

The Evolution of Command and Control Systems

Baruch Fischhoff
Stephen Johnson
Decision Research
1201 Oak Street
Eugene, Oregon 97401
(503) 485-2400

Decision Research Report 86-7

April 1986

Abstract

Complex technical systems go through a series of stages in their evolution from a concept of how to meet a possible challenge to an operational version responding to real-world crises. The present analysis offers a characterization of these stages and the factors that shape the transitions between them. It can be used to describe the status of a system, to characterize or anticipate developmental difficulties, and to diagnose the sources of disagreements among those involved with it. It is illustrated here in the context of a specific system for which all stages must be addressed successfully, the National Command Authority for control of nuclear weapons in the U.S. or U.S.S.R.

THE EVOLUTION OF COMMAND AND CONTROL SYSTEMS

INTRODUCTION

The United States and the Union of Soviet Socialist Republics face a variety of similar military problems. Both superpowers have vast nuclear arsenals, giant military establishments, global security interests, treaty commitments, and sophisticated technical systems for controlling military power. Both countries also have systems for ensuring civilian political control over the military. This system is known as a National Command Authority (NCA). Its goal is being in control of important military decision making, especially during crisis situations. The most dangerous crisis situations are ones threatening or actually involving a nuclear attack. Preventing or controlling such crises is, therefore, a major focus of both NCAs.

Many writers have tried to describe these NCAs, with most of their writing aimed, primarily, at explaining how each is meant to work and, secondarily, at how it actually will work in crisis situations (Carter, 1985; Hemsley, 1982; Young, 1982; Zraket, 1984). One common concern in these latter analyses is that the NCAs will not work well enough to ensure constant civilian control over military power (Blair, 1985; Bubrow, 1977; Bracken, 1983; Ford, 1985; Tucker, 1983).

When potential inadequacies are identified, two responses are possible. One is to take the limitations as a fact of life and explicate their implications for the vulnerability of the country and for the stability of the international system that links it with other countries (e.g., can such failures lead to accidental war or to the unintended escalation of crises?). Alternatively, one can accept the

reality of current limitations, but not their inevitability. In this light, critical analysis of the present system is a stepping stone to proposals for improving and perfecting it.

In either case, recognizing limitations means doubting, at least implicitly, the validity of strategic policies that assume the existence of more smoothly functioning systems. Whether those policies are so unrealistic that they should be abandoned entirely depends, in large part, upon what one believes the opportunities for change to be. If change is feasible, then the challenge becomes to keep the system from being tested before it can be perfected. If change is unlikely, then there is a need for strategic policies that acknowledge reliance on imperfect command-and-control systems.

One obvious precondition for meaningful change is having an articulated plan for a better system, supported by a reasonable expectation of being able to remove the technical obstacles to making it operational (e.g., producing the electronic components, training the operators). However, change also requires generating the political support needed to mandate and finance an initiative, getting those responsible for the current system to accept change, and making the new system work in the hands (and minds) of the fallible humans who sit behind the command-and-control panels.

A comprehensive assessment of the chances for change requires a comprehensive conception of the obstacles to it. We offer here such a conception. Appearing schematically in Figure 1, it describes the stages through which a command-and-control system evolves, at each stage looking at the social institutions responsible for a system's

evolution and the factors influencing their activities. As elaborated below, these stages take a system from the recognition that some response is needed to cope with an external threat, through the shaping influence of national ideologies, partisan politics, and interservice rivalries, and on to a real-life system as it attempts to respond to an immediate challenge.

A wholly new system would be expected to traverse this path from the beginning, with the realities and pressures of each stage contributing to its eventual form. However, once a system is in place, pressure for additional change could come at any level, including changes in the external threat, new thinking on what constitutes the "perfect system," tinkering with the actual system, or improvisation when it is under stress. Having such an evolutionary model in mind can, we believe, help to produce a fuller description of existing systems, one which is sensitive to how they have fared in the evolutionary process, what developmental conflicts have yet to be resolved, and how close the actual system is to the intended one. Where this description (or life itself) shows the system to be unsatisfactory, such an analysis can suggest how to design a comprehensive strategy for introducing change and, by identifying obstacles, indicate the limits to change (and, hence, to system perfectability).

For example, introducing a new concept from the top offers the best chance of coherent, integrated change. Yet, each stage between that concept and the eventual product can force the changes in unintended directions or prevent any change at all. By contrast,

initiatives coming from people within the system should be better able to anticipate and overcome technical problems and human resistance. However, they are also more likely to disrupt any overall system "logic" which is only apparent from the top, while remaining captive of the ideologies and political compromises embodied in the current system.

An evolutionary model can also help clarify what people are talking about when they discuss command-and-control systems. Do they mean the system that they would like to have, the system that is promised, the best system that is likely to be produced, or the system that currently exists? Each level calls for a different degree of specificity and a different degree of realism. Knowing the intended level is essential to having appropriate expectations and interpretations (Bracken, 1980; Everett, 1982; Pate-Cornell & Neu, 1985).

There are perhaps several ways to describe the evolution of a system. Ours emphasizes the role of humans in it, whether as theoreticians, designers, supporters, critics, manufacturers, or operators. Certainly, the physical side of command-and-control systems is essential to their operation. They can do little without adequate telecommunications, remote sensing, ordnance, and so on. However, technology cannot do the job alone and may, indeed, complicate matters if its operators mistrust or misuse it (Cushman, 1983; Welch, 1982; Wohl, 1981). Gardenier (1976) coined the term "radar-assisted accidents" to describe the new class of mishaps that accompanied the introduction of radar to ships in inland waterways. Although radar

averted some accidents, it created others by disturbing established patterns for ship operation and communication. Relying on trial and error to reveal human factors problems is expensive at best and impossible for novel applications (which any real attack would be, to some extent). Nor can it deal with problems due to operators' confusion over their goals (or those of their opponents). Focusing on human factors should increase the chances of addressing these problems and avoid excessive attention to the tidier problems of technology (Sheridan, 1980).

NCAs are unique in their mission and in their details. Yet, they are also special cases of what are currently being called distributed decision-making systems, that is, ones in which the information and responsibility for decision making are distributed over individuals who are typically (although not necessarily) physically separated (Athans, 1982; Fischhoff & Johnson, in press; Tenney & Sandell, 1981). Such systems are common in both military and civilian arenas (e.g., multinational corporations, diplomatic services, fast-food franchises, networks of self-help groups). To varying degrees, advances in telecommunications have given them new capabilities and exposed them to new stresses, which are perhaps felt with particular acuteness in military applications (Coulam & Fischer, 1985; Fischhoff, Lanir, & Johnson, 1986). Thus, although it focuses on the evolution of NCAs, the present analysis could be extended by analogy to any command-and-control system, distributed decision-making system, or complex technological system influx. Johnson and Fischhoff (in preparation) narrow the focus to a comparative analysis of the US and USSR NCAs.

Whatever comparison is made, it is essential to know the level of description, in order to know the kind of reality to ascribe to it. The present analysis offers a general characterization of levels, applicable to any operational systems, but focused here on the two NCAs.

THEORY OVERVIEW

As conceptualized in Figure 1, a system evolves through six stages from a perception of the external reality that needs to be managed, to a theoretically derived ideal of the perfect response system, to a notion of the best possible system given current technological and economic capabilities, to plans for the best feasible system given the resources actually allocated to it, to the operational system produced by investing those resources, to the system that is mobilized in response to an actual challenge. Each transition between stages introduces a somewhat different set of constraints and pressures. Each brings somewhat different individuals and institutions into prominence, as the result of having pertinent power or expertise. Each is bound by the outcome of preceding stages; for example, the best engineering and training can do little to overcome an overall concept that is based on a misreading of the enemy or which is distorted to meet political purposes.

To some extent, performing well at any level requires an understanding of the realities at all other levels. For example, the theorists who develop the concept of a perfect system are primarily responsible for interpreting the external situation created by the

superpowers' capabilities and intentions in the light of the values expressed in sources like the US Constitution and latest Harris poll. However, they also need some idea of whether their country has the technology and personnel to implement the systems that they might propose. Conversely, the system's operators need some big picture if they are to interpret ambiguous situations and maintain their morale.

One goal of designing an orderly evolutionary process for system development is to limit how much personnel at each level need to know about the others in order to perform their tasks. If the process worked perfectly, then people in it would only need an understanding of the adjacent levels; these in turn would be responsible for comprehending the levels immediately above and below them. Thus, strategic theorists would not need more than a vague idea of what it is like to be in a bunker or an airborne command post with a finger on the button. As long as each pair of levels was suitably linked, then these theorists' general intentions would be reliably transmitted downward and faithfully interpreted in the context of actual situations. Conversely, pertinent features of the reality below will be transmitted upward, so as to prevent the theorists from making unreasonable demands or harboring unrealistic expectations.

An evolutionary process can fail if this transmission process goes awry, if those at any level fail to perform their tasks, if any level is omitted or short circuited, or if individuals whose expertise is appropriately exerted at one level attempt to influence other levels. Once the process begins to falter, then the problems at each stage cumulate to produce systems and policies that are not understood or

desired. A detailed look at the levels and transitions can serve as the basis for anticipating, preventing, and treating problems. After discussing the evolutionary process, we suggest some generic pitfalls and possible solutions.

THE STAGES

The Concept

At its most abstract level, an NCA begins from recognition of a reality that must be managed. That reality has an external component, the threat posed by the other side's current and potential political-military stances, and an internal component, one's own comparable stances. The resulting concept describes what needs to be done, bounded by some general awareness of how (and whether) it might be accomplished. Some of these goals and constraints may almost go without saying (e.g., ensuring maximum survivability of key leaders, along with an orderly transition in the event of casualties). Others, however, require reflection on a country's guiding values or prediction of what might reasonably be achieved. For example, the extreme time constraints of an actual attack might suggest centralizing command authority in the military, hoping that the responsible commanders will be sensitive to the opinions that civilian leaders would express were there time to consult them. However, the importance of civilian control over the military is so important to both countries that constant civilian involvement is incorporated in the basic concept of each's NCA, implicitly accepting a possible loss in efficiency. On the other hand, the same time constraints do not lead both countries to similar conclusions regarding the feasibility of having a flexible

response to the onset of an attack (Ball, 1983). The USSR's concept calls for restraining their forces until such point as all are released. By contrast, the US's concept calls for keeping its options open, allowing for varied responses to perceived USSR actions.

Like other aspects of a basic concept, the positions of the USA and USSR on these two features (civilian control and flexibility) reflect a complex mixture of culture, law, social history, and national ideology (e.g., Huntington, 1964; Leebaert, 1981). For example, US insistence on flexibility may reflect, in varying parts, an accurate appraisal of superior US technological capability, another expression of American faith in technical solutions, a psychological inability to accept the prospect of things quickly getting out of control were a nuclear conflict to begin, the legacy of President Kennedy's desire for better real-time control of nuclear forces after the Cuban missile crisis, or continuation of the relative autonomy traditionally given to US field commanders to improvise the best solution consistent with general orders (thereby emphasizing initiative and self-reliance). For present purposes, however, the sources of this belief are less important than its consequences. Which systems does it promote and which does it discourage? Which features of the international system does it highlight and obscure? What reality does it create for the USSR in designing its own NCA?

The Perfect System

The concept provides an assemblage of goals and constraints indicating what an NCA should and should not be. The next stage in system evolution translates these statements of what to do (or not to

do) into a general description of how to do it. These proposals should be more constrained by current technical capabilities, without relinquishing the possibility of forcing the development of new technologies. It should also be more realistic about economic limitations, but without relinquishing the hope of increasing budgets by creating sufficiently persuasive plans. These limits on realism make it important to distinguish descriptions of this level from those for subsequent ones.

The image of a perfect system might be seen in a briefing that showed how the US NCA is meant to operate, reaching down to how it should control naval task forces or unified and specified commands. It would show the consistency of the ideal with the American concept (e.g., Mearsheimer, 1983) of military autonomy within civilian guidance and periodic oversight, which implies a faith in the military to carry out its mission (as defined by civilian authorities). As a result, a design principle for perfect systems is minimizing civilian interference in routine military decisions, in the interests of efficiency. That means, in turn, delegating authority throughout the system, as well as the information needed to exercise that authority. Thus, the design for a perfect system would emphasize features such as how information and instructions flow, as well as checks for ensuring that it is, in fact, operating without constant supervision. However, the design must also allow for assertion of civilian control at any time that faith is lost, novel situations arise, or political conditions dictate.

The Soviet Union has quite a different history of civilian-military relations. Since the founding of the Soviet state, the military has been viewed as the ultimate source of political power. As a consequence, strict control over the military is seen as essential to maintaining political power. In the Soviet perfect system, a General Staff controls day-to-day operations, whereas a Supreme High Command, chaired by the (civilian) General Secretary of the Communist party, in his role as Supreme Commander in Chief, controls all crisis situations (Scott & Scott, 1983). Although this arrangement would resemble the US perfect system in some respects, it would also have greater concentration of power (and less delegation of authority), so as to facilitate monitoring in routine situations. As a result, ensuring civilian control in crisis situations should require less transfer of power.

An orderly transition to this stage from its predecessor means deriving general design guidelines from the very general principles embodied in the concept. Doing so may mean confronting uncomfortable conflicts between ostensibly inviolate principles. What happens, for example, when civilian control comes at the expense of survivability? If the system theorists (at this level) make these tradeoffs differently than would the social theorists (at the concept level), then the system will embody unresolved conflicts in its objectives. Executing the perfect-system stage explicitly increases the chances of recognizing and addressing such problems, insofar as it is hard to discern the design philosophies of systems at the later, more concrete levels. Having an explicit image of the perfect system also provides

clearer directives for those subsequent stages.

The Best Possible System

The next set of constraints is set by scientific, industrial, and economic base of the country building the system. It determines what could be done were those resources turned single-mindedly to this particular task. Although the system at this level is still an abstraction, it is now one requiring detailed design work, sensitive to the capabilities and limitations of the people and machines available to do the job.

As before, both technical and social obstacles threaten adequate completion of this stage. The main technical threats are bad science, in the sense of applying current knowledge ineffectively, and bad science policy, in the sense of misestimating what knowledge can be developed and applied in the time allotted. The result might be either a bad plan or an impractical one. The main political threats are failing to generate the resources needed to realize the plan, or being forced to make compromises that disrupt its technical coherence or its fidelity to the concepts underlying the perfect system. In facing the conflicting pressures of faithfulness to that concept and promising adequate performance at a reasonable price, the system plan becomes the uncomfortable meeting ground for the conflict between national myths and objective constraints.

Current American efforts to develop a defense against ballistic missiles provide a good example of a system in transition from perfect system to the best possible system. The Strategic Defense Initiative (SDI) takes its inspiration from the perception that the situation

might arise in which the USSR would launch a ballistic missile attack against the US, against which no defense is currently available. The national value articulated by President Reagan as a response to this perception is that Americans respond to threats by using military technology to provide needed protection. One perfect system embodying this concept would be a shield protecting the entire United States with laser weapons and the like. As evidenced by its political success, this concept appears in harmony with popular ideology, culture, and history. SDI is now in the transition to a best-possible system, reflecting the limits of science, economics, and industrial capacity. During this phase, scientists and engineers are examining the limits of current (and hoped for) technologies to see what can actually be accomplished and what it will cost. The evolving system concept should offer the best possible configuration within reasonable (but optimistic) budgetary limits.

The resulting proposals could fail politically if they were perceived as expressing a view of superpower relations that is inconsistent with currently accepted notions of detente and mutual coexistence. As of this writing, however, a more imminent threat is inability to make a plausible case for SDI's technical operability. That is, can we create the science to bring it off? A productive national debate would ask such questions, as well as how development of a perfect SDI system would influence international relations during its construction period and what would it mean were an imperfect system to be the ultimate result.

The Best Feasible System

Before a system confronts the harsh reality of international affairs, the plans for it confront the harsh reality of national politics. To proceed, the plans must recruit the support of the political leaders (including legislators, publicists, etc.) who must appropriate the funds for its development, the scientists and engineers who must create the knowledge needed for its operation, and the military figures who must ensure its adoption. Each of these groups may have different interpretations of national goals, against which to compare the evolving system concept. Each is likely to have its own vested interests. Thus, the plan will be scrutinized by elected officials for how voters would view support for the system, by the military for how it would affect service capabilities, missions, and interservice balances, by contractors for how production would affect their balance sheet, by special interest and citizen groups for how well it fits their goals. In American terms, the emerging best feasible system will be "what comes out of committee."

To planners, these pressures may seem like needless complications, useful, at best, for getting the system's message "out to the broader public." There are, however, a number of constructive roles that such scrutiny can play. One is to check that the theoreticians and technical people have, in fact, produced a concept consistent with contemporary national values. A second is to cast a lay eye on the realism of the entire project. Non-technicians may have a relatively good feeling for how complex human systems actually work, as well as for the realism of the promises made by system proponents. A third

role is forcing the plan to defend itself against diverse critiques, each hungry for weaknesses that are inimical to its own best interests. A plan that weathers the attacks of the best technical experts that these interests can secure should be the better for it.

Perhaps more of a risk than ignorance is lack of coherence in the pressures exerted at this level. If the political and technical critiques are unbalanced (e.g., due to the dominance of a particular vendor or armed service), then the design may be confused or distorted. It is easier for critics to prevent bad designs than to promote good ones. Their aggregate effect should depend, in part, on the health and balance of the overall political process (e.g., Coulam, 1977).

The Actual System

Once resources have been allocated, the technical specialists must make good their promises, producing an operational system within the given budgetary and conceptual constraints. Doing so requires realizing the potential solutions to technical problems, interpreting with greater specificity the tradeoffs implied by general societal values, and managing an ever-expanding cast of individuals having or wanting a piece of the action (including government bureaucrats, corporate executives, construction personnel, and potential operators). It requires addressing any incoherence introduced by the budgetary process' allocation of funds and assignment of responsibility. It requires maintaining a delicate balance in communications with the outside world, simultaneously assuring critics and supporters that the work is going well, while acknowledging enough problems to explain delays and justify requests for additional funds. This balancing act

may be particularly difficult when a project needs some hyperbole to get funded at all or where it requires unproven technology (which can only be explored once funds are allocated). The Divad and Roland air defense systems are two recent examples of projects where the actual system failed to fulfill promises made during the best-feasible stage (Easterbrook, 1982). Whether their development, nonetheless, represented reasonable gambles when undertaken is a question for retrospective technology assessment (Tarr, 1976).

When a system stumbles, its proponents may have promised too much or they may have missed an attainable goal by mismanaging the construction and implementation stage. For example, unacceptable costs may reflect unrealistic initial estimates or uncontrolled defense contracting (e.g., profit pyramiding, cost-plus pricing). If the battle for adequate funding is lost, then the result may be purchasing only part of the system (threatening its coherence), purchasing fewer spare components than intended (threatening its reliability), or stretching out the development and production process (threatening the other systems and policies that depend on having an operational system).

However developed their hardware, systems are but abstractions until they work with their actual operators behind the controls. Historically, "human factors" issues have been faced rather late in systems development, well after most fundamental design decisions have been made. The result is more pressure on the operators to adapt to the machines than vice versa. If they cannot adapt, then the system's

performance may be disappointing or unpredictable (Brewer & Bracken, 1984).

For a new system to work as intended, its operators must individually overcome any inappropriate habits and expectations acquired during their experience with other, older systems. Collectively, they must coordinate their perceptions into a "shared model" of the new system and its environment, in order to perform effectively and cooperatively. With a system of any complexity, perceptions will likely be different for those approaching it from different angles. For example, in a hierarchically organized system, those at the "top" are likely to have an encompassing view that is sparse regarding details, whereas those at the "bottom" are likely to have detailed local knowledge without a complete understanding of the overall problem. Although formal training may try to inculcate the designers' view of the system, the operators' own hands-on experience is likely to create a shared model of the system with a life of its own, creating unexpected problems and solutions. For example, the robustness of NORAD reportedly owes much to its operators' ingenuity in diagnosing the sources of minor malfunctions (U.S. Government, 1981).

A significant part of any actual system is its command and reward structure. Any innovation must either accept the current command structure (e.g., regarding the autonomy of local commanders) or include a revision of that structure in its implementation plan. A system that had succeeded in all other respects could easily founder on institutional resistance in this regard.

The Mobilized System

The final version of a system emerges when it is put to the test, which for an NCA would mean a national emergency. The mobilized system that emerges in response to such tests is but a special case of the actual system, whose crisis performance can be predicted somewhat from observation of normal operation and responses to simulated crises. However, the actual transition to crisis footing must remain as something of an unknown, with the degree of unpredictability depending on the uniqueness of the crisis. For repetitive "low-grade" crises, such as hijackings, false alarms, and minor military engagements, some learning from actual experience is possible. For the most extreme event, nuclear war, no wholly realistic training or testing is possible.

Some unpredictability is also due to the incompletely understood physical effects accompanying this ultimate test. Although there is some theory regarding electromagnetic pulse, thermal pulse, shock waves, and other communications disturbances, their effect on NCAs is largely a matter of speculation (Ball, 1981). At least as speculative are predictions of how the system will "rewire" itself after physical destruction of units or links between them (Bracken, 1983; Tucker, 1983). How it could rewire itself depends on physical circumstances and technological capabilities. How well any rewiring works will depend on how accurately operators can interpret totally unique patterns of information and communication. That is, can they figure out who is still out there and what they know? Finally, there are non-cognitive psychological questions, regarding how operators will respond

emotionally to crisis situations, after doing their best to comprehend the facts about them. One trains and hopes for calm, analytic appraisals and responses, yet those can never be entirely assured.

DISCUSSION

Complications of the Scheme

Figure 1 describes an idealized process for the development of a novel system. Clearly, life is often more complex. One set of complications comes from the existing system that the new one replaces. Unless replacement is complete and instantaneous, the two systems must be compatible to avoid chaos during the transition. That constrains how innovative the new technology can be. Ignoring these constraints risks clashes between the two generations of software and hardware, and between operators' muddled mental models of them (Cushman, 1983; Lewis, 1983).

Analogous problems arise when a system changes during the development and implementation process, juxtaposing different versions of the system. Such changes in design may follow from advances in technology (perhaps stimulated by the system itself), from changes in how the external world is perceived, or from resistance by the organization receiving the system. All of these pressures are likely contingencies with complex systems having wide ramifications and long lead times. Experience with the system, as a concept, plan, proposal, or reality, will produce further pressure for change, introducing feedback channels from each stage to those preceding it.

Utilization of Scheme

Used statically, a scheme such as that in Figure 1 would

characterize the stage that a particular system description is meant to capture. More dynamically, it would follow, or perhaps even predict, the evolutionary process, showing its origin and future, as well as how its reality deviates from its ideal. Either statically or dynamically, the scheme can contrast how individuals at different levels perceive the system. Each transition can cause perceptions to diverge, with each discrepancy between perceptions having different practical importance. For example, differences between perceptions at the top and bottom levels would show where operators need education about the overriding national mission they are meant to execute or where theoreticians need some education about how the real world works. When discrepancies emerge between adjacent stages, remedial steps more explicit than general education may be undertaken. For example, it is possible, in principle, to reform those aspects of the military procurement process that make the actual system differ from the best feasible one. Conversely, the design procedures leading to characterization of the best feasible system could, again in principle, be brought into line with the realities of maintaining technical expertise (and morale) at defense contractors, as well as with the extra costs and seeming inefficiencies that that concession to corporate realities entails.

A scheme could like Figure 1 can also facilitate interpreting the meaning, at each level, of changes that occur in the system. Consider, for example, the recent redeployment of the National Emergency Airborne Command Post from Edwards Air Force Base to a location in the Midwest. As a belated and reluctant recognition of the impossibility of

evacuating the President in time of acute crisis, the move may mean relatively little for the actual operation of the US NCA. However, it would symbolically mean a great deal for the underlying perfect system, with its ideal of civilian control of the military.

Choosing a Level of Description

The purpose of a description should dictate its level. For example, the initial public debate over a system with implications for national policy debate should not focus on questions of technical feasibility. Where that happens, fundamental policy issues may get lost in a welter of technical details, increasing the risk of massive investment in systems that are incompatible with society's values. Conversely, allowing theoreticians to comment on a system's technical details threatens its operational viability, without much chance of achieving their social and political goals.

The present evolutionary scheme is moot regarding the viability of specific command-and-control systems. Rather, it makes the global prediction that an orderly evolutionary process increases the chances of system success. From that follows the prescription that everything should be done to ensure that the various relevant parties fulfill their designated roles, according to the scheme.

Where the process is out of balance, then an evolutionary description can help clarify what kind of system has been created. For example, knowing the limits to an NCA's operability may vitiate or validate some strategic doctrines, as can knowing that the system works, but in ways incompatible with the perfect system's intent. There are several points of leverage for bringing a system into line

with hopes and intentions. One might attempt to sharpen the perceptions of those participating in the development process, to refine the balance among competing groups, to refine the procedures used at each level, or to alter the evolutionary process (e.g., by creating feedback loops between the stages). Where flaws remain, the ability to diagnose them may allow planners to prepare for problems. For example, the interaction between the NCA's of the two superpowers might be stabilized by exchanging information about how to avoid situations likely to produce false positives.

To be meaningful, the description of a command-and-control system must specify the intended level of system evolution. That allows one to distinguish realities from promises, to judge informants' competence to make pronouncements at that level, and to anticipate at least some miscommunication. A system is unlikely to function well, if those concerned with it do not understand the limits of their own and others' understanding.

References

- Athans, M. (1982, September). The expert team of experts approach to command and control (C²) organizations. Control Systems, pp. 30-38.
- Ball, D. (1981). Can nuclear war be controlled? (Adelphi Paper No. 169). London: International Institute for Strategic Studies.
- Ball, D. (1983). Soviet strategic planning and the control of nuclear war (Reference paper no. 109). Canberra: Australian National University, The Strategic and Defence Studies Centre.
- Blair, B. G. (1985). Strategic command and control: Redefining the nuclear threat. Washington, DC: The Brookings Institution.
- Bubrow, D. D. (1977). Communications, command and control: The nerves of intervention. In P. Stern (Ed.), The limits of military intervention. Cabin John, MD: Seven Locks Press.
- Bracken, P. (1980, April). Command and control for a long war. Air Force Magazine, pp. 50-52.
- Bracken, P. (1983). Command and control of nuclear forces. New Haven, CT: Yale University Press.
- Brewer, G. D., & Bracken, P. (1984). Some missing pieces of the C³I puzzle. Journal of Conflict Resolution, 28, 451-469.
- Carter, A. B. (1985). The command and control of nuclear war. Scientific American, 252(1), 32-39.
- Coulam, R. F. (1977). Illusions of choice: The F-111 and the problems of weapons acquisition reform. Princeton, NJ: Princeton University Press.

- Coulam, R. F., & Fisher, G. W. (1985). Problems of command and control in major European war. In R. F. Coulam & R. A. Smith (Eds.), Advances in information processing in organizations: Vol. 2 Greenwich, CT: JAI Press.
- Cushman, J. H. (1983). Command and control of theater forces: Adequacy. Washington, DC: AFCEA International Press.
- Easterbrook, G. (1982, October). Divad. Atlantic Monthly, pp. 29-39.
- Everett, R. R. (1982, January). Yesterday, today and tomorrow in command, control and communications. Technology Review, pp. 66-68.
- Fischhoff, B., & Johnson, S. (In press). The possibility of distributed decision making. In R. Weissinger-Baylon & J. March (Eds.), Political-military decision making. New York: Pitman.
- Fischhoff, B., Lanir, Z., & Johnson, S. (1986). Military risk taking and modern C³I (DR Report No. 86-2). Eugene, OR: Decision Research.
- Ford, D. (1985). A reporter at large: The button (parts I and II). The New Yorker, pp. 43-91 (April 1), 49-92 (April 8).
- Gardenier, J. S. (1976, April). Toward a science of marine safety. In Symposium on Marine Traffic Systems. Symposium conducted at The Hague: The Netherlands.
- Hemsley, J. (1982). Soviet troop control: The role of command technology in the Soviet military system. New York: Brassey's Publishers, Ltd.
- Huntington, S. P. (1964). The soldier and the state. New York: Vintage.

- Johnson, S., & Fischhoff, B. (In preparation). National command authorities as distributed decision making systems. Eugene, OR: Decision Research.
- Leebaert, D., (Ed). (1981). Soviet military thinking. Boston: George Allen and Unwin.
- Lewis, K. H. (1983). On the appropriate use of technology. Orbis, 27, 274-285.
- Mearsheimer, J. J. (1983). The military reform movement: A critical assessment. Orbis, 27, 285-300.
- Pate-Cornell, M. E., & Neu, J. E. (1985). Warning systems and defense policy: A reliability model for the command and control of U.S. nuclear forces. Risk Analysis, 5, 121-138.
- Scott, H. F., & Scott, W. F. (1983). The soviet control structure: Capabilities for wartime survival (Strategy paper no. 39). New York: National Strategy Information Center, Inc.
- Sheridan, T. B. (1980). Human error in nuclear power plants. Technology Review, 82(4), 23-33.
- Tarr, J. (1976). Retrospective technology assessment.
- Tenney, R. R., & Sandell, N. R. (19?) Strategies for distributed decision making. IEEE Transactions on Systems, Man, and Cybernetics, SMC-11, 527-538.
- Tucker, J. B. (1983). Strategic command and control vulnerabilities: Dangers and remedies. Orbis, 26, 941-263.
- U.S. Government. (1981). Failures of the North American aerospace defense command's (NORAD) attack warning system. Washington, DC:

U.S. Government Printing Office, Committee on Government
Operations.

Wohl, J. G. (1981). Force management decision requirements for Air
Force tactical command and control. IEEE Transaction on Systems,
Man, and Cybernetics, SMC-11, 618-639.

Young, R. E. (1982, June). Crisis control. Wireless World, pp. 36-
39.

Zraket, C. A. (1984, June 22). Strategic command, control,
communication and intelligence. Science, pp. 1306-1311.

Evolutionary Theory of Command and Control Systems

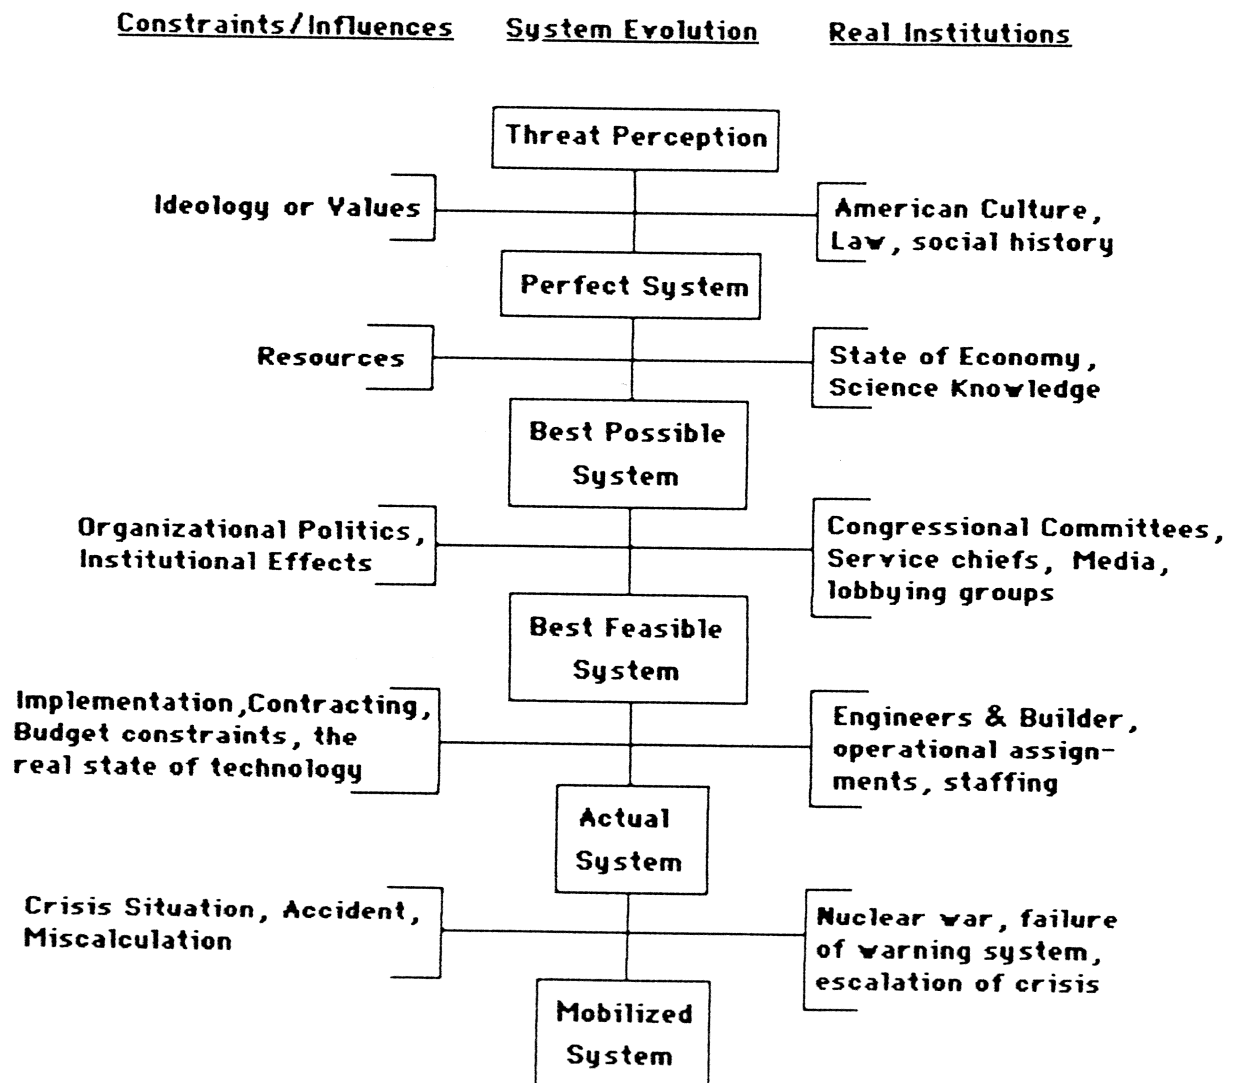


Figure 1