Multi-stratum Detection Probability Calculations for IAEA Safeguards: Foundations and Early Progress

Aaron M Bevill*, Thomas Krieger†, Robert Binner*, and Claude F Norman*

*International Atomic Energy Agency, Vienna, Austria

†Forschungszentrum Jülich GmbH, Jülich, Germany

a.bevill@iaea.org

Abstract:

IAEA safeguards experts use detection probability (DP) as the primary effectiveness metric for nuclear material inventory and flow verification activities. Most commonly, the DP is defined as the probability of identifying at least one defective item in a population, assuming that one significant quantity (SQ) of material has been diverted. The DP is calculated over a spectrum of diversion scenarios (from a few items with gross defects to a large number of items with bias defects), and the worst-case (lowest) DP is reported. Traditionally, safeguards verification activities are performed on the basis of stratified inventories or flows, whereby the material is grouped into strata on the basis of similar physical and chemical characteristics. On the basis of the measurement tools available to verify the material in the stratum, a separate sampling plan is established for each stratum with the aim of achieving a defined DP of detecting the diversion of 1 SQ from the stratum. However, at facilities, sites, and sectors in which multi-stratum diversion scenarios are plausible, it is desirable to calculate the worst-case DP for scenarios in which 1 SQ is diverted from among the various strata. This paper shows that the multi-stratum DP problem can have multiple local minima, so a global search algorithm is appropriate. The solution is estimated by discretizing the amount of material diverted from each stratum. The discretized problem is solved using a brute force approach and a dynamic programming approach. The two approaches calculate similar answers, but dynamic programming is faster for problems with many strata. This work provided background and early approaches that Annadevula et al. have refined and expanded upon in their concurrent publication [2]. Future work is planned to improve the trade-off between discretization error and calculation time.

Keywords: IAEA; safeguards; diversion detection; detection probability; constrained optimization

1 INTRODUCTION

The International Atomic Energy Agency (IAEA) has established the “timely detection of diversion of significant quantities of nuclear material … and deterrence of such diversion by the risk of early detection” as a safeguards objective [1]. In practice this objective is met through nuclear material accountancy, backed by inspections verifying the accountancy declarations’ accuracy. Although it is impossible to verify every declared item with perfect sensitivity, the inspections achieve the safeguards objective—risk of early detection—using random sampling. The risk is statistically quantified using Detection Probability (DP) and other metrics.

The DP considered here is the probability of inspectors identifying at least one defective item from the declared population, assuming that one significant quantity (SQ) of material has been diverted. A “defective item” would be one from which nuclear material has been diverted, with the diversion masked by over-statement of its contents. (Alternatively, the contents would be under-stated to mask a future diversion from a received item.) DP calculations are used for both inspection planning (to calculate sample sizes that are sufficient to achieve a target DP) and post-inspection effectiveness evaluations.

In safeguards inspections, declared inventories and flows are grouped into strata on the basis of similar physical and chemical characteristics. The goal of stratification is to cluster items with similar metrological properties and similar amounts of nuclear material. For example, large items are ideally stratified separate from small items, and heterogeneous material separate from homogeneous material. After stratifying the declared inventory or flow, inspectors establish a sampling plan for each stratum. The sampling plan should achieve a defined DP for all plausible scenarios in which 1 SQ is diverted from the stratum.

In some facilities, material could plausibly be diverted from multiple strata, combined to constitute one SQ, and processed along a single acquisition path. In fact, diversion spanning multiple facilities, sites, and sectors is possible. Since the single-stratum diversion scenarios are special cases among the multi-stratum scenarios, the worst-case multi-stratum DP can never be higher than any included single-stratum DP. It is desirable to have an approach to calculate the DP for these multi-stratum scenarios, especially if the calculation can be performed on a modern laptop in a few minutes.

This paper describes the foundations and early progress on multi-stratum DP calculations. The single-stratum DP calculation used by the IAEA is described as background. Then the multi-stratum DP calculation is defined, including a description of why the calculation is challenging. Brute force and dynamic programming solutions are demonstrated for small example problems and compared in accuracy and computation time.

The foundations and early progress documented in this paper are developed further in two concurrent publications [2][3].

2 BACKGROUND: Single-stratum Detection Probability Calculation

The formulas for calculating the single-stratum detection probability are briefly re-derived here as background for the multi-stratum calculation. Notation and formulas are adapted from relevant references [4, Annex B][5, Ch. 6][6].The single-stratum DP calculation is typically explained as the combined probability of selection and identification events.

2.1 Identification Probability

The identification probability is the probability that a defective item would be identified as defective, if it were selected for verification. The N items within a stratum are assumed to have the same amount of declared nuclear material, $\bar{x}$. (Refer to the equal mass assumption in Reference [6].) The divertor's goal is to divert an amount M material from the stratum, usually 1 SQ. Therefore $\frac{\bar{x}}{M}$ can be referred to as *goal per item* or *SQ per item*.

Items within the stratum are also assumed to be verified under similar measurement conditions. Verifications by method m are generally assumed to follow a Gaussian distribution with mean $\bar{x}$ and relative standard deviation δ_m. Usually m ranges from 1 to 3 because up to three measurement methods may be applied per stratum. The use of a **relative** standard deviation implies a multiplicative error model.

To divert M, exactly r items must be defected by an amount $\frac{M}{r}$. (Refer to the equal diversion hypothesis of Reference [6].) The number of defects r could be any integer from $\lceil M/\bar{x} \rceil$ to N (or some large number beyond which concealment would be implausible). The defects could be concealed by over-statement of declared inventories and outgoing shipments or by under-statement of received shipments. In the case of over-statement, defective items have true content of $\overline{x} - \frac{M}{r}$. Therefore verifications of defective items will follow a Gaussian distribution with mean $\overline{x} - \frac{M}{r}$ and standard deviation $\delta_m \left(\overline{x} - \frac{M}{r}\right)$. Per IAEA procedure, the defect is identified if the measured value is less than $3\delta_m$ below the declared value $\overline{x}$.

For under-statement, the defective items have true content of $\overline{x} + \frac{M}{r}$. Therefore verifications of defective items will follow a Gaussian distribution with mean $\overline{x} + \frac{M}{r}$ and standard deviation $\delta_m \left(\overline{x} + \frac{M}{r}\right)$. Per IAEA procedure, the defect is identified if the measured value is greater than $3\delta_m$ **above** the declared value $\overline{x}$.

From these Gaussian distributions, the identification probability is given by [6]

$$P_{\text{ident}}(\delta_m, M, \overline{x}, r) = 1 - \Phi\left(\frac{3\delta_m - \frac{M}{\overline{x}r}}{\left(1 \mp \frac{M}{\overline{x}r}\right)\delta_m}\right),$$

where Φ is the cumulative distribution function of the Gaussian distribution and $\mp$ indicates subtraction for over-statement and addition for under-statement. Often the *non*-identification probability $\beta_{\text{ident}} \equiv 1 - P_{\text{ident}}$ is preferred.

2.2 Selection Probability

Selection probability refers to the probability of selecting at least one defective item for quantitative verification. It is possible to select multiple defective items, thereby increasing the chance of detecting the diversion. Therefore the concept of selection probability is slightly broadened to consider the distribution of how many defective items could be selected.

If a single verification method is used to verify n of N declared items, sampled randomly without replacement, then the selection probability follows a hypergeometric distribution. The probability of selecting exactly i of r defective items is

$$P_{\text{select}}(i \mid N, r, n) = \frac{\binom{r}{i}\binom{N-r}{n-i}}{\binom{N}{n}},$$

where $\binom{p}{k}$ is the binomial coefficient with domain $0 \le p \le k$. Outside this domain, $P_{\text{select}} = 0$.

Generalizing to three methods, if n_m items are verified with method m, then the joint probability of selecting exactly i_m defective items for verification with method m follows a joint distribution:

$$P_{\text{select}}(i_1, i_2, i_3 \mid N, r, n_1, n_2, n_3) = P_{\text{select}}(i_1 \mid N, r, n_1)\, P_{\text{select}}(i_2 \mid N - n_1, r - i_1, n_2) \ldots$$
$$\ldots P_{\text{select}}(i_3 \mid N - n_1 - n_2, r - i_1 - i_2, n_3)\,.$$

This formula could be expanded for an arbitrary number of methods.

2.3 Detection Probability

The single-stratum detection probability is the probability that at least one defective item will be quantitatively verified and the defect identified. It is simpler to express the *non*-detection probability: the probability that none of the defective items in the sample is identified as defective. This is a weighted average of $\beta_{\text{ident},m}^{i_m}$ over the P_{select}:

$$DP(M,r) = 1 - \sum_{i_1=0}^{n_1} \sum_{i_2=0}^{n_2} \sum_{i_3=0}^{n_3} P_{\text{select}}\,(i_1, i_2, i_3 \mid r)\; \beta_{\text{ident},1}(M,r)^{i_1}\; \beta_{\text{ident},2}(M,r)^{i_2}\; \beta_{\text{ident},3}(M,r)^{i_3}\,,$$

when a purely random error model is assumed; see Reference [6].

Note that all terms depend on r, the number of defects. (The other parameters have been omitted for brevity.) It is assumed that the diverting party will choose r (within the range specified above) so as to minimize the detection probability. Therefore the "worst-case" DP is the conservative choice:

$$DP^*(M) \equiv \min_r DP(M,r).$$

The $DP(M,r)$ function is neither monotonic nor convex with respect to r, so minimizing $DP(M,r)$ is not trivial. Usually $DP(M,r)$ is evaluated for many values of r. This is effective because r is an integer variable and $DP(M,r)$ is relatively smooth.

Usually $DP^*(M)$ is evaluated for each stratum independently, with the goal of $M = 1$ SQ for each stratum. This is used, for example, to calculate sample sizes or evaluate inspection effectiveness. In some cases however, a multi-stratum DP calculation would also be desirable.

3 PROBLEM DESCRIPTION: Multi-stratum Detection Probability Calculation

The multi-stratum DP calculation extends the single-stratum calculations by allowing a variable goal M_s for each stratum s. Multiple strata $s = 1, \ldots, S$ are simultaneously considered. The *aggregate* detection probability (ADP) is the probability that the diversion is detected in at least one of the strata. It is calculated as one minus the product of the non-detection probabilities:

$$ADP(M_1, \ldots, M_S) = 1 - (1 - DP_1^*(M_1)) \ldots (1 - DP_S^*(M_S))\,.$$

It is assumed that the diverting party would choose each M_s so as to minimize the aggregate DP:

$$ADP^* = \min_{M_1,\ldots,M_S} ADP(M_1, \ldots, M_S)\,.$$

The minimization is constrained by the need to divert some aggregate quantity (usually 1 SQ) and by the nuclear material available in each stratum. Stated more formally, the constraints are

$$0 \le M_s \le N_s \bar{x}_s \quad for\ all\ s = 1, \ldots, S \quad \text{and} \quad \sum_s M_s \ge 1\ \text{SQ}\,.$$

By construction, ADP^* cannot be higher than the smallest DP^* of the aggregated strata. (Scenarios taking 1 SQ from an individual stratum are included in the ADP^* calculation.) In concurrently published work, Krieger et al. investigated how much lower ADP^* can be than DP^* [3].

Like the single-stratum DP^* calculation, calculating ADP^* is a minimization of detection probability over a parameter that the diverting party controls. However, the ADP^* calculation is significantly more challenging—not least because it requires numerous evaluations of DP^*.

3.1 Calculation Challenges and Opportunities

This section discusses properties of the ADP function. Some of these properties raise the difficulty of calculating ADP^*; others present solution opportunities that are not possible for other optimization problems.

The single-stratum $DP^*(M)$ is a discontinuous function of M. It is also non-smooth, in the sense that its derivatives are discontinuous. As M increases, the optimal integer r abruptly changes. Discontinuities in DP^* and its derivatives occur where the optimal r changes. These properties are demonstrated in a later example. The aggregate detection probability ADP depends directly on DP^*, so it too is non-smooth and discontinuous.

Similarly, DP^* is not a convex function of M, and the non-convexity passes directly to ADP. Convexity means that the second derivative is strictly positive. It would be a useful property because it would imply that the constrained minimization problem has no local minima besides the global minimum.

Because ADP is neither smooth, continuous, nor convex, many off-the-shelf optimization algorithms will be susceptible to converging on a local minimum, if they converge at all [7, p. 954]. This suggests the use of global search algorithms. Unfortunately, global search tends to scale poorly as the number of optimized parameters increases. For example, the brute force approach (described below) scales exponentially as the number of strata increases.

Fortunately, ADP is a separable function of M. More precisely, ADP can be transformed to an equivalent aggregated non-detection probability ($ANDP$), which can be written in the form

$$ANDP(M_1,\ldots,M_S) = 1 - ADP(M_1,\ldots,M_S) = f_s(M_1)\ldots f_S(M_S)\ .$$

Maximizing $ANDP$ is equivalent to minimizing ADP. This separability property is useful because it means that global search approaches can cache various values of $DP_s^*(M_s)$ and then combine the cached values to calculate ADP. It also forms the foundation of the dynamic programming approach described below.

The ADP also increases monotonically as any M_s increases. This property is not exploited in any of the approaches used in this paper, but it is being considered in ongoing work.

To summarize, the aggregate detection probability ADP is a piecewise-discontinuous, non-convex function of the continuous variables $M_1,\ldots,M_S$. But it is separable and monotonically increasing with M, which enables accelerated calculation techniques. Early calculation approaches that leverage these properties are described in the next section.

4 MULTI-STRATUM CALCULATION APPROACHES

This paper describes and demonstrates two early approaches: brute force and dynamic programming. In concurrent and subsequent work, Annadevula et al. have investigated advanced approaches, including a greedy algorithm and a multi-frontier approach [2].

4.1 Brute Force Approach

The brute force approach discretizes the $M_1\ldots M_S$ parameter spaces and then evaluates ADP for every combination. (Combinations summing to less than 1 SQ are discarded.) To compute efficiently, $DP_s^*(M_s)$ can be precalculated at the discretized values of M_s. This precalculation scales linearly with the number of strata (S), and can typically run in a few seconds on a modern laptop.

The brute force combination step is much faster than the precalculation step when S is small, but its complexity scales exponentially with S. If G gridpoints are compared for each of S strata, then G^S combinations must be evaluated. For larger problems, the combinations cannot be calculated in reasonable time on a modern laptop.

4.2 Dynamic Programming Approach

The dynamic programming approach is similar to brute force, but it improves the scaling by combining the strata in pairs. Non-optimal combinations are discarded to reduce the number of combinations in the next iteration. This approach is adapted from dynamic programming commonly used in operations research [7 p. 438].

First $DP_s^*(M_s)$ is precalculated at the discretized values of M_s. Then every possible pair of discretized M_1 and M_2 values are tabulated. Their combined DP is calculated:

$$DP_{1+2}(M_1, M_2) = 1 - (1 - DP_1^*(M_1))\,(1 - DP_2^*(M_2))\,.$$

This results in a list of all of the ways that the (discretized) strata 1 and 2 could be combined.

There are multiple pairs of M_1, M_2 that sum to similar amounts of material. The calculation only needs to keep the pair with the minimum DP_{1+2}. The other pairs could not possibly result in a lower ADP because ADP is separable. (Holding $M_1 + M_2$ constant, the pair that minimizes DP_{1+2} also minimizes ADP, regardless of how material is diverted from the other strata.)

This removal is mathematically described as a Pareto frontier. A pair (M_1, M_2) is "Pareto dominated" by another pair (M_1', M_2') if two criteria are met: $M_1 + M_2 \le M_1{}' + M_2{}'$ and $DP_{1+2}(M_1, M_2) \ge DP_{1+2}(M_1{}', M_2{}')$. Any pair dominated by another pair is removed. By taking the Pareto frontier, the combinations of strata 1 and 2 are reduced into a single "super-stratum" $1 + 2$, with DP_{1+2}^* evaluated at discrete values of M_{1+2}.

This process of combining and culling is repeated inductively. Stratum $1 + 2$ is combined with stratum 3 to form stratum $1 + 2 + 3$; stratum $1 + 2 + 3$ is combined with stratum 4 to form stratum $1 + 2 + 3 + 4$; and so forth. The number of combinations is similar at each step if the grids for each stratum are identical and regularly spaced. Therefore this approach scales linearly with respect to the number of strata.

This dynamic programming approach appears to find the same solution as the brute force approach, only faster. Formal proof that the dynamic programming approach always identifies the global minimum ADP remains for a longer follow-up publication.

Both approaches suffer from the same limitation: the discretization of M_s forces a tradeoff between solution precision and calculation time. If M_s is discretized with G grid-points in each stratum, then the $DP^*(M_s)$ precalculation complexity scales proportional to G. For the faster dynamic programming approach, the precalculation dominates the calculation time. To improve this tradeoff, mesh refinement approaches will be considered in future work.

It is also notable that any discretization error from the finite grid size always results in an over-estimate of ADP^*. A conservative lower-bound on ADP^* will be pursued in future work as well.

5 EXAMPLES AND RESULTS

The brute force and dynamic programming approaches are demonstrated on two problems below. Informal calculation times are noted to indicate order of magnitude. Precalculations of DP^* are parallelized on two cores; combination approaches are performed on a single core.

The uranium enrichment example includes details of the various calculation steps. The plutonium example is solved with varying number of gridpoints (G) to demonstrate how G affects the solution accuracy and calculation time.

5.1 Uranium Enrichment Plant

Material at a large uranium enrichment plant could be stratified into three strata: depleted (UFD), enriched (UFE), and natural (UFN) UF_6 cylinders. Parameters for each of these strata are specified in Table I. The parameters are synthetic but are representative of real-world calculations.

For single-stratum calculations, the diverting party would obtain 1 SQ of high enriched uranium by diverting and enriching 20000 kg U from the UFD inventory, 75 kg ^{235}U from the UFE inventory, or 10000 kg U from the UFN inventory. The multi-stratum calculation considers combining fractional diversions from among these strata, e.g. 0.5 SQ = 10000 kg U from UFD combined with 0.5 SQ = 37.5 kg ^{235}U from UFE. This is a drastic simplification of the acquisition paths, but it creates a clear example of the complexity of even small multi-stratum DP calculations.

TABLE I. Stratum specifications for the uranium enrichment example.

Variable	***Definition***	***UFD***	***UFE***	***UFN***
s	Stratum	1	2	3
M	Goal quantity, 1 SQ	20000 kg U	75 kg ^{235}U	10000 kg U
$\overline{x}$	Declared nuclear material content per item	5333 kg U	60 kg ^{235}U	8000 kg U
N	Number of declared items	700	500	600
r	Number of defective items	4 to 700	2 to 500	2 to 600
δ_1	Relative standard deviation of method 1	15%	15%	15%
δ_2	Relative standard deviation of method 2	10%	2%	5%
δ_3	Relative standard deviation of method 3	0.5%	0.1%	0.2%
n_1	Number of method 1 measurements	14	25	40
n_2	Number of method 2 measurements	10	12	20
n_3	Number of method 3 measurements	5	7	4

5.1.1 Single-stratum Detection Probability

Sample size calculations would normally be based on individual calculations of $DP^*(M_s, r)$, with $M_s = 1$ SQ. These curves are plotted in Figure 1. To reduce calculation time, some values of r are omitted; up to 50 logarithmically spaced integer values of r are used in each stratum. Calculation of these curves required 0.6 seconds.

5.1.2 Precalculation

To begin the multi-stratum DP calculation, $DP(M_s, r)$ was calculated at $G = 100$ equally spaced values of M_s. The results are plotted in Figure 2. This precalculation required 25 seconds.

The precalculation results demonstrate some of the key challenges with the multi-stratum DP calculation. Here we refer to $DP^*(M_s)$, the minimum $DP(M_s, r)$ for any given M_s. First, as clearly seen for stratum UFE, $DP^*(M_s)$ is non-convex; the second derivative is both positive and negative. Second, it is not smooth where $DP^*(M_s)$ transitions from one value of r to another. There, the derivatives abruptly change. Third, it is discontinuous at $M_s = 0.8$ SQ because the $r = 1$ curve ends. (Indeed it is not possible to divert more than $\overline{x} = 60$ kg ^{235}U from a single item.)

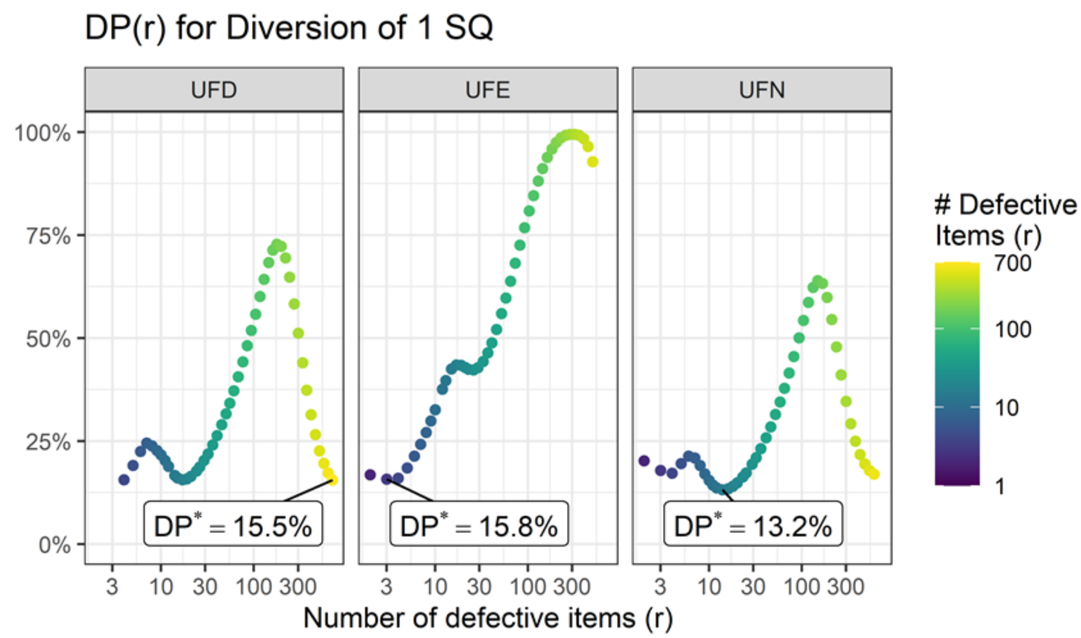

Figure 1. The single-stratum detection probability for 1 SQ diversions from each of the three strata in the uranium enrichment example. The worst-case DP^* is marked for each stratum.

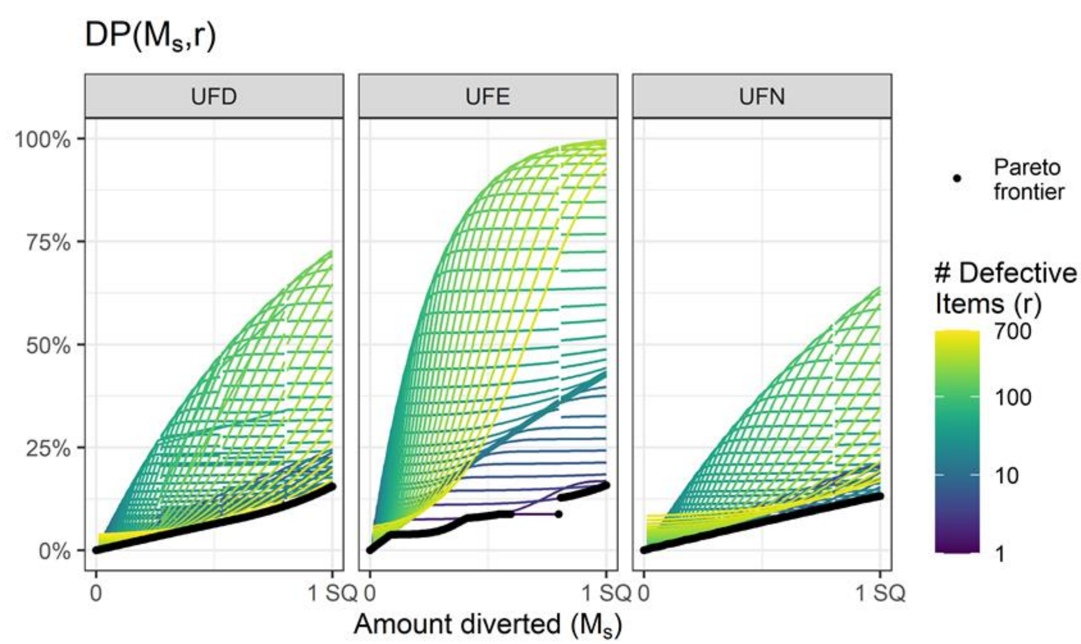

Figure 2. The precalculation results for the uranium enrichment example. Note that the horizontal axis is now the quantity diverted (M_s), and the number of defective items (r) is only on the color scale. The Pareto frontier (lowest $DP(M_s, r)$ for a given M_s) is marked with black dots (many of which overlap). The Paerto frontier at 1 SQ is DP^* in Figure 1.

From each stratum, the points with the lowest $DP(M_s, r)$ for a given M_s are identified. Calculating these Pareto frontiers required less than 0.05 seconds. The Pareto frontier points are marked as black dots in Figure 2.

Note that there are no frontier points for the UFE region $M_s = 0.5$ to 0.8 SQ. In this region, DP^* increases less than machine precision. Given the negligible increase in DP^*, it is reasonable to assume that the diverting party would choose the largest M_s. The other values are excluded from the frontier.

The frontiers were then combined using the brute force and dynamic programming approaches.

5.1.3 Brute Force Solution

The brute force solution evaluated every combination of the strata's Pareto frontiers. Since the frontiers contain 100, 81, and 100 points (respectively), the brute force calculation requires 810000 evaluations. The evaluations required 15 seconds.

The resulting ADP^* is 11.2%, achieved by diverting 0.202 SQ from UFD and 0.798 SQ from UFE. Refer to detailed results in Table II.

Table II. Brute force results of the uranium enrichment example.

Stratum	***Amount diverted, $\boldsymbol{M_s}$***	$\boldsymbol{DP^*(M_s)}$
UFD	0.202 SQ = 4040 kg U	2.6%
UFE	0.798 SQ = 59.8 kg ^{235}U	8.8%
UFN	0 SQ = 0 kg U	0%
Aggregate	1 SQ	11.2%

5.1.4 Dynamic Programming Solution

Using the dynamic programming approach, strata 1 and 2 are combined into stratum $1 + 2$. Every gridpoint from stratum 1 is combined with every gridpoint from stratum 2; see Figure 3. Combinations that are not on the Pareto frontier are removed. Pairs above 1.1 SQ are also removed, because they are unlikely to be relevant for calculating the 1 SQ ADP^*. This effectively calculates DP^*_{1+2} for the combined stratum $1 + 2$.

Then stratum $1 + 2$ is combined with stratum 3 to form combined stratum $1 + 2 + 3$. The process could be repeated for any additional strata. The final ADP^* is the smallest DP^*_{1+2+3} for a diversion of $M_{1+2+3} \geq 1$ SQ. The results match the brute force calculation, but the calculation time is less than 0.3 seconds.

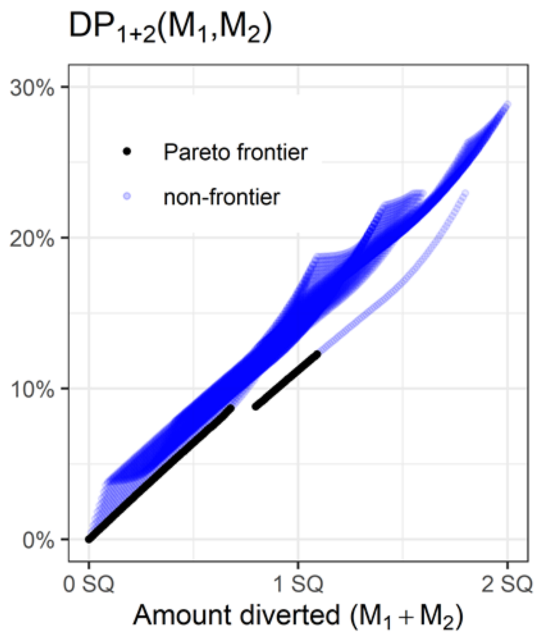

Figure 3. The dynamic programming approach solves the uranium enrichment example by evaluating all combinations from strata 1 and 2 (left, blue) and keeping only the Pareto frontier (black). Combinations are drawn as dots, but many overlap. This process is repeated to combine the $1 + 2$ frontier with stratum 3 (right).

5.1.5 Discussion

The two approaches calculate identical results, but the dynamic programming approach combines the strata 50 times faster. In this particular problem, the precalculation takes longer than either combination approach.

As expected, the $ADP^* = 11.2\%$ is lower than even the lowest single-stratum DP^*, 13.2%. Interestingly stratum UFN has the lowest DP^* but did not contribute any material to ADP^*. By combining UFD and UFE, the detection probability drops from 15.5% to 11.2%.

5.2 Plutonium Example

Table III lists parameters for a plutonium example with four strata. Each stratum is verified with one or two methods.

The single-stratum $DP^*(M_s)$ are calculated for all integer values of r in the specified range. For a diversion of $M_s = 1$ SQ, the single-stratum DP^* is at least 90.8% for every stratum.

The multi-stratum detection probability ADP^* was calculated with various grid sizes (G). The calculated ADP^* and calculation times are plotted in Figure 4. The brute force and dynamic programming approaches calculated the same ADP^* in all cases. However, the calculated ADP^* decreases from 90.8% (the lowest single-stratum DP^*) to 84.6% as the grid is improved. Unfortunately the grid cannot be refined infinitely because the calculation time becomes impractical.

Table III. Parameters for the plutonium example.

Variable	*PU1*	*PU2*	*PU3*	*PU4*
s	1	2	3	4
M	8 kg Pu	8 kg Pu	8 kg Pu	8 kg Pu
$\overline{x}$	0.4043 kg Pu	8.929 kg Pu	3.501 kg Pu	3.391 kg Pu
N	23	7	94	11
r	20 to 23	1 to 7	3 to 94	3 to 11
δ_1	1.48%	1.26%	3.76%	1.47%
δ_2	—	0.86%	0.54%	—
n_1	3	6	48	6
n_2	—	1	3	—

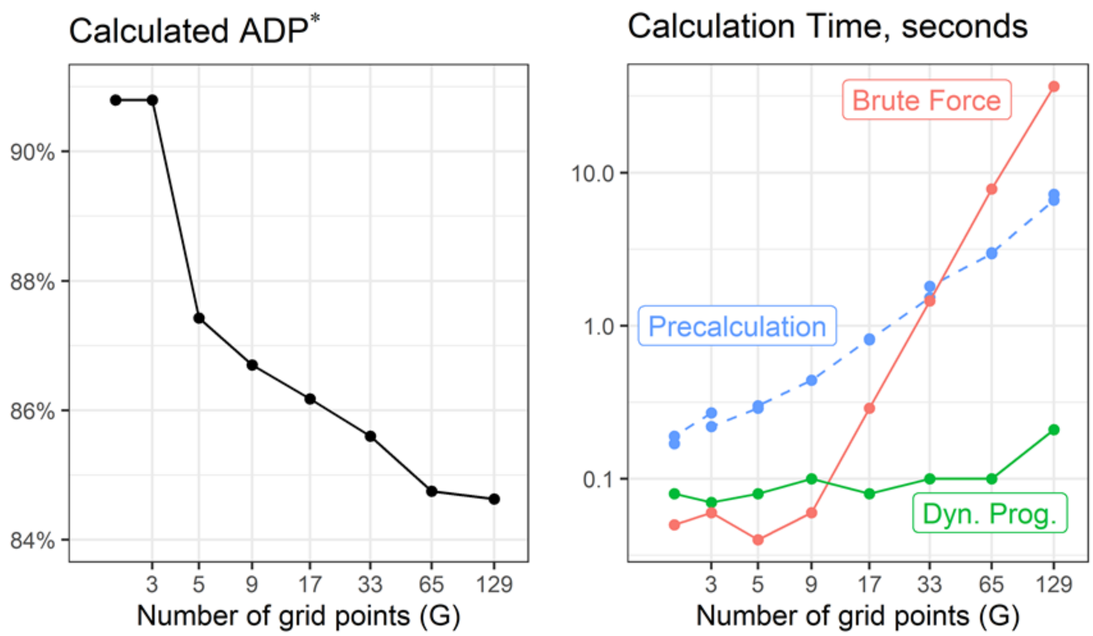

Figure 4. The aggregate detection probability ADP^* and calculation times for the plutonium example. Precalculation was performed twice for each value of G.

As expected, the precalculation time increases proportional to the number of gridpoints, G. Section 4.1 predicts that the brute force calculation would scale proportional to G^4 because there are four strata. However, in this example many points are eliminated by keeping only the Pareto frontier after precalculation. The calculation time is observed to scale proportional to the number of combinations of gridpoints, which is effectively $G^{2.4}$ in this example.

The dynamic programming combination is several orders of magnitude faster than the precalculation. Therefore the precalculation currently limits the scalability of ADP^* calculations.

6 CONCLUSIONS

Safeguards evaluations must assume that the diverting party would minimize the probability of detection. For multi-stratum calculations, this means minimizing the aggregate detection probability ADP by adjusting the quantity diverted from each stratum ($M_1, \ldots, M_S$).

The examples demonstrate that ADP is a non-smooth, discontinuous, non-convex function of $M_1, \ldots, M_S$. Many off-the-shelf algorithms would be susceptible to converging on local minima, if they converge at all. In this work, global minimization approaches were used to ensure accuracy.

A brute force approach and a dynamic programming approach were demonstrated. Both approaches precalculate components of ADP for discrete values of each variable $M_1, \ldots, M_S$. The brute force approach then evaluates ADP for every combination of discrete values. This approach requires intractably long calculations for large problems. The dynamic programming approach finds the same solution as the brute force approach, only faster. A publication is planned to formally prove that the dynamic programming approach is always accurate. In more recent work, Annadevula et al. have investigated advanced combination approaches [2].

One key limitation of both approaches is the poor tradeoff between solution precision and calculation time. Both approaches over-estimate ADP^*, and the precision can only be improved by refining the precalculation grid. It is unclear what grid size to prescribe for practical calculations. Research is ongoing to improve the tradeoff and to estimate the precision.

The test problems in this publication are modest in comparison to proposed multi-facility and multi-sector calculations. Future publications are planned to specify larger example problems.

Once ADP^* can be efficiently calculated, it could potentially be used to enhance safeguards evaluations and sample size calculations. For example, fixed inspection resources could be reallocated among strata to maximize ADP^*. Alternatively, sample sizes could be calculated to reach a prescribed ADP^* with minimal use of inspection resources. These applications are being considered as part of an ongoing collaboration for continuous improvement of IAEA safeguards.

7 REFERENCES

1. IAEA, “The Structure and Content of Agreements Between the Agency and States Required in Connection with the Treaty on the Non-Proliferation of Nuclear Weapons,” (*INFCIRC/153, Corrected*), Vienna: IAEA, 1972.
2. L. Annadevula *et al.*, “Review of Methods to Aggregate Diversion Detection Probabilities Across Multiple Material Strata,” submitted concurrently to *Proc. of the INMM & ESARDA Joint Annual Meeting*, 2021.
3. T. Krieger *et al.*, “Multi-stratum Detection Probability: Worst Case Scenarios,” submitted concurrently to *Proc. of the INMM & ESARDA Joint Annual Meeting*, 2021.
4. J. L. Jaech and M. Russell, “Algorithms to Calculate Sample Sizes for Inspection Sampling Plans,” IAEA STR-261, 1991.
5. IAEA, “Statistical Concepts and Techniques for IAEA Safeguards,” IAEA SG-PR-2016 Rev. 5, 1998.
6. T. Krieger *et al.*, “Statistical Methods for Verification Sampling Plans,” IAEA STR-381 draft, 2017.
7. F. S. Hillier and G. J. Lieberman, *Introduction to Operations Research*, 10th ed., McGraw Hill, 2014.