

CAL
EA725
86R03

DOCS

Number 3

Report

CANADIAN INSTITUTE
FOR INTERNATIONAL
PEACE AND SECURITY

The Risk of Accidental Nuclear War

A Conference Report
by
Andrea Demchuk

Proceedings of the Conference
on the
Risk of Accidental Nuclear War

Vancouver, 26-30 May 1986

The Canadian Institute for International Peace and Security was established by Parliament on 15 August, 1984. It is the purpose of the Institute to increase knowledge and understanding of the issues relating to international peace and security from a Canadian perspective, with particular emphasis on arms control, disarmament, defence and conflict resolution.

Part of this mandate is to encourage public discussion of the issues. The Institute, therefore, participates in and also sponsors conferences on related topics and may publish reports on their proceedings.

The views contained in this report do not necessarily represent those of the Institute.

Information on the Institute and its publications may be obtained by
writing to

The Canadian Institute for International Peace and Security
307 Gilmour Street,
Ottawa, Ontario K2P 0P7

Publié également en français

Preface

The Act of Parliament which created the Institute stated that part of its purpose was to encourage public discussion of issues of international peace and security. One of the ways in which the Institute seeks to achieve this objective is by playing an active role in conferences which address these subjects. It is interested in attending such conferences, in helping to sponsor them financially, and, on occasion, in organizing them and preparing reports of their proceedings.

This, the third such report, summarizes the proceedings of the conference on the Risk of Accidental Nuclear War which was held at the University of British Columbia in May 1986. The meeting, which raised considerable interest, was attended by specialists from a wide range of disciplines and from several countries including the United States and the Soviet Union.

As with the first two reports in this series, "Negotiations for Peace in Central America", and "Challenges to Deterrence", the problems raised in this account are of great interest to many Canadians. We hope, therefore, that this will be of use not only to those who attended the conference but to a much wider public which is concerned about the dangers which the conference discussed.

Geoffrey Pearson

Geoffrey Pearson
Executive Director

Dept. of External Affairs
Min. des Affaires extérieures

JUN 30 1995

RETURN TO DEPARTMENTAL LIBRARY
RETOURNER A LA BIBLIOTHEQUE DU MINISTERE

43 273 004

TABLE OF CONTENTS

Executive Summary	1
Introduction	3
The Risk of Accidental Nuclear War	5
A Conference Report	11
by	
Andrea Demchuk	17
II. Automating Command and Control and Risks	17
III. Nuclear Command and Control and Performance	17
IV. Behaviour in Crisis: Acting to Reduce the Risk	21
V. The Human Factor in Accidental War	26
VI. Conclusions and Recommendations	34
Appendix: List of Speakers and Participants	37

**Proceedings of the Conference
on the
Risk of Accidental Nuclear War
Vancouver, 26-30 May 1986**

EXECUTIVE SUMMARY

TABLE OF CONTENTS

Executive Summary	1
Introduction	3
I. Quantitative Assessment of the Risk	5
II. Automating Command Systems: Problems and Risks	11
III. Nuclear Command and Control: Goals and Performances	17
IV. Behaviour in Crises: Acting to Reduce the Risk	21
V. The Human Factor in Accidental War	26
VI. Conclusions and Recommendations	34
Appendix: List of Speakers and Participants	37

EXECUTIVE SUMMARY

Recently public attention has been directed towards the risk of accidental nuclear war. The two superpowers have been discussing ways of reducing the risk of war by miscalculation, and members of the US Congress have put forward bills promoting the establishment of joint crisis management centres. A number of articles in newspapers and magazines have asserted that the risk of war by accident is great and is increasing year by year.

Researchers are trying to analyze and to quantify this danger. What is the magnitude of the risk? Is it growing? What factors contribute to increasing the risk and how can it be reduced? A number of methods have been applied to the study, including modelling of alert/reaction systems, and statistical analyses of past international crises. This conference brought together representatives from different disciplines, using different research tools, so that they might exchange data, develop some form of consensus about the severity of the risk and suggest ways of reducing it. Their discussion covered five subjects: modelling, computer programming, command-and-control systems, crisis behaviour, and the role of human error. The mathematical models presented varied in their underlying assumptions and in the methods used for manipulating key parameters, but there was general agreement on the findings. The modellers present argued that the risk of an accidental nuclear launch due to unresolved false warning was increasing, especially with the deployment of weapons systems very near the borders or coastlines of the opposing power; short delivery time meant very little time to determine whether an alert was false. It was recommended that the nuclear powers agree to withdraw those weapons systems which have driven down decision times.

Short warning times are also related to concerns about the fallibility of computer programs and the vulnerability of command and control systems. Reduced decision times provide an incentive towards the automation of the alert/launch response. The vulnerability of command centres means that the military planners put high priority on getting the missiles out of the silos quickly while the communications channels are still functioning. It was recommended that C³I systems be made more "survivable" and that planners and policy-makers resist any pressure to move to an automated launch-on-warning posture in times of crisis.

The last two days of the conference examined the effects of human behaviour and human error on the risk of inadvertent nuclear war. Such a war would occur not because either side thought it could gain from launching a nuclear attack but because the crisis had inexorably escalated as a result of misperception and bad judgement, aggravated by a lack of diplomatic skill and military restraint.

INTRODUCTION

An international conference on the Risk of Accidental Nuclear War was held on 26-30 May 1986, in the Chapel of the Vancouver School of Theology at the University of British Columbia. The conference was sponsored by Science for Peace, the Vancouver branch of the United Nations Association in Canada, and the Peace Research Committee of the International Political Science Association. Major funding for the conference was provided by the Canadian Institute for International Peace and Security, the Disarmament Fund of the Department of External Affairs, the Gordon Foundation, and Science for Peace. Additional support and assistance was provided by the deans of Arts and Science of the University of British Columbia, the Centre for Continuing Education, the Vancouver School of Theology, and the Vancouver branch of the United Nations Association in Canada. Twenty-four speakers addressed the conference; a list of their names, current work and institutional affiliations is appended to this report. In addition, thirty-two academics, professionals, and representatives of peace groups registered for the entire week, while many others participated in one or more of the public sessions.

The Background and Rationale of the Conference

For a number of years, there has been growing concern in academic and policy-making circles alike, about the possibility that a strategic nuclear exchange might be initiated inadvertently or accidentally, or as a result of the mistaken belief that an attack from the other side was imminent or already underway. The reasons for this anxiety are threefold.

First, a growing number of military strategists have begun to worry lest the deployment of new strategic and theatre weapons systems, together with consequent changes in strategic doctrine and operations policies, prove highly destabilizing. The short flight-times of the new intermediate-range nuclear forces (INF) and forward-based submarine-launched ballistic missiles (SLBMs) render both sides' command and control vulnerable to "decapitation" and thus reduce the decision time available during crises. In response to these developments both sides may have adopted strategic postures and operating policies which amount to a virtual launch-on-warning* in time of crisis.

A second source of anxiety, about which computer scientists are expressing concern, is the growing tendency to automate decision-making within nuclear command and control systems. This process of automation, which has been developed in response to shortened warning times, carries with it the twin dangers of an increased number of errors and failures within command and control systems, and a decreased ability to check and rectify these errors while maintaining operational control.

* The definition of launch-on-warning is currently the subject of considerable controversy.

Several public interest groups have drawn attention to a third area of concern, by making public hitherto classified data on errors and false alarms within the US strategic warning system. The data they have produced suggest that false alarms (even quite serious ones) are a good deal more common than generally believed, and have apparently increased in frequency. Even more alarming, when the particular events leading to serious instances of false alarms are examined in detail, is the evidence that increasing automation of command and control systems gives rise to a greater number and a wider range of failures and false alarms.

In response to these trends, academics from a wide variety of disciplinary backgrounds and national origins have begun to study the magnitude of the risk that nuclear war might occur by accident. Is this risk (as the official view would have it) still infinitesimal, or has it reached an alarming level as a result of the factors mentioned above? The conclusion seems to be that in all probability the risk has increased substantially, but we lack the data, methodology and conceptual tools to assess the danger precisely.

In view of this situation there was felt to be a need for an international, interdisciplinary conference to exchange information and insights on the risk of accidental nuclear war. A committee was therefore struck to organize this, comprising representatives from the sponsoring organizations, and academics engaged in research in the field.

I.

QUANTITATIVE ASSESSMENT OF THE RISK

The first day's deliberations were devoted to the presentation and assessment of several mathematical models of the risk of accidental war, and to a general discussion of the utility of a quantitative analysis of risk in this context. Presentations were made by Ms. Barbara Leonard, Dr. Linn Sennott, Dr. Brian Crissey, and Dr. Daniel Frei; the commentator was Dr. Anatol Rapoport.

Barbara Leonard's paper contained a mathematical model of accidental nuclear war which she and Bill Rosenberg had developed. The model's equations simulate what would happen if a *de facto* "launch-on-warning" policy were in effect during a crisis alert. The paper took as its starting point the assumption that, if an unresolvable alarm were to occur when the system was in a launch-on-warning state, a war would be triggered.

According to the model the probability of an unresolvable alarm occurring during a crisis is the product of the interaction of four parameters: (1) the number of days per year that the system is in *de facto* launch-on-warning; (2) the length of time available to make a decision; (3) the number of false alarms per year; and (4) the time required to resolve a false alarm. Using unclassified American data to estimate these parameters, Leonard and Rosenberg made five key predictions. First, on the basis of a small number of crisis days per year, a 7-minute decision time, and the existing rate of increase in false alarms, there is a high probability of accidental nuclear war occurring within the next decade. Second, if the average time allowed for the resolution of a false alarm is no more than 3 minutes, there is a 95 per cent chance of accidental nuclear war within 6 years. Third, if the average resolution time is 2 minutes and there are as few as 5 crisis days per year, there is still a 50 per cent chance of accidental nuclear war within 11 years. Fourth, even if it is assumed that an unresolved alarm will not trigger a launch without a second confirming signal, there is still a 50 per cent probability of accidental nuclear war within 16 years. Their final and most crucial finding is that the probability of accidental nuclear war is dramatically affected by the amount of decision time available. With only one minute of decision time, there is a risk of accidental nuclear war occurring in less than a year. With a three-minute decision time, it is likely to occur within 3 years and with a four-minute decision time within 14 years.

The paper presented by Linn Sennott outlined two models, one of overlapping alarms and another of dual phenomenology. Both of these models build upon an earlier one developed by Wallace, Crissey and Sennott, which was published in the *Journal of Peace Research*. Each incorporates salient characteristics of NORAD's Early Warning System.

The overlapping alarms model is based on the assumption that a missile would be launched if a second false alarm went off before the previous alarm had been resolved. The model estimates the probability that this juxtaposition of alarms would occur, during a given time period, on the basis of three parameters:

- a) the average frequency of false alarms;
- b) the average time taken to resolve each alarm;
- c) the time interval under consideration.

The model shows that the average length of time until a lethal juxtaposition of false alarms occurred would be inversely proportional to the square of the frequency of false alarms. That is, if the number of false alarms per unit time doubles, the average amount of time before an overlapping alarm occurs is cut by a factor of four. Assuming a 3.5-minute false-alarm resolution time, 100 false alarms per year would give a 6.4 per cent chance of an overlapping alarm during that year. A doubling of alarms to 200 per year would increase this probability to 23.4 per cent and tripling the number to 300 per year would increase the probability of overlapping alarms to 45.1 per cent. (This figure is considerably greater than that provided by the US Department of Defence.)

Sennott's dual phenomenology model examines the military's key assertion that false alarms are not as dangerous as they might appear because of the policy of dual phenomenology, which requires that any indication of attack by one family of sensors, such as infra-red sensing satellites, must be confirmed by another family, such as radars. Sennott's model evaluates the claim that this form of redundancy would drastically reduce the likelihood of a "false positive" detection of incoming missiles. It shows that if "each stream" (satellites, radars) has 200 false alarms per year, the average time until an alarm occurs simultaneously in both systems is less than four years. (It is appropriate to note here that Bruce Blair asserted later in the conference that the "dual phenomenology" doctrine sometimes uses what he calls "strategic warning," that is, independent information from political or intelligence sources that an attack appears to be imminent.)

Sennott concluded that command, control and communications systems cannot be made completely secure by technological means and that detection errors *cannot* be eliminated. In attempting to eliminate as many detection errors as possible the military sensors must try to strike a balance between Type 1 errors, that is, failure to detect an incoming missile or missiles, and Type 2 errors, that is, detecting a non-present target. Sennott argued that the proportion of Type 2 errors will increase as decision time shortens. Most troubling to Sennott was her conclusion that:

We are reaching a situation of contradiction, namely, the time available for human intervention in the decision-making process

is shrinking to zero, yet having time for human intervention is absolutely imperative if we are to avoid catastrophic error – this problem doesn't have a technological fix.

Daniel Frei's paper provided a critical examination of the existing literature on accidental nuclear war. Rather than construct an empirically derived theory of the risk of accidental nuclear war, he identified some unexamined issues. One of the main issues with which he dealt was the probability of accidental nuclear war, probability being defined as the combined result of a number of risk factors. Some of these risk factors, would be independent; that is, the probability of their occurring simultaneously is very much lower than their occurring individually. Examples of a pair of independent risks, which could in tandem lead to accidental nuclear war, would be a faulty computer chip setting off a false alarm and a mentally ill submarine commander misinterpreting the data. On the other hand, some of the risk factors leading to accidental nuclear war would be interdependent since the failure of one factor may cause the failure of another. An example of interdependent risks would be an unintentional missile launch by the United States leading the Soviets to conclude that they were under a massive attack. The more interdependent risks there are in the system, the higher the probability of an accidental nuclear war. The provision of safeguards is intended to create redundancies within the nuclear-weapons system, and thus make interdependent risks independent. Among such measures, Frei lists double-key systems, permissive action links and detonation locks.

Dr. Frei pointed out, however, that in real world systems risks are neither fully interdependent or independent. Thus far the literature has failed to identify the level of interdependence between the risk factors leading to accidental nuclear war. Frei argued that in order to achieve credibility with decision-makers the literature must identify:

. . . precisely which risk factors on the level of weapons technology and command and control systems are affected by each other and the nature of an acute international crisis as compared to situations of 'normalcy', and in what ways are these causal interrelations structured.

In addition to emphasizing the need for the clear identification of interdependence among risks, Frei also warned that risks cannot be assessed in a political vacuum, since the current nuclear dilemma is very much the product of East-West confrontation:

The differences in values separating East and West should not be neglected nor must it lead to premature and superficial identifications of common interests. Today's situation is much more complex. Also the fact that both East and West wish to avoid a nuclear war does not necessarily imply that they share identical ideas about how to control and reduce this risk.

Political considerations had made Frei cautious about the prospect of measures to prevent accidental nuclear war. He argued that some unilateral measures, such as the West's removal of Pershing II's from Europe, or the adoption of a 'no-first-use' policy, would exchange the risk of accidental nuclear war for other risks, such as the Finlandization of Europe or the break-up of US security guarantees to Europe. In addition, such unilateral measures might

. . . invite the Soviet leadership to push forward again by a policy of *fait-accompli*, as it has done by deploying the SS-20 missiles, in order to create "leverage" and produce a "bargaining chip".

Frei seemed to treat bilateral or co-operative measures with even more caution. He pointed to the very real cultural differences between East and West, which might tragically impede successful co-operation. Making use of a classification developed by Joseph Nye, Frei envisioned three types of co-operative measures: (i) crisis management, (ii) crisis prevention, and (iii) long-run stabilization. Frei judged categories (i) and (ii) to be highly workable, but suspected that ". . . the goal of long-run stabilization may very probably already go beyond the confines of US-Soviet cultural community." He suggested that rather than abandon the goal of long-term stability, the West should develop a variety of measures to prevent or to manage crises with the deliberate intention of placing them in a "stabilization framework."

Brian Crissey's paper, which was accompanied by a live demonstration of the computer simulation, developed mathematical accidental nuclear war models by linking them to models of the arms race. This shows that the probability of accidental nuclear war "evolves" over time as the arms race proceeds. The model is built on the assumptions that growth in technological complexity, in the amount of space debris and in the number of weapons is directly linked to accidental nuclear war: that the probability of superpower crises is constant over time; and that the strategic window of decision time remains constant at 8-10 minutes. Even though models of this sort are not designed to produce reliable values, but rather to allow the researcher to "play" with alternative assumptions and parameters, two results of Crissey's simulation are of interest. First, the model "predicts" a sharp increase in the probability of accidental nuclear war in the early 1980s and a slower increase thereafter. Second, the model predicts that the probability of an accidental *Soviet* launch is far greater than an American one.

Anatol Rapoport provided a critique of Leonard, Sennott and Frei's papers. He emphasized that in analysing accidental nuclear war, equal attention should be paid to probabilities and what he termed, "utilities." The purpose of risk assessment is to allow decision-makers to make informed choices between alternatives. This is straightforward only if numbers can be assigned to the different alternatives in order to reflect their relative degrees of desirability or undesirability. Utilities are determined by compounding the values of possible (foreseen) outcomes of our

choices and the probabilities assigned to the conditions which lead to those outcomes. Thus, the utilities assigned to alternatives in risky situations can be statistically calculated, if positive or negative numerical values are assigned to the possible outcomes of such action according to their relative desirability or undesirability.

Rapoport did not share Frei's view that probabilities represent the most crucial dimension of the accidental nuclear war risk problem. He believed, rather, that utilities are of equal importance and that both must be numerically expressed if their product, risk, is to be defined. Risk must be defined numerically if decision-makers are to make unambivalent choices, since rational choice requires that the alternatives can be rated statistically.

The probability of a non-repeatable event, such as nuclear war, can only be estimated. It cannot be precisely defined. The probability of *accidental* nuclear war can, however, be more readily assessed because it can be assumed that the occurrence of accidental nuclear war can be related to the occurrence of other events which may recur and whose frequency can be observed. Rapoport noted that Leonard and Rosenberg had related the occurrence of accidental nuclear war to that of other repeatable events, such as Missile Display Conferences, Threat Assessment Conferences and frequencies of crises. As probabilities can be assigned to these events on the basis of their occurrence in a given span of time, and since an unfortunate coincidence of such events could trigger nuclear war, a subjective but meaningful probability can be assigned to accidental nuclear war.

Utilities, as defined by Rapoport to embody desirability, are by their very nature a reflection of personal values. Thus, even more than with probabilities, any assessment of them is inevitably subjective. Yet if action is to be taken on the basis of risk assessment, utilities must be determined. Rational choice of action in a risky situation entails comparison of the expected utilities corresponding to the various courses of action. Taking no action is itself considered to be an action; only Sennott's paper made this point. Her discussion of Type I and II errors in missile detection showed that Type I error is to disregard a real attack and the other is to respond to a false alarm as if it were a real attack.

Rapoport pointed out that Type I and II errors are always complementary, as there is a trade-off between risks. One risk can be made smaller only at the expense of making the other larger; reducing the risk of disregarding a real attack will increase the chances of responding to a false alarm. In dealing with the risk of accidental nuclear war Rapoport used the analogy of capital punishment. The abolition of capital punishment ensures that every accused person's life will be spared, including those of brutal murderers, in order to eliminate the chance of executing even one innocent person. The irreversibility of capital punishment is important to those who wish to have it abolished, since because of this they assign an infinite negative utility to executing an innocent person. The

risk of miscarriage of justice thus becomes unacceptable, no matter how low the probability of its occurrence, since the product of the finite probability multiplied by the infinite negative utility yields an infinitely negative product.

When this reasoning is applied to accidental nuclear war it follows that no matter how small the probability of such a war, the negative utility of wrongfully executing millions of innocent people makes the risk unacceptable. Rapoport argued, therefore, that nuclear weapons should be abolished on the same basis as capital punishment.

To the abolitionist, the cost of accidental nuclear war is always infinite and the benefits are always zero. However, an assessment of the probability of accidental nuclear war may still serve the abolitionist cause and there is even a possibility that the danger of accidental nuclear war may provide common ground between decision-makers and abolitionists. Rapoport suggested that military planners do not like the idea of accidental war because if war occurs, they want it to happen in the way and at the time which they have designated.

In further comment on the papers Rapoport said that abolitionists can draw attention to the way in which various technologies and policies affect the probability of accidental nuclear war. Leonard and Rosenberg's paper had shown how a launch-on-warning policy magnifies the risk of accidental nuclear war, while Sennott's paper had linked overlapping false alarms to accidental war. Frei had pointed out the difference between the independent and interdependent risk factors in the way they contribute to the risk of accidental nuclear war. While adding independent redundant safeguards to the nuclear system may decrease the probability of nuclear war, the growing complexity of the nuclear system and the consequent unforeseen interdependence of its components may increase the risk.

According to Rapoport the abolitionist would do well to stress the two-sided nature of every risk. Using the analogy of business, where every investment entails a risk but every non-investment entails an opportunity lost, he argued that military planners emphasize the risks of dealing with the Russians but neglect to mention the opportunities lost by not dealing with them:

What if, just *what if* Mr. Gorbachev's proposals for carrying out the abolitionist's programme are *not* a bluff. If they are not, what is the opportunity loss associated with dismissing them? If they are a bluff, what is the risk associated with calling it?

II.

AUTOMATING COMMAND SYSTEMS: PROBLEMS AND RISKS

The second day's deliberations dealt with the dangers posed by the increasing use of computers in the nuclear command and control process, and with their inherent inability to provide an "error-free" command system. Presentations were given by Dr. Severo Ornstein and Dr. Henry Thompson. The commentator was Dr. Joseph Weizenbaum.

Severo Ornstein's paper gave a detailed account of the limitations of computers in all their applications, including military decision-making. He pointed out that, since computers process most information in today's weapons systems, computing has taken over from physics as ". . . the technology at the cutting edge of the arms race." He suggested that computing software presents ". . . the most serious and least tractable problems in using computers in critical applications", and asserted that there is no way to completely rectify those problems or to ensure that "residual problems will not have catastrophic consequences."

As far as Ornstein was concerned, people had been effectively eliminated from military decision-making in time of crisis. He believed that people couldn't think fast enough to cope with the current level of "decision density" and that

. . . if decision-making is based only on simplified summaries presented by computers, then the ingredients necessary for informed judgement are absent and you might as well let the computers make the decision.

He also warned his audience that computers, while extremely fast at some types of tasks, were slow at others, especially pattern recognition. Central to Ornstein's argument was his description of how computers work. Basically computers execute software which may be defined as ". . . successive simple instructions which, when put together in meaningful sequences, can perform complex tasks." A computer's greatest power is not to move quickly but to make choices based on designated external or internal conditions while the program is running. Computers both react to and act upon the outside world. Most computers do not compute but process non-numeric information; they apply *ad hoc* rules more often than they implement equations. Thus, most computer programs are not founded on mathematically-precise models, but on other less formal models, the validity of which it is difficult to test.

As yet computers are capable of limited performance. According to Ornstein they do not have the ability to think for themselves, or to go beyond prescribed rules. Although some computers have been given a few rudimentary skills, their progress, when measured against a "human

yardstick”, has not been impressive. Researchers have discovered that the human thought processes are much more complex than originally surmised. Humans make use of a wide range of knowledge in functions such as pattern-recognition, generalization, learning, and strategy development. This knowledge is so wide-ranging that it has still to be fully catalogued by researchers, and so computers cannot yet come close to human performance in this area.

Another limitation is that when a computer breaks down, more often than not it becomes totally useless:

The notion that it will merely be a little sluggish or inaccurate derives from experience with other kinds of devices where a worn wheel or axle may indeed degrade performance. When something goes wrong with a computer, however, it tends to go berserk, not degrade.

Ornstein challenged the assertion that, even if Star Wars computers broke down, 95 per cent of incoming missiles would still be intercepted. He argued that in fact none would be intercepted in the event of a breakdown. Some systems are designed to prevent this sort of total breakdown. In these systems the effects of errors are somewhat alleviated by isolating various parts of the program and providing a minimum of restricted interconnections between the parts. Unfortunately, even these modular systems cannot be completely protected from breakdown. Furthermore, computers are notoriously unpredictable in interpreting unanticipated data. For example, sensors have interpreted the moon rising as an incoming missile. An additional limitation is the personality of the human operator. Although humans are an integral part of any operating system, human reactions are frequently ignored by systems designers and superseded by military protocol.

Ornstein maintained that the most crucial limitation of computers, however, arises from the inherent imperfectability of the software. Most of a computer's design is contained in the software and when a computer breaks down this is frequently the source of the problem. While computers are almost infinitely flexible because of their software, they are also easy to program incorrectly. Superficially it is easy to fix software problems, but the more complex a system becomes the more difficult and expensive it is to deal with these. In fact, a whole technique of “software maintenance” has arisen to deal with problems as they occur. John Shore, the author of *The Sachertorte Algorithm*, has argued that if a car needed as much attention from the manufacturer as computers do, then the car would probably be called a lemon!

In order to explain the fundamental problems that software imposes on systems, Ornstein described the two basic steps of software construction. First, the researcher studies a problem in the real world, decides which aspects of it are relevant and devises rules, then formal specifications, to govern the behaviour of the computer. Second, the researcher takes these

specifications and writes a program in accordance with them. In large systems there will be some interaction between these two steps but they are essentially separate procedures.

Software errors are not always the result of carelessness (for example, a missing comma). Frequently, errors cannot be discovered until after a program has been implemented:

The “unforeseen events” that cause trouble are less often unforeseen external events than perverse concatenations of perfectly normal events.

Complex interaction between the various parts of a program can occasionally produce wholly unexpected, erratic results. There are too many possibilities for interaction in modern computer systems for it to be feasible to test all of them, so numerous techniques have been developed to find errors.

Two classes of errors occur in writing programs from specifications. The first is the “typo” – A Mariner space probe was lost because a period was put where a comma should have been. The second, the “thinko,” involves minor errors in reasoning. Most of the latter result either from the fact that the specifications themselves contain ambiguities or lapses in reasoning, or from the problem that English, with all its ambiguities, has become the language most often used for specifications.

In theory, it should be possible to check a program against its specifications, but the techniques of “program verification” are not as accurate as the name would suggest. Ironically, both Greg Nelson and David Parnas, two top software engineers in this area, have become vocal opponents of the Strategic Defense Initiative because they believe it is impossible to build trustworthy software. Although program verification is an important and powerful tool, it is not a foolproof solution. In large systems the full program is never run, since full-scale testing can be embarrassing or truly impracticable. Consequently many possibilities are never investigated. Smart designers and programmers are the best antidote for bugs, but they cannot provide a complete solution. Techniques of simulation also help to identify errors. Here person A writes the specifications for the original program and then person B writes specifications for a program designed to imitate the environment with which the original program is to deal. This method is of great help in identifying errors but it, too, is affected by shared oversights and misconceptions. Having identified errors, fixing them, especially in large systems in which they are deeply imbedded, can be a difficult process and may itself create further problems. For example, the Advanced Research Projects Agency (ARPA) computer network needed a major redesign after it grew beyond a certain size.

To cope with the persistence of errors, humans have developed “fault-tolerance” techniques, based on the idea of building spare parts into a

system in order to allow it to keep functioning in the event of a single component breakdown. In computers, however, “fault-tolerance” is more easily implemented in hardware than in software. If a faulty program is replaced by an identical copy, the same problem will eventually arise again.

To meet this difficulty the human programmer is duplicated, and separate teams write separate programs from the same specifications, on the theory that,”

... Various teams are unlikely to introduce exactly the same errors and that by comparing the answers and doing software voting (the computers compare their own answers) you can eliminate errors.”

This technique of software “fault-tolerance” involves four major difficulties. First, computers cannot always compare “answers” since much information is carried along implicitly rather than explicitly. Second, the teams are not truly independent, as they are using the same, inevitably incomplete, specifications and therefore may very well build in the same wrong assumptions. Third, it is difficult to synchronize machines in order to compare “answers”. Fourth, specifications are often plainly incorrect and this technique does not address that.

Ornstein considered specifications to be “the least tractable part” of the software problem. Very often people don’t know what they want computers to do, and so cannot define the problem for them to solve. We can’t say to computers “Protect me!” We must tell them precisely how to do it.

In problem domains there is an excess of detail, some of which it is easy to overlook. Attempts to deal with this, by having independent teams write several programs, founder on the difficulty of making any objective comparison of the results. According to Ornstein, computers are extensions of human intellect and are therefore as limited as their originator, if not more so:

The point is that we must not expect that computers will safely solve the problem of the nuclear hair-trigger for us. They won't. They give the appearance of relieving human burden, but at the cost of increased chances of catastrophic blunder.

Henry Thompson outlined his concerns regarding automatic decision-making systems and accidental war. He noted that as yet it is human beings who would decide to use nuclear weapons, but there have recently been some suggestions that this should be changed. The proponents of automatic launch-on-warning systems maintain that such systems would increase security through improved reliability and greater credibility. Thompson argued against the implementation of these systems on technological grounds. Automatic decision-making is subject to the present

limitations of artificial intelligence, and these are likely to hinder progress in the design of automatic systems until well into the next century.

Artificial intelligence requires a two-step process: the representation of a domain of knowledge in computational form, followed by the design and implementation of algorithms to solve problems. Thompson sees the way in which decision-making is based on knowledge as imposing a crucial limitation on the further development of artificial intelligence.

Basically he sees knowledge as lying on a continuum from specialist to causal. Data at the specialist, structural, compositional or evidential end of the spectrum has been successfully utilised by artificial intelligence. Expert systems in word processing, mineral geology, bacterial infection diagnosis and computer installation are examples of artificial intelligence applications using specialist knowledge. In each of these systems, knowledge is represented in system specifications by "if-then" rules, and is manipulated by a variety of mechanized formal inference procedures. The other end of the spectrum, however, is general, functional, contextual, causal, situated or "common sense" knowledge, which entails understanding of use or significance. As yet the only decision-making system which uses general knowledge is the human being; artificial intelligence methodology has been unable to do so. Researchers do not know how to encapsulate the world in the machine. Of course there is no absolute division between the two forms of knowledge, but the further artificial intelligence systems diverge from the specialist end of the spectrum, the less successful are the results.

The artificial intelligence systems described thus far can be used for decision-making only in an advisory capacity. Since they utilise only the narrow specialist type of knowledge, they would not be good candidates for automatic launch-on-warning. Nonetheless, proposals to use them in that way are currently under consideration. Thompson was concerned that automatic launch-on-warning would empower inadequate machines to make the most crucial decisions facing humanity. He believed that what was essentially a political problem should not be delegated to technology. Flawed as humans are, they are far more capable of making sensible decisions than the best machine.

Joseph Weizenbaum suggested in his presentation that Ornstein and Thompson, if anything, probably erred on the side of overestimating the reliability of computers. He also pointed out that the unreliability of computers could cause unpredictable consequences; the recent failure of the computer system of a small US bank had had disastrous effects on the precious metals market.

He observed that Ornstein and Thompson had failed to note that large computer systems are incomprehensible in their entirety, and that no amount of study, simulation, or missile conferencing can overcome this. Larger computers are not so much deliberately designed as evolved, and

if their history is lost so is our ability to fully comprehend their logical nuances. This contributes to the impossibility of understanding the final workings.

Weizenbaum argued that artificial intelligence cannot make decisions which are acceptable to human beings. Instead of providing a solution to computer fallibility, as some claim, it would probably only compound many of the difficulties of computer systems. He emphasized that grandiose predictions concerning artificial intelligence often flout well-recognized principles. Even the advisory artificial intelligence systems described by Henry Thompson should be approached with caution, since human operators are wary of overriding the computer's results and tend to accept its judgements as automatically correct. The reason for this is that an employee whose decision to override a computer leads to a problem is often fired, whereas an employee who follows the advice of the computer has a convenient scapegoat in the event of error. Weizenbaum closed by pointing out that while technology has had many impressive achievements, it has also had an infinite number of unforeseen consequences; any proposal to make nuclear weapons systems more automated should be resisted on the grounds that it might have unanticipated catastrophic consequences.

III.

NUCLEAR COMMAND AND CONTROL: GOALS AND PERFORMANCES

The third day of the conference turned to what was perhaps the most central issue in an assessment of accidental war risk: the nuclear command systems themselves. There was a discussion of what should be required of nuclear command systems, and of the degree to which the existing systems of both superpowers met these standards. Presentations were given by Dr. Bruce Blair, General Mikhail Milstein, and Mr. Marco Carnovale; the commentator was Dr. Douglas Ross.

Bruce Blair's paper outlined the extent to which current US procedures for dealing with crisis alerts increased the risk of accidental nuclear war. He began his presentation with the assertion that the superpowers had overemphasized a cardinal principle of crisis stability, threat, at the expense of another complementary principle, reassurance, in designing procedures to reduce the risk of deliberate or accidental nuclear war. The ways in which threat increases the risk of accidental nuclear war are vividly illustrated by operational procedures for a crisis.

In peacetime, Blair argued, negative control measures act as safeguards to prevent the unauthorized or accidental use of nuclear weapons. In time of crisis, however, the emphasis would clearly switch from negative to positive control; the military would be more concerned with co-ordinating their forces and implementing their war plans accurately, than with maintaining safeguards against the accidental or unauthorized use of nuclear weapons.

The co-ordination of forces by a legitimate authority would be difficult due to the variety of weapons systems within the nuclear arsenal. Furthermore, the command system would be so vulnerable that positive control would be difficult to re-establish after a nuclear attack. Because of this vulnerability central authority would inevitably be weakened. National policy officials control the terms of the alert only in a legal sense; there is too much detail to be handled by central authorities and, in practice, alert authority resides as low as the level of commander. Because the commander is responsible for the safety of his troops, he is allowed to take those alert measures which he deems necessary or prudent.

Compounding the authorities' problem of dealing with a virtually incomprehensible amount of detail is the difficulty of maintaining communications with the field, once alert procedures have been put into effect. In order to avoid enemy detection, there are progressively stricter rules against radio transmission the higher the level of crisis alert. Blair noted the irony of this situation in which the higher the level of crisis, the more concerned national officials would be about operational interactions, but the less they would be able to control these interactions. Central authorities would become insulated from the realities of the field. However

well combat units followed procedures, they would probably not be able to strike a good balance between negative and positive control. In all likelihood, heightened alerts would translate into escalating mutual suspicions; alert measures would be met with countermeasures, thus creating a cycle of reinforcing alerts and suspicions.

The superpowers have exacerbated the inadequacies of central control by the adoption of *de facto* launch-on-warning. Launch-on-warning appeals to both sides because it appears to compensate for the vulnerabilities of their command systems, but Blair argued that this policy is at the heart of the problem of accidental nuclear war.

Launch-on-warning entails a rapid shift in priorities measured in seconds. The NORAD Commander must decide whether to prepare for, or prevent, a launch. He must base this judgement on dual phenomenology, that is, a combination of strategic warning indicators (classic intelligence sources), tactical warning indicators (from sensors such as satellite, infrared and ground radar) and confirmation by human operators of the data provided by the sensors. Too much must happen in too short a time for there not to be a high risk of mistake and so dual phenomenology has not eliminated the possibility of error.

Furthermore, those who would have to decide whether or not to retaliate and if so against which target would be under great stress. The few minutes allowed for a decision and the scant information available would not provide a clear picture of the attack. There would be no room for political, moral or even military reasoning, "and in a drill-like atmosphere, the risk of inadvertent war due to false alarm, misperception, or miscalculation can only be heightened." Blair would eliminate the perceived need for the hair trigger by designing and implementing a command system which would survive and could be reconstituted after an attack. Such a system would bolster deterrence more ably than would the ability to fire quickly. While the creation of a survivable command would not be cheap, it is feasible.

Douglas Ross commented on Blair's paper and on his 1984 book, *Strategic Command and Control*. As far as the paper was concerned, he concurred with Blair's emphasis on the trade-off between maintaining negative and positive control and with his major conclusion that the risk of accidental war would be reduced by devising and deploying survivable command systems.

Ross believed, however, that Blair's assertion that the superpowers had emphasized threat at the expense of reassurance was an overstatement. He noted that this assertion contradicted Blair's doubt that the vulnerability of its command system would allow the United States to respond to an attack. The American capacity for "overkill" on a first strike does not guarantee it the capacity for retaliation. In addition to command system considerations, the US "threat" capacity had strategic shortcomings.

Ross suggested that Blair had not addressed the beliefs of many strategists. First was their concern that the Soviets want to develop the ability to make a pre-emptive strike in time of crisis. If the Soviet Union were attempting to achieve this strategists would not accept Blair's aim of a secure second-strike. Second, many conservatives wanted to have a secure counterforce capability. Blair's prescription would not allow for this objective. While Ross did not espouse these views himself, he believed that they must be dealt with in any effective political debate.

Ross also noted that having vulnerable command systems undermined the stability of the relationship between the superpowers. Furthermore, he felt that there was a lack of political awareness of how urgent this problem had become.

General Milstein described Soviet strategic policy. He noted that the Soviet Union had declared a policy of no-first-use and maintained that it did not entertain the possibility of a limited nuclear war. He rejected the view of conservative strategists, as described by Ross, that the Soviet Union was planning to achieve the capacity for launching a pre-emptive strike and argued that if military targets were attacked, civilians would also be affected. Milstein asserted that the deterioration in the relationship between the United States and the Soviet Union was extremely dangerous from a nuclear standpoint. He believed that addressing the political problems was a necessary pre-requisite to solving the military and technical problems of nuclear war. He also insisted that security could not be one-sided and that for one side to have nuclear superiority was destabilizing. He cited Gorbachev's proposal to remove nuclear weapons, in three stages, by the year 2000.

Marco Carnovale drew attention to the differences between the intercontinental strategic situation and that in Europe. He pointed out that there were several irreversible reasons why the European command system was uncontrollable, and he gave a European perspective on what he considered the shortcomings of Bruce Blair's argument for a survivable, secure, command system.

According to Carnovale the European nuclear strategic situation differs from the intercontinental situation in eight essential ways. First, the European defence system has several dual capable systems, in which most weapons can be used for both short-range and medium-range purposes, and dual-key systems, where it is difficult to determine which authority has the final say over launching, the host country or the United States. Second, command and control in Europe must inevitably be more decentralized than at the intercontinental level. Third, the level of complexity of European arsenals is necessarily high. Fourth, Europe's geographic proximity to the Soviet Union requires short reaction times. Fifth, the political and military goals of the allies frequently diverge, even in peacetime. Sixth, the European weapons systems are characterized by what Paul Bracken referred to as "uncontrollability", in his 1983 book *The Command and Control of Nuclear Forces*. While Bracken had argued that this

ambiguity afforded the system some deterrent value, Carnovale did not think that this was a good basis for deterrence. Seventh, the European system has a high degree of operational complexity because a multitude of actors play roles in its day-to-day management. The European defence system is also constantly evolving, and even if there were total agreement on objectives it would still be difficult to keep constant track of all aspects of the system. Eighth, for geographical reasons, European weapons are mainly concentrated in a small number of sites. Even when mobile, they are highly vulnerable to Soviet targetting.

He believed that there were four reasons why the European command and control system was basically uncontrollable. First, nuclear matters are difficult for experts, let alone relatively transient politicians, to understand, and this is compounded by the fact that European politicians have a high rate of turnover: Italy, for example, has had forty prime ministers since World War Two. Second, it would be impossible to foresee all of the contingencies which European command systems might have to face. Given their vulnerability and short reaction times, much would have to be left to spontaneously-devised solutions. Third, the collective control of forces by NATO creates problems. If each NATO member has a trigger, then no one finger is on the safety catch. Conversely, if even one country has a safety catch, then no one really controls the trigger. There is no middle way. Fourth, strategists must decide at what level to attack the Soviet command system. Current NATO policy does not call for attacking the Kremlin, but for attacking lower levels of Soviet command. This would cause a problem, for there is no perfect level at which to target the Soviet command system that would both impair military capability and still leave a leadership with which to negotiate. Administrative, physical and informational controls had been in place in Europe for forty years and had evolved to high levels of sophistication, but nonetheless the aforementioned problems still remained intractable.

While Bruce Blair had argued for a safe, survivable, flexible command system, with tight negative and positive control, Carnovale argued that, from a European perspective, this solution would not be satisfactory. He agreed that such a system was desirable in peacetime, but suggested that it was not so attractive during a crisis. He was concerned lest survivable command made nuclear war thinkable from an American perspective, because it would enable the United States to survive conventional or limited nuclear war in Europe.

Carnovale proposed as an alternative that NATO should publicly announce a threshold which, if it were passed by the Soviets, would release a countervalue strike. He advocated automating the system and eliminating conventional weapons. These measures, he argued, would put in place a system which would have catastrophic consequences if it were to fail, but which would be extremely unlikely to do so.

IV.

BEHAVIOUR IN CRISES: ACTING TO REDUCE THE RISK

Speakers on the fourth day of the conference examined the various strategies used by national actors in conflict and crisis situations, and attempted to assess which of these were most likely to lead to war, whether by inadvertence or miscalculation. A major focus of discussion was the extent to which the problem of accidental nuclear war could be meaningfully distinguished from the problem of war in general. There were presentations by Dr. Russell Leng, Dr. Martin Hellman, and Dr. Johan Niezing; the commentators were Dr. John Barrett and Mr. John Lamb.

Russell Leng's paper summarized a series of five studies on bargaining strategies between states in times of crisis. The way leaders behave in a crisis is salient to any discussion of accidental nuclear war, since many experts have postulated that accidental war would most likely occur during a crisis. These studies suggest that *realpolitik* considerations, or more simply, concern for power, prestige and a national reputation for resolve, seem to be the chief factors motivating leaders during crises. The findings also indicate that when national leaders take these *realpolitik* considerations to their logical extreme, and ignore similar motivation on the part of their adversaries, this may result in the undesired outcome, war. Leng was disturbed by the finding that leaders may respond more to the realist prescription to show resolve than to the equally important prescription to act with prudence.

Leng was both encouraged and concerned by how the superpowers had behaved in past crises. Both sides had been able to exercise prudence and restraint during the two Berlin crises. During the Cuban Missile Crisis, the United States at first employed a purely coercive bargaining strategy against the Soviets and then later precipitated the end of the crisis by applying a carrot-and-stick strategy. Soviet responses to the initial coercive tactics revealed the danger of utilising brinkmanship in crisis. Leng added that coercive tactics had become even riskier because of the shrinking gap between US and Soviet capabilities and in light of the finding that the loser in one dispute is likely to behave more belligerently in the next dispute with the same opponent. The findings suggested that world leaders are motivated both by a rational calculation of strategy and by factors such as pride and personal status. These latter considerations may explain why statesmen react strongly to overt threats from states of comparable power and why a loser endeavours to regain face.

Two observations combine to give cause for alarm, especially in confrontations between nuclear powers: first, leaders tend to show resolve rather than prudence; second, such an unrestrained demonstration of resolve can result in escalation of a crisis. In conflicts between nuclear powers both sides are somewhat restrained by their awareness of the dangers of brinkmanship but, paradoxically, each side is also aware of the restraints that the nuclear reality imposes upon the other. Thus, nuclear

adversaries may be tempted to “test” each other by using coercive tactics. This leads Leng to conclude that:

Because these are also likely to be conflicts where there is a good deal of symmetry in usable military capabilities and, perhaps most important, motivation, such temptations could also too easily lead to a disastrous miscalculation, an accidental nuclear war caused by human error.

Johan Niezing’s paper was concerned primarily with the methodology of determining risk. It centred on the argument that risk analysis could be a useful tool in clarifying the basic shortcomings of nuclear strategy, especially in the assessment of accidental nuclear war. Niezing defined risk as the product of chance and damage, and maintained that investigation of both these factors would enable peace researchers to argue more persuasively. He asserted that while risk analysis cannot allot precise values in complex areas such as accidental nuclear war, it can identify some tendencies which increase or decrease the element of risk and thus it can counteract intuitive arguments which unduly minimize chance. From this point of view, studies of the structure of accidental nuclear war are a necessary complement to traditional nuclear strategy. Such studies should focus on the lack of redundancy in command and control systems and on the increasing inadequacy of international procedures to prevent accidental nuclear war. Nuclear strategy would benefit from an assessment of the probable damage of nuclear war in general, and particularly from studies of indirect or non-military effects, such as nuclear winter.

Niezing pointed out that risk analysis of accidental nuclear war was likely to evoke opposition in the form of silence or scepticism. Researchers should anticipate such opposition and seek to address it. Niezing asserted that risk analysis must analyze not only risk itself, but also the pattern of *risk acceptance*. Some people are prepared to accept a high risk of nuclear war because they perceive that it offers “benefits.” How these benefits are perceived will frequently be based on intuition or ideological assumptions. It is imperative that arguments against such views should not themselves be based on ideological assumptions, but should rather confine themselves to pointing out inconsistencies in the analysis and showing the ways in which such perceptions are arrived at. He suggested that the psychology of cognition might be relevant to such a study.

Cognitive theory shows that people block out information which they do not wish to take in, by various strategies ranging from ignoring to totally re-interpreting unpleasant information. “Hard” or “non-social” information, especially isolated facts, tend to be ignored, and complex sets of hard facts are trivialized by splitting them into isolated facts. “Soft” or “social” information can easily be transformed into less-threatening messages. These denials would not be so disturbing were it not for the fact that a failure to appreciate the importance of chance may cause great damage.

Risk assessments must be repeated frequently in order to uncover existing patterns of risk-acceptance. In doing this, attention should be paid to the ignoring of hard, non-social information; the fragmentation of complex situations of risk; the misleading reinterpretation of soft, social information; and the addition of new knowledge to the base of "hard" and "soft" information. Many individuals, especially policy-makers, have virtually ignored facts that contradict the "basis of nuclear strategy." For example, the electro-magnetic pulse from high-altitude nuclear detonations would probably cripple current US command, control and communications procedures during attack, and the phenomenon of nuclear winter could render irrelevant the concept of "limited" nuclear war.


In summary, risk analysis will not produce a definitive assessment of the risk of accidental nuclear war, but it will allow peace researchers to give some estimate of its likelihood and the damage it would cause. Such an assessment may have a persuasive effect on those who would otherwise deny the magnitude of the risk.

Martin Hellman's paper modelled the inevitability of nuclear war under current circumstances. Central to his argument was the concept of stability:

Stability has always been a critical consideration in the design of systems ranging from aircraft to national economies. But nowhere is stability more important than in the design of national defense systems in the nuclear age. Proponents of the MX and other counterforce weapons systems often justify their proposals by noting that deterrence in the form of Mutual Assured Destruction is unstable: even a small-scale Soviet attack would require an American response which would lead to all out nuclear war. Similarly, opponents of these same weapons systems point to a crisis instability engendered by a "use them or lose them" mentality.

Hellman noted that each argument had some merit but that the apparent paradox could be resolved only by differentiating between short- and long-term stability. Hellman introduced three figures representing physical models of stability:

Figure (I): Stable, unstable, and metastable systems.



a) Stable System

b) Unstable System

c) Metastable System

Figure Ia is used by physicists to denote a stable system. There is only one stable state of the system, which is modelled by the location of the ball at the bottom of the “potential well.” By contrast, figure Ib denotes an unstable situation. That is, the ball will only stay on the top of the hill if it is perfectly centred and completely unperturbed.

The only figure of the three which adequately models the current stability of the world’s nuclear system is the metastable figure Ic. “In the short-term deterrence is stable but in the long term it will surely fail. On our current path, nuclear war is inevitable.”

Hellman used a probabilistic model to show that the “. . . many random events constantly perturbing the state of the world: coups, civil wars, natural disasters, regional wars, misinterpretations, C³I false alarms, etc.” would cumulatively lead to nuclear war. He drew an analogy between nuclear strategy and officers’ roulette. On the first round your chances of surviving are good, but if you continue to fire the trigger round after round your chances of not hitting the chamber with the bullet are negligible.

He used the Cuban Missile Crisis as an example of how a single decision could destabilize the world’s nuclear system. President Kennedy’s advisors had recommended that the US military make a “surgical strike” to remove Soviet missiles from the Western Hemisphere, but later it was concluded by these same advisors that, far from correcting the problem, the strike would have led to a catastrophic world war. Such a strike would have been enough to dislodge a metastable ball from its perch. In closing, Hellman argued that the only way to move from a metastable to a stable world situation would be to abandon war because, by definition, a nuclear world could not be stable if war remained an accepted part of international relations.

John Lamb provided a critique of both Hellman’s and Leng’s papers. He began by noting that the conference had frequently returned to the proposition that the greatest risk of accidental nuclear war was posed by crises. The purpose of the present session, which was considering Leng, Niezing, and Hellman’s papers, was to go further and relate crisis behaviour to the prospect of accidental nuclear war. Lamb disagreed with Hellman’s analogy of nuclear strategy to officers’ roulette. He argued that one could not ascribe equal weight to each disturbance in the state of the world, which was what he thought Hellman had done. Crises between nations are not analogous to turns in officers’ roulette which in each case has the same set of circumstances and the same probability of leading to a catastrophe. He did not believe that the danger caused by crises was cumulative but thought that crises might be cyclical. Lamb did not believe that nuclear war was inevitable, but he did agree that the world situation was urgent and that new attitudes and institutions were needed to circumvent man’s apparent need to resort to force in solving disagreements.

Lamb stated that Leng's paper had addressed the issue of crisis behaviour most directly, but he suggested that none of the papers had defined precisely the sort of crisis which would most likely lead to accidental nuclear war. Did the greater risk lie in a US-Soviet crisis in Europe or in a US surrogate-Soviet surrogate crisis? Were geographically limited crises inherently less dangerous than widespread ones? In what ways did particular kinds of crises interact with the contemporary command system?

Lamb's major conclusion was that crisis control and crisis prevention, as well as technical measures to enhance command systems and nuclear disarmament, should be dealt with in any study of accidental war. In this area smaller powers could also have a role. Talks on arms transfers should be rejuvenated, and talks on geographical hot spots where a superpower confrontation might arise should be regularized. Finally, Lamb thought it was an attractive idea to make crisis simulation available to top decision-makers.

John Barrett's presentation dealt with Leng and Niezing's papers. He began by admitting his interest in arms control and policy analysis as opposed to events data analysis, of the sort carried out by Leng. He wondered whether findings from historical data could be usefully extrapolated to the present and if US-Soviet nuclear parity would affect Leng's findings. He was particularly interested in Leng's finding that the use of threats as a bargaining strategy by comparably powerful states would result in war.

Barrett found Niezing's focus on perceptions to be fruitful. He thought there might possibly be differences in the cognitive apparatus employed by those holding different views of contemporary nuclear reality. The whole theory of deterrence is based on analyses of risk and cost/benefit, and these should be explicitly explored.

Barrett also appreciated Niezing's focus on perceptions, especially when applied toward presenting advice for disarmament in a palatable form. Here the symbolic as well as the military importance of weapons systems were significant. Arms control experts could offer more credible alternatives if they succeeded in exposing the adverse psychological effects of certain weapons systems quite apart from any questions about their military value. Barrett suggested that workers in the field of arms control would benefit from having many of the academic works presented at the conference translated into more practical terms.

V.

THE HUMAN FACTOR IN ACCIDENTAL WAR

The fifth and final day of the conference was devoted to a discussion of the role played by human error in creating a risk of accidental war. A central theme was the difficulty experienced by military and political institutions in weeding out substandard performers from positions of nuclear responsibility. A second theme was the problem of leadership performance in cases of extended deterrence.

Presentations were given by Dr. Bruce Russett and Mr. Paul Huth, Dr. Lloyd Dumas, Dr. Dean Babst, and Drs. Luc de Seguin and Michel Haag. The commentators were General Mikhail Milstein and Dr. Russell Leng.

Lloyd Dumas argued in his paper that human operators were an integral part of the nuclear military system and that they behaved in specific ways which introduced an element of unpredictability.

Individual components of the system can be tested fairly readily, but as they are combined into complex systems it becomes less and less feasible to test them under realistic operating conditions. We cannot start a nuclear war in order to make predictions about the system's reliability.

Dumas pointed out that the behavioural sciences are notoriously weak in their ability to predict; he used his own field, economics, as an example. Economics is supposed to be the study of the most rationally-driven of all human activities, yet economists have difficulty in predicting the behaviour of the economy. Thus, disciplines such as psychology and sociology, which deal with emotional as well as rational behaviour, are even less likely to make firm predictions. Yet it is they, rather than economics, which are most relevant to the behaviour of the human operators of nuclear weapons. No behavioural science has yet approached the natural sciences or engineering in its ability to make predictions.

Dumas outlined several aspects of the problem of human reliability in the nuclear military. He showed that alcoholism, drug addiction and mental illness were present among military personnel, described how working conditions could adversely affect the behaviour of employees, and gave details of the way in which hierarchy and bureaucracy could reduce human reliability. While he used mainly US examples, since Soviet examples were rarely available, he suggested that there was every reason to believe that Soviet problems were as bad, if not worse.

Because alcoholism, drug addiction and mental illness are widespread in the population as a whole, it is impossible to completely screen these problems out of a large organization such as the military. A 1982 US Department of Defense worldwide survey of drug and alcohol abuse among US military personnel estimated that:

- 1) 14% of military personnel were heavy drinkers of alcohol; 9% were alcohol dependent.
- 2) 18% of all military personnel experienced one or more "serious consequences of alcohol use" within the past 12 months (these include work impairment, physical damage, need for detoxification, etc.)
- 3) 42% of DoD personnel had used one or more drugs for nonmedical purposes.
- 4) 9% had used drugs other than marijuana in the past 30 days.

Applying these percentages to the total number of US armed forces personnel led to the conclusion that in 1982, 295,000 personnel were heavy drinkers, 190,000 alcohol dependent and approximately 190,000 monthly users of drugs other than marijuana.

Alcoholism is known to be a widespread problem in the Soviet Union and reported estimates of alcohol dependency in the Soviet military vary from 18 to 30 per cent, which would put them on a level comparable with their US counterparts. In addition, it seems likely that drug problems in the Soviet military have increased since the invasion of Afghanistan, for reasons similar to those which caused increased drug use in the US military during the Vietnam War. The Soviet troops are fighting against relatively popular Afghan guerrillas who are not directly attacking the Soviet Union. They are fighting a drawn-out war with ill-defined objectives, in a foreign country, which, like Vietnam, is an opium-producer.

Specific data exist showing the extent of human reliability problems in the US *nuclear* military. From 1975 to 1984, approximately 5100 people per year were removed from ongoing nuclear duties, and roughly two-thirds of these removals were for drug and alcohol abuse or psychological reasons.

Working conditions in the nuclear military entail stress, boredom and isolation which adversely affect normal human beings, let alone those with drug, alcohol or psychological problems. Workers are placed at electronic consoles or in silos for hours on end, or sent out to submerged submarines for months at a time. They are isolated from humanity at large, and trained to destroy it. Although they are constantly practising for destruction, they are never allowed to follow through. Each of these conditions produces stress.

The "dulling effects of routine" are dangerous when operators have to deal with situations which suddenly become critical. For example, during the Three-Mile Island nuclear power accident, operators continued their routine tasks in the face of unusual conditions. Moscow party chief Boris Yeltsin said that human error was the cause of the Chernobyl accident.

Social psychologist Irwin Altman, has observed laboratory simulations of isolation. Group isolation, especially on long missions, has been shown "to

produce performance-reducing stress, increased possessiveness, defense of space and objects . . . and a greater tendency to personal isolation." Overall, these findings reveal that social isolation can be detrimental to human reliability both on and off the job. This conclusion has been confirmed by at least one former member of the nuclear military.

A final source of stress is the "familiarity syndrome." The more accustomed to operating a system humans become, the less attention they tend to pay to ensuring that no serious problems arise. The January 1986 Challenger space shuttle accident, the April 1986 Titan 34D rocket explosion, and the April 1986 Delta rocket failure seem to have "shattered" the complacency that years of "normal" missions had created. These accidents, however, do not rule out future missions in the way that nuclear war would almost certainly rule out future nuclear wars.

The final dimension of human reliability dealt with by Dumas is the effect of bureaucracy. He finds the most serious problem here to be the transmission of valid information from the bottom operational levels of the military hierarchy to the top. The reluctance of subordinates to point out their own or their superiors' mistakes, personal beliefs, rigid world views and concepts of loyalty have been shown to weaken the transmission of accurate information to the top. In this way dangerous problems may go unrecognized.

In business it has been recognized that the reluctance to impart bad news to the top is a widespread problem.

The consensus seems to be that there are basically two ways to handle this problem: create an organizational culture, an atmosphere of open communication and trust with informal channels of communication available; sharply reduce the degree of hierarchy in the organization.

Unfortunately, military organizational practices run counter to these sorts of reform. In addition to inhibiting upward communication, organizational barriers can also cause downward directives to be distorted, diverted or ignored:

One spectacular example of this command and control problem with obvious relevance to the nuclear arms situation, played a key role in the Cuban Missile Crisis of 1962, probably the closest the world has yet come to intentionally initiated full-scale nuclear war. At that time Soviet Premier Nikita S. Krushchev offered to withdraw the Russian missiles from Cuba if, among other things, the United States removed its nuclear-tipped missiles from Turkey. According to Attorney General Robert F. Kennedy, the President had asked the State Department to reach an agreement with Turkey "on several occasions" over the preceding 18 months, to withdraw the US Jupiter missiles from its territory. Apparently, on the last of these occasions (summer

1962) the President was told by the State Department that it was unwise to press the matter, but he disagreed and told them that he wanted the missiles removed. In Robert Kennedy's words, "the President believed that he was President and that, his wishes having been made clear, they would be followed and the missiles removed . . . The State Department representatives discussed it again with the Turks and, finding they still objected, did not pursue the matter. And so the international situation was seriously aggravated at a crucially dangerous point in human history, not by a conspiratorial plot, not by a mentally-deranged or drug-addicted military officer, not even by stress, monotony, or familiarity-induced failure of vigilance, but simply by bureaucratic inertia.

As far as Dumas was concerned, no aspect of the human reliability problem can ever be fully eliminated: "We cannot circumvent this dilemma by turning control over to machines, by somehow automating the human element out of the nuclear forces. For who designs machines and who will build them?" This being the case, he argued that we should exercise two other aspects of our humanity, namely, our wisdom and our instinct for survival, in order to recognize that ". . . the only effective military strategy for increasing national security is general nuclear disarmament."

Paul Huth and Bruce Russett's paper reported their research on sixty cases of "extended deterrence" which had occurred in the international system since 1880. Deterrence was defined, in this instance, as one nation threatening the use of force to prevent the first use of force by another nation. "Extended deterrence", on the other hand, meant to prevent an attack on another party such as an ally, client state, or friendly neutral. These cases of extended deterrence were analyzed statistically, in order to determine the circumstances under which deterrence was likely to succeed and the circumstances under which, if it failed, the crisis was likely to escalate to full-scale war.

The study examined four factors which affected the success of deterrence. First, deterrence tended to be successful in cases where the immediate balance of forces favoured the defender. This suggested that the attacker was probably dissuaded if there seemed little chance of being able to accomplish a quick *fait accompli*. However, the long-term balance of forces, which might ensure a defender's ability to prevail in a war of attrition, was not relevant, nor was a defender's possession or non-possession of nuclear weapons. Second, ties between the defender state and its protégé (for example, geographic proximity, alliance, military sales and assistance, and trade) seemed on the whole to be irrelevant, though military ties were of some importance before actual interstate bargaining began. Third, firm-but-fair and tit-for-tat diplomatic military behaviour was associated with successful deterrence, while either disproportionately bellicose or disproportionately concessionary behaviour were not. Fourth, clear-cut victory or defeat in previous encounters seemed to embolden

attackers to defy the next deterrent threat, whereas previous stalemates did not.

The study also examined how crises were resolved once deterrence had failed and the attacker pressed ahead. The crisis was more likely to escalate to war, and the defender more likely to fight, to the extent that (1) the defender was geographically close to the protégé and was its ally, (2) the short-term (and to some extent the long-term) military balance was in its favour and (3) it had previously followed a firm-but-fair diplomatic strategy.

Huth and Russett argued that various elements, including the balance of power, crisis behaviour and a state's reputation, all affect the success of deterrence and whether or not crises escalate to war. They asserted that avoiding war was not simply a matter of possessing great relative strength or behaving in a tough, inflexible manner.

They also discussed some differences between the first and second phase, and their impression that the variables which affected the attacker's decision to press ahead in the first phase seemed to have less influence on the defender's decision to fight once deterrence had failed. However, this did not necessarily mean that the attacker was generally wrong in its judgement of the defender; in some instances, the attacker might have inferred from those variables that the defender was likely to fight, and backed off. First, the firm-but-fair strategy is associated both with successful deterrence and with escalation to war in a crisis. This suggests that attackers who ignore firm and fair threats are likely to provoke a response. Second, long-term military balance of power is much more important to the second phase suggesting that whereas an attacker may press ahead looking forward to a *fait accompli*, it may be met with a defender capable and willing to sustain a long war. Third, the relatively greater role of alliance and geographical proximity in the second phase suggests that, while the attacker may pay little attention to the fact that the defender has both a material investment and its reputation at stake in its protégé, these factors can motivate the defender to fight back, resulting in escalation to war. Interestingly enough, the defender's possession or non-possession of nuclear weapons seems irrelevant and does not affect either the decision of an attacker to defy deterrence or that of a defender to resist attack.

To sum up, Huth and Russett's paper showed the influences which affect the development or non-development of crises after the employment of extended deterrence. Since contemporary major power confrontations are frequently fought through surrogates, this study was important in providing some insight into the probable setting of accidental nuclear war.

General Milstein commented on the papers presented by Russett, Huth and Dumas.

Concerning Russett and Huth's paper, Milstein argued that one should be careful in making extrapolations from conventional to nuclear conflicts. He believed that "deterrence" had not developed a strategic meaning before the advent of nuclear weapons, and noted that it would be impossible to give examples of cases in which the onset of nuclear war provided a test of the effectiveness of deterrence. He also suggested that a crucial element in extended deterrence was whether or not it was requested by the protégé; some cases of extended deterrence might be part of an imperialistic plan.

Milstein categorically denied Dumas' assertion that there was any drug or alcohol problem in the Soviet nuclear strategic forces, though he could not make any such statement on the Soviet forces in general. He observed that many of the working conditions in the nuclear forces, such as secrecy and isolation, were found in other parts of the service.

He agreed with Dumas that the only way to avoid nuclear war is to abandon nuclear weapons, and repeated his citation of Gorbachev's proposal to rid the world of nuclear weapons in three stages, by the year 2000. He concluded that no matter how much command systems are improved, no matter how well nuclear force personnel are screened, or how well leaders are educated, as long as there are nuclear weapons there is a danger of nuclear war.

Russell Leng also provided a critique of Dumas' and Russett and Huth's papers. He believed that Dumas' work, while important, was still at the stage of forming hypotheses, and observed that Dumas based his conclusions on anecdotes rather than on a systematic statistical overview. This first stage of research would have to be fully completed before empirical research could go on to a second, statistical stage.

Leng-observed that Dumas' opinions converged with those of other researchers. He noted that Dumas' conference paper, the Russett/Huth paper and the work of some social psychologists all indicated that a firm-but-fair bargaining strategy is the best way to avoid war. Neither bullying nor appeasement are as effective. In addition, finding that the loser in one crisis is likely to be belligerent in the next crisis is particularly applicable to chronic crises such as those between India and Pakistan, Israel and Egypt, or the United States and the Soviet Union. This finding suggests that leaders get a progressively stronger message that they must play tough. Using different methods and analyzing different cases, researchers have come to markedly similar conclusions. Leng found statistical analyses exciting because they lead to similar conclusions even though they lack the spontaneous appeal of anecdotal works.

The presentation by Haag and de Seguin focussed on the French experience with nuclear weapons, and attempted to use that experience to illustrate some general points with regard to the command and control of nuclear forces. They began by pointing out that the problem of nuclear

accidents is not confined to that of general war; there are also several other sorts of accidents to be considered.

First, there is the significant risk of an accident involving a single nuclear weapon. This danger is not without significance in the French case, because of the large number of accidents which have occurred involving French nuclear-capable aircraft and missiles. This is not a specifically French problem; there is some evidence that the newly-deployed American Pershing II and cruise missiles are far from completely reliable.

Second, there is the risk of an accident in the stockpiling process. We have no data concerning this risk, but we do know that weapons-grade plutonium stockpiles are subject to inventory "shrinkage", indicating a danger at this stage as well, since the whereabouts of what is unaccounted for is unclear.

Third, there are the risks involved in the actual manufacture of weapons-grade plutonium itself, especially at the fast-breeder reactor at Malville. It is well known that such reactors are much more dangerous than ordinary pressurized light water reactors.

The French experience also revealed a risk not yet discussed at this conference – the possibility that civil conflict might result in the illicit seizure of nuclear weapons by a rebel group. During the "Generals' Revolt" of 1961 some of the rebel forces apparently attempted to seize a nuclear weapon which was being readied for testing in the Sahara. To forestall this, the weapon was detonated prematurely.

A further problem for France is that of control. The French command system is totally centralized; only the personal order of the President of the Republic can release the "force de frappe." But in any conceivable scenario, the President would have less than ten minutes to decide whether to give the order. To counter with the hollow boast that the President has "nerves of steel" is to fly in the face of everything we know about human psychology. The French nuclear doctrine is that of *presidential infallibility*, analogous to the doctrine of papal infallibility, except that in our case "God" is the H-bomb.

Dean Babst's paper recommended the creation of accidental war assessment centres around the world to be funded in a way commensurate with the degree of the péril. He asserted that non-nuclear nations could contribute to the world's arms reduction efforts by publicizing their own studies on accidental nuclear war. He noted in particular that computerized analytical models are good tools for assessing the dangers of accidental nuclear war. They can anticipate dangers, provide quick answers and build increasingly sophisticated scenarios; and they cost relatively little.

A first step in fostering assessment centers has been the creation of the Accidental War Information Exchange by the Nuclear Age Peace Foun-

dation in California and the Richardson Institute of the University of Lancaster in England. The Exchange would act as a co-ordinating service, build on existing information centres and encourage the release of information. Accidental war research would supplement, not supplant, other efforts to reduce and eliminate nuclear weapons.

VI.

CONCLUSIONS AND RECOMMENDATIONS

The conference participants agreed unanimously that the week's deliberations had provided a valuable forum for an exchange of views and information. As someone put it, the problem of accidental nuclear war reminds one of the parable of the four blind men and the elephant. The computer scientist, the strategist and the social scientist, the European, the Russian and the North American, each see the problem through the prisms of their various disciplinary and national perspectives. Bringing the different viewpoints together was not only instructive, but generated an unexpected convergence of ideas. Two examples of this are noteworthy.

First, it was clear that there was considerable agreement between the mathematicians and computer scientists on the one hand, and the strategists on the other, that the reduced warning times arising from the deployment of short-flight time weapons systems had had a severe impact on the ability of the command system to cope with false alarms. Three streams of research – one based on mathematical modelling, a second founded on the known limits of computers and artificial intelligence systems in duplicating human reasoning, and a third grounded in a detailed examination of the command systems and nuclear alert procedures – resulted in very similar conclusions.

Second, there was also a great deal of convergence concerning the type of national behaviour that was most likely to avert the risk of inadvertent war. A “firm-but-fair” (or “tit-for-tat”) strategy, combining firmness with conciliation, appeared to be the most efficacious in avoiding crises. Once a crisis had developed, however, threats were found to be dysfunctional in avoiding escalation to war. These conclusions were buttressed by three different sets of research findings: by the Russett/Huth studies of extended deterrence, the Leng studies of crisis behaviour, and Rapoport's research into the computer simulation of conflict behaviour.

But agreement at the conference extended well beyond specific research findings and conclusions. Believing the danger of accidental war to be critical, the conference participants drafted a statement in which they sought to voice their concerns. They unanimously concluded that

the danger of accidental nuclear war is substantial and increasing for the following reasons among others:

1. deteriorating global political relations coupled with lack of real progress in disarmament and arms control, and the high frequency of international crises;
2. escalation of the arms race leading to the development and deployment of destabilizing weapons systems and technologies;

3. increasingly complex and unmanageable command and control systems with reduced warning times demanding decisions and actions on a time scale exceeding human capabilities;
4. increasing reliance on automated decision-making systems leading to a greater likelihood of catastrophic error. Measures must be taken to halt this drift toward unparalleled catastrophe.

The statement goes on to assert that a number of steps can be taken to avert the danger, but cautions that "purely technological measures will not eliminate the risk."

There was insufficient time to forge a consensus concerning specific policy measures, and many felt that such an effort was inappropriate without the participation of officials responsible for policy implementation. However, a number of participants did make general recommendations which met with considerable agreement.

One recommendation (made most forcefully by Leonard and Sennott) was that both superpowers should eliminate or prevent the deployment of weapons systems which have the effect of increasing the load on warning systems. Deemed to be especially dangerous were short flight-time systems, such as Pershing II missiles and SLBMs, deployed in the opponent's coastal waters. ASAT systems also fall into this category, as their use would cripple each side's ability to resolve false alarms quickly.

A second recommendation, emphasized by Blair, was the need for command systems which could survive a nuclear exchange. The achievement of such a survivable command capability would remove the danger of decapitation, thus reducing the incentive to adopt a policy of launch-on-tactical-warning in time of crisis.

A third recommendation, strongly endorsed by both the strategists and computer scientists at the conference, warned against the increasing reliance on automated decision-making in nuclear command systems. Fallible and uncertain as human behaviour often is, the use of automated systems in the unpredictable circumstances of a severe crisis would inevitably result in a still greater likelihood of failure and more uncertainties, and thus a greater risk of unintentional war.

A fourth recommendation, made by Crissey, called for additional unclassified data on false alarms and warning system failures to be made available to researchers. Until recently, a considerable amount of data had been made available, but in the last two years or so there had been new attempts to restrict access to this information. This policy change severely inhibits a thorough scientific examination of the problem.

A final recommendation called for the creation of some institutional mechanism to engage in an ongoing study of the risk of war by accident. A

specific proposal to this effect was put forward by Babst, who suggested the formation of accidental war assessment centres in different nations around the world. Others proposed that new studies be funded, and that a follow-up conference be held with the specific aim of bringing the issues discussed at the conference to the attention of the officials directly responsible for policy in these areas.

APPENDIX: LIST OF SPEAKERS AND PARTICIPANTS

Dean BABST	Chairman, the Accidental Nuclear War Project, the Nuclear Age Peace Foundation.
John BARRETT	Research associate, the Canadian Centre for Arms Control and Disarmament.
Bruce BLAIR	Research associate, the Brookings Institution; former member, US Department of Defense Task Force on Nuclear Command and Control
Marco CARNOVALE	Graduate student in Political Science, M.I.T.
Brian CRISSEY	Head, Department of Computing Science, Linfield College, Oregon.
Lloyd DUMAS	Professor of Political Economy, the University of Texas at Dallas.
Daniel FREI	Professor of Political Science, the University of Zurich; author of the United Nations' Report on the Risks of Accidental War.
Luc de SEGUIN	Paris physician who has researched French nuclear accidents.
Michel HAAG	Paris physician who has researched French nuclear accidents.
Martin HELLMAN	Professor of Electrical Engineering, Stanford; known for work in cryptography and military information systems.
Paul HUTH	Graduate student, Political Science, Yale.
Fred KNELMAN	Former director of the Science and Technology Program at Concordia University; author of <i>Reagan, God and the Bomb</i> .
John LAMB	Executive Director, the Canadian Centre for Arms Control and Disarmament.
Russell LENG	Professor of Political Science and former Dean of Arts, Middlebury College, Vermont; specialist in international crisis behaviour.

Barbara LEONARD

New Zealand computer professional;
General Secretary of Computer People
Against Nuclear War.

**Colonel-General
Mikhail MILSTEIN**

Retired officer attached to the Institute
for the Study of the US and Canada,
Moscow.

Johan NIEZING

Director, the Institute of Polemology, the
Free University of Brussels; author of
several works on nuclear accidents.

Severo ORNSTEIN

President, Computer Professionals for
Social Responsibility, Palo Alto,
California.

Anatol RAPOPORT

Professor of Peace Studies, the University
of Toronto; past President of Science for
Peace.

Douglas ROSS

Associate Professor of Political Science,
the University of British Columbia.

Bruce RUSSETT

Dean Acheson Professor of International
Relations, Yale University; editor of the
Journal of Conflict Resolution.

Linn SENNOTT

Professor of Mathematics, Illinois State
University; author of articles on the
mathematical modelling of nuclear
control and system failure.

Henry THOMPSON

Lecturer in the Artificial Intelligence
Programme, University of Edinburgh.

Joseph WEIZENBAUM

Professor of Computer Science, M.I.T.; a
pioneer in the field of artificial
intelligence; author of *Computer Power
and Human Reason*.

CONFERENCE CHAIRMAN

Michael WALLACE

Professor of Political Science, University
of British Columbia.

CONFERENCE CO-ORDINATOR

Debbie MCCLELLAND

United Nations Association of Canada,
Vancouver Branch.

CONFERENCE ASSISTANT

Kim ARNOTT

Political Science student, University of
British Columbia.

The Canadian Institute for International

LIBRARY E A/BIBLIOTHEQUE A E




3 5036 20041226 3

Board of Directors

William Barton,
Chairman
Norman Alcock
Christoph Bertram
George Bell
Lise Bissonnette
David Braide
Harriet Critchley
Kalevi Holsti

Richard Mercier
Geoffrey Pearson,
Executive Director
John Sigler
Brian Urquhart
Jean Guy Vaillencourt
Paul Warnke
Lois Wilson
Gerald Wright



The Canadian Institute for International Peace and Security
307 Gilmour Street
Ottawa, Ontario K2P 0P7