An in-depth case study: modelling an information barrier with Bayesian Belief Networks

Paul Beaumont[1], Edward Day[2], Neil Evans[2], Sam Haworth[2],
Michael Huth[1], Tom Plant[2], and Catherine Roberts[2]

[1] Department of Computing, Imperial College London, London, SW7 2AZ, UK
{paul.beaumont, m.huth}@imperial.ac.uk
[2] AWE Aldermaston, Reading, Berkshire, RG7 4PR , UK
{Neil.Evans, Tom.Plant}@awe.co.uk

Abstract. We present in detail a quantitative Bayesian Belief Network (BBN) model of the use of an information barrier system during a nuclear arms control inspection, and an analysis of this model using the capabilities of a Satisfiability Modulo Theory (SMT) solver. Arms control verification processes do not in practice allow the parties involved to gather complete information about each other, and therefore any model we use must be able to cope with the limited information, subjective assessment and uncertainty in this domain. We have previously extended BBNs to allow this kind of uncertainty in parameter values (such as probabilities) to be reflected; these *constrained* BBNs (cBBNs) offer the potential for more robust modelling, which in that study we demonstrated with a simple information barrier model. We now present a much more detailed model of a similar verification process, based on the technical capabilities and deployment concept of the UK-Norway Initiative (UKNI) Information Barrier system, demonstrating the scalability of our previously-presented approach. We discuss facets of the model itself in detail, before analysing pertinent questions of interest to give examples of the power of this approach.

1 Introduction

Arms control verification processes do not in practice allow for the parties involved to gather complete information about each other.Instead, each must make decisions about whether or not other parties are complying with their obligations, on the basis of limited information. They must also make decisions during negotiation of a verification regime about the measures to be used, and during implementation of that regime about how and when to use the inspection tools at their disposal. Decision-making under uncertainty is therefore a core element of the arms control verification problem. Our work, aims to extend and combine mathematical modelling and verification approaches such that they can cope with the inherent lack of available data in this domain, and potentially be used to support policy-makers in practice of arms control treaty design and implementation.

In our work, we present a quantitative model that is useful in decision support: this should additionally express both the modeler's degree of belief over a proposition, say that "the item under inspection is a weapon", and their confidence in that degree of belief subject to perturbations of parameter values in the model. Working in this way reduces the risk of bias on the part of the decision maker or inspector having a material impact on verification activities. The model encodes a belief-revision framework that is updated as a function of the evidence available: the basis for each decision is recorded, auditable and testable. Further, the intrinsic sensitivity of the model to potential bias in its design can be tested in advance of the conduct of verification activities.

Our unique approach is to model different facets of the beliefs of each party and the various inspections control processes in a software tool known as a Satisfiability Modulo Theories (SMT) solver [11]. This offers a general purpose approach to the automated analysis of mathematical models that can be expressed as a set of mathematical constraints over theories such as those of the real numbers. In our case we use SMT to deal with uncertainty in (or absence of) data in the model by expressing such uncertainty as the under-specification of probabilities of events in model verification processes. In other words, we don't have to choose values for probabilities - such as the likelihood of the other party tampering with the measurement equipment, for example. If we don't know such likelihoods we can pick a range of possible likelihood values, or leave the value totally unconstrained, to gain confidence by analyses that are valid for all such values.

For this approach of under-specified parameter values, we have so far considered Bayesian Belief Networks (BBNs) [1], Game theoretic models [2] and dynamical systems [3]. BBNs allow the representation and analysis of multiple variables, the causal relationships between them, and their associated conditional probabilities. They are particularly useful for making judgments under uncertainty, in representing probabilistic reasoning, and in handling objective and subjective data. This makes them attractive tools in principle for describing and analyzing arms control processes; however, it can be extremely difficult to set appropriate and defensible values for all conditional probabilities, which challenges their use in practice. At the 20th European Symposium on Research in Computer Security (ESORICS2015), four of the authors proposed a methodology to address this by extending the BBN framework to allow representation of sets of BBNs [1]. We refer to this as a *constrained* BBN (cBBN). A cBBN represents a set of Bayesian Belief Networks by symbolically expressing uncertainty about probabilities and scenario-specific constraints that are not representable by a conventional BBN, but can still be symbolically evaluated. This extension retains the advantages of BBNs, allows the incorporation of any scenario constraints, and assesses the robustness of models and their analyses when little or no contextual data are available. Technically, we achieve this by specifying cBBNs in the Satisfiability Modulo Theories (SMT) solver Z3, and by formulating and running confidence and optimization queries using Z3.

Our previous work used a simple BBN representation of an information barrier in an arms control context. We can now present this new case study, using a much more detailed and nuanced model, based on the technical capabilities and concept of operations of the UK-Norway Initiative (UKNI) Information Barrier [8]. Here we focus on the model facets and the technical changes needed to turn this project from a research tool into a system that can demonstrably cope with realistic sized models. This application of our methodology to a more detailed and realistic case illustrates the power of this approach; it can for example compute optimal BBNs for measures of interest (as shown in this paper), judge the influence of uncertainty about correct conditional probabilities on the findings of a given network, or assess the worst-case sensitivity of findings to a particular variable.

2 The Model

Consider the situation of two fictitious nation states (referred to as "nations" below to avoid confusion with the technical notion of "states" of a BBN), **N1** and **N2**.

N2 is tasked with identifying whether an item belonging to **N1** and available to 'inspect' in a controlled inspection facility is a nuclear weapon. The purpose of this inspection within an arms control agreement may be that the item is on its way for decommissioning, destroying, storage, etc. Our mathematical model of the scenario does not reflect what may happen to the material post-inspection, but more detailed models may well reflect this.

The nations' non-proliferation obligations and national security concerns dictate that the design details of the item must be protected, and therefore the inspecting party will have no visual access to the item. Instead the parties agree that the objects that they declare to be weapons contain Plutonium with the isotopic ratio 240Pu:239Pu below a certain threshold value, which they set at 0.1. In order to draw conclusions about whether an item presented for

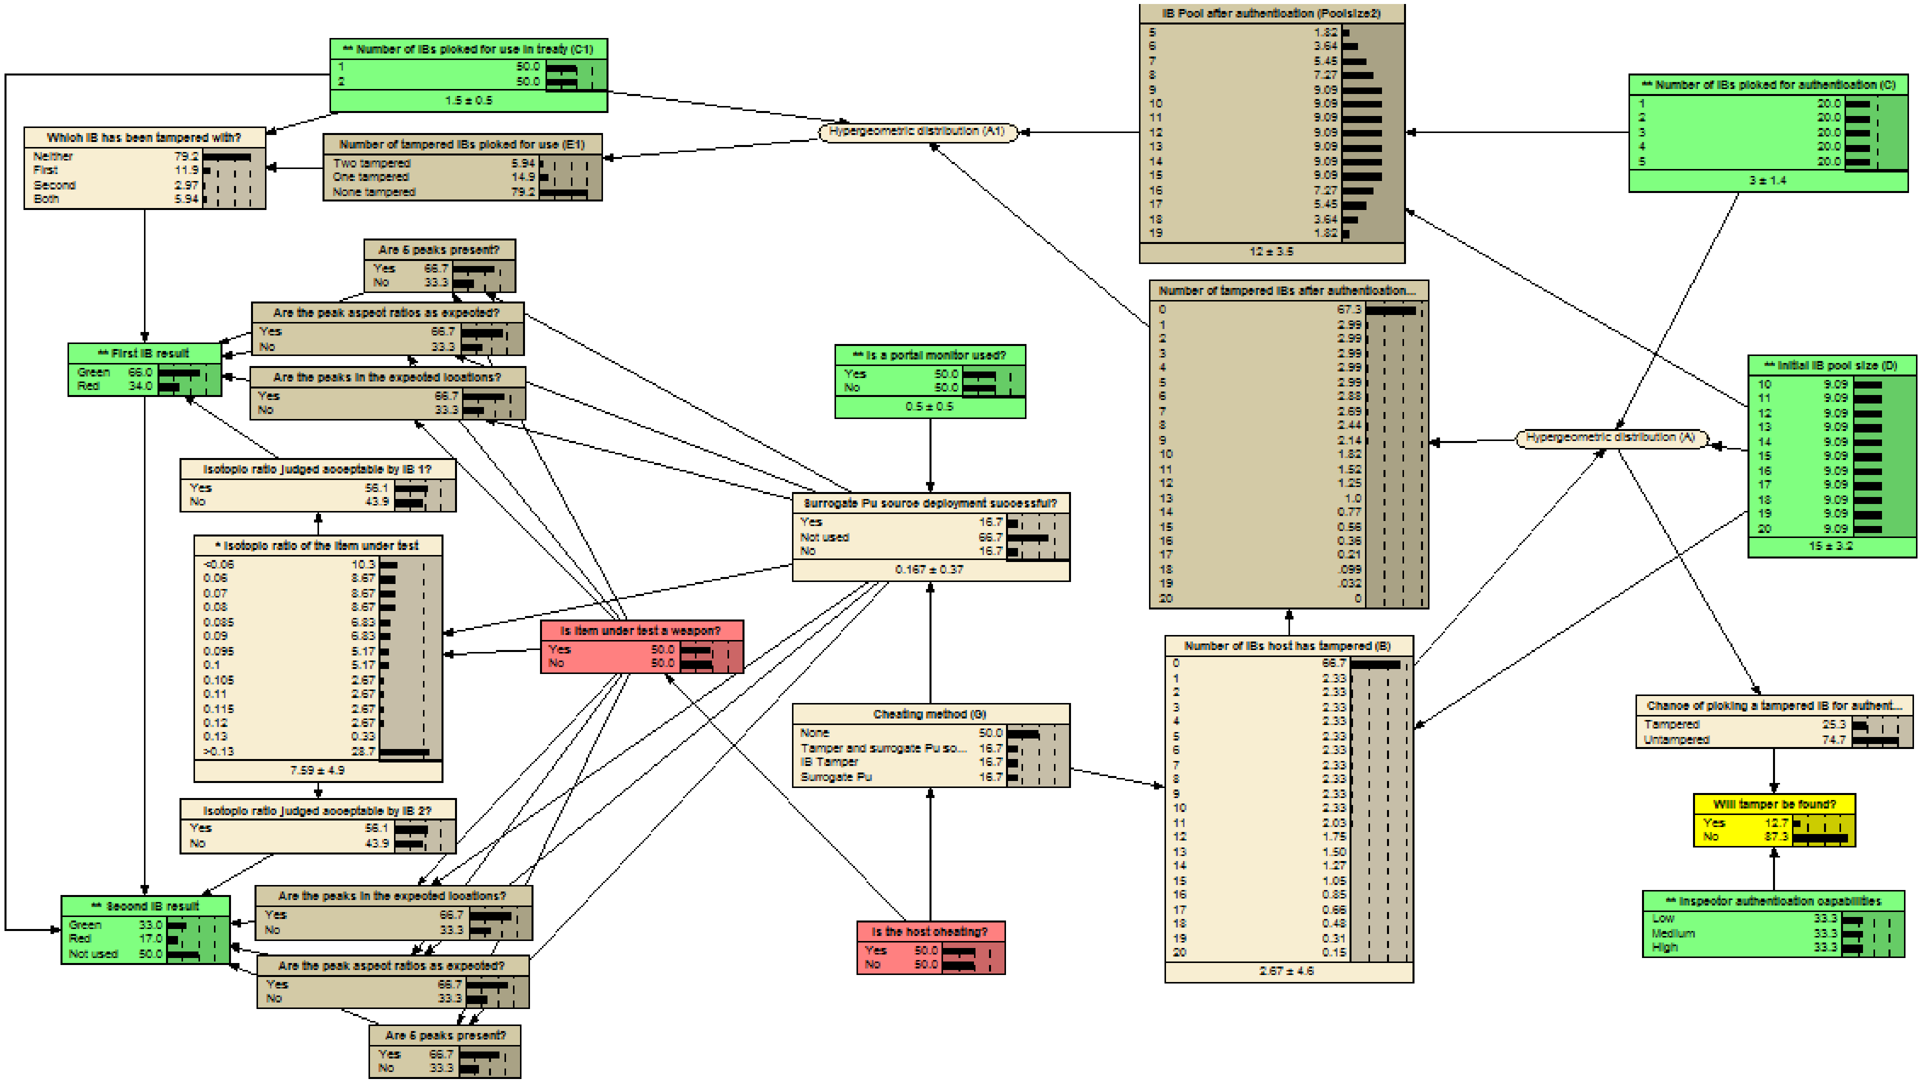

Fig. 1: Detailed BBN model of an arms control inspection that uses an information barrier for obfuscation

inspection is a weapon, the inspecting party uses an information barrier (IB) system comprising a HPGe detector and bespoke electronics with well-understood performance characteristics (see Figure 2a, [9]) to conduct measurements on the object while the object is concealed in a box. The IB system displays a green light if it detects a gamma spectrum indicative of the presence of Plutonium with the appropriate isotopic ratio; if it does not detect this spectrum for whatever reason then it shows a red light. No other information is provided, and therefore weapon design information is protected [8].

Nation **N2** believes that it may be possible for nation **N1** to spoof a radioactive signal (or in some way provide a surrogate) to fool the detector, or that **N1** may have just placed Plutonium with the appropriate isotopic ratio in the box rather than a weapon. These subjective assessments should be reflected in the model alongside the error rates of the IB system. In order to deter cheating, **N2** is allowed to choose the IBs used in the verification from a pool of machines. They are also able to take some unused IBs away for authentication (activities designed to check whether or not the IB has been tampered with) from the same pool, although they are not allowed to take the ones used in verification activities in case there is any residual information to be gained following their use. This selection process is designed to ensure that a nefarious host is held at risk of detection, because in order to tamper with the IBs used in verification it would have to run the risk of one or more tampered IBs being selected for authentication. Although authentication cannot be assumed to be perfect - and this too should be modelled - the prospect of detection may still give pause to such a host.

There are of course multiple ways of modelling this scenario. One advantage of our approach in general is that it is able to compare different such models analytically: as seen in [1] where we compare multiple models of varying levels of abstraction and assess how the abstraction process may lead to different insights and decision support. It should be noted that the models of this case study are created by nation **N2** in order to assess this scenario, but where **N2** is as objective as possible.

Bayesian Belief Networks (BBNs) are probabilistic models, based on conditional probability that encode events of interest in a node-structure. BBNs are dependency graphs, displayed diagrammatically, as a directed acyclic graph (DAG) [6]. Dependencies of events are captured and expressed in probability tables, with observations of events influencing the chances of related events by "updating" the probabilities using 'hard evidence' passed through nodes using algorithms such as the Junction Tree Algorithm (JTA) [6].

Nodes with 'parents' have probability tables that list the probability distributions conditionally on the aforementioned parent events. These have been used to assess the true value of beliefs held after observations of certain events in order to provide an untainted logical viewpoint on a situation.

In our model, nodes will have different types and numbers of outcomes (True/False, Yes/No, Green/Amber/Red, etc.) for which we can compute marginal probabilities. These outcomes then in turn feed down into the possible parent conditions of their child nodes, as seen in Figure 2b for the Host cheating node's output Is cheating and Is not cheating. Nodes that are not joined by a directed edge (indicating parent/child relationship) are considered conditionally independent of each other (for instance Number of tampered IBs and Isotopic ratio of item).

BBNs can be extended to what are formally known as a Credal Network: a family of BBNs following the same structure, but with at least one node or event offering an - often infinite - number of different probability distributions of events [5]. We create these networks by replacing one of the numerically-expressed probabilities with a formal parameter, say, x and asserting that x lies within a range of values, say, x in $[0.01, 0.15]$. As reported in [1], we refer to this as a *constrained* BBN (cBBN). A cBBN extends a conventional Credal Network by also allowing us to assert scenario-specific logical constraints that are not representable by a conventional

BBN, but that can still be symbolically evaluated. Such logic may be of use when running two models of the same scenario in two different BBNs, say at different levels of abstraction, or by two different modellers. In such cases we would wish to ensure that the models were considering the same situation to ensure useful results; for instance that if a Portal Monitor is used in one, it is used in the other too.

(a) A prototype information barrier for use in arms control. This would be built by the host and expected to return either a green or red light to the inspectors based on presence of nuclear material.

	Cheating Method	
	Is cheating	Is not cheating
None	0	1
IB tamper only	1/3	0
Surrogate source only	1/3	0
IB tamper & surrogate source	1/3	0

(b) A sample node table for the node Cheating Method prior to any under-specification. The output options (none, IB tamper only, surrogate source only or both) are conditioned the parent node Host cheating?

Fig. 2: To help explain our BBN model of Figure 1, we include samples of both the objects being modelled, and the format of the probability nodes tables that do the modelling.

Our newly proposed model, designed using the Netica Application [7], of the inspection scenario is seen in Figure 1. It depicts different aspects of the verification procedure in four key areas: it considers the *selection of the IBs* for use and authentication purposes, the *workings of the IB* in the inspection itself, *authentication* of (other) IBs, and these aspects are tied together to ascertain whether we believe there is *any possibility of cheating overall*, be it through tampering with the IB, surrogate nuclear sources, etc. Throughout, we use a sans serif font to denote nodes of our model, e.g. IB Pool size.

The selection of the IB starts with the IB pool size; a selection of IBs built by the host **N1**, from which there will be a Number of IBs picked for authentication and Number of IBs picked for use by the inspecting nation, **N2**. Should a Number of tampered IBs exist, then the selection process (blind to such a tamper), following a Hypergeometric distribution will probabilistically see whether such tampered IBs make it in to use in the verification process, authentication or neither. The choice of distribution reflects that once 'chosen' from the selection process, IBs won't be used for any other purpose.

The IBs picked for use help the inspecting party to judge whether the item under inspection Is a weapon. Whether it is a weapon, or whether there is a Surrogate Pu source determines physical nuclear properties about the Isotopic ratio of Plutonium elements. Here we capture a possible inspector judgement that a surrogate source would only be used if the host felt that it was extremely likely to pass the IBs tests, and therefore that any surrogate source would have isotopic properties as least as good as those of a real weapon; the conditional probabilities chosen for each isotopic ratio given that the test object is or is not a weapon are not derived from real-world weapons data, but instead reflect in broad terms that Plutonium with a higher isotopic ratio than the chosen threshold is less likely to be found in a nuclear weapon. A bespoke algorithm is used by the IB system on the collected gamma spectrum to test whether both the Peaks are in the expected locations and the Peak aspect ratio is as expected. If all 5 peaks are present and the Ratio of 240/239 isotopes is acceptable then the one or both of the First IB result or Second IB result are reported (conditional on any tampering), depending on whether or not two IBs are used to test the same object .

We cannot hope to model each potential tamper. Therefore, we model authentication as an assessment of the Inspector's authentication capabilities; for which the better these are, the more likely the Tamper will be found. Tampers can only be found however, if the selection of IBs for authentication included one tampered IB. This is controlled by the parent nodes: the aforementioned 'Hypergeometric distribution', and a node Chance of picking a tampered IB for authentication.

The model is drawn together by the overarching question of is the Host cheating?. If so, we then determine a Cheating method, which reflects nation **N2**'s understanding of the possible ways that **N1** could try to cheat, as outlined above, and **N2**'s prior beliefs about the relative likelihood of each method if it were cheating (we have set this as a uniform prior in our model, but this could be readily varied to take account of information from other sources or subjected to sensitivity analysis if required). Finally we check to see if a Portal monitor is used to stop the transportation of radioactive material in and out of the facility which could be used as a surrogate source, although we do not model this system in detail.

Probabilities used in the model and available from the online code come from a variety of sources. Some are arbitrarily selected, as described above, and therefore make perfect candidates for underspecification, as this allows exploration of a range of values over which there is uncertainty. Probabilities relating to the performance of the IB system are derived from experimental analysis of the UKNI IB [8, 9].

3 Queries of Interest

To illustrate the advantages of modelling the inspection procedure in a BBN, and underspecifying it to a cBBN, we consider four queries of potential interest to a decision maker, potentially in Nation **N2**.

Question 1. We are unsure about how a host nation is most likely to cheat (Node Cheating method). There are three possibilities modelled: IB tampering only, x; surrogate source only, y; or both IB tampering and use of a surrogate source, $1-x-y$. The probabilities must add to 1, so we use x, y and $1-x-y$ with the constraints that $0 < x < 1, 0 < y < 1, 0 < 1-x-y < 1$. What ranges of x and y do we detect the cheating for (Node Will tamper be found?) with highest probabilities (within 0.02 of the highest probability)?

Question 2. Given we definitely observe no tamper abnormalities on the authenticated IBs (Node Will tamper be found?), what effect does uncertainty over whether the portal monitor functioned correctly (Node Portal monitor used?) have on the likelihood of the IB reporting positively (Node First IB result) that it believes nuclear material is present? Here we model a non-zero probability of system failure, either through tampering or through a fault of some kind, distinct from the nominal experimental error in successfully detecting surrogate source deployment. How big a difference can it make? With what certainty would we believe the portal monitor to be working assuming we were observing a positive IB result 60% of the time?

Question 3. The Node Isotopic ratio of item models the inspectors' assessment of the probabilities of various possible isotopic ratios being present in a weapon and in a surrogate source; here these ratios are presented as a discretised range from 0.06 to 0.13, with extremal values represented as < 0.06 and > 0.13. These probabilities would in reality be based on expert judgement, including some assessment of how likely the inspecting party felt it would be that a nuclear weapon could be made employing a particular isotopic ratio (although the probabilities encoded in our model, available online, are purely nominal and do not draw on any such assessment). Given the subjectivity associated with this judgement it is natural to consider potential means of conducting sensitivity analyses on it. In this case we're interested in the effect on the First IB result of changing these assessed probabilities: specifically, given one or more uncertainty bounds for some or all of the assessed probabilities, how much does this expectation change?

Question 4. A decision on how much to prioritise research into IB authentication capabilities needs to be taken. We under-specify the likelihood of the inspector's capabilities in finding a tamper (Node Inspection authentication capabilities). The number of IBs picked for authentication has also yet to be determined (Node Number of IBs to pick for authentication), and so we seek to test our capabilities in scenarios of this being either 1 or 5 machines picked for testing. In a situation where the host was cheating (Node Host cheating?), how large a difference in Node Will tamper be found? can the different number of devices authenticated make, over different possible capabilities for authentication?

4 Analysis of Queries

This paper embodies two separate streams of work: that of the accurate mathematical modelling of the inspection scenario, and of modelling BBNs in a constraint solver. The main technical workings of modelling a much smaller cBBN in a constraint solver [1] considered the engineering behind under-specifying the inputs to the Junction Tree Algorithm (JTA) that propagated hard evidence through the cBBN, and the assertion of equations describing the resulting marginal probabilities from the nodes in our Satisfiability modulo theories (SMT) [4,10] constraint solver of choice, Z3.

We explore the results from our queries of Section 3 here. It is worth noting that our representation of cBBNs in Z3 uses internal "code names" for mathematical entities of interest, for example, `v_Find_S1` for the marginal Will tamper be found? = Yes.

Question 1

We start by maximising the marginal Will tamper be found? = Yes. We encode $1 - x - y$ as u and constrain $x + y + u = 1$ for simplicity in the assertions: x represents tampering with the IB; y models using a surrogate Pu source, and u represents the possibility of doing both.

The result comes back as `u = 0.0, y = 0.0,  x = 1.0, v_Find_S1 = 0.19714`, which gives us a target 'range' for being within 0.02 of the highest probability of $0.17714 \leqslant$ `v_Find_S1`. Through a system of maximising and minimising variables in turn, we calculate that this target range is achievable for values of x in $0.0 \leqslant x \leqslant 1.0$ and y in $0.0507 \leqslant y \leqslant 0.1015$.

These target range results say that the node of interest is insensitive to changes in the node underspecified for x, which is allowed to vary over its full range without having much impact on the results. Value y is more constrained, meaning that $1 - x - y = u$ is allowed to vary more too: u picks up most of the 'slack' and variation as we move away from the maximum point. We thus remain in the area of highest probability for detecting tampering as long as x or u are large.

Due to the way the model and node are set up, our analysis sensically shows that the method of cheating that gives us the highest chance of finding a tamper is when the cheating method is 'tamper'. However, it also points out that unless both tamper and surrogate source, or tamper on its own are used, there are limited ways in which to detect cheating through these nodes. It therefore suggests that the use of a portal monitor is advisable, as any increase in y moves us out of the region of highest probability of detecting cheating, and decreases the chances of cheating being detected otherwise.

Perhaps the most interesting part of these results is the range of y, which gives potential insight into future work required on tamper detection for the inspecting nation. Despite not contributing to an IB tamper or detection, y can vary by over 0.05 - over two times that of our end probability range of only 0.02. This suggests there are other limiting factors on the tamper detection, such as capability, that could be much improved on.

Question 2

Uncertainty in whether the portal monitor functioned correctly in the manner described above (where we use x to represent this probability) would have a considerable effect on whether a positive or negative result was expected on the IB. The maximum chance we'd have of seeing a positive IB result is computed as 0.789 and occurred when most likely the portal monitor wasn't working (i.e., $x = 1$). The rate of positive IB results we'd expect to see when the portal monitor was definitely working ($x = 0$) coincided with the minimum rate of seeing a positive IB result, computed as 0.561. This is because deployment of a surrogate source to spoof the IB, which is most likely to be successful if any deployed portal monitor isn't working, would increase the chance of a positive IB result.

We can use the model to infer likelihood of the portal monitor working, and use the symbolics of the model to work 'backwards' from the rate at which we observe the positive IB result and any other knowledge that we acquire - in this case, for example, we model that a tamper hasn't been found on the authenticated IBs. If we observed a positive IB result 60% of the time, the appropriate assumption calculated by our tool would be that the portal monitor wasn't working 17.1% of the time. While we should treat the numerical result with some caution, given the uncertainties in the data, the relatively high likelihood that there may be a system failure - either because of the actions of a nefarious host or a fault of some kind - should give rise to further investigation of the portal monitor system.

The implementation of any such investigation or additional check would lower the expected rate positive IB results in this scenario. This process could be incorporated in our tool and analysed in a similar fashion to the IB authentication process.

Question 3

Where a surrogate source is used, we parameterise the lower end of the isotopic ratio range with x, such that high x indicates a greater prior belief that any surrogate source used would have a lower isotopic ratio; similarly, where no surrogate source is used, we parameterise using y such that high y indicates a similar belief about nuclear weapons. To set up a specified tolerance on some of these assessments on x and y independently, we replace the assessed probabilities of 0.2 in four of the states with x and constrain that x such that $0.15 \leqslant x \leqslant 0.25$; we conduct a similar procedure (which is detailed in full in the code available online) for y and constrain y such that $0.1 \leqslant y \leqslant 0.2$. For the minimisation, the witness returned was `x = 0.15, y = 0.1, v_Result1_S1 = 0.65824`; and for maximisation `x = 0.25, y = 0.2, v_Result1_S1 = 0.710302`. For comparison, we also record the result with no error, `x = 0.2, y = 0.15, v_Result1_S1 = 0.68427`.

We see that x plays only a small part in the result; indeed when $x = 0.25, y = 0.1$, then `v_Result1_S1 = 0.66254`, which is very close to the minimum value. This reflects that any sensible use of a surrogate source would seek to ensure a positive IB result (and therefore avoid any potential further investigation that might result in detection of the source). Varying y has a greater effect, because - although a sensible host will wish to negotiate the threshold such that their weapons pass this test with a very low false negative rate - we cannot be so certain and therefore we admit a greater probability of the object in question having an isotopic ratio closer to the threshold, where the IB system performs least well.

Question 4

We represent the two BBN models as a combined set of constraints in Z3. Capturing the situation of there being 1 IB picked for authentication as model 1 with variable `v_Find_S1`, and 5 IBs picked as model 2 with variable `v_Find_S1_2nd`, we underspecify the likelihood of an inspector having high capability to find a tamper as x, medium capability as fixed 0.333 and low as $0.667 - x$ in model 1, and using y in an identical way in model 2. We define a new variable, `DIFF` as the value `v_Find_S1 - v_Find_S1_2nd`. Iterating over all possible x and y, $0 \leqslant x \leqslant 0.667$ and $0 \leqslant y \leqslant 0.667$, we can see how `DIFF` changes in Figure 3.

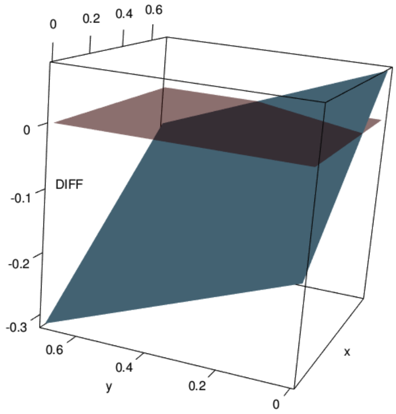

Fig. 3: The plot (in solid blue) of the differences between the two models' marginal probabilities for a tamper being found, as x and y vary in the models (with values taken at a 0.01 granularity). Plane $\texttt{DIFF} = 0$ (translucent red) cuts the blue plane, and marks the boundary of where the changes in capabilities for the 1 and 5 IB authentication test models favour running 1 IB test (above the red plane) and 5 IB tests (below the plane). The equation of the boundary line is approximately $-0.213y + 0.355x = 0.1652$.

The result is a flat surface, reflecting the linear relation of the under-specification in the capability node to its child node of 'will the tamper be found'. The mostly negative surface shows that the case of testing 5 information barriers for tampers is nearly always better, irrespective of the confidence we have in our ability to find a tamper. This is true, other than for the most extreme cases when we have least confidence in our authentication capabilities when testing five barriers ($y = 0$) and most confidence when we are testing one ($x = 0.667$) - which makes sense.

We investigate this extreme by considering a situation where the inspector believes their authentication capabilities are high, such that $x = y = 0.567 \pm 0.1$, and investigating the maximum positive and negative differences between the two models.

Maximising the `DIFF` variable (over this more limited range), the model returned:

```
x = 0.667,  y = 0.467, DIFF = -0.02817, v_Find_S1_2nd = 0.17715, v_Find_S1 = 0.14898
```

Whilst when minimising, the witness was:

```
x = 0.467,  y = 0.667, DIFF = -0.14163, v_Find_S1_2nd = 0.24804, v_Find_S1 = 0.10641
```

The largest absolute difference between the two models in this case is just over 0.14. This is the largest difference we could expect under an error tolerance of 0.1. The largest absolute value happens to be the smallest numerical value (minimum). This reflects a situation where the inspector chooses five IBs and has particularly high authentication capabilities ($y = 0.667$, `v_Find_S1_2nd` $= 0.24932$). This results in a considerably greater payoff than only inspecting one IB with more moderate capability ($x = 0.467$, `v_Find_S1` $= 0.10696$). As a direct comparison of the authentication capabilities (i.e., with $x = y$ fixed), the difference between models is felt most when authentication capability is lowest. We confirm this using the surface of Figure 3 and for the cases $x = y$ and $x = \{0.0, 0.667\}$ where `DIFF` $= \{-0.165156, -0.07075\}$ respectively.

A decision maker could use the information here to make an informed decision about how many IBs to authenticate, weighing the cost of IB production against the cost of developing and employing more advanced authentication capabilities to conduct a cost-benefit analysis. They could also query in detail how the results of these cost-benefit analyses might change as new information is learned or new techniques deployed. This capability might help decision-makers to balance their priorities and to gain the best assurance possible without excessive cost that the verification regime they implement is effective.

5 Discussion & Conclusions

Overall, we have produced a comprehensive mathematical model of a nuclear arms inspection scenario. This carefully captures different facets of the verification process in order to present an unbiased assessment for use in helping decision making. This methodology could either be applied in the planning stage of a treaty, or be used to reduce or trace biases in reports based on observations in the field.

Using our under-specification methodology, we have considered queries of the form *"How large a difference can the number of devices being authenticated make in a tamper being found, given uncertainty over a nation's authentication capabilities?"*. We used our methodology to exploit new and interesting queries such as these through the under-specification of variables and automated analysis. Using this approach we hope to provide more information to a decision maker so that more comprehensive and fully informed decisions can be made in the field in the face of many unknown quantities.

Acknowledgements: The authors from Imperial College London would like to thank AWE for sponsoring a PhD studentship under which the research reported in this paper was carried out.

Open Access of Research Data and Code: The Python and SMT code for the queries and models of this paper and raw SMT analysis results are reported in the public data repository `https://bitbucket.org/pjbeaumont/inmm2016/`

References

1. Beaumont, P., Evans, N., Huth, M., Plant, T.: Confidence analysis for nuclear arms control: SMT abstractions of Bayesian Belief Networks, Computer Security – ESORICS 2015, Lecture Notes in Computer Science, Springer, 2015
2. Beaumont, P., Evans, N., Huth, M., Plant, T.: Confidence analysis for nuclear arms control: SMT abstractions of Game Theoretic Models, INMM57, 24-28 July, Atlanta, USA, 2016
3. Beaumont, P., Evans, N., Huth, M., Plant, T.: Bounded Analysis of Constrained Dynamical Systems: A Case Study in Nuclear Arms Control, INMM57, 24-28 July, Atlanta, USA, 2016
4. Barrett, C., de Moura, L.M., Ranise, S., Stump, A., Tinelli, C.: The SMT-LIB initiative and the rise of SMT, Hardware and Software: Verification and Testing - 6th International Haifa Verification Conference, HVC, 2010
5. Cozman, F.G.: Credal networks. Artif. Intell. 120(2), 199–233 (2000)
6. Fenton, N., Neil, M.: Risk Assessment and Decision Analysis with Bayesian Networks. CRC Press (2013)
7. NorSys Netica: https://www.norsys.com/netica.html
8. UK MoD: The UK/Norway initiative: report on the UKNI nuclear weapons states workshop (March 2010)
9. UKNI: http://ukni.info/
10. de Moura, L.M., Bjørner, N.: Z3: An efficient SMT solver. In: TACAS. pp. 337–340 (2008)
11. De Moura, L., Bjørner, N.: Satisfiability Modulo Theories: Introduction and Applications, Communications of the ACM, Vol. 54 No. 9, Pages 69-77, 2011