ARL-SR-0390● Dec 2017

US Army Research Laboratory

2015 Army Science Planning and Strategy Meeting Series: Outcomes and Conclusions

by Joseph Mait, Dawanne Poree, John Prater, Peter Reynolds, David Stepp, Bruce J West, Alex Kott, Ananthram Swami, Arwen H DeCostanza, Piotr J Franaszczuk, Kaleb McDowell, Brett Piekarski, Brian M Sadler, Rob Carter, and Jeff Zabinski

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer’s or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

ARL-SR-0390 ● DEC 2017

US Army Research Laboratory

2015 Army Science Planning and Strategy Meeting Series: Outcomes and Conclusions

by Joseph Mait
Office of the Director

Dawanne Poree, John Prater, Peter Reynolds, David Stepp, and Bruce J West
Army Research Office

Alex Kott and Ananthram Swami
Directorate, ARL

Arwen H DeCostanza, Piotr J Franaszczuk, and Kaleb McDowell
Human Research and Engineering Directorate

Brett Piekarski
Sensors and Electron Devices Directorate

Brian M Sadler
Vehicle Technology Directorate

Rob Carter and Jeff Zabinski
Weapons and Material Research Directorate

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YYYY)* December 2017	**2. REPORT TYPE** Special Report	**3. DATES COVERED (From - To)** November 2015–January 2016

4. TITLE AND SUBTITLE 2015 Army Science Planning and Strategy Meeting Series: Outcomes and Conclusions	**5a. CONTRACT NUMBER**
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S) Joseph Mait, Dawanne Poree, John Prater, Peter Reynolds, David Stepp, Bruce J West, Alex Kott, Ananthram Swami, Arwen H DeCostanza, Piotr J Franaszczuk, Kaleb McDowell, Brett Piekarski, Brian M Sadler, Rob Carter, and Jeff Zabinski	**5d. PROJECT NUMBER**
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) US Army Research Laboratory ATTN: RDRL-HR Aberdeen Proving Ground, MD 21005-5425	**8. PERFORMING ORGANIZATION REPORT NUMBER** ARL-SR-0390
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	**10. SPONSOR/MONITOR'S ACRONYM(S)**
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Under the direction of the Assistant Secretary of the Army for Acquisition, Logistics, and Technology [ASA(ALT)], the US Army Research Laboratory (ARL) hosted a series of meetings November 2015–January 2016 to develop a strategic vision for Army Science. Meeting topics were vetted through the ARL Director and approved by the ASA(ALT). Their selection was based on their potential to dramatically impact military capabilities in the long term. This report is a summary of those meetings and their outcomes. The 6 topics selected were the Internet of Battlefield Things, Cyber Fog, Individualizing Technology for Effective Teaming, Distributed and Collaborative Intelligent Systems, Microscale Adaptability, and Expeditionary On-Demand Manufacturing. These 6 areas are conveniently thought of in terms of 3 broader thrusts—Cyber, Human, and Materials Sciences—which map onto 3 three domains of conflict: Virtual, Human, and Physical. Questions considered at these meetings included, “Within this technical area, what capability can it deliver to the military 25 years from now? What technical hurdles exist that limit our ability to realize this capability? What research does the Army need to support now to overcome these hurdles and enable the desired capability?”

15. SUBJECT TERMS

technology, strategy, planning, cyber, robotics, manufacturing

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON** Joseph Mait
a. REPORT Unclassified	**b. ABSTRACT** Unclassified	**c. THIS PAGE** Unclassified	UU	90	**19b. TELEPHONE NUMBER (Include area code)** 410-278-1453

Contents

List of Figures

Acknowledgments

The authors gratefully acknowledge the support of the Assistant Secretary of the Army for Acquisition, Logistics, and Technology, the US Army Research Laboratory (ARL) Director's Office, the on-site contractors who helped with registration, and the many participants, all of whom contributed to an intriguing, productive, and successful set of meetings.

Mr Gabe Smith and Ms Tammy Christenson, ARL Office of the Director, deserve special recognition for their administrative support to all of the organizers and all of the meetings.

Executive Summary

Under the direction of the Assistant Secretary of the Army for Acquisition, Logistics, and Technology (ASA[ALT]), the US Army Research Laboratory (ARL) hosted a series of meetings in fall 2015 through winter 2016 to develop a strategic vision for Army Science. Meeting topics were vetted through the ARL Director and approved by the ASA(ALT). Their selection was based on their potential to dramatically impact military capabilities in the long term. This report is a summary of those meetings and their outcomes.

The 6 topic areas selected were the Internet of Battlefield Things, Cyber Fog, Individualizing Technology for Effective Teaming, Distributed and Collaborative Intelligent Systems, Microscale Adaptability, and Expeditionary On-Demand Manufacturing. These 6 areas are conveniently thought of in terms of 3 broader thrusts—Cyber, Human, and Materials Sciences—that map onto the 3 domains of conflict: Virtual, Human, and Physical. Questions considered at these meetings included, "Within this technical area, what capability can it deliver to the military 25 years from now? What technical hurdles exist that limit our ability to realize this capability? What research does the Army need to support now to overcome these hurdles and enable the desired capability?"

Meeting organizers were all senior members of ARL's technical staff and included ARL Branch and Division Chiefs, elected Fellows, and Army STs. For each meeting, ARL invited a small number of world-class experts as speakers who have a long-term, broad view of a specific area and an awareness of its trends. The meetings were structured to obtain a variety of viewpoints, not just near-term, Department of Defense-related expertise. Target attendance per meeting was roughly 25. Attendees spanned a large variety of research and development organizations, both civilian and defense.

- The objective of the "Microscale Adaptability" workshop was to explore the integration of new molecular building blocks and novel assembly and processing techniques to generate new materials with unprecedented responsive and reconfigurable properties and manipulate light–matter interactions. The conclusion was that advancements in computational modeling and nanoscale characterization tools will be key to enabling each of these opportunities. Energy management within materials systems, as a means of driving the assembly and reconfiguration processes, is ripe for exploration.

- The objective of the “Expeditionary On-Demand Manufacturing” workshop was to determine what research is needed to enable the Army to rapidly build complex systems in austere environments. Four areas were explored: Rapid Certification, Moving Logistics Forward, Hybrid Manufacturing, and Bioproduction. The conclusion was that a fresh start is needed; new manufacturing processes, new specifications and standards, new data standards for engineering drawing and data packages, and new design methodologies are needed to take full advantage of the foreseeable new capabilities. Additional research is needed to discover design principles needed for programming advanced genetic circuits to provide the desired responses.

- The objectives of “The Internet of Battlefield Things” (IoBT) workshop were to examine the theoretical foundations of the Internet of Things concept in the context of Army battlefield operations. IoBT potentially leads to significant improvement of situational awareness. The workshop focus was on communications and processing in the massively complex and dynamic IoBT network and the effects of IoBT on the warfighter and the enemy. The conclusion was that success in fighting the battle of IoBT cannot be assured without new theoretical explorations; it requires major, new results in large-dimensional game theory and new theory to formalize and normalize diverse definitions and conceptualizations of risks and uncertainty. Deception should be integral to this theoretical analysis, for example, in the level of complexity necessary to deceive successfully.

- The objective of “The Fog of Cyber War” workshop was to examine the theoretical foundations of the concept of fog in cyberspace for Army battlefield operations and how one might create a Cyber Fog that is transparent to friends and opaque to foes. The conclusion was that due to the extreme challenges and complexities inherent in Cyber Fog, fundamental advances in theories of deception and counterdeception, novel game theoretic approaches, information semantics, machine learning, and formal methods would be required—for example, using formal methods to execute and manage successful obfuscation of friendly information in a cyber fog.

- The objective of the “Individualizing Technology for Effective Teaming” workshop was to challenge beliefs about the development and design of military-technology teams and to identify novel adaptive and individualized approaches to such teaming. One conclusion was that shifting to the human–agent team paradigm will yield performance in the human–technology unit that has potential to advance exponentially, with critical implications for talent management and organizational effectiveness. Research is needed to develop theoretical understanding of the underpinnings of human variability, as well as the reliable and valid technological capabilities to estimate and predict real-world individual physiology and behavior.

- The objective of the “Distributed and Collaborative Intelligent Systems” workshop was to explore the key focus areas: Distributed Awareness, Distributed Intelligence, Adaptable and Resilient Control, and Scaling Experimental Complexity. Among the conclusions drawn was the need for a control architecture that simultaneously deals with both global and localized control of single agents, multiagent teams, and localized swarm behavior. In addition, current game-theory-based formulations are computationally intractable when dealing with teams involving hundreds of agents and must be adapted to distributed decision making.

Based on the output of the meetings and subsequent discussions among meeting organizers, ARL developed the following specific recommendations.

Research Recommendations

The ASA(ALT) should invest in the following programs:

- Energy management within materials systems, as a means of driving the assembly and reconfiguration processes, is ripe for exploration.
- Computational modeling and nanoscale characterization tools to enable efficient design of hybridized manufacturing; realtime, multiscale computational capability to enable predictive analytics for expeditionary on-demand manufacturing
- Discovery of design principles to enable programming advanced genetic circuits with specified functionality
- Formal theories of deception, counterdeception, discovery or rejection of deception, in the context of IoBT
- Novel approaches to information theory, appropriate for large, dense, dynamic heterogeneous networks, with nonergodic transient information flows
- Major new program in large dimensional (stochastic) game theory, where the number of players and the action space is very large and dynamic
- Distributed intelligence (e.g., via novel approaches to learning in an adversarial environment and under resource constraints)

- A program to develop formal methods to execute and manage successful obfuscation of friendly information in a cyber fog. Deception should be integral to this theoretical analysis, for example, in the level of complexity necessary to deceive successfully.
- Research is needed to develop theoretical understandings of the underpinnings of human variability as well as the reliable and valid technological capabilities to estimate and predict real-world individual physiology and behavior.
- Novel approaches to enable efficient high-fidelity simulation of a large number of heterogeneous entities (including mixed human–machine teams) operating autonomously in complex environments

1. Introduction

In 2013, the Assistant Secretary of the Army for Acquisition, Logistics, and Technology (ASA[ALT]) charged the US Army Research Laboratory (ARL) to develop a strategic research plan for the Army for the next 20 to 30 years. In accepting this charge, ARL recognizes that armed conflict remains a contest of wills and that actors in conflict fundamentally seek to persuade others to accept their perspective. ARL also recognizes that, historically, investments in science and technology focused on combat in the physical domain. However, as stated in the report of the first Army Science Planning and Strategy meetings in 2014, the means for this persuasion now and into the near future exist in 3 realms:

- physical, the domain of activities defined in space and time by the laws of physics;
- virtual (or informational), the domain of activities defined by thought and perception; and
- human (or cultural), the domain of activities defined by the interaction of people and societies.

See Fig. 1.

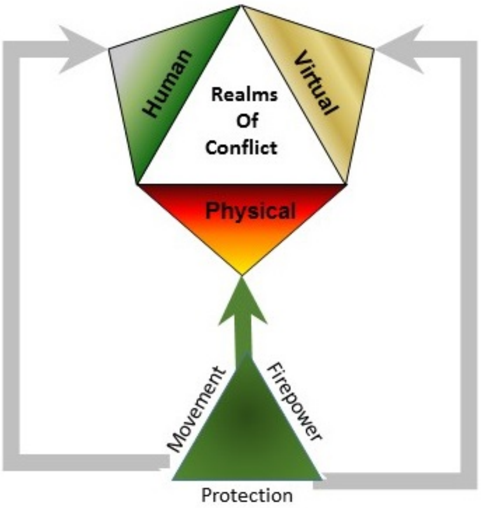

Fig. 1 Three realms of future conflict

ARL has focused its strategic intentions on these 3 realms and, in recent years, especially on overlapping areas such as mixed teams of humans and intelligent systems and operations in cyber space whose effects are realized in physical space. Four of the meetings held in November 2015–January 2016 focused on this overlap area. The 2 that did not instead addressed materials and material research, the traditional realm of the physical domain.

As described in the next section, ARL chose 6 meeting topics to consider in 2015–16, 2 each in materials (Microscale Adaptability and Expeditionary On-Demand Manufacturing), cyber operations (The Internet of Battlefield Things and Cyber Fog), and human–machine teaming (Distributed and Collaborative Intelligent Systems and Individualizing Technology for Effective Teaming). See Fig. 2.

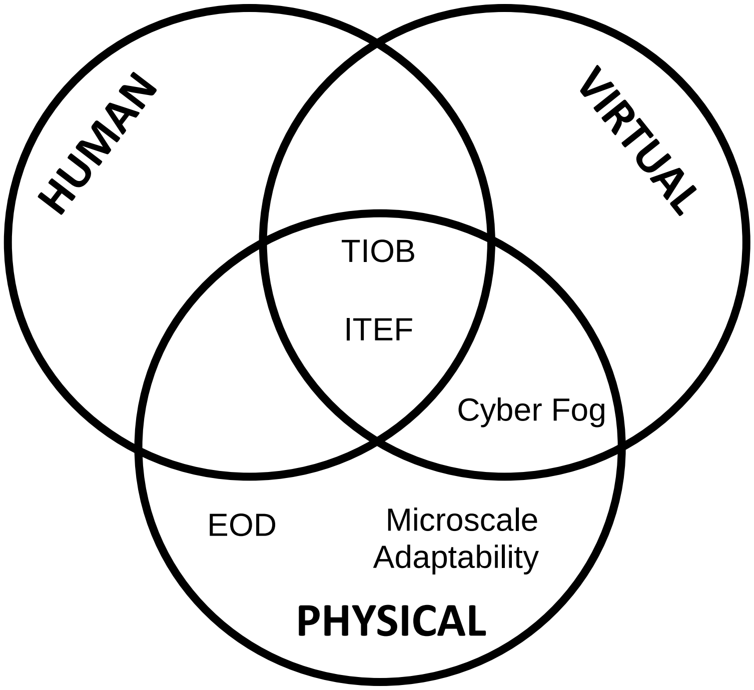

Fig. 2 Topic areas for planning meetings held in fall 2015; overlaps are representative and not exhaustive. Key terms are EOD: Expeditionary On-Demand Manufacturing; TIOB: The Internet of Battlefield; and, ITEF: Individualizing Technology for Effective Teaming.

The sustained focus on materials is justified based on the intimate link between materials and mankind's ability to exploit them. This is evident in the names of cultural epochs (e.g., stone, bronze, and iron ages). This relationship has not changed for millennia. Materials remain of fundamental importance to all aspects of the Army enterprise, especially lethality and protection. ARL believes the future of materials processing lies in adaptive materials and pushing materials processing and manufacturing closer to the point of application.

More recent cultural epochs reflect man's desire not just to exploit materials but also to harness energy in various forms. This includes steam and electricity, which drove industrialization. If one accepts that the transfer of information requires energy transfer, the information age continues man's desire to harness energy. The information age was enabled by advances in electronics and other microtechnologies, physics, and mathematics, and spawned hand-portable internetted sensors, communications, and computers. It is these technologies that provide the nexus for the 3 realms and justify our focus on capabilities in cyber space and on teaming between humans and intelligent systems.

During fall 2015 ARL conducted six 2-day meetings in these topic areas. Each meeting addressed the following questions: What capability can this technology deliver to the military 25 years from now? What technical hurdles exist that limit our ability to realize this capability? What research does the Army need to support now to overcome these hurdles and enable the desired capability?

The meetings were structured to obtain a variety of viewpoints, not just near-term, Department of Defense (DOD)-related expertise. Attendance was invitation-only and the number of attendees per meeting was roughly 25 but varied from meeting to meeting. ARL invited a small number of world-class experts as speakers from government, academia, and industry who have a long-term, broad view of a specific area and its trends.

2. Meeting Background

Meeting topics have been honed over several years. The process began in November 2012 with ARL's Visioneering 2050 workshop. Visioneering 2050 identified several research areas that had the potential to impact dramatically military capabilities in the long term. ARL chose 6 of these to explore in depth during the first Army Science Planning and Strategy Meetings (ASPSMs) in 2013. The first 6 meetings were Materials in Extreme Environments, Biological Sciences, Quantum Information and Sensing, Intelligent Systems, Information at the Tactical Edge, and the Human Dimension. Participants, primarily from academia, were asked to identify capabilities desired in the future, hurdles to their realization, and the research required to overcome these hurdles.

In a report to the ASA(ALT), ARL recommended multidisciplinary investigations in the following areas: hybrid biological and nonbiological systems, the integration of neuroscience and training, trust in information and intelligent systems, the mathematics and cognitive transformation of data to information, and intelligent platforms as personal advisors and personal assistants. Further, the report recommended increased in-house efforts in quantum information and sensing in conjunction with the establishment of an off-base lab or a joint research institute with a university. Consequently, ARL's investment in Quantum Information Science is now close to $15 million and includes collaborative efforts among ARL, the US Naval Research Laboratory, the National Institute of Standards and Technology, and the University of Maryland.

In the 2014 set of meetings, as opposed to considering technologies ARL instead considered future military capabilities in logistics, maneuver, and cyber. Although the goal of the meetings remained to identify research required to enable the desired capability, the participants were primarily from the Training and Doctrine Command's Army Capabilities Integration Center and Centers of Excellence and sought to identify capability "pull" rather than technology (science) "push." See Fig. 3.

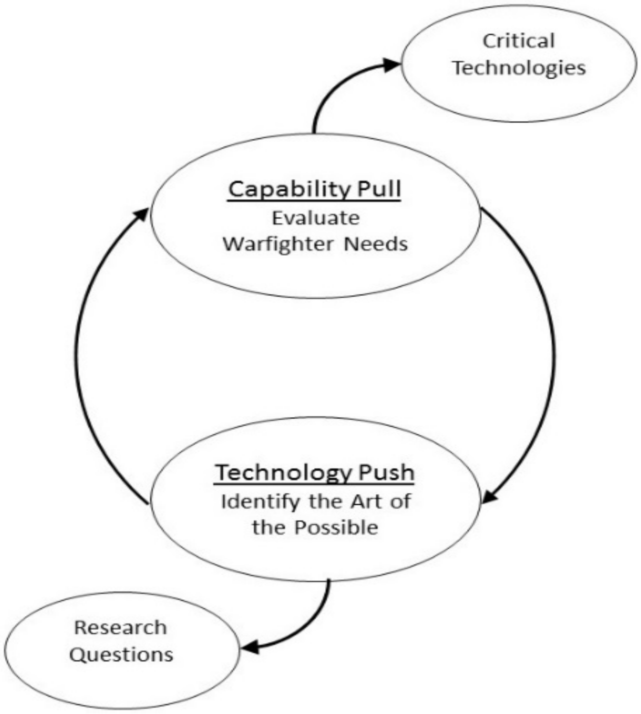

Fig. 3 Research-strategy-definition process delineating the difference between capability push and technology pull

The overwhelming message from these meetings was the need to develop the capacity for operations in the 3 realms. In addition, attendees across all meetings identified small autonomous platforms operating in a coordinated fashion as a critical technology for the Air–Land–Sea battle. Collaborative operation of such systems can enhance the military's capabilities in lethality, sensing, and communication. The meetings spawned other notional concepts of operations including "plug into the city" (i.e., exploit available infrastructure in urban terrain for resources and data collection), "embrace our vulnerabilities" (i.e., accept that we will be vulnerable but insure that we are less vulnerable than our adversaries), and "just-add-water protection" (i.e., increase logistics capacity without sacrificing protection by shipping high-density, low-volume raw materials, or using indigenous materials, that are prepared on site to provide protection).

These notions were distilled to arrive at the topics addressed in 2015. For example, "plug into the city" evolved into the Internet of Things and "embrace our vulnerabilities" into Cyber Fog. The emphasis on autonomous systems led ARL to consider both collections of mobile intelligent systems and a heterogeneous collection of intelligent systems and humans. "Just add water protection" evolved into Expeditionary On-Demand Manufacturing.

The next section summarizes the discussions that occurred within each meeting. In Section 3 we highlight topics that overlapped 2 or more of the meetings. We also consider areas not covered in this initial set of 6 meetings and end with summary recommendations.

3. Discussion

3.1 Microscale Adaptability

Recent breakthroughs in materials, chemistry, and materials assembly now enable opportunities for revolutionary processing approaches (i.e., control, speed, and cost) and unique materials (i.e., bulk responsive and adaptive properties that have not previously been possible). These breakthroughs provide a superb opportunity to create new molecular building blocks and manufacturing techniques that will generate materials with new responsive and reconfigurable properties. In particular, the recent demonstration of rationally designed mechano-responsive optical materials suggests new paradigms for optical materials that exhibit "smart" properties and robust tunability are now possible. The objective of the Microscale Adaptability meeting was to explore the integration of new molecular building blocks and novel assembly and processing techniques to generate new materials with unprecedented responsive and reconfigurable properties and manipulate light–matter interactions.

To achieve this objective, the meeting sought to bring together 2 disparate research communities—responsive self-assembled materials and optical metamaterials—to identify breakthrough strategies (both theoretical and experimental) that would facilitate the design and robust self-assembly of multicomponent 3-D structures with precisely engineered electronic and optical properties. A diverse group of academic and government scientists was invited to participate in the workshop. They were tasked to identify new scientific opportunities that may be possible by converging the fields of responsive self-assembled materials and optical metamaterials and the key barriers to achieving these opportunities and explain how success of convergence may enable future Army capabilities. Engaging in interactive, cross-disciplinary, small-group discussions, the workshop participants identified the following promising areas of opportunity.

3.1.1 Self-Assembly of Reconfigurable Elements

Recent studies have demonstrated the feasibility of using tailored interactions with small molecules to drive the assembly of inorganic nanoparticles into larger structures spanning multiple-length scales. Similarly, dynamic assembly processes have been studied in which the bonding and network structures can undergo disassembly and reassembly in route to their final structures. By extension, one can imagine the insertion of reconfigurable elements (building blocks capable of dynamically altering their configuration), varying their coupling interaction, and even switching between physical states potentially enabling one to synthesize "smart" materials that possess the ability to change their properties in response to external command signals or environmental cues. Such systems will require the development of techniques for staging individual attachment steps into a sequential assembly that leads to complex and/or hierarchical architectures with unique properties. Additionally, avenues for the capture, conversion, and/or transduction of various forms of energy should be incorporated into the design as a means of driving the assembly and/or reconfiguration processes. This concept of energy management within a materials system was noted by the workshop participants as an area ripe for exploration. It was also noted that self-assembly approaches seem to be inherently susceptible to the formation of defects that can greatly degrade the mechanical and functional performance of the materials. Over the millennia, living systems have evolved mechanisms for countering the effects of defects. Similar approaches to defect tolerance and functional redundancy need to be perfected for man-made materials.

3.1.2 Responsive Behavior

A daunting challenge toward synthetic "living" systems is predictably propagating a molecular-level change, generated through the selective sensing of a trigger, into a readily discernible macroscopic change in a material's fundamental properties. This can only be addressed by developing a fundamental understanding of the chemical processes that occur at multiscale levels to enable active control that spans from molecular to nanoscale to macroscopic length scales and from nanoseconds to hours. The inherent complexity involved in connecting these length scales, and the transduction (detection, amplification and propagation) of external signals into macroscopic responses, requires a cohesive, multidisciplinary approach.

3.1.3 Engineered Interfaces

Access to the interfaces at all levels in these materials systems was called out by the group as a particularly compelling opportunity for the field. Opportunities for manipulating the assembly process by using aspects of shape, intermolecular interactions, induced conformation changes, functionalized adduct and site-specific binding groups, molecule-to-substrate interactions, and external fields need to be fully developed. An intriguing new development is the concept of pluripotent matter in which changes in bonding are driven by the surface functionalization of the basic building blocks, which is sensitive to external signals including changes in the external environment. Similarly, liquid infiltration into voids within percolating porous systems (e.g., opal and inverse opal systems) is another broad approach available to radically alter the optical properties of a system via changes in index matching. Another interesting possibility is the incorporation of particle jamming at interfaces in 2-phase liquid systems, which could impart on-demand rigidity into these systems. Finally, methods for effectively interfacing various biological and inorganic elements into complex functioning architectures capable of maintaining biological functionality need to be fully investigated.

Advancements in computational modeling and nanoscale characterization tools will be key to enabling each of these opportunities. More specifically, new theoretical tools and computational methods capable of modeling the self-assembly process and identifying valid self-assembly pathways that lead to stable hierarchical architectures and desired functionality are needed. Additionally, the theory needs to incorporate robust reconfigurability and ultimately predict the range of dynamic behavior that can be achieved in these systems; furthermore, these predictions need to be experimentally validated. Regarding analytical characterization, the development of quantitative, high-resolution techniques for spatiotemporal characterization and methods that will advance our understanding of the fundamental structure–function relationships at the molecular and nanoscales as well as across scales are a high priority.

3.2 Expeditionary On-Demand Manufacturing

The advances in materials and manufacturing sciences are generating new capabilities for producing materiel via new production routes that will require the Army to evaluate and change how it operates. The impact will enable threat-responsive evolution of materiel and produce rapidly responsive logistics. Expeditionary manufacturing is the area of research using materials and manufacturing technologies, to include rapid certification, synthesis, processing, and microstructural control of advanced materials, multimaterial and multifunctional structures, and biological production, to fabricate materiel with new capabilities—and includes the ability to harvest materials from the battlefield environment.

To achieve these goals, significant research is needed to mature the fundamental materials science, processing and manufacturing sciences, design methodologies, data management, and specifications and standards to enable them to be useful to the soldier. Four topics were chosen to focus the discussion on what research is needed to enable: Rapid Certification, Moving Logistics Forward, Hybrid Manufacturing, and Bioproduction.

3.2.1 Rapid Certification

Rapid certification is essential to being able to use the materiel being fabricated. Current 3-D printing, which also is known as additive manufacturing (AM), has captured the imagination of engineers and tinkerers; but, in reality, there very few commercialized structural components being produced by this method. There are several reasons for this with the first 2 being cost and certification of the parts from the new process. Cost is high since the processes uses expensive feedstocks; also, there are high equipment costs and slow production rates for high-volume manufacturing. AM's strength is the complexity now enables new design spaces that were previously unmanufacturable. Certification is being driven by advanced materials and computational science's ability to predict properties for a given additive manufacturing process and to certify the part will be able to function as intended. Significant research is needed to enable in-line (where the certification is concurrent with the fabrication process) certification. New material feedstocks designed for AM processes, new manufacturing processes, process models, verified computational models, new specifications and standards, new data standards for engineering drawing and data packages, and new design methodologies are needed to take full advantage of the new capabilities that are foreseen in this area.

3.2.2 Moving Logistics Forward

The implications of expeditionary manufacturing to reduce logistics were discussed. An analysis of current AM technologies conducted by the Army Logistics Innovation Agency shows that building new parts does not compete with traditional manufacturing on cost or on reducing logistical footprint, but AM does provide a more timely response. This trade, based on current equipment and materials, is anticipated to change significantly in the next 30 years.

Logistics implications for AM are muddied by the complexity of the current infrastructure and business practices of the Army and DOD as a whole. The Defense Logistics Agency currently maintains millions of parts to support DOD materiel, but 99% of those parts have paper drawings, which are of limited use to AM technologies. Each of those drawings would have to be converted to a digital technical data package. That assumes each of those drawings contains the technical data necessary, such as tolerances, materials, material-processing requirements, and specifications. New science and technology (S&T) efforts to rapidly convert millions of parts to digital data, whether through image-recognition software or other information technologies, are needed.

It was noted the use of AM to make parts is limited to making existing parts and structures that are designed for traditional manufacturing processes. The idea of making more efficient parts is attractive but would require certification and a configuration-change exemption from the original equipment manufacturer.

3.2.3 Hybridized Manufacturing

Hybridized manufacturing is a general term for the fabrication of multiple materials simultaneously to produce a composite or hybrid structure that has increased functionality. This functionality could be structural, mechanical, magnetic, electrical, and so forth. Discussion in this area focused on manufacturing electronics of sufficient complexity in austere environments. While it was felt that high-end computational systems (central processing units) would be beyond the reach of manufacturing in austere environments, the idea of pick-and-place insertion of prefabricated components into a build of high-complexity systems would alleviate that issue.

Another area of discussion was on design tools. The new manufacturing technologies are creating new design spaces that are more complex than traditional design practices. Humans are not good at rapidly converging on efficient designs. New tools are needed to optimize the designs. Soldiers know what they need but lack the knowledge and skill to design and make it. A “PhD in a box” is needed to take the interfaces of what is needed and be able to design the appropriate part. Artificial intelligence or at least significant computational assets will be needed to enable this capability.

Other areas discussed were voxel-by-voxel printing in which each voxel is a functional element. An example was where each voxel was a battery, IR emitter, computational chip, or structural member. When properly assembled a TV remote control could be created. While a simplistic example, this shows that continued miniaturization of functional elements could lead to complex functionality from a toolkit of building blocks that could be assembled pending the desired application. The drawback, from this area of discussion, was that the Army currently is maintaining systems for 50 years—so, would these repair/replacement technologies be able to fix a current system when it is serviced 30 years from now?

3.2.4 Bioproduction

Research is needed to harness biology to convert indigenous feedstocks to materials, pharmaceuticals, chemicals, and even body parts. Synthetic biology will transition from a discovery phase to a readily programmable engineering technology that can rapidly evolve nutritional, vaccination, and biocompatible materials. Discussions indicated that advancements in synthetic biology will allow biosystems to be reprogrammed at the point of need to produce a wide range of products. Bioproduction is capable of producing molecules that are not tractable using standard chemical means; this is expected to extend polymer science beyond current limitations as novel polymer precursors and chemistries are explored. Controlled printing of cells will allow precise placement of cell-programmed cell types to produce 3-D gradient biomaterials such as organ replacements or complex manufacturing platforms. To realize the potential of on-site reprogrammable biomanufacturing, further development is needed to speed the time it takes to engineer microbes, select for desired properties, ensure stability of the system, and design systems to produce chemicals at large scale. The current state of the art is that synthetic DNA can be readily made and programmed, but future research is needed to discover the design principles needed for programming advanced genetic circuits to provide the response we desire.

3.2.5 Topics of Interest Not Covered

While much of the discussion focused on the potential for new technologies, a few key points were not actively covered in the meeting. One key area was power and energy. It was identified that all of these expeditionary on-demand manufacturing activities require the processing (often by melting) of materials to create new structures. This will require significant amounts of energy. This will not be an issue for the Navy, which will house 3-D printers on aircraft carriers and other large ships with large power plants to drive the processes. The Army will have to consider and develop energy sources to power manufacturing equipment to enable this process.

3.3 The Internet of Battlefield Things

The rapid emergence of the Internet of Things (IoT) is propelled by the logic of 2 irresistible technological arguments: machine intelligence and networked communications. Things are more useful to the human warfighters when they possess more intelligence and more ways to coordinate their actions among themselves. We call this the Internet of Battlefield Things (IoBT) and 20–30 years from now it will be a dominant presence in warfare.

To become a reality, this bold vision will have to overcome a number of major challenges; it will require organizing and managing a large number of dynamic assets (devices and channels) to achieve changing objectives with multiple complex tradeoffs. Such adaptation, management, and reorganization of networks must be accomplished almost entirely autonomously—to avoid imposing additional burdens on the human warfighters—and without much reliance on support and maintenance services. Secondly, human warfighters, under extreme cognitive and physical stress, will be strongly challenged by the massive complexity of the IoBT and the information it will produce and carry. The IoBT will have to assist the humans in making useful sense of this massive, complex, confusing, and potentially deceptive ocean of information while taking into account the ever-changing mission as well as the social, cognitive, and physical needs of humans. Finally, the most important feature of the battle, the enemy—besides being a lethal physical threat to the humans and IoBT—will be infiltrating the IoBT networks and its information. The IoBT itself will be a battlefield between its owners and defenders and its uninvited part-owners: the attackers.

3.3.1 Managing and Adapting the IoBT

By virtue of its exceptionally large scale, IoBT will require new theoretical results, models, concepts, and technical approaches. Indeed, IoBT's number of nodes for a future Army brigade might be several orders of magnitude greater than anything that has been considered in current practice. This is particularly true in the environments where such a brigade will find it advantageous to make use of networked devices and channels that it does not own; such as, when making use of existing, local civilian IoT (networking infrastructure and things) in military operations in a megacity. New theoretical results are needed to understand the degree of determinism resulting from very large ensemble of things and data. One such concept is the recent proof of Wiener's Rule, which stipulates that complex networks in the social and life sciences experience control resulting from the flow of information, not the flow of energy, as is the case in physical networks. The control of such networks is determined by information forces, not physical forces.

A military force that uses an existing IoT of a local society, for example, a megacity, will have to learn (and update automatically and dynamically) about behaviors and performance characteristics of any parts of its IoBT, during the operation. Behaviors and intents of humans—friendly warfighters, adversaries, and neutral civilians—will have to be dynamically detected, identified, characterized, and projected to operate the IoBT.

3.3.2 Making IoBT Information Useful

As important as communications bandwidth is for effective operation of IoBT, it is the human-cognition *bandwidth* that will emerge as the most severe constraint. Human warfighters do not need and cannot process the enormously large flows of information produced and delivered by IoBT. Instead, humans seek well-formed, reasonably sized, essential information that is highly relevant to their cognitive needs, such as effective indications and warnings that pertain to their situational awareness (SA) and mission. To make its information useful, IoBT technologies will have to deal with a quantity and complexity of information that are truly unprecedented in their extent. Consequently, the very foundations of information theory will need to be reconsidered; for example, ensemble probability densities are foundational for information theory and require the underlying process to be ergodic. However, the IoBT is expected to have nonlinear dynamic processes that are sufficiently complex to generate events with nonergodic statistics. The information entailed by the occurrence of such events must be based on single time series and not on an ensembles of time series. Furthermore, nonintuitive, novel phenomena may emerge in the transfer of information between dissimilar large networks—say, between the IoBT and the network of warfighters.

The IoBT's colossal volume of information must be reduced to a manageable level, and to a reasonably meaningful content, before it is delivered to humans and intelligent things. One approach to such a challenging fusion task is to populate IoBT with a layered hierarchy of information brokers, or "concierges", to aggregate, fuse, interpret, and deliver appropriate information.

3.3.3 Dealing with Deception and Adversarial Nature of IoBT

The adversarial nature of the environment is the primary concern in the life of the IoBT. The enemy threatens physical survival and functioning of IoBT as well as the confidentiality, integrity, and availability of the information within IoBT, in addition to attacks on the cognition of human warfighters. Within the IoBT, humans are the most susceptible to deception, particularly to those based on cognitive and cultural biases. The enemy will attempt to violate the information's integrity by modifying it with cyber malware, inserting rogue things into IoBT, intercepting and corrupting it while in motion, and presenting wrong information to the information-acquiring things. Machine learning approaches will be developed to deal with data as big and dynamic as IoBT will possess and will be challenged by the possibility that the enemy will adapt and evolve faster than the learning process can. Learning normal patterns and detecting anomalous deviations, however, does not work well against a well-designed deception. The IoBT may help defeat deception because "lying consistently is difficult"; it may be particularly difficult when the available sources of information are so numerous and are as heterogeneous as in IoBT. In general, much research is needed on approaches to counterdeception, discovery, or rejection of deception for the uniquely complex environment of the IoBT.

3.4 Cyber Fog

The fog of war describes the uncertainty that envelopes a commander engaged in conflict. It is the medium through which he or she can only perceive, not see. This analogy is just as apt in the new conflict domain of cyberspace as it is a physical battlefield. However, it is conceivable in cyberspace to exploit uncertainty to one's advantage and use mathematical methods that increase the security of communications with friends yet further confuse and disrupt an adversary. How one might create a Cyber Fog that is transparent to friends and opaque to foes was considered. The results of the workshop are summarized in this section.

3.4.1 Feasibility, Value and Challenges of Dispersion

The first strategy discussed was the utility of radically fragmenting friendly data into a large number of shards, which would then be dispersed across multiple devices within the battlefield network to increase Cyber Fog. An in-depth discussion of the technical and practical difficulties associated with this method of making information more secure ensued. Research and successful products exist that use some forms of fragmenting, dispersing and frequently repositioning data shards. Other ways to disperse data, which could accomplish the same end, include diversification of channels, protocols, and media. The challenges to dispersion and coherent regathering of data are formidable, but many can be cast as well-defined research topics.

3.4.2 Dispersion and Effective Regathering

Dispersed information will be eventually requested by users and will have to be regathered—in a timely and efficient fashion—and regenerated into a useful form. This could be helped by intelligent dispersion: put shards where they are more likely to be accessible at the time when they are more likely to be needed by the users. One way to achieve improved regathering of information is to account for the semantics of information while splitting it into shards. This could help intelligent prepositioning of related shards. Data provided to the user must be not only relevant but timely; such real-time tradeoffs of security versus timeliness are complex and lack formalization.

3.4.3 Situational Awareness and Information Semantics

Enhanced SA is the goal of information. In addition to regathering data, SA requires discovering information: Where and how do I find what I need to achieve the required SA; which shards need to be collected to get high-quality SA? Semantics of information—including the dispersed data, the semantic context of the friendlies' mission, and the background knowledge of the users—are critical for effective and accurate "defogging". Semantic information theory seems highly relevant to challenges of Cyber Fog.

3.4.4 Risk and Mission

Risk could serve as a comprehensive framework for characterizing the "goodness" of Cyber Fog and should be analyzed in terms of impact on the mission. The consequences of failures of Cyber Fog should be best assessed in terms of consequences to the mission objectives. This implies a need for an adequate model of a mission (including its dependencies on network and computing assets)—a modeling problem that is known to be highly complex. Other complexities arise in seeking ways to measure (quantify) consequences to the mission (e.g., the timing of information determines its importance to the mission). Understanding the risk to mission in an adversarial environment could clearly benefit from a game-theoretic treatment.

3.4.5 Deception and Obfuscation (D&O)

Both dispersion and obfuscation share the key idea: Increase uncertainty (to the adversary) through increased diversity. Arguably, dispersion helps to perform obfuscation and possibly its stronger form, deception. However, very little rigorous, quantitative research has been directed at either deception or obfuscation. Within the Cyber Fog concept, D&O may take multiple forms: In addition to dispersion and frequent repositioning of information, there is the presentation to the adversary of false software and hardware vulnerabilities; diversity of channels helps D&O, as do honeypots and honeynets. But perhaps the most difficult is the construction of a believable deception: creating false battle plans and other unstructured documents (which is very challenging); placing false shards into the fog; designing believable feint attacks that effectively support a real attack; and generating complex multistep deceptions. Machine learning techniques might be applicable to generating believable deceptions, as well as being useful in the no-less-challenging area of counterdeception.

3.4.6 Applicability of Formal Methods

Given the extreme challenges and complexities inherent in the world of Cyber Fog, designing tools and planning specific activities within such an environment may greatly benefit from formal methods. The present state of formal methods suffers from lack of insights into formulating the right questions to ask; that is, which property to verify. Formal methods developed for one domain do not transfer well to another domain (e.g., methods developed for verification of hardware do not transfer well to verification of software and even less so to verification of deception plans).

3.5 Individualizing Technology for Effective Teaming

Since Helmuth von Moltke established the concept of interchangeable officers and organizations in the 19th century Prussian Army,[1] the concept has remained a cornerstone of effective military teaming. The initial dramatic effectiveness of von Moltke's implementation has been linked to an extraordinary training and selection processes, which enabled a high degree of consistency and rigor to behaviors of both officers and organizations. This concept continues to directly influence training, selection, and organizational design today and extends to the integration of US Soldiers with military systems. That is, military systems are designed for use by interchangeable operators with the widest range of aptitudes possible with the belief that appropriate training will enable effective Soldier–system performance. This sets up an inherent trade-off between expanding capabilities to meet increasingly complex operational challenges and designs that maximize the range of potential operators. Inevitably, this leads to the simplification of system designs. This will limit system capabilities, often to an unacceptable extent. It also likely constrains high-performing individuals from using the full extent of their own and the system's capabilities. Moreover, the explosion of technological advancements over the past decades brings with it an understandable desire to insert new technologies to expand operational capabilities. This exacerbates the trade-offs among operator selection and technology complexity as well as with competing priorities for reduced training time, cost, and burden on the Soldier. The goals of the "Individualizing Technology for Effective Teaming" workshop were to challenge our beliefs about the development and design of military human–technology teams and to identify new paths toward providing solutions that will maximize the capabilities of the Army.

As technologies advance toward higher forms of machine intelligence there are dramatic increases in pervasiveness, interconnectivity, and coordination; add to these the greater knowledge of the dynamics of an operator's knowledge, skills and abilities, and the concept of interchangeability takes on new meaning. Technologies are enabling a future of adaptive and individualized systems that function with an individual's capabilities and limitations to achieve greater human–system performance. This individualized human–technology approach enables greater variety in human behavior, while having the ability to maintain consistent, robust outcomes when viewing the human–technology behavior as a system. That is, the base interchangeable element can be conceived of as shifting from an individual Soldier or squad to a joint human–technology element or heterogeneous human–agent team. This approach has 3 major benefits. First, in allowing individuals to behave in manners that are consistent with their own strengths rather than imposing uniform behaviors, individual performance should be dramatically improved. Shifting to this paradigm also enables technological solutions to overcome limitations of individuals, which raises performance in general as well as potentially significantly reducing constraints on personnel selection and assignment. Second, individualized and adaptive designs have the potential to reduce aspects of current training requirements and enable novel training focused on the technological complexity and pervasiveness expected to dominate the next generation of warfare. This approach directly enables human–technology designs as well as heterogeneous human–agent organizational designs that maximize the potential capabilities that future technologies offer rather than restricting capability to ensure maximal interchangeability of operators. Third, the synthesis of human–agent teams working symbiotically should not only enhance individual capabilities, but augment both group-level coordinative processes and technological/system performance, as well. Shifting to this paradigm, performance in the human–technology unit has potential to advance exponentially, with critical implications for talent management and organizational effectiveness.

The individualized human–technology approach offers a vision of teaming defined by Soldier interactions with adaptive and interactive technologies to achieve superior performance. In this vision, teams could be considered as ranging from an individual Soldier with multiple technologies through many Soldiers with many heterogeneous agents. The workshop identified several critical research challenges to meet this aggressive vision. Underlying each of these challenges are common themes, including 1) embracing and advancing machine learning with theoretical considerations and historical knowledge, 2) understanding how to effectively leverage data-driven approaches for effective theory and application development, 3) enabling effective approaches to develop and share data, and 4) emphasizing the use of computational theories and quantitative approaches in human sciences as well as the longitudinal study of humans, human–technology interactions, and heterogeneous teams.

3.5.1 Linking Individuals to Emergent Team Behaviors

A fundamental gap exists in our understanding of how to individualize or adapt technologies to effectively influence team performance, particularly in more complex teams with heterogeneous agents (e.g., human, virtual, robotic). Perhaps the most critical scientific issue underlying this gap is the lack of understanding of the linked relationships between an individual team member's characteristics and dynamics and the emergent, dynamic properties and behaviors of the group. Examples of limitations in current research include a limited ability to infer individual states from observable group measures; limited insight into goals, capabilities, and constraints at the team level; and only an elementary understanding of the emergent and dynamic properties of heterogeneous human–agent groups. Research focused on multiscaled, systems-based approaches to link high-resolution intrapersonal assessment to team and organizational performance; modeling team performance, including critical factors, such as the effects of social isolation in distributed, heterogeneous human–agent teams; and characterizing groups as biological or cognitive metasystems offers promising opportunities to overcome these limitations.

3.5.2 Estimating and Predicting Individual Soldier Capabilities, States, and Behaviors

Humans exhibit high variability, both physiologically and behaviorally, which is understood to be a function of intrinsic dynamics, task demands, environmental influences, and social factors. A significant barrier to the successful development of technologies that adapt to an individual or clusters of individuals is the current paucity of reliable, valid, real-world human estimation/prediction technologies that account for human variability. Research is needed to develop theoretical understandings of the underpinnings of human variability that accurately describe the dynamic, time-dependent causal interactions between domains critical to real-world human performance. Potential avenues of research include the integration of theoretical perspectives across levels, from hormonal/genetic and neuronal through behavioral and societal; and translational theories that focus on complex real-world behaviors, specifically moving beyond explaining behaviors from laboratory settings that are incapable of eliciting an adequate range of human states. Research is also needed underlying the reliable and valid technological capabilities to estimate and predict real-world individual physiology and behavior. Important research topics underlying estimation and prediction involve approaches that integrate scientific knowledge over disparate disciplines, models of human behavior, and ongoing data streams over various time scales—including specific topics focusing on resilience to sparse and incomplete data streams; integration across the numerous factors that influence human performance, including mission, task, and environmental context; integration of high-resolution, multiaspect historical data of individuals, groups, and context; and the capability to detect critical events and factors that influence behavior across multiple time scales.

3.5.3 Enabling Joint Human–Technology Team Capabilities and Performance

A fundamental gap also exists in our understanding of what capabilities and mechanisms will most effectively influence team performance. Technologies are increasingly capable of modifying human cognition, memory, affect, perceptions, metabolism, personality, physical actions, and social behaviors. Alternatively, technologies are increasingly capable of replacing aspects of team function that are currently performed by human team members. From the perspective that complex teams are integrated systems, research is needed to characterize system-level requirements; explore novel human–technology roles and relationships, states, and processes and their impact on team performance; and identify and pursue specific potential technological capabilities, including human–technology-interaction technologies to augment team performance, rather than just individual performance. One promising avenue of research focuses on the notion of self-expansion, illustrated by a squad being composed of a single individual with many virtual and/or physical components or systems. A second promising avenue of research focuses on "super team" technologies aimed at coordinating emotion, affect, and cognition across multiple constituent elements to enable better, faster communications and task diversity across heterogeneous teams that include multiple human members. Critical research topics in these avenues include control-theoretic approaches focusing on the emergent states and behaviors of interdependent groups, mutual adaptation, and dynamic resource allocation; hybridizing human–agent intelligence to outperform either human or agent teams; increasing socio-cultural intelligence in agents; approaches to infer, develop, and maintain group intent; multidimensional prescriptive theories of team trust; and understanding biochemical/neurochemical influences on Soldier emotional, cognitive, and social function.

3.5.4 Evolving Joint Human–Technology Teams

Underlying the potential for dynamic reconfiguration of teams is a fundamental gap in the understanding of how to rapidly develop those teams within dynamic environments and scenarios including rotating or transitional members. This gap forms a critical limitation in the ability to leverage potential technologies and realize optimal team performance. One promising research avenue is to promote the rapid human–technology integration through the evolution of self-organizing systems, which would allow for adaptation to specific task and contextual demands.

Consider the notion of self-expansion (in Section 3.5.3) for an individual integrated with numerous component technologies or prototype “building blocks” technologies that, when combined in potentially unique ways, can lead to novel human–agent interactions. Critical research topics within accelerated evolution and rapidly adapting human–agent system capabilities are expected to provide new and flexible technology solutions. Specific research areas include investigating dynamic adaptation with both bottom–up (trial and error) and top–down (model) approaches; high-frequency experimental paradigms mixed with other mechanisms such as crowdsourcing; characterizing dynamic changes in human agency when teams are fluid; replacing traditional training for specific systems with training for adaptable systems; and understanding how to develop systems for improved adaptability, flexibility, and training transfer—yet avoid maladaptive and far-from-optimal outcomes.

3.6 Distributed and Collaborative Intelligent Systems

Advancements in intelligent and autonomous systems have been primarily focused on individual agents or small teams. Future systems are envisioned to be highly heterogeneous, collaborative, and distributed to provide key capabilities and tactical advantages in future, complex operational scenarios. Four key focus areas were identified for the advancement of Distributed and Collaborative Intelligence Systems: 1) Distributed Awareness, 2) Distributed Intelligence, 3) Adaptable and Resilient Control, and 4) Scaling Experimental Complexity. Some specific topics generated during the ASPSM workshop on Distributed and Collaborative Intelligent Systems are detailed in Sections 3.6.1–3.6.8.

3.6.1 New Sophisticated Planning and Control Architectures

Traditional Human–Robot Interface has focused on ways to assist humans in controlling robots via interfaces; but with large numbers of robots, the cognitive load is simply too high for this traditional approach. How to do collective multiagent coordination at this level or to include the human as another intelligent agent within the network is an entirely unsolved task. Depending on scale and complexity, current implementations for intelligent systems typically rely on either centralized or decentralized control architectures and do not simultaneously deal with both global and localized control of single agents, multiagent teams, and localized swarm behavior. More research is needed to develop the general science or architectures for large numbers of distributed heterogeneous agents and for finding optimal plans, which are often computationally hard, especially for systems in complex environments.

3.6.2 Resiliency

The resiliency of large multiagent systems needs to be considered based on realistic communication links and localization uncertainties. The most popular distributed control and planning paradigms today simply assume that communication links are a given, or can be modeled with simplistic "on/off" linkages with unlimited bandwidth, or that localization errors are similar across the multiagent network. While this enables the use of convenient analytical tools for proving stability/optimality guarantees, such assumptions do not accommodate important sources of error and uncertainties. Efforts are needed to know what information is needed and not given, to determine what is missing or corrupted, and to collect missing data. Future paradigms should also consider morphing, reconfigurable, and adaptable platforms and systems performance as ways to increase overall mission resiliency.

3.6.3 Multiagent Learning

Multiagent learning is an attractive alternative to directly coding teams or swarms of agents or robots, particularly since it could be used in the field by nonexperts to train groups of robots to do tasks on the fly, or to model and respond to threats or unexpected environments. A fundamental challenge with multiagent learning is the "Multiagent Inverse Problem". The standard way to overcome inverse problems is to use optimization. While some optimizers, such as reinforcement learning or policy search, can be used in simple scenarios, these methods are not likely to scale when agents become complex and heterogeneous and have complex interactions. More research is needed to find approaches to teach large swarms of complex agents how to do nontrivial collective tasks in real time and in the physical world.

3.6.4 Level and Mix of Heterogeneity

Computers, phones, and other devices have moved toward homogeneity in design rather than heterogeneity. New simulation tools are needed to explore and optimize needed levels of heterogeneity. The current game-theory-based formulations are computationally intractable when dealing with teams involving hundreds of agents. The expected speed-up in computational power will be of very limited use in addressing this issue because of the exponential nature of these formulations.

3.6.5 Adaptability and Reconfigurability

Complex missions require multiple teams to simultaneously carry out multiple tasks, and agents may need to play multiple roles that may span across teams. As contingency situations arise, rapid reconfigurations in teams, both locally and globally, will be needed across the distributed architecture. How to synthesize new behaviors on-line to deal with the unexpected contingencies will be a must for the system to exhibit resilient behavior. Doing on-line synthesis of behaviors in a distributed architecture in a fast-paced mission will be very challenging.

3.6.6 Operational and Experimental Complexity

There have been recent examples of operating fully autonomous systems in complex environments and large numbers of homogeneous agents/swarms in simple environments; but, to make these demonstrations tractable, researchers typically reduce the complexity along several axes: 1) number of agents, 2) degree of heterogeneity among the agents, 3) agent behavior complexity, autonomy, and adaptability, 4) the degree of interactions and communication among the agents, 5) speed of operation, and 6) the complexity of the environment and available infrastructure. Performing large-scale demonstrations and simulations that push the degree of complexity along each of these scales is not currently possible today.

3.6.7 Distributed Decision Making

How to make decisions in a distributed manner that is considered acceptable is an open question for which we need to develop methods that work with arbitrarily complex cost functions in complex environments. We also envision that access and use of the cloud, big data, social media, real-world complex models (i.e., weather), and other knowledge bases can be included and leveraged to support intelligent/semantic routing of valuable information or answer critical questions that are unknown beforehand due to the rapid situation change on a battlefield. When communication between agents is limited or even completely disrupted, the only way to counter such an adverse situation is to perform reasoning and prediction. Reasoning and prediction are also critical when missions and objectives are not clear or are changing rapidly in dynamic and complex environments.

3.6.8 Soldier–Multiagent Collaboration and Decision Making

How best to express operational intent and SA to a large group of heterogeneous platforms to process directives or observations provided by a human is a challenge. It is also a challenge to ensure that human users/designers can even comprehend the scope or scale of large-scale robotic network capabilities or be sufficiently aware of risks involved in completing certain tasks. Alternative programming and software-design paradigms may be needed to encode, simulate, and validate very large-scale robotic systems, particularly if each agent is expected to be highly adaptable and possibly capable of complex autonomous behaviors. How humans should interact with robots in large heterogeneous systems is another open question.

Humans will be basically interacting with robot "crowds". Language-based communication will be useful, but we should examine other modalities as well.

4. Observations

A review of the individual meetings by the organizers revealed common themes, even across the materials-focused and information-focused meetings. Organizers observed that enabling a future military to embrace its vulnerabilities requires materiel capable of responding to rapidly changing threat surfaces. From a very broad perspective, many of the desired capabilities required adaptively reconfiguring a heterogeneous collection of entities in response to external stimuli—and to do so with constrained resources. In some instances the resources may even be contested, unreliable, and untrusted. This applies as much to networks and distributed processing systems as it does to teams of humans and intelligent systems and to materials. Reconfigurability spotlights not just the entities but also the interconnections between them and on how energy and information flow among the entities.

At a basic level of understanding, one can characterize entities by their number, their function, and their degree of homogeneity across the collection; further, one can characterize connections by their density, degree of nonlinearity, and "noisiness" (i.e., how much energy or information they inject into a signal due to random fluctuations). But, to implement reconfigurability, ultimately one requires rules to design a system that adapts intelligently to external stimuli.

In the areas we considered, we do not know these rules because we do not understand how energy flows, or how information flows, between interconnected elements. Predictive models, whether based on physics or empirical data, do not exist.

In physical systems such as reconfigurable materials, models must reflect the physical mechanisms to capture, transform, and transfer energy in feed-forward and feedback networks. Systems driven by information also require models for capture, transformation, and transfer. However, the models for the 2 systems differ considerably. Whereas physical systems produce forces when there exists an energy gradient (a difference in energy between 2 points), in information systems the forces and movement are generated by gradients in entropy, a measure of information. Such gradients are a product of system topology.

Another feature shared by the material and information systems we considered is extreme scale. For materials, our ability to control structure at a nanometer scale gives us the potential to control 10^{15} volume elements or a petavoxel of material.

In information systems, although a small tactical unit may contain only a handful of soldiers and tens of intelligent systems, if it is “plugged into the city” the small unit is connected to a much larger network whose total number of connections is also on the order of 10^{15}. Systems with extreme scale require extreme modeling and simulation tools.

In summary, an expeditionary fighting force that is willing to exploit unconventional logistics (“plugged in”) and embrace its vulnerabilities requires rapidly reconfigurable materiel. This includes collaborative teams of humans and autonomous agents, networks of sensors and intelligent devices, and materials. The desire for reconfigurability requires a deeper understanding of forces, both energetic and entropic, and a deeper understanding of the role of topology in these forces; to complement analytic studies, we require extreme modeling and simulation tools.

5. Recommendations

Based on the output of the meetings and subsequent discussions among meeting organizers, ARL developed recommendations specific to each area.

The ASA(ALT) should invest in the following programs:

- Energy management within materials systems, as a means of driving the assembly and reconfiguration processes, is ripe for exploration.
- Computational modeling and nanoscale characterization tools to enable efficient design of hybridized manufacturing; real-time multiscale computational capabilities to enable predictive analytics for expeditionary on-demand manufacturing
- Discovery of design principles to enable programming advanced genetic circuits with specified functionality
- Formal theories of deception, counterdeception, discovery or rejection of deception, in the context of IoBT
- Novel approaches to information theory, appropriate for large, dense, dynamic heterogeneous networks, with nonergodic transient information flows
- Major new program in large dimensional (stochastic) game theory, where the number of players and the action space is very large and dynamic
- Distributed intelligence (e.g., via novel approaches to learning in an adversarial environment and under resource constraints)

- A program to develop formal methods to execute and manage successful obfuscation of friendly information in a cyber fog. Deception should be integral to this theoretical analysis, for example, in the level of complexity necessary to deceive successfully.
- Research is needed to develop theoretical understandings of the underpinnings of human variability as well as the reliable and valid technological capabilities to estimate and predict real-world individual physiology and behavior.
- Novel approaches to enable efficient high-fidelity simulation of a large number of heterogeneous entities (including mixed human–machine teams) operating autonomously in complex environments

6. Summary and Conclusion

The 6 meetings in fall 2015 reflect those areas ARL considered critically important to the Army. The recommendations within each area are intended to show where investment will most provide significant new understanding to advance capabilities desired by the Army.

The strategic recommendations reflect cross-cutting thrusts that encompass fundamental aspects of the individual areas and, if followed, should lead to new, broad-based capabilities. It is telling that so many of these recommendations are concerned with humans and their interface to engineered systems, both physical and cybernetic. A particular emphasis in the cybernetic interface is on the flow of data and information between humans and engineered systems.

The Army's recognition of parallel convergences in the technical domains of physics, biology, and cybernetics and in the warfare domains of the physical, virtual (informational), and human (cultural) should allow it to prepare for an uncertain future.

7. References

1. Howard ME. The Franco-Prussian war: the German invasion of France, 1870–1871. New York (NY): The Macmillan Company; 1961.

INTENTIONALLY LEFT BLANK.

Appendix. Individual Meeting Summaries

Microscale Adaptability

January 11–12, 2016
Aberdeen Proving Ground, Maryland

Organizers: Dr Dawanne Poree (US Army Research Laboratory [ARL] Army Research Office [ARO]), Dr John Prater (ARL–ARO), and Dr David Stepp (ARL–ARO)

Introduction

The bottoms-up design and assembly of materials offers the opportunity for achieving exquisite control over the local chemistry and properties of a material, enhancing our ability to engineer greater functionality and design complexity into future material systems. Nature provides clear examples of the efficacy of this approach, as well as some of its limitations. In particular, research has demonstrated the feasibility of using tailored interactions between small molecules to drive the assembly of inorganic nanoparticles and biological entities (e.g., DNA and proteins) into larger assemblies that can extend over multiple length scales. Similarly, dynamic assembly processes have been studied in which the bonding and network structures can undergo disassembly and reassembly in route to reaching their final state. Finally, research is underway to develop a fundamental understanding of how to propagate and amplify molecular-level detection events over extended length and time scales to drive a macroscopic material-property change. The major long-term objective of the research is the design and synthesis of systems that exhibit adaptable, responsive living traits and materials that display emergent properties. However, the design and sequential assembly of complex 3-D structures culminating in hierarchically structured materials with specifically targeted properties and/or engineered dynamic responses still remains well beyond our grasp.

The "Microscale Adaptability" Army Science Planning and Strategy Meeting (ASPSM) sought to bring together 2 disparate research communities— responsive self-assembled materials and optical metamaterials—to identify breakthrough strategies (both theoretical and experimental) that would facilitate the design and robust self-assembly of multicomponent, 3-D structures with precisely engineered electronic and optical properties. More specifically, the overall objectives of this meeting were 3-fold: 1) to assess the current state of research in each field, 2) to identify milestones necessary to achieve these long-term outcomes, and 3) to identify Army-relevant science and technology (S&T) capabilities that might emerge from the convergence of these fields of study. To meet these objectives, a diverse group of academic and government scientists was invited to participate in the workshop. The expertise of these researchers spanned a wide range of disciplines including organic/polymer chemistry, materials science, physics, chemical engineering, and multiscale theory.

The workshop used a nontraditional, interactive format relying almost entirely on small-group breakout sessions to identify new collaborations and drive discussion. The workshop began with 2-minute introductory presentations from the academic and bench-level government scientists to briefly discuss their research interests and capabilities relevant to the topic area. Following the presentations, the participants were organized into small groups and tasked to identify 1) new scientific opportunities that may be possible by converging the fields of responsive self-assembled materials and optical metamaterials, 2) key barriers to achieving these opportunities, and 3) how success of convergence may enable future Army capabilities. The dynamic nature of the workshop allowed it to evolve in real time, as the results and feedback from each breakout session were used to determine the objectives and focus of the subsequent sessions.

Scientific Recommendations and Opportunities

The bottoms-up assembly of materials is broadly accepted as a key enabler for material designers to establish exquisite control over local chemistry as a means of building functionally complex systems and boosting overall material performance. In addition, this approach may lead to new avenues for realizing smart systems capable of collecting information about their surroundings and modifying specific functions in response to its environment, akin to living systems. While many of the pieces required to achieve such an advance exist, integrating the parts to deliver a robust capability remain problematic. Inspired by recent demonstrations of new paradigms for optical materials that exhibit “smart” properties and robust tunability, this workshop set out to explore the feasibility of integrating new molecular building blocks and novel assembly techniques to generate new materials with unprecedented responsive and reconfigurable properties to manipulate light-matter interactions. Engaging in interactive, cross-disciplinary, small-group discussions, the workshop participants identified a number of promising areas of opportunity, such as these 3:

Self-Assembly of Reconfigurable Elements

Recent studies have demonstrated the feasibility of using tailored interactions with small molecules to drive the assembly of inorganic nanoparticles into larger structures spanning multiple length scales. Similarly, dynamic assembly processes have been studied in which the bonding and network structures can undergo disassembly and reassembly in route to their final structures. By extension, one can imagine the insertion of reconfigurable elements (building blocks capable of dynamically altering their configuration), varying their coupling interaction, and even switching between physical states, potentially enabling one to synthesize "smart" materials that possess the ability to change their properties in response to external command signals or environmental cues. Such systems will require the development of techniques for staging individual attachment steps into a sequential assembly that leads to complex and/or hierarchical architectures with unique properties. Additionally, avenues for the capture, conversion, and/or transduction of various forms of energy should be incorporated into the design as a means of driving the assembly and/or reconfiguration processes. This concept of energy management within a materials system was noted by the workshop participants as an area ripe for exploration. It was also noted that self-assembly approaches seem to be inherently susceptible to the formation of defects that can greatly degrade the mechanical and functional performance of the materials. Over the millennia, living systems have evolved mechanisms for countering the effects of defects. Similar approaches to defect tolerance and functional redundancy need to be perfected for man-made materials.

Responsive Behavior

A daunting challenge toward synthetic "living" systems is predictably propagating a molecular-level change, generated through the selective sensing of a trigger, into a readily discernible macroscopic change in a material's fundamental properties. This can only be addressed by developing a fundamental understanding of the chemical processes that occur at multiscale levels to enable active control that spans from molecular to nanoscale to macroscopic length scales and from nanoseconds to hours. The inherent complexity involved in connecting these length scales, and the transduction (detection, amplification and propagation) of external signals into macroscopic responses, require a cohesive, multidisciplinary approach.

Engineered Interfaces

Access to the interfaces at all levels in these materials systems was called out by the group as a particularly compelling opportunity for the field. Opportunities for manipulating the assembly process by using aspects of shape, intermolecular interactions, induced conformation changes, functionalized adduct and site-specific binding groups, molecule-to-substrate interactions, and external fields need to be fully developed. An intriguing new development is the concept of pluripotent matter in which changes in bonding are driven by the surface functionalization of the basic building blocks, which is sensitive to external signals including changes in the external environment. Similarly, liquid infiltration into voids within percolating porous systems (e.g., opal and inverse opal systems) is another broad approach available to radically alter the optical properties of a system via changes in index matching. Another interesting possibility is the incorporation of particle jamming at interfaces in 2-phase liquid systems, which could impart on-demand rigidity into these systems. Finally, methods for effectively interfacing various biological and inorganic elements into complex functioning architectures capable of maintaining biological functionality need to be fully investigated.

Nanoscale Characterization Tools, Computational Modeling

Advancements in computational modeling and nanoscale characterization tools will be key to enabling each of these opportunities. More specifically, new theoretical tools and computational methods capable of modeling the self-assembly process and identifying valid self-assembly pathways that lead to stable hierarchical architectures and desired functionality are needed. Additionally, the theory needs to incorporate robust reconfigurability and ultimately predict the range of dynamic behavior that can be achieved in these systems; furthermore, these predictions need to be experimentally validated. Regarding analytical characterization, development of quantitative, high-resolution techniques for spatiotemporal characterization and methods that will advance our understanding of the fundamental structure–function relationships at the molecular and nanoscales, as well as across scales, are a high priority.

Potential S&T Impact

A robust research program that facilitates the design and fabrication of reconfigurable matter, capable of autonomously responding to its environment and circumstances, would have profound implications in enhancing the capabilities of the future warfighter.

Potential applications include materials for adaptive optics, conformal antennas, reversible adhesives, tunable negative-index materials, laser protection, autonomous sensors, agile communications and electronic systems, temperature stabilization, synthetic immunogens and clotting agents, microrobotic systems, energy harvesting, hydrogen storage, and radio frequency (RF) ID.

Expeditionary On-Demand Manufacturing

January 21–22, 2016
Aberdeen Proving Ground, Maryland

Organizer: Dr Robert Carter (ARL–Weapons and Material Research Directorate [WMRD])

Summary and Introduction

Whereas the potential of additive manufacturing (AM) and manufacturing on-demand is obvious, many fundamental questions related to their practical implementation remain unanswered. The objective of this meeting was to identify technical hurdles that limit the applicability of on-demand manufacturing in austere military environments. A group of more than 50 scientists and engineers reviewed the current global state of the art and projected the needs of expeditionary manufacturing technologies to include rapid certification, synthesis, processing and microstructure control of advanced materials, multimaterial and multifunctional structures, and biological production and identified key areas of medium- and long-term scientific opportunity within this theme.

Recent advances in manufacturing technologies have created new potentials for expeditionary and responsive production of materiel. These technologies could have significant impact on the Army's capability to be threat responsive, reduce the logistics footprint, and create new capabilities to build complex systems rapidly in austere environments.

This ASPSM group reviewed the current and discussed projected future capability in 4 areas: Rapid Certification of AM Structures, Moving Logistics Forward, Hybridized Manufacturing, and Bioproduction. The ARL hosted more than 50 attendees from the Department of Defense (DOD), other government agencies, academia, and industry to discuss strategic research needs to best develop advanced manufacturing technologies enabling expeditionary manufacturing capabilities for the Army of 30 years from today.

Rapid Certification of AM Structures

This session focused on the current state of the art in reducing the time needed to qualify materials and certify new processes. Currently, it takes several years to decades to get a new material or manufacturing process matured sufficiently to transition to widespread adoption. Additive technology is 30 years old and is only now finding use in some application. Rapid certification is only in its infancy and has demonstrated the potential to bring new parts in a greatly compressed timeline.

The goal would be to make in-line certification possible so that any part made could be put into service directly from the manufacturing process.

Moving Logistics Forward

This session focused on the needs and future applicability of distributed manufacturing to impact the Army's logistics needs. This topic's discussions ranged from the effectiveness of deploying advanced manufacturing capability to minimize logistics, to "build new" versus "repair existing", to indigenous manufacturing (using materials from the surroundings or waste streams as feedstocks).

Hybridized Manufacturing

Battlefield dominance hinges on the ability to handle known and unknown threats, situations, and environments. Therefore, capability for threat responsiveness (e.g., creating a weapon or a coordinated–distributed configuration of intelligent mobile platforms for Intelligence, Surveillance, and Reconnaissance for offensive/defensive weaponry) requires building with multiple materials and integrating multiple functions (electronic and optical integration onto a flying/ground-mobile intelligent system). We cannot create components composed of multiple materials. We do not have the means nor do we understand the relevant interfaces or how to control interface properties. The S&T to create/build with multimaterials sets is critical, as well as S&T to integrate electronic, senor, and actuator functions into the build. Various approaches, from creating circuits and building function on the fly (like printable electronics and "pick and place") to creating functional voxel building blocks and assembling were discussed. Ideas are out there, but there is a clear need for S&T to enable them.

Bioproduction

The Army requires chemical feedstocks to create materials and "factories" to purify water, create energy, create structure, and create useful biological function (e.g., gut biome). S&T is required to harness the power of biology (synthetic and systems biology).

Recommendations and Opportunities

Advanced manufacturing brings the potential for an adaptable, threat-responsive capability. Primary interest for current research is focused on developing AM for logistics.

Scientific Recommendations and Opportunities

- Materials by design capability to create the modeling tools for in-line certification of parts. Development of a “closed-loop” additive system would need Integrated Computational Materials Engineering-based computations running orders of magnitude faster than currently possible. Also, in situ sensing of complex matter–energy interactions at the point of fabrication. Component performance (macro) relies on management of microstructure, composition, and defects (among others) and nano/micro scale. High spatial and temporal resolution are required to control and create complex structure, provide the information necessary to certify the parts, and predict performance. Computation in its current embodiment takes far too long to be used in the process for multimaterial sets and different designs. Gaps include
  - Real-time computational predictive capability
  - Rapid measurement of local material properties, state (e.g., temperature), and defects
- Additive-process development—High spatial control of deposition of multiple materials and higher throughput are needed to create higher complexity and more useful structures. The higher rate would increase the return on investment for logistics impact.
- Development of new design tools to take advantage of the increases in complexity enabled by AM. This is more than just topology optimization—but where are there potentials for merging topology optimization, artificial intelligence, and big data to rapidly optimize designs? There should be a very different approach to defining the application space and allowing computers to optimize the solution.
- Need better life-prediction models to predict the service life of optimized parts and parts produced from AM processes.
- Need to manufacture complex multimaterial–multifunctional systems in the field. This would include sensors, electronics, power sources, motors, and actuators.
- Polymer chemistry (and thus the available set of materials properties) are limited by the chemistry of petroleum feedstocks. Biology may be exploited to produce designed (molecular weight, branching, functional groups, configurations, chirality, etc.) chemical precursors to produce exceptional products. The chemical precursors can be pharmacologically active as

well—antidotes, vaccines, human enhancement, and so on. S&T is required to enable this capability and move it toward “expeditionary”.

Programmatic Recommendations and Opportunities

1) The Army needs a large and sustainable financial commitment to acquire and sustain the expertise and resources needed to solve the extraordinarily challenging problems that will arise, given the necessity of complex theoretical, experimental, and modeling linkages.

2) To stay scientifically competitive, the Army *must* invest in long-term, on-site collaborative research efforts with national institutes and facilities.

3) Underpinning each of the scientific recommendations is the inherent ability to characterize phenomena at relevant length scales, from subatomic to mesoscale, and at relevant time scale.

4) Similarly, multiscale computational efforts are needed to predict properties and provide relevant and real-time control of advanced manufacturing methods to develop closed-loop control systems for in-line certification of parts. This will require physics-accurate simulations of matter–field interactions, far-from-equilibrium process models to incorporate nano- to microstructure evolution, and process property correlation.

5) “Partnerships” should reach across fields including computational mathematics, multiscale science, multiphysics applications, optimization and uncertainty quantification, and computational multiphysics.

6) Seed resources should be used to form interdisciplinary teams to tackle subproblems in a truly collaborative way, including periodic discussions and yearly focused meetings.

Speakers

Scot Seitz (Army Logistics Innovation Agency)

Raymond Clinton (NASA)

Jordan Brandt (Stanford University)

Rob Ivester (Department of Energy’s Advanced Manufacturing Office)

Mick Maher (Defense Advanced Research Projects Agency)

James Neumann (Honeywell)

Jason Sebastian (QuesTek)

Andrew Baker (Boeing)

Pete Collins (Iowa State University)

LJ Holmes (ARL)

Mark Schlien (Army Edgewood Chemical Biological Center)

Dean Hutchins (Defense Logistics Agency)

Neal Orringer (3D Systems)

Caroline Scheck (Naval Sea Systems Command)

Eric Forsythe (ARL)

Ken Church (nScrypt)

Hod Lipson (Cornell University)

David Roberson (University of Texas at El Paso)

Christian Sund (ARL)

Gerald Grant (University of Louisville)

Pamela Peralta-Yahya (Georgia Tech)

Adam Safir (Zymergen)

Alex Tobias (Du Pont)

The Internet of Battlefield Things

November 9–10, 2015
Adelphi Laboratory Center, Maryland

Organizers: Alexander Kott (ARL–Computational and Information Sciences Directorate [CISD]), Ananthram Swami (Scientific Professional [ST], ARL–CISD), and Bruce J West (ST, ARL–ARO)

Introduction

Organized by ARL, "The Internet of Battlefield Things" ASPSM took place on November 9–10, 2015, at the ARL's Adelphi Laboratory Center (ALC). The rapid emergence of the Internet of Things (IoT) is propelled by the logic of 2 irresistible technological arguments: machine intelligence and networked communications. Things are more useful and effective when they are smarter and even more so when they can talk to each other. Exactly the same logic applies to things that populate the world of military battles. They, too, can serve the human warfighters better when they possess more intelligence and more ways to coordinate their actions among themselves. We call this the Internet of Battlefield Things (IoBT). In some limited ways, IoBT is already becoming a reality,[1] but 20–30 years from now it is likely to become a dominant presence in warfare.

The battlefield of the future will be densely populated by a variety of entities ("things")—some intelligent and some only marginally so—performing a broad range of tasks: sensing, communicating, acting, and collaborating with each other and human warfighters.[2] They will include sensors, munitions, weapons, vehicles, robots, and human-wearable devices. Their capabilities will include selectively collecting and processing information, acting as agents to support sensemaking, undertaking coordinated defensive actions, and unleashing a variety of effects on the adversary.[3] They will do all this collaboratively, continually communicating, coordinating, negotiating, and jointly planning and executing their activities. In other words, they will be the Internet of Battlefield Things.

To become a reality, however, this bold vision will have to overcome a number of major challenges. As one example of such a challenge, the communications among things will have to be flexible and adaptive to rapidly changing situations and

[1] Seffers GI. Defense department awakens to internet of things. Signal Mag. 2015 Jan 1.

[2] Kott A, Alberts DS, Wang C. Will cybersecurity dictate the outcome of future wars? Computer. 2015;48(12):98–101.

[3] Scharre P. Robotics on the battlefield part II: the coming swarm. Washington (DC): Center for a New American Security. 2014.

military missions. This will involve organizing and managing large number of dynamic assets (devices and channels) to achieve changing objectives with multiple complex tradeoffs. Such adaptation, management, and reorganization of the networks must be accomplished almost entirely autonomously—to avoid imposing additional burdens on the human warfighters—and without much reliance on support and maintenance services. How can this be done?

Secondly, human warfighters, under extreme cognitive and physical stress, will be strongly challenged by the massive complexity of the IoBT and the information it will produce and carry.[4] The IoBT will have to assist the humans in making useful sense of this massive, complex, confusing, and potentially deceptive ocean of information, while taking into account the ever-changing mission as well as the social, cognitive, and physical needs of humans.

Finally, nobody can discount the most important feature of the battle: the enemy. Besides being a lethal physical threat to the humans and IoBT, the enemy will be lurking in and around the IoBT networks and its information. The IoBT itself will be a battlefield between its owners and defenders and its uninvited part-owners—attackers. How will the IoBT manage risk and uncertainty in this highly adversarial, deceptive environment?

Topics of Discussion

These are some of the questions that were discussed at the strategic planning meeting organized by ARL (http://www.arl.army.mil) on November 9–10, 2015. It brought together a number of scientists from academia and industry and military experts. The topics discussed at the meeting included the following:

- Quantifying information gain including uncertainty a) in a given adversarial context and b) for the purposes of improved situational awareness
- Learning in the environment of adversarial deception/misinformation
- Dynamic discovery and allocation of heterogeneous, potentially composable information-gathering resources to optimize situational awareness (SA) of the IoBT, in the context of the mission, and adversarial deception/misinformation
- Theoretical foundations for detection of anomalies and adversarial deception/misinformation, approaches to correcting the gathered

[4] Kott A, Wang C, Erbacher RF, editors. Cyber defense and situational awareness. New York (NY): Springer; 2014.

information, and maintaining fault tolerance and integrity of information and cyber–physical resources

- In-network, distributed and asynchronous processing, analysis, and fusion of multimodal heterogeneous information, in the context of the mission, including quantification of risk and uncertainty
- Theoretical foundations for information discovery, processing, and delivery, leading to an understanding of tradeoffs (amount of information collected, opportunity for tampering, resource consumption, latency, etc.) and, thus, predictive resource allocation (sensing, computing, communications, etc.) taking into account risk and uncertainty
- Semantics-oriented approaches to represent, organize, summarize, and reason about (the potentially large volume of) information generated by the IoBT—"big data" issues specific to IoBT
- Foundational approaches to efficient communications and networking, in the context of dense IoBT networks, and unusual traffic patterns

The suggestions and concerns that emerged at the meeting coalesced into a rich and ambitious research agenda, summarized in the following sections.

Managing and Adapting the IoBT

In spite of voluminous, current, and past research on related topics in network science and engineering, merely by virtue of its exceptionally large scale IoBT will require new theoretical results, models, concepts, and technical approaches. Indeed, IoBT's number of nodes for a future Army brigade might be several orders of magnitude greater than anything that has been considered in current practice. This is particularly true in the environments where such a brigade will find it advantageous to make use of networked devices and channels that it does not own, (e.g., when making use of the existing, local civilian IoT: networking infrastructure and things) in military operations in a megacity. In this case, the meeting's participants suggested, IoBT scale on the order of a million things per square kilometer is not an unreasonable target for exploration.

On the other hand, the massive scale of IoBT can be advantageous in practice and even for theoretical purposes. For example, availability of very large and densely positioned number of things, such as sensors, can help eliminate currently common concerns about availability of any of them at a given time. To this end, theoretical results are needed to understand the degree of determinism resulting from very large ensemble of things and data.

For example, Norbert Wiener, the author of Cybernetics, speculated that the complex networks in the social and life sciences experience control emanating from the flow of information, not the flow of energy.[5] He considered a system high in energy coupled to one low in energy, but extremely high in information (i.e., of great negative entropy). He goes on to conjecture the coupling is such that the information, negative entropy, passes from the system at low energy to the system at high energy and subsequently determines its organization. Wiener's speculation —recently proven[6]—implied the force laws and control in social phenomena do not follow the negative gradients of energy potentials (if they could be defined) but, rather, they follow gradients due to information imbalance. The force involved might be viewed as an information force.

Quite apart from its large scale, extreme heterogeneity of IoBT will call for new research and approaches. Not only the local IoT will consist of a broad range of commercial things and networks, but even the equipment the warfighters will bring with them into the battle will likely rely on commercial offerings.[7] It is probable that future commercial IoT will continue to exhibit a lack of standards, partly driven by desires of individual manufacturers to control its market, and will be generally chaotic. The military will have to adapt rapidly—and have suitable technologies and techniques for such an adaptation—to use a broad variety of things, protocols, and communication technologies from multiple manufacturers.

In such a heterogeneous, highly dynamic, and largely unpredictable environment, new approaches will be needed to facilitate discovery, characterization, and tracking of relevant, available, and useful things dynamically in time and space. In particular, a military force that uses an existing IoT of a local society (e.g., a megacity) will not able to make reliable assumptions about behaviors and performance characteristics of any parts of its IoBT; instead, such behaviors and characteristics will have to be learned and updated automatically and dynamically during the operation. Speaking of complex and unpredictable behaviors, one must not forget that humans—whether we call them "things" or not—are crucial and highly influential elements of the IoBT. Behaviors and intents of humans—friendly warfighters, adversaries, and neutral civilians—will have to be dynamically detected, identified, characterized, and projected to operate the IoBT.

Communications between things will also be challenged by the high complexity, dynamics, and scale of IoBT. Finding, sharing, and managing communication

[5] Wiener N. Time, communication, and the nervous system. Annals NY Acad Sci. 1948;50:197–220.

[6] Aquino G, Bologna M, Grigolini P, West BJ. Beyond the death of linear response theory: criticality of the 1/f-noise condition. Phys Rev Lett. 2010;105:040601.

[7] Downing C. The internet of things for defense. White paper. Wind River; 2015.

channels among large numbers of competing, heterogeneous, and often unpredictable things will require novel approaches. Highly intelligent automation will be required to continually allocate and reconfigure the resources of the communications network. Information-sharing strategies and policies—who talks to whom, when, about what, and how long—will have to be automatically designed and modified dynamically.[8] Highly scalable architectures and protocols will be needed along with rigorous methods to determine and validate properties of protocols and architectures. In extreme situations when the IoBT experiences catastrophic collapse or becomes largely unavailable, or is untrustworthy due to enemy actions,[9] the autonomous management of the IoBT will need to provide at least a "get me home" capability that will enable the continuation of operations, even if at a limited level of functionality.

Additional complexity will arise from the wide range of timing constraints on communications. Some communications can wait for hours, while other communications will pose real-time requirements; for example, for sensing and actuating. The channels will be constrained in highly heterogeneous ways as well. It is expected that 30 years from now consumers will use wireless channels typically for only a few meters before the data enter fiber or other high-capacity channels; at the same time, the military will require at least a few kilometers of wireless channels before encountering fiber.

To enable the dynamic management of IoBT, situational awareness[4] of the IoBT as a whole will be formulated and updated rapidly and automatically; therefore, new approaches will be desired and directed toward the ability to measure relatively few variables of the complex system while thereby obtaining or inferring sufficiently complete information about the system.

While managing the IoBT its purposes and uses must be taken into account, and these will be diverse. Some of its purposes will be relatively well understood, such as tactical military logistics or distributed computing. Others will be novel and will emerge from the availability of the IoBT itself, such as perhaps the use of it for Position, Navigation and Timing needs and as a supplement to, or replacement for, GPS.

[8] Misra S, Xue G. Efficient anonymity schemes for clustered wireless sensor networks. Int J Sens Net. 2006;1(1/2):50–63.

[9] Køien GM. Reflections on trust in devices: an informal survey of human trust in an Internet-of-Things context. Wire Pers Comm. 2011;61(3):495–510.

Making IoBT Information Useful

As important as communications bandwidth is for effective operation of IoBT, it is the human-cognition *bandwidth* that will emerge as the most severe constraint.[10] Human warfighters do not need and cannot process the enormously large flows of information produced and delivered by the IoBT.[11] Instead, humans seek well-formed, reasonably sized, essential information that is highly relevant to their cognitive needs, such as effective indications and warnings[12] that pertain to their current situation and mission. Responding to each thing that demands the human's attention, and to each piece of data that seems vaguely interesting, is not a feasible option in the context of the IoBT. In fact, a key risk of the IoBT is providing human warfighters with inappropriate information that leads—or misleads—to an action with an outcome worse than what would occur without that "information". Besides humans, somewhat similar concerns apply to all intelligent things in the IoBT; for them, too, unless information is useful it is likely to do more harm than good.

To make its information useful, IoBT technologies will have to deal with a large volume and complexity of information that are truly unprecedented in their extent.[13] Arguably, the quantity of data within IoBT will far exceed any likely advances predicted by Moore's Law (exponential increase with a doubling rate of 18 months) and exceed the ever more efficient use of bandwidth in the future. Besides the sheer volume, the complexity of the information will be formidable. For example, levels of abstraction of the information (produced or consumed) will vary drastically between different things. Similarly, trustworthiness and value of information arriving from different things will be highly variable.

The very foundations of information theory will need to be reconsidered[14]; for example, ensemble probability densities are foundational for information theory and require the underlying process to be ergodic. However, the IoBT is expected to have nonlinear dynamic processes that are sufficiently complex to generate events with nonergodic statistics. The information entailed by the occurrence of such

[10] Kranz M, Holleis P, Schmidt A. Embedded interaction: interacting with the internet of things. IEEE Inter Comput. 2010;14(2):46–53.

[11] Kott A, editor. Battle of cognition: the future information-rich warfare and the mind of the commander. Westport (CT): Greenwood Publishing Group; 2008.

[12] Salerno J, Hinman M, Boulware D. Building a framework for situation awareness. Rome (NY): Air Force Research Laboratory Information Directorate (US). 2004.

[13] Miorandi D, Sicari S, de Pellegrini F, Chlamtac I. Internet of things: vision, applications and research challenges. Ad Hoc Net. 2012;10(7):1497–1516.

[14] Li S, Li DX, and Xinheng W. Compressed sensing signal and data acquisition in wireless sensor networks and internet of things. IEEE Trans Ind Infor. 2013;9(4):2177–2186.

events must be based on single time series and not on an ensemble of time series.[15] Furthermore, nonintuitive, novel phenomena may emerge in the transfer of information between dissimilar large networks. An example would be in how SA is modified by the information exchanged back and forth between the IoBT and the social network of human warfighters.[16] Such unexpected phenomena may also influence—in yet-unknown ways—the ability of warfighters to control, inform, and be informed by the IoBT.

On the other hand, humans are able to adapt to the complexities and dynamics of real-world operational environments to a degree unmatched by current physical or virtual forms of autonomy. As a result, system integrators have developed a wide range of approaches to allow or even require humans to make critical decisions and adjustments within complex, dynamic environments. With relatively few exceptions, these approaches generally situate the human at the apex of a hierarchy and/or as operators of systems. Conceptual and technological advancements in fields such as Human Computation are challenging these prototypical approaches to modeling the role of humans in decision-making processes. Future research may fundamentally change how humans and autonomy are integrated to make decisions. Such areas of research will likely include, but not be limited to, distributed decision making, human computation, and nonhierarchical approaches to heterogeneous-agent systems that will underlie dramatic transformations in how future warfighters interact with intelligent systems and the critical roles they perform in the joint decision-making process. Transformations should be envisioned that recast our perceptions of the capabilities of humans within systems.[17]

Still, at the very least, the IoBT’s colossal volume of information must be reduced to a manageable level, and to a reasonably meaningful content, before it is delivered to humans and intelligent things. A likely target for compression and fusion of data into information, the meeting’s participants conjectured, would be by a factor of 10^{15}. One approach to such a challenging fusion task is to populate the IoBT with a layered hierarchy of information brokers,[18] or “concierges”, that would aggregate, fuse, interpret, and deliver appropriate information. The fusion process[19] should begin at the lowest possible level; for example, whenever possible, all information-

[15] West BJ, Grigolini P. Complex webs; anticipating the improbable. Cambridge (UK): Cambridge University Press; 2011.

[16] West BJ, Turalska M, Grigolini P. Network of echoes; imitation, innovation and invisible leaders. New York (NY): Springer; 2014.

[17] McDowell K. ARL–HRED. Private communication. 2016 Feb 1.

[18] Korzun DG, Balandin SI, Gurtov AV. Deployment of smart spaces in internet of things: overview of the design challenges. Internet of things, smart spaces, and next generation networking. Berlin (Heidelberg): Springer; 2013. p. 48–59.

[19] Kott A, Singh R, McEneaney WM, Milks W. Hypothesis-driven information fusion in adversarial, deceptive environments. Info Fus. 2011;12(2):131–144.

producing things should be equipped with the means to perform locally a degree of filtering, interpretation, and fusion before sending data to the network. Although such layers of intermediaries do complicate or restrict the discovery of underlying data, it may be a necessary price to pay for arriving at useful, manageable, and meaningful information.

However, for information brokers to do their job, they need to know what constitutes useful information. Where would such knowledge come from? One source could be mission planning and rehearsal that could help determine what information is required by the mission-performing agents (human and artificial) and what is the likely available information. To capture the resulting knowledge, a machine-interpretable, formal, broadly applicable and military-relevant language will be needed to express informational needs in a highly heterogeneous IoBT.[20] Moving beyond the inevitable limitations of mission planning and rehearsal, IoBT approaches will need approaches to self-learning (and appropriate relearning)[21] of what information is needed for particular warfighter(s) and particular missions. Such approaches will likely require a form of integration of machine learning and semantic knowledge-based techniques.

More generally, executable models of the IoBT and its surrounding world are needed to enable validation, interpretation, fusion, and assessment of trustworthiness of the information.[22,23] Large-scale simulation may help large-scale sensing and interpretation of information in a targeted, purposeful manner. The research on formulating and automatically creating (and dynamically maintaining) such models is in its infancy. Effective solutions to this challenge will likely involve distributed self-modeling, self-calibration, and self-validation of the IoBT.

Dealing with Deception and Adversarial Nature of IoBT

Nothing differentiates IoBT from IoT more than the battle—the B in IoBT—against a determined and lethal enemy. The adversarial nature of the environment is the primary concern in the IoBT's life. The enemy threatens physical survival and functioning of the IoBT with kinetic, directed-energy, and electronic attacks against its things, by jamming the RF channels, by destroying fiber channels, and by

[20] Barnaghi P, Wang W, Henson C, Taylor K. Semantics for the internet of things: early progress and back to the future. Int J Sem Web Info Sys. 2012;8(1):1–21.

[21] Ning H, Liu H. Cyber-physical-social based security architecture for future internet of things. Adv Inter Things. 2012;12(01):1.

[22] Cho JH, Swami A, Chen IR. A survey on trust management for mobile ad hoc networks. IEEE Comm Surv Tutor. 2011;13(4):562–583.

[23] Kwok SK, Ting JSL, Tsang AHC, Lee WB, Cheung BCF. Design and development of a mobile EPC-RFID-based self-validation system (MESS) for product authentication. Comp Ind. 2010;61(7):624–635.

depriving the IoBT of its power sources. The enemy also threatens the confidentiality, integrity, and availability of the information within the IoBT by electronic eavesdropping and by deploying malware into it.[24] Finally, and perhaps most importantly, the enemy attacks the cognition of human warfighters. Humans will be elements within the IoBT that are most susceptible to deceptions, particularly to those based on cognitive and cultural biases.[25] Humans' IoBT use will be handicapped when they are concerned (even if incorrectly) the information is untrustworthy[9,26] or that some elements of the IoBT are controlled by the enemy. Similar susceptibilities, in part, apply to artificial intelligent entities.

Among the top priorities will be to minimize the enemy's opportunities to acquire information about the IoBT and the warfighters it serves. While many of the applicable measures are the same as those for conventional battlefield networks, the IoBT's exceptional scale, heterogeneity, and density offer additional opportunities for friendly information protection.[27] The sheer quantity of things (especially in those cases when friendly forces leverage the local IoT) permits the use of "disposable" security: devices that are believed to be potentially compromised by the enemy are simply discarded or disconnected from the IoBT. To defeat the enemy's eavesdropping, the defenders may want to take advantage of the plentiful availability of things and inject misleading information into a fraction of them.[28] The density, complexity, and diversity of message traffic within the IoBT will make it more difficult for the enemy to perform the traditional traffic analysis that could reveal details of the friendly command and control structure. Similarly, with a large number and density of things, it may be less expensive and more efficient to stymie the enemy's cyber intrusions by creating large, believable honeypots and honeynets, which are currently expensive to produce and to maintain dynamically—although in the long run a honeynet may be less expensive than the devastation wrought by an adversary's cyber intrusion.

Besides acquiring friendly information (i.e., violating its confidentially), the enemy will attempt to violate the information's integrity by modifying it with cyber

[24] Babar S, Mahalle P, Stango A, Prasad N, Prasad R. proposed security model and threat taxonomy for the internet of things (iot). Recent trends in network security and applications. Berlin (Heidelberg): Springer; 2010. p. 420–429.

[25] Heckman KE, Stech FJ, Thomas RK, Schmoker B, Tsow AW. Cyber denial, deception and counter deception: a framework for supporting active cyber defense. New York (NY): Springer; 2015.

[26] Yan Z, Zhang P, Vasilakos AV. A survey on trust management for internet of things. J Net Comp App. 2014;42:120–134.

[27] Misra S, Tourani R, Majd NE. Secure content delivery in information-centric networks: design, implementation, and analyses. ACM SIGCOMM Information-Centric Networking Workshop; 2012. pp. 73–78.

[28] Bisdikian C, Sensoy M, Norman TJ, Srivastava MB. Trust and obfuscation principles for quality of information in emerging pervasive environments. IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops); 2012.

malware, inserting rogue things into the IoBT, intercepting and corrupting it while in motion between the things, and presenting wrong information to the information-acquiring things (e.g., sensors). The IoBT will likely fight back with anomaly detection that can highlight unexpected data patterns, unexplained dynamic changes, or lack of expected events (*the dog that does not bark*).[29] To enable the anomaly detection, machine-learning approaches will be developed to deal with the data as big and as dynamic as the IoBT will possess. Such a continuous learning process will be computationally and bandwidth-wise expensive. It will be further challenged by the possibility the enemy will adapt and evolve faster than the learning process can. To prevent the enemy from acquiring physical or software modifications of friendly things, approaches will be needed to achieve large-scale physical fingerprinting (e.g., collection of power-consumption patterns) of things and continuous IoBT-wide monitoring of such patterns.[30] More generally, there will be means for active "stimulative intelligence"—ongoing physical and informational probing of IoBT that could help reveal the structure and behavior, including anomalous and suspicious ones, of the IoBT.[31]

Learning normal patterns and detecting anomalous deviations, however, does not work well against a well-designed deception.[32,33] In fact, learning can be a very dangerous double-edged sword with respect to deception. A common approach to deception is for the enemy to cause the friendly forces to learn a certain normal pattern and then perform actions that blend into that pattern but result in an unanticipated outcome. Any measure of normalcy can be defeated by effective deception. Still, the very large scale and heterogeneity of the IoBT may help defeat deception because "lying consistently is difficult"; it may be particularly difficult when the available sources of information are so numerous and are as heterogeneous as in the IoBT. In general, much research is needed on approaches to counterdeception, discovery, or rejection of deception for the IoBT's uniquely complex environment.[24] And, considering that a friendly IoBT will be necessarily connected with the local civilian IoT and thereby to the enemy's IoBT, approaches are needed to execute offensive operations within the intertwined space of friendly and enemy networks.

[29] Raza S, Wallgren L, Voigt T. SVELTE: real-time intrusion detection in the internet of things. Ad Hoc Net. 2013;11(8):2661–2674.

[30] Roman R, Najera P, Lopez J. Securing the internet of things. Computer. 2011;44(9):51–58.

[31] Ashraf QM, Mohamed HH. Autonomic schemes for threat mitigation in internet of things. J Net Comp App. 2015;49:112–127.

[32] Vidalis S, Angelopoulou O. Assessing identity theft in the internet of things. IT CoNVergence PRActice (INPRA). 2014;2(1).

[33] Teixeira A, Dan G, Sandberg H, Johansson KH. A cyber security study of a SCADA energy management system: stealthy deception attacks on the state estimator. Proceedings of 18th IFAC World Congress. 2011;44(1):11271–11277.

Such advanced capabilities will not be possible without new theoretical explorations. Fighting the battle of IoBT may require major new results in game theory, particularly focused on problems with very large numbers and very diverse game moves; near-infinite opportunities for probing; high complexity of utility functions; and partial observability of the game board limited to a very small fraction of the overall space. New theory is needed to formalize and normalize diverse definitions and conceptualizations of risks[34] and uncertainty. Deception should be integral to this theoretical analysis. For example, theoretical results should help predict the appropriate (or counterproductive) degree of complexity for a successful deception.

[34] Kott A, Arnold C. The promises and challenges of continuous monitoring and risk scoring. IEEE Sec Priv. 2013;11(1):90–93.

Cyber Fog

January 7–8, 2016
Adelphi Laboratory Center, Maryland

Organizers: Alexander Kott (ARL–CISD), Ananthram Swami (ST, ARL–CISD), and Bruce J West (ST, ARL–ARO)

Introduction

Organized by ARL, "The Fog of Cyber War" ASPSM took place on January 7–8, 2016, at ALC. The focus of this meeting was to examine the theoretical foundations of the "fog of cyber war" concept for Army battlefield operations. Clausewitz's fog of war spoke of uncertainty in information at a time in history when information was synonymous with knowledge—a situation that no longer exists. More recently, the development of Internet technologies has led to cloud computing which, depending upon the situation, some refer to as a fog rather than a cloud. These seemingly disparate notions of fog merge when one considers how cyberspace is now and will in the future be used in conflict. One possibility to assure friendly networks and information is to maximize the "fogginess" of the friendly information as it appears to the adversary. Networks at the tactical edge—and the tactical information they carry—must be resilient to cyber and electromagnetic operations by a capable adversary; even when partly compromised, they should remain opaque to the adversary and effective for friendly forces.

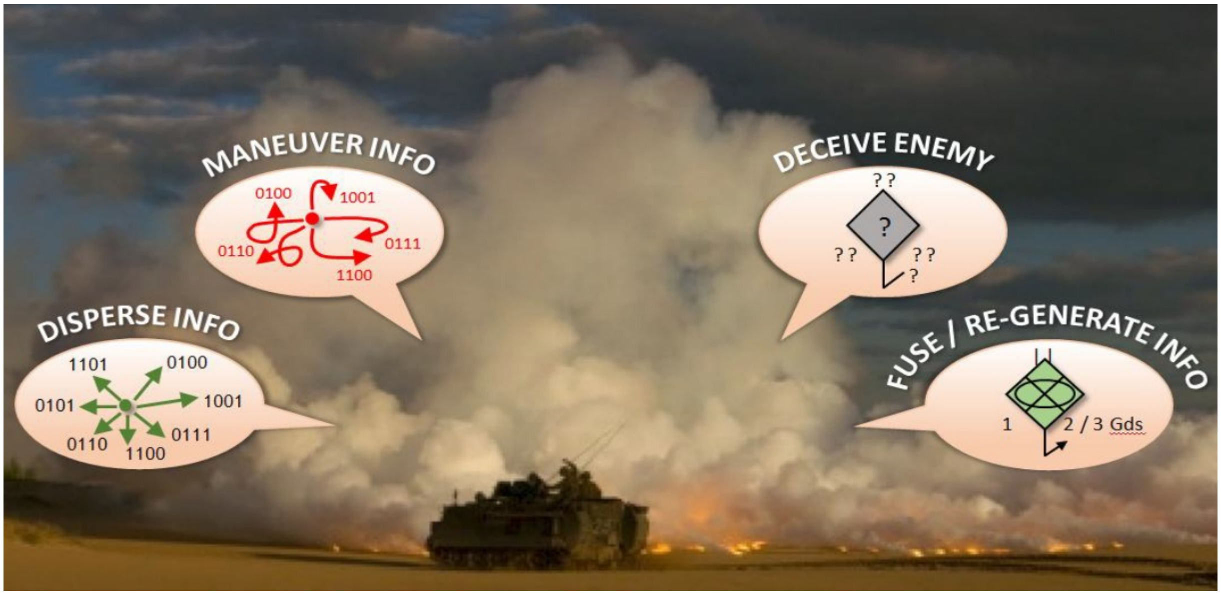

One concept for achieving such an opaqueness—and this was the focus of the ALC meeting—is to determine the consequences of performing radical fragmentation (splitting) of friendly data into a large number of fragments (cyber "fog") and to continually maneuver them across multiple devices of the battlefield edge (i.e., end-user) networks. Data-splitting for security and scalability is practiced in many modern commercial data stores (e.g., Voldemort of LinkedIn, Dynamo of Amazon), but not for tactical, edge devices and networks. With growing interests in Fog Computing and Fog Networks[35] and maturing of edge-network distributed data stores (e.g., GaianDB, product of ARL's UK–US International Technology Alliance), the Army should explore the first tactical use of data-splitting.

While potentially offering a number of military-relevant benefits (e.g., resiliency to adversary's electronic-warfare, cyber, and kinetic attacks and intercept; agile maneuvering of data, rapid recovery, obfuscation, and deception), concepts of this nature also present formidable challenges of complexity of data management and reassembly; demands on bandwidth, storage, and battery power; latency of reassembly; and impact of intermittent connectivity. The excessive amount of data also results in information uncertainty, which becomes crucial in the reaggregation of strategically fragmented data.

The goal of the meeting was to identify fundamental research issues that need to be addressed, which may enable future military-relevant capabilities. Participants were asked to identify gaps in scientific understanding and describe how to apply existing scientific understanding to establish bounds on performance. The meeting encouraged structured yet open and broad-ranging discussion and exploration of multiple perspectives on the issues.

Multiple Topics and Perspectives

The meeting's topics included the following:

- Methods and underlying theoretical models for securing information by fragmenting it and dispersing and moving it in fragmented form across multiple, tactical heterogeneous devices
- Analysis, synthesis, and prediction of behaviors, structure evolution, and emergent phenomena in such highly dynamic systems of information and networked devices; phase transitions; controllability and system identification and state estimation; and role and behavior of human elements in such a system, including the dynamics of human comprehension, trust, and confidence in the system
- Approaches to characterizing tradeoffs of potential benefits and added vulnerabilities—lower vulnerabilities to capture of devices, keys, and data in tactical environments; data exfiltration; loss of availability due to cyber–electromagnetic activity (CEMA) effects; increase in replication without greater danger of data loss to adversary; obfuscation of friendly Electronic Order of Battle (EOB) and portrayal of deceptive EOB; complexity of

[35] Chiang M. Fog networking: an overview on research opportunities. Princeton (NJ): Princeton University. 2015 Dec. [accessed 2016 Mar 1]. https://arxiv.org/ftp/arxiv/papers/1601/1601.00835.pdf.

information management; tradeoffs among computing power, storage, communications bandwidth, energy consumption, and latency

- Formal languages for representation, analysis, synthesis, and provably correct construction of deceptions; techniques for effective, near-automated execution of deceptions against a near-peer adversary who eavesdrops and otherwise attacks the friendly information via CEMA; and theory and methods of control for data flows that allow deceptive modifications of apparent communications patterns
- Computational methods and underlying theory for analyzing and quantitatively managing risk to friendly information, particularly the survivability of information in terms of maintenance of confidentiality, integrity, and availability; and characterizing uncertainty of such assessments
- Approaches to continually assess informational needs of Soldiers from context and mission information; potential enhancements, such as reduction in latency, by learning the user's information-demand model and using the model to modulate the degree of splitting, distance of dispersions, and prepositioning of data
- Approaches and supporting theory for fusion and (re)generation of needs-relevant information from highly fragmented and dispersed data; ensuring high-quality, fused information to friendly forces; and maintenance of SA for the Soldier in spite of extreme volume, dynamics, and dispersion of the information

The following sections capture some of the discussions and findings of the meeting.

Feasibility, Value, and Challenges of Dispersion

Research and practical successes of the database-security community have already demonstrated, to a large extent, the feasibility and value of data dispersion and, to a lesser extent, frequent repositioning of data fragments. Research and successful products exist that use some forms of fragmenting, dispersing, and frequently repositioning data "shards." For example, an industry publication[36] presents a distributed algorithm that uses replication and fragmentation schemes to allocate the files over multiple servers. The file confidentiality and integrity are preserved, even in the presence of a successful attack that compromises a subset of the file servers. Further exploration of the Cyber Fog concept would benefit from interactions and collaboration with the database-security community.

However, dispersion of data is but one of many ways of increasing diversification—and, thereby, uncertainty to the adversary—within a communication system. Other examples include diversification of channels, protocols, and media. Software Defined Networks (SDNs) are

[36] Mei A, Mancini LV, Jajodia S. Secure dynamic fragment and replica allocation in large-scale distributed file systems. IEEE Trans Para Distr Sys. 2003;14(9):885–896.

potentially effective mechanisms for increasing such a diversification. Diversification is helped by applying it over the large scale of the network of devices and channels, which suggests there are benefits in dispersing information not only over friendly networks but also civilian networks and even the adversary network.

The workshop participants frequently mentioned Shamir's Secret Sharing[37] scheme as either a metaphor or an actual component of a Cyber Fog approach. Roughly, a Shamir-like scheme of dispersion may enable sharing of information in such a way that even if the adversary captures a significant fraction of shards, s/he will not be able to reconstruct any meaningful information from it. It was also speculated such a scheme might help balance the bandwidth requirements over time; that is, a bulk of shards could be distributed during the lull in communications demands, while only the final and critical shards—a few—would be sent over the network during the busy periods. Secret sharing, particularly when verification of reconstructed secrets is required, could be made computationally efficient.[38]

Granted, challenges of dispersion are formidable. First, there are the obvious challenges of developing, validating, and managing the complex mechanisms required to perform dispersion with desired effects. Increased diversification (e.g., dispersion of data) and the complex mechanisms required to manage the diversification also create new attack surfaces and venues for cyber attacks. For example, if SDN is used, a centralized SDN is a single point of failure; thus, a more complicated, distributed SDN will be needed. In particular, a Cyber Fog approach may increase a network's vulnerability to availability attacks, even as it improves its resilience to confidentiality attacks. Therefore, a complex tradeoff between availability and confidentiality may need to be managed in real time depending on mission and circumstances of the friendly forces.

Consistency, too, is complicated to achieve (e.g., updates across the system) in this scheme, although local consistency may be easy enough. Increased diversification also makes it more difficult to ensure that friendly users obtain all of the information they need; this could be mitigated by relying on the background knowledge that friendlies have and adversaries probably do not have. We subsequently discuss this point in detail.

Dispersion and Effective Regathering

Dispersed information will be eventually requested by users and will have to be regathered—in a timely and efficient fashion—and regenerated into a useful form. This could be helped by intelligent dispersion: put shards where they are more likely to be accessible at the time when they are more likely to be needed by the users. One way to achieve improved regathering of information is to account for the semantics of information while splitting it into shards. This could help

[37] Shamir A. How to share a secret. Comm ACM. 1979;22(11):612–613.

[38] Subbiah A, Blough DM. An approach for fault tolerant and secure data storage in collaborative work environments. Proceedings of the ACM Workshop on Storage Security and Survivability; 2005. p. 84–93.

intelligent prepositioning of related shards. While doing so, care must be taken not to introduce some regularities into the dispersion scheme that would make it easier for the adversary to find and gather that information. In fact, this creates a tradeoff between data security and the anticipated availability of the data. (A Client-defined privacY-protected Reliable cloUd Service [CYRUS],[39] for example, ensures user privacy and reliability by scattering files into smaller pieces across multiple clouds so that no one cloud can read users' data; an algorithm selects clouds from which to download user data so as to minimize latency.)

To determine which data are more likely to be required by the user and when, it is important to have means of automatically determining relevance of information to the user. Cyber Fog complicates determination of relevance: Unlike in a conventional files system where File A is likely to be relevant to the same issue as File B in the same folder, colocation of 2 shards tells us nothing about their common relevance.

Data provided to the user must be not only relevant but also timely. Timing issues are also complex: The way a collection of information is dispersed (e.g., how small the shards are and how far they are dispersed) depends on when and how rapidly these bundles of information will be needed by the user. Not only such real-time tradeoffs of security versus timeliness are complex, the timeliness is even difficult to define—*I need message M by time T* or *How much of message M do I need to have by time T and still derive sufficient value from M?* The timeliness-versus-security tradeoff is dependent on the nature of the mission: If security only needs to be maintained for a short period of time, it may be acceptable that an adversary has a higher chance of obtaining the information. Researchers[40] have considered 1) how to geographically distribute fragments and replicas so as to minimize expected latency for retrieving data and 2) how to optimize a utility function, which incorporates both aggregation latency and storage overhead.

Consideration of timeliness also depends on the intended or likely purposes of the data: whether the data are needed for real-time execution (in which case the data need to be dispersed in a way that allows rapid and reliable regathering) or the data are intended for postoperation analysis, in which case they can be dispersed with less care for rapid regathering. Network structure and characteristics—such as the network's profile of connectivity and the network's diameter—also influence the optimal ways of dispersion. In some cases, timeliness can be improved by avoiding regathering (e.g., using distributed analytics to obtain the desired answers without regathering the shards).

[39] Chung JY, Joe-Wong C, Ha S, Hong JWK, Chiang M. CYRUS: towards client-defined cloud storage. Proceedings of the Tenth European Conference on Computer Systems (EuroSys'15); ACM. 2015.

[40] Bilbray K, Sigelbaum D, Blough DM. GeoShare: experience with a geographically diverse cloud data storage service. Georgia Tech CERCS Technical Report; 2015. Report No.: 15-02.

Situational Awareness and Information Semantics

Ultimately, it is SA that is the goal of information, and even a timely and relevant information delivery does not guarantee high-quality SA. For one, not all shards are equally valuable from the SA perspective. A shard could be used for creating multiple different pictures or drawing multiple conclusions, depending on how the shards are “glued” together for SA purposes—does it make that multipurpose shard more or less valuable? To a large extent, SA presents us with the problem of not merely regathering but also discovering information: Where and how do I find what I need to achieve the required SA; which shards need to be collected to get the right SA? A centralized or distributed indexing scheme may be required to help this process; however, this too introduces security concerns. Novel methods of information fusion will be required to achieve adequate SA, especially when regathering is incomplete due to an adversarial action or network failures.

Semantics are significant for successful SA formation. Consider a game where the players are given a few letters and asked to guess a phrase to which the letters belong. As the players are given more and more letters in the phrase, they eventually recognize the phrase. The lowest fraction of letters, when recognition becomes possible, is called here the “phase transition”. Note that the phase transition is significantly lowered when the phrase is familiar to the players or when they know something about the phrase. Thus, background information or context matters. Knowledge of the semantics of the information and the semantic context of the information are highly influential on how correctly the information is understood by the recipients. Ideally, phase transition should occur rapidly for the friendlies (who possess background information) and less rapidly for the adversary (who presumably does not possess such background information).

To reiterate, semantics of information—including the dispersed data, the semantic context of the friendlies’ mission, and the background knowledge of the users—are critical for effective and accurate “defogging”. Semantic information theory seems highly relevant to challenges of Cyber Fog. Sheaf theory was mentioned as relevant in this context.

Mission context is particularly important because the success of the mission is the true measure of goodness. An adversary may need only very little information to disrupt a key element of the mission. Thus, understanding of the value of information is critical. The dispersion, the regathering, and SA-formation processes must be designed and executed in a way that information has high value for the friendlies and low value for the adversary. This implies, inter alia, the need for a thorough knowledge (model) of the adversary’s intent and prior knowledge.

Risk and Mission

Risk could serve as a comprehensive framework for characterizing the “goodness” of Cyber Fog. It is recognized that Cyber Fog scheme could potentially increase risk in certain aspects and decrease it in others. Because poorly understood and modeled phenomena like obfuscation and deception play important roles in Cyber Fog, new risk models are unquestionably needed.

Although it is tempting to formulate risk in Cyber Fog in terms of data (e.g., a fraction of data captured by the adversary), it would be misleading. Rather, risk should be analyzed in terms of impact to the mission. Consequences of failures of Cyber Fog should be best assessed in terms of its consequences to the mission objectives. This implies a need for an adequate model of a mission (including its dependencies on network and computing assets)—a modeling problem that is known to be highly complex. Other complexities arise in seeking ways to measure (quantify) consequences to the mission. Some may be indirect and involve impact on the adversary: how much the adversary lost or invested; how long our deception story holds, and so on. Time plays a role: the same event can have very different consequences depending on its timing, and time decay of the importance of the information may be involved (loss of dated information could be less important than that of freshly obtained information). Additive properties of failures, such as information losses, are important too; such as, if you know A and B—high value; if you know A or B—zero value. Uncertainty of failure increases risk: If I know I lost data A, I can decide and do something about it; but if I am uncertain, my effectiveness is impaired.

Understanding the risk to mission in an adversarial environment could clearly benefit from a game-theoretic treatment. Risk is highly dependent on the decisions and actions of the opponents, who are interdependent. The game here is far from classical. It deviates strongly from the traditional zero-sum game; conducted under partial information, bounded rationality, and so forth. In fact, even the mission itself—that is, the goals of the game—can be subject to change if some supporting assets fail or are captured by the adversary. Further, this is a game involving deception.

Deception and Obfuscation

Both dispersion and obfuscation share the key idea: increase uncertainty (to the adversary) through increased diversity. Arguably, dispersion helps to perform obfuscation and possibly its stronger form: deception. The workshop participants discussed possible differences between obfuscation and deception. One interpretation suggests that while obfuscation intends to present the adversary with information that leads to multiple, seemingly equally possible interpretations, the deception aims to present the adversary with information leading to a specific interpretation beneficial to the friendlies. Very little rigorous, quantitative research has been directed at either deception or obfuscation. In the following, we use the term deception implying both deception and obfuscation, unless only obfuscation is discussed.

Within the Cyber Fog concept, deception may take multiple forms. Merely the dispersion and frequent repositioning of information by itself presents the adversary with uncertainty as to where s/he could find information relevant to their interests and how to reconstruct it from the shards s/he captured. Examples of other types of deception include presentation to the adversary of false software and hardware vulnerabilities, thereby inducing the adversary to expend efforts and resources on unsuccessful attacks. Diversity of channels helps deception (e.g., the deceiver could use one channel for real communications and another for deception). SDN could be used to present the adversary with a misleading view of the network. Honeypots and honeynets deceive the adversary as well. This may include cyber–physical honeynets, such as a honeynet that looks to the adversary like a friendly tank.

Still, even with multiple ways to create a deception, it is hard to create a deception that is believable. For example, creation of believable battle plans and other unstructured documents is very challenging. Also very challenging is the problem of creating believable network traffic that the adversary's traffic-analysis mechanism would perceive as a particular EOB of the friendlies. Other examples include placing false shards into the fog; designing believable feint attacks that effectively support a real attack; and generating complex multistep deceptions. Creating a believable deception is harder when the adversary observes the friendlies along multiple dimensions (physical movements, cyber activities, etc.) and when the friendlies are uncertain what the adversary can actually observe. Machine-learning techniques might be applicable to generating believable deceptions. Parenthetically, because in Cyber Fog the adversary is likely to spend more efforts discovering the desired information of the friendlies, the deceiver might benefit from observing these efforts and determining better ways to formulate the deception.

Counterdeception (discovery of a deception) is no less challenging. Much research is needed to determine fundamental limits on counterdeception, as well as actual techniques for performing counterdeception. Detection of deception might be assisted by the fact that a deception—purposeful human creation—is likely to be far less complex and rich in details than real-world information. Lessons might be learned from work on code deobfuscation, such as truth-maintenance approaches. Machine learning might also be applicable to detection of anomalies indicative of a deception. However, sophisticated adversaries may specifically target machine-learning techniques to defeat them. If so, research is needed on limits and verification of how a particular classifier (machine learning) can be fooled by particular inputs.

As mentioned earlier, deception requires game-theoretic approaches. Examples of highly challenging and poorly studied issues are payoff function or metrics of goodness for a deception; modeling of deceivers' behaviors; and modeling of humans (and humans with computational tools) who are targets of deceptions. Considering that the battlefield of the future will be populated by many artificially intelligent systems, it is important to study how artificial intelligence (AI) and human differ (or not) with respect to perceiving a deception.

Applicability of Formal Methods

Given the extreme challenges and complexities inherent in the world of Cyber Fog, designing tools and planning specific activities within such an environment may greatly benefit from formal methods. If successful, such formal methods would assure the friendlies that their environment and plans are guaranteed to exhibit certain properties. Unfortunately, the current state of capabilities in formal methods presents a number of limitations. For example, formal methods suffer from lack of insights into formulating the right questions to ask; that is, which property to verify. Formal methods developed for one domain do not transfer well to another domain (e.g., methods developed for verification of hardware do not transfer well to verification of software and even less so to verification of deception plans). Some difficulties can be mitigated by designing structures that lend themselves to formal methods; for example, some language primitives lend themselves to verification by formal methods, check points in software make it easier to verify formal methods, and so on. Perhaps it might be possible to create a Cyber Fog that would lend itself to formal methods.

Furthermore, it is unknown how well, if at all, formal methods apply to human factors such as the role of cognitive factors in deception. It may be possible to prove formally the consistency of a deception story but it may not be possible with respect to the cognitive aspects of that deception. If formal proof of deception may not be possible for a human receiver, one might conjecture that it might be possible for an AI system that is a receiver of a deception. A possible starting point in research that explores applicability of formal methods to deception could be a problem of proving that a deceiver is producing and delivering to the receiver a picture that the deceiver intended and that meets the deceiver's specification.

Other References Suggested by the Workshop Participants

Subramanian N, Zalewski J. Quantitative assessment of safety and security of system architectures for cyberphysical systems using the NFR approach. IEEE Sys J. 2014;10(2):397–409.

Subramanian N, Drager S, McKeever W. Designing trustworthy software systems using the NFR approach. In: Akhgar B, Arabnia H, editors. Emerging trends in ICT security. Cambridge (MA): Elsevier Inc.; 2014. p. 203–225.

Médard M, Sprintson A, editors. Network coding: fundamentals and applications. Oxford (UK): Academic Press; 2012.

Di Mauro A, Mei A, Jajodia S. Secure file allocation and caching in large-scale distributed systems. Proceedings of the International Conference on Security and Cryptography (SECRYPT 2012); 2012 Jul 24–27; Rome, Italy, p. 182–191.

Individualizing Technology for Effective Teaming

December 1–2, 2015
Adelphi Laboratory Center, Maryland

Organizers: Dr Kaleb McDowell (ARL–HRED), Dr Brett Piekarski (ARL–Sensors and Electron Devices Directorate [SEDD]), and Dr Brian Sadler (ARL–CISD)

Introduction

On Tuesday and Wednesday, December 1 and 2, ARL hosted an ASPSM on "Individualizing Technology for Effective Teaming". This meeting focused on the incorporation of individualized and adaptive approaches to optimize human–human and human–system teaming. It is understood that humans are unique and dynamic, technologies can be individualized, and groups have emergent states and behaviors. In this forum, we asked if technologies can be individualized to dynamic human states and behaviors in a systems-based context to optimize the emergent properties of teams. The objective of the meeting was to identify critical midterm (approximately 2020–2030) basic research issues that would enable such an "individualized teams" concept and the potential impact of successfully addressing those issues.

A fundamental gap exists in our understanding of how mechanisms to individualize or adapt technologies influence the emergent properties of teams, particularly in more complex teams including diverse, rotating, and/or distributed members and heterogeneous agents (e.g., human, virtual, robotic). The meeting was developed to provoke thought and spur discussion of the future critical research issues to address this gap and the potential impact of such future research. Specifically, we asked a select number of experts spanning a wide range of fields to share their perspectives and have an open discussion of ideas. To meet the meeting's objectives, we took 2 different approaches: first, identify critical research issues for the 2020s timeframe and articulate the potential of the research to overcome specific technical barriers or enable specific technologies; second, identify potential technologies of 2040, the technical barriers to the development of those technologies, and link research issues to the barriers and technologies.

2020s Research

On the first day of the meeting, the conversation was initiated with presentations of potential research topics from experts spanning a wide variety of backgrounds. These presentations were intermingled with small-group "2020s research" breakout discussions. For these breakout sessions, 3 teams of researchers each discussed 1) identifying a future perspective or vision, 2) identifying midterm (approximately 2020–2030) basic research issues that would support individualized and adaptive approaches to optimize human–human and human–system teaming, and 3) articulating the potential of the research to overcome specific technical barriers or enable specific technologies. The following summarizes the discussion of the 3 teams.

Blue Team

Perspective

The vision of future teaming was characterized by several critical factors that are expected to influence team performance, including teams comprised of members spread across large distances and facing multiple distractions; an increase in nonhuman agents of increasing capabilities; and an increase in communication modalities between humans and between humans and technology. The vision of highly distributed teams faced with large informational challenges has a significant potential to threaten optimal team performance, intensifying effects that include social loafing, social isolation and withdrawal, diffusion of responsibility, lowered executive function, and sleep fragmentation, among others.

Research

Two specific approaches were proposed:

- A human-autonomy interaction line of research was proposed that focused on evaluating the effects of agents and modalities on social isolation in human team members. The goal of this research focused on developing an understanding of social isolation in distributed, heterogeneous human–agent teams that could underlie the development of early social-isolation detection systems, prevention approaches, and mitigation strategies.
- The second approach focused on driving research through a human-autonomy “Social Olympics” to evaluate the effects of agents and modalities on threats to team performance and on mutual adaptability among humans and agents. The approach was envisioned to take a much stronger human-centric approach than is observed in current robotic grand challenges, including examining team performance under real-world conditions such as stress, sleep deprivation, long periods of task performance, and a range of operator experiences and training, among others.

Barriers

These research approaches are expected to overcome 1) the lack of systems for evaluating of nonhuman-agent integration on social brain mechanisms, 2) the lack of systems for detecting early threats to team performance, particularly socially based threats, and 3) the lack of interventions to act on detection of problems in either of the 2 previous barriers.

Yellow Team

Perspective

Several interrelated perspectives of future teaming were described:

- The first focused on the notion of self-expansion, illustrated by a squad being composed of a single individual with many virtual and/or physical systems. This perspective includes concepts of mutual adaptation between human and intelligent agents, systems' understanding of cognition and effect on team performance, and intelligent team composition and accelerated learning. The perspective also assumes the human–intelligent-agent relationships will progress from their current state (e.g., pilot/autopilot) to ones more similar to that of human dyadic relationships.

- Future Soldiers or "expanded" Soldiers would be enabled by a networked, virtual "Coach" or "XO" ("executive officer") that had the capabilities to develop and maintain enhanced individual and team performance. These digital assistants would combine sophisticated models of humans and team behavior, task requirements, and environmental factors with feedback and real-time, historical, and third-party data streams to support and guide team coordination and activities.

- In a highly technologically immersed future, a shift was envisioned in the factors that future system designers, industrial/organizational psychologists, and others will consider when influencing team development for either human–human (mitigated through technology) or human–agent teams. To reflect today's focus, one can imagine a pyramid with organizational factors (coordination, interdependence of tasks) on the base; cognitive considerations (mental models plus anticipation, cognitive states) in the midlevel; and affective and social factors (social belonging, evolution) on the tip. In the future, the focus on these factors is envisioned to be flipped with affective and social factors as the base. This shift reflects the expected proliferation of virtual and physical intelligent agents and the need to for greater buy-in for these forms of autonomy.

Research

A long list of research challenges was identified, including

- control theoretic views on team functions;
- understanding how/what to communicate from human to machine and vice versa;
- expand content of models from organizational to cognitive and cognitive to socio-emotional;
- increased sensemaking out of sensor-driven data with a focus on interpretation and translation to augmentation;
- sensing technology to reliably assess deeper cortical structures to move toward a joint understanding of cognitive and socio-emotional processes;

- formal theories of trust that are multidimensional and prescriptive, have quantitative value metrics, account from mutual trust in cohesive heterogeneous teams, and include timescales of trust development/adaptation;
- developing intelligent agents that are transparent to humans and capable of metacognition;
- the longitudinal assessment of human–agent interaction;
- affective aura including brain–machine interfaces and machine-mediated human–human communications; and
- understanding how to effectively leverage data-driven approaches to augment theory, optimize team building, and improve learning.

Red Team

Perspective

Dual perspectives were offered. A human-centric vision focused on developing deep, rich insights into human behavior through lifelong individual models and data that were able to integrate and interpret information across multiple sources including brain, internal to body, external to body, and situation and context. Individual models were envisioned to feed team models in support of accelerated team training, team maintenance, and team repair. An autonomy-centric vision depicted a nonlinear advancement in the human-like nature of autonomy (image, movement) increasing in the 2020s, dipping in the 2030s in response to negativity associated with the uncanny-valley issue, and dramatically increasing in the 2040s. These perspectives were merged in future heterogeneous teams of humans and intelligent agents, which were guided by models and data, to coordinate on a subsecond time scale to enable action, attention, and skill.

Research

Several research areas, potentially beneficial fields of research, and recommendations were proposed, including

- Research areas included understanding how to "engineer" brains to improve team performance/training transfer; uncovering the neural basis of why it takes N years to learn, why people have off days and slumps, why people "choke"; and understanding how to engineer team interactions to improve performance.

- Potentially beneficial fields of research discussed included parenchyma, neurophysiology, neuromorphology, whole brain network-connectivity analysis, brain area maps, microcircuits patterning, micromuscular resource control, and genetics.

- Recommendations included testing the interpretive boundaries using current sensors (e.g., EEG, EMG, gaze, facial expression, body pose, intentions) over large datasets and long durations; examine neuroscience in the wild; embrace machine learning and publicly available databases; focus on longitudinal learning using multidimensional (e.g., brain, behavior, performance) data; and refocus to computational neuroscience.

2040s Vision of Teaming

On the second day, the teams were mixed and 3 new discussion groups were formed that focused on envisioning the future of human–human and human–system teaming and how technology may mediate that teaming. Specifically, the discussions were focused on 1) identifying technologies of 2040, 2) discussing underlying, specific technical barriers or enabling technologies, and 3) linking research issues to those barriers and enabling technologies.

Orange Team

2040 Technology

An AI team "Coach" was envisioned that could function as a life coach. The AI was foreseen to have capabilities to facilitate team bonding and interactions, find and track "sweet spots" for perturbation and interventions to optimize team performance, optimize human–machine interactions through techniques such as mimicking human communications, and tailor guidance and interactions to the team makeup and account for critical factors such as experience and background, personality, human/agent capabilities and states, and culture. An AI Coach would enable large teams to optimize over competing and common goals (e.g., Paris Climate Summit or US Congress) and enable the rapid formation and reformation of teams (e.g., dynamically create multidisciplinary scientific teams).

A second future technology was dubbed "KITT" after the AI embodied in a highly advanced, autonomous automobile in the "Knight Rider" TV series. This application illustrates a dyadic relationship between a human and agent. The core 2040s technology envisioned was the ability of the agent, "KITT," to interpret the situation; the past, present, and predicted state of the human; the human's intent; and the goals of the multiple ongoing tasks—and then engage in joint human–agent decision-making processes that lead to actions, which optimize across short- and long-term team performance. Achieving this technology will require overcoming current barriers in effective, efficient communications; mutual-trust formation; and robust prediction.

2040 will also see alternative team makeups and structures compared to today. For example, the concept of a Facebook "friend" has led to a reconceptualization of peoples' relationships compared to 2 decades ago. As technology advances, it is expected that new, unforeseen, teaming relationships will form between diverse teams of humans and technology.

Research

Critical research areas to all 3 of the technology areas discussed include 1) the evolution of AI as an effective teammate of humans, with an emphasis on emotional bonding and social interactions;

2) development of predictive technologies that integrate models and data; 3) the need for the longitudinal study of team development and dynamics; and 4) the need to move beyond current conceptions of human–robot interaction to the study of human–agent dyadic relationships (2020s) and larger teams (further in the future).

Light-Blue Team

2040 Technology

Four technology areas were highlighted:

- Agents functioning as specialized assistants to identify and augment where Soldier skills are lacking (training focus). Technologies that integrate the biological system (e.g., genetic, biochemistry, neural) are expected to have capabilities such as enhancing team affective and cognitive processes, reducing the time to develop super learners and super performers, and increasing the effectiveness of transfer of knowledge.

- Virtual assistants designed for automated human–agent team composition and effectiveness (operational focus). Concepts such as virtual XOs, information sources, or Commanders illustrate the types of technologies that can assess knowledge, skills, abilities, and other characteristics (KSAOs) across human and agent team members and automatically develop the right mix of KSAOs for any task in real time; implement technological augmentations or employ social agents in operational settings; amplify information (affective, cognitive) internal to team that is critical for performance and cohesion; and merge affective, cognitive (intentions), behavioral, and environmental information to build, maintain, and arbitrate team SA.

- Technology aimed at enabling persuasion and influence for small teams to larger groups.

- A set of super-team technologies aimed at synching emotion, affect, and cognition to enable better, faster communications, and task diversity. Example technologies include increased fully autonomous agents, autonomy with increased socio-cultural intelligence, reduction or elimination of verbal communications (brain–brain communications), and integration of biochemistry/neurochemistry signatures of Soldier emotional and cognitive function.

Enabling Technologies

Three categories of enabling technologies were identified: 1) technologies for influence—machine language translation, natural language translation, natural language processing, communications translation, multiaspect human interpretation (e.g., language, affect, posture, voice, eye), and cultural interpretation; 2) technologies for team resilience (adaptation to failures, localization of resources); and 3) technologies for connectedness (embedded, networks of multiple humans and multiple machines).

Research

A long list of relevant research areas was identified, including

- neurobiology of expertise;
- the interaction between expertise and teams;
- team interactions and dynamics through a multidisciplinary approach;
- computational theory of teams that accounts for neural, biochemical, behavioral, and social components;
- characteristics critical to teams;
- the relationship between biochemistry and neuroscience specifically in socio-emotional/affective contexts;
- the effects of human stimulation on learning and augmentation;
- managing human attention;
- human-autonomy interactions for combative aspects of war;
- team interactions with populace;
- brain-to-brain interfaces (mitigated by technology);
- dynamic resource allocation models; and
- social neuroscience.

Green Team

2040 Technology

Accelerated evolution of new “Hybrid” organisms: Driven by task selection, hybrid organisms would be formed by initially assembling individual Soldiers with small, building-block artificial agents with specific KSAOs (imagine intelligent LEGOs). These initial assemblies would form a fuzzy, intuitive starting point from which rapid self-organization would allow for adaptation to the specific task demands. The goals of these technologies would be to provide new “organisms” (i.e., a human augmented by a set of intelligent building blocks) capabilities that do not currently exist (e.g., new appendages for physical tasks; super-predictive powers for high-level decision makers; instantaneous cultural knowledge for ambassadors). These organisms could be considered future teams themselves, or the principles of self-organization and adaptation could be extended to teams of organisms.

Research

Two components were emphasized: research underlying the base building blocks and research into accelerated evolution. The first recommendation focused on bringing together complex, multidisciplinary teams of researchers to create novel human–computer interactions for creative, small building blocks. The second recommendation focused on research in "organism evolution" from the base building blocks. Suggestions included investigating dynamic adaptation with both bottom–up (trial and error) and top–down (model) approaches; and a high-frequency experimental paradigm mixed with other mechanisms such as crowd–sourcing.

Attendees

Speakers

Chris Atkeson (Carnegie Melon University); John Cacioppo (University of Chicago); Stephanie Cacioppo (University of Chicago); Jamie C Gorman (Georgia Institute of Technology); Tzyy-Ping Jung (University of California, San Diego); John Krakauer (Johns Hopkins University School of Medicine); Stephen Macknik (State University of New York Downstate); Michael Rosen (Johns Hopkins University School of Medicine); Paul Sajda (Columbia University); and Ronald Stevens (University of California, Los Angeles; The Learning Chameleon, Inc.)

Additional Participants

Bill Casebeer (Lockheed Martin); Jay Goodwin (Army Research Institute); Tim Mullen (co-founder and CEO, Qusp/Syntrogi Inc.); Daniel Serfaty (founder and CEO, Aptima, Inc.; chairman, Aptima Ventures, LLC); Marissa Shuffler (Clemson University); Jessica Wildones (Institute For Cross Cultural Management, Florida Institute of Technology); and Diego Zapata-Rivera (Educational Testing Service)

Moderators

Joseph Mait (Army Research Laboratory) and Kaleb McDowell (Army Research Laboratory)

Discussion Facilitators

Arwen DeCostanza (Army Research Laboratory); Piotr Franaszczuk (Army Research Laboratory); Brett Piekarski (Army Research Laboratory); and Brian Sadler (Army Research Laboratory)

Observers

Nora Pasion (Office of the Assistant Secretary of the Army); Dr Thomas Russell (Army Research Laboratory); and Jeff Singleton (Office of the Assistant Secretary of the Army)

Support

Tammy Christenson (Army Research Laboratory) and Gabe Smith (Army Research Laboratory)

Distributed and Collaborative Intelligent Systems

December 3–4, 2015
Adelphi Laboratory Center, Maryland

Organizers: Dr Brett Piekarski (ARL–SEDD), Dr Brian Sadler (ARL–CISD), and Dr Kaleb McDowell (ARL–HRED)

Introduction and Background

Now in the third year of Army Science Planning and Strategy Meetings, ARL hosted a 2-day meeting on December 3-4, 2015, that focused on distributed, collaborative intelligent systems and aimed to identify grand challenges and technical expertise needed to ensure cutting-edge research and scientific impact for Army 2050.

Each day included coarsely defined topic sessions with a series of talks followed by an extended discussion period. It is broadly recognized that intelligent systems are critical and will become pervasive across all Army R&D.

Objective and Scope

Advancements in intelligent and autonomous systems are rapidly changing the state of the art through fundamental research programs within the DOD services as well as commercially, as evidenced by the emergence of driverless cars, the nearly ubiquitous nature of small drones, and the recent defeat of the European Go champion by the Google DeepMind program AlphaGo through the use of deep-learning techniques. These advancements are primarily focused on individual systems and agents. Future systems are envisioned to be highly heterogeneous and collaborative and distributed both spatially and temporally.

Future complex operational scenarios are envisioned such as megacities and subterranean environments with increased restricted- and denied-access areas. It is easily anticipated that distributed intelligent systems might provide key capabilities in these and other cases, with advanced reasoning and very fast operational tempo to maintain a US tactical edge in complex environments. The meeting focused on determining new research areas for the long term to achieve convergence of heterogeneous platforms, distributed computing, advanced networking, sensing, and intelligence.

Overview

Three technical sessions broadly addressed 1) swarm formation and control, 2) large heterogeneous systems, and 3) distributed intelligence. Across the 11 technical talks, speakers articulated emerging trends in the field of intelligent and autonomous systems research, and many accepted a challenge from the workshop organizers to conclude with a few gutsy futuristic predictions to spark a hearty discussion among the approximately 41 attendees from DOD, other government research agencies, academia, and industry research laboratories. From these lively and invigorating discussions, several technical trends and challenges emerged that coalesced into themes identified by the group as research areas likely to have the greatest impact (particularly on DOD and Army) over a 10–30-plus-year time frame.

Topic Descriptions

Swarm Formation and Control

This area focused on determining new scientific opportunities/capabilities that may be possible by the convergence of the fields of air and ground cooperation and collaboration; networking and network security; control and reconfigurability with uncertainty and risk for high-speed swarm formations; human-in-the-loop and human supervision; covertness; environmental complexity; environmental representation, navigation, and planning; and sensing with a variety of modalities and overall system configuration.

Human-in-Loop Control

Traditional Human–Robot Interface has focused on ways to assist humans in controlling robots via interfaces; but with large numbers of robots, the cognitive load is simply too high for this traditional approach. Instead, we are forced to offload more and more responsibility and autonomy onto the individual robot and multiagent team/swarm. This enables the Soldier to focus on more abstract commands, with the swarm collectively following mission command, perceiving and adapting to local variations in the local environment, and anticipating and responding to the Soldier’s expectations. In addition, the Soldiers can act as sensors and advisors within the network of agents providing environmental information and SA when available as well as acting as a coach or intelligent agent within the overall system to provide limited guidance, globally or locally, to improve overall system efficiency. How to do collective multiagent coordination at this level or to include the human as another intelligent agent within the network is an entirely unsolved task—much research lies ahead to make it a reality.

New Sophisticated Control Architectures

The majority of current research in AI and autonomous systems is on single agents, small teams of heterogeneous agents, or large swarms of homogeneous agents. Depending on scale and complexity, current implementations for intelligent systems typically rely on either centralized or decentralized control architectures. There is no general science or architectures for large numbers of distributed heterogeneous agents. Flocking is reasonably well understood, but this is a small piece of the distributed-intelligence problem. More research is needed in new sophisticated, hybrid control architectures for large, highly heterogeneous teams that may include both global and localized control of single agents, spatially and temporally distributed small and large teams, and localized swarm behavior. A key issue will be the abstraction of localized behaviors and local controls for global control; this includes many research topics such as representation, formal methods, finite-state automatons, dimensionality reduction, uncertainty, noise, and adaptive and resilient behaviors. For large heterogeneous teaming, it can be assumed that not all communications will be bi-directional and must be understood in the context of abstraction, roles, and heterogeneity. Some paths forward include coupling abstraction with learning, information theory, and the value of information.

Fast Suboptimal Planning

Finding optimal plans for autonomous agents is often computationally hard, especially for systems in complex environments. For most military operations this is also desired to happen in real time and at the operational tempo of the squad. This problem is only exasperated when it is expanded to large heterogeneous multiagent systems in which planning is coordinated across many heterogeneous systems with varying mission objectives, where individual agents may or may not have the same goals, where some agents may not be able to complete their tasks due to failures, and there exist uncooperative players or adversaries.

One method to address this is the use bounded suboptimal plans. This approach provides a promising alternative that can speed up planning while resulting in plans that still provide reasonable quality guarantees for success. This makes an interesting area to study versus optimal guaranteed planning in multiagent systems and the associated trades in planning speed, mission speed, and probability for operational success.

Resiliency

The resiliency of large multiagent systems needs to be considered based on realistic communication links and localization uncertainties. The most popular distributed control and planning paradigms today simply assume communication links are a given or can be modeled with simplistic “on/off” linkages with unlimited bandwidth. While this enables the use of convenient analytical tools for proving stability/optimality guarantees, such assumptions do not accommodate important sources of error and uncertainties, such as 1) data processing and the reliability and assurance of no loss of information; 2) availability of sufficient bandwidth or ability to work with limited bandwidth; 3) determination of usefulness of data to support decision making; 4) data analysis and integration of analysis with control; 5) data reduction (not limited to dimensionality reduction for analytics, should not process all data by all sensors/platforms); 6) information integration and network availability and robustness (including system level); and 7) signal interference due to the environment or multiple sources, signal jamming, or multipath signaling. Efforts are needed to know what information is needed and not given, to determine what is missing or corrupted, and to collect missing data.

Furthermore, control/planning architectures require accurate inertial or relative localization/pose information that cannot be guaranteed with limited sensing capabilities on size–weight–power-constrained platforms, which may need to operate in naturally GPS-denied environments such as indoors or in urban canyons or which may consist of microvehicles that can only allocate a limited amount of processing power to perception. In essence, scalable control and planning strategies are needed to holistically model “messy” vehicle state and communication uncertainties, rather than separate/ignore these as is typically done in conventional distributed-control/AI paradigms. Some limited success has been achieved using “loosely coupled” approaches to distributed control and sensing (e.g., approximating the main statistical effects of packet erasures along certain communication channels). However, approximations and control/estimation architectures used in these approaches could be generalized to relax assumptions that may hold for specific kinds of platforms or tasks. Future systems will need to address this potential variation and uncertainty or risk completely missing localization and pose information across all of the agents.

Mission and operational resiliency of heterogeneous systems was another topic of discussion. A critical issue is the lack of design methods and models for such systems. Heterogeneous teams will likely require multiscale optimization. In many cases this is also called a “resource allocation” problem—optimize use of available resources, which may be very disparate (e.g., time, energy, computational resources, communications resources, degree of motion required). Morphing, reconfigurable, and adaptable platforms’ and systems’ performance were ways discussed that could offer increased resiliency. For these to be effective, behavior synthesis should be rapid and scalable (“online behavior synthesis”). Learning methods could be applied to reduce needed synthesis, but both of these are complicated by the potential use of many small platforms with low capability.

Multiagent Learning

Multiagent Learning is an attractive alternative to directly coding teams or swarms of agents or robots, particularly since it could be used in the field by nonexperts to train groups of robots to do tasks on the fly (so-called “Learning from Demonstration”) or to model and respond to threats or unexpected environments. A fundamental challenge with multiagent learning is the “Multiagent Inverse Problem”. Imagine you are trying to teach a swarm of robots to collectively storm the castle. Even if it were possible to quantify exactly what you wished the agents to achieve collectively (which is a challenge in and of itself), in order to do learning each agent must instead know what “it” needs to do, not what the collective must achieve. Put another way, even if you can explain the desired macrolevel phenomenon, the agents instead need to know what their individual microlevel behaviors should be. Unfortunately, while it is possible to build a function that provides the macrolevel phenomenon that emerges from many agents performing specific microlevel behaviors (a simulator), it is very difficult or maybe impossible to provide the “inverse” function that provides the microlevel behaviors necessary to achieve a given macrolevel phenomenon. The standard way to overcome inverse problems is to use optimization. While some optimizers, such as reinforcement learning or policy search (essentially hill-climbers), can be used in simple scenarios, when agents become complex and heterogeneous and have complex interactions, these methods are not likely to scale. Instead, we must fall back on last-ditch methods, notably metaheuristics such as evolutionary computation techniques. These methods are very costly, ad hoc, and not well formalized or studied in the context of swarm or multiagent optimization, requiring significantly more research. Additionally, optimization methods require large numbers of trials and samples: this is feasible with a simulator, but in real-world (physical-world) training or learning from demonstration, a simulator is not reasonable and other approaches to gathering data are simply not possible in the physical world. We simply cannot wait for a human to put a robot swarm through 30,000 trials until he or she has figured things out. Thus, more research is needed to find approaches to teach large swarms of complex agents how to do nontrivial, collective tasks in real time and in the physical world—a problem area with only a few recent breakthroughs.

Large Heterogeneous Systems

The focus this session was on determining new scientific opportunities/capabilities that may be possible by the convergence of the fields of tasking and distributed control; networking and communications; dealing with adversarial networks and network security issues; scenarios and operational concepts; extended reach and robustness in complex environments; system architectures consisting of a variety of scales in computation, communication, and mobility (air and ground); and power and energy constraints and management.

Level and Mix of Heterogeneity

Computers, phones, and other devices have moved toward homogeneity in design rather than heterogeneity. Some open questions, then: Would standardization enable advancements in R&D in distributed intelligence and needs for future distributed and collaborative intelligent systems? What is the right mix of heterogeneity in sensing, computation, platforms, levels of autonomy, and human–robot teams? And, what is the right ontology (robotic ontologies exist, but are missing coupling with reasoning, with cognition, and with task allocation)?

New simulation tools are needed to explore and optimize needed levels of heterogeneity. The current game-theory-based formulations are computationally intractable when dealing with teams involving hundreds of agents. The expected speed-up in computational power will be of very limited use in addressing this issue because of the exponential nature of these formulations. How to make sound, distributed decisions in the presence of an intelligent adversary is another open question. It is also unclear what the potential action set of an intelligent adversary might be. We will need to figure out how to develop formulations that are computationally tractable.

Adaptability and Reconfigurability

What is the best organizational structure to offer a balance between resiliency and operational efficiency, and, how to reconfigure teams in the middle of a mission using a distributed architecture? Complex missions require multiple teams to simultaneously carry out multiple tasks. Agents may need to play multiple roles that may span across teams. As contingency situations arise, rapid reconfigurations in teams, both locally and globally, will be needed across the distributed architecture. Another related question is, what kind of heterogeneity is desired? The answer to this question will help in arriving at the right balance of specialized and versatile platforms in the team. High diversity in platforms will create challenges in service and acquisition. These aspects will have to be considered in the system performance.

How to synthesize new behaviors online to deal with the unexpected contingencies? Dealing with intelligent adversaries will force the team into unforeseen contingencies. The ability to generate new behaviors online will be a must to deal with contingencies and for the system to exhibit resilient behavior. Online behavior synthesis is a challenging problem even when done using a central architecture. Doing online synthesis of behaviors in a distributed architecture in a fast-paced mission will be very challenging.

Operational and Experimental Complexity

Given the highly multidisciplinary and complex nature of this research area, experiment-driven research is critical to explore and discover the brittle connections and interdependencies among perception systems, interactions with external data sources, efficient data sharing and processing methods, intelligence and decision-making algorithms, multiagent navigation and collaborative behaviors, and the collective performance of spatially and temporally relevant missions.

There have been recent examples of operating singular, fully autonomous systems in complex environments; small heterogeneous teams with moderate complexity and interactions; and large numbers of homogeneous agents/swarms in simple environments and with limited autonomy (aerial systems with no obstacles and with GPS coverage). To make these demonstrations tractable, researchers typically reduce the complexity along several axes: 1) number of agents; 2) degree of heterogeneity among the agents; 3) agent-behavior complexity, autonomy, and adaptability; 4) degree of interactions and communication among the agents; 5) speed of operation; and 6) complexity of the environment and available infrastructure (e.g., GPS). Performing large-scale demonstrations that push the degree of complexity along each of these scales is not currently possible. Research in ways to simultaneously push the complexity along each of these axes is necessary for success. A lack of design methods and models for such systems is a remaining critical issue as well to reduce the time cycle for technology development and costs related to iterative field testing of large, complex systems. As the degree of heterogeneity increases, so does the design and task allocation complexity. Metrics and roles for heterogeneous elements must be understood.

Some critical issues identified with hundreds to thousands of intelligent systems (these relate to numbers, heterogeneity, communications, and computation):

- Assumption of highly resourced nodes (or, when won’t this be true?)

- Many more interactions than in small numbers
- A massive design space that is combinatorially complex
- The "emergent behavior problem" may be a killer issue
- We know large numbers for some cases, at least for flocking/formations
- Hierarchical architectures make sense, but also may be problematic
- A metacontroller is likely needed
- Adversarial swarms, simple counter-swarm ideas
- Understanding group behavior from observation, and the ability to camouflage intent ("stigmergy spoofing")

Another issue in large-scale operations and experiments, especially for small aerial systems, is that there is no Moore's Law for power and energy. This implies the need for energy awareness and efficient, rapid deployment of small nodes while exploiting heterogeneity in system design (e.g., marsupial platforms, duty cycling of small platforms, hybrid systems, collaborative tasking).

Distributed Intelligence

The focus of this session was on determining new scientific opportunities/capabilities that may be possible by the convergence of the fields of rapid and distributed decision making; distributed perception and uncertainty; human-related aspects such as human–machine interaction, human supervision, and real-time crowdsourcing; tradeoffs in computation, communication, mobility, and power consumption; and exploiting tactical-cloud resources, such as a knowledge base.

Distributed decision making

How to make decisions in a distributed manner that are considered acceptable based on the cost–benefit preferences of the mission commander is an open question. Currently, most robotics researchers develop their own cost functions and show that their distributed decision-making algorithms produce good decisions with respect to these cost functions (in relatively simple environments). These cost functions have certain nice mathematical properties. Unfortunately, a mission commander's preferred cost function might be far from these "nice and well-behaved" cost functions. In such cases the existing algorithm might produce highly suboptimal (and hence undesirable) results. We need to develop methods that work with arbitrarily complex cost functions in complex environments.

Access and use of the cloud, big data, social media, real-world complex models (i.e., weather), and other knowledge bases can be included and leveraged to support intelligent/semantic routing of valuable information or answer critical questions that are unknown beforehand due to the rapid situation change on a battlefield. In future systems, this knowledge base could be very distributed; a key step in intelligent/semantic routing is reasoning based on knowledge base and the current situation. Routing and knowledge base are the 2 aspects of the same problem: Routing relies on local knowledge bases and routed information enriches knowledge bases.

When communication between agents is limited or even completely disrupted, the only way to counter such an adverse situation is to perform reasoning and prediction (i.e., based on the existing local knowledge base, to predict the situation and future movements/decisions of allies and adversaries). Reasoning and prediction are also critical when missions and objectives are not clear or change rapidly in dynamic and complex environments in order for the agents to predict human and mission intent and maintain operational tempo. Similar techniques will also help determine the most critical information to send out when the network bandwidth is limited. A few key technical components/challenges in distributed intelligence are Knowledge Representation, Reasoning/Prediction, Routing and Easy Access for People.

Soldier–Multiagent Collaboration and Decision Making

Aside from difficulties in determining how best to express operational intent and SA to a large group of (possibly heterogeneous) platforms, how should a large network of such platforms even begin to organize themselves to process directives or observations provided by a human? Another challenging problem arises in the form of ensuring that human users/designers can even comprehend the scope or scale of large-scale robotic network capabilities: How does a human commander know what such a network is actually capable of and what he or she should expect in terms of acceptable performance/behavior under different operating conditions? Cognitive science informs us that humans are notoriously poor at making rational decisions involving extremely small/large quantities; so, will users of large swarms be sufficiently aware of the risks involved in completing certain tasks and what the operational consequences of losses are? On the flip side, how will designers/engineers of such systems encode desired behaviors to such systems? Alternative programming and software design paradigms may be needed to encode, simulate, and validate very large-scale robotic systems, particularly if each agent is expected to be highly adaptable and possibly capable of complex autonomous behaviors (e.g., platforms that can "fill in" for each other at a moment's notice, as the remainder of the network reconfigures itself appropriately). Probabilistic programming techniques, for instance, can help drastically reduce the time needed for machines to learn and adapt to new patterns extracted from various data streams (which may include human input) while helping to improve decision-making/perceptual transparency for complex problem solving.

How humans should interact with robots in large, heterogeneous systems is an open question. In many future scenarios we are likely to encounter situations where the numbers of robots are likely to be an order of magnitude higher than the numbers of humans. Humans will be basically interacting with robot "crowds" (e.g., a human might be working with 50 robots). Language-based communication will be useful, but we should examine other modalities as well. Robots can transmit information to humans by showing synthesized scenes in the virtual world. If we make sufficient progress in the brain–computer interface, this might be a possibility. Humans can also send commands to robot crowds using augmented-reality interfaces. It might be much easier for humans to directly manipulate visual representation of information in robots' brain.

Summary

The 4 key focus areas were identified for the advancement of future Distributed and Collaborative Intelligence Systems:

- **Distributed Awareness—**use of distributed perception; use of humans as sources of information; efficient and resilient data sharing; and access to knowledge databases for increased SA across the network of agents;
- **Distributed Intelligence** using deep-learning methods; real-time on-board and off-board simulations; self-awareness; reasoning and prediction; and human awareness and guidance for efficient local and global decision making;
- **Adaptable and Resilient Control**—hybrid control architectures for simultaneous local and global control of large multiagent systems; resiliency in dealing with failures, uncertainties, and uncooperative agents; real-time adaptive and reconfigurable formations and dealing with soft versus hard rules; efficient human interactions; and multimode communications; and
- **Scaling Experimental Complexity** for large-scale experimental evaluation of heterogeneous multiagent networks (scaling complexity in number of agents; level of heterogeneity; agent-behavior complexity, autonomy, and adaptability; degree of interactions and communication among the agents; speed of operation; and complexity of the environment and available infrastructure).

Acknowledgements

While all the participants contributed in a significant way to the discussions and thoughts captured in this report, the authors would like to acknowledge the significant, additional contributions by Nisar Ahmed, Dai Liyi, Satyandra Gupta, Sven Koenig, Sean Luke, and Xifeng Yan toward the writing of this report.

Speakers

Nisar Ahmed (University of Colorado), Stephano Carpin (University of California, Merced), Kostas Daniilidis (University of Pennsylvania), Eric Frew (University of Colorado), Volkan Isler (University of Minnesota), Sven Koening (University of Southern California), Vijay Kumar (University of Pennsylvania), Brian Sadler (Army Research Laboratory), Mac Schwager

(Stanford University), Xifeng Yan (University of California, Santa Barbara), and Tansel Yucelen (Missouri University of Science and Technology)

Participant and Contributors

- Ahmed, Nisar (University of Colorado Boulder)
- Belta, Calin (Boston University)
- Bornstein, Jonathan (Army Research Laboratory)
- Carpin, Stefano (University of California, Merced)
- Cortes, Jorge (University of California, San Diego)
- Dai, Liyi (Army Research Laboratory)
- Daniilidis, Kostas (University of Pennsylvania)
- Fink, Jonathan (Army Research Laboratory)
- Fregene, Kingsley (Lockheed Martin)
- Fresconi, Frank (Army Research Laboratory)
- Frew, Eric (University of Colorado)
- Gini, Maria (University of Minnesota)
- Gupta, Satyandra (University of Maryland)
- Isler, Volkan (University of Minnesota)
- Iyer, Purush (Army Research Laboratory)
- Kearns, Kristen (US Air Force)
- Koenig, Sven (University of Southern California)
- Kroninger, Christopher (Army Research Laboratory)
- Kumar, Vijay (University of Pennsylvania Engineering)
- Luke, Sean (Department Of Computer Science, George Mason University)
- Mait, Joseph (Army Research Laboratory)
- McDowell, Kaleb (Army Research Laboratory)
- Nothwang, William (Army Research Laboratory)
- Parker, Lynne (National Science Foundation)
- Pasion, Nora (Office of the Assistant Secretary of the Army for Acquisition, Logistics, and Technology)
- Petersen, David (Defense Threat Reduction Agency [DTRA])
- Piekarski, Brett (Army Research Laboratory)
- Rao, Raghuveer Army Research Laboratory)

- Reed, Michael (ARA, Inc.; representative for DTRA)
- Russell, Thomas (Army Research Laboratory)
- Sadler, Brian (Army Research Laboratory)
- Sadowski, Robert (Army Research, Development and Engineering Command–Tank Automotive Research, Development and Engineering Center)
- Scheidt, David (Johns Hopkins University, Applied Physics Laboratory)
- Schwager, Mac (Stanford University)
- Scutari, Gesualdo (Purdue University)
- Smith, Gabe (Army Research Laboratory)
- Steinberg, Marc (Office of Naval Research)
- Yan, Xifeng (University of California, Santa Barbara)
- Young, Stuart (US Army Research Laboratory)
- Yucelen, Tansel (Missouri University of Science and Technology)
- Zhang, Pei (Carnegie Mellon University)

List of Symbols, Abbreviations, and Acronyms

3-D	3-dimensional
AI	artificial intelligence
ALC	Adelphi Laboratory Center
AM	additive manufacturing
ARL	US Army Research Laboratory
ARO	US Army Research Office
ASA(ALT)	Assistant Secretary of the Army for Acquisition, Logistics, and Technology
ASPSM	Army Science Planning and Strategy Meeting
CEMA	cyber–electromagnetic activity
D&O	deception and obfuscation
DOD	Department of Defense
GPS	global positioning system
ID	identification
IoBT	Internet of Battlefield Things
IoT	Internet of Things
IR	infrared
KSAO	knowledge, skills, abilities, and other characteristics
R&D	research and development
RF	radio frequency
SA	situational awareness
SDN	Software Defined Network
S&T	science and technology
ST	Scientific Professional
XO	executive officer

1 DEFENSE TECHNICAL
(PDF) INFORMATION CTR
DTIC OCA

2 DIR ARL
(PDF) IMAL HRA
RECORDS MGMT
RDRL DCL
TECH LIB

1 GOVT PRINTG OFC
(PDF) A MALHOTRA

1 DIR ARL
(PDF) RDRL HR
J MAIT