Impediments to Systems Thinking: Communities Separated by a Common Language

William L. BAHN
Leemon C. BAIRD III
Department of Computer Science
United States Air Force Academy
Colorado Springs, Colorado 80840

and

Michael D. COLLINS
Department of Defense

ABSTRACT

Systems-thinking requires participants to view problems, and their solutions, within the context of the overall system. When this involves participants from diverse professional communities, several barriers to effective communication can arise. As in all human communities, shorthand representations, such as specialized jargon, that aid in efficient and precise communications have developed. These shorthands rely on the common background of the participants, including common knowledge, perceptions, and assumptions about the field and its concepts, capabilities, and limitations. These shorthands can become a source of miscommunication when participants are from different communities.

The development of concurrent coding theory has brought members of the network, information assurance, and communications communities together in ways unfamiliar to most. This is due to the manner in which the theory is applied to omnidirectional communications to mitigate hostile jamming without the use of shared secrets. Being novel and unconventional, several conflicting assumptions have been exposed.

Here we discuss our experiences in recognizing, analyzing, and overcoming many of these misconceptions, including our own, despite how surprisingly resilient some of them have proven to be. However, by adopting a systems-thinking approach and drawing upon most peoples' innate ability to think analogically, we have significantly improved our ability to convey our key concepts.

Keywords: Systems thinking, Jam-resistant communications, interdisciplinary communication.

INTRODUCTION

It is not surprising that the same word or phrase can mean different things to different people. There is the old joke that when told to ”secure” a building, Marines will initiate a frontal assault, soldiers will occupy the building and prevent all others from entering, sailors will turn off the lights and lock the doors, and airmen will take out a long-term lease with an option to buy. While humorous, each interpretation is largely consistent with how that term is commonly used within each of those communities. To be fair, the context within which the term is used is very important and, in practice, members of each community would likely extract the correct intended meaning based on that context. Humans are actually quite adept at sensor fusion - taking input from may different sources and cues and integrating them into a reasonably accurate overall interpretation.

We humans communicate vast amounts of information and, in doing so, we naturally seek to be as efficient as possible by conveying the most information with the fewest symbols. Our arsenal consists of many different media ranging from the spoken and written word to facial expressions, gestures, pictures and many others. In technical communications, however, we are frequently limited to just a few media, such as words, equations, and graphs.

Different technical communities are not unlike different societal communities and they almost invariably develop ways of conveying significant amounts of information using few words by drawing upon a common background and understanding of the concepts in that field. By-and-large, this is a very positive thing, but it frequently leads to miscommunications when people not as deeply immersed in those concepts are involved in the conversation. This is aggravated when field-specific jargons include terms and concepts that are used by other fields, or even in common parlance. For example, cryptographers discussing “shared secrets” and “secret sharing” would be dealing with two very different and almost unrelated concepts, but “outsiders” would likely have difficulty following the conversation because, on the surface, the two phrases appear largely equivalent.

Making matters worse is that many professional fields have overlaps involving concepts common to both fields but the

This work was sponsored in part by the Air Force Information Operations Center (AFOIC), Lackland AFB, TX, and was performed principally at the Academy Center for Information Security (ACIS) at the United States Air Force Academy.

level of understanding by the members of one community as far as how those concepts are dealt with by the other community is frequently quite superficial.

In developing the theory of concurrent codes and, more to the point, its application to omnidirectional jam-resistant communications that do not rely on shared secrets, we have discovered several key terms and concepts that the various communities involved in the implementation of secure and reliable ad hoc wireless networks do not completely agree on and, more importantly, do not recognize that significant disagreement even exists. To paraphrase George Bernard Shaw, they have become communities separated by a common language.

THE BACKGROUND DISCUSSION

The Problem - In a Nutshell

The framework discussion that led to this paper involves the need to provide jam-resistance in the ad hoc wireless networks that will form critical portions of the Global Information Grid (GIG)[1], [2]. The GIG is anticipated to play a vital role in nearly all aspects of military operations as the United States pursues a doctrine of net-centric warfare.

Inevitably, the ”edge of the GIG” will be largely wireless since a primary goal is to network fielded elements so that real time information and feedback can flow between them and higher echelons. In the more ambitious models, individual pieces of equipment (e.g., aircraft, vehicles, and rifles) and perhaps even ammunition (e.g., missiles and long range artillery shells) will be nodes on the network. But even without this, such a network will be ad hoc, mobile, highly dynamic, and involve many omnidirectional wireless links.

Constructing and administering such a network is a daunting task under ideal circumstances and the GIG will not be so favored. As the United States becomes more reliant on the GIG, our adversaries will become more motivated to disrupt it and will devote more resources to that cause. While the type and variety of attacks the GIG will face will certainly be numerous, our adversaries will seek the Achilles' Heel of the system - the attack that provides the most damage for the least effort. It's impossible to say what that Achilles' Heel will be or how vulnerable it will be. One possibility that suggested itself to us is the reliance on spread spectrum technologies to provide jam resistance in such a large, diverse network. While spread spectrum can provide significant jam resistance to an omnidirectional radio link [3], [4], [5], it does so via a *shared secret*. A shared secret, also known as a *symmetric key*, is information that authorized parties must have but that must be denied to unauthorized parties. One example is the exact sequence of frequencies that will be used in a frequency-hop spread spectrum system.

Unfortunately, virtually all of the jam-resistance afforded by a spread spectrum system disappears if the adversary obtains the key. Hence an integral part of analyzing the jam-resistance of a fielded system must include assessing how likely the adversary is to succeed in obtaining the keys. Consequentially, the question is not only a signal and systems question, but also an information assurance question and, ultimately, a network question since key management will have a significant impact on the network implementation.

The management of symmetric keys in a system involving hundreds of thousands, if not millions, of nodes is untenable. This is especially true when many of those nodes, such as distributed sensor networks scattered throughout a theater of operations, will be located amongst the enemy where it is inevitable that many will be captured and compromised.

If omnidirectional wireless links are to play a major role in the GIG, then those links must possess an adequate and robust level of jam-resistance. If spread spectrum is going to provide that jam resistance, then either a reliable means of providing symmetric key management on the scale involved must be devised, which seems unlikely, or a form of spread spectrum that does not use shared secrets must be found. The latter requires the ability to exchange information in the face of hostile jamming even when the jammer knows everything that is common to both the sender and receiver.

The Proposed Solution - In a Nutshell

An analogical approach to this problem naturally recommends itself because, in many ways, the situation in the cryptographic world prior to the introduction of asymmetric cryptography in the early 1970's [6] is highly parallel. In that case, an alternative to centrally managed symmetric keys arose via the development of a Public Key Infrastructure (PKI) to exchange temporary symmetric keys (known as session keys) between nodes that have no prior shared secret. The session key exchange is done using only publicly available information; this is the way that secure internet (https://) transactions are performed countless times each day.

With this recognition, the problem is then immediately reduced from finding a way to carry on an extended conversation without shared secrets to finding a way protect a conversation just long enough to exchange a session key. Just as in the existing PKI analog, key exchange thus does not have to be particularly efficient given its short duration. This is important since symmetric ciphers are typically a couple of orders of magnitude faster than their asymmetric brethren.

As with all analogies, this one can't be taken too far. The asymmetry in PKI may permit two people having no prior relationship to carry out a conversation such that an adversary can't understand it, but the conversation itself can still be easily jammed. The need here is for a system whose asymmetry permits two people to carry out a conversation such that the adversary can't (easily) prevent the message from being conveyed.

Traditional spread spectrum offers no solution because, without the key, there is no asymmetry to exploit; the adversary can easily jam the key-exchange transmissions. Tempting though it might be, turning to error correcting codes also yields no solution. While such codes perform very well in the presence of non-malicious noise, they are easily defeated by malicious noise; all an attacker must do is superimpose one or more legitimate messages on top of the genuine message and the genuine message will be blocked.

Because it is the physical medium that is being jammed, it is the the physical medium that must provide the necessary asymmetry; this places constraints on the adversary that can then be exploited in the coding scheme. The simplest asymmetry is the fact that transmitters can easily add energy to the medium but can't easily remove it or mask it. To exploit this asymmetry a new coding theory, namely concurrent codes, has been developed that strives not to correct for errors in a transmission, but to separate concurrently transmitted messages so as to successfully recover them all.

A meaningful presentation of concurrent coding theory is well beyond the scope of this paper; details can be found in the original tech report[7]. Here we will describe only the most basic concepts and how they require a significant departure from traditional transmission schemes.

Central to the success of a concurrent code is the notion of an "indelible mark". By this we mean a transmission of some type, such as a short burst of high power noise, that is very difficult for the attacker to mask. It is fine if they distort it, for there is no need to extract data from the pulse - the existence of the pulse *is* the data. In and of itself, this does not represent a significant departure from traditional transmission schemes, in fact it is the oldest scheme and dates back to the days of spark-gap generators. What is different is that concurrent coding theory can tolerate large numbers of similar indelible marks inserted by the attacker. The transmission can still be blocked - nothing is jam proof - but a genuine sender might transmit several hundred marks to send their message while an attacker might have to transmit hundreds of thousands or even hundreds of millions to block its reception. Such an asymmetry in energy expenditure places the attacker at a significant disadvantage - they must not only have the capabilities to commit sufficient energy to the attack, but they must do so without attracting unwanted attention from the network defenders.

But this "old fashioned" approach is at odds with modern digital transmission schemes that seek to transmit large amounts of data using techniques that minimize the bit error rate (BER). In doing so, they almost universally strive to achieve symmetric bit error probabilities. However, the use of concurrent codes requires a major departure from this approach in that we are largely unconcerned with how little data can be transmitted (the primary goal is exchanging a relatively short key at the beginning of a conversation) and we are specifically seeking highly asymmetric error probabilities - we can tolerate a large number (up to half) of the zeros being turned into ones but if more than a few ones are turned to zeros the message will not get through.

THE ORIGINAL APPROACH

Initial presentations on concurrent codes and their application to the problem of jam-resistance used an approach that varied only slightly depending on the audience. Technically oriented audiences received a more detailed explanation of the algorithm whereas audiences with a managerial leaning received a presentation that emphasized the high-level problem. In doing so, we allowed one of our own preconceptions to guide our actions. We assumed (correctly, in most cases) that people from a non-technical background would not find the low-level details of our algorithm particularly interesting or helpful; these tended to be "big picture" people that wanted to understand the high level nature of the problem and that our work offered a potential solution. But we had also assumed that people with a strong technical background related to computer networks, information assurance, or communication systems possessed an essentially homogeneous ability to work with the basic concepts of all three areas with perhaps the need to fill in a few gaps. So when we began collaborating with members from various technical communities, we tended to discuss concurrent codes as a jam-resistant alternative to traditional spread spectrum that could be used even without a shared secret. Our entire viewpoint, directly or indirectly, was presented from that premise. We generally provided a short background on the growing reliance on large wireless ad hoc networks and the problems of managing symmetric keys in such an environment. As part of this we would point out that traditional spread spectrum loses its jam resistance if the keys are compromised. At that point we believed that adequate groundwork had been laid to allow the discussion to proceed to our theory and algorithms.

THE BARRIERS TO SYSTEMS-THINKING

What we discovered was that our approach was obviously flawed and that this not only resulted in poor communications, but were a threat to continued collaboration. This forced us to reflect on why we weren't getting our ideas across. First we identified the major barriers to communication. This was done by noting which concepts had to be revisited frequently and what connections between concepts weren't being made in the manner we had expected. It also became apparent that we were not keeping audiences focussed on the main problem. While we expected and needed discussions to delve into the fine points, we found that they did not remain faithful to the constraints of the whole system and frequently became counterproductive.

We identified several recurring misconceptions that were coloring the discussion and were surprised how resilient many of them were; even after appearing resolved, they tended to play significant roles in follow-on conversations and had to be revisited multiple times. We concluded that the difficulty wasn't that a particular concept was being repeatedly forgotten or misunderstood, but rather an awareness of its relationship to and consequences on other concepts

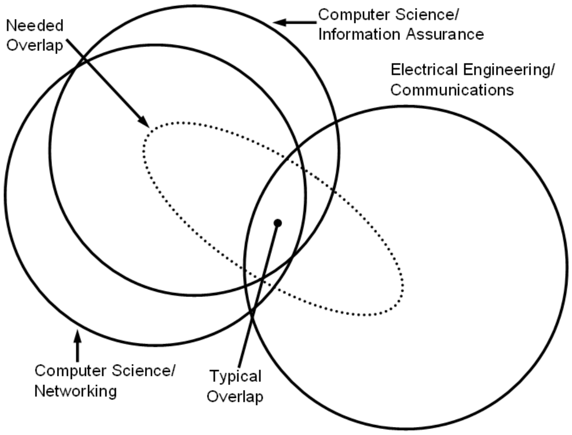
Needed Overlap
Computer Science/ Information Assurance
Electrical Engineering/ Communications
Computer Science/ Networking
Typical Overlap

Fig. 1. Overlap of the technical communities involved.

was faulty or missing altogether. We further noticed that these misconceptions followed a fairly predictable pattern and were highly correlated to which of the three communities the particular person was from.

Figure 1 is a very rough depiction of the degree to which the three communities overlap and the degree to which they need to overlap to work on the keyless jam resistance problem. As the figure shows, and as would be expected, there is a very large overlap between the networks community and the information assurance community since they are both specialties under the computer science umbrella. It should be noted that our observations and conclusions are highly subjective and generalized; they most certainly do not apply equally to everyone.

The Network Community

The barriers stemming from the computer network community appear to be unintended consequences of the somewhat arbitrary partitioning of network systems along the lines of the OSI Seven Layer Network Model [8]. The lowest layer in this model is the Physical Layer (PHY) which is responsible for conveying a binary bitstream from one node in the network to another via some appropriate physical means, be it a piece of wire, a laser beam, or a radio signal. The majority of people in the network community seldom work with the PHY and tend to view it as a black box that feeds them bit strings from other users and sends their bit strings to other users in turn. The inner workings of PHY layer devices are largely seen as the realm of the communications community. It's understood that the bit strings may or may not arrive at their destinations and may have errors that higher level layers in the model must deal with. The use of error correcting codes can help with a certain level of random corruption and sophisticated media access control (MAC) protocols have been developed to cope with interference amongst competing users of the physical media. The unstated assumption made by most network designers, however, is that users, even malicious users, of the network will play by the MAC protocol rules. The real assumption being made is that the designers of the physical layer have successfully accomplished their task and kept those that don't play by the MAC rules out of the network. Consequently, the need for jam-resistance and participation in ways to achieve it seldom appears above the horizon. They "know" that jam-resistant links exist and are content to leave it in the PHY black box. We have even been told that jam-resistance has nothing to do with either computer science or computer networks, that it's a electrical engineering problem. Others, including some that are responsible for overseeing the security aspects of implementing ad hoc wireless networks, are only interested in working "higher in the stack."

The Information Assurance Community

The information assurance community is well aware of the scalability limits imposed by the key management problems associated with symmetric-key based systems. In fact, it is so apparent to them that there appears to be a tendency to expect it to be as apparent to people in both the network and the communications communities as well, which we have observed is far from a universally valid assumption. Like the network community, they tend to see the physical communications process as being unrelated to keys and key management. They *also* "know" that jam resistant communication systems exist and assume that they are adequate. Worse, because of the assumption that the members of the other communities are aware of the symmetric key management problems, they tend to assume that other communities will avoid implementing large networks that require symmetric keys.

Another factor at play is that the information assurance community does not always give equal weight to each of its security goals: confidentiality, integrity, and availability. In particular, the availability goal is frequently treated only in a very narrow sense, such as dealing with specific types of denial of service (DoS) attacks. Sophisticated schemes are routinely discussed concerning problems caused by someone injecting counterfeit messages into a system; virtually no time is spent talking about someone preventing any messages from flowing by means of jamming the transmission links. When such problems are discussed it is often as a footnote pointing out that such problems are the purview of the communications community and that they have various jam-resistant ways of communicating data from one point to another.

The Communications Community

Of the three communities involved, the communications people have had the hardest time incorporating the new concepts into their thinking. This appears to be the result of three factors: (1) the use of concurrent codes requires a significant step backward in communications technology, (2) the awareness of what constitutes the system "key" appears to be too superficial, and (3) the basic need for something like concurrent codes is not obvious because it conflicts with certain ingrained assumptions about the capabilities and limitations of key management systems.

As mentioned previously, concurrent codes require highly asymmetric channels the likes of which the communications community has invested a great deal of time avoiding. They are considered “a bad thing” because they increase overall bit error rates in most modulation schemes. As a result, it was repeatedly explained to us why we shouldn’t use them even before we had the opportunity to explain that the asymmetry is critical.

A second barrier that took a long time to even recognize, because of the assumptions we were making, is that many people in the communications community view the “key” in a very narrow and literal sense. To them, the key is information such as where in the spreading code the sender begins the transmission. When we spoke of the “key being compromised” we meant that any and all secret information has fallen into the hands of the attacker. Yet on numerous occasions the people we were talking with would assert that even with a compromised key a spread spectrum system still retains a significant amount of jam resistance. It wasn’t until one person specifically stated that this was because even if the attacker has your key they still won’t know all of the “side information” such as the time of day that a transmission will be made. In some cases it significant effort to establish that any and all such side information is part of the key and that managing it is exactly the same as managing what they had previously understood the key to be.

The third barrier is essentially the same as the PHY layer barrier facing the other two communities, only in reverse. The communications crowd “knows” that the network and information assurance folks can encrypt and decrypt information and since the secret keys are simply information, transferring keys to the people that need them is a solved problem. This it is not a consideration in analyzing or designing a spread spectrum system. It is simply taken for granted that the attacker will be denied the key.

Perhaps more accurately, the view is that key security is simply an intrinsic part of using spread spectrum securely and that there is no way around it. This is reinforced by the apparent fact that many, if not most, people in the communications field are not even aware that asymmetric cryptography and PKI exist, despite the fact that they use it whenever they make an on-line purchase. At first this startled us until, upon reflection, we again realized that it was us making false assumptions regarding how widespread topics that seem ubiquitous to us really are.

The Major Misconceptions Between Communities

Because the overlap between the networks and information assurance communities is high, the misconceptions between them were fairly minor, so we will lump them together from this point on. The major misconceptions described above can possibly best be summarized in the following oversimplified and overgeneralized statements:

Viewpoint of network and information assurance folks: We won’t worry about using ad hoc networks, because the communications folks will take care of jam resistance.

Truth about communications folks: We can’t do jam resistance in omnidirectional links without a shared secret.

Viewpoint of communications folks: We won’t worry about using shared secrets, because the network and information assurance folks know all about key management.

Truth about network and information assurance folks: We can’t keep a secret, not when it’s shared by half a million people.

The underlying problem revealed by the above viewpoints is that none of the communities involved are operating from a systems-thinking perspective. They are trying to operate in isolation within their own specialties and throw things over the proverbial wall to each other; their view is too parochial.

BREECHING THE BARRIERS

A fair degree of success in overcoming these barriers has been achieved using the following approach: (1) use a “system focus” to frame the discussions; (2) speak the native tongue; (3) educate each community about the relevant limitations of the others; (4) use analogies where possible; (5) use simple or extreme examples that force people out of their comfort zone; and (6) meet face-to-face in interactive settings whenever possible.

Maintaining a focus on the overall problem is critical. Technologists tend to get caught up in the details, it’s why we exist. But we can lose sight of the big picture and how those details have to come together in the end. By frequently tying the discussion back to the system-level view, we overcome that tendency and establish a common framework of reference for all of the communities involved in the discussion.

Obviously, “speaking the native tongue” involves much more than just using the right jargon. A concerted effort must be made to understand what those words mean to the people of each community and, in particular, how those meanings differ. This is not easy and it can require a careful examination of what we, ourselves, mean when we use particular terms. In practice it comes down to keeping an eye out for miscommunications and then analyzing the approach and adjusting it accordingly.

The need to educate each community on the limitations of the others results from the superficial understanding we all tend to have regarding fields removed from our own. We generally overestimate the capabilities of other fields because we have not been exposed to the fine print.

The use of analogies can be quite helpful, but any analogy can be stretched too far; always be on the lookout for people drawing more from an analogy than is warranted. From our perspective, the very fitting analogy to the development of asymmetric cryptography and PKI arose early on, but we didn't leverage it nearly as much as we do now.

We have found that using either simple or extreme cases and examples can force people to question their own assumptions. A particularly good example of this occurred when we were having trouble pointing out the fundamentally different nature of the jamming problem when the adversary can transmit waveforms that are just as valid as what the genuine sender would generate. At one point, somewhat in desperation, we pointed out that unintentional jamming of a similar nature occurs all the time in narrowband systems, such as walkie-talkies, when two parties transmit at the same time. Everyone immediately agreed, but when we suggested that an analysis of that situation was particularly relevant, the response was that such an analysis is never done, to which our reply was, "Why not?" Someone claimed that it wasn't applicable since we weren't talking about a narrowband system. We pointed out that it had already been agreed that, with a compromised key, a spread spectrum system is effectively reduced to a narrowband system where the jammer's capabilities are concerned. The question, "So if that is an effective means of jamming, why shouldn't it be analyzed as such?" resulted in a very fruitful discussion expected to result in at least one joint publication.

Finally, we have found that miscommunications and misconceptions that have survived a forty-one page tech report are also likely to survive days of e-mail exchanges but that a fifteen minute phone call or five minute meeting in person usually resolves them. An interactive setting allows us to draw upon additional communications media, such as tone of voice, facial expressions, or even posture. With those as a guide, we can frequently recognize that a known misconception is coming into play. For example, in one meeting with a dozen very senior communications engineers we recognized that we had strayed into unfamiliar territory for them and quickly determined that none of them had ever heard of public key cryptography. Immediately upon recognizing the poor assumption we had made in preparing our briefing, we were able to step back and provide a five minute overview of the subject. That was all that was needed to bring this group of very bright and competent people to a point that they not only grasped the fundamentals of public key cryptography, but also of digital signatures and how they can be used to provide authentication. Had we not had an interactive setting or had we chosen to "stick to the script" an extremely effective presentation probably would have been an extremely disappointing one instead.

THE SYSTEMS THINKING APPROACH

With everything we have learned, we now, first and foremost, push hard for a face-to-face meeting when presenting our work to a new audience. We then being our presentations with a high level overview, much as we previously did, but we emphasize certain concepts, namely key management, the scale of the networks being planned, and the motivation for hostile players to seek any effective means of disrupting those networks. As we discuss the lower level details, we frequently re-introduce those central systems-level concepts making clear how they relate to those details. We then make a point of discussing how the limitations of various parts of the system affect other parts, taking into account the makeup of the particular audience we are addressing. For instance, when talking to communications folks, we make a point of emphasizing that any and all information that the attacker is not supposed to know is part of the key and that the information assurance guys are not capable of managing keys on that scale because the problem is just too massive. To make this point clear, we use an example wherein the enemy captures a vehicle, its radio, and its radio operator and then applies a rubber hose to the soles of the operator's feet until all secrets have been revealed. In contrast, when talking to the network and information assurance crowd we emphasize that, in order to provide the jam-resistance that they are taking for granted, the communications folks will end up saddling them with a key management problem from their worst nightmares. Perhaps most important, we try to be very diligent in detecting when any miscommunication or misconception is coming in to play and go to significant lengths to quash it as quickly as possible.

CONCLUSION

Adopting a systems-thinking approach, particularly when involved in an interdisciplinary discussion, and diligently looking for ways to accommodate the diverse terminology and ways of thinking that typify the various communities involved can substantially reduce the number of miscommunications that occur and make overcoming the ones that do occur proceed more smoothly.

REFERENCES

[1] http://www.nsa.gov/ia/industry/gig.cfm?MenuID=10.3.2.2.

[2] Defense acquisitions - the global information grid and challenges facing its implementation. U.S. Government Accountability Office, 2004.

[3] A. Belouchrani and M. Amin. Jammer mitigation in spread spectrum communications using blind source separation. *Signal Processing*, 80(4):723–729, April 2000.

[4] L. B. Milstein. Interference rejection techniques in spread spectrum communications. *Proc. IEEE*, 76(6):657–671, June 1998.

[5] G. J. Saulnier, Z. Ye, and M. J. Medley. Performance of a spread-spectrum ofdm system in a dispersive fading channel with interference. In *Proc. MILCOM Conf.*, pages 679–683, 1998.

[6] W. Diffie and M. E. Hellman. New directions in cryptography. *IEEE Trans. Inform. Theory*, IT-22(6):644–654, 1976.

[7] L. Baird, W. Bahn, and M. Collins. Jam-resistant communication without shared secrets through the use of concurrent codes. Technical Report USAFA-TR-2007-01, United States Air Force Academy, 2007.

[8] A. Tanenbaum. *Computer Networks*. Prentice Hall PTR, 2003.