Top crop mark   Center image left to right   Top crop mark

NUREG/CR-3926
Vol. 1

Strategic Analysis for Safeguards Systems: A Feasibility Study

Main Report

Prepared by D. A. Seaver, A. J. Goldman, W. H. Immerman,
F. L. Crane, M. S. Cohen

The MAXIMA Corporation

Prepared for
U.S. Nuclear Regulatory
Commission

NOTICE

Availability of Reference Materials Cited in NRC Publications

Most documents cited in NRC publications will be available from one of the following sources:

1. The NRC Public Document Room, 1717 H Street, N.W.
   Washington, DC 20555
2. The NRC/GPO Sales Program, U.S. Nuclear Regulatory Commission,
   Washington, DC 20555
3. The National Technical Information Service, Springfield, VA 22161

Although the listing that follows represents the majority of documents cited in NRC publications, it is not intended to be exhaustive.

Referenced documents available for inspection and copying for a fee from the NRC Public Document Room include NRC correspondence and internal NRC memoranda; NRC Office of Inspection and Enforcement bulletins, circulars, information notices, inspection and investigation notices; Licensee Event Reports; vendor reports and correspondence; Commission papers; and applicant and licensee documents and correspondence.

The following documents in the NUREG series are available for purchase from the NRC/GPO Sales Program: formal NRC staff and contractor reports, NRC-sponsored conference proceedings, and NRC booklets and brochures. Also available are Regulatory Guides, NRC regulations in the *Code of Federal Regulations*, and *Nuclear Regulatory Commission Issuances*.

Documents available from the National Technical Information Service include NUREG series reports and technical reports prepared by other federal agencies and reports prepared by the Atomic Energy Commission, forerunner agency to the Nuclear Regulatory Commission.

Documents available from public and special technical libraries include all open literature items, such as books, journal and periodical articles, and transactions. *Federal Register* notices, federal and state legislation, and congressional reports can usually be obtained from these libraries.

Documents such as theses, dissertations, foreign reports and translations, and non-NRC conference proceedings are available for purchase from the organization sponsoring the publication cited.

Single copies of NRC draft reports are available free, to the extent of supply, upon written request to the Division of Technical Information and Document Control, U.S. Nuclear Regulatory Commission, Washington, DC 20555.

Copies of industry codes and standards used in a substantive manner in the NRC regulatory process are maintained at the NRC Library, 7920 Norfolk Avenue, Bethesda, Maryland, and are available there for reference use by the public. Codes and standards are usually copyrighted and may be purchased from the originating organization or, if they are American National Standards, from the American National Standards Institute, 1430 Broadway, New York, NY 10018.

GPO Printed copy price: ______________

Top crop mark Center image left to right Top crop mark

NUREG/CR-3926
Vol. 1
RS, 1S

Strategic Analysis for Safeguards Systems: A Feasibility Study

Main Report

Manuscript Completed: September 1984
Date Published: December 1984

Prepared by
D. A. Seaver, M. S. Cohen, The MAXIMA Corporation
W. H. Immerman, F. L. Crane, International Energy Associates Limited
A. J. Goldman, Consultant, John Hopkins University

The MAXIMA Corporation
7315 Wisconsin Avenue - Suite 900N
Bethesda, MD 20814

Subcontractor:
International Energy Associates Limited
2600 Virginia Avenue, N.W.
Washington, DC 20037

Prepared for
Division of Risk Analysis and Operations
Office of Nuclear Regulatory Research
U.S. Nuclear Regulatory Commission
Washington, D.C. 20555
NRC FIN B6763

ABSTRACT

Strategic analysis (game theory) is a formal method for modeling adversary situations that, when solved, yields an optimal strategy that maximizes the expected payoff to the player. As such, it appears to be potentially applicable in the nuclear material accounting context in which there is potential for an adversary attempting to divert special nuclear material. The NRC has previously supported research to develop preliminary strategic analysis models which has been considered to be only partially successful. This study reviewed previous efforts and other game theory research and assessed the feasibility of: 1) applying strategic analysis in a regulatory framework, 2) making strategic analysis understandable by licensees, and 3) assuring that strategic analysis can effectively be enforced. This report includes a discussion of the role of strategic analysis in material control and accounting, and of the mechanisms by which the NRC could implement strategic analysis. A set of feasibility criteria are described including both technical feasibility and organizational/implementation feasibility. Alternative strategic analysis model options are evaluated with respect to these criteria, as is the current material accounting practice. The assessment determined that the development of a payoff function that adequately represented the NRC's (and therefore the public's) values with respect to the consequences of diversion and the actions taken to prevent it is the most serious impediment to implementation. Given the limited role of material accounting in safeguards and the uncertainty regarding the development of a payoff function, the NRC should not proceed with full-scale implementation of strategic analysis. It does, however, have sufficient potential to warrant further development, with first priority going to the development of an appropriate payoff function.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS

This research benefitted considerably from the support and assistance of the two NRC Technical Monitors, Andrea Kuffner and Sandra Fratalli. Dan Lurie of the NRC also contributed useful background information and insight into the technical details of game theory in material accounting and current material accounting practices. In addition, several NRC staff members in NMSS and I&E met with project staff to provide needed information and expertise on practical aspects of implementing strategic analysis. The contributions of all of these individuals are gratefully acknowledged, though the interpretations, conclusions, and recommendations remain those of the authors.

EXECUTIVE SUMMARY

This report describes the results of research undertaken to investigate whether strategic analysis (game theory) can and should be employed in the treatment of inventory differences (ID's) of special nuclear material (SNM) at processing facilities. Strategic analysis is a formal method for modeling decision contexts with more than one decision maker--often adversaries. When solved, a strategic analysis model yields an optimal decision that maximizes the player's expected payoff with payoffs previously defined for the player. Previous NRC-supported research developed two versions of a preliminary model (NUREG-0290 and NUREG/CR-0490) which have been subsequently reviewed by a group of experts (NUREG/CR-0950). Drawing on this past research, other literature on game theory, discussion with NRC staff, and analysis by the project team, this report assesses the feasibility of: 1) applying strategic analysis in a regulatory framework, 2) making strategic analysis understandable by licensees, and 3) assuring that strategic analysis can effectively be enforced.

As background, the role of material control and accounting (MC&A), including the limited role of material accounting in diversion detection, is discussed. The most feasible instruments that the NRC could potentially use to implement strategic analysis include: regulations, regulatory guides, and license conditions.

To assess the feasibility of strategic analysis, a framework was created consisting of two parts: 1) a breakdown of feasibility into major component criteria that are conceptually distinct and may be independently addressed, and 2) a breakdown of methods for applying strategic analysis into a set of roughly independent features or options. The feasibility criteria include both technical feasibility, which includes validity and practicality and which are in turn further subdivided, and organizational feasibility which is subdivided into subcriteria concerning both implementation and enforcement by both the NRC and the facility operator. The methods correspond to alternatives with respect to each of the components of the strategic analysis model and to NRC organizational/implementation options. The model components include the players, the setting, the defender's (NRC) strategies, the diverter's strategies, the defender's payoff function, and the diverter's payoff function. The complete set of technical model options and organizational options was evaluated with respect to each of the feasibility criteria. Then, a composite of the most promising strategic analysis options was compared on the feasibility criteria with current practice.

This assessment process provided indications of the cost and feasibility of resolving four specific questions that had been posed by the NRC regarding: 1) the use of probabilistic mixtures of strategies, 2) the formulation of the payoff function, 3) the practicality of a model that is both valid and solvable, and 4) the cost of using strategic analysis. This process also assessed the value, cost, and likelihood of success of four research questions: 1) optimal tradeoffs between false alarms (Type I errors) and missed diversions (Type II errors), 2) development of measures of effectiveness for the establishment of alarm thresholds, 3) evaluation of selection procedures for alarm thresholds at reprocessing and other breeder reactor cycle facilities, and 4) the costs and benefits of implementing strategic analysis.

As a result of this assessment, two issues were identified as the most significant potential impediments to the implementation of strategic analysis: 1) the formulation of a valid and acceptable payoff function would be a difficult and probably costly endeavor, and 2) the role of material accounting, particularly analysis of ID's (and therefore strategic analysis when applied to ID analysis) is quite limited for most safeguards functions, particularly diversion detection, and therefore, any improvements in ID analysis may have only limited effects on overall safeguards. Therefore, we cannot recommend that the NRC proceed with full-scale implementation of strategic analysis. It clearly, however, has sufficient potential to warrant further development. Research on the payoff function without the full strategic analysis model could provide an initial payoff function and at the same time could resolve several research issues that are relevant to safeguards even if strategic analysis is not implemented. Such research is indeed worth pursuing.

1.0 INTRODUCTION

This report describes research undertaken to investigate whether strategic analysis (game theory) can and should be employed in the treatment of inventory differences (ID's) of special nuclear material (SNM) at processing facilities. The NRC has previously supported research that developed two versions of a preliminary model (NUREG-0290 and NUREG/CR-0490). Subsequently, the NRC formed a group of experts on game theory and nuclear safeguards (subsequently referred to as the "Peer Review Group") to review this work and to make recommendations regarding the future of strategic analysis in NRC's safeguards function (NUREG/CR-0950). That review identified several technical difficulties with the particular model formulated, but found game theory to be a promising tool for safeguards. This study builds on the findings of the Peer Review Group to address the feasibility of resolving several technical and practical issues that must be addressed before strategic analysis could be implemented.

1.1 Overview of Game Theory

Before describing this effort further, a brief, non-technical description of game theory is given to provide a basis for subsequent discussions. A more complete description is provided in the Appendix. Game theory is a formally appropriate technique for modeling situations where two or more actors (players) are controlling some of the variables of the situation and the value or utility received by each player depends at least in part on the actions of the other players. Generally, the objectives of the players will be at least partially conflicting so that increased rewards for one player can result in decreased rewards for another. Thus, generally, game theory is applicable in adversary contexts. Once the rewards for each player are quantitatively defined for the various combinations of actions available, game theory defines and (assuming sufficient computational resources) can solve for the "optimal" action for a player, taking into account the actions of other players.

A game theory model consists of four elements:

- the players,
- the setting,
- the players' strategies, and
- the players' payoff functions.

Although game theory models need not be limited to two players, two player games have been the focus of much of the game theory research, including previous NRC-sponsored efforts. Two person games often provide a reasonable representation of the actors involved, for example, a "defender" (e.g. NRC representing the public), and a diverter.

The setting allows for numerous variations in game theory models. For example, the NUREG-0290 and NUREG/CR-0490 model deals with ID's in a single material balance areas, and (in other literature) have been developed for multiple time periods. Game theory models could be developed for multiple material balance areas and for multiple time periods.

The players' strategies refer to the actions that each player has available to take. For example, in the previously developed model the defender's actions involve setting an alarm threshold for the level of ID, as well as varying levels of response depending upon whether the threshold is exceeded and the amount of SNM estimated to be missing. The diverter's strategy involves selecting the amount to be diverted, which can range from nothing up to a fixed amount which is necessary to provide a credible threat.

The payoff functions are the quantitative representation of the costs and benefits to each player that would result from each possible combination of actions by all the players. In many two-player situations, the objectives of the players can be assumed to be diametrically opposed so that any loss by one player is a gain by another. Such games are called zero sum games, and the previously developed model includes this assumption. The payoff function in that model included terms representing the cost of a clean-out inventory, the cost of the diverted SNM, the cost of searching for the estimated amount of missing SNM, the cost of an incorrect estimate of the amount missing, and the utility of recovering the missing material.

1.2 Objective

The present research was undertaken to assess the feasibility of integrating strategic analysis into current NRC material control and accounting (MC&A) regulations. Specifically, the objective was to investigate whether it is feasible to:

- apply strategic analysis in a regulatory framework,
- make strategic analysis understandable by licensees, and
- assure that strategic analysis can effectively be enforced.

1.3 Scope

Our approach to assessing feasibility has been a broad one, encompassing both the technical feasibility of strategic analysis models, and the organizational feasibility of implementing strategic analysis within the current NRC regulatory environment. In addition we have addressed, to some extent, whether strategic analysis should as well as could be implemented.

At the outset of the effort, several limitations on the scope of the effort were agreed upon.

- Strategic analysis was to be compared only with current regulations including formal NRC management positions as written, and not with current practices beyond these regulations, nor with changes in regulations currently under consideration.
- The focus of the feasibility study should be on Category I facilities.
- Strategic analysis should be explicitly considered only in the context of material accounting and its concern with inventory differences.
- Overall assessment of feasibility should be generic, rather than dependent on particular features of specific strategic analysis models.

Based on this approach, this research was carried out drawing on the expertise and experience of the project staff, reviews of relevant literature and documents, and discussions with appropriate NRC staff.

1.4 Organization of the Report

In addition to this introduction, the main body of this report contains four major sections. Section 2.0 provides a background discussion on the potential role of strategic analysis in safeguards and on NRC options for implementing strategic analysis. Section 3.0 describes the framework that was developed for assessing the feasibility of strategic analysis. Then, in Section 4.0, this framework is applied to evaluate strategic analysis alternatives, and strategic analysis more generally, with respect to current practice. Finally, in Section 5.0, we summarize our findings in terms of specific questions and issues developed by the NRC in its Request for Proposals for this effort. This section also includes specific recommendations.

In addition to the main body, this report includes an appendix as Volume II prepared by a member of the project team, Professor Alan J. Goldman. This appendix provides additional detail regarding the study, providing an especially helpful review of related literature and many insights into feasibility issues that are addressed at best in a limited fashion in the main report.

2.0 BACKGROUND

One critical point, perhaps the major one in determining the feasibility of a strategic analysis approach to acting on measured inventory differences is, its effectiveness. It is apparent that the usefulness of strategic analysis depends on its leading to an expected net increase in effectiveness of whatever changes would be required in material accounting practice. Furthermore, this increase must be large enough to offset the costs (economic and otherwise) of whatever changes would be required in material accounting practice. If strategic analysis were used strictly as an adjunct to current material accounting, then the combination would perform no worse than current practice, but cost (including acceptability) would be higher than if the change substituted strategic analysis for some current material accounting.

Measuring effectiveness of safeguards is not a trivial matter. In fact, our inability to accurately assess the probability and type of potential adversary actions ultimately limits our ability to measure safeguards effectiveness. Various safeguards effectiveness methods have been developed and used, but they generally apply only to some limited portion of safeguards system To analyze the incremental effectiveness of strategic analysis relative to current material accounting, it is necessary to outline the objectives and goals of nuclear safeguards and security, and the elements of the safeguards system designed to achieve these goals. This information will help to identify: the role of material accounting; its place among other parts of the safeguards system; some questions which must be answered to assess safeguards system effectiveness and the value of changes to that system; and other potential areas of applicability for strategic analysis in nuclear safeguards and security.

2.1 The Role of Nuclear Material Control and Accounting at the NRC

2.1.1 Scope. This section provides an overview of areas in which NRC must make decisions to discharge responsibilities in nuclear material control and accounting (MC&A). The emphasis is on fairly general responsibilities, rather than on decisions regarding the adequacy of a particular safeguards subsystem. A particular goal of the discussion is to provide perspective on the role and potential contribution of material accounting with respect to overall NRC safeguards responsibilities.

Our approach to identify MC&A decisions is to characterize safeguards in several ways and to associate relevant questions with each characterization. Each question relates to the adequacy of NRC safeguards; answering it represents an opportunity or requirement for NRC decisionmaking. Since each characterization represents an alternative view of safeguards structure, the questions are not independent of each other: though not duplicative. In general, although not exclusively, the questions relate to adequacy of the measurements or criteria of safeguards regulation.

Within this structure, we may further characterize the role of material accounting and evaluate its potential contribution to safeguards. The monumental assessment of strategic analysis in material accounting relative to conventional accounting methods should be considered in the perspective of the overall role of material accounting. An appropriate question may be, "Is a modification of requirements to use strategic analysis of inventory differences (IDs) cost-effective given the bounded contribution of accounting to overall nuclear materials safeguards?" On the other hand, in stressing the relative contribution of an accounting system change to safeguards, we should not bias the outcome by assuming limitations on the role of accounting which are based on inadequacies of current techniques not shared by strategic analysis.

The overview presented in this report should also provide a basis for identifying alternative safeguards applications for strategic analysis. Although this study did not examine such alternatives in detail, we note--as did the strategic analysis Peer Review Group (NUREG/CR-0950)--that this basic methodology could perhaps be fruitfully applied to other questions of safeguards adequacy.

2.1.2 NRC Safeguards Responsibilities. NRC safeguards responsibilities stem from an overall objective, which is to protect the public health and safety from death, injury, or property damage which could result from malevolent activity directed toward, or utilizing licensed nuclear materials or facilities. These activities would hypothetically be conducted by an individual referred to generically as an "adversary." There are several alternative decompositions of the safeguards responsibilities which help to illuminate their scope, and to focus the identification of safeguards decision criteria. Some of these alternatives represent distinct dimensions of safeguards (e.g., safeguards system functions versus type of facility under safeguards). The following represents the top levels of a hierarchy categorizing NRC safeguards decision areas:

R. Regulatory Considerations

- R.1 Measurement of Safeguards and Safeguards Criteria
- R.2 NRC Functions (regulation, licensing, inspection, other)

S. Safeguards (SGs) Systems

S.1 SGs Program Components (physical protection, material control, material accounting, safeguards management)

S.2 Facility Types (reactor, fuel cycle, transportation)

S.3 SGs System Functions (deterrence, prevention, response, assurance)

S.4 SGs Risk (frequency of attempt, conditional probability of success, consequence, adversary action sequence)

At this point, we desire that any area which might illuminate a pertinent SGs decision be examined. However, many such decisions are not relevant to material accounting, per se, and so will be eliminated from consideration for this study. Thus, we will not attempt a unique, partitioned decomposition of these decisions. Rather, we will temporarily treat each item at the same level of the hierarchy as complementary. Specifically, we will regard the decisions and questions which are meaningful for our purposes as being constructed as "cross products" of items from different branches of the hierarchy, as long as one is not dependent on another. For instance, it would be sensible to ask if assessment techniques were adequate (R.1) to assess the suitability of a proposed regulation (R.2) for prevention of relevant acts (S.1) directed at sabotage of a power reactor (S.2). This procedure simplifies the identification of such decisions by reducing them to separate concerns which can be combined. The determination of the answer to such a question is an NRC safeguards decision.

Combining these areas will not make equally good sense, but nonsensical cases produced by this systematic exhaustive procedure can readily be identified and discarded as they arise. For the balance of this section, questions implied in each area will be identified and discussed separately, with the intent that they be combined for future application freely within the limits of meaningful juxtaposition.

NRC safeguards decisions may be immediately divided into those components related to the regulatory role and those related to safeguards system performance. It should be noted that NRC has no statutory responsibility for design, implementation, or operation of safeguards system. Its responsibility and authority is to <u>regulate</u> and <u>license</u> activities. The subject of NRC safeguards regulatory actions is, however, systems and operations. Thus, the system characterization must be understood to assess safeguards systems, and their cost effectiveness are also relevant inputs to regulatory decisionmaking. In some cases, systems or measurements are "conceptually designed" by NRC and specifically required or strongly recommended. In these cases, NRC assumes a more direct role in system implementation than it does when establishing performance requirements. Finally, certain NRC activities are not really regulatory, but rather are a corollary of other agency capabilities and responsibilities. In these instances, NRC may fulfill a need not adequately met elsewhere. For instance, NRC is involved in assessment of communicated threats involving licensed material although such an event occurs outside the framework of licensing and regulation. For all these reasons, NRC safeguards responsibilities are integrally tied to the detailed characteristics of safeguards systems as well as those of regulation.

R. Regulatory Considerations

In terms of the upper level of the hierarchy discussed above, the regulatory aspects of NRC safeguards decisions are as follows:

R.1 Measurement of SGs and SGs Criteria

A fundamental distinction in the regulation of safety is that between the measurement of safety level and the establishment of criteria for adequacy of safety. Some questions related to each are:

R.1.1 Safeguards Performance Measurement

- Have detailed objectives for SGs systems been established which, if met, would assure that the overall SGs objectives are met?
- Have the parameters necessary to measure safeguards system performance been identified?
- Are methods available to assess SGs system performance in terms of these parameters and relate them to SGs (sub-) objectives? Are they comprehensive and accurate? Is the modeling of uncertainty acceptable?
- Are the appropriate data available for input to the assessment methods?
- Are regulatory and license requirements appropriately tied to the performance assessment capability? (e.g., will licensee reporting requirements provide the necessary data for assessment?)

(Note that the assessment methods referenced here may range from subjective and informal evaluation by a qualified person, to use of formal, mathematical tools for modeling system performance or for analyzing data. In fact, achievement of adequate tools for safeguards system assessment has so far required a mix of methods including both of these extremes.)

R.1.2 SGs Criteria Development

- Have criteria been developed which are compatible with the available assessment methods?
- Is satisfaction of these criteria necessary and sufficient for achieving the safeguards (sub-) objectives? If not, is there a less stringent criterion whose satisfaction would still be sufficient to meet (sub-) objectives with "reasonable assurance"? (This is a complex issue--but briefly, two important points are the acceptability of regulatory conservatism implied by "reasonable assurance," and the relative nature of "less stringent," i.e., less stringent for NRC or for the licensed industry, etc.)
- Does the regulatory entity satisfy the criteria? (By entity we mean regulation, license, system, etc.)
- If not, what alternative entities could be developed or adopted to the required performance?

R.2 NRC Functions

The discussion here is further amplified by the description later in this section of NRC regulatory instruments. In particular, that description outlines major points relevant to the choice of instruments with which to implement a requirement. The following will focus on certain additional issues:

R.2.1 Regulation

- Are the regulations or proposed legislations legal?
- Can they be implemented?
- Are they adequately clear to licensees and license-applicants? If not, are supporting documents available tp elaborate and clarify their intent?
- Are they consistent internally, with other sections, and with other regulated activities?
- Will the regulations, if implemented, result in licensee safeguards systems which satisfy all regulatory criteria?

R.2.2 Licensing

- Are requirements for license applications clear and available to prospective licensees?
- Are licensing criteria consistent with regulations?
- Is an effective process established for license review?
- Are license terms and conditions clear to the licensee?
- Are license terms and conditions inspectable and enforceable, legally and practically?
- Do license terms and conditions, if implemented, satisfy the intent and letter of the regulations, and any other regulatory criteria?

R.2.3 Inspection

- Will NRC inspection activity adequately verify licensee compliance with the terms and conditions of the license, and with the regulations?
- Will the inspection program alert NRC to other potential safety problems?
- Will the inspection program verify that the activity, as implemented, agrees with the design or plan as submitted to NRC? (This is actually part of the first question, since license application information is generally incorporated by reference into a license.)

R.2.4 Other NRC Activity

- Are sufficient NRC resources available to implement regulatory decisions?
- Are NRC mechanisms utilized to identify additionalor incipient problems related to safeguards system performance?
- Are NRC mechanisms utilized to improve understanding of safeguards performance and to reduce uncertainty?
- Is NRC prepared to accomplish its interagency responsibilities in safeguards? (Important instances include emergency response and communicated threat-assessment functions.)
- Is the public confident that NRC safeguards criteria are: sufficient to protect them; not unduly intrusive; representative of the achieved safeguard level?

S. Safeguards Systems

S.1 NRC SGs Program Components

At this level of generality, the questions and decisions associated with each of the components below essentially are covered by asking whether all other applicable questions have been answered with respect to each part of the program individually. Relevant definitions and explanations are summarized below. They are paraphrased from the MC&A Task Force Report (NUREG-0450):

S.1.1 Physical Protection

This is the safeguards program which encompasses equipment and procedures to control physical access to nuclear materials or facilities, to detect unauthorized access, and to respond in a timely manner to such intrusion.

S.1.2 Material Control

This portion of the safeguards program monitors nuclear material location and quantity, controls its movement, and exercises responsibility for its appropriate use and for cognizance of its status.

S.1.3 Material Accounting

This portion of the safeguards program measures, records, reports, and analyzes the quantities of nuclear material from the time it initially comes under regulatory or management controls, until it is permanently disposed of, or transformed into uncontrolled material. The material accounting responsibilities shift among different parties as they take possession of the material. The intent is to maintain a continual "audit trail" of the material quantities. Sufficient control should be maintained to be able to balance the records to show that no net loss or gain of material has occurred. In practice, this is possible for unit accounting (i.e., accounting for discrete items of material), but not for material in process. For the latter material (for instance liquids in a process stream at a reprocessing plant), process variations prevent perfect material balance. In these cases, statistical techniques are used to determine the acceptability of the material balance.

S.1.4 Safeguards Management

The MC&A Task Force distinguishes the "NRC activities necessary to assure the effectiveness and efficiency of the safeguards program" by this name. It includes all regulatory and licensing activities, intelligence activities, safeguards research, etc.

S.2 Facility Types

These categories are also self explanatory. They are, however, relevant to the implementation of safeguards systems and requirements as indicated below:

S.2.1 Nuclear Reactors

These facilities are of safeguards concern mainly from the point of view of sabotage. The potentialto initiate a reactor accident by malevolent means, thereby using the energy generated by the accidentto disperse hazardous material, is the principal safeguards consideration. Fresh fuel for power reactors is generally low enriched and therefore not suitable for building weapons (unless means for further enrichment are available, a possibility outside the spectrum of adversary capability which NRC feels is prudent to consider). Irradiated fuel contains special nuclear material (SNM) which is extremely radioactive, and which would require chemical processing even to achieve concentrated low grade (not weapons grade) fissionable material. The hazard and low utility of this irradiated fuel makes it an unlikely target for theft. Certain research reactors use (smaller) quantities of high enriched fuel which could be more attractive targets for theft. Thus, reactor safeguards consist largely of physical protection.

S.2.2 Fixed Site Fuel Cycle Facilities

These are all static facilities which use SNM. They include such facilities as enrichment plants, fuel fabrication plants, and reprocessing facilities. Of these, only fuel fabrication facilities are currently of concern with regard to NRC safeguards. (There are no enrichment or reprocessing plants currently subject to NRC safeguards.) Although criticality accidents caused by sabotage could pose some safeguards concerns, the major safeguards interest in this type of facility is to protect against theft or diversion of material. Physical protection, material control, and material accounting are all important in achieving this protection.

S.2.3 Transportation Safeguards

SNM in transit, including spent fuel, is of concern to safeguards due to the possibility of theft including that of the entire vehicle. In the case of spent fuel, this could be accompanied by later sabotage in a populated area, or the threat to do so. Theft of other forms of SNM is of concern. Material in transport provides special problems for physical protection. Since material is packaged in specially designed containers, MC&A considerations are limited to material controls and unit accounting.

S.3 SGs Functions

The four safeguards functions identified by the MC&A Task Force are defined as follows:

(1) Deterrence: measures to discourage a potential adversary from attempting a malevolent act.

(2) Prevention: measures to impede or stop a malevolent act, including a hoax.

(3) Response: measures to detect and assess losses, and to act in response to a real or alleged malevolent act.

(4) Assurance: measures to enhance NRC and public confidence that effective safeguards are in place and that nuclear materials are controlled and accounted for.

The MC&A Task Force associated detailed objectives and goals for the MC&A system with these functions. Each of these goals and objectives can be turned into a question or questions--i.e., is the specified measure in place and is it effective? We do not repeat these here. A more detailed analysis can do so and would find the Task Force Report extremely useful for this purpose. It should be noted that the goals of that report are not part of current safeguards requirements. Below, we summarize some general questions appropriate to each function, based on the Task Force recommendations (note that these are particular to MC&A SGs):

S.3.1. Deterrence

- Is detection and subsequent response sufficient to achieve adequate deterrence?

- Are material controls adequate to ensure a timely detection of unauthorized removal of SNM with high enough probability to create desired deterrence?

- Will material controls trace and identify adversaries with sufficient resolution to effect adequate deterrence?

- Is material accounting adequate to augment detection and to identify where loss may have occurred, or records falsified, to aid in deterrence of potential adversaries?

S.3.2 Prevention

- Will material controls detect unauthorized removal or diversion of material in sufficient time to allow effective response?

- Will material accounting augment detection of unauthorized removal of diversion, in certain areas? Will it help provide a baseline knowledge of material holdings?

- Can this function affectively interface with physical security to stop or impede a malevolent act in progress?

S.3.3 Response

- Are response mechanisms planned and reliable to identify loss of material not previously detected?

- Are source materials adequate to identify the source and amount of loss?

- Is material accounting adequate to aid after a loss of material in assessing the amount of loss, the mode of removal, the location of the removal or diversion, possible locations to which removal was diverted, and potential insider suspects?

S.4 Safeguards Risk

Safeguards risk concepts were developed and expressed in "Societal Risk Approach to Safeguards Design and Evaluation" (ERDA-7). One author of that report was also on the Strategic Analysis Peer Review Group; another was the head of an initial MC&A Task Force which preceded the one which generated the Task Force Report. Safeguards risk, generally, is expressed as:

$$R_{SG(Total)} = \sum_{ij} \Pi_{ij} P_{ij} C_i$$

where Π_{ij} = frequency of attempt using adversary action sequence j to achieve event i

P_{ij} = probability of success given an attempt

C_i = consequence of event i

Measurement of the total R_{SG} is probably impossible, due to our inability to estimate Π_{ij}.

It is commonly believed that Π_{ij} depends on P_{ij} (which is directly influenced by safeguards), and on C_i or more precisely, on the potential adversary's perception of these terms. Furthermore, the common assumption is that this dependence is monotonic increasing so that the greater the probability of success, the greater the frequency of attempt would be. On this basis, the risk equation has been used to priority-rank SGs concerns. Some pertinent considerations of the risk-based approach are:

- Is the overall risk acceptable? (It is unlikely that this can be known precisely.)
- Are all significant (e.g., risky) adversary action sequences known and protected against.
- Are all significant events known and protected against?
- Are measures to decrease frequency of attempt adequate?
- Are safeguards systems adequately controlling P_{ij}?
- Is any given event of such consequence, C_i, that special measures should be used to assess its risk and to safeguard it?
- Is the variability in our knowledge of R_{SG}(Total) (i.e., uncertainty) acceptable?

- Are risks balanced between events and adversary action sequences (i.e., is $R_{ij}=R_{i'j'}$ for each i,i',j,j'?).

- Is total SGs risk balanced against other societal risks from other sources?

- Is public (and adversary) perception of the components of SGs risk accurate? If not, what effect does this have on Π_{ij}?

2.1.3 Inventory Differences (IDs). The goal of the preceding discussion was to generally identify safeguards system decisions by describing them in a form amenable to elaboration. The format provides some structure to assure a reasonably comprehensive identification.

We do not intend at this time to identify in detail all of the particular decisions regarding material accounting, and in particular, inventory differences. These could be developed as refinements of the existing hierarchy. Based upon the preceding discussions, however, we note several points pertinent to IDs.

First, it is possible to characterize material accounting using our structure. Obviously, it is not physical protection, materials control, or safeguards management (S.3). The subject of interest in this study, inventory difference, would be a further refinement of hierarchy element S.1.3. It is mainly of interest for a fixed fuel cycle facilities, and in the case of this study we are only concerned with a small number of Category I facilities (S.2).

The MC&A Task Force concluded that material accounting is pertinent principally to the assurance function, although it has a lesser role in deterrence, prevention, and response as well. In particular, the Task Force Report minimizes the role of material accounting in detection. It is not expected to contribute significantly to "timely detection" (S.3). Existing SGs risk approaches do not really treat the "assurance" function as described by the MC&A Task Force. Thus, the risk approach (S.4) has little to say about material accounting, except perhaps that by implication it may not be a large contributor to SGs risk reduction. In terms of regulatory concerns, consideration of strategic analysis in inventory difference is connected with both measurement and criteria. (Even though IDs and LEIDS currently have a role in measurement of SGs effectiveness, the application of a different model to those statistics may affect the interpretation of what is being measured (R.1).) Finally, application of strategic analysis to IDs could influence all regulatory functions (regulation, licensing, inspection, and other), although it need not.

Thus, it appears that measurement of IDs represents only a portion of MC&A decisions, which in turn play a relatively small role in NRC safeguards. In particular, its role in detection and response is less significant than its contribution to assurance. The game theory model being evaluated, however, is primarily directed at the former two functions. For these reasons, we recommend that the feasibility and cost effectiveness of modifying safeguards requirements include the use of strategic analysis applied to IDs be considered in light of net contribution to safeguards objectives. On the other hand, it is possible that the <u>qualitative role</u> of material accounting could change if more sensitive techniques were used. This possibility should be addressed in any research related to use of strategic analysis for timely detection. Furthermore, we note that there are many other potential applications of game theory to safeguards not directly addressed by this study. Finally, we suggest that the following specific areas pertinent to ID measurement and analysis be explored.

- S.3 Elaborate specific pertinent decisions in deterrence, prevention, and response.
- S.4 Try to quantify the role and contribution of accounting in risk reduction.
- R.2 Identify specific approaches to incorporating this use of strategic analysis in regulatory instruments. Identify specific pros and cons of each, and if possible estimate costs of each development. Focus on useof regulations, regulatory guides, license conditions, and informal reports (e.g., NUREGs).

2.2 <u>Regulatory Requirements--NRC Options for Implementing Requirements</u>

2.2.1 <u>Scope</u>. The premise of this section is that NRC has, in theory and in practice, a variety of means by which it may implement requirements. Our use of the term "requirements" for the purposes of the discussion is rather imprecise and not in keeping with its generally accepted regulatory denotation--which is to say a set of legally binding rules. This departure from practice is intentional, since the discussion will try to convey that in a very real sense NRC discharges its regulatory duties utilizing a variety of what we will term "regulatory instruments" (documents and procedures), including but not limited to regulations.

This section attempts to characterize the regulatory process via an examination of the various regulatory instruments. The purpose of doing so is to establish a basis for later investigation of the organizational effects of implementing strategic analysis in material accounting. The discussion here will elaborate on the differences in effects of each option for regulatory implementation, as well as the factors influencing cost of implementation. The presentation is qualitative; neither measures of effectiveness nor absolute or relative costs nor times to implement are given. In fact, reliable quantitative estimates of these factors will be difficult to achieve. Such estimates would be, however, more appropriately produced in the context of choice among alternative instruments as a means for achieving some specific objective. This discussion will address only the global view of alternatives for regulation implementation.

2.2.2 Overview. One way to categorize regulatory practice is by differentiating the distinct functional activities performed by a regulatory agency. In the case of NRC, the main regulatory functions have generally been identified as:

- promulgating standards
- licensing
- inspecting

There are, as well, additional NRC functions including performing research; enforcing compliance; conducting business with other Federal, state, and international agencies; and administering the agency itself. It is clear that any number of functional characterizations are possible. The above three, however, are considered typical of regulatory as distinct from the other agencies. It is interesting to note that the Energy Reorganization Act of 1974 (P.L. 93-438), which created the U.S. Nuclear Regulation Commission, created three statutory offices: The Office Of Nuclear Reactor Regulation (NRR); the Office of Nuclear Material Safety and Safeguards (NMSS); and the Office of Nuclear Regulatory Research (RES). Clearly, Congress did not organize the agency along functional lines (or perceived a different view of function than outlined above).

Associated with two of the three regulatory functions are instruments which directly impose legally binding requirements on NRC licensees. That is, regulations contained in Title 10 of the Code of Federal Regulations (10 CFR) are the rules promulgated by NRC. Licenses are, of course, the regulatory instruments associated with licensing. Licensees are inspected (the third function) to determine compliance with both the regulations and with their license, so this activity does not introduce any new primary regulatory instruments. We will return to the question of inspection briefly later.

Licensees are obligated to comply with both the set of applicable regulations and with the conditions of their license. Penalties for non-compliance include having their license revoked or suspended, have operations suspended, and being fined. These penalties imply substantial costs, especially in the event operations are suspended. The Energy Reorganization Act, in addition, provides NRC the authority to conduct inspections as necessary to ensure compliance (Section 206(d)). Thus, these two instruments, i.e., the regulations in 10 CRF (Parts 0-199) and individual licenses, constitute a legal, enforceable framework of NRC requirements with the opportunity for meaningful sanction in the event of non-compliance. Furthermore, these two are directly authorized by the NRC's enabling legislation. The Energy Reorganization Act of 1974, in Section 201(f) transfers most of the Atomic Energy Commission's (AEC's) "licensing and related regulatory functions" to NRC. These functions derive in turn from the authority provided by the Atomic Energy Act of 1954 (as amended), Chapter 10 and Chapter 14, Section 161(b). Loosely speaking, these laws provide that all production, acquisition, possession, use, transfer, import, or export of nuclear facilities or material in the U.S. will be governed by regulations to be promulgated by the Commission and will be licensed in accordance with those regulations and with the Act itself.

NRC regulations (10 CRF Parts 0-199) and NRC-issued licenses, then, are similarly rooted in law going back to the earliest days of commercial use of nuclear materials. They differ mainly in their target audiences. In the case of a license, the coverage extends only to a designated activity and licensee. Regulations apply more generally to a class of licensees, which may be empty or may have a single member (e.g., proposed high-level waste disposal regulations). Most often, multiple licensees are affected.

The choice in practice between these two mechanisms for imposing requirements is instructive in identifying some of the pragmatic concerns which have created the other regulatory instruments which this discussion addresses. The key would appear to be in the distinction between the two, namely generic vs. facility- or activity-specific concerns. The decision to use regulation or licensing to establish a criterion could accordingly be decided depending on whether the requirement applied to many potential licensees or just one. In fact, this doesn't fully explain certain instances and phenomena. Why, for instance, has there been extensive effort to write regulations for a high-level waste disposal facility which may be the only one of its kind, and will almost surely belong to the sole licensee (i.e., the Department of Energy)? On the other hand, why are hundreds of licenses for certain nuclear materials written with the same, or highly similar, license conditions? Why do some licenses include or paraphrase certain regulations and not others, while other licenses do not repeat the regulation at all?

One answer to these questions is that there is no specified criteria which determines what should go into a regulation and what should not. The regulations in 10 CFR have been written and modified over many years, according to prevailing concerns and standards.

2.2.3 Regulatory Instruments. This section identifies a variety of regulatory instruments and characterizes some aspects of each recognizing that consistency is not necessarily a characteristic of any of them.

2.2.3.1 Regulations. Regulations were discussed in the previous section to some degree. They are (usually) generic to a group (possibly empty) of licenses, are authorized by law, and impose legally binding requirements on licensees. These features all may act as selection criteria if a change in (effective) requirements is desired. That is, a regulation change or addition might be a useful approach if NRC judged that a requirement were necessary in all but exceptional cases within some readily identified groups of licensees. It should be noted that regulations themselves do not normally apply absolutely. Rather, they may be conditioned and exceptions may be granted at NRC's discretion. A material control regulation in 10 CFR 70.58(a) is an instance of a conditional requirement. It applied only to licensees who are authorized to possess at any particular time and place no more than one effective kilogram of special nuclear material. Here, the condition is fixed and identified in the regulation. Part 70.14, on the other hand, provides for exemption from Part 70 requirements based on no fixed condition.

Despite this possibility for regulatory flexibility, exceptions to 10 CFR requirements are rare. Therefore, if site-specific concerns are known to be important, the regulator might have reason to prefer using a regulatory instrument which can be tailored on a case-by-case basis (e.g., a license).

Consistent with the view of regulations as generic is the notion that they tend to be general. This characterization is useful, but again has its exceptions. Part 30, containing generally applicable rules for domestic licensing of byproduct material which governs a substantial number of licenses, provides both an example and the counterexample. Most of the part defines what materials and facilities are subject to these requirements and how licenses may be obtained and maintained. There are two brief sections defining criteria for granting a Part 30 license. The first, 30.33, "General requirements for issuance of specific license, "states criteria at a level of detail typified by the following:

> (a) An application for a specific license will be approved if: (1)...(2)The applicant's proposed equipment and facilities are adequate to protect health and minimize danger to life or property...

Similarly, 10 CFR 30.34, "Terms and Conditions of Licenses," states, in effect, that the Commission may add to the conditions of any license at the time it is issued or subsequently "as it deems appropriate or necessary in order to: (1) promote the common defense and security; (2) protect health or to minimize danger to life or property:..."

In the same section, however, we find 30.34(g):

> Each licensee preparing technitium-99 in radiopharmaceuticals from molybdenum-99/technitium-99m generators shall test the generator eluates for molybdenum-99 breakthrough in accordance with 35.14(b)(4)(i) through (iv) of this chapter.

The examples represent extremes of the relationship between regulation and licensing practice. The very general regulations cited here support detailed license application criteria requiring specific equipment and facilities. The detailed requirements, on the other hand, is repeated almost verbatim in licensing guidance, with little additional elaboration in some cases. Thus, we can say that although regulations may be a good choice to implement generic or general requirements it is not necessary that this be done. Conversely, a regulator can suitably condition a requirement or rely on exemption to utilize a regulation which may not apply to all cases.

Pragmatically, it is less frequent to see a specific requirement implemented in the regulations than it is to see general requirements expressed in another form of instrument. This observation is based on some additional characteristics of regulations: they have a long lead time for introduction, they are difficult and costly to implement, and often stimulate more controversy than lower level changes. Since regulations are law, there is a prescribed process through which significant modifications must go. The process includes advance publication in the Federal Register, public comment, response to public comment, and likely modification of the proposed rule. The cycle may be iterated several times, and the result may be delay or indefinite postponement of a final rule. This period may include, in complex cases, involved public and industry meetings. Note that this process is interleaved with a broad internal production and review cycle within the NRC, often spanning several major offices at the staff level, separate review by the ACRS (Advisory Committee on Reactor Safeguards) and, of course, ultimate review by the Commission itself (i.e., the five Commissioners). Finally, it should be noted that 10 CFR 51 delineates requirements for the Commission to satisfy criteria of the National Environmental Policy Act of 1969 (NEPA). These include the requirement for an Environmental Impact Statement for substantive changes to 10 CFR Parts 20, 30, 40, 50, 70, 71, 73, and 100 (from an environmental impact point of view). Adopting a new approach to material control and accounting could fall in this category. If not, an analysis would be required to demonstrate that the change was not substantive (10 CFR 51.44(c)). It is likely, though, that a negative declaration could be based on arguments that no impact was expected since safeguards requirements had not been relaxed.

Finally, two additional characterizations of regulations as opposed to other regulatory instruments may be appropriate. First, although the enabling legislation does not appear to explicitly require it, all NRC licensed activity is also covered by the provisions of some part of the regulations. The coverage may be very broad, adding no significant criteria for licensability or operation, but to the best of our knowledge the applicable regulations exist in each case. Second, the process of promulgating a formal requirement, although painstaking, is valuable for communication among interest groups, including the various agencies of the Federal Government, the state governments, the regulated industry, and the public at large. It would be unacceptable in certain cases to circumvent the process by failing to issue a new or modified regulation.

2.2.3.2 Regulatory Guides. Regulatory Guides are a series of documents providing elaboration of criteria in 10 CFR Parts 0-199, evaluation methods or assumptions which the staff has found acceptable in assessing compliance with a requirement, acceptable methods of implementing certain systems or functions, or guidance in composing a license application. (They were initially called "Safety Guides," and some early ones still bear this title.) They are categorized into ten divisions, as follows:

Division	Area
1	Power Reactors
2	Research and Test Reactors
3	Fuels and Materials Facilities
4	Environmental and Siting
5	Materials and Plant Protection
6	Products
7	Transportation
8	Occupational Health
9	Antitrust and Financial Review
10	General

These Regulatory Guides do not have the "force and effect of law" as do regulations. They are, however, widely used in establishing interpretations of rules and acceptability criteria for them. A previous study by members of this study team found that despite the admonition that the Guides are only guidance and that a licensee can propose acceptable alternatives, it is highly unlikely that an applicant would have an alternative accepted. In the area of reactor licensing, the applicant generally takes a much more active role than do those small material users; therefore, they may be more successful in establishing alternatives to the positions in the applicable Regulatory Guides. In those cases, however, the analysis is much more complex, and deviation from the norm more risky. Thus, in practice, Regulatory Guides are a common source of "requirements" in the sense of this discussion. That is, in a non-deterministic sense, they are responsible for a large degree of the uniformity seen in licenses which is unexplained by common requirements in the regulations. Once again, though, we must express caution regarding inconsistency. Some Regulatory Guides are not even as explicit as the related regulation. In many cases, however, Regulatory Guides provide substantially more detailed criteria than do the regulations.

In terms of impact, production of Regulatory Guides tends to be much less expensive and more flexible than is promulgation of a regulation. This results directly from the fact that they are not law nor issued by the Commission, but are in fact issued by the Regulatory Staff. The review they undergo varies, but for a consequential guide the review can be considerable, internally. A licensee is free, in principle, to deviate from the practice or evaluation presented in the Guide, so that much of the negotiation associated with public comment on rules may be deferred in the cases of Guides to individual licensing case review rather than wholesale conflict over the Guide.

As mentioned before, reactor-related RGs differ from others somewhat in practice and utilization. In the nuclear materials area, the issuance of a Guide can be a relatively lengthy process due to the relatively low priority of non-reactor Guides and the limited visibility of an RG compared with a regulation. In some cases, draft Regulatory Guides have existed for years without being finalized. Since they are not binding, they are still usable as guidance even in draft form if they express the position of the licensing staff. In this form, however, they may be less than adequate. Writing and approving a Regulatory Guide can take resources from several offices, including the relevant program office and the Office of Research. They are also often subject to ACRS or other review. Finally, one must consider the impact of the dichotomy between the nominal use of an RG as "guidance" and its likely use to define criteria without the level of review and comment afforded regulations. Qualitatively, it appears that NRC staff would be much more likely to prefer to use guidance than to promulgate a regulation and a formal Regulatory Guide is perhaps the most common form of guidance.

2.2.3.3 Standards. Standards here refers to the use of industry standards, as opposed to the work performed by NRC's former Office of Standards (now part of Research). NRC is often represented on various standards committees and, in turn, often incorporates such standards by reference in other instruments. For instance, the American Society of Mechanical Engineers (ASME) Boiler and Pressure Vessel Codes are incorporated by reference in 10 CFR 50, and in several Regulatory Guides as well. NRC is free in doing so to adopt only portions of the standard or otherwise to modify it for use in a nuclear facility. These standards have no independent effect on nuclear regulation apart from use in another instrument. In practice, however, if a broad, technically qualified group can reach a consensus, the resulting standard is likely to have some value to a variety of users, including NRC.

Standard setting outside NRC is generally a long, expensive process in its own right. It is generally of value when a large enough group has commercial interest in a common standard. It is unlikely to be pertinent to use of strategic analysis in MC&A.

2.2.3.4 Standard Review Plans. A standard review plan is a blueprint for the license review process, including items to be evaluated and criteria to be applied. It is therefore of considerable value to an applicant as well as to a reviewer. It has roughly the status of a Regulatory Guide, although it is focused more on what the review process should be than on what evaluation method is acceptable. A standard review plan, in the sense intended here, exists only for reactor licensing. Elsewhere in NRC licensing, review checklists and other less ambitious or formal plans exist. The development of particular equipment would not indicate that a standard review plan was needed. However, if a plan existed, a criterion could be effectively implemented by incorporating it in the plan as standard practice.

2.2.3.5 License Conditions. Licensing has already been discussed to some extent. License conditions are the instruments which justify the premise that many non-binding "requirements" are actually binding. As pointed out in the discussion of Regulatory Guides, a de facto standard is often extremely influential across many licenses, even if it is not legally a requirement. A license condition allows a license reviewer to "close the loop" on such "requirements," by stipulating that the licensee must abide by the terms of his final license application. This application would reference or use the criteria found in a Regulatory Guide or other instrument. The net effect is that of a legally binding requirement.

Some license conditions establish common requirements without any other instrument. In these cases as well, a common requirement is incorporated in a condition to the license. Thus, even if a criterion is not written elsewhere, it may act as a requirement by virtue of appearing consistently in many licenses.

The impact of conditioning a single license is minimal. It may be expensive if contested by a license applicant, but one would presume that the applicant would object at least as strongly to the requirement if published in some other form (e.g., as a Regulatory Guide). The cost of repeatedly evaluating and incorporating conditions on a case by case basis for each license can be expensive, especially where there are many licensees (e.g., materials licenses under Parts 30 and 40). However, in the case of the Category I licensees, who are subject to the stricter material control and accounting requirements, individual license modifications could be a realistic approach to establishing a new practice.

2.2.3.6 Other. Several other instruments are used or have been used to impose criteria on licensed activities. One such example is the Branch Technical Position, used at one time within NRR. It has largely been superseded b_ use of Regulatory Guides and the Standard Review Plan, although the regulations still refer to some. The use of Branch Technical Positions points out, however, that a position taken by even a small organizational unit responsible for a given type of licensing effectively imposes requirements with little overhead associated with review and promulgation. Since the Branch is largely responsible for that portion of licensing review and have considerable discretionary power, the written position merely codified a standard practice. By itself, it has no legal stature, but in conjunction with the license review and approval process it can effectively implement "requirements" in the sense used within this paper.

NUREG reports are informational documents produced by NRC, often reporting the results of a research program or other Commission activities. Rather than rewriting them into a rule or a guide, however, they are sometimes distributed and used as another form of guidance. The incremental impact of this use (compared with the original cost of producing the report) is low, although its effectiveness relative to another instrument may be more variable, depending on mutual acceptance of the approach by licensing staff and the license applicant.

Finally, we should briefly note that although inspection is not ostensibly a criterion-setting activity to the extent that regulation and licensing are, it is, of necessity, judgmental. Therefore, criteria are in effect established according to what an inspector does or does not notice or utilize. There are criteria and training for the inspection program which attempt to standardize what "good inspection practice" should be. It might be possible to make use of the inspection program as an implementation vehicle for a new regulatory practice, especially one like the use of strategic analysis.

2.3 The Use of Strategic Analysis for Material Accounting

The existing strategic analysis model for material accounting (NUREG/CR-0490) is principally directed toward selecting a defender strategy for estimating actual material losses, and possibly an alarm threshold, given measured inventory differences. It assumes two players (defender and diverter) with opposing interests (zero sum game) for a particular payoff function. It is used as if material accounting were the only component of safeguards evaluation, and as if timely detection were the only function of material accounting. In particular, NUREG/CR-0490 (pp. 2-4) states that "nuclear material accounting safeguards in the nuclear industry relies on material balance accounting. This accounting signals the occurrence of losses, if any, and may be the basis for subsequent inspection and recovery actions.... The decision-making problem is, given a MUF [material unaccounted for] reading, what action should be taken to verify possible theft and/or recover material that may possibly have been diverted.... The essential safeguards material accounting decision problem is how to establish the alarm threshold in a manner to satisfy safeguards objectives."

As noted earlier, this may not reflect the most accurate impression of the role of material accounting. For instance, the Material Control and Accounting Task Force identified assurance, not detection, as the most significant role of material accounting. Assurance is directed toward verifying whether the safeguards system is under control and functioning as intended. It also provides one means by which alarms may be assessed to aid in estimation of the probability, time, and amount of loss. Functioning as an assurance mechanism, the material accounting system may be used to respond to a variety of "signals," including alarms (possible detections) resulting from inventory analysis or from other parts of the safeguards system (e.g., material control system). In some cases, assurance is required to assess non-alarm signals (e.g., communicated threat messages, intelligence information). Finally, assurance plays a role in safeguards management unrelated to response and alarm assessment; i.e., demonstrating that the system is under control and providing a baseline for future measurements. Certainly, a greater degree of effectiveness of material accounting as a timely detection mechanism could shift the emphasis from accounting primarily for assurance to accounting principally for detection. The inadequacy of material accounting as the principal means of timely detection, however, is at least partially based on inevitable process noise in the measured data itself. Since this noise is unrelated to adversary action, it would affect the efficacy of strategic analysis as well as conventional material accounting. In any case, the model as formulated emphasizes what is currently perceived by NRC as a secondary material accounting function. If material accounting using an improved strategic analysis model is to provide an adequate detection mechanism, some research may be indicated to resolve a natural skepticism as to its absolute effectiveness in that role (as distinguished from its effectiveness relative to current practice).

In addition to the possible overemphasis on the importance of material accounting as a contributor to effective diversion detection, the existing strategic analysis model appears to underemphasize the contribution of the rest of the safeguards system to safeguards decisions. Such an assumption is not necessary to a strategic analysis model, and should be avoided in any implementation-worthy model. NUREG/CR-0490 (p. 11) says, "The defender's decisions are: given a MUF reading, whether to alarm or not and what preliminary estimate of unauthorized diversion to make, based upon the MUF reading. This preliminary diversion estimate will influence the defender in the resources he allocates for the post-alarm search. A final estimate of diversion should be based upon not only the results of the post-alarm search, but also relevant information from other safeguards system (e.g., the physical security system and the material control system) as well as pertinent external intelligence information, police reports, etc. This study, however, is concerned only with the material accounting system. Consequently when we talk about the estimate of the amount diverted, we mean only the preliminary estimate based upon the MUF reading in and of itself." Despite the acknowledgement of a role for other parts of the safeguards system, this assumption does not account for the realistic role played by these other subsystems in signaling, assessing, and responding to an alarm.

Subsequent sections of this report discuss certain technical assumptions of the existing model (e.g., two-person, zero sum assumptions). It should be noted here, however, that material accounting roles are relevant to several of these considerations as well. For instance, the framework for safeguards objectives, goals, and functions discussed above from NRC's perspective is actually part of a larger structure, with similar but not identical features for operators, and for diverters as well. Inventory data are recorded and material balances drawn at least in part to satisfy the operator's payoffs that may reflect different attributes and utilities than NRC's. Thus, the defender must be carefully specified if a useful payoff function is to be constructed.

In each of the above cases, discrepancies between the model's assumptions and actual or perceived material accounting roles may affect both the technical adequacy and the acceptability of this application of strategic analysis. Some of these points will be discussed further in the balance of this report.

3.0 FRAMEWORK FOR FEASIBILITY EVALUATION

3.1 Need for a Framework

Assessing the feasibility of applying strategic analysis in a regulatory context involves the examination of a large number of analytic and organizational options with respect to numerous criteria or dimensions of feasibility. It is critical, for example, that strategic analysis as such not be pronounced infeasible due to flaws in a single proposed model, or, conversely, that important obstacles to its implementation not be overlooked. Confidence in the conclusions of such an assessment depends on some assurance of completeness, both with respect to options and criteria. To provide such assurance, a systematic framework has been created within which both options and criteria of feasibility may be placed. The framework is meant to include all, or at least most, of the major options facing the NRC and all, or at least most, of the features relevant for selection of an option for actual implementations.

Utilizing this framework, each option can be analyzed with respect to each relevant evaluative criterion. The goal, however, is not a final pronouncement regarding feasibility in each case. Rather, the advantages of structuring the discussion in this way are threefold: first, it provides a powerful means of organizing currently available information about the relative merits and feasibility of different approaches to the analysis of inventory differences; second, we can pinpoint areas where the feasibility of an approach has not been proved or disproved, but where additional evaluative research is required; finally, we can suggest directions for the development of new analytic or organizational approaches to strategic analysis that retain the advantages but avoid the shortcomings of currently proposed methods.

3.2 Overview of the Framework

The framework for feasibility evaluation consists of two parts:

- a breakdown of feasibility into major component criteria that are conceptually distinct and may be independently addressed
- a breakdown of methods for applying strategic analysis into a set of roughly independent features or options

Feasibility criteria (Exhibit 3-1) fall under two main headings: technical and organizational. Technical feasibility, in turn, breaks down into validity and practicality; organizational feasibility is subdivided into subcriteria concerning both implementation and enforcement with respect to both the NRC and to the relevant facility.

Options for the implementation of a strategic analysis model also fall under technical organizational categories. Technical options (Exhibit 3-2) involve the common set of elements possessed by all game-theoretic models: players, setting, strategies (of the defender and diverter), and payoff functions (of the defender and diverter). The NUREG/CR-0490 model is only one point in a rather large space of strategic analysis approaches which might be considered (along with current practices) as methods for analyzing inventory differences. Organizational options (Exhibit 3-3) concern methods of implementation and of enforcement.

3.3 Role of Material Accounting Objectives

In evaluating the options (laid out in Exhibits 3-2 and 3-3) against the criteria (in Exhibit 3-1), an important additional concern is the objective or set of objectives which strategic analysis and material accounting are thought to be serving (see Section 2.0 above). The Material Control and Material Accounting Task Force (NUREG-0450) identifies four potential objectives: deterrence, prevention, response, and assurance. Exhibit 3-4 describes in broad terms the functions that material accounting would have to perform to contribute to each of these objectives. Thus, any contribution of material accounting to prevention (i.e., the disruption of an on-going diversion attempt) would require either timely detection based on material accounting or timely use of material accounting to verify (i.e., back-up) the results of other systems (e.g., physical security or material control). In fact, as noted in Section 2.0, a timely detection role for material accounting seems highly improbable.

A contribution of material accounting to the response objective (where the extent of an already-accomplished diversion must be assessed) requires either that it be the initial source of the detection or that it verify or rectify the conclusion of some other system. Material accounting contributes to assurance by demonstrating that the process of measuring inventory differences falls within statistically acceptable limits.

The framework or merit of a particular method of analyzing inventory differences may depend in a dramatic fashion on which of these functional requirements is thought to be preeminent. The objectives and requirements establish the relative importance of the feasibility criteria laid out in Exhibit 3-1. For example, prevention emphasizes practicable methods that can be applied quickly, even if at some expense in accuracy, while assurance sets the opposite priority. As another example, both prevention

Exhibit 3-1

Criteria for Feasibility Evaluation

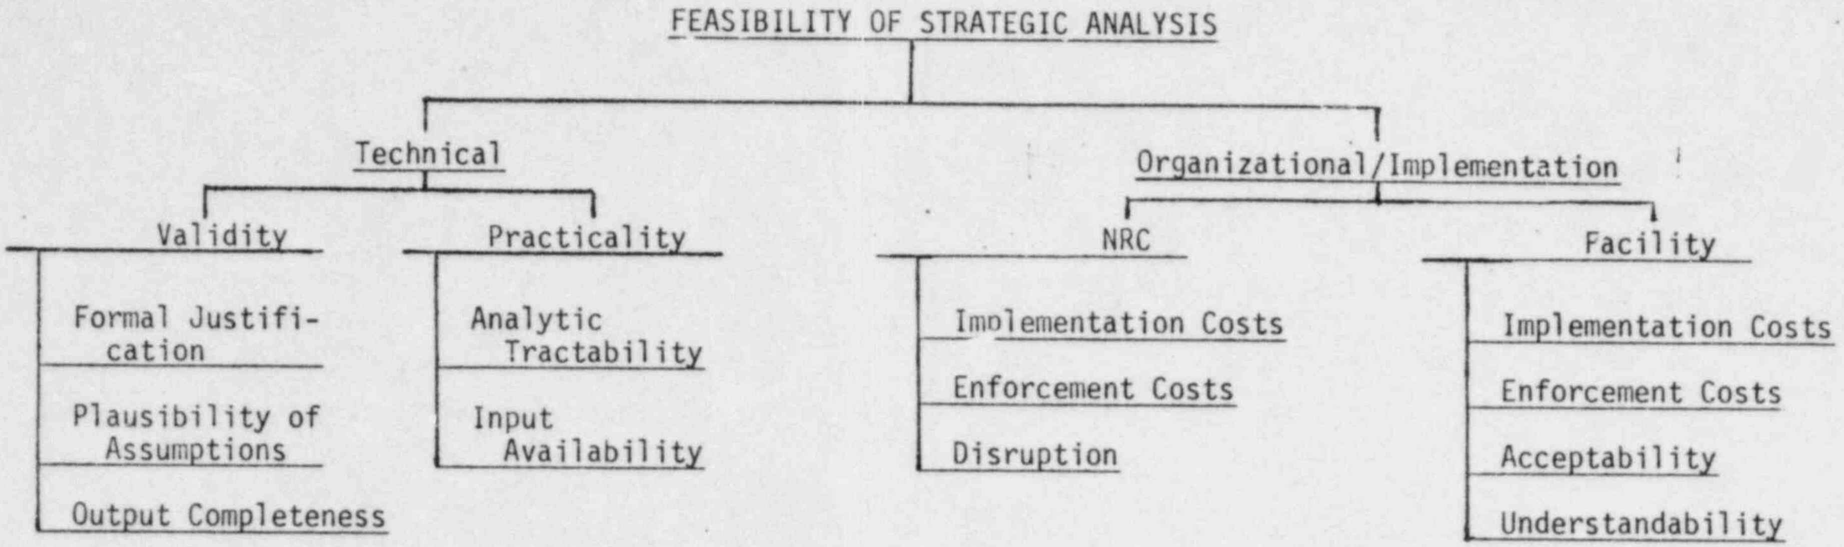

Exhibit 3-2

Technical Options for Strategic Analysis

- OPTIONS
  - Technical
    - Players
      - Multiple Players
      - No Diverter
    - Setting
      - Multiple Time Periods
      - Multiple Facilities
    - Defender's Strategies
      - Fixed Versus Variable Versus Imperfectly Known Fixed Threshold
      - Multiple Thresholds (and responses)
    - Diverter's Strategies
      - Upper Limit On Amount Diverted
    - Defender's Payoff Function
      - Accounting Accuracy Versus Detection Versus Deterrence
      - Functional Form (e.g., linear versus non-linear versus ordinal)
      - Social Versus Financial Costs
      - Adjustable to Individual Facilities?
    - Diverter's Payoff Function
      - Zero Sum Versus Non-Zero Sum
  - Organizational

Exhibit 3-3

Organizational Options for Strategic Analysis

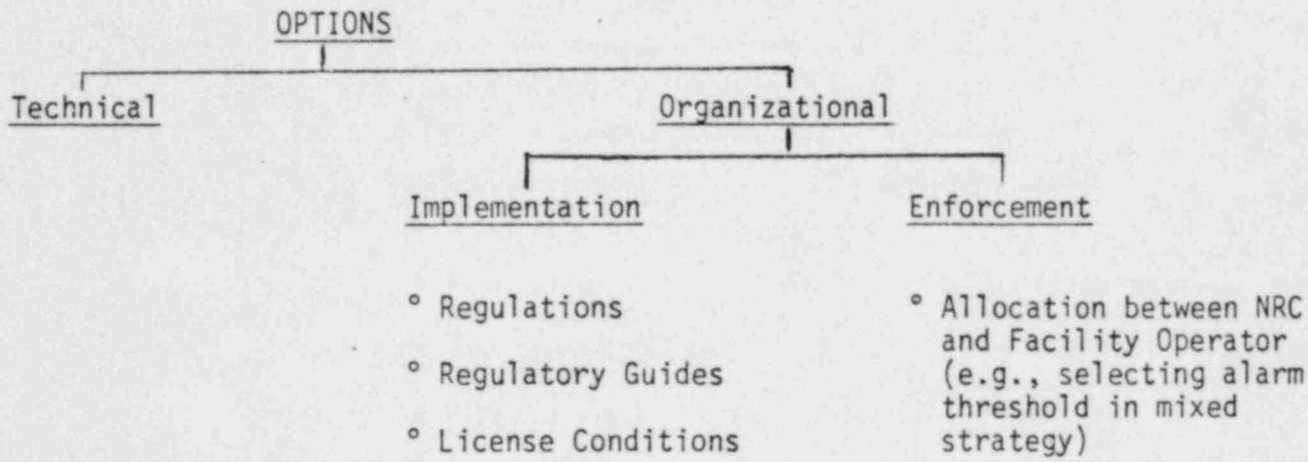

Exhibit 3-4

Material Accounting Functions and Objectives

Functions Required	To Achieve		Safeguards Objectives
Timely Detection or Timely Verification	→	Prevention	Deterrence
Detection or Verification	→	Response	Deterrence (of long-term diversions)
Statistical Control	→	Assurance	

and response (in contrast with assurance) require that conclusions regarding the possible occurrence of a diversion be utilized in potential actions against the diverter (not just in actions to improve the measurement process). As a result, the importance of conditions of formal validity based on decision theory and strategic analysis is much heightened. Finally, verification (as opposed to detection) requires a capability to provide, as part of the method's output, an integrated assessment incorporating the outputs of the systems, such as physical security and material control, into the conclusions of material accounting.

3.4 Technical Feasibility

In subsequent sections of this report, technical and organizational options, respectively, for implementing strategic analysis will be reviewed in terms of the feasibility criteria in Exhibit 3-1. In this section, we discuss each technical criterion briefly, mentioning the technical options upon which it has particular impact.

3.4.1 Formal Justification. Formal justification refers to the prescriptive adequacy of an analytic technique as a method for drawing conclusions or recommending action in the face of uncertainty and competing objectives. Perhaps the most significant application of this criterion is to the comparison of current practice (based on classical statistics) and strategic analysis *per se* (as a general approach). The classical hypothesis testing method indicates when an observed inventory difference exceeds thresholds that would only rarely (e.g., 5% of the time) be exceeded by chance if there were no diversion. But unlike strategic analysis, no explicit justification is offered concerning the choice of a particular threshold for "alarm" or for taking action (e.g., clean-out inventory). Strategic analysis, by contrast, offers a systematically justified method for balancing the costs of a false alarm against the benefits of detection.

A second important application of this criterion is in the choice among specific strategic analysis models. For example, in NUREG-0290 the defender's alarm threshold is regarded as a fixed parameter within the solution of the game; recommendations regarding the threshold are made by optimizing defender payoffs with respect to that parameter. NUREG/CR-0490 shows that defender payoffs may be substantially improved by treating the alarm threshold as a strategic variable which is optimized within the game solution. (The latter approach, for that particular model, yields a random mix of thresholds as the optimal strategy.)

3.4.2 Plausibility of Assumptions. Analytically tractable models, with readily available inputs, typically involve a variety of simplifying assumptions. The issue raised by this criterion is whether those assumptions are adequately realistic within the intended scope of application of the model, either because they are strictly true or because when they fail to hold, model output and defender payoffs are insensitive to violations.

Among the assumptions incorporated in the NUREG/CR-0490 model, and to be discussed in greater detail in Section 4.0, are the following:

- There is at least one, and only one, potential diverter;
- Diversions take place within a single inventory period;
- Diversion attempts are focused on a single facility;
- Diverter payoffs are exactly opposite to the defender's (the zero sum assumption);
- The cost of diversion is linear in the amount diverted;
- The diverter has not interfered with records or with the measurement process.

3.4.3 Output Completeness. This criterion refers to the extent to which desired information is included within the output of a model. For example, NUREG/CR-0490 model posits a single alarm threshold, hence, can make recommendations regarding only a single action alternative (e.g., inventory or no inventory). A more complete model might specify a variety of thresholds corresponding to the severity of the situation and the disruptiveness of the recommended action. The NUREG/CR-0490 model is also incomplete in not providing an integrated output capable of incorporating information that might, in some circumstances, be available from other components of the safeguards system. Finally, in many instances non-zero sum strategic models fail to specify a stable optimal action alternative at all, or specify several "optimal" actions.

3.4.4 Analytic Tractability. To be successful, a model must not only be valid (i.e., formally justifiable based on realistic assumptions, providing appropriate output), but it must be practicable. Analytic tractability is relative to the computational resources available and to the urgency with which a solution is required. Nonetheless, the tradeoff between validity and practicality can be severe. Generalizations of strategic analysis which drop the simplifying assumptions mentioned previously entail a cost in workability; e.g., provision for multiple or no diverters, multiple defenders, multiple time period, and multiple facilities, non-linear utility functions, and diverter strategies which include interference with the measurement process.

The same is true of models which include interference with the measurement process. The same is true of models which provide more complete output; e.g., multiple action thresholds and output integrated with information from other safeguard systems. Each recommended extension of the strategic analysis model must involve a careful weighing of costs and benefits.

3.4.5 Input Availability. A second aspect of practicality is the practical obtainability of inputs at the level of reliability and accuracy required by a model. Inputs for which obtainability is an issue include:

- Errors in the estimation of inventory differences in the absence of diversion, including such potential sources of error as interference with records, holdover of material, and bookkeeping errors.
- The probability of no diverter and the probabilities of various different types of diverter groups.
- Payoff functions for different types of diverter groups.
- A comprehensive listing of plausible diverter strategies.

3.5 Organizational/Implementation Feasibility

In addition to the technical criteria, feasibility must also include criteria reflecting the costs (not just financial) of implementing strategic analysis. Even if strategic analysis is shown to be both valid and technically practical, the effort involved in implementing it may be such that the potential safeguards gains are not sufficient to offset the organizational costs required. In the following, we briefly describe the organizational/implementation criteria that must be considered. Again, various options will be evaluated with respect to these criteria in the subsequent section.

3.5.1 NRC Implementation Cost. This criterion refers to the monetary costs associated with developing and implementing a strategic analysis model (or models). In particular, it includes personnel costs for any NRC staff that would be involved in the effort and costs for outside contract research. Obviously such costs would not be incurred if strategic analysis were not pursued. The extent of these costs will vary depending upon the strategic analysis model developed (particularly research costs) and the type of instrument (e.g., regulation, Regulatory Guide) used for implementation.

3.5.2 NRC Enforcement Costs. These costs are associated with NRC enforcement of the use of strategic analysis. They reflect differences between costs associated with enforcement of strategic analysis and costs of enforcing the currently used material accounting procedures. Both routine inspection costs and costs of responding to alarms are included. The latter costs will depend in part, for example, on the number of false alarms produced by using strategic analysis versus current practices.

3.5.3 NRC Disruption. This criterion refers to the "hassles" involved in changing or expanding from the current procedures to a new procedure based on strategic analysis. It goes beyond financial effects to incorporate what might be called organizational stress. It includes, for example, any difficulties in negotiating with facility operators, the non-monetary administrative difficulties of interacting with and responding to the public, and any requirement to interact with Congress, as well as internal NRC discussion and and negotiation. This criterion is affected primarily by the instrument used to implement strategic analysis if it is implemented.

3.5.4 Facility Implementation Costs. These costs are those incurred by the facility operator in implementing strategic analysis. They include primarily personnel costs, and costs associated with training staff in new procedures, and costs resulting from the need to collect and/or analyze new information (e.g., costs for incorporation into the payoff function).

3.5.5 Facility Enforcement Costs. The facility will also incur costs in the on-going use of strategic analysis which may differ from costs associated with current practice. These costs include both routine costs such as record keeping and data collection, and costs associated with responding to alarms.

3.5.6 Facility Acceptability. While the NRC clearly does not have to have the voluntary agreement of facility operators to implement strategic analysis, industry support is still desirable and would probably lead to more efficient implementation and more effective use of strategic analysis. Thus, this criterion is included.

3.5.7 Understandability by Facility. This criterion is not independent of the previous one, but separate consideration is desirable to point out potential difficulties. Again strategic analysis need not be fully understood by the facility operator to be implemented, but if it is not, additional guidance from the NRC may be required as may be training programs. This criterion is included because it is particularly relevant to certain options (e.g., use of mixed strategies).

4.0 EVALUATION OF STRATEGIC ANALYSIS OPTIONS

As noted in section 3.0, many options for implementing strategic analysis are available, including both alternative game-theoretic models (technical options) such as the NUREG/CR-0450 model, and alternative NRC implementation approaches. The feasibility of strategic analysis as a safeguards system is highly dependent upon which of these options are selected for development and possible implementation. Therefore, we have chosen to evaluate the feasibility of independent parts of a total strategic analysis system using the framework described above. This evaluation will provide the necessary analysis to suggest which parts would be combined to produce the most feasible strategic analysis approach, and whether such an approach is to be preferred over the current practice.

Our feasibility evaluation here follows the organization of options presented in Exhibit 3-2. The following discussion describes these options in more detail and evaluates their feasibility using the criteria described in Section 3.3. In addition, more technical detail regarding the technical options and their feasibility can be found in the Appendix, along with a discussion of relevant literature. After assessing the feasibility of the various options, both technical and organizational, we again apply the feasibility criteria to compare strategic analysis more generally with current practice.

4.1 Technical Options

The technical options are divided into the common elements of game-theoretic models: players, setting, strategies, and payoff functions. Options available for each of these elements are described and evaluated below.

4.1.1 Players. A primary technical question is whether the material accounting process should be modeled using a two-person game-theoretic model or a multiple-player (more than two players) model. The previously developed model is a two-person game with a single diverter and a single defender. Initial consideration of the nuclear safeguards context might suggest that such a simplified game does not adequately represent this complex context. In particular, some might argue that both the NRC and the facility operator should be incorporated because their objectives do not necessarily coincide as is implied by a single defender. In addition, multiple diverters may be present, possibly with different objectives and therefore payoff functions. The public, or even different segments of the public, could also be included explicitly as players.

Several of the feasibility criteria distinguish between two-person and multiple-player games. While both have adequate formal justification, the assumptions required for the two-person game are somewhat less plausible. All other things being equal, multiple-player games clearly have the ability to more accurately model the safeguards context. The two-person game, however, is generally an adequate representation of the context, if payoff functions are properly developed. For example, the defender's payoff function can include potential effects on the public that make the need for the public to be a separate player less important. A similar argument suggests that the facility operator's payoffs can also be incorporated into the defender's payoff function. While multiple diverters is clearly a credible assumption, the use of a two-person, zero sum game in which the diverter's objectives are assumed to be diametrically opposed to the defender's represents a worse-case, and therefore conservative assumption. (This is discussed more extensively in the subsequent discussion of the diverter's payoff function.)

Since the plausibility of assumptions is the only criterion on which multiple player games are more feasible than two-player games, the effects on other criteria indicate that a two-person game model is more feasible overall. With respect to technical criteria, two-player games are to be preferred because of output completeness, analytic tractability, and input availability. One of the important output capabilities of any such model is sensitivity analyses. Sensitivity analyses, at least certain types, are relatively straight-forward with two-player games. They can be very difficult, on the other hand, with multiple-player games. In addition, multiple-player games may be quite difficult to solve, requiring relatively onerous numerical methods. Finally, multiple-player games obviously require considerably more inputs than two-player games, including the strategies and payoff functions for each player. Given the difficulties in developing such inputs, particularly the payoff function, this is a critical advantage for two-player games.

Two-player games are also more feasible in terms of the organization/implementation criteria. Because of their added complexity, multiple-player games would have higher implementation costs for both NRC and facility operators. While enforcement costs would not be appreciably different for the two models, a two-player game is likely to be more understandable to facility operators. This will help with the acceptability, particularly if the facility operator can see its own values (costs) represented in the defender's payoff function.

A second option with respect to players is whether the diverter should be considered surely to be present. The NUREG/CR-0490 model assumes, as would any traditional two-player game, that the diverter is indeed present with intent to obtain SNM. This may be an overly conservative assumption. In a model which gives no "credit" to the capabilities of other safeguards, this further degree of conservatism may be excessive rather than prudential. Modifications could be made in the strategic analysis model to introduce a probabilistic representation of the diverter's presence, i.e. the diverter can be assumed to be present with some probability p. While this is a more realistic representation, the determination of p would be difficult (reducing input availability) and the assumption that p is an exogenous quantity independent of the defender's strategy would itself be controversial. The criticality of precision in the determination of this parameter, however, could be assessed rather easily using sensitivity analysis.

This option would have little effect on organizational/implementation criteria. Therefore, with respect to overall feasibility, if strategic analysis is pursued, this option could be investigated at least to the point of developing the appropriate model modifications and performing the sensitivity analyses.

4.1.2 <u>Setting</u>. The primary options available in modeling the setting are to extend the single time period, single facility game to multiple time periods and/or multiple facilities. Neither of these options particularly affects the formal justification of the model, but both make the underlying assumptions more plausible. A multiple time period game in which the diverter's and defender's actions at any time are linked to actions at previous times better represents the true situation. For example, it is reasonable to assume that the amount of SNM a diverter attempts to obtain depends upon whether any SNM has been previously obtained. Similarly, actions at one facility may depend upon occurrences at other facilities. An organized diverter group could attempt to spread its efforts to obtain SNM across several facilities, thus affecting its objectives and actions at any single facility.

The advantages of multiple time period and multiple facility models in terms of their plausibility of assumptions and realism are offset to some extent by some increases in the difficulty of developing and solving the models (NRC implementation costs and analytic tractability). Both types of models should generally prove solvable (with heightened effort) if the comparable single period, single facility model is solvable. The additional costs involved in developing such modifications do not appear to be large, suggesting that they should be investigated. If strategic analysis is to be further developed, however, a moderating strategy could include first the development of the basic single period, single facility model, and subsequently its extension after initial testing.

4.1.3 <u>Defender's Strategies</u>. Options to be considered for the defender's strategy include the use of variable (possibly mixed random) alarm thresholds, the withholding of information regarding thresholds, and the use of multiple alarm thresholds with appropriately varying responses. In a context such as the SNM safeguards context, information regarding other actors' strategies is clearly valuable. The more a diverter knows about the defender's strategy, the better the diverter's chances are of obtaining desired objectives. One way to deny information to the diverter is for the defender to use a variable strategy--one that is not the same in each situation even though the situation is comparable (e.g., the same ID is obtained). In particular, a randomized mixture of strategies in which each of several strategies, s_i, may be selected with probabilities, p_i, provides a variable strategy that obviously denies the diverter information regarding the defender's strategy since even the defender does not know which strategy will be implemented until the random selection process is completed. Such a mixed strategy is typically the optimal strategy in game-theory models. The advantages of including such randomization in strategies were demonstrated by the modification to the original game theory model (NUREG-0290) made in NUREG/CR-0490 which included the randomized selection of the alarm threshold.

With respect to formal justification, use of mixed strategies is to be preferred to use of only a pure (not mixed) strategy. The player using a mixed strategy will always do at least as well in terms of the payoff function as if a pure strategy is used. This technical feasibility, however, must be balanced against effects on implementation criteria including enforcement costs and possible acceptability.

Although the use of a random device to determine what action for the defender to take may initially seem inappropriate and therefore appear to affect acceptability; a careful explanation of the advantages of mixed strategies and a demonstration of their improvement over pure strategies should be possible. This could be expected to reduce the effect of mixed strategies on acceptability to a negligible level.

On the other hand, the use of a mixed strategy will increase enforcement costs. Mixed strategies require that several different actions be possible. In the safeguards context, this would require that plans be made for each action, rather than for only a single action as required for a pure strategy. Development of these plans would lead to additional expense. The amount of this expense would depend on the number of possible actions included in the mixed strategy, and their similarity. The fewer and the more similar the actions (i.e., those that can benefit from some joint planning), the less the expense. However, few and similar actions also limit the advantages of mixed strategies since they produce less uncertainty in the diverter. The optimal number and type of actions to be included in a mixed strategy is one topic that should be addressed if strategic analysis is to be further developed by the NRC.

Some of the advantages achieved by using mixed strategies can also be obtained by using other methods to deny information regarding the defender's strategies to the diverter. For example, if a single fixed alarm threshold were to be used as a pure strategy, denying this information to the diverter would improve the defender's payoff. Such information could be classified, though that does not provide complete assurance that the diverter does not have the information. In general, regardless of the strategic analysis model used, as much information as possible should be kept from the diverter.

Another option is the use of multiple thresholds each with different responses required. Such a model would reflect the reality that different levels of IDs can lead to fundamentally different actions on the part of the NRC and the facility operator. This option would also provide a mechanism for incorporating other aspects of the total safeguards system into the strategic analysis model. For example, response actions could be based on physical security information as well as measured IDs.

This option represents a more realistic model of the safeguards problem, thus rendering the assumptions more plausible. The output is more complete because the response actions are more comprehensively determined and are used explicitly in this model. Multiple thresholds will, however, require additional input. In particular, the responses must be adequately defined and the costs associated with them must be determined for inclusion in the payoff function.

For these reasons, the multiple threshold model would also be somewhat more expensive to develop and implement. The added expense does not appear to be excessive, however, since it involves primarily the development of needed input to the payoff function. This input, which involves primarily NRC and facility operator costs for various actions, is a relatively inexpensive part of the payoff function to develop, relative to the scale of possible societal costs resulting from successful diversion. Furthermore, analysis of alternative response levels and their costs is a natural NRC activity whether or not strategic analysis is adopted. Therefore, the added expense is probably justified.

4.1.4 <u>Diverter's Strategies</u>. The previously developed model assumed an upper limit on the amount of SNM that a diverter would attempt to obtain. While this assumption may approximately represent the situation where a diverter would like to obtain SNM up to a fixed amount (e.g., the amount necessary for a bomb) above which additional SNM may have little value, it may be unnecessarily restrictive. This fixed limit assumption improves the analytic tractability of the model, making it easier to solve. At this point, it is unclear how much more difficult to solve a model would be without this restriction. If, for numerical convenience, such an upper limit is used in future models, model developers and potential users should be suspicious if the optimal action of the diverter is driven to this limit or if the model's solution is particularly sensitive to the amount of the limit. For example, the optimal strategy in the previously developed model was not sensitive to the assumed upper limit.

4.1.5 <u>Defender's Payoff Function</u>. In the previously developed model, the payoff function was one of the most controversial parts of the model (see e.g. NUREG/CR-0950). This is not surprising since it is undoubtedly one of the most difficult parts of the game to formulate in the safeguards context. Rather than discussing particular difficulties with the specific payoff function in NUREG-0290 and NUREG/CR-0490, that have been reviewed in NUREG/CR-0950, here we address the feasibility of more general options in the development of the payoff function.

In developing any payoff function, the specific objectives of the use of strategic analysis must be determined. In particular, is strategic analysis to be used for accounting accuracy, timely detection, deterrence, assurance, or some combination of objectives with some specified relative emphases. Different answers to this question imply different payoff functions. For example, if accounting accuracy is the only objective of the material accounting system in which strategic analysis could be deployed, then strategic analysis is probably inappropriate. Accounting accuracy does not involve a diverter and therefore should not be in the payoff function even if accounting accuracy is part of the relevant safeguards objectives. If a deterrence objective is taken to imply that the defender would be hurt by an attempted diversion, not just an actual diversion, then diversion attempts even with no loss of SNM should be considered in developing the defender's payoff function.

Once the objectives of the strategic analysis have been determined, a specific payoff function representing these objectives can begin to be developed. Generally, such a function will take the form of a multiple objective or multiattribute utility function (e.g., Keeney and Raiffa, 1976). The development of such functions has proved to be feasible in numerous social utility contexts (Keeney and Raiffa, 1976), but whether such a function can be developed in the safeguards context depends on several issues.

In general, a linear payoff function (one in which all terms are linear function of certain variables) will be the most analytically tractable and the easiest on which to perform sensitivity analyses. Sensitivity analyses can be performed, for example, in a two-person zero-sum game without re-solving the model. Alternatively, nonlinear functions may more accurately represent real values, and therefore be more plausible. For example, if the function contains a term representing the cost of searching for SNM potentially diverted, and if this cost is not proportional to the amount of SNM estimated to be missing (e.g., the cost goes up at a faster rate than the amount missing), then a nonlinear term is required. Thus, in order to adequately represent defender values (i.e., the public as represented by the NRC), a nonlinear function is almost certainly needed. Such a function comes with some negative effects on output completeness in terms of sensitivity analyses and on analytic tractability. As long as the model used is a two-person zero-sum model, however, the game is solvable at least by using numeric methods and certain types of sensitivity analyses are possible.

An alternative to use of linear and nonlinear payoff functions is the use of ordinal functions. These functions simply provide a preference rank order on the value of the outcomes of various actions rather than a numerical representation of the value. To solve a game with an ordinal payoff function requires that some level of confidence (e.g., 90%) be selected a priori. The solution then identifies a strategy (possibly mixed) that assures achievement of the highest possible outcome-rank attainable with the selected level of confidence regardless of the opponent's strategy. (See Appendix for additional discussion and relevant literature.)

The advantage of using ordinal functions to represent payoffs is that the required input may be easier to obtain, particularly to the extent that judgment must be used. It may be easier and more acceptable for policy makers, using whatever additional expertise is available, to rank order different outcomes than it would be to specify certain tradeoffs (e.g., search costs versus social costs of bomb explosion), even though such tradeoffs might be implicit in the rank ordering. Also, no assumptions are required on what the opponent's preference ranking is, or how it relates to one's own.

The completeness of output of the ordinal model would be less than for a linear or nonlinear model. Sensitive analysis techniques have not been developed for such models. Also, the optimal strategy only achieves the highest rank with some probability not with certainty, and no attention is paid to how bad, bad outcomes are or the chances they will occur.

Given the current status of research on ordinal game theory models, they should be pursued only as a research adjunct to the development of a nonlinear model. The research effort could prove to be very valuable, if appropriate payoff functions cannot be otherwise developed.

Probably the most critical aspect of the payoff function is the incorporation of social costs. This, of course, is also likely to be the most difficult. For example, some measure of the social cost of not recovering diverted material is needed. Such a measure will most likely be based, at least in part, on expert judgment. It will require some determination of what will happen if a diverter has some amount of SNM. These inputs will be difficult to obtain since most people will not be willing to be responsible for such a determination. However, these inputs may possibly best be obtained using a group of experts with subsequent review and concurrence by other experts and policy makers.

Perhaps even more difficult inputs are the parameters which would be needed to specify tradeoffs between various terms in the payoff function. For example, even if only implicitly, somehow the social effects of detonating a bomb must be equated with search costs. Again, such inputs are primarily subjective and should be specified by senior policy makers. Although policy makers are willing to make such tradeoffs in some social contexts (e.g., Keeney and Raiffa, 1976), in a situation as controversial and open to public reaction as the safeguards context, it is not clear that such tradeoffs can be obtained. The NRC has previously attempted to specify similar types of tradeoffs in the development of safety goals. While some quantitative tradeoffs were specified, e.g. radiation exposure related to dollars, the process was difficult, time consuming, and expensive; and the results were not comprehensive in terms of all possible safety effects and remain controversial. One can appear to finesse such controversies by seeking to optimize a payoff not involving "big bang" or "big threat" costs, subject to a limit on the allowable "missed alarm" probability. But then the controversy has merely been concealed within the issue of how that limit should be set.

4.1.6 Diverter's Payoff Function. The development of strategic analysis in the safeguards context to date has focused on zero sum game where the diverter's payoff function is assumed to be diametrically opposed to the defender's. This is a conservative assumption that allows the defender to do the best possible (in terms of the payoff function) against the worst possible opponent. The price for such protection is that such a model may not be realistic in its assumptions and as a result may lead to a lower expected payoff to the defender than would a model that did not include the zero sum assumption.

Non-zero sum situations are by no means implausible. For example, the diverter's chief fear may well be apprehension. But the defender's chief goal need not be apprehension of diverter, but rather deterrence of diversion. (If defender succeeded totally in the latter, the defender would, quite happily, not have any more diverters to apprehend!) Thus, the defender may exploit the diverter's tendency to rationally pursue his goal (avoid apprehension) in order to achieve his own goal of reducing the frequency of diversions. Deterrence as an objective is impossible in a zero sum game.

Exhibit 4-1 illustrates this argument with a highly simplified non-zero sum games. We suppose that the defender can either conduct a clean-out inventory (I) or not (I); and the diverter may either divert material (D) or not (D). Defender and diverter preferences for outcomes are ordered as follows:

<u>Defender</u>	<u>Diverter</u>
$\overline{ID}$	$\bar{I}D$
$I\bar{D}$	$I\bar{D}$
ID	$\bar{I}\bar{D}$
$\bar{I}D$	ID

Note the defender most prefers the outcome of no inventory and no diversion ($\bar{I}\bar{D}$), while the diverter most dislikes the situation in which the diversion is detected (ID). We consider two ways of playing this game:

(i) Defender plays a zero sum game based on his own payoffs, ensuring himself a minimum payoff against a rational and totally hostile adversary. This game has a saddle point (always inventory) with a security level of zero.

(ii) The defender seeks an equilibrium pair of strategies in the full non-zero sum game. Such an equilibrium for the defender involves a mixed strategy of inventorying with probability $P_I = 1/(1+a+c)$ and a game value of $db/(1+b+d)$ for the defender. Thus, defender can improve payoffs over his security level by playing against the actual, rather than the hypothetical diverter. However, choosing the correct strategy requires the defender to know the values of a and c in the diverters payoff function (or at least in their sum). Misestimation of these quantities could leave the defender worse off than in (i).

It might be expected that the use of non-zero sum models would reduce the sensitivity of payoffs to detection threshold location. But the presence and magnitude of any such effect would have to be examined for particular types of models.

To change the zero sum assumption, however, would require the explicit modeling of diverter objectives. Even with the zero sum assumption, diverter's objectives must be considered. For example, although the existing strategic analysis model represents the defender's payoffs, not the diverter's, those payoffs are a function of the defender's perception of potential adversaries. For instance, the model has a term representing

Exhibit 4-1
Simplified Non-Zero Sum Game

Defender

Diverter	I	$\bar{I}$
D	-a,0	1,-d
$\bar{D}$	c,b	0,1

Note: First entry is Diverter payoff, second is Defender payoff.

$0 < a,b,c < 1$

the value to the defender, of recovery of the material diverted. This value, if it is to include the value of withholding material of strategic value to an adversary, will depend on the defender's perception of the adversary's motivations, intentions with regard to utilization of the material, capabilities to carry out these intentions, and the defender's estimation of the potential consequences resulting (directly or indirectly) from this implementation. Thus, adversary characteristics and motivations are at least implicit in any reasonable choice of defender payoff function.

Aside from their presence in the defender's payoff function, adversary characteristics and motivations may be pertinent in evaluating the appropriateness of certain modeling choices. For instance, the zero sum assumption implies that adversary payoffs are opposite to those of the defender. Information on the credible spectrum of adversaries could help determine the degree of conservatism inherent in this (zero sum) assumption. Unfortunately, despite the need for this understanding to design, implement, and assess a strategic analysis safeguards system, it is difficult or impossible to accurately and authoritatively assess these adversary characteristics and motivations. Potential adversaries span a tremendous range, from trivial to incredible. None of the generally utilized assessments of adversary characteristics and motivations proposes an unbounded adversary profile. Clearly, given any safeguards system one could postulate an adversary with sufficient capabilities to breach the system, and sufficient motivation and capabilities to utilize the material so as to create great harm. What, then, is acceptable as a design basis safeguards threat? Several studies have been performed and used to attempt to answer this question. (Most of those used and recognized by NRC are listed in the bibliography of the Appendix). The methods used in these studies help to illuminate the difficulties inherent in projecting adversary characteristics and motivations.

Even if payoff functions for a non-zero sum model could be developed, they would create certain additional technical difficulties. Non-zero sum games can have multiple solutions which do not necessarily give the same payoff to the defender. Finding the solution or solutions for a non-zero sum game is much more difficult and finding all the possible solutions can be very difficult (expensive and time consuming). And once multiple solutions are identified there is no clear cut process by which to select a single solution for implementation.

A more general point is that potential diverters may vary in their motivations and objectives, and a variety of zero sum and non-zero sum models might be required to model such diversity. Strategic techniques described by Harsanyi (1967) may be useful in this context, but the difficulty of applying game theoretic techniques successfully would be expected to grow considerably. The gain in simplicity and solvability achieved by using a single zero sum model should be carefully weighed against the expected loss in payoffs due to excessive conservatism.

In addition to the technical difficulties, implementation of a non-zero sum model would be considerably more expensive because of the need to develop a separate payoff function for the diverter. In addition, the reduction in conservatism from the zero sum model as well as the availability of multiple solutions may make the model less acceptable to facility operators. This loss in acceptability may be offset somewhat by an increase in acceptability produced by the relatively higher payoff of using a non-zero sum game and decreased risks to the public.

4.2 Organizational Options

As with the development of a strategic analysis model, there are several options with respect to how strategic analysis would be implemented and enforced by the NRC and facility operators. These options do not affect the technical feasibility of strategic analysis, so in terms of the framework described in Section 3.0, they are evaluated only with respect to the organizational/implementation criteria.

While we recognize that the effort called for in this contract did not explicitly request consideration of how strategic analysis could be implemented and enforced organizationally, the objective of this study implicitly suggests this should be considered. In particular the feasibility of "applying strategic analysis in a regulatory framework" and "assuring that strategic analysis can be effectively enforced" cannot be assessed without this type of consideration. Determining that strategic analysis is technically feasible would be pointless unless the NRC has some feasible instrument for implementing and enforcing it.

4.2.1 Implementation. Various instruments that the NRC could use to implement strategic analysis were described in Section 2.0. Three deserve explicit consideration: regulations, license conditions, and regulatory guides.

With respect to implementation costs for the NRC, regulations are clearly the most costly, followed by license conditions, with regulatory guides being considerably less costly. The development of regulations and the necessary public review and comment would be a relatively costly and time consuming process requiring several staff-years of effort. Recent experience with the development and promulgation of regulations at the NRC indicates that such an effort would be spread over several years. This is in addition to any research time and support required for the development of the specific model(s). (We have assumed here for the purpose of a comparative evaluation among the three options, that the research required would not be affected by the instrument chosen for implementation.)

Use of license conditions would shitt NRC costs from developing regulations to reviewing license applications and/or amendments. Given the small number of Category I facilities currently licensed or expected to apply for licenses in the near future, the number of such reviews would not be large (i.e. probably no more than a half dozen). Some guidance would have to be provided to facility operators so that use of license conditions might also include the development of a regulatory guide. Once the necessary research has been conducted, a regulatory guide could probably be developed with less than one staff-year of NRC effort. Review of all immediate applications could probably also be accomplished with less than a staff-year of effort.

NRC enforcement costs should not be significantly affected by the implementation instrument. Once the appropriate model is implemented, its costs for enforcement should be comparable, regardless of how it was implemented. The exception to this could be the use of only a regulatory guide if a licensee chose to try to meet regulatory requirements using a procedure other than that suggested in the Guide. This would then require the NRC to have alternative enforcement procedures, thereby increasing enforcement costs. Experience with other regulatory guides suggests, however, that this situation is very unlikely to occur.

The effects on NRC disruption are similar to costs. Development of regulations produces the most disruption through the public comment process and the effort to negotiate some agreement among groups with diverse viewpoints. License conditions reduce the disruption by making any negotiations primarily between the NRC and facility operators. Regulatory guides are the least disruptive, probably requiring some negotiations with facility operators, but a lesser degree.

For the facility operators, implementation costs would vary only slightly among the options. Regulations and license conditions would be somewhat more expensive because of the formal approval required and the accompanying paperwork needed. Costs involved in obtaining any needed data and specifying appropriate parameter values would be similar for each option.

Similarly enforcement costs would vary little. Acceptability, however, could be better with regulatory guides because they allow the operator somewhat more freedom. No particular differential effects on understandability are obvious.

4.2.2 Enforcement. Enforcement options refer primarily to how enforcement efforts would be allocated between the NRC and facility operators. One option in particular could be used to improve facility operator acceptability--have the NRC make any random selections needed for a mixed strategy. This may help facility operators to feel less capricious in their response actions and may also make safeguards based on strategic analysis less susceptible to inside adversaries. This option would also put enforcement of safeguards under more NRC control.

Numerous other options for enforcement allocation could also be considered. The primary effect, however, of any allocations will be on NRC and facility operator enforcement costs. Specific cost data are not presently available. From a broad, but reasonable, perspective, the distribution of these costs may not very salient, particularly in determining the feasibility of alternative enforcement options, since most Category I facilities are operated under contract to the U.S. Government with operating costs being reimbursed. One practical implication is, however, that funding for such costs may be easier to obtain for the facility operators than for the NRC.

4.3 Strategic Analysis Versus Current Practice

With the evaluation of the feasibility of strategic analysis alternatives presented above, it is now possible to evaluate the overall feasibility of using strategic analysis by comparing it with current practice. It is important to note at the outset of this comparison that strategic analysis and the approach to material accounting based on classical statistical hypothesis testing are not intrinsically competitive and in practice could address different parts of the material accounting functions. This potential complementarily has been previously emphasized by the Peer Review Group (NUREG/CR-0950). Thus, while for evaluation purposes here we compare strategic analysis with current practice, if strategic analysis were to be implemented it would probably not displace statistical hypothesis testing, but instead would be implemented in addition to statistical hypothesis testing.

The statistical hypothesis testing approach really has nothing to say about whether a diversion has occurred given a particular ID. Rather, it primarily provides evidence regarding whether the process is in control, i.e. whether the ID is within the range of variation that is to be expected given errors in measuring inventory. Strategic analysis, on the other hand, proceeds as if a diverter were present, and optimizes the defender's actions against this presence.

Current practice grows out of the statistical hypothesis testing approach. It has, however, gone beyond this classical approach to incorporate practical considerations. In particular, because the measured ID is assumed to be true ID plus some error, it is necessary to estimate the error associated with determining ID. This error can arise from many sources of which only measurement error is currently known with any reasonable degree of certainty.

Thus, a classical hypothesis testing approach based on only measurement error, which is an underestimate of total error, is likely to produce too many false alarm (measured ID is above threshold when true ID is zero).

Current NRC practice as spelled out in 10 CFR 70 and NRC management action (Page, 1974) attempts to alleviate this potential false alarm problem by establishing limits on the limit of error for inventory differences (LEID). The LEID is meant to define a 95 percent confidence limit around ID measures and is consistent in spirit with the statistical hypothesis testing approach, implying that when ID=0, the probability that ID will exceed LEID (false alarm) is .05. 10 CFR 70 defines limits on the LEID (LEID-limit) based on specified percentages of additions to or removals from the material in process. The so-called "Page letter" (Page, 1974) then specifies actions that facility operators are to take based on comparing ID with LEID and LEID-limit.

In order to compare the feasibility of strategic analysis to this current practice, we specify some details of the strategic analysis model. These details are based on the evaluation of the feasibility of various strategic analysis options discussed above. While certain comparisons are dependent upon the specific model assumed, much of the feasibility is generic, though feasibility limitations discussed above should be kept in mind. Specifically, strategic analysis as we envision its operational form includes:

- only two players,
- capability to assume diverter is not always present or free to act,
- multiple time periods,
- multiple facilities,
- mixed strategies,
- multiple thresholds,
- upper limit on the amount diverted,
- nonlinear payoff function based on realistic objectives including social costs, and
- the zero sum assumption.

The following discussion assesses the feasibility of strategic analysis with respect to the criteria presented in Section 3.0.

4.3.1 <u>Formal Justification</u>. Both statistical hypothesis testing and game theory are well-based theories. Current practice, however, as noted above has been extended beyond standard hypothesis testing to incorporate limits of LEID. Thus, some of its formal justification has been reduced. Current practice is based on maximizing the number of correct detections of loss for a fixed false alarm rate, with no explicit consideration given to the relative costs associated with false alarms, correct detections, and undetected losses. Strategic analysis, however, takes explicit account of these costs in determining appropriate actions. Thus, in this regard, strategic analysis is considered to have a somewhat better justification although current practice does have a theoretical basis.

4.3.2 Plausibility of Assumptions. The primary assumption underlying current practice is that measured ID is equal to the true ID plus error and that the error is measured at least approximately by the LEID-limit. While the measurement error in the error term is known reasonably well, other components of error are not and the LEID-limit is a rather arbitrary term. NRC has, however, undertaken research to improve the determination of the error term that can eventually be used to improve the assumptions regarding error. In addition, the current practice does not explicitly consider the presence of a diverter. While in most situations with an ID, this assumption may be reasonable, in any circumstance where a diverter is present, it may be very misleading and result in high potential costs.

Strategic analysis, as typically applied, would assume that a diverter is always present, a very conservative assumption that may lead to excessive costs. This difficulty can be alleviated to some extent, by the specific options discussed above. With the option to incorporate the probability of a diverter being present, the strategic analysis model can include the most appropriate assumption regarding the presence of a diverter.

The strategic analysis model also assumes that payoffs can be measured at least to a reasonable degree explicitly by a multiple objective utility function. In addition, the zero sum assumption treats the diverter's payoffs as diametrically opposed to the defender's. As discussed above, this latter assumption, while clearly not realistic, probably is an adequate and feasible representation. The former assumption, while reasonable in theory based on multiattribute utility theory, has some practical drawbacks as is discussed below for input availability. Thus, with respect to plausibility of assumptions, strategic analysis again appears to be feasible and to be preferred to current practice.

4.3.3 Output Completeness. Current practice, while specifying actions for various ID levels (Page, 1974), does so on a rather arbitrary basis. Strategic analysis, however, specifies actions based on optimal considerations with respect to the costs of different actions and possible outcomes. Both approaches allow for sensitivity analyses although because of the relative simplicity of current practice, sensitivity analyses are much easier than for strategic analysis. Strategic analysis also, in theory, could include some consideration of the uncertainty with respect to diverter objectives, although the zero sum assumption needed to help ensure feasibility would inhibit such modeling. Current practice does not explicitly consider possible diverter objectives. Neither approach as presently formulated, takes into account other aspects of the safeguards system. On the basis of these factors, strategic analysis appears to be neither more nor less feasible than current practice in terms of output completeness.

4.3.4 Analytic Tractability. Current practice is obviously sufficiently tractable to be used, having been in practice for many years. A strategic analysis model, with the features described above is also analytically tractable, though requiring considerably more computation than current practice. Given the current state-of-the-art in solving games and the capabilities of existing computing facilities, this additional computation should prove a relatively negligible effect.

4.3.5 Input Availability. The primary inputs needed for current practice are estimates of the measurement error in IDs needed to derive LEIDs and the throughput of facilities need for LEID-limits. Both inputs are readily available. However, as noted above, current practice would be much more soundly based if it were based on including all sources of error, not just measurement error, in LEIDs and not using LEID-limits. Current research supported by the NRC is developing these needed error estimates.

This error estimate (or more precisely, the distribution of ID in the absence of diversion) is also needed for strategic analysis. The need for this input, thus, does not distinguish between strategic analysis and current practice, so strategic analysis can also benefit from current research.

For strategic analysis, the major technical obstacle to feasibility is the need for a payoff function. While the outputs of strategic analysis are more directly relevant to decision making, they are, of course, only as good as the extent to which the payoff function represents the actual values of the defender. Given the complexity of these values for the defender (e.g. the NRC representing the public), an explicit formulation of the payoff function is extremely difficult.

Among the NRC staff and outside experts such as the Peer Review Group (NUREG/CR-0950), there is near-unanimous consensus that the payoff function formulated in NUREG-0290 and NUREG/CR-0490 is not adequate. This viewpoint, however, does not all all imply that strategic analysis is not feasible. Rather, it only suggests that this particular part of strategic analysis is critical and should be given special consideration in assessing feasibility.

The problem faced in developing an appropriate payoff function is how to measure the value of non-quantitative outcomes such as the reaction of society to a nuclear incident; and how to trade off such values against quantitative measures such as financial costs of searching for potentially missing SNM. The development of such a function will require judgmental inputs from experts and from policy and decision makers. It will also require a broad acceptance of these judgmental inputs. Techniques have been developed to obtain such needed judgments in the context of multiattribute utility theory and have been successfully applied to develop functions measuring social values (Keeney and Raiffa, 1976). Thus, we can say with considerable assurance that experts and decision makers can make the needed judgments. Whether they will is another question, and the acceptability of such judgments is far from assured, and can only be answered through the actual process of obtaining concurrence. Progress along these lines would obviously be very valuable to NRC safeguards systems analysis in general, not only to a strategic analysis approach.

The recent development of the "safety goal" by the NRC is the most analogous situation. A similar process of several workshops in conjunction with contractor and NRC staff efforts could be used to develop a proposed function. This function could then be subjected to public review and ultimately to Commission approval.

In summary, we feel that the development of an appropriate payoff function is technically feasible through a process that has some precedent--though a controversial one--within the NRC. Such development would, however, entail considerable costs, as is discussed below. Discussion regarding this and other tradeoffs to determine the overall feasibility and desirability of strategic analysis is reserved for Section 5.0.

4.3.6 Implementation Costs. Our discussion here does not distinguish between NRC and facility operators' costs, because the allocation of these costs would be up to the NRC and the facility operators. The costs of implementing strategic analysis within the NRC regulatory framework fall into two areas: research and operations. The costs to licensees are for all practical purposes in the operational area. In NRC's case, the greatest costs appear to be in the research area. The status of the existing model is such that a significant level of activity is required to bring strategic analysis to a point where it could play an active role in MC&A regulatory programs. Estimating research costs is difficult and particularly so when the program entails the kind of basic research needed for this application of strategic analysis. The still unanswered questions concerning the "best" model, and the benefits of the approach suggest that considerable additional investment would be required. New investments, as with any research program, will entail a risk that no truly useful approach will be found.

Operational costs for the NRC would begin with a decision that research has been successful and that the staff should proceed with implementation. The extent of these costs depends on which regulatory function (see Section 2.2) NRC selects. If the full regulatory process is the choice, then costs begin with dedication of safeguards staff members to the development of a new rule and supporting documentation. The cost of rule development is not easily estimated, but if recent efforts in the safeguards and security area are proper indicators, the process will entail at least several person-years of professional staff time. The new safeguards "Insider" rule has been in the development process for seven years and the new MC&A rule has been in development for five years. Each of these two rules has involved the dedication of several NRC staff members and significant support funding for technical assistance contracts. Rule development activities for strategic analysis can be expected to be of comparable costs.

Should NRC decide on an approach less demanding then the full regulatory program, the costs will be correspondingly less. If the approach is at the level of publishing a NUREG for industry's information, the costs would be minimal. If the Safeguards staff undertakes an internal program supporting its oversight responsibilities, the costs will be considerably less than the full regulatory approach, but the assets required will still be significant. As staff involvement increases, NRC operational costs, including training in the application of strategic analysis, increase.

From these discussions of NRC costs, it is clear that research and operations are closely linked and dependent. The decision as to which regulatory approach to pursue affects the level of confirmatory research required to support that approach and conversely the success of confirmatory research in large part determines the level of regulation that is appropriate and workable. Thus it is an interactive process that should be re-evaluated as new information becomes available, if the NRC makes an initial decision to pursue a strategic analysis program.

Industry costs would be minor compared to NRC costs, particularly if the strategic analysis approach selected does not increase inventory or shutdown frequency. Furthermore, within the small industry of interest to strategic analysis, the Category I fuel cycle facilities, safeguards and security costs are in most cases borne by the Government.

Should strategic analysis for MC&A take a form consistent with the existing game theory model, there would be little in the way of increased staffing or data collection required of licensees. Some training of existing MC&A staff would be required so that licensees understand the model, but otherwise no requirement for additional significant expenditures is evident.

4.3.7 Enforcement Costs. Again we do not distinguish here specifically between NRC and facility operator costs. Strategic analysis would have some effect on enforcement costs. Additional costs would be incurred in order to regularly update strategic analysis models so as to ensure particularly that the payoff function remains an appropriate representation of current costs and values.

It is unclear at this time whether the use of strategic analysis would lead to higher costs because of response actions required (e.g. more clean-out inventories). It is impossible to tell how close current practice is to the optimal approach that would be indicated by strategic analysis. Response costs under current practice could be lower than they would be using strategic analysis simply because we have been lucky, i.e. there have been few, if any, diversions. If, however, strategic analysis does lead to higher response costs, we can be assured that these costs are more than compensated for by increased safeguards effectiveness, since all such costs are included in the payoff function that is used to optimize responses.

4.3.8 Acceptability. Acceptability of new regulations in general is a broad issue, involving the reaction within the NRC, the licensee management and professional staff, industry as a whole, and the public, including special interest and intervenor groups. The concerns of all these groups are important considerations during the process of developing and implementing any regulations or regulatory approach. In the case of strategic analysis for MC&A, the challenge of acceptability may be greater than normal. The innovative nature of the approach will present special problems in an area where innovation and changes have been infrequent. Analytic techniques and mathematical models have experienced difficulty becoming a formal part of NRC's regulatory programs. The reactor "safety goal" and probabilistic risk assessment (PRA) are mathematical concepts that only recently are gaining acceptance. Even these concepts, which have been in development for many years, gained prominence only in response to the difficult questions concerning the NRC's safety objectives and programs that grew out of the accident at Three Mile Island. Furthermore, these techniques are still experiencing acceptability problems, both inside and outside the NRC.

Strategic analysis can expect at least the same acceptability problems as PRA and the safety goal, if it becomes a formal regulatory instrument. Less formal implementation can be expected to mean fewer acceptability problems, primarily because the impact of the new approach on the NRC and industry will not be as great. Historically, the use of strategic analysis for MC&A has been an internal issue within the NRC, where the majority of the safeguards staff has not been enthusiastic about its potential. Comments written during the development of the existing model generally recognize the theoretical superiority of strategic analysis for the adversarial situation that underlies and is the reason for MC&A safeguards programs. However, the difficulties that most have had with the formulation of the payoff function and with the results that the existing model has produced have led to a general skepticism within the NRC that the technique can become a formal regulatory instrument. The acceptability of the existing model has been hampered by the structure of the payoff function, which many involved in safeguards (as well as game theory experts in the Peer Review Group) view as arbitrary. Also the actions the model entails have been questioned as arbitrary and counterintuitive. The functional relationships established by the payoff functions between the size of a theft and such parameters as the amount-searched-for and the various "costs," that are expressed in unusual units, have contributed to the general hesitance to accept the approach. The defender actions required by the models, particularly the mixed strategy model, represent a unique departure from standard regulatory methods. Searching for "lost" material when there is no indication of a loss, or when there is even an apparent "gain" of material, and utilizing a random process to select a mandatory regulatory action may be demonstrably the optimum theoretical approach. To gain acceptability, however, will require a continuous involvement of NRC and industry staff in the development of the appropriate strategic analysis model so they can learn to appreciate the value of particular components (e.g. mixed strategies) and that all previous game theory results are not necessarily inherent to all strategic analysis models. If strategic analysis is demonstrably better than the current practice, which it has the potential to be, it can be expected to be accepted. If, however, the uncertainty and controversy surrounding the previously developed models cannot substantially be dispelled by a further research program, then it is unlikely that the NRC staff, industry, or other interested groups will find strategic analysis acceptable. Our belief is that additional research can dispell the uncertainty and controversy.

4.3.9 Understandability. At times questions have been raised regarding whether industry will be able to understand the approach. However, there is no reason to believe that industry will experience any problem understanding the approach if the NRC staff understands it. The tenets of game theory are relatively simple, although the mathematics associated with optimization and the logic needed to formulate appropriate payoff functions can be very involved. Both the mathematics and the logic must be rigorous and tractable before NRC can ever develop a regulatory program. Industry and other groups will assure this. Therefore, the ability of industry or any other knowledgable group to understand strategic analysis is not an issue and should have no effect on feasibility if strategic analysis is properly developed.

5.0 CONCLUSIONS AND RECOMMENDATIONS

In this section we draw conclusions regarding the feasibility of strategic analysis and make recommendations regarding further development by the NRC. To help structure the conclusions, we draw on the specific questions and issues raised by the NRC in its RFP for this work. Thus, the first subsection addresses four questions that were raised regarding the feasibility and cost of strategic analysis. Then the second subsection discusses four research activities that are relevant to the development of strategic analysis. The third subsection expands the discussion of feasibility to include some additional important issues that were identified in the course of this work. Finally, the last subsection draws together these conclusions and presents recommendations with respect to how the NRC should proceed.

5.1 Cost and Feasibility of Implementing Strategic Analysis

Our investigation of the feasibility of strategic analysis has attempted to take a very practical approach. This approach should help to provide the NRC with the information it needs to make decisions regarding the future of strategic analysis. This investigation has been guided in part by the NRC-identified need to determine the cost and feasibility of resolving the following four issues.

5.1.1 Can a Probabilistic Mixture of Strategies, the Solution to a Game-Theoretic Model, Be Made Acceptable as a Condition of Day-to-Day Plant Operation? Based on our assessment, the key points of which were discussed in Section 4.1.3, a probabilistic mixture of strategies is feasible provided NRC and plant personnel are actively involved or at least regularly informed in the development of the strategic analysis model. With careful explanation, the advantages of a mixed strategy are intuitive. Denying the diverter information regarding the defender's strategy, as is done with a mixed strategy, improves the payoff to the defender. Making the explanation for mixed strategies intuitive and carefully pointing out the economic as well as social cost advantages (in terms of the payoff function) of mixed strategies should make mixed strategies acceptable.

Acceptability may also be improved once the strategic analysis model is developed, by a careful examination of the optimal strategy and the use of sensitivity analyses. Non-intuitive strategies in the probabilistic mixture can be thrown out with changes in payoffs computed for the remaining mixture. If only small changes in payoffs are obtained, it may be desirable to throw out some strategies in the model that is implemented in order to improve its acceptability.

The increased costs associated with use of a mixed strategy result from the need to be prepared for any one of the several possible actions included in the mixed strategy. Complete response plans would be needed for each action rather than a single plan for a pure strategy. Thus, the development of response plans for a mixed strategy could be several times as expensive as for a pure strategy. With adequate plans, however, the cost of actually responding with a mixed strategy should be only slightly higher than with a pure strategy with the added costs coming from the need to maintain the capability for several responses.

These additional costs should be offset by the improvement in the expected payoff to the defender achieved by using a mixed strategy. If sensitivity analyses suggest that improvements in payoffs are not large, then a pure strategy could be used to reduce these costs.

5.1.2 Can the Present Formulation of the Payoff Function (NUREG-0290 and NUREG/CR-0490) Be Modified In Composition and Behavior To Be a Suitable Vehicle for Licensee Application? If an appropriate payoff function is to be developed, it should be developed from scratch rather than by modifying the function in the previously developed model. Both the Peer Review Group (NUREG/CR-0950) and our review of this payoff function have identified numerous deficiencies. The theory and technology (multiattribute utility theory) do exist, however, to develop the necessary payoff function. The difficulty in applying this technology is the need for considerable judgmental inputs. These inputs, which must come from policy and decision makers as well as from other experts, would be rather controversial and subject to extensive scrutiny and discussion. Thus, the people who would make these inputs can be expected to be reluctant to do so. Judgments such as how many dollars is it worth to avoid potential terrorists obtaining five kg. of SNM are not ones with which anyone would feel comfortable. Many methods can be employed by the model developers to make the required judgments as easy as possible, and many people can be involved in making the judgments to help diffuse the responsibility, but still the judgments must be made.

The NRC has some precedent for specifying the types of judgmental tradeoffs that would be required in the payoff function. The development of the safety goal included development of tradeoffs, for example, between dollars and radiation exposure. A similar process could be used to develop the payoff function, but would be very expensive and time consuming, requiring several person-years of NRC staff and contractor effort, and several years to develop.

Alternatively, the payoff function could be developed as a research effort by some combination of NRC staff and contractor effort with little outside participation except some consultation with facility operators regarding costs. This would result in a payoff function, probably as appropriate as one developed through the more public process, but more difficult to defend and open to extensive criticism without the broader consensus of support. Such development would require approximately one to two person-years of effort. While the payoff function developed in such a research project might not provide the function to ultimately be used in strategic analysis because of the problems just mentioned, it would provide a very reasonable starting point and would also provide support for other safeguards problems including those discussed below in Section 5.2.

5.1.3 Can Game Theory Yield a Fully Operational Model That Is Both Valid Enough To Be Useful and Simple Enough to be Solved? While predicting the success of a research and development effort always involves some uncertainty, the answer to this question appears to be a solid "yes." The extension of the two-person zero sum game to include multiple time periods and multiple facilities increases the validity of the model, yet should still lead to a solvable model, as should the inclusion of multiple thresholds and responses. Again, the primary validity issue is the payoff function. Assuming a suitable payoff function is developed, which the previous subsection indicates is feasible, then the resolution of this question is certainly feasible.

In addition to the cost for developing the payoff function, the costs for resolving this question would entail research efforts to develop the model with features as outlined above. As with most research efforts, the costs could vary over a wide range depending upon the depth and breadth of research undertaken. The effort here could range from one person-year to several.

5.1.4 Will a Game-Theoretic Model Be More Costly To Use Than a Classical Statistical Model? A preliminary assessment of the relative costs of the two approaches was discussed in Section 4.3.7. The two areas identified there in which costs could be higher for strategic analysis were the need to regularly update parameters in the payoff function and the costs of response actions. The former type of costs could be estimated fairly accurately as part of the research to develop the payoff function. This effort, with very little additional cost, could identify which parameters would need regular updating, how the updating would be done, and what it would cost.

Since the response costs should be included in the payoff function, any increased response costs should be more than offset by improvements in other objectives represented in the payoff function. Once the strategic analysis model is developed, however, simulation could be used to determine differences in costs. A sample of typical ID's could be run through the strategic analysis model. Based on the indicated responses, the cost of responding could be estimated using information from the payoff function. These costs could then be compared to those that would be incurred if responses were based on a classical statistical model, or some variation of it such as current NRC practice. Most of the information necessary to estimate these costs would also be readily available because of its use in the strategic analysis payoff function. The cost for such a simulation would not be large, once the strategic analysis model is developed, certainly well-less than a person-year of effort.

5.2 Value, Cost, and Likelihood of Success of Research Activities

In addition to directly addressing the feasibility of strategic analysis, the NRC is also interested in several related research areas. These research areas while not absolutely necessary for the development of strategic analysis could lead to results that would enhance its use. For each of the following research activities, we assess its value, cost, and likelihood of success.

5.2.1 Investigation of Optimal Tradeoff Between Type I and Type II Errors. In setting alarm thresholds, consideration should be given to two types of errors. False alarms, in which the alarm is sounded (i.e. ID exceeds threshold) when the true ID is zero are called Type I errors; and misses, in which the alarm is not sounded even though the true ID is not zero are called Type II errors. Different costs are associated with these two types of errors and should be appropriately balanced to set a threshold that determines the relative likelihood of the two types of errors. Clearly as the threshold goes up, the likelihood of false alarms decreases while the likelihood of misses increases.

Classical statistical hypothesis testing sets the threshold so that the probability of a false alarm is fixed (e.g., .05). Strategic analysis, on the other hand, finds an optimal threshold so that the maximum payoff is received based on the costs of the various errors. Thus, if the payoff function in strategic analysis is appropriate, it specifies the optimal tradeoffs between the two types of errors. Avenhaus and Fricke (Avenhaus and Fricke, 1977; Fricke, 1978) have developed game theory models in the international nuclear safeguards context with solutions that minimize the miss rate for a fixed false alarm rate. This variation on game theory partitions the problem into the analytics of solving the game and the cost/benefit type considerations that are appropriate to set the desired false alarm rate. This neat partitioning can be somewhat misleading, however, because the miss rate must be known in order to set the optimal false alarm rate, and it cannot be known until the game is solved. Thus, determining the optimal false alarm rate would be an iterative process. The cited models lack the relative valuations of Type I and Type II errors needed to drive such a process.

The value of finding an optimal tradeoff between Type I and Type II errors lies in improving the expected payoff to the defender in a strategic analysis model, or to the NRC and/or facility operator under current practice. Thus, the value of any such research would not depend upon the implementation of strategic analysis. The classical statistical hypothesis testing approach could be modified into a statistical decision theory approach with the Type I/Type II error tradeoffs specified. For strategic analysis, however, the development of this tradeoff would be an integral part of the model development.

Thus, research on this tradeoff for strategic analysis would not add any cost to the development effort. As an independent research effort not tied to the development of strategic analysis, this research would be comparable to that needed to develop the payoff function for strategic analysis as discussed in Section 5.1.2, costing approximately one to two person-years of effort. The discussion in that section regarding the likelihood that such a function could be developed also applies to the likelihood of success of this research. Our belief is that it can be done, but it would require some willingness on the part of the NRC to make required judgments.

5.2.2 <u>Development of Measures of Effectiveness for Establishment of Alarm Thresholds</u>. Again this research activity is very closely tied to the development of the payoff function for strategic analysis, though it too could be undertaken independently. The components of the payoff function would specify appropriate measures of safeguards effectiveness. As we noted in Section 4.1.5, the first step in the development of the payoff function is the specification of objectives, and in particular which safeguards objectives are to be achieved at least in part through the setting of alarm thresholds and resulting actions. The payoff function is then the translation of these objectives into specific measures. Therefore, the value of such research is that it would provide better measures of the results of safeguards action, and thus allow the actions to be defined in a more nearly optimal manner.

This research activity, if undertaken as part of the development of strategic analysis, would not add any cost. If undertaken independently, it would have approximately the same cost as the development of the payoff function of the Type I/Type II tradeoff research activity, approximately one to two person-years, with the same likelihood of success. Because the basic research required for this research area is mostly the same as that for the Type I/Type II tradeoff research, both areas could be addressed simultaneously with very little additional cost.

5.2.3 Evaluation of the Workability of Solutions for Selection of Alarm Thresholds at Reprocessing Facilities and Other Breeder Reactor Cycle Facilities. Our assessment of this research activity determined no reason why strategic analysis would not be as workable for these types of facilities as well as for the current Category I facilities. The primary modification that would be required to implement strategic analysis for a reprocessing or other breeder reactor cycle facility is in the payoff function. This modification would supply appropriate parametric values (e.g. clean-out inventory costs) to represent the specific facility.

The potential value of using strategic analysis to improve the setting of alarm thresholds at such facilities is probably greater than for existing facilities, because of the higher potential damage of diversions from such facilities. This value could be estimated once an appropriate payoff function was developed. Thus, the value of the research, consisting primarily of specifying parameters in the payoff function, would be to determine whether the improvements in safeguards produced by using strategic analysis have value to society. This issue is discussed more generally in the following subsection.

The cost of the research needed to modify the payoff function for a reprocessing or other breeder reactor cycle facility, would be less than for the initial development of the payoff function--somewhat less than a person-year of effort. If the more general payoff function can be developed, there is little chance that these needed modifications could not be successfully made.

5.2.4 Evaluation of the Costs and Benefits From Implementation of Strategic Analysis. This research activity is also strongly dependent upon the development of an appropriate payoff function. If such a payoff function was an adequate representation of safeguards objectives and costs, once it was developed it could be used to determine the expected payoff of using strategic analysis and also of current practice or other possible approaches. Dividing payoffs into the various components of the payoff function will also help determine who pays the costs and receives the benefits, e.g. the NRC, licensees, or the public more generally.

Ultimately, however, the value of strategic analysis (or any other safeguards approach) is dependent upon the rate at which diversions would "naturally" (i.e. independent of deterrence effects) be attempted. For example, if we knew there would never be any diversion attempts then strategic analysis could not possibly have any value. (If there were no attempts, but we did not know, strategic analysis could have value in assuring against threats and hoaxes.) One of the components that we suggested is feasible and desirable for a strategic analysis model is the capacity to allow for the possibility that a diverter is or is not present. With such a model, sensitivity analyses could be run to compare expected payoffs for strategic analysis and current practice with varying diverter presence. If such analyses showed, as presumably they would, that the difference between the payoffs of the two approaches was highly dependent upon the extent of diverter presence, then more careful consideration would have to be given to how likely the presence of a diverter is in reality.

Such a question is not an easy one to resolve. Precise estimates of the probability of a diverter being present cannot be obtained. Rather, a determination will have to be made regarding what likelihood of presence should be protected against and then the relative costs and benefits for the two (or other) approaches could be determined.

This research activity, in contrast to the previous three, cannot be undertaken until a full strategic analysis model is available, since the model must be applied to determine what costs and payoffs are expected. The value of this research would be in its resolution of whether strategic analysis should be implemented; and if it were implemented, to demonstrate why it is better and to enhance its acceptability.

Including the development of the full strategic analysis model, this research would probably require two to three person-years of effort. Much of this would be in model development with much less effort devoted to a comparison of the expected costs and benefits of strategic analysis.

All of the caveats discussed previously regarding the feasibility of strategic analysis of course apply to the likelihood that this research activity would be successful. If an appropriate strategic analysis model is developed, the subsequent determination of its relative costs and benefits will most likely be achievable.

5.3 The Role of Strategic Analysis in Safeguards

In Section 3.0 we presented a framework for evaluating the feasibility of strategic analysis which included several criteria. While these criteria do determine whether strategic analysis _can_ be implemented (is feasible), they do not determine whether it _should_ be implemented (is desirable). This desirability of implementing strategic analysis cannot be determined without careful examination of the role of strategic analysis, in its suggested use, in the overall safeguards program. As presently conceived, strategic analysis is viewed as being potentially applicable to setting alarm thresholds for ID's. While strategic analysis may be formally appropriate for this function, if for practical reasons the primary objective of this function is assurance, then implementing strategic analysis may not be desirable, since ultimately it may have little, if any, effect on diversion.

The attractiveness of strategic analysis is that it addresses the underlying adversarial nature of the safeguards and security programs. Safeguards and security planning by definition is an adversarial function, wherein two parties plan so as to "defeat" each other. Strategic analysis in the form of game theory grew out of a desire to optimize decisions in such an adversarial situation. The NRC's current MC&A practices do not explicitly incorporate the optimization of adversarial payoffs in the formal sense that game theory does.

Since strategic analysis does include optimization in an adversarial situation, theoretically it offers an attractive approach for most safeguards and security decision making situations. Therefore, the answer to the question of whether the technique can be applied to areas other than evaluation of inventory differences is theoretically "yes." However, this affirmative answer must be conditioned by the important caveat that successful application in any area may be no less difficult than it is for inventory difference analysis. Successful application in other areas of MC&A or in physical security may be feasible, but significant research and development efforts are a prerequisite to any definitive judgments--just as they are for inventory differences analysis.

Perhaps the most reliably promising applications of strategic analysis are those that address relatively narrow issues that can be isolated from the larger picture. Inventory difference analysis and regulatory actions presented a particularly difficult area for the initial application of a relatively novel technical approach, an area in which its reception has been severely hampered by the difficulty of justifying the optimality of the results. The broad implications of the regulatory actions that result from optimization of the total spectrum of defender's objectives, i.e., the payoff function, have been attacked not only because of the way those objectives are formulated, but because they do not account for safeguard and security measures other than material accounting. The approach has been very ambitious.

Examples of less ambitious, but practical and intuitively appealing applications of strategic analysis are limited actions such as the randomization of the patterns of security force patrols, the frequency of audit, and the frequency of personnel and vehicle searches at portals. The defender can obviously benefit from randomization in these areas by increasing the level of uncertainty in the adversary planning. Randomization in such limited areas can be optimized by well established strategic analysis techniques. Thus, narrow applications of strategic analysis to relatively simple and well-defined situations will with all likelihood be successful. Broad applications to problems with significant regulatory implications have greater risks (effort and uncertainty) associated with their success, but of course offer higher payoffs from success. There is a whole spectrum of applications between these two extremes with success potential proportional to their position within the spectrum. This spectrum is dimensioned both by the complexity of the decision to be made and the position and responsibilities of the decision maker.

5.4 Recommendations

On the basis of this study, we cannot recommend that the NRC proceed with full-scale implementation of strategic analysis. It clearly, however, has sufficient potential to warrant further development. Since the necessary research (in particular to develop a suitable payoff function) can be undertaken without large costs, we conclude that such research is indeed worth pursuing.

As we have noted above, development of an appropriate payoff function is the critical element of strategic analysis. A relatively low level research effort (i.e. on the order of one to two person-years) could attempt to develop the payoff function. If this effort were successful, the development of strategic analysis including a comprehensive, public process for reviewing and modifying the payoff function could proceed. If it were not, strategic analysis, at least in this application, could be dropped. With relatively little additional effort, this research could also include the research activities discussed in Sections 5.2.1, 5.2.2, and 5.2.3. Such an investment in research by the NRC would appear to have a relatively high payoff. Even if strategic analysis were not to be further developed, the results of this research with respect to finding optimal solutions to tradeoffs between Type I and Type II errors and suitable measures of effectiveness based on safeguards objectives would be of value for current safeguards practice.

We also suggest that consideration be given to other uses of strategic analysis in the safeguards and security program. Although an examination of where use might be made is beyond the scope of this effort, such applications appear to be appropriate, valuable, and less costly and controversial than use in the present context.

6.0 REFERENCES

Avenhaus, R. and Fricke, H. Game theoretical treatment of material accountability problems, Part I. International Journal of Game Theory, 1977, 5, 117-135.

Fricke, H. Game theoretical treatment of material accountability problems, Part II. International Journal of Game Theory, 1978, 6, 41-91.

Harsanyi, J. Rational Behavior and Bargaining Equilibrium in Games and Social Situations. London: Cambridge University Press, 1977.

Keeney, R. and Raiffa, H. Decisions With Multiple Objectives: Preferences and Value Tradeoffs. New York: Wiley & Sons, 1976.

NUREG-0290. A Study of Nuclear Material Accounting, June 1977.

NUREG/CR-0450. Report of the Material Control and Material Accounting Task Force, April 1978.

NUREG/CR-0490. The Application of Game Theory to Nuclear Material Accounting, November 1978.

NUREG/CR-0950. A Review of the Application of Strategic Analysis To Material Accounting, August 1979.

Page, R. Letter to Gentlemen, Material and Plant Protection Branch, U.S. Atomic Energy Commission, 1974.