Detection Systems Information Fusion

William G. Bath, Christopher M. Boswell, Suzette Sommerer, and I-Jeng Wang

The transformation of the U.S. military from Cold War to current operations includes a greater reliance on distributed systems and on sensor data fusion as a force multiplier. The highly mobile fighting force will depend on the distributed operation and fusion of data to achieve net sensor performance greater than could be practically achieved by any single sensor. Detection systems information fusion (DSIF) is the fundamental enabling technology for this force multiplier. APL has a long history in designing, prototyping, and transitioning DSIF applications. In this article we summarize past APL accomplishments in DSIF and present a science and technology vision on how we should move forward to ensure APL's leadership position in this important area, building on our past successes.

BACKGROUND

Detection systems information fusion (DSIF) is a process that combines information from multiple sources to improve detection, tracking, characterization, identification, estimation, and entity/situation assessment to affect human or automated responses. Successes of many DoD and homeland protection applications critical to national security rely on the efficient and timely execution of the DSIF process to provide key information for effective decision making. APL has long been involved in DSIF activities that have led to significant advances in the performance of major DoD systems. To motivate our science and technology (S&T) vision in DSIF technologies, we discuss APL's past accomplishments with two prime examples and summarize the key enabling components that led to our past successes.

APL's Accomplishments in DSIF

Among the many APL accomplishments in DSIF are the Cooperative Engagement Capability (CEC) and the Trident programs. In both cases, APL was involved from the initial design phase, through overall system design and optimization, and finally to the deployment and testing of the prototype systems. These two programs demonstrate the significant impact APL can have when addressing the design of DSIF applications from an overall systems perspectives.

CEC

The CEC was conceived by APL in the early 1970s. Requirements development and critical experiments were performed primarily by APL as Technical Direction Agent for a Navy program to explore air defense coordination called Battle Group Anti-Air Warfare (AAW) Coordination. The first critical at-sea experiment with a system prototype occurred in 1990. CEC became an acquisition program in 1992. Additional system tests followed, including a trial deployment in 1994–1995 with a battle group of the Sixth Fleet. CEC was fielded in 1998.

CEC is an AAW system designed to maximize the utility of existing air defense systems in a battle group. By sharing high-quality measurement data from all the sensors among all the participants, each ship or aircraft gains access to a common air picture of the battle space of sufficient quality and timeliness that the data can be used by fire control systems to control weapons—just as though the information had been generated by the platform's local sensors.

Such a system has many advantages, for both the individual participants and the battle group as a whole. With all participants working from the same information, confusion is minimized, and information about which participant is engaging what can be shared. Rather than relying on tracks from a single sensor, composite tracks are formed based on measurements from a number of platforms. All sensors have limitations, and for the radars that provide the majority of the data to this system, sensor-target geometry and environmental factors have a strong influence on target detectability and measurement quality. By building up tracks from the measurements of multiple sensors on different platforms, target tracks in the air picture are much more continuous, errors in estimated target position are reduced by including measurements on targets from multiple angles, and perhaps most importantly, all participants have access to all tracks, even those on targets completely undetectable by the local sensors. For platforms with older and less capable sensors, this also gives them access to the capabilities of the most modern sensors in the battle group without upgrading older systems.

CEC improves the common picture by cueing individual sensors that could acquire targets but have not yet done so. When such a sensor applies the additional resources needed to acquire the target (concentrating only on the known location and not expending resources over a broader search area), an additional data source becomes available to all of the participants. CEC has logic for joining pieces of track for a target for which there are no measurements during a gap in data scans. This is particularly valuable in cases where the target is identified before the gap. One of CEC's most impressive accomplishments is the ability of one platform to engage a target using data supplied entirely by other platforms. This can be a major advantage when the ship with the best geometry for detecting and tracking the target is not the one best positioned to engage it.

CEC builds an identical high-quality air picture on each platform, interfaced with the combat system. The Cooperative Engagement Processor (CEP) receives unfiltered sensor measurements (range, bearing, elevation, sometimes Doppler) from the platform's sensors through the local combat system. Depending on the platform, either the local combat system or the CEP associates each measurement with a particular track using statistical measures to determine the best association. CEC shares information with other participating platforms through its own data distribution system, a high-capacity jam-resistant point-to-point communication system. The communication system provides very low latency data from every member to every member of the network. On arrival at the other CEPs, each measurement is incorporated into the track with which it is associated using a Kalman filter to estimate the target's position and velocity. By sharing all the information at the measurement level, and by providing identical data association and filtering algorithms at every platform, CEC allows each participant to build its own nearly identical copy of the same air picture. The minor differences among the databases resulting from delays in communication and limitations in bandwidth are very small compared with those among the completely independent platforms' track pictures that would occur without CEC.

Fusing data at the measurement level is key to CEC's success. By combining all the measurements of a target into a single composite track, the target becomes visible to and engageable by all participants with appropriate weapons. By maintaining track continuity through maneuvers and crowded environments that would hopelessly muddle the track picture from a single sensor, CEC provides a cleaner, much less ambiguous picture with more persistent target identification. Coordination of engagements based on data from all of the AAW sensors transforms a battle group from a collection of individual combatants into a much more powerful networked air defense system.

Trident

Beginning in 1975, as a result of APL's involvement in the Polaris and Poseidon programs and the foresight of the Navy sponsor Strategic Systems Programs, the Laboratory was able to play a significant role in the Trident Strategic Weapon System Improved Accuracy Program.[1] Methods for understanding the causes of inaccuracy during flight tests of the Trident I weapon system were developed that included advanced instrumentation and telemetry as well as system accuracy error estimation techniques. For the Trident I weapon system, these allowed system-level model validation.

To improve the accuracy of the more advanced Trident II weapon system, mere model "validation" was inadequate to effect significant precision enhancement; rather, system performance would need to be estimated with sufficient fidelity. In addition, flight tests were limited in number and constrained by expense and other logistical concerns, resulting in the commitment to maximize the information used from these tests.

The current Trident II accuracy instrumentation suite includes an acoustic-based submarine velocity and position reference system, a GPS translator on the missile, reentry body inertial measurement instrumentation, and an acoustic sea-based portable impact location system, among others. In other words, it is a constellation of sensors using various modalities or types of physical measurements to enable observation of a complex phenomenon (Trident II accuracy performance). The proper fusion of the data gathered by the instrumentation is required to ensure a high-confidence accuracy evaluation capability. To accomplish this, stochastic models of accuracy contribute at a fundamental level, independent of the test environment, and are based on physics, first principles, and engineering test experience. Complete mathematical descriptions of these stochastic and physical processes, which include, for example, gyroscopic and accelerometer hardware misalignments, scale factors, cross coupling, and functional dependence on acceleration, were developed. The structure of these models provides the fundamental parameters that can affect accuracy. These parameters are then estimated, the estimates (and uncertainties on the estimates, correlated where appropriate) being derived from flight test data from the test sensor constellation and other system tests.

Information theory provides the groundwork for estimating the parameters and uncertainties or states of the model. Based on the data collected using maximum likelihood methods, the estimation process solves the nonlinear equations for the means, variances, and Markov parameters defined by the model structure. The confidence is quantified by the uncertainty on the parameters. This model-based evaluation allows a further benefit: Since the parameters are at the fundamental level (suppose they include a bias and scale factor that do not depend on acceleration, plus terms that depend on acceleration linearly, plus terms that depend on acceleration squared), their effect on accuracy can be predicted on untested scenarios (for instance, with a different acceleration history). The models and detailed physical and engineering simulations are used to "propagate" the model parameter states and uncertainties (at the fundamental level) to the untested scenario of interest. The resulting uncertainty encapsulates both the measurement uncertainties and the parameter sample statistics through the states of the model and quantifies the ability of the model to make estimates in the new scenario.

The propagation of estimates with uncertainties in untested regions is a powerful result of this type of model-based test and evaluation (T&E) and has led to many important contributions, including the quantified understanding of system performance and confidence in system performance predictions. Anomalies in flight tests can be more easily detected and their cause more easily isolated. Improved understanding of the accuracy has also allowed optimization of controlling software, feeding the information gained from tests into system performance improvements. In addition, the development of the accuracy T&E system, along with a rigorous analytical understanding of information fusion techniques, and the development of a proper instrumentation suite with information fusion as a key part of the systems engineering design, has enabled maximal use of limited flight test assets, allowing the reduction of follow-on testing.[2,3]

Key Enabling Components to APL's Past Accomplishments

Careful studies of the above (and many other) past successes in DSIF applications at the Laboratory show that these outcomes would not have been possible without the breadth and depth of relevant expertise at APL, our commitment to an approach driven by the overall system perspective, and our deep operational understanding of the applications. The breadth and depth of our expertise enables us to address technical challenges arising in multiple facets of a complex system in an integrated framework. The system-perspective thinking focuses our effort on optimizing designs of a system based on system-level performance (instead of independent optimization of each component in isolation). A solid understanding of the operational view ensures the definition of relevant performance metrics and the dynamic optimization of trade-offs among them. As we address new and increasingly more complex challenges in DSIF, we will leverage our strengths and investigate strategic enabling research and technology areas where those strengths can best be used. At the same time, we will explore new and emerging applications for DSIF technologies, continuing to direct our research and development toward tackling longer-term critical challenges faced by the nation. Figure 1 summarizes the S&T vision in DSIF as defined by the APL DSIF Working Group (an S&T cross-enterprise initiative with representatives from multiple business areas, including Air and Missile Defense, Infocentric Operations, Strategic Systems, S&T, and Undersea Warfare).

STRATEGIC TECHNOLOGIES AND RESEARCH AREAS

Three technologies and research areas are well aligned with our strengths (Fig. 1) and also critical to the success of future DSIF applications: distributed human decisions, model-based fusion algorithms, and fusion-driven resource management. We recognize that fundamental research areas such as sensor technology, estimation and optimization techniques, signal processing algorithms, and human-machine interface technology are also key to the success of DSIF applications. APL, as well as many peer research organizations, has strong expertise in all of these areas. These three areas have been selected because of their strategic importance. We

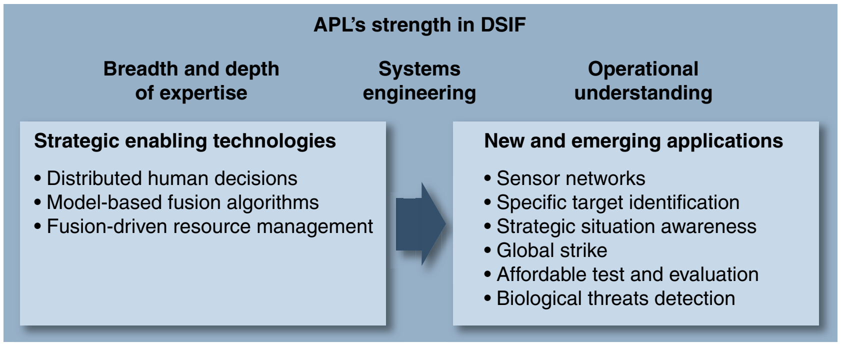

Figure 1. S&T vision on DSIF defined by the Laboratory's DSIF Working Group.

believe that the Laboratory is well and uniquely positioned to make critical contributions in these areas that will lead to significant advances in DSIF technologies and ensure APL's continuing leadership in this arena.

Distributed Human Decisions

Defining the role of the human in sensor data fusion has always been a challenge. Almost 40 years ago, during the Vietnam War, APL instrumented destroyers in the Gulf of Tonkin to determine how effectively humans could track targets on radar screens. The limited capacity of the observer to handle large numbers of simultaneous targets was documented and led to the development of automatic detection and tracking systems at the Laboratory in the 1970s. These automated systems eventually removed human involvement in sensor data processing and fusion for many Navy sensors. However, human involvement in sensor data processing remains key in many applications. Airborne sensor platforms such as the Navy's E-2C and the Air Force Airborne Warning and Control System (AWACS) have traditionally relied heavily on human interpretation of the radar screen before making a target declaration and on human interaction to keep reliable track of targets over time. Most intelligence data are subject to human interpretation before use, and most identification (ID) decisions in the military remain human decisions. The Laboratory has developed automated ID systems such as Auto-ID and CEC; however, these systems generally are not the last word in ID but instead provide an input to a final human decision. The recent Defense Science Board investigation of the Iraqi Freedom Campaign called for more human involvement in ID decisions—not less.

One common attribute in these uses of human decision making is the human's ability to gather and organize unstructured, *a priori* information about a problem and then mix that information with measured sensor data, making inferences that could not have been made using the sensor data alone. E2-C and AWACS operators work with infrequent sensor measurements. During data gaps they often can extrapolate the measured sensor data using their *a priori* knowledge of what this particular target can and is likely to do, whereas automatic algorithms are generally unable to know enough about the "situation" to do this.

Sensor data on a target (envision an SUV traveling through a city) can be spotty and eventually the target disappears, leaving no systematic way to predict its movement. The human, however, can use unstructured *a priori* information (the existence of a roadblock, the history of past target road choices) to make the best extrapolation.

Although automatic systems can certainly be developed to use *a priori* information, applying the well-known principles of Bayesian inference, this requires putting the information into a structured mathematical framework, which is difficult with a voice radio report of a roadblock or a collection of observations by other operators on past target tendencies. But the human mind excels at combining divergent types of information arriving in very unstructured formats. The key challenge arises when one wishes to fuse together inputs from different nodes at which humans are collaboratively making inferences (Fig. 2 illustrates a possible scenario). While techniques and algorithms for automatic fusion of sensor data are well understood and have been researched extensively, the necessary tools to support effective collaborative human decisions remain an open and potentially fruitful area for the Laboratory.

Model-based Fusion Algorithms

Many fusion approaches are limited to the correlation of different sources of data, resulting in a multiply co-registered, yet still complex and disparate,

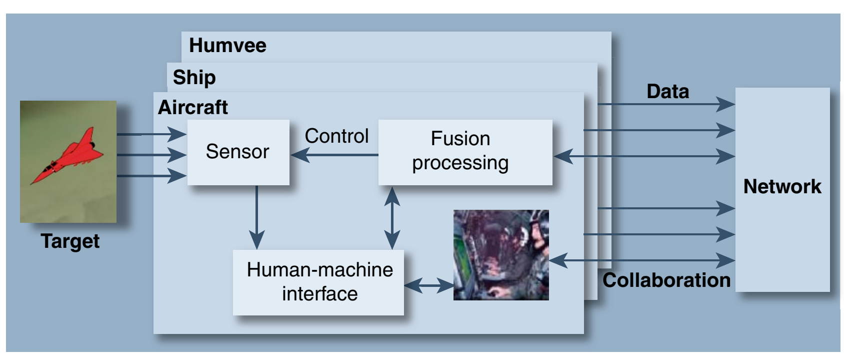

Figure 2. Distributed human decisions scenario.

collection of data points. These approaches have some use in relatively data-sparse tracking applications but may not employ the majority of the information content available to the fusion process. Some more complex approaches use a feature-level fusion based on estimates of the feature fusion methods (probabilities). These estimates may be empirically derived, or at best derived for a limited portion of the operational envelope of the detection systems whose data are intended to be fused. Any deviation from that setting (for example, a change in sensor/object relative geometries, velocities, and accelerations or new data from an untested scenario) would at least necessitate a rederivation of the feature fusion result, and feature-level fusion may not work at all for raw measurement-only methods.

However, there are methods of fusion theory, including methods of model-based fusion, that allow the simultaneous operation of changes in sensor geometry and measurement effectiveness, changes in the object, and environmental or other "nuisance" parameters. The word "model" here does not connote a cardboard or plastic object with which to view aspects of the problem, nor is it meant to be limited to a family of computer-aided design and engineering results. Rather, the term is understood to be a mathematical abstraction for the purpose of representing parameters of interest and the structure of the dynamics of the system, which includes the system model, measurement model, state parameters to be estimated in the state vector, initial conditions, physical constraints, other assumptions (e.g., measurement noise is usually statistically uncorrelated with process noise), etc.[4] Such a mathematical abstraction may be based on the physics of the target common to all sensor modes and must model states that can potentially be observed.

Information fusion processes that can be performed using likelihoods of data fusion (or data prediction, with probabilities conditioned on the previous data through the states of the model) permit a much broader range of applicability when the model states describe a sufficiently fundamental or invariant representation of the sources of the data being fused. Models that can be used range in fidelity from simply statistical (with little regard for the reasons behind the states) to intricate and accurate stochastic representations of any observable physics of the target or observed phenomena, depending on the purpose for the estimation and the observability that can be operationally expected. Most often a model is a combination of the two, with some crucial states being modeled with measurement sensitivities of high fidelity and other states estimated in the aggregate.

Model-based fusion—with evaluations of probabilities (including the multiple hypothesis methods necessary to nearly optimally combine multiple measurements and the careful Bayesian computations of likelihood ratios) as well as estimates of the parameters of model states and the uncertainty in those parameters—can yield the most accurate account of the actual amount of information available from all measurements. In addition to estimating the parameters, the uncertainty in the estimates can show where new data can make the greatest impact on the fused estimate, allowing feedback to tasking future measurements for a full use of quantitative inductive logic (see, e.g., Ref. 5).

Fusion-Driven Resource Management

Efficient management of resources (i.e., sensing, computing, communications) is important to the success of DSIF applications. Traditionally, resource management in a data fusion process is treated as a separate component that is part of the so-called level-4 process refinement as defined in the JDL data fusion process model[6] (Fig. 3). Consequently, the designs of resource management techniques are focused on ensuring the timely delivery and processing of data/information to support the data fusion processes for source processing and level 1–3 fusions. As DoD systems become more net-centric and rely on shared resources, this simplistic view of resource management as part of the DSIF process is no longer adequate. One challenge is the increasing need for distributed information fusion over limited and dynamic network connectivity for operations in urban and complex environments. Management of communication resources for distributed fusion requires effective dynamic topology management and information sharing strategies that are tightly coupled with the design of fusion algorithms.[7] Furthermore, for many emerging

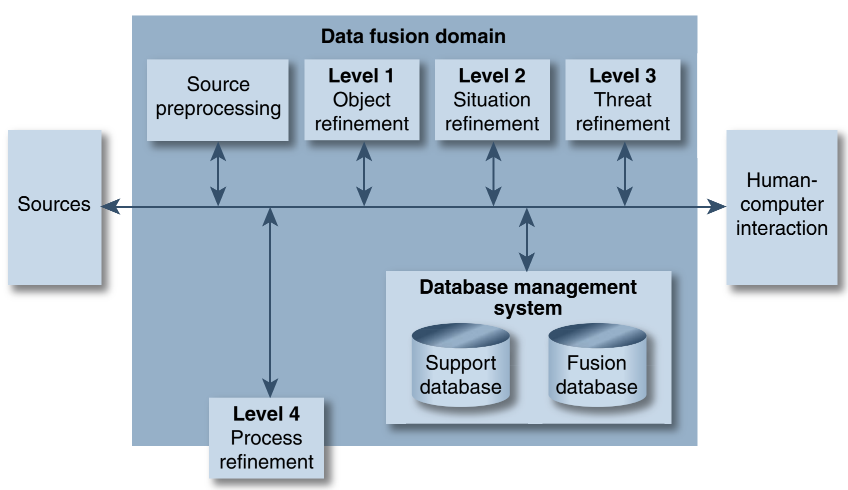

Figure 3. JDL data fusion process model.

DSIF applications involving large networks of low-power devices, resources need to be tightly managed to ensure the necessary operating system lifetime while maintaining strong detection and classification performance. Figure 4 illustrates the strong couplings and interdependencies among the three functional components in a network-centric DSIF application: sensing control, data processing, and communications.

Existing techniques for handling limited resources for distributed fusion applications are focused on local optimization of individual resources. For example, power constraints may be dealt with by sleep cycles that are planned locally, with a goal of minimizing power usage within a subsection of the system, without reference to the needs of the detection system's information state. However, a small amount of (locally) nonoptimal power usage resulting in an increased information gain (and thus performance) for the entire system is globally preferable. We believe that an effective approach is to focus on global resource management driven by system-level performance defined by the information fusion application.

Working with the Undersea Warfare and Infocentric Operations business areas, the S&T business area is investigating fusion-driven resource management techniques to improve the detection performance and energy efficiency of autonomous sonobuoy networks for the ASW application under independent R&D funding (preliminary concept design and analysis are reported in Ref. 8). The emphasis of this effort is on providing an example of optimization of global resource utilization for a system in which such optimization is not currently being done. The successful result would be a quantification for increased performance (measured by the detection goal of the entire system or network) using the fusion process as the standard for optimization, even when a subsystem or aspect of the system in question appears to be nonoptimal.

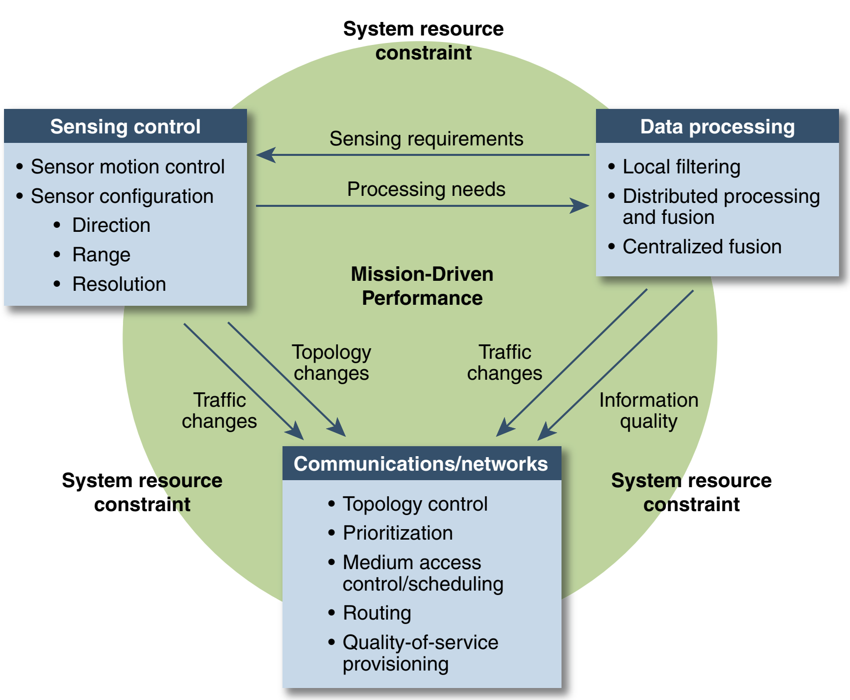

Figure 4. Coupling of three high-level functions in net-centric DSIF.

NEW APPLICATION AREAS

In this section, we briefly describe a number of important new application areas for the DSIF technologies. These applications define novel and significant technical challenges being considered as we formulate our S&T vision in DSIF. One new application critical to national security is the distributed surveillance and detection of biological threats. APL has played a key role in the development and deployment of large-scale syndromic surveillance infrastructures. The S&T vision for this important area is discussed in this issue in the article by Demirev et al. and references therein.

New Sensor Networks

Sensor networks are poised to revolutionize the way we interact with the environment. Fueled by Moore's law, we can now manufacture inexpensive computing nodes equipped with sensors that can be deployed very close to the physical phenomena we want to observe and possibly alter. These nodes are battery operated and can communicate with each other using low-power radios, enabling easy deployment, low-maintenance, and long-duration observation of physical phenomena. In the civilian sector, sensor networks have been deployed for environmental monitoring, remote monitoring of the behavior of animal species, and structural monitoring of buildings and equipment. Sensor networks have already seen active duty deployments in the military. They have been used to detect enemy movements and locate snipers in urban environments.

Although these early applications indicate the potential offered by sensor networks, they only scratch the surface in terms of complexity of the applications and the size of the network deployments. Specifically, most sensor networks today are used to transfer measurements taken by individual nodes to a centralized location for post-processing. In this way, the sensor network acts as a wireless telemetry network passing measurements from the field to the lab for analysis. However, the full benefit of sensor networks is realized by in-network processing, where sensor nodes collaborate to process collected information inside the network and pass the results to the remote site, thereby reducing the amount of data transmitted, and in turn reducing power consumption and increasing the network's lifetime.

DARPA's Sensor Information Technology[9] (SensIT) and Networked Embedded Systems Technology (NEST) programs have made great strides in developing collaborative signal processing algorithms, power-efficient communication protocols, and middleware and programming tools for wireless sensor network applications. Nevertheless, successful deployment of any large-scale wireless sensor network application remains a significant technical undertaking. The key challenge lies in the complexity of developing such applications with the software tools available today. Even with nesC,[10] the state-of-the-art programming language for sensor network applications, application programmers are forced to write low-level system code that resembles the assembly code used in the 1950s and 1960s for writing mainframe applications because sensor nodes are resource constrained and therefore the application must control low-level details to optimize resource use. However, many of these resource constraints are temporary. We expect that the next generation of sensor nodes will have an order of magnitude higher processing speed and memory capacity. Nevertheless, domain experts in relevant fields are not expected anytime soon to learn to write low-level code for their sensor networks.

The Laboratory has engaged in several important DoD initiatives that involve deployment of large wireless sensor networks, including the Giant Shadow Exercise, Operation Silence Hammer, and the sonar buoy network for the ASW application. Our strengths in addressing end-to-end system optimization driven by operational understanding will provide key insights into the development of effective tools and enabling algorithmic techniques for the deployment of future sensor network applications.

Specific Target Identification and Modeling for Ground Targets

After the measurements from a set of sensors indicate a number of objects, there may follow a gap in the collection of data scans. Once new data arrive after the gap, the determination of the likelihood that the objects measured before and after the gap are identical is called the process of specific target identification (STI). Consider a scenario where the objective is the surveillance of a vehicle traveling close to similar vehicles as they enter a tunnel. Later, a vehicle emerges: Is it the one of interest? Another example is a set of unattended aerial vehicles watching some military threats across some border when priorities change and the sensors must be retasked. After a while, the same border is under surveillance again, and similar military threats are observed. Are these the same threats? What is the total number of such threats? Questions such as these are addressed by STI. Target classification may help the STI process but is not required. It may be sufficient in some applications, even when the classification is unknown, to know that, whatever its classification, the target is the same before and after.

Most real-world applications for STI will be in the context of multiple sensor data fusion. Any sensors that can detect the target or signature of interest will likely also have other tasks, so they cannot simply "stare and keep staring" at the object of interest. In addition, if the importance of the object is high enough, disparate types of sensors may be configured to provide scans of the data (occasional "looks," "listens," etc.) that would include the object of interest. The need to use data from multiple scans and multiple sensor types, with potentially varying ability to observe the appropriate signatures, would argue for a type of model-based fusion. If the STI likelihoods could be computed using multiple hypothesis methods and model-based fusion with appropriate fidelity, the results would contain the maximal data available for STI.

Strategic Situation Awareness

A great deal of the actual data that eventually inform decisions by the National Command Authority are likely in some kind of electronic form before any human actually evaluates the data. In addition, developing a response to any situation of strategic importance involves the work of many people, most of whom are probably not located where different specific courses of action are being developed or decisions are being made. With the advent of global networks, it is possible to instantly share data and collaborate on evaluation and possibilities. How much of this has gone beyond "chatroom"-type discussions and Internet "phone lines"? Using the various DSIF technology areas, it would be possible for computer networks to fuse and correlate data into a clearer picture of the actual situation, allowing the humans to spend more of their resources doing what they do best—providing creative alternatives, being ready for any *coup d'oeil* resulting in a successful plan, and making decisions based on complex and sometimes poorly quantified human factors. This information fusion bridge between and among intelligence/surveillance/reconnaissance (ISR), strategic command and control, and all levels of national response may prove to be a crucial tool in future strategic calculations.

Global Strike

The Strategic Command Global Strike mission requires sufficient information to assess a potential threat anywhere in the world, define a course of action, and implement the action in an extremely expedited time frame. The resultant needs for fast battle-space awareness, fast decision making, and fast damage assessment can all benefit from information fusion technologies and expertise. Multiple sensor "strategic targeting" employing multiple hypothesis methods and optimized ISR allocation tools, all using model-based fusion and validated sensor performance models, can allow the best application of available assets to find, locate, and track the types of threats involved in the shortest time possible. The understanding of this fused data, correlated with other needed information in a collaborative environment, will be aided by the distributed human decisions technologies of the type discussed previously. No operation is complete until an assessment of success shows that no further action is necessary, and this can be demonstrated in the most timely manner by the optimizing and quick methods similar to those used for the initial ISR picture.

Affordable Test and Evaluation

Constraints on testing budgets and range safety and other logistical issues make operational T&E difficult at times. One of the main purposes of operational T&E is to provide the quantified assurance that the system under test will perform as required in an operational setting. In model-based T&E, common fundamental models of those factors that can affect system performance are developed that are independent of exact mode of operation, location in the design envelope, or environment across multiple potential missions. Sufficient instrumentation and test systems engineering allow a data fusion approach for estimating the scenario-independent common performance models to accumulate information on model parameters from data across both traditional "operational tests" and to properly constructed laboratory and simulated environment tests, analogous to the Trident II T&E outlined earlier.

With high-fidelity fundamental models, regions of performance in many environments (acceleration, vibration, heat, etc.) are estimated with the terms that show the effect of the environment but are themselves independent of them. Enhanced ground tests, plus properly instrumented flight tests, will show the predicted performance of the system to untested scenarios and also provide the quantification of how well that performance is known, pointing to areas of need for future testing. This results in the best use of all test data and an objectively quantified confidence in performance predictions. Experience has shown that this combination can dramatically reduce the assets required to perform operational tests to attain a given confidence and provide a robust, analytical means to deal with the fact that not all potential operational scenarios are possible in a test environment.

CONCLUSION

In this article we have reviewed past APL accomplishments in DSIF and presented an S&T vision on how we should move forward to ensure APL's leadership position in this important area, building on our past successes. Motivated by the new and emerging DSIF applications for DoD and homeland protection applications, we see the three key research areas covered in this article that are well aligned with the Laboratory's strengths.

REFERENCES

[1]Coleman, D. R., and Simkins, L. S., "The Fleet Ballistic Missile Accuracy Evaluation Program," *Johns Hopkins APL Tech. Dig.* **19**(4), 393–397 (1998).

[2]Levy, L. J., "Systems Analysis and Test and Evaluation at APL," *Johns Hopkins APL Tech. Dig.* **24**(1), 8–18 (2003).

[3]Levy, L. J., and Porter, D. W., "Large-Scale System Performance Prediction with Confidence from Limited Field Testing Using Parameter Identification," *Johns Hopkins APL Tech. Dig.* **13**(2), 300–308 (1992).

[4]Gelb, A. (ed.), *Applied Optimal Estimation*, M.I.T. Press (2001).

[5]Hall, D. L., and Llinas, J., *Handbook of Multisensor Data Fusion*, CRC Press, Boca Raton, FL (2001).

[6]White, F. E., "A Model for Data Fusion," in *Proc. 1st Nat. Symp. on Sensor Fusion* (1988).

[7]Llinas, J., Bowman, C., Rogova, G., Steinberg, A., Waltz, E., and White, F., "Revisiting the JDL Data Fusion Model II," in *Proc. 7th Int. Conf. on Information Fusion* (Jun 2004).

[8]Doshi, B., Benmohamed, L., Chimento, P., and Wang, I-J., "Sensor Fusion for Coastal Waters Surveillance," in *Proc. Multisensor Multisource Information Fusion Conf., SPIE Defense and Security Symp.*, Sarasota, FL (Mar 2005).

[9]Gharavi, H., and Kumar, S. P. (guest eds.), Special Issue, *Sensor Networks and Applications*, *Proc. IEEE* **91**(8), 1151–1153 (Aug 2003).

[10]Gay, D., Levis, P., von Behren, R., Welsh, M., Brewer, E., and Culler, D., "The nesC Language: A Holistic Approach to Networked Embedded Systems," in *Proc. Programming Language Design and Implementation (PLDI) 2003*, pp. 1–11 (Jun 2003).

THE AUTHORS

William G. (Jerry) Bath received B.E.S., M.S.E., and Ph.D. degrees from JHU in 1974, 1975, and 1980, respectively. Since joining the Laboratory in 1974, Dr. Bath has worked on the signal processing, automation, and integration of radar systems for Navy ships (AN/SYS-1, AN/SYS-2, MK-92 CORT, and SPS-48C DDC) and the early networking of Navy radars using the Surface Gridlock System. For 14 years he led the Engineering Analysis Group that developed the concepts, requirements, and algorithms for sensor netting in the CEC. Since 1997, he has led the Sensor and Weapon Control Integration Branch in the Air and Missile Defense Department (AMDD). Dr. Bath has published over 30 papers on radar signal processing, tracking, and networking. He is a senior member of the IEEE and the first recipient of the IEEE Radar Systems Panel Award for outstanding contributions to radar. **Christopher M. Boswell**, a member of APL's Principal Professional Staff, completed a B.S. from Seattle Pacific University in 1986 and M.A. and Ph.D. degrees from The Johns Hopkins University in 1988 and 1994, respectively, the latter based on work in experimental particle physics. Since joining APL in 1997, Dr. Boswell has been involved in ballistic missile analysis, terrain-aided navigation evaluation, sensor development, statistical modeling, and the application of optimization techniques. His current research involves ISR sensor fusion, including the development of models and algorithms for multisensor data fusion to provide precise target alerts. He is a member of the American Physical Society and the IEEE. **Suzette Sommerer** is an engineer and Assistant Department Head in the AMDD. She received a B.S. in chemical engineering from Washington University in St. Louis in 1979, and M.S. degrees in numerical science and in applied physics from The Johns Hopkins University in 1985 and 1988, respectively. Before joining APL, she was involved in chemical process research and development for Monsanto Corporation and environmental engineering for Atlantic Research Corporation. Since joining APL in 1982, Ms. Sommerer has analyzed satellite observations of submarine-launched ballistic missiles, developed algorithms to correlate and track sensor data over a large area, studied real and proposed systems to characterize wide-area sensors, and developed a multiplatform simulation to evaluate coordinated tactical ballistic missile defense concepts. She led the multi-organization network-control algorithm design team for the CEC's point-to-point communication system from its beginning in 1985 through several major fleet demonstrations and CEC's deployment in the fleet. She also led the Strategic Systems (now Global Engagement) Department's Guidance and Navigation Control Group for 3 years, starting in 2001. In 2004, Ms. Sommerer became Assistant Branch Supervisor for the Strategic Systems Branch, and then returned to the AMDD in her current role in 2005. **I-Jeng Wang** received an M.S. degree from Penn State University in 1991 and a Ph.D. degree from Purdue University in 1996, both in electrical engineering. He joined APL in 1997 and is a member of the System and Information Sciences Group of the Research and Technology Development Center. Dr. Wang's research interests include stochastic optimization, reasoning under uncertainty, and large-scale network management and control. He is a member of the IEEE, ACM, and INFORMS. Further information on the topics discussed in this article can be obtained from Dr. Bath. His e-mail address is jerry.bath@jhuapl.edu.

William G. Bath

Christopher M. Boswell

Suzette Sommerer

I-Jeng Wang