Operational and Technical Implications of the Robotics Revolution

Theodore Bially
Director, Information Exploitation Office
Defense Advanced Research Projects Agency, 3701 N. Fairfax Drive, Arlington, VA 22203

ABSTRACT

The premise of this talk is that more sophisticated robotics hardware and smarter software alone won't suffice for achieving the full potential afforded by robotic assets. We will also have to develop entirely different mechanisms for planning and executing unmanned missions and new bandwidth-efficient approaches to sharing imagery and other wideband sensor data among large numbers of collaborating platforms. Our current incredibly detailed and complex mission planning methods will bog down if applied to operations involving large numbers of platforms – yet we will be deploying many more robotic assets in the future than we do today. In addition, inter-platform and reachback RF communications resources will be stretched thin if the sensor outputs of all these platforms need to processed off-board. A futuristic ISR mission involving multiple unmanned platforms and a broad range of wideband sensors is presented as an example to demonstrate the transformational potential of the large-scale deployment of robotic assets, and also the need for revolutionary new approaches to mission planning, command and control, sensor processing and exploitation.

Keywords: robotics; unmanned systems; global information grid; remote operations; intelligence, surveillance and reconnaissance (ISR) systems; automatic target classification and/or recognition (ATC/ATR).

INTRODUCTION

Unmanned systems have been proliferating rapidly and will continue to do so at an accelerating rate as we find new ways of using them in place of manned assets in critical – and often dangerous – military operations. Along with our own improved understanding of how to exploit, manage and control robotic platforms, the unmanned systems themselves are becoming "smarter" in their own right as we endow them with more sophisticated software capabilities. And the procurement and operating costs of these systems will begin to diminish as greater numbers of robotic assets are produced and as we learn how to incorporate them more efficiently into our operations. All this adds up to the fact that the expected military robotics revolution is well underway – but it's also got a long way to go.

The premise of this talk is that more sophisticated hardware and smarter software alone won't be enough to propel the robotics revolution to its ultimate transformational state, but that in order for the revolution to achieve its full potential we will have to develop entirely different mechanisms for planning and executing unmanned missions. This need is precipitated in part by the observation that our current incredibly detailed and complex mission planning methods will bog down if applied to operations involving large numbers of assets – yet we will be deploying many more robotic assets in the future than we do today. The motivation for changing how we do business is also driven by the necessity of communicating in a shared resource environment. Our wideband communications infrastructure will be servicing a large population of simultaneous users, many of whom will have compelling need for bandwidth and fast response under warfighting conditions. So we anticipate that bottlenecks will develop as the number of deployed assets rises. Thus reachback communications from unmanned systems and time-sensitive command and control messages may sometimes experience serious latency problems at inopportune moments.

Our focus here is on the inter-relationship between two transformational developments - the continuing evolution of robotic systems and the emerging wideband global information grid - and between these and a third continually evolving discipline, namely intelligence, surveillance and reconnaissance (ISR) systems. Ongoing ISR developments include new sensors and sensor modalities, and improved techniques for automatic target classification and recognition. Robotics, communications and ISR each present major technical challenges in their own right and vigorous research is underway separately in each discipline. But considerably less attention appears to have been paid toward understanding the systems-level issues and tradeoffs inherent in exploiting these capabilities in combination. We speculate in this presentation on how future unmanned ISR and strike missions may be structured using robotics, communications networks, sensors and automatic exploitation capabilities in combination, and we try to anticipate some of the inter-

Defense Transformation and Network-Centric Systems, edited by Raja Suresh,
Proceedings of SPIE Vol. 5820 (SPIE, Bellingham, WA, 2005)
0277-786X/05/$15 · doi: 10.1117/12.608435

disciplinary compatibility issues that are likely to arise. Emphasis is on major paradigm shifts that may overcome some of the technical constraints and limitations imposed by each of the disciplines upon the others.

THE REVOLUTION HAS BARELY BEGUN

While substantial progress has been made in designing and building a wide variety of unmanned platforms and we've learned to use these to great advantage in warfighting situations, the anticipated robotics revolution has not yet fully come to pass. The word "revolution" implies transformational change and entirely new ways of doing business. But from an operational perspective we're managing much of the present generation of unmanned systems in essentially the same ways we manage manned platforms. In many of today's operational UAV and UGV systems the vehicles are driven and controlled by humans, albeit remotely. Human operators are responsible for mission success. They exercise real-time control of the vehicles and their payloads based on evolving conditions as recorded by on-board sensors and transmitted to the remotely located operators. In addition, the number of robotic vehicles that are deployed in typical present day operations tends to be limited because unmanned assets aren't yet all that plentiful, failure rates are relatively high, and numerous operators are required for remote operation and mission support.

The fact that humans are in the loop in so many of our contemporary unmanned operations indicates that the anticipated robotics revolution isn't yet fully underway. The same conclusion follows from the observation that today's unmanned platforms are usually deployed singly - or at best in moderately small groups. While tele-operation is a highly effective and entirely reasonable way of controlling small numbers of unmanned assets, the approach doesn't scale well with increasing vehicle count. But simultaneous deployment of many low-cost unmanned platforms will be critical to mission success across a wide range of future military applications. In fact one of the important attributes of the upcoming robotics revolution is that because it will provide wider diversity and greater numbers of platforms and payloads than we've previously enjoyed, this will lead to entirely new problem solutions and operational concepts. However these will need to be compatible with human resource limitations and with the evolving global information infrastructure.

Progress is being made on the manpower issue, with the goal of changing the remote operations approach from many-on-one to one-on-many. That is to say, we're working toward allowing a small number of remote operators (ideally, one) to manage and control several robotic platforms simultaneously – while accomplishing complex and risky tasks under hostile conditions. This will be a lot easier to accomplish when the participating platforms travel together in close formation, than if their mission plan required them to travel somewhat different routes. A small number of remote operators may be able to maintain good situational awareness for a tight grouping of vehicles - for instance by thinking of the group as a single large entity. But when the platforms disperse or separate as part of their planned operation, it will be necessary to maintain detailed, up-to-the-second situational awareness for each one independently of the others. This may require high-resolution video coverage of the space surrounding each platform and possibly radar and IR as well, and all of this will need to be presented to mission controllers in real time. A one-on-many control methodology is unlikely to work well in this situation. The number of people needed to operate and coordinate a large population of unmanned vehicles will become impractically large as the number of deployed vehicles grows and as the complexity of their missions increases.

Even if we solve the people resources problem, considerable RF bandwidth will be needed to communicate needed sensor data plus telemetry information to remote control points from every platform in a large, dispersed group. And this raises the issue of network resources. The limited availability of RF bandwidth, especially under battlefield conditions, will impose serious constraints on the quantity of sensor data that can be communicated from multiple unmanned platforms to remote control facilities.

COMMUNICATION CONSTRAINTS

Success in the modern battlefield hinges greatly on having reliable communications connectivity among large numbers and all categories of warfighting assets, both manned and unmanned, plus reach-back to in-theatre and CONUS command centers. The evolving global communications network is expected to provide this capability. Anticipated attributes of the network are that it will be robust in the face of enemy actions such as jamming and physical attack while providing bandwidth sufficient to accommodate all mission-critical transactions in a reliable and timely manner. Considerable resources are being expended to develop and implement this infrastructure, with many segments of the network already deployed and more to come.

As with any new capability, the advent of the global network presents new technical challenges – *e.g.,* prioritizing and managing its resources, maintaining necessary quality of service and providing multi-level security. Network resource management, especially under stressful and demanding warfighting conditions, will occasionally force some network users to settle for reduced quality of service - often leading to degraded performance of user systems. Although available network bandwidth and overall quality of service will be considerably greater than what we've had before, anticipated usage is also much higher. In other words we're building a high capacity infrastructure but it's very likely that we'll eventually wind up using it to the fullest extent possible. So we will inevitably continue to operate in a trade space in which the performance of systems using the network will have to be balanced against the need to provide reasonably acceptable network services to all who need them.

A logical conclusion from the foregoing is that as the robotics revolution advances, the concept of remote operation will have to be replaced by a less resource-consumptive alternative. Otherwise the revolution will stall as a result of people limitations and communication constraints. In other words the robotics revolution will accelerate in earnest when unmanned air and ground vehicles gain the ability to execute tasks autonomously and in self-organized teams, with minimal human involvement and commensurately less dependence on reachback communications. Good progress is being made in this direction but there's still a long way to go. For example we have technology that allows UAVs to execute certain complex maneuvers with minimal operator intervention. However, situational assessment, especially in hostile territory and in the face of unanticipated events still requires operator participation and this in turn implies reliable, real-time telemetry of wideband data back to a remote facility.

Unfortunately our experience to date has been that as more capability for autonomous operation is introduced into our unmanned platforms, they're able to handle more complex tasks, but the number of remote operators needed to manage and control them hasn't gone down appreciably. One reason for this may be that, in accordance with tradition, we're carefully planning every platform's route and maneuver in advance, and we're trying to maintain compliance with that plan throughout the mission. And this requires continual man-in-the-loop monitoring and oversight. So it may be that we'll need an entirely different mission planning and monitoring paradigm if we want to conduct more complex, multi-platform missions with fewer humans in the loop.

We imagine a future in which most of the human participation in robotic military operations focuses almost exclusively on mission definition and planning, and nearly all of the mission execution and real-time decision making is handled by the robots themselves. Considerable new technology development will be needed for this to happen, especially in the areas of cognitive software, signal and image understanding and distributed autonomous control. Absent these capabilities it's hard to imagine deploying robotic systems of any significant size without requiring inordinately large numbers of support personnel and/or seriously straining limited RF bandwidth resources.

In the remainder of this talk we explore some of the technical issues and operational implications of the impending robotics revolution, with emphasis on future unmanned ISR and strike missions. For reasons described below, missions of this type will require large numbers of geographically dispersed platforms, each equipped with one or more high data-rate imaging sensors - exactly the combination that will be highly impractical for remote, man-in-the-loop analysis and control.

EVOLVING ISR NEEDS

Present day ISR practice generally emphasizes individual platforms carrying a limited number of complimentary sensors *e.g.,* a Global Hawk or a Predator UAV outfitted with radar and electro-optics. These platforms typically transmit raw or partially processed sensor data at high continuous bit rates to remote exploitation facilities via dedicated point-to-point communications links. In this regard current ISR practice is highly analogous to that of robot tele-operation, *i.e.,* there's a strong dependence on a remote facility and a critical need for reliable, high-speed communications between the sensor platform and the exploitation center.

While the single platform ISR approach is effective for many situations, it doesn't provide true persistent surveillance capability - an acknowledged critical, but presently unsatisfied, military requirement. Terrain variations, vegetation and cultural objects all can limit visibility at various times – and enemy forces are able to exploit these obscuration effects to their advantage. A multiplicity of sensors carried on a variety of manned and unmanned platforms will need to work collaboratively in order to achieve the needed persistence. But if all the participating ISR assets have to transmit their wideband sensor data back to a remote exploitation center, this will impose the same sort of loading on the RF communications infrastructure as would simultaneous tele-operation of the various platforms – and probably more. Note that while imagery and video bandwidth compression methods generally work well when the image products are viewed by human analysts, it's frequently the case that the performance of automatic detection and classification algorithms degrades significantly with high compression ratios. So image compression techniques may not offer sufficient relief.

By analogy with the concept of distributed, autonomous on-board control of robotic platforms, and for the same reason, future ISR systems will have to employ automatic on-board exploitation processing to a much greater extent than is currently practiced or even feasible. The objective of the on-board processing of course would be to perform automatic cueing and prioritization of possible targets or areas of interest on the sensor platforms themselves. Network bandwidth would be conserved by communicating only the exploitation results or, on occasion, mission-relevant subsets of the sensor data to a command facility. Automatic on-board exploitation is by no means a new idea; it's been an elusive goal of the ISR community for decades. But thus far we have been unsuccessful in achieving real-time automatic target classification and/or recognition (ATC/ATR) with acceptable detection performance concurrent with sufficiently low false alarm rates, in all but the most benign situations. Current techniques are far from adequate under conditions of complex terrain, evasive or deceptive deployments and in urban environments. So we conclude that that, among other things, a technical breakthrough in ATC/ATR may well be a necessary precursor to the large scale use of unmanned platforms for military ISR applications.

In view of the fact that reliable ATC/ATR has eluded adequate solution for many years it would be unrealistic to believe that a breakthrough is imminent unless a previously unexploited dimension was to manifest in the solution space. Interestingly, this indeed may be the case – where the newly available dimension is the spatial diversity made possible by the growing population of unmanned platforms – that is to say, by the advancing robotics revolution.

Until recently most efforts to develop ATC/ATR capability have been conducted in the context of single platforms, and often with a single sensor. Fusion of data from multiple sensors for purposes of classification or recognition has generally focused on combining results from different sensors located on a common platform. The idea of using several sensor-carrying platforms to view the same object simultaneously or near-simultaneously from different spatial vantage points hasn't been a practical option because most of our ISR assets have been of the so-called "low-density, high-demand" variety. The added spatial dimension afforded by more numerous platforms, in conjunction with networked communications connectivity, will bring a new twist to the ATC/ATR problem and perhaps a long-awaited breakthrough. And this in turn may help to relieve the communications bottlenecks that will otherwise occur as a result of the proliferation of assets. But we need a methodology by which data from multiple sensors of various types located on different platforms at different spatial positions can be exploited jointly to yield more accurate, reliable and robust ATC/ATD performance than is possible today.

Assuming we can achieve reasonable ATC/ATR performance from some form of diversity-based exploitation, and if the technique were to be implemented in a distributed manner on-board the participating platforms, it could relieve the pressure on network resources caused by the need to transmit large volumes of sensor data to a common exploitation point. However, that advantage will be thoroughly wiped out if, in order to achieve adequate performance, the various platforms had to access each other's signal- or pixel-level data. So a challenge we face is to develop ATC/ATR mechanisms that exploit the new resource-rich ISR environment in ways that minimize the volume of data transfer between participating platforms and between the platforms and remote analysis facilities.

While the most satisfying approach from a sensor utilization perspective would be to perform multi-sensor signal- or pixel-level fusion by combining the raw data from all available sensors prior to ATC/ATR analysis, this clearly imposes a very high communications penalty and appears to be fundamentally incompatible with concept of a distributed, resource-rich ISR architecture. So it behooves us to look for alternative approaches that are inherently conservative of communications resources.

One way to proceed would be by deemphasizing optimality of sensor utilization and focusing instead on getting the right answer, albeit inefficiently. In other words when there's an abundance of assets we can think about using those assets inefficiently – for example by performing imperfect ATC/ATR processing independently on each platform and subsequently combining individual results at the symbolic or soft-decision levels. This would be highly efficient with respect to network bandwidth utilization but it clearly doesn't take advantage of all the information that may be contained in the source sensor data. It's analogous to a sub-optimal – but often adequate – target tracking strategy that combines state vectors from multiple independent trackers instead of updating a single integrated track filter with detection metrics from all contributing sources. Other approaches to communications-efficient multi-asset exploitation are clearly possible and should be explored.

To summarize this ISR discussion, we observe that as robotic ISR platforms proliferate it will become increasingly difficult to conduct ISR missions using the traditional remote exploitation approach because of the heavy loading this will place on the RF bandwidth resources. In order to depart from this tradition we will need to have reliable and accurate automatic target classification and recognition capabilities on-board the platforms themselves. This has been an elusive goal that has defied solution for many years but for which a new approach may now be possible, offering good possibility of success. Interestingly, the enabler of this new approach may be the very proliferation of robotic platforms that gives rise to the need for it in the first place. Hmmm…..

MISSION PLANNING AND CONTROL

As noted earlier, current mission planning methods are incredibly detailed and complex, and extending the present approach to a substantially greater population of assets would seem to be impractical at best. To compound the issue, as a mission evolves it is invariably necessary to re-plan the operation based on the developing situation. For unmanned assets this means that some or all of the sensor data will have to be communicated to a manned platform or remote command center in real time. As in the exploitation case, current practice implies that large volumes of sensor data will have to be transmitted at real-time rates from multiple platforms to a central command facility. And this will additionally seriously stress network bandwidth resources.

Radically different mission planning and command and control mechanisms will need to be developed for managing the multiplicity of unmanned platforms and sensors that are likely to be involved in future persistent ISR operations. One possibility is outlined below, again motivated by the need to conduct unmanned operations in ways that conserve limited network resources.

With lots of assets available it may no longer be necessary to perform the kind of exquisitely detailed mission planning and dynamic re-planning that is typical with today’s low-density, high demand platforms. Applying the same logic we used for the ATC/ATR problem, suppose we were to deploy our assets in a less efficient manner than we’ve become accustomed to, and depend instead on there being a sufficient number of them in an area to satisfy a reasonable range of unforeseen needs as they arise. This can be thought of as a statistical form of planning, in which detailed, deterministic plans are replaced by selecting the right statistical mix and density of assets to deploy over and within given ground areas in order to maximize the likelihood of having the right assets available when and where they may be needed. The selected assets can be launched with general guidelines such as “wander at will within an area defined by the following geospatial coordinates, don’t collide with anything moving or stationary, respond as appropriate to sensor coverage and/or data requests from other system components, and listen for instructions from the command center.”

This isn’t easy to implement even with manned platforms, and it will be substantially more difficult to accomplish for unmanned vehicles. But assuming we succeed at some sort of statistical mission planning, our next challenge will be to cause multiple platforms to self-organize into autonomous ad-hoc teams in response to high level command inputs that define tasks and mission objectives. Individual platforms and cooperating groups of platforms will need capability to develop situational awareness with considerably less dependence on remote command and control centers than is current practice. They will have to maintain cognitive awareness of their objectives and their environment, and they will have to adapt on their own to unplanned and unforeseen circumstances. Some of this might be accomplished in the near term by forming mixed teams of manned and unmanned platforms, but unless many of the coordination and distributed situational awareness tasks are automated, additional personnel may be required in the manned platforms.

A long term goal of course would be to accomplish most, if not all, future ISR – and strike – missions using only unmanned assets. While the technological foundations to enable this have been active research topics for many years, much more progress will have to be made before practical systems of this type will be possible. Here’s a hypothetical view of how the future might look for a search and destroy type of mission in which the objective is to seek out enemy targets over a specified patch of terrain and destroy those that are assessed by the mission commander to be of high priority:

ISR/STRIKE SCENARIO

We assume this mission is conducted entirely by a collection of autonomous unmanned assets that form into collaborative teams in response to mission-level tasking directives from a remote command center. The composition of the collection (*i.e.,* types and quantities of unmanned platforms, sensors and weapons) is selected by the commander based on mission objectives, terrain characteristics, the enemy order of battle and anticipated enemy countermeasures. Platforms are initially deployed with instructions to distribute themselves more-or-less uniformly over the area of interest. Because of the large number of deployed platforms it's reasonable to expect that a suitable set of assets will be at or near the places they will be needed when they are needed. Once the mission begins the robotic assets function mostly autonomously, with minimal human intervention.

From the commander's perspective the entity he communicates with is the overall system of assets – not specific platforms, sensors or weapons. High-level commands are issued to the system as a whole. In response to these commands individual platforms communicate among themselves at relatively low data rates to form ad-hoc teams that will collaborate together to execute a task.

For example suppose the task were to perform wide area surveillance over a given patch of terrain and report back the locations and types of candidate targets. Assets required to accomplish this would include one or more high altitude, wide area surveillance platforms carrying synthetic aperture radar imaging sensors and possibly one or more optical instruments. In addition we postulate a fairly large number of low altitude UAVs and/or ground-based robots equipped with high resolution imaging sensors (e.g., Video, IR, MS/HS, LIDAR, *etc.*) and on-board ATC/ATR software. Each ATC/ATR is assumed to have capability to classify targets based on local sensor data alone, and also from a combination of local data and summary information or soft decisions received from other assets observing the same object.

When a possible target is detected by one of the wide area surveillance assets that platform transmits the geospatial coordinates of the detection to all platforms in the system that are within line-of-sight communications range. Those platforms that are able to view the target with sufficient resolution to perform ATC/ATR processing form themselves into an ad-hoc team for the purpose of classifying or recognizing the detected object. They independently task their on-board sensors to view the object and they each process their collected data to produce single-platform estimates of the target type. These estimates are then communicated to all other team members. The team converges on its best estimate based on inputs from all the observing platforms and sensors and a short message is transmitted to the remote command center indicating target location and estimated type, and a listing of the sensor data that could be sent back to the commander for eyes-on verification. Supporting data as requested is then transmitted to the command center directly from the collecting platforms. Once a detected target is identified and registered with the command center the various ISR platforms that constituted the ad-hoc ATC/ATR team for that target can return to the general pool where they will be available to participate in other ad-hoc teams in response to new tasking or new target detections from the wide area surveillance platforms.

Given large numbers of robotic assets it will be possible to distribute a variety of weapons throughout the battlespace in addition to a rich mix of ISR assets, and this should result in considerably reduced strike latency. Once a commander determines that a weapon should be launched against a candidate target he can issue an appropriate command to the system and, as in the ISR case, the system ought to organize itself to execute the process. This would include selecting appropriate weapons for the target at hand and identifying a specific platform that is equipped with one of these and is advantageously positioned to deploy it.

Bomb damage assessment can be handled by coordinating an ad-hoc team of appropriate platforms and sensors with a weapon release event. In fact the same multi-platform team that determined target type might be asked to linger in viewing range until after the weapon impact. A process similar to the distributed ATC/ATR calculation might be invoked to provide the commander with an accurate assessment of battle damage.

The above-described scenario is of course hypothetical, incomplete and highly speculative, but it's not unreasonable to believe that technology advances can enable future systems to evolve along these or similar lines. In fact it's possible to accomplish much of this with today's technology except for the fully autonomous part.

EXAMPLE SYSTEM ATTRIBUTES

By way of summary and with some added detail, here are some preliminary concepts regarding the properties of a future system like this:

1. Mission planning will be statistical. The planning process will be aimed at selecting specific types, numbers and densities of assets to deploy in a given area, with the planning objective being to maximize likely availability of required assets when and where they might be needed.
2. Each platform will have "knowledge" of its own capabilities, both in terms of mobility and endurance, as well as its sensors, weapons and on-board processing resources. It will also require knowledge of its "neighbors" capabilities and current states. Ad-hoc team formation may well proceed by having a team leader (*e.g.,* the first platform on the scene) poll all platforms in the vicinity for their ability to support the task in question, based on their inherent capabilities and the priorities of other tasks in their queues.
3. Assets should be tasked at a high level (*e.g.,* observe a specified ground area) and should self-organize into ad-hoc teams to perform tasks efficiently with little or no external control. Marketplace techniques such as auctions may prove useful for team-forming under constraints of time, task priority, fuel reserves, *etc.*
4. Each platform will also maintain situational awareness with respect to its immediate environment and the progress of its task. It will have ability to modify its behavior in response to unexpected events, such as loss of a teammate or enemy actions. This may involve the deployment of countermeasures, finding new routes to its objective or assuming a different role in the mission execution.
5. Platforms will subscribe to a common set of "interaction" rules that allow them to collaborate with their neighbors to form teams, execute tasks and react to new situations. These rules should as simple as possible, and should manifest in message exchanges among the platforms and between the platforms and remote command facilities.
6. Commanders should have preemptive privileges over system components, and should always be aware of the states of the platforms and the overall system.
7. Sensing will involve multiple platforms, multiple sensors and multiple sensor modes. Automatic target classification and recognition will be facilitated by viewing targets from different spatial vantage points both simultaneously and at different times. It may turn out that the spatial diversity afforded by large numbers of unmanned platforms allows imperfect ATC/ATR algorithms of the sort we currently have, to perform considerably better as a group than as individuals.
8. Exploitation processing will be accomplished locally and in a distributed manner, among the platforms constituting an ad-hoc team tasked with target ID. Distributed processing should be accomplished without transmitting large data volumes (*e.g.,* video or high resolution wide-area imagery) among the participating platforms. By way of example, here's an incomplete and over-simplified set of rules for this function:
   - Wide area surveillance assets report the geo-coordinates of new detections to their "neighbors." A neighbor is any sensor that can observe the subject detection. The platform that initially detects the object will serve as team leader until the ATC/ATR phase of the process is completed.
   - Sensors receiving detection reports run local ATC/ATR algorithms on their most recent data of the subject geo-coordinates and broadcast prioritized lists of possible target type to their neighbors and the team leader.
   - Sensors that are able to take immediate new looks at the detection coordinates queue themselves to do so at appropriate (mission-specific) priority levels. New data is then processed locally and prioritized results are broadcast to neighbors and the team leader.
   - Platforms that are outside the viewing area but which can move quickly into viewing positions prepare (plan) to do so, and execute upon command from the team leader.
   - The team leader collects all prioritized ATC/ATR results from the participating platforms and determines the most likely target type. This determination includes a confidence metric which, if insufficient, causes the leader to request additional data from the team, and from any additional assets that may be nearby and available.
   - *etc…*

With respect to planning, it should be possible to develop guidelines for the planning process via computer simulation by "playing" software implementations of the system against specific missions in a Monte-Carlo manner, using different numbers of assets, different types of assets, and different pseudo-random platform deployment strategies and motion patterns. The simulation ought to embody the situational awareness and understanding aspects of the individual platform types, plus the inter-platform message exchanges that facilitate team-forming and collaborative behavior. Because of the statistical spread of cases to be tested, it may not be necessary for the simulation to be absolutely accurate. It may suffice for it to generate reasonable approximations of what could transpire in actual operations – provided all of the important behavioral characteristics of the aggregate system emerge in the simulation environment.

Simulation environments offer a good way of both designing the interaction protocols among the platforms and discovering any emergent behavioral aspects of the system that these protocols might spawn. It should be possible to model a simplified system-of-systems implementation and simulate its behavior for different missions and with different asset types and numbers. One could start by implementing very simple interaction rules and, by experimentally applying these to several representative mission scenarios, slowly evolve to more sophisticated and inclusive rule sets. When a reasonably complete set of interaction rules begins to emerge the simulation can begin to evolve into a mission planning tool.

SUMMARY

The impending proliferation of robotic platforms coupled with the advent of reliable wideband, distributed communications networking brings enormous opportunity for transformational military capability. But it also precipitates a need to rethink and restructure the ways in which we organize and control our unmanned assets and it forces us to seek new ways of conducting intelligence, surveillance, reconnaissance and strike operations. We need to consider robotic platforms and the wideband communications infrastructure jointly as integral components of systems-of-systems that include sensors, weapons and command and control assets. And we need to develop entirely new ways of doing business that will allow us to achieve critical mission objectives better and faster than ever before without straining the capabilities of our emerging shared resource infrastructure.