ARE RANDOMIZED INSPECTIONS AT Pu OR HEU STORAGE FACILITIES SUFFICIENTLY EFFECTIVE UNDER INTEGRATED SAFEGUARDS?

Russell Leslie, Annette Berriman and John Carlson
Paper presented at INMM 2005 symposium Phoenix, USA – July 2005
Australian Safeguards and Non-Proliferation Office (ASNO)
RG Casey Bldg, John McEwen Crescent, Barton ACT 0221 Australia

Abstract:

Unirradiated Direct Use (UDU) material (Pu, including MOX, and HEU) is the most proliferation-sensitive nuclear material. Safeguards measures must take into account that access to significant quantities of UDU would allow a potential proliferator the quickest route to developing nuclear explosive devices.

In the move towards integrated safeguards (IS) it will be necessary to develop IS inspection schema for facilities that contain UDU. This paper will examine the possibility of IS inspection schema taking full advantage of the flexibility allowed under IS including increased reliance on randomised inspections (whether unannounced or allowing for shortened notice periods).

Issues addressed by this paper include:

- Average inspection frequency;
- Calculation of mean time to detection under a variety of assumptions;
- The possibility of determining optimal inspection strategies

1. INTRODUCTION

With the establishment of Integrated Safeguards (IS) the IAEA is faced with the task of finding the optimum combination of available measures in formulating its safeguards approach for any given state. In its choice of methods the IAEA is bound by a series of important constraints, including (*inter alia*):

- the need to be effective
- the need to be efficient
- the need to work within available resource constraints (including both human and financial resources)
- the need to be non-discriminatory
- the need to maintain the confidence of the international community in the effectiveness of the safeguards system.

One proposal for the development of IS safeguards approaches to increase the use of Short Notice Random Inspection (SNRI). These are inspections timed in a way that is unpredictable to both the facility operator and the state. The use and potential value of SNRI will be examined in context of potential safeguards approaches.

In choosing any particular safeguards approach over others that are available it is helpful to be able to develop a method to assist in the evaluating which approach is best suited to the IAEA's needs (and which best fits within the above listed constraints). One useful methodology is to determine a "mean time to detection" (MTTD) for each approach and use the calculated figure as a guide to decision making. It is important to emphasise that such a methodology has a limited range of applicability – it is useful for comparing different types of inspection schemes (e.g. fixed v. SNRI, X inspections per year

v. Y inspections per year) but it would not be useful for comparing schemes based on widely different underlying approaches (e.g. inspection schemes v. information analysis).

Unirradiated Direct Use (UDU) material (Pu, including mixed oxides (MOX), and high enriched uranium (HEU)) is the most proliferation-sensitive nuclear material. Safeguards measures must take into account that access to significant quantities of UDU would allow a potential proliferator the quickest route to developing nuclear explosive devices.

Two parameters that are key to the understanding of this discussion, the first is referred to as μ_i – the mean time between inspections and the second is referred to as μ_d – the mean time between detections. These two μ parameters are closely related to each other and in this discussion if the term μ is used without a subscript then the discussion is meant to apply to both μ_i and μ_d. The number of days between inspections is referred to as the "inspection interval". The number of days between a diversion and its detection is referred to as the "detection interval".

Key Assumption 1 (KA1): It is assumed throughout this discussion that diversion is not a random act, the diverter is actively seeking to defeat the safeguards system and is prepared to take all necessary steps to prevent detection (in effect to maximise the time to detection). One key element of any strategy to maximise the time to detection is to divert material immediately after the completion of an inspection. Key Assumption 1 (KA1) holds throughout the discussion that follows.

This paper is a short summary of a longer report that has been produced by the Australian Safeguards Support Program under Task AUL C 01208: "Re-Examination of Basic Safeguards Implementation Parameters". A more complete explanation of the calculations undertaken in the paper and the terms used can be found in report AUL Report 2005-02 "Randomized inspections at UDU storage facilities under Integrated Safeguards, Mean Time to Detection", available from the authors.

This paper will look at the issue of the adequacy of random inspection schema for addressing the most sensitive material subject to safeguards. The methodology used will be via the use of calculated mean time to detection parameter.

2. DISCUSSION

Mean time between inspections

The calculation of a "mean time between inspections" parameter μ_i for either fixed or randomised inspections is trivial. It is dependent only upon the length of the inspection year[1] and the target number of inspections. If there is an average of only one inspection per facility per year, then the long term average interval between inspections per facility must tend towards 365.25 days. Shorter or longer average inspections intervals per facility may be observed over relatively small collections of inspection years (1 to 100 inspection years), but the long term average must tend towards 365.25.

Demonstrating the logic behind this simple result is straightforward. Assuming, for example, that the long term average inspection interval per facility is any number smaller than 365.25 (for example 300

1. The term inspection year is used to mean one year of Agency inspection effort with the year starting at the facility Physical Inventory Verification (PIV).

days) then over one thousand inspection years[2] there will be approximately 1220 inspections. In this example, the shorter average inspection interval would result in 22% more inspections than required over the period. The average would not be 1 inspection per facility per year but would be 1.22 inspections per facility per year instead. Similarly if the average inspection interval is any number larger than 365.25 (for example 460 days) then over one thousand inspection years there will be approximately 800 inspections. The longer average inspection interval would result in 20% fewer inspection than required over the period.

The calculation of an average inspection interval is not a matter that requires recourse to complex mathematical algorithms, it is a matter of simple arithmetic.

$$\mu_i = \psi / \acute{\iota}$$

= mean time between inspections

where

ψ = the length of the inspection year[3].

and

$\acute{\iota}$ = the target number of inspections per year.

Probability of inspection on any day

If the number of inspection opportunities is equal to the number of days in the inspection year, then the probability of an inspection taking place on any given day (referred to as λ) is given by the following formula:

$$\lambda = 1/\mu$$

As the value of μ_i and of λ are both independent of the number of inspection opportunities per year it follows logically that changes in the number of inspection opportunities per year (e.g. via limiting inspection opportunities to facility working days or imposing minimum intervals between inspection) do not change the value of μ_i.

Probability of inspection at any inspection opportunity

It is obvious that the probability of inspection per inspection opportunity can not be independent of the number of inspection opportunities, and so a new parameter is required, $\bar{\lambda}$ (referred to as "lambda bar" or "lambda adjusted"). The new variable $\bar{\lambda}$ represents the adjusted value of the daily probability of inspection that is necessary to maintain the target number of inspections per year when the number of inspection opportunities in less than the number of days in the inspection year.

For example, imposing a minimum inspection interval (I_{min}) on a randomised inspection scheme removes a minimum of $\acute{\iota}$ x (I_{min}-1) inspection opportunities from the inspection year (as no inspection can take place within I_{min} days of any other inspection). This adjustment removes the simple numeric relationship between λ and μ, with $\bar{\lambda}$ defined as:

$$\bar{\lambda} = \acute{\iota}/(\psi - \acute{\iota} \times (I_{min}-1))$$

2. 1000 inspection years can occur as any combination of numbers of facilities and numbers of inspections per year that sum to 1000 inspection years (for example, 5 inspection years at each of 200 facilities, 10 inspection years at each of 100 facilities or 1000 inspection years at one facility).

3. The value of ψ places an upper bound on the number of inspection opportunities per year. In many practical cases the facility is not available for inspection every day of the year.

Mean time between detections

Once a value for the mean time between inspections, μ_i, has been determined, there are a number of different methodologies that can be used to calculate a mean time to detection (MTTD) parameter μ_d. If we were to assume that all inspections are perfect (probability of detection = p = 1 and probability of non-detection = 1- p = ß = 0) it is clear that the mean time to detection is exactly the same as mean time between inspections (i.e. $\mu_i = \mu_d$). Restating this point for clarity – assuming perfect inspections leads logically to the point that any diversion will be detected at the next inspection. The key assumption of this paper (KA1) is that diversion is not a random act, it is a deliberate attempt by a diverter to defeat the safeguards system. Under KA1 it is assumed that the diverter will seek to maximise μ_d hence it is conservatively assumed that the diverter will divert immediately after an inspection and $\mu_i = \mu_d$.

Allowing for more realistic, imperfect inspections represents the more reasonable case in which p<1 and ß>0. In that case the calculation of μ_d has to be re-examined ($\mu_i \neq \mu_d$). If p<1 then μ_d is no longer simply equal to ψ / ί. The new formula becomes:

$$\mu_d = (\psi / ί) / p$$
$$= \mu_i / p$$

If p=0.5 with 12 inspections per year out of 365.25 inspection opportunities then:

$$\mu_d = (\psi / ί) / p$$
$$= (365.25/12)/0.5$$
$$= 60.9 \text{ (cf 30.4 for p=1)}$$

Average Probability of Detection

In simple terms – the detection interval is the number of days from a diversion to its detection. As detections can only take place at inspections, the detection interval is a function of the inspection interval and the probability that the diversion will be detected at any given inspection (i.e. the average probability of detection).

Allowing for the use, in the safeguards criteria, of one set of quite restrictive values for p and ß for PIV inspections (p_{piv}) and a less restrictive set of values for p and ß in the case of interim inspections for timeliness purposes ($p_{interim}$) and assuming that there is never more than one PIV inspection per year, the average detection probability at any one inspection can be approximated over all inspection years as:

$$p_{av} = (p_{piv} + (ί-1) \times p_{interim})/ί$$

Cumulative Detection Probability

The variable p_{av} is intended to approximate the average detection probability for each inspection across all inspection years – it is not intended to represent the cumulative detection probability across the inspection year p_{cum}. The value for p_{cum} for any one inspection year is given by the following formula:

$$p_{cum} = 1 - (1-p_{piv}) \times (1- p_{interim})^{(ί-1)}$$

It is possible to approximate the value for p_{cum} in terms of p_{av} as follows

$$p_{cum} = 1 - (1-p_{av})^{ί}$$

Traditional Safeguards

As a basis for comparison for the discussion that follows it should be noted that traditional safeguards inspection schemes for facilities with UDU have an expected mean time to detection of 61 days – this figure has a standard deviation of 30.4 days[4]. The cumulative probability of detecting a diversion within 61 days is equal to (roughly) 78% (using the formula for p_{cum} as a function of p_{av}). Figure 1 provides a plot of the cumulative probability of detection and detection probability for idealised inspection years under traditional safeguards (the inspection year used in the Figure 1 has been rounded to 360 days to ensure integer results).

It can be seen from Figure 1 that, under the idealised case[5], detections are only possible at the fixed inspection interval of 30 days and that there is an average 47% cumulative probability that the detection interval will exceed the 30-day timeliness goal. There is also an average 22% cumulative probability that detection interval will exceed the design detection goal as expressed in the defined detection probability for the inspection (61 days).

Figure 1 - Plot of cumulative probability and % detection probability at a given inspection for an idealised traditional safeguards inspection scheme for a UDU facility (plot of % probability vs. detection interval (in days)).

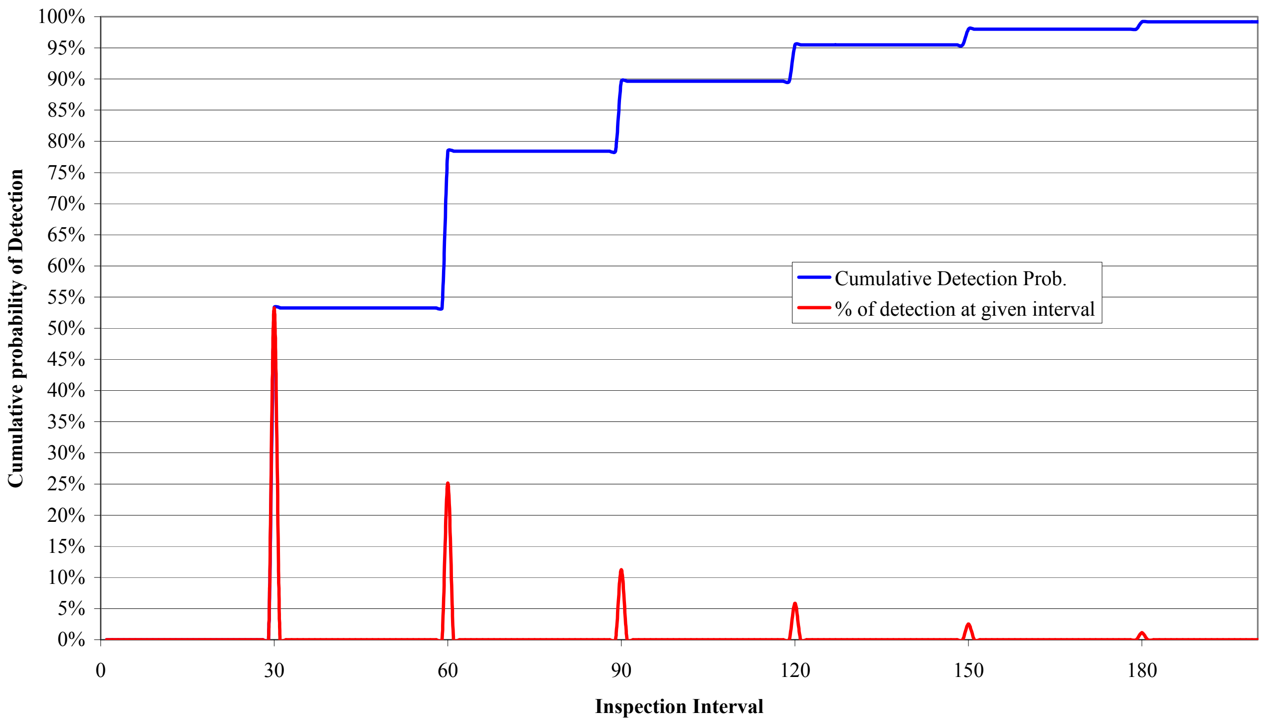

4. Based on a 365.25 day inspection year with a 12 inspections per year with a fixed inspection interval averaging to 30.4 days and an average detection probability p_{av} of 53.3% (based on 1 PIV with p=90% and 11 interim inspections with p=50%).
5. The idealised inspection year used in Figure 1 assumes: a timeliness goal of 30 days; inspections intervals of exactly 30 days, a PIV detection probability of 90%; and an interim inspection detection probability of 50% on an idealised 360-day inspection year.

Figure 2 - **Plot of <u>inspection intervals</u> for various inspection schemes with fixed date for PIV and SNRI for all other inspections, both with (λ) and without ($\bar{\lambda}$) a defined lower boundary condition of a 7 day minimum interval.**

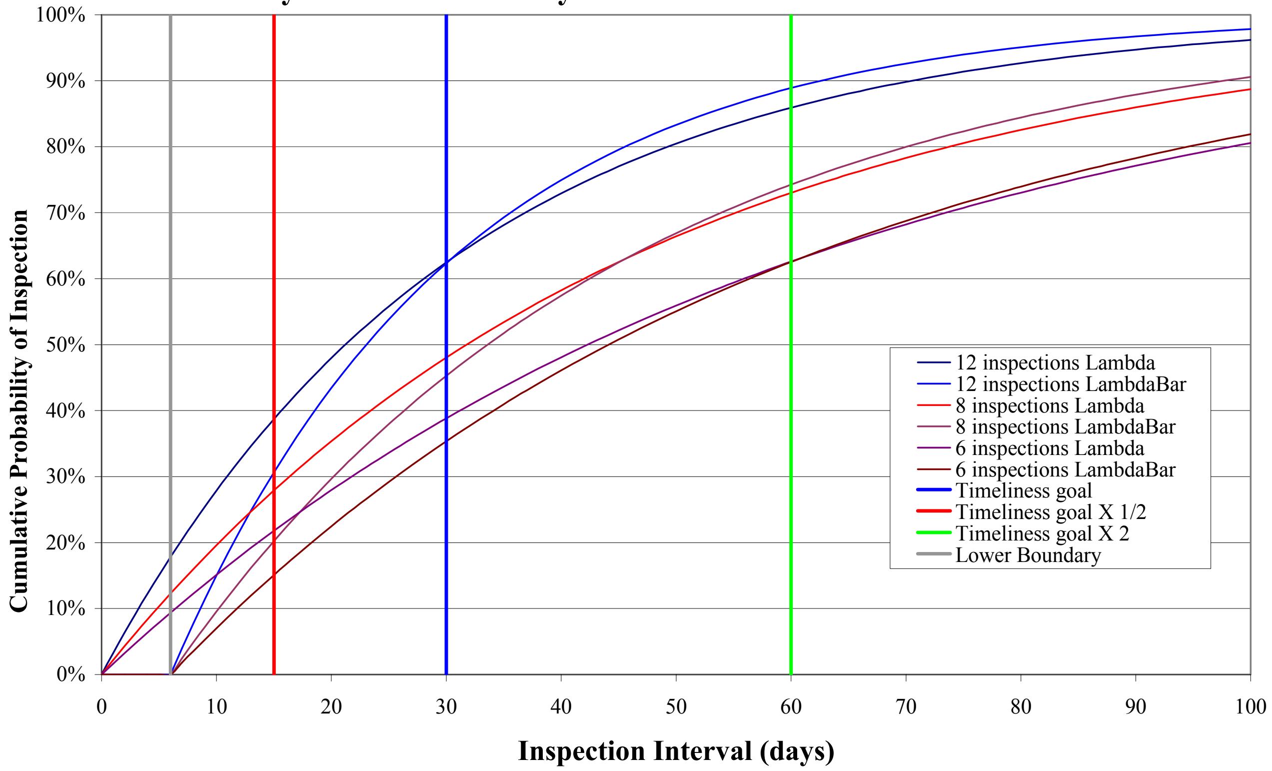

Figure 3 - **Plot of <u>detection intervals</u> (based on p_{av}) for various inspection schemes with fixed date for PIV and SNRI for all other inspections, with (λ) and without ($\bar{\lambda}$) a defined lower boundary condition of a 7 day minimum interval.**

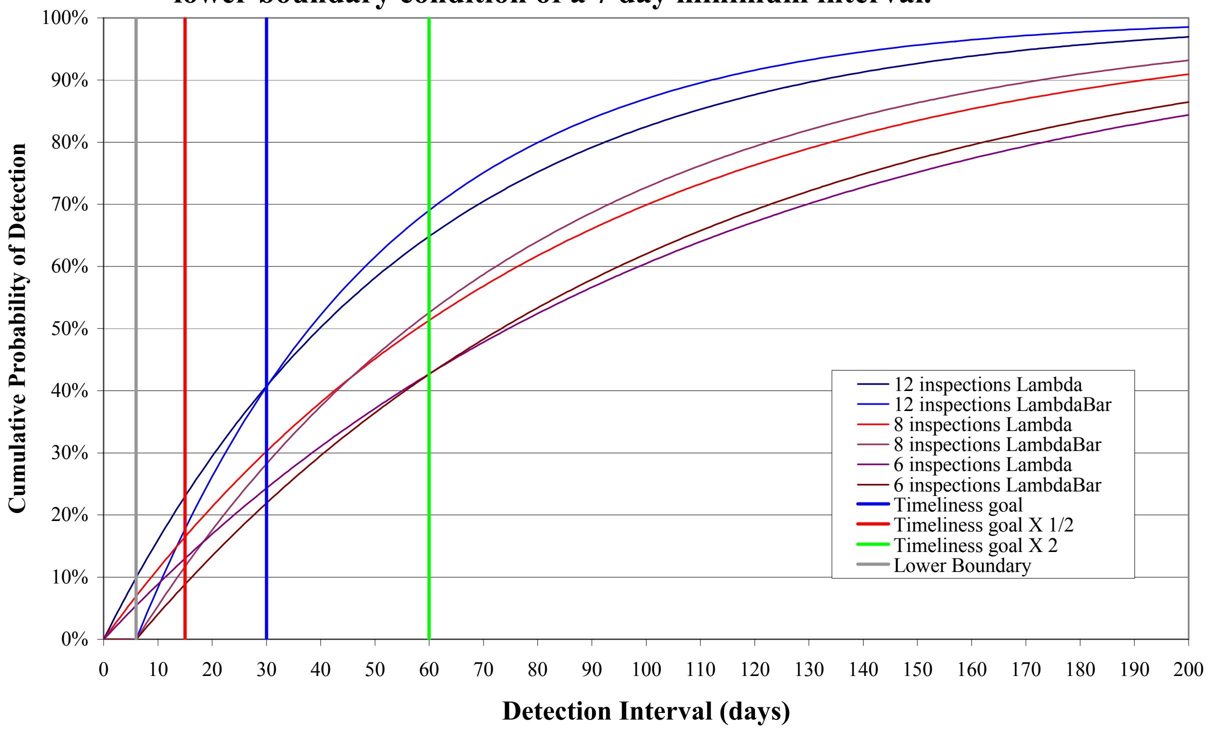

When comparing SNRI schemes with traditional inspection approaches it is important to acknowledge that there are limitations on the effectiveness of traditional safeguard approaches. Under traditional safeguards approaches to UDU facilities, there is a significant probability (approximately 22%) of any diversion being undetected for periods longer than 61 days.

Short notice random inspections

The attractions of randomised inspection schemes include that:

- they offer the possibility of inspection intervals shorter than traditional timeliness goals;
- they offer the possibility of detection intervals shorter than traditional timeliness goals; and
- they offer the possibility of significant resource savings when a reduced number of randomised inspections is used to replace traditional fixed inspection schemes.

As can be seen from Figure 2, with SNRI, there is a significant probability of inspection intervals of shorter than one half of the traditional timeliness goal, even in cases where there are, on average, half the traditional number of inspections conducted per year (i.e. 6 inspections per year).

Figure 3 also shows that there is a significant probability of achieving detection intervals equal to or less than the timeliness goal with each of the inspection schemes shown.

Table 1 - Table of cumulative detection probabilities for: a traditional safeguards approach; an approach with 8 inspections per year (1 fixed PIV and 7 SNRI) and an approach with 6 inspections per year (1 fixed PIV and 5 SNRI)) for a UDU facility.

Inspection Interval (days)	Traditional Safeguards	1 fixed PIV + 7 SNRI	1 fixed PIV + 5 SNRI
15	0%	11%	9%
30	53%	28%	22%
60	78%	53%	43%
90	89%	69%	58%
120	95%	79%	69%
150	98%	86%	77%
180	99%	91%	83%

Table 1 illustrates the different cumulative detection probabilities for three different approaches. While the traditional safeguards case provides a higher probability of achieving a detection interval of 61 days or less (based on an average inspection interval of 30.4 days and a p_{av} value of 53%), it does not provide any capability of achieving a detection interval of less than 30 days. In choosing among the options presented above, a policy decision would be needed as to the weight that should be given to the probability of detecting a diversion of UDU at intervals less than the traditional timeliness goals for the material in question.

3. CONCLUSIONS

Conclusion 1: The use of SNRIs provides the possibility of detecting diversion of UDU at intervals shorter than the timeliness goal for the material in question.

Conclusion 2: When combined with policy decisions as to acceptable levels of the probability of achieving specific timeliness goals, mean time to detection calculations provide a useful guide to inspection decision making.