are delivered to a common operating picture (COP) where a human teammate can use this information to make decisions and issue commands given the current state of the environment and mission.

During these scenarios, target detection is performed using a variety of sensing modalities, showcasing the ability to address the sensing heterogeneity inherent in the battlefield environment. In addition to evaluating target detection performance with these approaches, field evaluation provides opportunities to evaluate (i) energy usage, where less energy can extend battery life of the platforms to engage in longer missions, (ii) reduced data transmission sizes which could allow for longer range communication between platforms and the command post, (iii) detection speed improvements through unconventional, opportunistic, or non-line of sight sensing that stems from multimodal processing, and (iv) how fast information can be shared and the additional certainty associated with to detection to speed up and improve commander decision making.

5.0 SUMMARY AND FUTURE WORK

The research threads discussed above collectively resulted in significant advances in performance and resilience of executing sensor-to-effects loops. On the performance side, these advances included (i) several fold reductions in detection time of objectives using sensor-side innovations that reduce the cost of running deep neural networks on low-end devices, (ii) new feasibility limits on quickest change detection in adversarial contexts, (iii) a 3-4x reduction in the cost of intelligent data fusion uncertainty estimation, and (iv) algorithms for fast speculative inference in the presence of multiple sensing modalities. A new paradigm was developed for deep neural network inference in the frequency domain that reduces inference cost. Over 1000x improvement was achieved in compression to reduce distributed learning cost.

On the resilience side, innovations included (i) algorithms to increase tolerance to attacks on sensors, (ii) risk aware placement to increase resilience to resource outages, (iii) techniques for improved closed-loop robustness to variability in run-time computational platform latency, (iv) Byzantine tolerance and provably robust distributed optimization, and (v) various advances in secure deep learning inference, robustness to out-of-distribution data, safe reinforcement learning, and high-confidence generalization.

In accordance with the above three guiding principles (time is a weapon, IoBT is a fighting network, and intelligence at the point of need), further advances are sought along three general thrusts, with distributed system architecture and integrated experimentation for evaluation as a cross-cutting issue:

Hyper-performance. Novel computational paradigms are sought that allow the IoBT to break performance and scale barriers, giving rise to highly resource-economical, hyper-specialized, adaptive, and stealthy functions to be implemented in connecting sensors and triggers at minimum resource cost while embracing distribution, heterogeneity, and scale.

- **Resilience in Distributed Persistently Transient Systems.** Novel solutions are investigated for offering resilience in dynamic environments, where systems operate in a perpetual state of change,

while at the same time offering quality guarantees.

- **Knowledge-directed Distributed Intelligence.** Bringing machine intelligence to the field requires significant improvement in the efficiency of machine interference. This improvement can be attained by developing new paradigms for bridging scientific physical models with neural-network-based machine inference.

The content described in this paper are critical to dominance in the future battlespace, which will consist of active enemy, friendly, and civilian information-driven resources capable of affecting the physical world, where adversarial deception and cyber and electromagnetic activities will be the norm, resource ownership and other boundaries will be diverse and transient, and the operational environment – whether megacities or austere environments – will be dynamic. Militaries must be able to counter this proliferation of intelligent battlefield technologies, exploit them where possible, and dynamically deploy IoBT technologies to support warfighters conducting diverse missions.

The key scientific goal of the Internet of Battlefield Things Collaborative Research Alliance is to develop a fundamental understanding of dynamically composable, adaptive, goal-driven IoBTs to enable distributed analytics and intelligent battlefield services, exploiting 1000's of disparate assets. This will provide the theoretical foundations for understanding complex tactical, intelligent systems-of-systems composed of sensing, actuators, and analytics that are capable of continuous machine learning, counter-detection, self-protection, and performance-assuring behaviors on behalf, and in defense, of warfighters.

REFERENCES

[1] Abdelzaher, T., Taliaferro, A., Sullivan, P., & Russell, S. "The multi-domain operations effect loop: from future concepts to research challenges." In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications II*, vol. 11413, Int. Society for Optics and Photonics, 2020.

[2] Russell, S., Abdelzaher, T., & Suri, N. "Multi-Domain Effects and the Internet of Battlefield Things," In *IEEE Military Communications Conference*, pp. 724-730. 2019.

[3] Rovatsos, G., Moustakides, G.V., & Veeravalli, V.V. "Quickest Detection of a Dynamic Anomaly in a Sensor Network." In *IEEE Asilomar Conf. on Signals, Systems, and Computers*, pp. 98-102 2019.

[4] Rovatsos, G., Zou, S., & Veeravalli, V.V. "Quickest Detection of a Moving Target in a Sensor Network." In *IEEE Int. Symposium on Information Theory*, pp. 2399-2403. 2019.

[5] Rovatsos, G., Veeravalli, V.V., Towsley, D., & Swami, A. "Quickest Detection of Growing Dynamic Anomalies in Networks." In *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*. 2020.

[6] Zou, S., Veeravalli, V.V., Li, J., Towsley, D. & Swami, A. "Distributed Quickest Detection of Significant Events in Networks." In *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*. 2019.

[7] Li, J., Towsley, D., Zou, S., Veeravalli, V.V., & Ciocarlie, G. "A Consensus-based Approach for Distributed Quickest Detection of Significant Events in Networks." In *IEEE Asilomar Conf. on Signals, Systems, and Computers*. 2019.

[8] Bu, Y., Lu, J., & Veeravalli, V.V. "Model change detection with application to machine learning." *In*

IEEE Int. Conf. on Acoustics, Speech and Signal Processing, pp. 5341-5346. 2019.

[9] Wilson, C., Bu, Y., & Veeravalli, V.V. "Adaptive Sequential Machine Learning." *arXiv preprint* arXiv:1904.02773 (2019).

[10] Bu, Y., Lu, J., & Veeravalli, V.V. "Active and Adaptive Sequential Learning with Per Time-step Excess Risk Guarantees." In *IEEE Asilomar Conf. on Signals, Systems, and Computers*, pp. 1606-1610. 2019.

[11] Papernot, N., McDaniel, P., Goodfellow, I., Jha, S., Celik, Z.B., & Swami, A. "Practical black-box attacks against machine learning." In *Asia Conf. on Computer and Communications Security*. 2017.

[12] Jha, S., Raj, S., Fernandes, S., Jha, S.K., Jha, S., Jalaian, B., et al. "Attribution-Based Confidence Metric for Deep Neural Networks." In *Advances in Neural Information Processing Systems*. 2019.

[13] Jha, S., Raj, S., Fernandes, S.L., Jha, S.K., Jha, S., Verma, G., Jalaian, B., & Swami, A. "Attribution-driven causal analysis for detection of adversarial examples." *arXiv preprint* arXiv:1903.05821 (2019).

[14] Yao, S., Zhao, Y., Zhang, A., Hu, S., Shao, H., Zhang, C., Su, L., & Abdelzaher, T. "Deep learning for the internet of things." *Computer* 51, no. 5 (2018): 32-41.

[15] Yao, S., Zhao, Y., Zhang, A., Su, L., & Abdelzaher, T. "DeepIoT: Compressing deep neural network structures for sensing systems with a compressor-critic framework." In *ACM Conf. on Embedded Network Sensor Systems*, pp. 1-14. 2017.

[16] Yao, S., Zhao, Y., Shao, H., Liu, S., Liu, D., Su, L., & Abdelzaher, T. "Fastdeepiot: Towards understanding and optimizing neural network execution time on mobile and embedded devices." In *ACM Conf. on Embedded Networked Sensor Systems*, pp. 278-291. 2018.

[17] Deshmukh A., Liu J., Veeravalli, V.V., & Verma, G. "Information Flow Optimization in Inference Networks," In *Int. Conf. on Acoustics, Speech and Signal Processing*, 2020.

[18] Mao, Y., Mitra, A., Sundaram, S., & Tabuada, P. "When is the Secure State-Reconstruction Problem Hard?" In *IEEE Conf. on Decision and Control*, pp. 5368-5373. 2019.

[19] Lee, J., Marcus, K., Abdelzaher, T., Amin, M.T.A., et al. "Athena: Towards decision-centric anticipatory sensor information delivery." *Journal of Sensor and Actuator Networks* 7, no. 1 (2018)

[20] Neglia, G., Calbi, G., Towsley, D., & Vardoyan., G. "The Role of Network Topology for Distributed Machine Learning." In *IEEE Conf. on Computer Communications*, pp. 2350-2358. 2019.

[21] Bu, Y., Zou, S., & Veeravalli, V.V. "Tightening mutual information based bounds on generalization error." In *Int. Symposium on Information Theory*, pp. 587-591. 2019.

[22] Bu, Y., Gao, W., Zou, S., & Veeravalli, V.V. "Information-Theoretic Understanding of Population Risk Improvement with Model Compression." *arXiv preprint* arXiv:1901.09421 (2019).

[23] Basu, D., Data, D., Karakus, C., & Diggavi, S. "Qsparse-local-SGD: Distributed SGD with Quantization, Sparsification and Local Computations." In *Advances in Neural Information Processing Systems*. 2019.

[24] Singh, N., Data, D., George, J., & Diggavi, S. "SPARQ-SGD: Event-Triggered and Compressed Communication in Decentralized Stochastic Optimization." *arXiv preprint* arXiv:1910.14280 (2019).

[25] Data, D., Song, L., & Diggavi, S. "Data Encoding Methods for Byzantine-Resilient Distributed Optimization." In *Int. Symposium on Information Theory*, pp. 2719-2723. 2019.

[26] Data, D., & Diggavi, S. "Byzantine-Tolerant Distributed Coordinate Descent." In *Int. Symposium on Information Theory*, pp. 2724-2728. 2019.

[27] Zhang, X., Zhang, K., Miehling, E., & Basar, T. "Non-Cooperative Inverse Reinforcement Learning." In *Advances in Neural Information Processing Systems*, pp. 9482-9493. 2019.

[28] Kaelbling, L.P, Littman, M.L., & Moore, A.W. "Reinforcement learning: A survey." *Journal of Artificial Intelligence Research* 4 (1996): 237-285.

[29] Ng, A.Y., & Russell, S.J. "Algorithms for inverse reinforcement learning." In *Int. Conf. on Machine Learning*, pp. 663-670. 2000.

[30] Abbeel, P., & Ng, A.Y. "Apprenticeship learning via inverse reinforcement learning." In *Int. Conf. on Machine learning*. 2004.

[31] Rovatsos, G., Veeravalli, V.V., Towsley, D., & Swami, A. "Quickest detection of anomalies of varying location and size in sensor networks." *IEEE Transactions on Aerospace and Electronic Systems* (2021).

[32] Yao, S., Li, J., Liu, D., Wang, T., Liu, S., et al. "Deep Compressive Offloading: Speeding Up Edge Offloading for AI Services." *GetMobile: Mobile Computing and Communications* 25, no. 1 (2021).

[33] Liu, R., Garcia, L., Liu, Z., Ou, B., & Srivastava, M. "SecDeep: Secure and Performant On-device Deep Learning Inference Framework for Mobile and IoT Devices." In *Int. Conf. on IoT Design & Implementation*. 2021.

[34] Li, T., Huang, J., Risinger, E., & Ganesan, D. "Low-latency speculative inference on distributed multi-modal data streams." In *Int. Conf. on Mobile Systems, Applications, and Services*, pp. 67-80. 2021.

[35] Huang, J., Ganesan, D., Marlin, B., & Kwon, H. “CLIO: Enabling automatic compilation of deep learning pipelines across IoT and Cloud.” In *Int. Conf. on Mobile Computing and Networking*. 2020.

[36] Liu, S., Yao, S., Fu, X., Tabish, R., Yu, S., Bansal, A., Yun, H., et al. “On removing algorithmic priority inversion from mission-critical machine inference pipelines.” In *Real-time Systems Symposium*.” 2020.

[37] Hu, Y., Liu, S., Abdelzaher, T., Wigness, M., & David, P. “On exploring image resizing for optimizing criticality-based machine perception.” In *Int. Conf. on Embedded and Real-Time Computing Systems and Applications*. 2021.

[38] Jha, S., Rushby, J. & Shankar, N.. “Model-centered assurance for autonomous systems.” In *Int. Conf. on Computer Safety, Reliability and Security*. 2020.

[39] Vadera, M., Jalaian, B., & Marlin, B. “Generalized Bayesian Posterior Expectation Distillation for Deep

Neural Networks." In *Uncertainty in Artificial Intelligence*. 2020.

[40] Suri, N., Nilsson, J., Hansson, A., Sterner, U., Marcus, et al. "The Angloval Tactical Military Scenario and Experimentation Environment." In *Int. Conf. on Military Comm. and Information Systems*. 2018.