Bridging the gap between radiation protection and safety: The control of probabilistic exposures

A review of the evolution of common safety approaches

by A.J. González and G.A.M. Webb

The operator of the food irradiation facility was relaxed at the console behind a thick concrete wall. He knew that the shielding provided adequate protection against the radiation beam crossing it. Such protection was governed by the "system of dose limitation" recommended by the International Commission on Radiological Protection (ICRP) and adopted by most national authorities and international organizations. An enormous number of standards, guides, and recommendations have been developed worldwide to implement the ICRP system. The IAEA alone had issued more than 100 of these documents. The operator was aware that, through the application of these interrelated regulations, the main features of the ICRP system were being respected worldwide.

Even such a well-developed system, however, did not eliminate all potential radiation risks. Based on assumptions about the health effects of low radiation doses delivered at low dose rates, the radiation risk — estimated from radioepidemiological data from people exposed to rather high doses and dose rates — was of the order of 1 in 100 000 per millisievert of dose incurred. The dose rate outside the shielding was such that the dose to anyone permanently camped there would be a few millisievert at most. The operator's own exposure was lower. During the past year, he had incurred an accumulated dose of a fraction of one millisievert and therefore his risk during that year was less than 1 in 100 000. He realized that this was negligible, much lower than other risks he was facing daily. The operator was satisfied with his working conditions and confident that the protection system adequately protected him against the radiation source.

*****

As the operator watched the console, a flashing light showed that something was going wrong inside the irradiation enclosure: the system transporting the food was blocked and he would have to go into the enclosure to repair it. Following his written operating instructions, he pressed the controls to shut the source down, submerging it into the water pool which served as its shielding. An indicator at the console told him that the

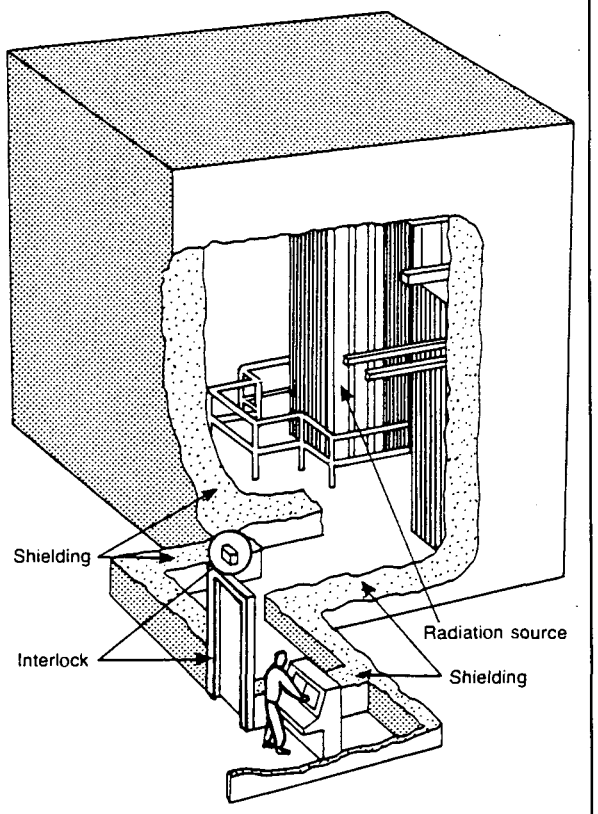

Schematic of irradiation facility.

Mr González is Head of the IAEA Radiation Protection Section, Division of Nuclear Safety. Mr Webb is the Secretary of the National Radiological Protection Board of the United Kingdom. Views expressed in the article are their own and do not necessarily reflect those of the organizations to which they belong.

operation had been successful. He was now clear to open the door to the labyrinth, go into the enclosure and get on with the repair. At this moment, however, some questions crossed his mind — Was the source really down in the pool? Had the mechanism operated correctly as expected? Even if the shutdown system had failed, he knew that when he opened the door another safety system based on an interlock connected to the door would shut the source down anyway. But what if that safety system also failed? He would then be severely injured by the very high exposure he would receive.... He shrugged his shoulders — why worry about something so unlikely. The operator confidently entered the enclosure. After all, he had been told the safety systems were of the latest proven design, built to good engineering standards. An accidental overexposure was near enough impossible.

*****

"Impossible" is a word that scientists are reluctant to use. They prefer to qualify the possibility of occurrence of phenomena by their certainty (or uncertainty) and to measure such certainty by means of the quantity called probability. Consider the two possible situations of radiation risk described above: In one, there is a *certainty* that an *exposure* will occur and a subsequent probability of *radiation harm* caused by such exposure. In the other, there is just a *probability* that the *exposure* will be received, but should it actually occur, depending on the dose level, there may be a *certainty* of *radiation harm*. In both situations, it is feasible to assess the radiation risk:* while in the first case the risk is proportional to the dose, in the second case it is proportional to the product of the exposure probability times the dose. It should be feasible to control the level of risk by means of technological systems.** For example, for the first situation, the control system is the radiation shielding, and the control parameter is the shield's thickness. For the second situation, the control system is the interlock and the control parameter is its reliability.

The discipline called radiation protection usually deals with the first type of situation. The second type is usually studied by safety experts.

This paper describes the evolution of common safety approaches for dealing with these two types of situations. If a common, coherent, and consistent approach is achieved, a balanced partnership will exist for the radiation protection and safety disciplines.

Situations involving radiation exposures

Three types of situations can be envisaged when forecasting possible scenarios of radiation exposure:

- anticipated situations in which the exposure of people is planned for and assumed to occur with certainty (i.e., with probability unity or very near to it)
- situations that can be anticipated but whose occurrence is not certain; if they occur, however, they would give rise to exposure of people
- (de facto) situations that may or may not have been anticipated but, in the event they do occur, allow only remedial actions to be taken.

Protection against ionizing radiation is usually based on the recommendations of the International Commission on Radiological Protection (ICRP). The recommendations are not specific in the exposure scenarios they cover. In practice, however, such recommendations have been used for the first scenario and partially for the third one. The ICRP recommendations are mainly used for *situations* that involve conditions of *exposures that are assumed to occur with certainty*. These will be called "certain" exposures in this article; the term roughly encompasses those which in practice are called "normal" exposures and exposures from "routine operational occurrences". The recommendations have not been used in practice for the second scenario, i.e., for exposures that may occur with a probability lower than one. These will be called "probabilistic" exposures in this article.

National and international standards implementing the ICRP recommendations have implicitly recognized this fact.*

"Certain" exposures: Protection policy

For "certain" exposures the ICRP recommends a system of dose limitation which includes the following interrelated requirements:** (a) no practice shall be adopted unless its introduction produces a positive net benefit (or *justification of the practice*); (b) all exposures shall be kept as low as reasonably achievable, economic and social factors being taken into account (or *optimization of protection*); and (c) the dose equivalent to individuals shall not exceed the limits recommended for

* The word risk is used to mean the probability of severe harm due to radiation exposure.

** The word control is used to mean exercising restraint rather than checking or verifying.

* For example, the *Basic safety standards for radiation protection* of the IAEA, International Labour Organisation (ILO), World Health Organization (WHO), and the Nuclear Energy Agency of the Organisation for Economic Co-operation and Development (NEA/OECD) recognize "two distinct conditions of exposure, which are: (i) conditions in which the occurrence of exposure is foreseen and can be limited by control of the source and by the application of the system of dose limitation...(normal exposure conditions); and (ii) conditions in which the source of exposure is not subject to control so that any subsequent exposure can be limited in magnitude, if at all, only by remedial actions...(abnormal exposure conditions)". Condition (i) clearly applies to what this paper refers to as situations involving "certain" exposures and may apply in part to anticipated situations. Condition (ii) applies to situations that have occurred. The *Basic safety standards* as a whole clearly apply to condition (i), while only a few general provisions can be extended to condition (ii).

** *Recommendations of the International Commission on Radiological Protection*, ICRP Publication 26, *Annals of the ICRP*, Vol. 1, No. 3, Pergamon Press, Oxford (1977).

the appropriate circumstances by the ICRP (or *individual dose limitation*).

It is worthwhile to analyse these three basic policy principles for "certain" exposures within the framework of a possible extension to "probabilistic" exposures.

Justification of the practice. This requirement simply specifies that, in order to permit the introduction of a radiation practice, more benefit than harm should be expected. Its implications, however, have not yet been analysed in full by the ICRP or by other organizations. Moreover, when moving from the limited framework of "certain" exposures to the wider scope of "probabilistic" exposures, the practical implementation of the principle of justification becomes more complicated. For some scenarios, the probability of occurrence can be very low but, if the scenario happened, the consequences could be high. It is not clear how these situations should be included in an assessment of justification.

Optimization of protection. The ICRP used this term to express its intentions to keep all doses as low as is reasonably achievable (ALARA), social and economic considerations being taken into account.* Unfortunately, this simple requirement has been misinterpreted by many people as a synonym for cost-benefit analysis. It has been clearly indicated that the implementation of optimization does not necessarily require the use of any particular decision-aiding technique, such as cost-benefit analysis, and that optimization is amenable to any suitable technique, including simple intuition and common sense.* This broader description of optimization is essential for extending the optimization principles beyond the limited scope of situations involving "certain" exposures.

Individual dose limits. They are recommended by ICRP for "certain" exposures from artificial sources of radiation. In setting the limits, however, it recognized the potential for "probabilistic" exposures. If this scenario had not been included, the numerical values of the limits would have been different.

Since dose limits apply to individuals and a single individual may be exposed to several sources, both the ICRP and the IAEA have also recommended the application of *upper bounds*. The upper bound for a single source is set at some fraction of the dose limit applying to total exposure to all sources.

"Probabilistic" exposures

As discussed before, the ICRP system of dose limitation covers all situations in which radiation exposure of people is planned to occur and the source can be controlled. Although the principles are universal, they cannot be used in their current form to control sources that may or may not give rise to exposures. For "probabilistic" exposures, therefore, the ICRP system is not directly applicable, but its underlying principles could be developed for use. To do so, a system would need to be based on probability and dose control rather than on dose control alone.

Effectively, all planning situations with radiation sources involve "certain" as well as "probabilistic" exposures. The relative importance of the two modes may differ enormously for different sources but in principle both should be considered for all sources. In an abstract situation it is not always obvious how the distinction between the two modes is to be made, but it is not too difficult to resolve the issue for any particular source. For example, a routine exposure may result from a number of separate incidents; if these incidents are reasonably frequent, there may be a tendency to designate the exposure as "certain", while if they are infrequent, the resulting exposure may be included in the "probabilistic" category.

The quantities that characterize the source from a safety viewpoint will vary according to the scenario being considered. For scenarios involving "certain" exposures, the relevant quantity is the dose distribution, usually characterized by both the dose to the most exposed individual and the collective dose delivered by the source. These quantities are usually sufficient for the low doses expected in normal operation: the incremental dose received by individuals is assumed to produce a proportional incremental harm and, therefore, as the individual dose is a measure of the individual harm, or risk, the collective dose became a measure of the total expected harm. For "probabilistic" exposures, it is also possible to identify a probability of individual harm, or risk. This is done by combining the probability of occurrence of the dose and the probability of inducing harm given the dose. It follows that a probabilistic distribution of consequences could also be identified.

Controlling "certain" and "probabilistic" exposures: Parallel developments

Some procedures for the assessment and control of "probabilistic" exposures have been developed in parallel to, and to some extent separately from, the basic principles of radiation protection. Radiation safety objectives have been developed at the national level for some sources of "probabilistic exposures" — notably for nuclear power reactors — and an international consensus on some nuclear safety principles seems to be emerging.* Procedures for assessment and control relating to waste disposal have also started to evolve separately but have now been tackled by extending and developing the ICRP basic recommendations to deal with the particular problems of wastes.** Suggestions have been made for a unified approach to control for all

* *Cost benefit analysis in the optimization of radiation protection*, ICRP Publication 37, *Annals of the ICRP*, Vol. 10, No.2/3, Pergamon Press, Oxford (1983).

* *Basic safety principles for nuclear power plants*, IAEA Safety Series 75-INSAG-3, IAEA, Vienna (1988).

** *Radiation protection principles for the disposal of solid radioactive waste*, ICRP Publication 46, *Annals of the ICRP*, Vol. 15, No. 5, Pergamon Press (1985).

Dose-risk relationship used for radiation protection purposes

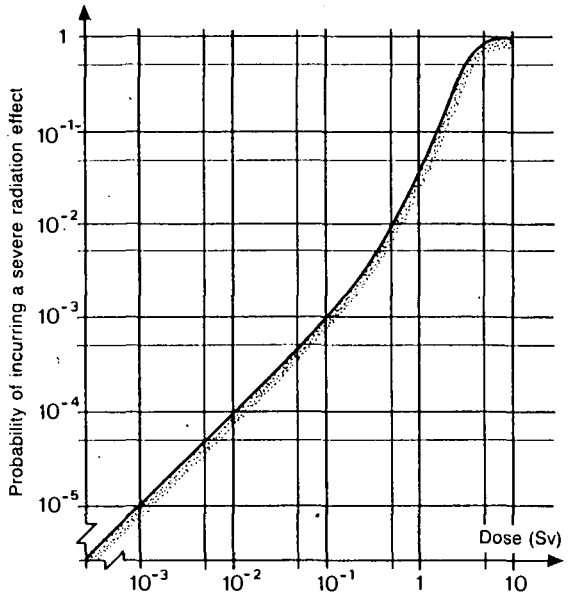

The figure illustrates the dose-risk relationship used for radiation protection purposes. Three regions can be recognized: (1) *At dose levels below a fraction of one sievert,* only stochastic effects occur, the probability of occurrence being directly *proportional* to the dose level. These include fatal cancers in the irradiated individual and severe genetic effects in the succeeding generations of descendants of the irradiated individual. In this range it is assumed that, following any increment of dose, there is a proportional increment in the probability of an effect. This is not necessarily an exact representation of radiobiological data, but should be seen as a simple practical assumption intended for planning purposes rather than as an assessment tool for individuals exposed. The relationship of probability of harm to dose is therefore assumed to be linear in this range. For radiation protection purposes, the slope of the line (i.e. the risk factor in this region) is currently taken to be 1.6×10^{-2} Sv^{-1}. New radioepidemiological information on atomic bomb survivors seems to indicate that the risk factor might change in the future. (2) *For doses of a significant fraction of one sievert,* delivered in a short period of time, non-stochastic effects may occur. The dose-risk relationship approximates to a sigmoid relationship. As in the stochastic range, the exact shape depends on a number of factors, such as the dose rate, which could be relevant for particular scenarios of exposure. For a dose of approximately 3 Sv the probability of death is about 0.5. (3) Finally, it is assumed that for *doses higher than about 5-10 Sv,* delivered in a short period of time, practically all irradiated individuals will suffer an acute radiation syndrome and eventually die as a consequence of the irradiation. Therefore, the relationship is assumed to approach asymptotically a probability of unity for doses higher than about 5-10 Sv.

these areas of concern, with common principles for dealing consistently and coherently with routine and potential exposures. For its part, the IAEA has recently produced a consultative document on the application of the principles of radiation protection to sources potentially causing exposure, with the intention of working towards a unified approach to radiation safety.*

* "The application of the principles of radiation protection to sources of potential exposure: Towards a unified approach to radiation safety" (a consultative document), IAEA, Vienna (1988).

Basis for a radiation safety policy

The basis for a policy on radiation safety that encompasses all scenarios with all exposure conditions, both certain and probabilistic, may be founded on the relationship between risk and dose used for radiation protection purposes, which is based on a number of radiobiological assumptions. A general radiation safety policy should admit some probability for doses exceeding limits and upper bounds and even entering the region of doses where severe, "non-stochastic" effects may occur; it is therefore particularly important to specify these assumptions. Such a dose-risk relationship can be used as a basis for a common safety policy. *(See accompanying figure.)*

Towards a converging policy

It seems reasonable to focus on the idea of a *limit on individual risk* as one *necessary*, although *not sufficient*, requirement for a unified approach to radiation safety in general and, particularly, to the control of probabilistic exposures. It is tempting to search for some compatibility with the current system of dose limitation and therefore to specify an overall limit for the risk to any individual, and to deal with the question of the extent to which safety should be improved below this limit by extending the concept of optimization to include the consideration of all risks to the exposed population. This is conceptually attractive, and the two ideas of individual limits and optimization under the context of probabilistic exposure are being developed and elaborated.

It is necessary to emphasize, however, that "societal risks" have also been considered by nuclear safety experts in terms of criteria framed as "societal risk limits", or "societal risk objectives".* These social criteria seem to be outside the direct extension of the requirements of individual dose limitation and optimization. They might, however, be associated with the requirement of justification, but this potential connection will not be analysed in this article.

Individual risk limits

The current dose limit of 1 millisievert (mSv) per year recommended by the ICRP for members of the public corresponds to a committed risk of approximately 10^{-5} and can be used as a reference value for developing an individual risk limit criterion. For routine exposure below the dose limit, the only health effects that have to be considered are cancers and hereditary effects. If the dose limit now is converted into a risk limit, the constraint on dose itself is removed; other health effects such as death from high doses must therefore be taken into account. It is not necessary in principle to give all types of health effects the same

* *Status, experience and future prospects for the development of probabilistic safety criteria,* report of a technical committee meeting in 1988, IAEA TECDOC (in publication).

weight; if, for instance, the loss of person-years were used as a weighting function, then acute death from high doses would have more weight than death from cancer after a latency period. However, for the sake of simplicity it seems reasonable in the context of limits (not objectives) to treat health effects as equally serious; it is then possible to specify a single numerical value as a risk limit to apply to all events of a probabilistic nature.

Therefore, for consistency with the general safety standards of dose limitation a risk limit of 10^{-5} in a year has been proposed for individual-related assessments of radiation safety. This limit would apply to the individual risk in the most highly exposed individual (except patients) from all sources of potential exposure (except natural radiation sources). It is important to be aware that this risk limit would be the lower boundary of the region of unacceptable risk; a risk below this limit should not necessarily be judged to be acceptable.

Since an individual can be at risk owing to more than one source, in addition to the risk limit (which refers to an individual) there needs to be a source-related *risk upper bound* (or even a scenario-related risk upper bound), which limits the individual risk coming from a single source (or single exposure scenario, respectively). The risk upper bound is apportioned from the risk limit (i.e. is chosen as some fraction of the latter) and may depend on the source or scenario being considered. A risk upper bound, allocated to a source, is to be used in design and regulation of a particular facility in the same manner as current dose upper bounds are used. The simplest method of incorporating probabilistic scenarios into a risk-based system of radiation protection is to define separate risk upper bounds for probabilistic exposure while retaining the current dose upper bound for normal operations.

Depending on the selected risk upper bound, a criterion curve for limiting individual risk per scenario can be formulated as a direct derivation from the risk-dose relationship.* (As shown earlier in the figure on risk-dose relationship.) *(Also see accompanying figure illustrating a criterion curve.)*

Safety optimization

Ensuring that no individual will incur an unduly high radiation risk is a necessary but not sufficient condition for ensuring the appropriate level of safety of a radiation source. The question remains whether that level should be improved further by taking into account, for instance, that a large number of individuals incurring an individually low risk may represent an unacceptably high overall expectation of harm.

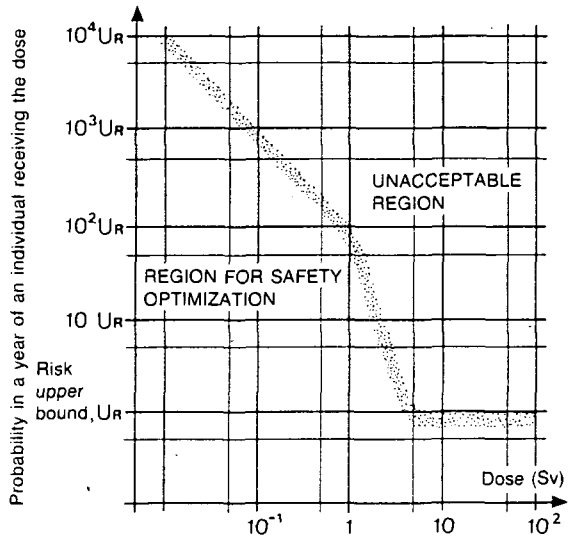

One procedure for applying individual-related requirements to probabilistic events is to express risk limits in a criterion curve. Such a criterion curve, of the maximum probability that can be permitted for an estimated dose from the initiating event, based upon an annual risk upper bound constraint, U_R, to the critical group, is shown here. The relevant features of the criterion curve are as follows: an inverse proportionality region; a non-proportional region for the dose range in which non-stochastic effects may also occur; and a constant probability for doses that are lethal. In the lethal dose range, the probability is constant irrespective of dose, because the consequence to the individual is the same regardless of the dose received. For the range of doses in which only stochastic effects occur, the relationship between probability and dose is inversely linear, with values representing the product of the probability of the dose, the annual dose, and the probability of a health effect per unit dose. Finally, in the dose range where non-stochastic effects may occur, i.e. individual doses exceeding a few sievert, the shape of the criterion curve is non-linear, in order to take into account the increasing probability of death. This portion of the curve should approximate a sigmoid relationship and would depend to some extent on the time over which the dose is delivered.

The proposed criterion curve can be used to indicate whether a given safety option complies with the risk-related requirements in the following manner. First, the events, or sequences of events, with the potential to cause exposure to individuals should be identified. An event or sequence of events may be selected as representative of a group of similar scenarios, so long as the maximum consequences are considered. Second, the probability of occurrence of each event, and the consequent exposures of the critical group should be assessed. Finally, the point representing the probability of occurrence of the initial event and all other environmental conditions and the corresponding maximum dose is plotted. If the point is in the unacceptable region, then the option should be rejected. However, even if all the points are in the acceptable region, the proposal being assessed may not be acceptable because it is not optimized. Therefore, the usefulness of a criterion curve is limited at this stage to that of a basic decision tool for checking whether an option is unacceptable. The next stage is to check whether the option meets the ultimate requirement that safety is optimized.

* "The regulatory use of probabilistic safety analysis in Argentina", by A.J. González, in *Proceedings of the international meeting on thermal nuclear reactor safety*, Chicago, USA, NUREG/CP-0027, (1982).

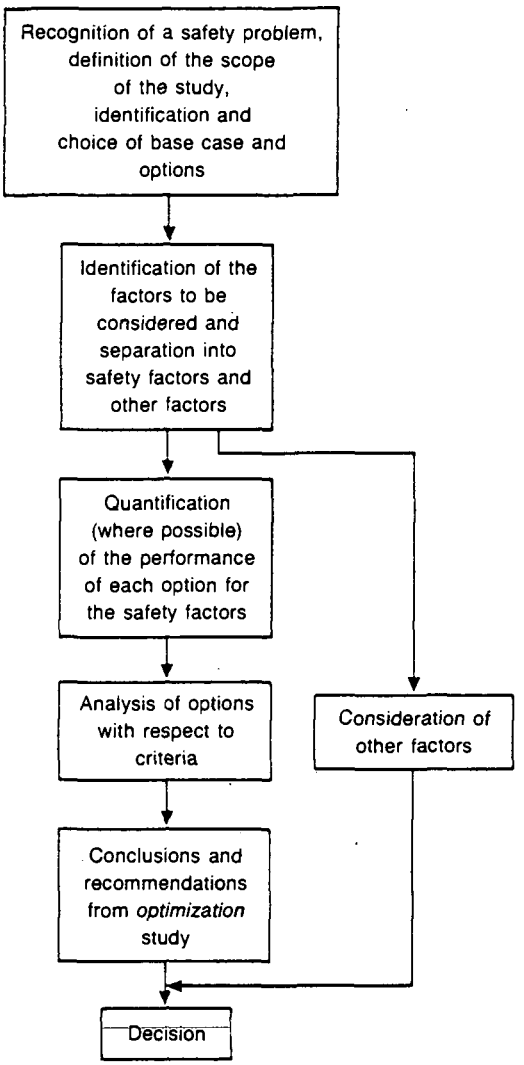

A structured approach to optimization of safety is important to ensure that no important aspects are overlooked and to record the analysis for information and for assessment by others. The essential elements of this structured approach can be set out as shown here. The terms are used as follows: *safety option:* a specified design or set of operational procedures; *base case:* the starting point from which changes are assessed; in a design study normally the cheapest option, in the case of operations the current set of procedures; *factor:* an identified measure of either the cost or the performance of an option; *safety factors:* those factors which are related to the level of safety achieved. These will include those factors describing the resultant radiation risk distribution in any way, and those factors describing the costs and other disadvantages incurred in modifying the risk distribution; *other factors:* those factors which are related to or describe the performance or costs of an option but are not related to the level of safety, e.g. costs incurred solely to increase process efficiency, for aesthetic or public relations reasons, differences between options expressing their likelihood of public acceptance, measures of process efficiency such as throughput, planning consideration; *performance of the option:* the results of applying a specified design option or set of operational procedures; expressed in terms such as the resultant doses and occurrence probabilities; *criterion:* a quantitative or qualitative measure of what is acceptable or desirable for one or more of the factors. This is then the basis against which the performance or cost of an option can be compared.

For "certain" exposure situations it is required that the radiation protection applied to the source must be optimized. This requirement generally leads to individual doses being well below the individual dose limits. The concept of optimization of protection involves the choice of the most appropriate level of protection, taking into account a number of factors, the major ones being: (1) the total harm to the exposed population, represented by the collective dose, and (2) the costs of protection. It is recognized, however, that other factors may be taken into consideration, such as the distribution of doses. Thus the process of optimization can be thought of as using a decision-aiding technique. Since there will also be other factors that enter the final decision, some of which have nothing to do with radiation protection, the result of optimization can be seen as a partial input to the final decision.

It is clear that the full assessment of expected harm from "probabilistic exposure" scenarios includes consideration of the number of people affected and the probability and level of doses to them, and of all efforts, including costs, required to improve safety. This aspect is sufficiently close to the ideas involved in optimization of protection that it is sensible to examine how the concept might be extended to a risk-based system.

In extending this system, it is most useful to focus on the "decision-aiding" concept of optimization and to extend the number of factors incorporated to include the probability and consequences of potential exposures. It is also useful to view optimization as a structured approach in the context of decision-making. *(See accompanying figure.)* In extending optimization it may not be helpful to make the assumption about the equivalence of various types of health effects that was adopted for simplicity for establishing the risk limitation criteria. In particular, it is probably better to deal separately with the consequences in terms of non-stochastic effects, particularly acute death, and with stochastic effects.

The relevant quantity suggested by the ICRP and adopted by the IAEA for optimization of protection against "certain" exposures is the "detriment", defined as the expectation value of harm for the group of people affected by the source of radiation. For probabilistic exposures, the use of the concept of detriment may not be straightforward, as the following example shows:* Consider an accident sequence which has a low probability P of occurrence and which has a high consequence C if it occurs (and of course no consequences if it does not occur). The expectation value of harm is given by the product PC. If P is very small and C is very large, the detriment will be of an intermediate value which does not adequately quantitatively represent the situation, which is that there are either no consequences or major consequences. In other words, the large uncertainty about the magnitude of the consequences is not

* "Critical views on the application of some methods for evaluating accident probabilities and consequences", by D.J. Beninson and B. Lindell, *Current nuclear power plant safety issues*, IAEA (1980).

evident to the decision-maker and therefore is not able to be included in the decision-making process. Thus, detriment may not be a useful quantity for evaluating options in such cases. For "probabilistic" exposures therefore, the quantities to compare should include — besides the safety efforts — the full distribution of probabilities and consequences.

In addition to the problem of what quantities to compare, there is the problem of how to include in the comparison process quantities or preferences which are not expressed in commensurate units. Such preferences, which should explicitly be accounted for, include the degree of risk aversion for higher consequence accidents; social costs for restrictions or inconveniences; morbidity and mortality of the various types of radiation effects; and the relative weighting (i.e., the degree of relative importance) of the effects.

The problem of comparing quantities which are not directly linearly comparable can be addressed using utility functions and decision theory.* Preferences for quantities of differing types are expressed using a utility function which prescribes how the different types of quantities are to be combined for the purposes of comparison. The resulting utility functions can then be combined using a decision-aiding mechanism to arrive at a "best under the circumstances" (i.e. optimized) option. The use of this class of decision-aiding technique is being discussed by a Task Group of the ICRP and has been described in outline form.**

Outlook: A unified policy; problems to be solved

In summary, a risk limitation system based on the principles of safety optimization under the constraint of individual risk upper bounds, together with the current system of dose limitation, may form the basis for a unified policy for radiation safety.

For a successful implementation of such policy, however, some practical problems remain which need further investigation; they include the following:

- There are many *uncertainties involved in probabilistic safety analysis*. The consequent lack of confidence in the result should be reflected either in the degree of conservatism to be used in establishing the relevant safety objective or when comparing the results with the objectives. For instance, the use of risk limits, or upper bounds, as a safety objective rather than as a constraint to the objective may present difficulties in this context. It should be clear, therefore, that risk limits and upper bounds can neither be interpreted as objectives nor as goals. Rather, they can be interpreted as a boundary of a forbidden region and they should incorporate the necessary conservatism to cope with the expected uncertainties.
- Currently, there are *no standardized tools available* for performing probabilistic assessments. Large variations in the results can be obtained for the same situations if the methodology and the boundary conditions for the analysis are not specified in sufficient detail. This may produce the undesirable situation that two analyses of the same situation might show either compliance or non-compliance with a defined risk constraint. The solution to this problem seems to be the development of standardized probabilistic safety assessment procedures, and their incorporation into the relevant regulations. However, this would run counter to the regulatory philosophy in some countries that adopt a non-prescriptive approach.
- Another problem is that of *measurability and accountability of risk*. For "certain" exposures the "effective dose equivalent" is used as an indirect measure of the risk incurred by exposed persons. Such a quantity is "measurable", albeit by way of other related physical quantities and various assumptions and hypotheses, and therefore can be accounted for in relevant records with legal status. The situation is rather different for probabilistic exposures: the effective dose equivalent to be incurred if the exposure is actually delivered does not measure the risk (since there is a chance that the exposure does not occur) and no other "measurable" and "accountable" quantity exists. The reliability of the safety systems, the consequent probability of the exposure or the combination of probability and dose cannot be "measured" in the instrumental sense; they may not be accepted as quantities for record purposes, so *a posteriori* compliance would not therefore be legally "demonstrable". There may be ways of solving this legal problem but they have not been explored yet.

In view of these problems, some safety experts have expressed caution concerning the application of a risk limitation system to nuclear power plants. They prefer to focus on general probabilistic objectives. Other experts have felt that in spite of practical difficulties the establishment of a basic philosophical framework of risk limitation should be encouraged. We feel the latter view has enough prospect of success to be worth pursuing.

* See two papers by D.J. Beninson: "Optimization of radiation protection as a special case of decision theory", *Optimization of radiation protection*, IAEA, Vienna (1986); and "Application of radiation protection optimization principles to potential exposures from accidents", *Nuclear power performance and safety*, Vol. 4, IAEA, Vienna (1986).

** "Decision-aiding techniques for radiological protection", by G.A.M. Webb and J. Lombard in *Radiation protection in nuclear energy*, proceedings of the IAEA conference in Sydney, April 1988 (to be published).