Dr. Tom Wagner
Program Manager
Information Processing Technology Office

We Are Not Alone

We live in a highly interconnected world, and these connections will only increase as computing and connectivity become ubiquitous. You can imagine networked processors in everything from doorknobs to refrigerators, from the warrior's helmet to his MREs. We are looking at a future of widespread, loosely coupled, distributed intelligent systems—with the possible exception of the MREs, which may or may not be intelligent. We need to create technologies that enable distributed intelligent systems to work together.

Why? We are not alone and we don't act alone. All intelligent action is interdependent. What's more, interdependent actions among multiple parties are complex to understand and control, and this has a fundamental impact on our military.

To illustrate how widespread and complex these interdependencies are, consider a simple, everyday example: buying a big-screen TV. At first glance, this may seem like a simple, independent activity. But think about it: you go to an electronics store and soon you wind up having to discuss different models with the salesperson. If you are a discriminating buyer, you might have more questions than one salesperson can answer so he might get others to help. Once you've decided which model to buy, a different clerk will ring you up.

But you're not done yet.

Now you have to arrange shipping, because you can't move the TV by yourself. You also have to be home when the TV is delivered. This requires coordination of dates and times with work and home schedules and may even force you to rearrange your plans, which can cause a cascade of changes on both fronts. Maybe there will be only one guy in the delivery truck, so you have to arrange for a neighbor to come over to help. He may have to rearrange his own plans, possibly causing another cascade of changes. You might also have to negotiate with him over compensation. Once your TV arrives, the three of you will have to work together to figure out how to get the box in the house, perhaps trying different doors or different orientations.

Of course, I'm leaving out the negotiation with your spouse, because now you're going to put a giant television in her living room. So back when you were standing there in the store, you probably had to pick up your cell, call your wife, and discuss things like the size of the TV, how deep it is, how much it's going to cost, and how long it's going to take you to pay for it. As you might imagine, this can lead to still more interconnected processes such as flower delivery. And let's not forget that you may have to take a trip to the bank to arrange financing, and you may have to contact your cable company to extend your service to high-definition and they may want to sell you a new package that changes the terms of your monthly payment. And on it goes.

Lo and behold, what appears to be a simple isolated act is actually a richly interconnected process. In this simple example alone, we touched on issues of coordination, negotiation, collaborative problem-solving, working together on physical tasks, exchanging information and artifacts, and what you might call, multi-agent planning.

Now consider large-scale military operations. If buying a TV is an interconnected process, imagine what occurs in situations where many forces are working on common goals and objectives. Getting a Special Operations team to the right place at the right time, equipped correctly, is a complex multiperson task. Scheduling and executing a joint operation—perhaps a hostage rescue—is a complex, multiperson task. In fact, I would suggest that virtually everything our military does involves cooperation, collaboration, and multi-agent planning.

The world is interconnected, and it's only going to become more so. These interconnections are not problems, but are opportunities we can leverage if we address some hard technical challenges—challenges you can help with. Let's consider some specific examples.

We're all familiar with basic parallel processing and linear speedup. Typically, if more people are performing a fixed set of tasks, they can be expected to do the same amount of work in less time.

But this view is oversimplified. Imagine I have two people performing a set of tasks in parallel, and suddenly I enable them to communicate and compare their tasks. In this case, they might be able to detect task overlap. For example, if you and I both need to go to the motor pool, you might carry the paperwork for both of us and reduce the amount of work we have to do collectively. Technically what we're really talking about here are autonomous agents that have interdependencies between their activities or have overlapping subproblems. Situations like this seem simple, but they generally require deeper analysis. For instance, you and I might both have to shuffle the rest of our tasks so our paperwork is ready at the same time.

If we looked over a longer interval, we might also determine that we're both planning to go to the operations center to pick up orders. I could trade you the motor pool run for the ops center run, and we'd both benefit directly. As before, the timing of the local and joint tasks may have to be adjusted, but here we're also talking about a trade. In these computations, the costs and benefits of the exchange must also be considered.

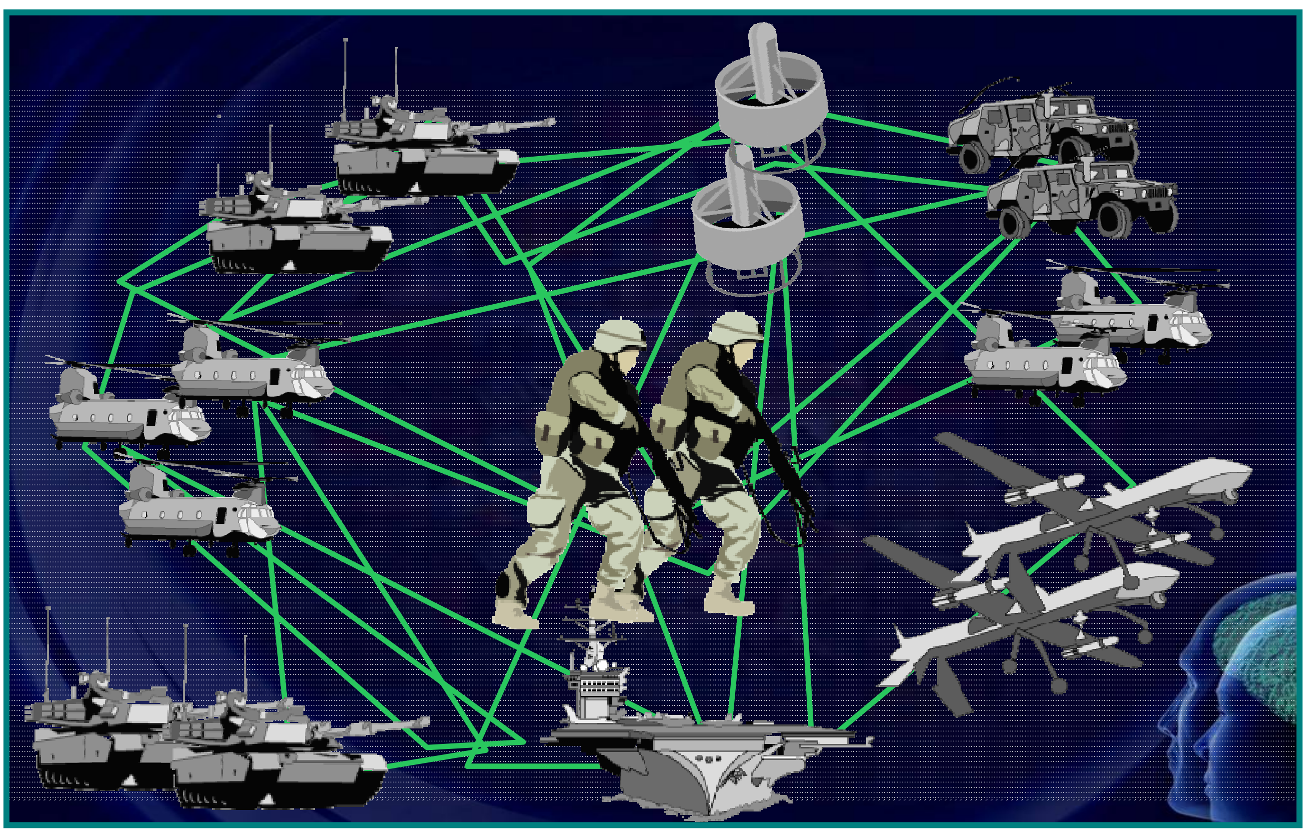

We Are Not Alone

By comparing our tasks, we might also be able to detect negative interactions. For example, you and I might both be planning to print large documents on the same printer at the same time and we might have deadlines that will be violated if we attempt to do this. The reasoning issues here are similar, although the goal is conflict avoidance.

The motivation for enabling intelligent systems to work together goes even deeper: a body of empirical data shows that on some tasks humans actually perform better in groups than we do if we work alone. Thus, it is not just that I can do more with more people, it is that the rate of processing changes. This is also a form of an overlapping subproblem, but in this case the overlap changes the characteristics of task performance.

Finally, in the human world we have to collaborate to get anything done; e.g., I have a piece of the puzzle, you have a piece of the puzzle, and someone else has a third piece. We often have to share both expertise and information in order to find solutions to tough problems.

We are not alone. We are interconnected. As are the intelligent systems of today and as will be the next-generation cognitive systems, the same systems described and envisioned in the rest of this volume. Such systems must be able to work with one another and with their human partners.

How do we accomplish this? In the past—in the recent past—we've assumed that we could simply program in coordination. This does not work. For example, in a previous DARPA program, we tried using distributed planners to control a set of simulated helicopters. The planners used hard-coded synchronization points to get the helicopters to function as a team. For instance, one helicopter would fly recon and then come back to the others and report before the team would continue their

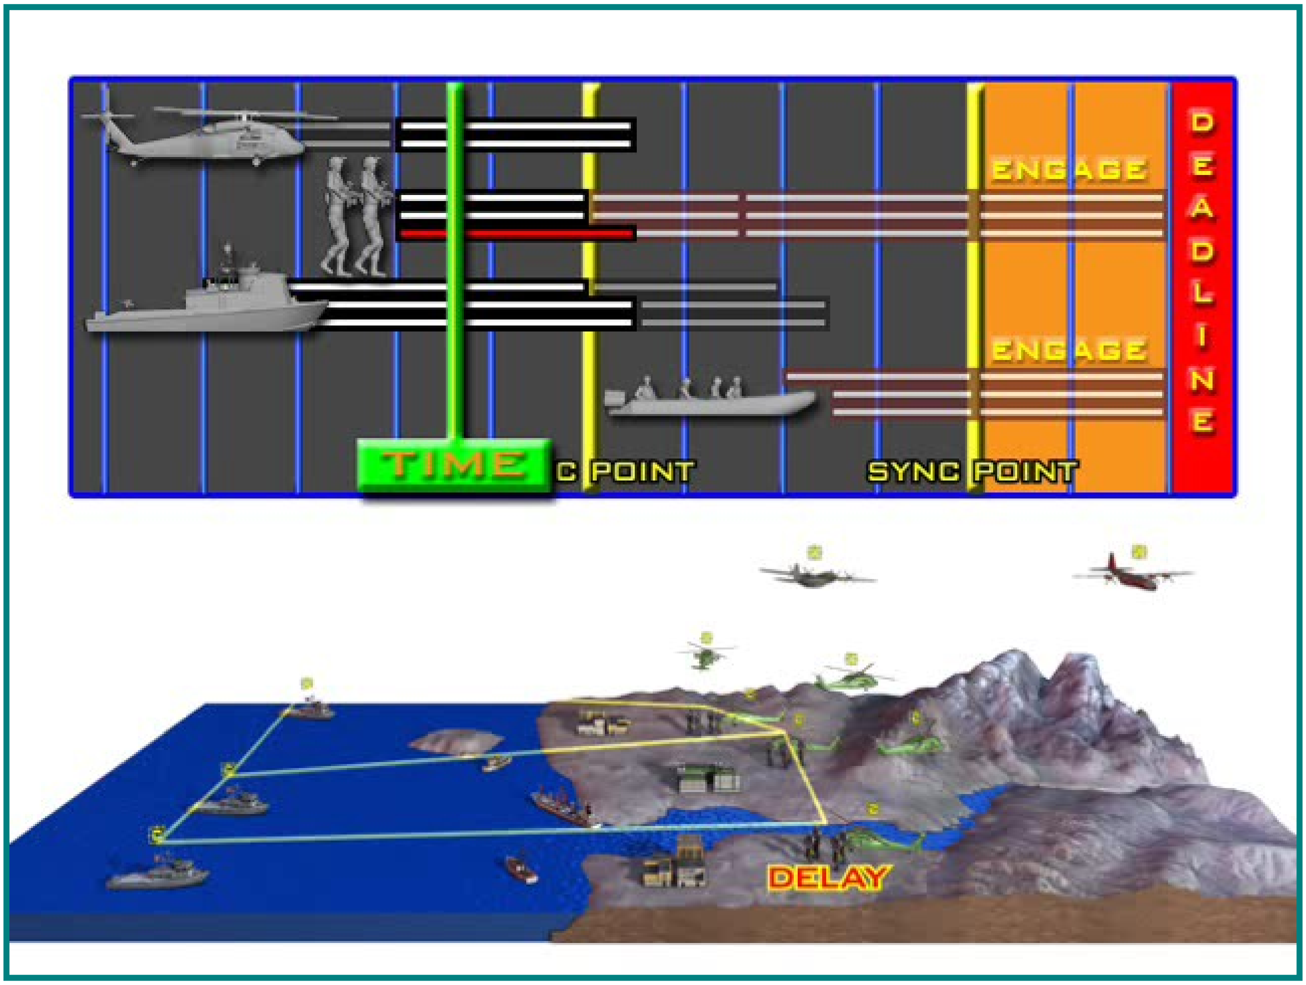

mission. Unfortunately, a problem developed during the demonstration, and the recon helicopter was shot down. Thus, it never returned to trigger the other helicopters into action, and they essentially just hovered and would still be hovering today if the researcher hadn't run out of excuses for the program manager and admitted there was a problem!

Why did the system fail? It failed because there was no deep reasoning about joint plans, no deep reasoning about the individual roles being performed and how they fit together, and no intelligent machinery enabling them to say, "Look, the recon chopper is down. Someone else has to assume his role and come back and report to us."

This is a key area of research where we are in need of a significant advance. Systems must reason explicitly about joint action in order to work together. We have some formalisms, but they don't address all the reasoning necessary to resolve task interactions or overlapping subproblems when they are detected. If you are familiar with parallel programming, this kind of reasoning would be akin to multiple instruction multiple data (MIMD) programming where the code has to be designed so it can figure out, online, at runtime, who to communicate with, what information to send, and what to do in response to any information that is received. To make matters more complicated, the what-to-do aspect of this is not a simple if/then/else block encodable in advance by a human, but is a reasoned or deliberative process that must be generated dynamically by the software. Research in software agents, like that done under the DARPA CoABS program, has made advances, but key problems remain unaddressed. We need your help in inventing new technologies for generalized collaborative reasoning—technologies that will enable many independent agents to organize and

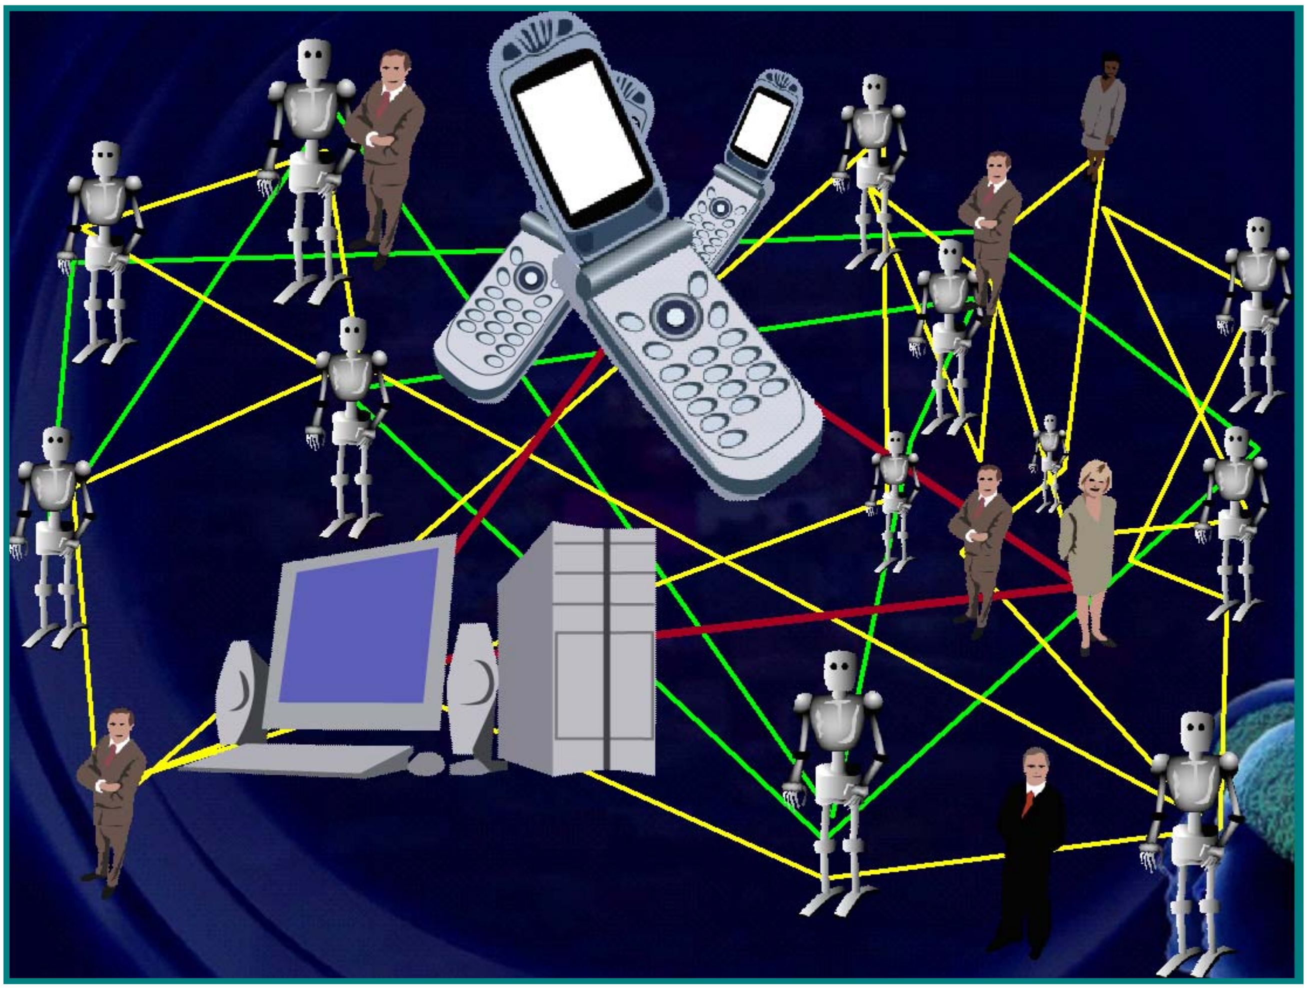

collaborate even when situated in open and dynamic environments with lots of unanticipated input from the real world.

We are not alone. We are interconnected and we can leverage these interconnections in support of our Armed Forces. The COORDINATORs program is a start in this direction. The idea is to help fielded units adapt their mission plans online in response to change. This is accomplished by having an intelligent coordination system partnered with each team. When change occurs and impacts the mission, the systems reason about the change, exchange information with each other, generate candidate response options, and even recommend a course of action. For example, if your support has been delayed, the COORDINATORs might enable you to delay by adjusting your timings and the timings of everyone else connected to you. They might also find an alternative support, for instance, exchanging air support for tank support and working out all the timings. Automating coordination reasoning will enable teams to respond more rapidly and will reduce the cognitive load on the humans. This technology will pay immediate dividends to the warfighter. It will also produce a technical foundation that can be used to coordinate intelligent computer systems as well as humans. It is one building block leading toward a larger vision. But we need to do more, and we need your help to do it.

Another building block will come from a new effort I'm envisioning. Imagine being able to harness the expertise of large numbers of humans and intelligent systems from the computer of a single Soldier or commander—possibly being able to engage in problem solving with national experts—in a sense renting their brainpower. The idea is like a problem-solving Internet combined with eBay. We would create a computational grid that connects human experts and intelligent systems that work together to solve problems. Note that participating in this grid isn't everyone's day job; instead, the humans and systems come and go at will. The idea is that customers of this grid would be able to submit problems to the network that would then work the problems and provide solutions. Of course, there are some complexities: the network has to reason about how to solve the problem, decompose it into smaller problems, find parties currently plugged in to work it, contract with them to do the work, negotiate costs, and schedule the process. The network also has to make sure progress is being made and reallocate work if it's not. There are many technical challenges in this area. For instance, how do you design a large-scale system to support collaborative distributed problem solving? How do we deal with the real-time, complex resource allocation demands that a national-scope grid would present? How do you advertise for, find, and connect the right kind of expertise at the right time—all without human intervention?

Another important building block that will play a major role in the future of our military is in cooperative efforts between humans and robots. I'm imagining a new effort in which robots and humans form teams and engage in joint physical action—possibly being trained together and coached to work more effectively as a team. Here the system is of a smaller scale and more tightly coupled than the one I just mentioned. Key research goals might include the formation and maintenance of common goals, dynamic team formation, the recognition of plans in teammates, and what you might call "social learning." If you have innovative ideas on these topics, we want to hear from you.

So, where are all these building blocks going? Where will the invention of new methods for collaboration and coordination take us?

To me, the end goal is to enable humans—our warriors and commanders in particular—to leverage their intellects to the nth degree. Imagine a future in which you can interact with other intelligent systems at will, at any distance, and share both low-level requests and high-level information, for instance, your hypotheses and half-formed plans, maybe even as you form them. Imagine another person or an intelligent system helping to refine your ideas, adding to them, or giving you new relevant information. Imagine being embedded in a cyberspace populated by humans, intelligent software systems, intelligent robotic platforms, semi-intelligent sensors, and being able to share ideas and information seamlessly. Harnessing the brainpower of large groups would certainly provide unprecedented leverage for our military decision-makers.

We are not alone. We are interconnected as are our cognitive systems. We can build technologies to leverage these interconnections—they are one of the greatest untapped enablers. With your help and ideas, DARPA will lead the way to a future where the combined intelligence of our cognitive systems will be much greater than the sum of their parts. Let’s do this together.