

The Technical Issues in Geographically Distributed Congressional Deliberation and Decision

Testimony given to the House Administration Committee of the
U.S. House of Representatives
May 1, 2002

Robert Thibadeau, Ph.D.
Director of Security Architectures
Seagate Technologies, Inc.
Pittsburgh PA

The House of Representatives or other deliberative lawmaking bodies may find it important to function during times when the voting members and supporting staff are geographically distributed. The purpose of this statement is to try to identify the principal technical issues that need to be effectively addressed in order to enable distributed deliberation and decision by the entire congressional membership.

The best approach to this problem is clearly not to look at available technologies first. If we look at communications technologies first we are apt to create an awkward, and perhaps unusable, patchwork of solutions. The preferred approach is to analyze the deliberative and decision-making processes that we wish to achieve, and then identify the necessary and sufficient technical means to achieve these ends. This brings us to the first issue:

Should we attempt to mimic the deliberative and decision-making processes as they exist in the current congressional context, or should we develop a separate set of processes more suited to the circumstances of distributed action?

There is no easy answer to this question. It is almost certainly true that if we make any physical change to the geographical context of congressional processes, the processes themselves will change. As a matter of practice, it will be impossible to perfectly mimic all forms of human interaction that take place in congress that can influence collective decision making. Trying to mimic may well be the wrong goal. A better goal may be to achieve outcomes that members feel comfortable would be the same outcomes achieved if the membership were meeting in normal congress. The telecommunications and computing tools needed to achieve this effect of natural outcomes may well not at all mimic the ways in which the membership interacts on congressional hill.

However, there are certain basics of what it means to be in congress that allow us to confidently define other issues that will most certainly need to be addressed. So, let's just take the dictionary definition of congress as a "formal assembly of representatives to discuss problems and legislate." This congress can be characterized scientifically as a dynamic matrix of communications among members, staff, and public. The communication is not constant, but ebbs and flows. A congress naturally incorporates provision for time to study and reflect, and for many kinds of special interactions in groups that precede the full congress assembled.

For those of us who have had many years of daily experience with computers, the technology of

chat, as may be found in the original Internet Relay Chat (IRC Chat), or AOL Instant Messenger, provides a natural way to permit members, staff, and public, to both publicly and privately deliberate. The Blackberry mobile email devices in use by many legislators also provides chat-like interactions among members, staff, and public. But, while Chat has many of the correct properties of congressional interaction, it is only for textual input and does not have easy means for audio-visual interactions. Chat may seem appropriate, but this is letting available technology drive our thinking about congressional processes. We can't really evaluate whether we want to change congressional processes until we can enumerate them. So, the second issue that we can identify is this:

Can we define, in precise terms, the model of communications necessary for the congress to deliberate and decide?

Clearly, this model of communications needs to describe the ebb and flow of private, semi-public, and public deliberations and decision-making. It needs to specify rules, timing, meeting, and authority. It also needs to specify all modalities of communication needed including textual, documentary, audio, visual, and perhaps others (such as gestural) in a context where everyone finds the communication natural and appropriate.

We know that there will be certain technical limits. For example, it would be impractical to have 435 live talking heads on a screen, first because that many talking heads cannot be put on a screen, second because it would lead to incomprehensible jabber, and third because the bandwidth limitations of our telecommunications infrastructure simply make 435 live talking heads impossible. Of course, it is possible to have all 435 Members on line simultaneously, just not in live video. This brings us to a third issue:

Given a desirable model of communications, how do we technically accomplish this with a natural, transparent, user interface?

So, we may, for instance, find that instead of 435 live talking heads, a given member of congress may choose to watch the Speaker's dais and follow one conversation at a time with only a few participants on screen at any given moment in time. Staff may track other events occurring in parallel or certain events may be stored for replay. Priority interrupts may be possible by senior members needing to intervene in the interest of timely decision making.

Supposing that we resolved all these issues to congressional satisfaction, many other issues still remain. Many of these involve information security. There is a fairly well understood technology associated with security. Consider any given security problem, we can measure security against six considerations: integrity, privacy, authentication, authorization, audit, and availability. For example, if we consider the security of the communication between two members of congress, we uncover new issues:

Integrity: Is the integrity of the communication preserved and not tampered with?

Privacy: Is the communication hidden from all those who are not authorized to receive it?

Authentication: Are the parties to the communication actually the people who they pretend to be. Is this really this particular member of congress?

Authorization: Is this party to the communication authorized to be a party to the communication? This may not seem like much if we are just talking about two individual members of congress, but

what about authorizing all democrats to a caucus, or all committee members to a meeting?

Audit: Is there a record of the communication that may need to be consulted if a violation of security is suspected but not detected at the time of the violation. We may note that many of the most damaging attacks against security are covert attacks that can only be revealed through forensic analysis of logs and audits.

Availability: Is the timeliness of the communication protected against denial of service attacks or, even, simple system failures. We can note that human communication can be easily disrupted by simply creating artificial delays.

These six issues need to be addressed for any aspect of the security problems that can be identified in the dynamic communications model and the systems needed to achieve a viable user interface. So, for example, during a vote it is necessary to determine that all votes are correct, they are anonymized or hidden (as with voice votes) when they need to be, they are coming from who they seem to be coming from, they are all authorized, a record is kept in case of suspected security failure, and the votes can be executed and completed in a timely fashion.

We can also raise these six issues for any component of any particular process. So, for example, we can simply examine the digital link between two locations and confirm that the integrity of the link is not compromised, that the information is hidden from anyone who does not have a right to see it, that end points of the digital links are indeed the end points that they pretend to be, that the communication along the link is authorized, that it is audited, and that it is available at all times required by the dynamic communications model and the user interfaces.

So, the above six security issues really multiply into quite a number of issues because of all the types of communications necessitated by the dynamic communications model and the user interfaces. There is no reason to enumerate all these types except to raise this enumeration, itself as an issue:

Can we enumerate all the security problems associated with the dynamic communications model and the user interfaces?

The answer is that we can probably do this. However, it is not a simple matter. One can take a seemingly simple act, like a single email message, and analyze this down through a very large number of potential security problems. For example, many people don't realize how easy it is to spoof an email server or to sniff email. It is also possible to create a 'man-in-the-middle' that can alter the email message in ways meant to change behavior. The integrity of seemingly instant actions can also be breached by system operators. How do people know that I've unplugged five Members of Congress just before a vote? The practical method to enumerate the security problems is to enumerate or identify only those security problems that are suggested by risk analysis to be worth analyzing. The risk is ultimately to a breakdown of the congressional decision making itself, but minor risks, such as the risk of a brief delay in email delivery, is not worth mentioning.

There are some other issues that I would term "special issues." I will take a few of these in turn:

How do we authenticate a member of congress?

This turns out to be a very interesting problem. The simplest notion of authentication is that the

member of congress logs in with his username and password. But we know that someone else may guess a password. We may require, then, that the member of congress use a physical, unique, token, such as a smart card, along with a password, to authenticate himself as being the actual congressman in question. Finally, we may go beyond what the congressman has and what he knows, to what he is. We may use a fingerprint scanner, voice recognition, or an iris scanner to identify that this is truly the congressman in question.

But there is more to authenticating the congressman than just this. In a deliberative, interactive, process we may wish to constantly know that this is really the congressman. For example, suppose an attacker knocks the congressman out after the congressman logs in, and now the attacker can appear to be congressman. The authentication may need to be continuous. This special issue of authentication can become quite important.

But the special issues of authentication do not end with this. Since a congressman can vote, how do we know when the congressman is 'just absent' versus if he has been incapacitated in a fashion that allows a secondary authority to vote in his stead? We need to both authenticate his presence and authenticate the nature of his absence. This would also be true of staff. Indeed, authenticating absence is at the root of much trust among members and also in establishing rights.

How does the public play a role?

The dynamic communications model will need to have a role for public input and public inspection of both process and decisions. Clearly, even the congress assembled on the hill today are concerned about misuses of public access – access that can distort proven deliberative processes. How can the public have input if the Members and staff are at locations where physical access is not possible? What kind of feedback is needed? Perhaps all that is needed is that the public can see the summaries of their input so that they know that the congress is taking note of the public debate. This is an issue needing serious study.

It also brings us to still another very interesting special issue:

How can we confirm that a congressman has actually reviewed the material we think, or hope, he has reviewed?

If we simply look at the technical communications model, it may not be sufficient that the mail has arrived in the congressman's inbox. We may need to know that he has read the mail. With public input, we may want to know that the congressman has at least looked at the summary statistics. A very effective security attack, which can also be 'socially engineered' by the bad guys, is to create the circumstances under which a decision maker does not review certain documents or material key to his deliberation. On the hill, a quick remark or a glance, can confirm that someone has read a document in question. In a geographically distributed system, there may need to be other methods. Ideally these other methods will be well engineered for ease of use and also protect privacy where appropriate.

Finally, let us bring up perhaps one of the most interesting of the special issues. This is the one that revolves around emotion:

Compromise and negotiation as often as not involves a clear understanding of emotional

commitments. How do we carry 'the fair fight' off the hill, and into a distributed framework?

This special issue really represents but one of a family of 'special human issues' that greatly facilitate the process of congress. It is also here that I believe we are most likely to find alternate human strategies that are compatible with the technical infrastructure. There are many ways to signify that one is mad, for example. The ways one may use in daily person-to-person contact may be different from those employed over the Internet.

This brings us, in our analysis, full circle to the original issue of whether we should seek perfect mimicry of process or simply perfect mimicry of results.

Even though this has not been a discussion of particular technologies, I believe it is important to emphasize that we do know something about the characteristics of the technologies that will go into resolving all these issues. The most important characteristic is that the solutions will involve telecommunications and computing, and that the solutions cannot simply be software solutions. We know, for example, that any purely software solution will need some special component hardware to harden the security and protect member privacy.

We also know that the only people who can actually judge the efficacy of a geographically distributed emergency system are the participants in it. Even after we think we have the issues resolved, there should be the expectation of practice runs, which will refine the process and eliminate oversights and errors.

This, then, is my statement. I have tried to provide a simple method for enumerating the issues in developing a geographically distributed congress that can nevertheless function as a congress assembled as dictated by our constitution.

Robert Thibadeau, Ph.D.

Dr. Thibadeau is presently the Director of Security Architectures for Seagate Technologies, one of the largest disk drive and storage solutions providers in the world. He is on leave from his position as Director of the Internet Systems Laboratory in the School of Computer Science at Carnegie Mellon University. He has been at Carnegie Mellon since 1978 and is one of the founding directors of the Robotics Institute in the School of Computer Science. Since 1999, he has taught Computer Security. In addition, he was a principal architect of the security firmware products for Phoenix Technologies, now licensed through Verisign, Inc., for ubiquitous global distribution. He also has a strong interest in personal information privacy and digital rights management. He is an invited expert to the World Wide Web Consortium in privacy. The code base created at Carnegie Mellon for this purpose has become the basis for the European Privacy Demonstrator developed by the European Commission. He is a leading proponent of unified and hardened security architectures that also deliver strong protections for personal information privacy and digital rights management.