DARPATech 2004
Cognitive Networks

Mr. Christopher Ramming
Information Processing Technology Office
Defense Advanced Research Projects Agency

If you're anything like me, you don't need to think back very far to remember having trouble using the Internet and all of its wonderful services. Network problems are beginning to seem like a fact of life, so we learn to live with them, and also to rely on expert assistance to resolve problems. But warfighters like the ones celebrated by Time Magazine this year have something important to do and there may not be anyone to help if the network isn't working.

I would like to speak today about a transformation in-progress, and the consequences of that transformation. It is a transformation of connectedness: one that brings us into contact with our families across distances, that is—bringing economies around the world to new levels of efficiency, and that (not incidentally) is shaping military concepts of operations in the defense of this nation.

That transformational connectedness is enabled by networks and by distributed applications such as the World Wide Web and Internet Telephony. Whether used to communicate holiday greetings, shorten supply chains, or share a common tactical operational picture, these networks are revolutionizing the way we live. Increasingly, networks touch every part of our lives. It should not be controversial that networks are a potent force for peace and security.

But there are also traps and pitfalls to avoid. We can see these pitfalls by reaching back into the early days of the telephone. In the early days of the telephone, switching was accomplished by highly capable humans. In an emergency, "Mabel" (the human personification of Ma Bell in all the old telephone company anecdotes) knew where the doctor was making house calls. Mabel knew whether or not to interrupt a conversation in an emergency. Mabel's speech recognition capabilities were outstanding. And if something wasn't working, Mabel could make it right.

The problem was that Mabel didn’t scale: there were early fears about the growth of the telephone operator workforce. Only the advent of mechanical switching allowed the system to scale and achieve its intended effects without causing the majority of the American workforce to serve as in the role of telephone operator.

Extrapolating from current reliability statistics, it is clear that in a few years we will all need to be trained as network technicians. But we cannot afford for this trend to continue. The problems of network management are frustrating at home.

In an enterprise they are costly—an hour of downtime can cost millions of dollars in some settings. In the military, IT problems are simply unacceptable; lives must not be lost because a soldier is debugging a DNS problem using "ping" and "traceroute". Nor can we send a technician into battle with every soldier in order to ensure state-of-the-art capabilities.

It is our clear responsibility in the technical community to create capabilities that allow networks to manage themselves so that people can focus on other things. People are good at managing the unexpected.

When we encounter a new situation we reason around it; when we find a shortcut or a dead end we remember it—we learn it. And our systems improve over time, or at least they do when we ourselves remember that history repeats itself.

But currently, network management software is not so smart; machines are only good at doing rote things quickly. When the unexpected is encountered, our computer system break. Even if a solution or a workaround is found, it takes place outside the system and cannot be recalled by the system later on.

In general, our networks are not self-aware in the sense that they cannot make configuration decisions in the context of a mission and a specific environment. Networks that manage themselves require a new kind of technology: cognitive technology, as Ron Brachman has discussed. The kinds of cognitive technology that will be useful in future network settings have specific characteristics. In contrast to present industrial trends, we must work toward distributed solutions rather than centralized solutions.

The appropriate cognitive technology must support privacy and security; and whatever we do must be general enough to evolve gracefully over time as the underlying technology evolves so that history does not see a parade of dead-end point solutions. The challenge of cognitive networking is very hard; it is "DARPA-hard".

Think for a moment about the future: envision embedded, ubiquitous computing; imagine seamless mobility; anticipate the advantages of the truly connected world. None of these visions will see the light of day without self-managing networks; we would all need to be network technicians in the future we want to avoid. At DARPA we propose to enable self-managing networks by developing a new kind of cognitive layer.

The "why" of cognitive networks may be something that you experience on a daily basis. The "how" is not so obvious. But there are several new Ideas (with a capital I) to bring into play. The first idea is that the boundaries of responsibility between the network and the application can be redefined.

It doesn't make sense to talk in isolation about networks with four and five nines of availability, because we have all seen too many situations where the network holds itself blameless and yet the user is frustrated by firewall configurations, DHCP problems, and other details. So, perhaps network management should encompass more of what we think of today as the application.

The network should understand what the application is trying to accomplish, and an application should be able to understand what the network is capable of doing at any given moment. This would allow a cognitive network to make use of new capabilities such as the ones being explored by my colleague and fellow program manager, Jonathan Smith.

Jonathan is interested in bridging the gap between network applications and network services by learning application requirements and dynamically choosing the network protocols that will meet those requirements. Another thrust is applying cognitive technologies to create radios which are capable of situational awareness and personalize themselves to warfighter needs.

We will both need your ideas, not only to enable those breakthroughs but to shape the application/network boundary and interfaces so that every knob and dial, and every new capability, can be used effectively in support of an application's requirements.

A second Idea is that we might explore high-level languages to represent both the intended and actual configurations of a network, as well as the essence of its mission as embodied in user and operator requirements.

What is important about a domain-specific language is not merely convenient syntax, but is rather the language's concrete reflection of deep thinking that comprises a domain model, an architecture, and carefully chosen tradeoffs between expressiveness and analyzability.

Domain-specific languages could enable users and operators to describe their goals and requirements, and it is the statements in these languages that the cognitive network would use to determine the proper balance of resources. Domain-specific languages for networking would also make it possible to generate low-level configurations automatically, and the models embodied in these languages would also be useful for runtime diagnostics as a form of logical redundancy.

The challenge will be to develop high-level domain-specific languages that can be used to effectively to express concepts that drive low-level configuration and ongoing diagnosis.

Third, we can develop new operational metrics that can be used on-line by a cognitive network management system. Where today we measure the availability of individual nodes or where we refer to "Service Level Agreements" defining throughput in some generic way, we might instead think about the aggregate satisfaction of individual end users. However, it is necessary to prioritize and give preference to certain individuals or applications, especially in a military context. We might think of a priority-weighted measure, such as the aggregate utility with respect to a particular mission.

Today, we complain to an administrator who apportions limited time and attention according to an internal measure of priorities; a cognitive network would make this process of gathering end-user complaints automatic.

The cognitive network would automatically seek to balance resources and respond to complaints according to an ever-changing mission-specification, in a dynamic environment, and would do so in timeframes that are far shorter than is currently possible.

If today a key measure of certain networks is aggregate application-layer throughput, e.g. "goodput", we might brainstorm the potential meaning of a new metric called "mission-informed goodput", or shall we say "bestput".

We need your ideas to understand what a bestput-related distributed utility function should look like, how it should be affected by a mission specification, and how to calculate and use it on an ongoing basis.

Finally, machine learning and reasoning techniques that are already available can be used in the short term to build cognitive networks. This might require that we derive distributed versions of today's centralized learning and reasoning algorithms. For example, reinforcement learning provides a framework for gradually improving systems geared toward action.

Model-based systems have allowed spacecraft to operate in unforeseen circumstances without the need for human intervention. We can exploit these results. But there are also more inherently distributed approaches to explore.

One is stigmergy, a biologically inspired concept that exploits the relationship between an agent and its environment. Another is the fertile ground of algorithmic game theory, which can provide mechanisms for understanding how distributed agents will arrive at a stable systemic optimum.

We need your ideas about how these existing breakthroughs can be applied to cognitive networks in the short run, and we look for breakthrough ideas as well to enable learning and reasoning about networks and distributed systems.

All of these research areas will need to be tied together in a principled architecture for network management: one that applies not just as a point solution to one problem, but as a general new structure that will evolve gracefully as the underlying network does and that will perhaps even apply to more than one kind of network.

This architecture and the work of cognitive networking cannot be enabled by any one community today, and as a reflection of that, cognitive networking ideas are being pursued jointly between two DARPA offices, IPTO and ATO. The cognitive community is absorbing the experience with large-scale distributed systems that is a core competency of the networking community.

Conversely, the networking community is absorbing history and experience with learning and reasoning and cognitive architectures. What emerges will be an invisible college that combines the insights of both fields to achieve remarkable new results.

In the area of cognitive networking we are staking out goals that intrinsically further our understanding of cognitive techniques. Ultimately we will arrive at general solutions for at least two major aspects of network management—fault management and configuration. And in the future, when everything is connected and automated, we will all be able to focus on our day jobs.

Concepts of operations will not rest on briTTle software that could fail catastrophically at any moment. Networks will be resilient, and they will be taken for granted as we all get on with the really important, transformational work that communications enable as we strive for peace, prosperity, and the defense of this nation.

Thank you.