Science & Global Security, 23:20–47, 2015

ISSN: 0892-9882 print / 1547-7800 online
DOI: 10.1080/08929882.2015.996076

Routledge
Taylor & Francis Group

Antineutrino Reactor Safeguards: A Case Study of the DPRK 1994 Nuclear Crisis

Eric Christensen,[1,2] Patrick Huber,[1] and Patrick Jaffke[1]

[1]Center for Neutrino Physics, Virginia Tech, Blacksburg, VA, USA
[2]Kennesaw State University, Kennesaw, GA, USA

In this article a case study of the application of antineutrino safeguards to a real-world scenario, the North Korean nuclear crisis in 1994, is presented. Detection limits to a partial or full core discharge in 1989 based on actual IAEA safeguards access are derived and it is found that two independent methods would have yielded positive evidence for a second core with very high confidence. To generalize these results, detailed estimates for the sensitivity to the plutonium content of various types of reactors, including most types of plutonium production reactors, are presented, based on detailed reactor simulations. A key finding of this study is that a wide class of reactors with a thermal power of 0.1–1 gigawatt can be safeguarded achieving IAEA goals for quantitative sensitivity and timeliness with antineutrino detectors adjacent to the reactor building. Antineutrino reactor monitoring does not rely on the continuity of knowledge and provides information about core inventory and power status in a timely fashion. The necessary detection systems do not exist yet but are expected to become available within two to five years.

INTRODUCTION

Neutrinos were postulated by Wolfgang Pauli in 1930 and have been experimentally discovered by Clyde Cowan and Fred Reines in 1956.[1] The Cowan-Reines experiment successfully detected antineutrinos emitted from the Savannah River reactor. A nuclear reactor is, to be precise, a source of electron antineutrinos. Antineutrinos are nearly massless, electrically neutral, spin half particles, and play a central role in the electroweak Standard Model of particle physics. Antineutrinos participate only in weak interactions and therefore possess unusual penetrating power. No practical means to attenuate

Received 22 January 2014; accepted 22 October 2014.
Address correspondence to Patrick Huber, Center for Neutrino Physics, Department of Physics, Virginia Tech, Blacksburg, VA 24061, USA. E-mail: pahuber@vt.edu
Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/gsgs.

or to shield antineutrinos are known. Antineutrinos are copiously produced in the beta-decays of fission fragments, making nuclear reactors the most powerful artificial antineutrino source.

The basic concept of monitoring nuclear reactors using antineutrinos was proposed by Borovoi and Mikaelyan in 1978.[2] More recently, there has been an effort to assess the possibility to accurately determine the plutonium content in a reactor core using antineutrino measurements.[3] If confirmed, such ability could open the development of antineutrino detection based reactor safeguards. However, the very different assumptions made about antineutrino detector capabilities and the various levels of statistical analysis in recent studies have produced a wide range of results leading to mixed perceptions of the feasibility of antineutrino safeguards.

Antineutrino based safeguards are often evaluated in standard scenarios where more conventional safeguards already perform well, limiting *de facto* interest in their development. On the other hand, inventing scenarios in which the standard methods fail brings about the criticism that these scenarios are artificial, unrealistic, and contrived for the sole purpose of demonstrating the usefulness of antineutrinos. There exist some historical cases where conventional safeguards did not perform or could not be implemented as expected. The first North Korean nuclear crisis of 1994 is a famous example.[4] The crisis revolved around whether or not the Democratic People's Republic of Korea (DPRK or North Korea) had produced and separated much larger quantities of plutonium in its Yongbyon nuclear reactor than it was claiming. However, the DPRK limited the amount of information available to the International Atomic Energy Agency (IAEA) to offer an independent assessment. The subsequent North Korean nuclear tests confirmed the existence of a military nuclear program.

After presenting the technical aspects of antineutrino safeguards, this article first shows that a wide class of reactors with a thermal power of 0.1–1 gigawatt (GW_{th}) can be safeguarded achieving IAEA goals for quantitative sensitivity and timeliness with antineutrino detectors outside the reactor building. It then presents what would have been improvements in the quantitative understanding of the DPRK's plutonium production and in particular the assessment of the veracity of its initial declaration to the IAEA if antineutrino safeguards had been in place.

The study concludes that under the historical constraints and boundary conditions of the DPRK 1994 nuclear crisis, the special capabilities of antineutrino safeguards could have provided a decisive advantage over conventional techniques by confirming the diversion of important quantities of plutonium.

ANTINEUTRINO REACTOR MONITORING

Neutrinos are not directly produced in nuclear fission but result from the subsequent beta-decays of the neutron-rich fission fragments. The total number

of emitted antineutrinos is proportional to the total number of fissions in the reactor. On average there are about six antineutrinos per fission emitted and each fission results in an energy release of about 200 mega electron volt (MeV). Thus, for one gigawatt of thermal power a flux of about 10^{20} antineutrinos per second (s^{-1}) is produced. Moreover, the distribution of fission fragments, and hence their beta-decays, are different for different fissile isotopes. Thus, careful antineutrino spectroscopy should provide information not only on the total number of fissions but also about the fission fractions of the various fissile isotopes contained in the core. The basic concepts of both power monitoring and observing the plutonium content of a reactor were experimentally demonstrated in pioneering work performed by a group from the Kurchatov Institute lead by Mikaelyan.[5] They deployed an antineutrino detector of about 1 m^3 volume at the Rovno nuclear power plant in Ukraine. For the power measurement, an agreement with the thermal measurements was found to within 2.5 percent,[6] and the effect due to a changing plutonium content was demonstrated,[7] and more recently the quantitative accuracy has been studied as well.[8] This allows one to determine the plutonium content and power level of the reactor core *in situ* at a standoff distance of 10's of meters.[9] The practical feasibility of reactor monitoring using antineutrinos has also been demonstrated using a small, tonne-size detector at the San Onofre power station, called SONGS,[10] in Pendleton, California.

ANTINEUTRINO DETECTION

Beginning with the discovery of the antineutrino, inverse beta-decay (IBD) has been the workhorse of reactor antineutrino experiments

$$\bar{\nu}_e + p \rightarrow n + e^+ \tag{1}$$

In IBD, an electron antineutrino interacts with a proton to produce a neutron and a positron, this process has an energy threshold of $(m_n - m_p + m_e)c^2 = 1.8$ MeV. The positron will go on to annihilate with a nearby electron producing a pair of 511 keV gamma rays. This energy deposition is typically detected together with the kinetic energy of the positron E_e and thus the visible energy in detector, $E_{\text{vis}} = E_e + 2 \times 511\,\text{keV}$. There is a one-to-one correspondence between antineutrino energy and the positron energy $E_\nu = E_e + 1.8\,\text{MeV}$. Therefore, a measurement of E_{vis} directly translates into a measurement of the antineutrino energy E_ν.

The reaction in Equation (1) also results in a neutron, which will slow down in collisions with the detector material and eventually undergo neutron capture. A careful choice of the nucleus on which the neutron captures allows tailoring this signature. Common neutron capture agents are gadolinium, e.g., Daya Bay experiment,[11] or lithium, e.g., Bugey experiment.[12] For instance, in the case of gadolinium, the signature of neutron capture is the emission of

several gamma rays with a total energy of 8 MeV. The slowdown and capture of the neutron requires a characteristic time, allowing for what is called a delayed coincidence: there is a primary energy deposition from the positron followed somewhat later by a neutron capture signal. This delayed coincidence is key to separating antineutrino events from backgrounds. The neutron capture cross sections of gadolinium is much larger than of any of the other detector materials. Therefore, even small concentrations will result in the majority of neutron captures occurring on those nuclei; similar arguments apply to lithium.

All signatures will result in ionization, which is detected by using organic scintillator, either in liquid or solid form. The organic nature of the scintillator provides the free protons for the interaction in Equation (1). Recently, there have been three experiments aimed at fundamental physics employing gadolinium-doped liquid scintillator at a large scale of several 10 tonnes (t) without any safety incidents and excellent long-term stability.[13] Specifically, this article considers a 5 t detector based on organic scintillator corresponding to 4.3×10^{29} target protons. A real detector will not have 100 percent efficiency, and to obtain the same number of events a larger detector will be needed.[14] Many antineutrino detectors with efficiencies above 50 percent have been built and thus even a realistic detector yielding the same event numbers would have a mass of less than 10 t. A standard 20 foot intermodal shipping container has a net load capacity of 28.2 t, thus even a 10 t antineutrino detector fits easily within such a container together with its support systems. The antineutrino spectrum is divided in energy from 1.8 MeV to 8 MeV in bins of 0.2 MeV width, which at 4 MeV approximately corresponds to $10\,\text{percent}/\sqrt{E}$ resolution, which is similar to the resolution of recent experiments.[15] A resolution half as good would yield virtually identical results, as shown by tests. For the IBD cross section, the result of Vogel and Beacom[16] are used and corrected for a neutron lifetime of 878.5 s.[17] For all measurements at reactors, the standoff is 20 m, which allows for deployment outside the reactor building. Such a detector at this standoff would typically register about 5,000 events per year for a reactor operating at 1 MW_{th} throughout that year.

Currently, no detector systems with all the required characteristics exist and in particular, operation at the surface with sufficiently high signal efficiency needs to be demonstrated (see also the section on the impact of backgrounds). A major step towards this goal was recently achieved by a Japanese group, which managed to detect reactor antineutrinos from the back of a van.[18] Given the significant international effort towards short-range reactor antineutrino detection to study the reactor antineutrino anomaly (see the following section on reactor flux models), it is likely that a first prototype having all required properties will be available within 12–18 months. These prototypes will

be in the million dollar range, but second generation systems have significant potential for improved cost effectiveness.

REACTOR FLUX MODELS

More than 99 percent of the power in reactors, in a uranium fuel cycle, is produced in the fission of four isotopes: uranium-235, plutonium-239, uranium-238, and plutonium-241. A reactor with fresh fuel starts with only fissions in the uranium isotopes and plutonium is produced via neutron capture on uranium-238 as the burn-up increases. The total neutrino flux from a reactor, ϕ, in units of antineutrinos per second and per MeV can be written as

$$\phi(E) = \sum_I f_I S_I(E)\,, \tag{2}$$

where f_I is the fission rate in isotope I in units of fissions per second and $S_I(E)$ is the antineutrino yield for the isotope I, in units of antineutrinos per fission and MeV. The thermal power of the reactor is also given in terms of the fission rates

$$P_{\text{th}} = \sum_I f_I p_I\,, \tag{3}$$

where p_I is the thermal energy release in one fission of the isotope I; the values for p_I are taken from Kopeikin.[19] To disentangle the contributions of the four isotopes, the antineutrino yields S_I needs to be calculated. The neutrino yields are given by the antineutrino spectra $\nu_k(E)$ of each fission fragment k and the cumulative fission yield for each fragment, Y_k^I,

$$S_I(E) = \sum_k Y_k^I \nu_k(E)\,, \tag{4}$$

where k typically runs over about 800 isotopes. In practice, the antineutrino spectrum of a given fission fragment is not known; only information, often insufficient, regarding the beta spectrum is available. Due to the underlying nuclear physics, a direct computation of the neutrino yields S_I via the summation of all individual antineutrino spectra will be inaccurate.[20]

A more accurate method is based on the measurement of the integral beta spectrum of all fission fragments[21] and subsequently the antineutrino spectrum can be reconstructed from those measurements.[22] This method is less dependent on nuclear data, but is not free from uncertainties related to effects of nuclear structure.[23]

Comparing the event rate predictions for various flux models one finds that a simple model based on direct summation of the antineutrino spectra and more sophisticated calculations like the one by Fallot et al.,[24] as well the

inversion method, yield very similar results in terms of rates and mean energies when normalized to uranium-235.[25] In other words, the difference in antineutrino yield and mean energy between the fissile isotopes is consistently predicted by the various flux models.

In practice, the current errors of any flux model are significant and a set of calibration measurements at reactors of known fissile content is required to mitigate the effect of these uncertainties, particularly in view of the reactor antineutrino anomaly.[26] A proof of concept at a theoretical level for these calibrations has been performed.[27] On the experimental side, the Daya Bay collaboration has demonstrated the ability to cross-calibrate a set of eight antineutrino detectors to within better than 0.5 percent.[28]

BURN-UP CALCULATIONS

The connection between fission rates and mass inventory requires a more detailed look at the reactor physics. For a neutron flux which is constant in time and space, the fission rate, f_I, and mass of a given fissile isotope, m_I, have a simple linear relationship

$$f_I = \phi_n \, \sigma_I \, m_I \,, \tag{5}$$

where σ_I is the energy averaged fission cross section and ϕ_n is the neutron flux. Throughout the evolution of the core, all factors on the right hand side of Equation (5) will change. Due to burn-up effects, the mass m_I will change and the neutron flux typically will be adjusted to compensate for changes in reactivity while maintaining constant power. The accumulation of fission fragments will change the neutron absorption, which, in turn, alters the neutron energy spectrum; the cross section σ_I will evolve as well. We have performed evolution or burn-up calculations for several reactor types using the SCALE software suite.[29] For the further discussion it is useful to introduce fission fractions F_I, which are defined by

$$F_I = \frac{f_I}{\sum_I f_I} \quad \text{with} \quad \sum_I F_I = 1 \,. \tag{6}$$

This definition has the advantage that the problem can be phrased independently of reactor power. For illustration, the time evolution of the F_I for a graphite moderated, natural uranium fueled reactor is given in the left hand panel of Figure 1, where the fission fractions are shown as a function of the burn-up. F_{Pu241} is very close to zero in this type of reactor and therefore is not visible in this figure. The fission rate in uranium-238 stays constant since the amount of uranium-238 in the reactor changes very little with time. There is a clear anti-correlation between the fission fractions in uranium-235 and plutonium-239. The anti-correlation is nearly exact as shown in the right panel

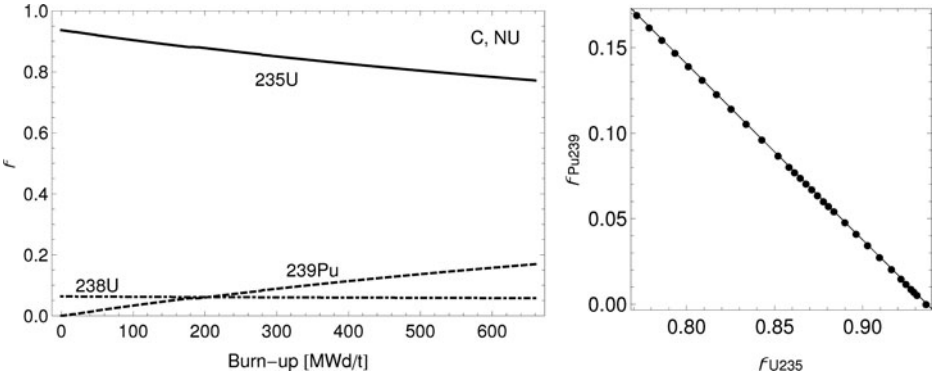

Figure 1: The left panel shows the evolution of the fission fractions in a graphite moderated natural uranium fueled reactor as a function of burn-up. The right panel shows the anti-correlation of the fission fractions in uranium-235 and plutonium-239.

of Figure 1. Burn-up measures the number of fissions which have occurred per unit of fuel mass or, in other terms, the amount of energy extracted; the unit for burn-up is MWd/t. Neglecting radioactive decays, the isotopic composition of samples with the same burn-up would be identical since the total number of fissions that took place is the same. The details of the power history and shut downs, have only a minor impact on the reactivity and fission fractions. For the purpose of simulation, a reasonably accurate model of the reactor power history is required, which in turn serves as input for a detailed reactor physics calculation. The ability to predict the antineutrino emission over time relies only on a model of the burn-up as a function of time. And conversely, the precise measurement of neutrino emission will allow to infer the burn-up as a function of time.

These burn-up calculations confirm that the relation between fission fractions and mass inventory as given in Equation (5) remains linear and, that the value of $\phi_n \sigma_{\mathrm{Pu239}}$ stays nearly constant.

PLUTONIUM CONTENT DETERMINATION

The difference in the antineutrino yield can be used to disentangle the contribution of each of the fissile isotopes. The event rate in each bin n_i is given as

$$n_i = N \sum_I f_I \int_{E_i-\Delta E/2}^{E_i+\Delta E/2} dE\, \sigma(E)\, S_I(E)\,, \tag{7}$$

where E_i is central energy of bin i, ΔE is the bin width and $\sigma(E)$ is the IBD cross section. N is an overall normalization constant set by the number of free protons and detection efficiency. To compute the event rates, n_i, the four fission

rates need to be specified $\mathbf{f} = (f_{\mathrm{U235}}, f_{\mathrm{U238}}, f_{\mathrm{Pu239}}, f_{\mathrm{Pu241}})$. The true or input values for the calculation are denoted by a superscript 0, i.e., the true fission rates are $\mathbf{f}^0$, and similarly the n_i computed for the true values $\mathbf{f}^0$ as n_i^0. The χ^2-function is define as

$$\chi^2(\mathbf{f}) := \sum_i \frac{\left(n_i(\mathbf{f}) - n_i^0\right)^2}{n_i^0}\,, \tag{8}$$

This χ^2-function will be zero for $\mathbf{f} = \mathbf{f}^0$. The allowed region for $\mathbf{f}$ is obtained by requiring that

$$\chi^2(\mathbf{f}) \leq \chi_c^2\,, \tag{9}$$

where the critical value χ_c^2 is determined from a χ^2 probability distribution with, in this case, four degrees of freedom. If only the total number of fissions in plutonium given by $f_{\mathrm{Pu}} = f_{\mathrm{Pu239}} + f_{\mathrm{Pu241}}$ is of interest, the following marginalized function has to be used

$$\bar{\chi}^2(f_{\mathrm{Pu}}) = \min_{f_{\mathrm{U235}}, f_{\mathrm{U238}}, \kappa} \chi^2(f_{\mathrm{U235}}, f_{\mathrm{U238}}, (1-\kappa)\, f_{\mathrm{Pu}}, \kappa f_{\mathrm{Pu}})\,, \tag{10}$$

and with a single parameter, f_{Pu}, the number of degrees of freedom reduces to one. Similarly, a corresponding single parameter function can be defined for the measurement of reactor power.

f_{Pu} will be proportional to the plutonium mass, m_{Pu}, in the reactor

$$\gamma = \frac{m_{\mathrm{Pu}}}{f_{\mathrm{Pu}}}\,, \tag{11}$$

where γ is the proportionality constant. Therefore, a measurement of f_{Pu} translates into a determination of m_{Pu}. γ, in turn, depends on the details of the reactor physics as well as the instantaneous reactor thermal power; note that according to Equation (5), $\gamma = 1/(\phi_n \sigma_{\mathrm{Pu}})$ and thus is inverse to the neutron flux density ϕ_n. The determination of m_{Pu} and its connection to γ is clearly illustrated in Figure 2, where the accuracy in the determination of m_{Pu} is shown for a variety of reactor types as a function of the thermal power. This figure is based on a full calculation of the reactor burn-up, where "C, NU" corresponds to a graphite moderated reactor running on natural uranium and the dot on this line is the five megawatt electrical power ($\mathrm{MW_e}$) reactor in the DPRK. "H_2O, HEU" and "H_2O, HEU + NU" correspond to the IRT, another reactor in the DPRK, with drivers only and to the IRT with drivers and targets, respectively. The case "H_2O, LEU" is computed for a typical pressurized light water reactor. A power history was taken from one such reactor, with a total fuel load of 72.4 MTU enriched to 3.7 percent.[30] The case "D_2O, NU" describes a heavy water moderated reactor running on natural uranium modeled on a CANDU design with a 8.6 MTU natural uranium fuel load and running at 40 $\mathrm{MW_{th}}$. The 40 $\mathrm{MW_{th}}$ point on this line resembles the Iranian reactor at Arak,[31] and

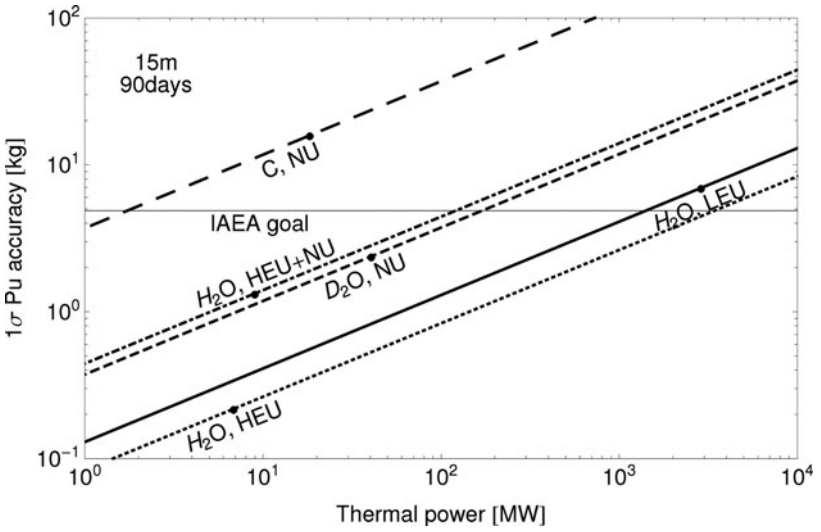

Figure 2: Absolute accuracy in the determination of the plutonium content based on the measurement of the antineutrino spectrum as a function of the design thermal power of the reactor. The different lines stand for different types of reactors as indicated by the labels: the first term indicates the type of moderator, whereas the second part denotes the fuel type, natural uranium (NU), low enriched uranium (LEU), and highly enriched uranium (HEU). This figure assumes a 5 t detector, a standoff of 15 meter (m) from the reactor core, and 90 days of data collection. The horizontal line labeled "IAEA goal" indicates the accuracy which corresponds to the detection of 8 kg of plutonium at 90 percent confidence level, which is the same as 5 kg at 1 σ.

the accuracy would be at the level of 2.7 kg within 90 days. The horizontal line corresponds to a sensitivity to 8 kg plutonium at 90 percent confidence level (or 5 kg at 68 percent confidence level) within 90 days.

For each different reactor type the accuracy of a m_{Pu} measurement can be described by the following simple relation

$$\delta m_{\mathrm{Pu}} = 1.9\,\mathrm{kg}\left(\frac{\gamma}{10^{-16}\mathrm{kg\,s}}\right)\left(\frac{L}{\mathrm{m}}\right)\left(\frac{P_{th}}{\mathrm{MW}}\right)^{1/2}\left(\frac{\mathrm{tonnes}}{M}\right)^{1/2}\left(\frac{\mathrm{days}}{t}\right)^{1/2}, \quad (12)$$

where L is the standoff of the antineutrino detector, P_{th} is the average thermal reactor power, M is the detector mass in tonnes (assuming 8.65×10^{28} protons per tonne), and t is the length of the data collection period. Table 1 lists the corresponding values of γ, and using those values, Equation (12) reproduces the results of the full calculation within a few percent.

Table 1: The values of γ for a number of reactor types.

reactor type	C, NU	H_2O, HEU	H_2O, HEU+NU	D_2O, NU	H_2O, LEU
γ (10^{-16} kg s)	2.9	0.064	0.34	0.30	0.11

For most reactor operating conditions, the variation in γ is very small and depends only very weakly, at the level of a few percent, on burn-up and reactor history.

The IAEA goal in safeguarding plutonium in irradiated fuel is the detection of the diversion of one significant quantity, corresponding to 8 kg of plutonium, within 90 days at 90 percent confidence level according to the IAEA safeguards glossary.[32] It is important to note, that "detection of diversion" not only requires a measurement of how much plutonium is present but also an understanding of how much plutonium is expected to be present. This expectation, for instance, can result from a prior inspection or from a calculation based on declared or measured parameters. Here, the assumption is that the expectation of how much plutonium should be present has been derived and the error in this expectation is small compared to the errors of the antineutrino measurement. Under this assumption, the accuracy of the antineutrino measurement and the detection limit for a diversion become the same and therefore, these terms will henceforth be used synonymously.

For reactors with a thermal power in excess of 1 GW_{th}, which represents the majority of all reactors globally used to produce electricity, this approach to safeguards will have difficulties meeting the IAEA goals. On the other hand, antineutrino safeguards are quite straightforward for research, small modular reactors, and plutonium production reactors.

The fission fractions and thus the fission rates are *not* independent from each other but are coupled by the physics inside the reactor (see the right panel of Figure 1). This allows a rephrasing of the fitting problem in terms of one independent quantity, the burn-up. The result of the analysis will be a value for burn-up and some error bounds and since the reactor model also provides all the mass inventories as a function of burn-up, a measurement of the burn-up translates into a measurement of the core inventory and the errors can be determined by standard error propagation. In the case of the graphite moderated reactor this reduces the error in plutonium mass determination by roughly 50 percent, see the section on the Yongbyon 5 MW_e reactor. In cases where there is reliable design information and the key operating parameters are known, the burn-up model will reproduce the core inventory to within the 5 to 10 percent range, yielding a small extra contribution to the overall error. If the reactor design and operating parameters have to be considered as unknown, a fit to fission fractions and power should be performed.

SUMMARY OF THE 1994 NORTH KOREAN CRISIS

On 26 February 1993, the IAEA called for special inspections to resolve the discrepancies indentified during the first safeguards inspections conducted in 1992. The issue of contention was the amount of plutonium the DPRK had separated from spent nuclear fuel. North Korea declared that it produced about

90 g,[33] but IAEA data allowed for the possibility of a much larger amount, perhaps as much as 14 kg.[34] On 12 March 1993, the DPRK declared its intention to leave the Treaty on the Non-Proliferation of Nuclear Weapons (NPT) by 12 June 1993 after being threatened with special inspections but was persuaded by the United States on 11 June not to do so.

North Korea started its 5 MW_e reactor at Yongbyon in 1986. In 1989 there was a 70-day shutdown, providing an opportunity to unload 50 to 100 percent of the spent fuel in the core. In its initial declaration to IAEA in 1992, North Korea indicated that it ran a one-time reprocessing campaign in 1990 that resulted in 90 g of plutonium from a limited number of damaged fuel rods removed during the 1989 shutdown. The results of IAEA environmental sampling conducted during the first safeguards inspection, however, indicated at least three campaigns of reprocessing in 1989, 1990, and 1991,[35] which in turn admits the hypothesis that a significant fraction of the spent fuel had been removed in 1989. As a result, a larger amount of separated plutonium may have been obtained by the DPRK, possibly sufficiently large to build two or more nuclear bombs. In particular, finding and sampling the reprocessing waste streams was a priority for IAEA, eventually triggering the request for special inspections.[36] The diplomatic exchange between IAEA and the DPRK dragged on in parallel with negotiations between the DPRK and the United States; the latter eventually leading to the *Agreed Framework*. In April 1994 North Korea forced the issue by beginning to unload spent fuel from the reactor core. An analysis of the gamma-radiation of spent fuel taken at known positions in the reactor core would have resolved the question of how much spent fuel was discharged in 1989. However, the unloading proceeded very fast and any information about the original position of each fuel element in the reactor was lost. As a result, crucial evidence was denied to the IAEA and on 2 June 1992 Hans Blix, director of the agency at the time, declared that the ability to resolve the issue had been "seriously eroded."[37] The fuel discharged in 1994 was put into storage and was under IAEA surveillance until 2003. The 1994 crisis was resolved by the so called *Agreed Framework* under which the DPRK halted any plutonium production and fuel reprocessing in exchange for the promise to obtain two pressurized light-water reactors at no cost.[38] The *Agreed Framework* unraveled in 2003 when the DPRK declared that it was leaving the NPT and eventually, in 2006, North Korea conducted its first nuclear test.

A comprehensive account of the history of the North Korean nuclear program is provided by Hecker.[39] There are two reactors, the IRT, a Soviet-supplied research reactor with a power of around 8 MW_{th}, and the 5 MW_e reactor, a graphite moderated reactor with a thermal power of approximately 20 MW_{th}, which generally is referred to by its electrical power, hence its name. There also is the Radiochemical Laboratory, which is a reprocessing facility which allows for the extraction of plutonium from the spent fuel from

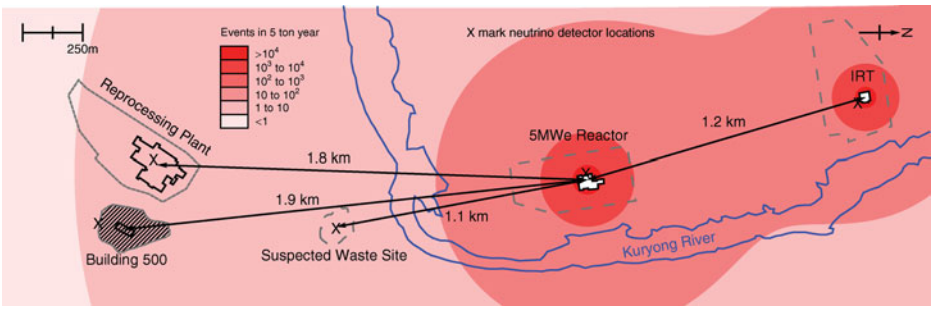

Figure 3: A map of relevant boundaries and geographies of the Yongbyon nuclear facility. Contours show expected inverse beta-decay event rates for a 5 t detector over the course of a year. X's mark the location of various antineutrino detectors used in this article. The satellite image on which this map is based was taken on 16 May 2013 by GeoEye-1. © Digital Globe. Reproduced by permission of Digital Globe.

the 5 MW_e reactor. These facilities and their relative locations are shown in Figure 3, as well as IBD event rate iso-contours.

In the 1960s, the IRT was supplied by the Soviet Union.[40] This reactor is a light-water moderated reactor running on HEU, with enrichment from 10 to 80 percent.[41] The Soviet Union also provided the HEU fuel until its own demise in the 1990s. The nominal power of this reactor is 8 MW_{th}.

North Korea started serious fuel cycle activities in the 1980s and the plan was to build and operate three gas-cooled, graphite moderated, natural uranium fueled reactors. The design followed the British Magnox design as this reactor type was well adapted to North Korean indigenous industrial capabilities. At the same time, Magnox reactors were originally designed as dual-use facilities to produce both electricity *and* weapons-grade plutonium.

The amount of plutonium produced in a reactor can be estimated if the integrated neutron flux, which is proportional to the total energy produced, is known, or equivalently if a *complete* history of the reactor power is available. To obtain the produced plutonium in usable form, the reactor has to be shut down,[42] the irradiated fuel rods removed, and the plutonium chemically separated from the spent fuel at the Radiochemical Laboratory. The location of the various facilities can be seen in Figure 3.

The time evolution of the burn-up for the 5 MW_e is shown in Figure 4 which has been adapted from *Nuclear Puzzle*.[43] The solid curve is based on the declarations made by the DPRK and, thus, the assumption is that no major refueling took place in 1989. The dashed curve is derived assuming that the full core has been replaced with fresh fuel in 1989 under the constraint of arriving at the same final burn-up. These numbers can be readily converted into reactor thermal power levels using the fact that there are approximately 50 tonnes of uranium in this reactor.[44] The power levels then form the input for a detailed calculation of the reactor isotopic composition and fission rates.

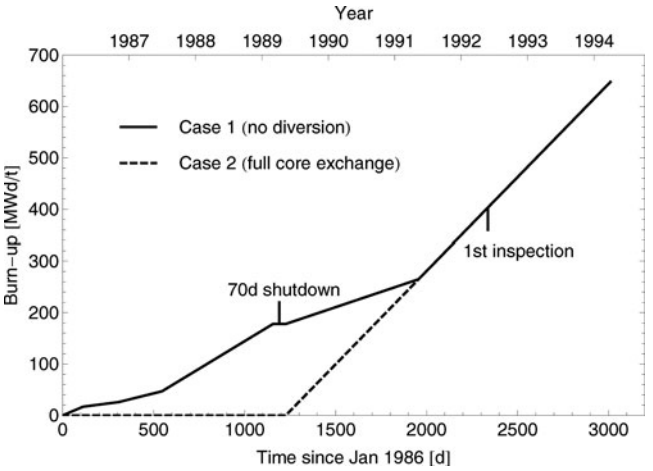

Figure 4: Burn-up of the fuel in the 5 MW_e reactor as function of time measured in days since 1 January 1986. The solid curve is based on the values declared by the DPRK, i.e., no major refueling had taken place in 1989. The dashed curve is derived assuming that the full core has been replaced with fresh fuel in 1989. From Albright and O'Neill, adapted with permission.

ANTINEUTRINOS SAFEGUARDS IN THE DPRK 1994 CRISIS

The central question for the international community, after the initial discrepancies appeared in 1992, was how much plutonium the DPRK had separated. The lower bound on this quantity is represented by assuming that the DPRK's initial declaration to IAEA was quantitatively correct, i.e., only 90 g of plutonium were separated. The upper bound on the amount of separated plutonium is obtained by assuming that the full core with a burn-up of approximately 200 MWd/t was discharged in 1989, containing 8.8 kg of plutonium and that the full core was subsequently reprocessed. The North Korean scientists could have produced additional plutonium over a long period of time by irradiating natural uranium targets in the IRT. Limitations in fuel availability combined with effective IAEA safeguards imply that not more than 1 kg of plutonium could have been produced in the IRT[45]; as it turns out, bounds from antineutrinos would be similar.

In the 5 MW_e reactor, at the time of the first IAEA inspection in 1992, the burn-up and reactor power were the same for both the extreme cases (see Figure 4). Therefore, this analysis will include the hypothetical scenario where antineutrino safeguards were applied before and after the 1989 shutdown.[46] The specific unique capability represented by antineutrino safeguards in this case derives from the ability to measure the power history and burn-up *independently*, and any mismatch indicates a fuel diversion.

In the PUREX process for reprocessing, the fission fragments remain in the aqueous phase and therefore will end up in the waste. Some of these fission fragments produce antineutrinos above IBD threshold even after a considerable time interval has elapsed. They are referred as long-lived isotopes (LLI), in particular: strontium-90 with a half-life of 28.9 years (y), ruthenium-106 with a half-life of 372 days (d), and cerium-144 with a half-life of 285 d. These three isotopes have large direct fission yields and are accurate tracers of burn-up. Detecting neutrinos from LLI is a direct method to find reprocessing wastes or spent fuel and, in principle, also yields an estimate of the amount of plutonium separated. Given the high penetrating power of antineutrinos, this method is equally applicable to buried wastes.

Finally, antineutrinos can travel arbitrary distances, and thus an antineutrino detector deployed for safeguarding the IRT would also be sensitive to neutrinos from the 5 MW_e, especially during times when the IRT is shut down. This signal will allow a remote power measurement which can distinguish the two cases shown in Figure 4.

The Yongbyon 5 MW_e Reactor

Sensitivities to power, burn-up, and plutonium content are determined based on the declared power history. This history is displayed as solid curves in the various figures in this section. Comparisons are made to a hypothetical undeclared core swap to a fresh reactor core during the 70 day shutdown period, displayed as dashed curves. The difficulty in determining the difference between the two curves lies in the fact that after 1992, power and burn-up are the same for both scenarios. After the first inspection, all the fission rates from the four primary fissioning isotopes are identical with or without diversion. For the following analyses, a standard 5 t detector at 20 m standoff from the reactor is used, which for a data taking period of one year corresponds to about 95,000 events.

A power sensitivity computation is first considered. The analysis is done using the following χ^2-function

$$\chi^2 = \sum_i \frac{1}{n_i^0} \cdot \left[\left(N P_{\text{th}} \sum_I F_I S_{I,i} \right) - n_i^0 \right]^2 , \tag{13}$$

where F_I is the fission fraction for isotope I, n_i^0 is the measured number of antineutrino events in energy bin i, and $S_{I,i}$ is the antineutrino yield in energy bin i for isotope I. P_{th} is the thermal power and N is a normalization constant. Moreover, the fission fractions F_I are subject to a normalization constraint as given in Equation (6).

This analysis assumes precise knowledge of the distance from the reactor to the detector and treats them both as points. Any uncertainty in the

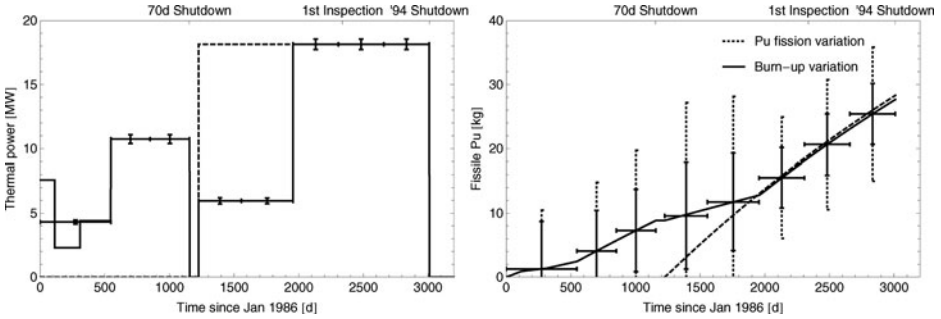

Figure 5: In the left panel, 1σ sensitivities to reactor power are shown for varying data collection periods using a 5 t detector at 20 m standoff from the 5 MW_e reactor. Fission fractions are free parameters in the fit. In the right panel, 1σ sensitivities to burn-up are shown as solid black error bars, where power is a free parameter in the fit. Black dashed error bars show the 1σ sensitivity by measuring the plutonium fission rates with uranium fission rates and reactor power free in the fit. Solid black error bars show the 1σ sensitivity determined by constraining the burn-up using a reactor model. The solid curve shows the history under the assumption of no diversion. The dashed curve shows history for the case of a full core discharge in 1989.

geometric acceptance will directly relate into an uncertainty of the normalization constant, N, and thus into an uncertainty in the power P_{th}. Neglecting this potential source of systematic uncertainty, a power accuracy of around two percent can be achieved.

A similar analysis can be done for the burn-up, BU, using Equation (13). Now, P_{th} is free in the fit and the fission fractions F_I are now functions of burn-up, determined by a reactor core simulation. The results of this analysis are shown in Figure 5. Burn-up across the history of the reactor has an error of approximately 100 MWd/t. Closely related to the burn-up is the amount of plutonium in the nuclear reactor. This analysis is done again using Equation (13). This time, P_{th} as well as F_{U235} and F_{U238} are free parameters as well as the relative contribution of the two plutonium fission rates, κ, and the resulting sensitivities are shown as dashed black lines in the right panel of Figure 5. Alternatively, one can use the burn-up sensitivity to constrain the plutonium content as well. After computing burn-up errors, a reactor model is used to compute the change in plutonium fissions. This is shown as the solid black error bars. These errors are given both in terms of raw plutonium fissions in the left panel as well as the corresponding plutonium masses in the right panel.

5 MW_e REACTOR POWER MEASUREMENT AT IRT

An additional benefit of having a antineutrino detector at the IRT reactor is that it would also be sensitive to antineutrinos from the 5 MW_e reactor. Note that resulting signal rates will be very small and our approximation of neglecting backgrounds is not well justified (see the section on the impact of

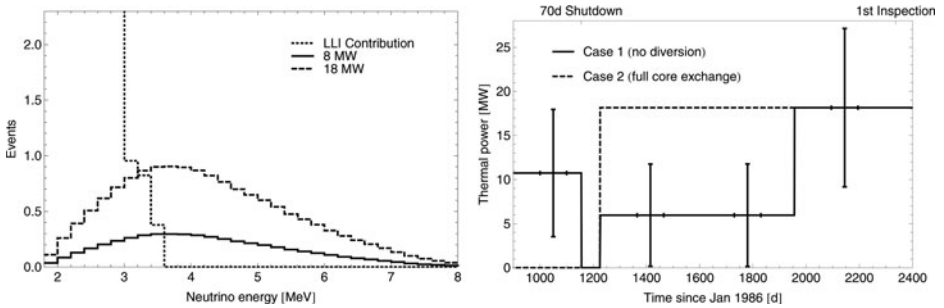

Figure 6: In the left panel events are shown for 200 days of data collection 20 m from the shut down IRT reactor and 1.2 km from the running 5 MW_e reactor. The IRT is assumed to only contribute to the detected antineutrino spectrum through its long lived isotopes shown as a dotted line. The 5 MW_e reactor is assumed to be running either at the declared 8 MW_{th}, shown as a solid curve, or at 18 MW_{th}, shown as a dashed curve. The right panel shows the 1σ sensitivities to reactor power resulting from this measurement. The solid curve shows the power history under the assumption of no diversion. The dashed curve shows the power history if there had been diversion.

backgrounds). This measurement is particularly useful during times when the IRT is shut down, which happens for approximately 100 days every year.[47] This will yield two measurement periods of 100 days each for the reactor power of the 5 MW_e reactor during the crucial time, after the 70 d shutdown and before the first inspection, where the declared power was low, around 8 MW_{th}, but would have been as high as 18 MW_{th}, in order to bring the second core to the same final burn-up (see Figure 4).

Data collection is assumed to start shortly after an IRT shutdown at a point where all but the long-lived antineutrino producing isotopes have decayed leaving only the LLI: strontium-90, ruthenium-106, and cerium-144. This occurs on the order of days. The number of atoms for each of the LLI is: 3.4×10^{23} for strontium-90, 2.8×10^{22} for ruthenium-106, and 2.5×10^{23} for cerium-144. As in the previous sections, a 5 t detector is used at 20 m standoff from the IRT and 1.2 km from the 5 MW_e reactor (see Figure 3). Data are collected over two 100 day periods and the detected spectrum is shown in the left panel of Figure 6. The signal event numbers are small. Thus, the appropriate Poisson log-likelihood is used to define the χ^2-function

$$\chi^2 = 2\sum_i \left[n_i \log\frac{n_i}{n_i^0} - \left(n_i - n_i^0\right)\right] \quad \text{with} \quad n_i = NP_{\text{th}} \sum_I F_I S_{I,i} + LLI_i\,, \tag{14}$$

where LLI_i is the long lived isotope contribution in the bin i. Resulting sensitivities are shown in the right panel of Figure 6. This corresponds to an uncertainty of about 3.8 MW_{th} during the periods of interest. The difference in reactor power for a second core would be detected at 3.2 σ.

WASTE AND SPENT FUEL DETECTION

In addition to directly monitoring reactors, antineutrino detectors can be used for detection of nuclear waste or spent fuel. With sufficient insight of where waste might be disposed, a nearby antineutrino detector can see the signature of LLI, even after years of storage. In contrast to conventional techniques, heavily shielded or buried material can also be detected. The numbers of atoms of each of the three primary LLI are 1.2×10^{24} for strontium-90, 1.4×10^{22} for ruthenium-106, and 3.7×10^{22} for cerium-144. These number reflect what would be expected in the waste (or spent fuel) at the point in time of the first inspection, roughly three years after the 70 day shutdown. In the following analysis, it is assumed that the complete core was removed during the 70 day shutdown and the resulting reprocessing wastes are stored together in one of three locations: the "suspected waste site," building 500, or the Radiochemical Laboratory.[48] All three locations are shown in Figure 3. For building 500, the detector cannot be deployed inside the hatched area, since this facility was declared to be a military installation exempt from safeguards access.[49] The resulting standoff distances range from 25–100 m.

Due to the low event statistics, a Poisson log-likelihood is used, as in Equation (14), with the difference that the reactor events from the 5 MW_e are now background and the signal are the LLI_i. Only for the smallest standoff around 25 m the signal of 8.9 events per year can be detected, either at the suspected waste site or the reprocessing plant with reactor backgrounds of 35.3 and 11.8 events per year, respectively. This leads to 95 percent confidence level detection times of 0.33 and 0.15 years.

In this calculation, all non-reactor sources of background are neglected, which clearly is not realistic (see the discussions of background impacts in the correspoding section). Also, the energy of the antineutrinos from LLI is below 3.6 MeV and therefore a detector which can detect antineutrinos with nearly full efficiency at the IDB threshold of 1.8 MeV needs to be assumed. This requirement may be in direct conflict with the need to control backgrounds.

CONTINUOUS REACTOR POWER MONITORING USING ANTINEUTRINOS

An antineutrino detector present for the lifetime of the reactor would measure reactor power precisely (see Figure 5) and the declared power would have to match the measured power at all times and, because the burn-up is the time integrated thermal reactor power, the burn-up could be inferred from the complete power history. At the same time, a burn-up measurement, in contrast to an inferred burn-up value, can also be derived from a antineutrino measurement. The diversion scenario that has been considered relies heavily upon the

ability to adjust the power relative to the declared power so that both the power and burn-up match at a later time. In the presence of a antineutrino detector, the difference in burn-up will be frozen between the declared burn-up and the actual burn-up of the new core. The fact that a antineutrino detector can simultaneously measure power, as well as fission fractions, is what allows it to detect this difference in burn-up. To determine sensitivity to such a situation, a modified version of Equation 13 is used

$$\chi^2 = \sum_t \sum_i \frac{1}{n^0_{i,t}} \left[(1 + \alpha_{\text{detector}})\, P^t_{\text{th}} \sum_I F_I(BU^t) S_{I,i} - n^0_{i,t} \right]^2 + \left(\frac{\alpha_{\text{detector}}}{\sigma_{\text{detector}}} \right)^2. \quad (15)$$

where t is indexing the time interval for each measurement. α_{detector} is a detector normalization parameter with uncertainty $\sigma_{\text{detector}} = 1\%$. P^t_{th} is the average reactor power in each time bin t. F_I are the fission fractions which are a function of the burn-up in each time bin t, BU^t. The burn-up as a function of time is given by

$$BU^t = \left(\sum_{\tau=1}^{t-1} \frac{P^\tau_{\text{th}} \Delta\tau}{M_{\text{core}}} \right) + BU^0 \quad (16)$$

where $\Delta\tau$ is the width of the time bin, BU^0 the initial burn-up at the start of data taking and M_{core} the mass of the reactor core in terms of fuel loading. If this initial burn-up BU^0 is well known, as it would be if data collection began at start-up, such an analysis greatly reduces the uncertainty in the total plutonium budget. In Table 2, the total error budget is given through the use of this method, labeled "method 2," and is shown compared to the results if only the burn-up but *not* the power history is measured based on the results of the previous sections, labeled "method 1." For method 2, the assumption is that the reactor starts with a well known composition, that is $BU^0 = 0$ and all the P^t_{th} are free parameters in the fit. The question is, what is the maximum change in BU^x during the 70 day shutdown. The value of BU^x is translated into the resulting plutonium mass sensitivity by using the reactor model. It is clear that method 1 is less accurate but does not rely on a continuous observation whereas method 2 is much more accurate but requires continuous observation. The 5 MW_e reactor plutonium error shown in Table 2 is a combination of removed plutonium that may have been removed during the 70 day shutdown and the final plutonium content in the reactor at the 1994 shutdown. The quantities are independent if data are only taken after the first inspection and correlated if taken from start-up. The flat burn-up analysis adds a fixed burn-up to each time bin and the final plutonium error is the final plutonium difference between the burn-up increased data and the expected data. The power constrained analysis assumes the starting fuel composition is known and the burn-up is given by the integration of the power with an assumed one percent

Table 2: Plutonium content and 1σ uncertainties are given for two analysis techniques for the 5 MW_e reactor.

		Final		Method 1, 1σ		Method 2, 1σ	
Reactor		Burn-up (MWd/t)	Pu (kg)	Burn-up (MWd/t)	Pu (kg)	Burn-up (MWd/t)	Pu (kg)
5 MW_e from	Core 1	178	8.83	178	9.5*	N/A	
1st inspection	Core 2	648	27.7	95	3.29		
5 MW_e from	Core 1	178	8.83	138 (83)	6.68 (3.76†)	43 (1.9)	2.12 (0.11)
start-up	Core 2	648	27.7	52 (66)	1.81 (2.30†)	6.7 (6.9)	0.23 (0.24)
5 MW_e Core 3		307	14.6	51	2.17	3.2	0.14
5 MW_e Core 4		255	12.3	53	2.36	2.7	0.12

*Using uncertainty from Albright and O'Neill, *Nuclear Puzzle*.
†These two numbers are anti-correlated with a correlation coefficient of −0.962.

detector normalization uncertainty. The plutonium error is the maximum plutonium difference attainable through power increases and fuel removal (in the case of the 5 MW_e reactor). Values are given for 1σ sensitivities for maximizing the plutonium available for core 1 or core 2 respectively. Parenthesis are for uncertainties in cores using only data from the respective section. Core 3 and core 4 are additional fuel loads that are irradiated in the 5 MW_e reactor post-1994 and are added for completeness.[50]

IMPACT OF BACKGROUNDS

So far in this analysis, backgrounds not related to neutrino emissions have been neglected. There are reactor related non-neutrino backgrounds, like gamma rays and neutrons directly from the nuclear reactions inside the reactor, which at a distance of 20 m, outside the reactor building, should not exceed natural radioactivity. The remaining main backgrounds in inverse beta-decay detectors are: accidentals, where two uncorrelated events caused by ambient radiation in the detector accidentally fulfill the delayed coincidence requirements in both time and energy; fast neutron induced backgrounds, where a fast neutron enters the detector without leaving trace and scatters off a proton, which then is confused with the primary energy deposition of a positron, and subsequently the neutron thermalizes and captures like a genuine neutron from inverse beta-decay; β-n backgrounds, where interaction with cosmic ray muons produces a short-lived radioactive isotope which decays by beta-delayed neutron emission, which mimics an antineutrino event. The rate of accidentals is determined by the rate of ambient radioactive decays. Fast neutrons are a result of cosmic ray interactions in materials surrounding the detector and

thus depend on the rate of cosmic ray muons; the same is true for β-n backgrounds. Therefore, the measured background rates due to those two sources have to be scaled from the underground location, where most current antineutrino detectors are located, to the surface.[51] Using the numbers measured in the Double Chooz experiment at a depth corresponding to 300 meter water equivalent (mwe),[52] and scaling the problem to a surface deployed detector, $1\,d^{-1}\,t^{-1}$ fast neutron events and $43\,d^{-1}\,t^{-1}$ β-n events are obtained, where one tonne is assumed to have the composition of CH_2. These rates exceed the accidental rates by a large factor and therefore the accidental backgrounds can be neglected. This scaling is tested against several data sets from different experiments spanning a depth range from 850–120 mwe and the scaling is found to be accurate within a factor of two.[53] At very shallow depths of less than 10 mwe, the hadronic component of cosmic radiation is non-negligible and the scaling is no longer valid. With this caveat in mind, the signal to noise ratio is between 10:1 to 2:1 for the direct reactor measurement at the $5MW_e$ reactor and the IRT. On the other hand, for the IRT parasitic measurement and the waste detection, the background is several hundred to thousand times larger than the signal. Clearly, very significant advances in detector technology are needed to move these secondary measurements into the realm of the feasible. The required rejection factor can be reduced significantly by providing a moderate overburden of 10–20 mwe, which, in principle, can be engineered into the detector support structure.

Fortunately, there is a significant on-going experimental effort in several countries to address the R&D for antineutrino detectors with greatly improved background rejection. These initiatives are motivated by the search for a new particle called a sterile neutrino through the use of antineutrinos from reactors with detectors placed within meters of the reactor core.[54] The close proximity to a reactor core results in a high-background environment which can include a significant flux of fast neutrons and high-energy gamma-rays from the reactor itself. Almost all reactor sites under consideration offer only very minimal overburden of 10 mwe or less. Therefore, these experiments face essentially the same magnitude of problems in terms of signal to noise conditions as safeguards detectors would under the conditions outlined in this article.[55]

TIMELINE OF EVENTS UNDER FULL ANTINEUTRINO SAFEGUARDS

The events in 1994 put a premium on understanding the actual history of the North Korean plutonium program and thus actual methods, not relying on antineutrinos, were devised. The basic concept of these methods is to map out the three dimensional burn-up distribution inside the reactor core.[56] One

possibility is to sample about 300 out the of about 8000 fuel rods and to measure their gamma-emission from mostly cesium-137, which is a good proxy for burn-up. According to an expert from Los Alamos National Laboratory who was closely involved in the 1994 DPRK issue, this technique would provide burn-up errors below five percent if good quality calibration data existed.[57]

Another method is based on sampling the reactor graphite and measuring the isotope ratios of certain trace elements to reconstruct a three dimensional map of neutron fluence. This graphite isotope ratio method (GIRM) was first proposed by Fetter in 1993[58] and subsequently developed in considerable detail at Pacific Northwest National Laboratory. The overall accuracy is expected to be in the 1 to 5 percent range.[59] The accuracy of neutrino measurements falls in the same general range as those achievable by conventional means. A marked difference is the level of intrusiveness, which in descending order goes from GIRM, which requires drilling sizable holes into the moderator; to the sampling of spent fuel, which requires considerable access during refueling; and then to antineutrino monitoring, which only requires access to the exterior of the reactor building. Also, antineutrinos are the only method providing information on the scale of weeks concurrent with reactor operation, a significant advantage in the context of break-out scenarios.

Based on the quantitative results and the time-line of events in 1994 the following scenario may have been put into effect:

- The IRT is under full antineutrino safeguards with a dedicated 5 tonne detector from 1978 on, which is located outside the IRT reactor building at the southern wall.
- The 5 MW_e is under full antineutrino safeguards with a dedicated 5 tonne detector from May 1992 on, which is located outside the 5 MW_e reactor building at the western wall.
- A search for antineutrino emissions from the reprocessing waste was initiated in November 1992. Three 5 tonne detectors are deployed: one at the reprocessing plant; one at the suspected waste site, located above the center of the waste site; and one at building 500, located right outside the southern fence.

This scenario is fully consistent with the *actual* safeguards access the IAEA had and, in particular, all detector deployment locations reflect actual physical access. As a result, the detectors at the IRT and 5 MW_e have a standoff of 20 m, the detectors at the suspected waste site and reprocessing plant have a distance of 25 m, and the one at the building 500 has a standoff distance of 80 m.

Furthermore, we assume that in 1989 the DPRK discharged the complete first core, which seems to be corroborated by the declaration of the DPRK in 2008 that it possessed 30 kg of plutonium.[60] After reprocessing of the spent fuel, the waste was stored somewhere in the reprocessing plant.[61] Finally, it is assumed that the burn-up declared by the DPRK in 1992 is indeed correct.[62]

The first relevant piece of data would be obtained by the IRT detector during periods when the IRT is shut down, about 100 days out of each year. The antineutrino signal stemming from the operation of the 5 MW_e is clearly detectable at this detector location and provides a measurement of reactor power. In 1989, this signal would have been recorded. Soon after the DPRK had submitted its initial declaration to the IAEA, in May 1992, these data would have resulted in a discrepancy which, in combination with the results from environmental sampling would have led to the conclusion that a large amount of plutonium had been separated in 1989. This measurement has a sensitivity which corresponds to 2.55 kg plutonium. Taking 4 kg of plutonium as the quantity needed for a nuclear bomb,[63] this result translates into a 12:13 confidence that the DPRK has at least enough plutonium for one bomb.

In November of 1993, after a year of data collection, the detectors at the suspected waste site would not have found anything nor would the detector at building 500. The detector at the reprocessing plant would have shown the presence of high-level radioactive waste[64] (or spent fuel), corresponding to a plutonium accuracy of 1.67 kg. Moreover, with high confidence it would have been known that enough plutonium for one weapon was processed (or at least removed from the reactor). Six months later, in May 1994, the 5 MW_e detector would have confirmed the burn-up declaration of the DPRK with an accuracy of 15 percent. In combination, these results would have implied a 56 percent chance of there being enough plutonium for two or more bombs.

Overall, had the DPRK allowed the detectors to be installed and operated, antineutrino safeguards would, in the scenario considered, have changed the state of information significantly. The existence of a separate first core would have been established with very high confidence and the fact that this core was reprocessed (or had been moved to the reprocessing facility) would have been known with very high confidence. It would have been known, with very high confidence, that at least enough plutonium for one bomb has been separated. There would have been some indication that there could be enough separated plutonium for two bombs; and all of this information would have been available by the end of 1993.

There are many arguments which can be levied against the scenario described in the preceding paragraphs. On the technical side, antineutrino detectors which achieve the required level of background rejection did not exist in the 1990s and do not exist now, at least not with demonstrated capabilities.

CONCLUSION

Antineutrino reactor monitoring offers unique capabilities which seem to make this method, as proposed more than 30 years ago, a useful tool for safeguards. Also, antineutrino detectors have been continually refined since the days of Cowan and Reines and can be considered a mature technology. Given the mechanisms of antineutrino production and detection, antineutrino safeguards provide bulk measurements of reactor core parameters like power or burn-up. This is to be contrasted with the current safeguards approach which largely relies on item accountancy and, in particular, neither power nor fuel burn-up are routinely measured by the IAEA[65] or verified through independent calculations for any reactor.

Using the 1994 North Korean nuclear crisis as a virtual laboratory, it was shown that antineutrino reactor monitoring could enable the IAEA to detect unreported plutonium production or to infer the diversion of declared plutonium at the one significant quantity level, i.e., 8 kg of plutonium within 90 days at 90 percent confidence level at light-water moderated reactors producing less than 1 GW_{th} power and at heavy-water moderated reactors producing less than 0.1 GW_{th} power. Antineutrino-based safeguards allow conclusions about the plutonium content of a reactor core and the detection of potential diversion, to be drawn in a timely fashion, whereas conventional methods provide the information only after the fact, once the reactor is shut down and defueled. For all of these applications, antineutrino detectors have to work with minimal or no overburden and the lower the residual background is, the more versatile the resulting system will be.

To conclude, the results presented in this article suggest that antineutrino safeguards would be relevant and an interesting choice to monitor various types of research reactors and for most of the planned commercial small modular reactors. For the same reasons, the Arak heavy-water reactor, located in the Islamic Republic of Iran, could be an ideal test bed for antineutrino safeguards.

ACKNOWLEDGEMENTS

We thank Robert Gallucci, Chris Gesh, Olli Heinonen, Howard Menlove, and Bruce Reid for their willingness to be interviewed for this project. We also thank A. Erickson, L. Kalousis, J. Link, C. Mariani, and in particular T. Shea for their expert opinions on many of the technical issues involved regarding nuclear reactors, antineutrino detectors, and safeguards. We thank M. Fallot for providing reactor antineutrino fluxes in machine readable format and we also acknowledge useful discussions about solid, segmented antineutrino detectors with A. Vacharet and A. Weber.

FUNDING

This work was supported by the U.S. Department of Energy under contract DE-SC0003915 and by a Global Issues Initiative grant by the Institute for Society, Culture, and Environment at Virginia Tech.

NOTES AND REFERENCES

1. C. L. Cowan et al., "Detection of the Free Neutrino: A Confirmation," *Science* 124 (1956): 103–104.

2. A. A. Borovoi and L. A. Mikaelyan, "Possibilities of the practical use of neutrinos," *Soviet Atomic Energy* 44 (1978): 589.

3. A. Bernstein et al., "Nuclear Reactor Safeguards and Monitoring with Anti-neutrino Detectors," *Journal of Applied Physics* 91 (2002): 4672; M. M. Nieto et al., "Detection of Anti-neutrinos for Nonproliferation," *Nuclear Science and Engineering* 149 (2003): 270–276; P. Huber and T. Schwetz, "Precision Spectroscopy with Reactor Anti-neutrinos," *Physical Review* D70 (2004): 053011; A. C. Misner, "Simulated Antineutrino Signatures of Nuclear Reactors for Nonproliferation Applications" (PhD thesis, Oregon State University, 2008); A. Bernstein et al., "Nuclear Security Applications of Antineutrino Detectors: Current Capabilities and Future Prospects," *Science and Global Security* 18 (2010): 127–192; V. Bulaevskaya and A. Bernstein, "Detection of Anomalous Reactor Activity Using Antineutrino Count Rate Evolution Over the Course of a Reactor Cycle," *Journal of Applied Physics* 109 (2011): 114909; A. C. Hayes et al., "Theory of Antineutrino Monitoring of Burning MOX Plutonium Fuels," *Physical Review C85* (2012): 024617; P. Huber, "Spectral Antineutrino Signatures and Plutonium Content of Reactors," in Proceedings of the 53rd Annual Meeting of the Institute for Nuclear Materials Management (INMM) (Institute of Nuclear Materials Management, 2012), 7797.

4. J. S. Wit, D. Poneman, and R. L. Gallucci, *Going Critical: The First North Korean Nuclear Crisis* (Washington, DC: Brookings Institution Press, 2007).

5. Borovoi and Mikaelyan, "Possibilities of the Practical use of Neutrinos."

6. V. A. Korovkin et al., "Measuring Nuclear Plant Power Output by Neutrino Detection," *Soviet Atomic Energy* (1988): 712–718.

7. Y. V. Klimov et al., "Measurement of Variations of the Cross Section of the Reaction $\bar{\nu}_e + p \rightarrow e^+ + n$ in the $\bar{\nu}_e$ flux from a Reactor," *Soviet Journal of Nuclear Physics* 51, 2 (1990): 225–258.

8. Huber and Schwetz, "Precision Spectroscopy with Reactor Anti-neutrinos."

9. Y. V. Klimov, V. I. Kopeikin, L. A. Mikaélyan, K. V. Ozerov, and V. V. Sinev, "Neutrino Method Remote Measurement of Reactor Power and Power Output," *Atomic Energy* 76, 2 (1994): 123–127; Huber, "Spectral Antineutrino Signatures and Plutonium Content of Reactors."

10. A. Bernstein, N. S. Bowden, A. Misner, and T. Plamer, "Monitoring the Thermal Power of Nuclear Reactors with a Prototype Cubic Meter Antineutrino Detector," *Journal of Applied Physics* 103 (2008): 074905.

11. F. P. An et al., "Observation of Electron-antineutrino Disappearance at Daya Bay," *Physical Review Letters* 108 (2012): 171803.

12. Y. Declais et al., "Search for Neutrino Oscillations at 15-meters, 40-meters, and 95-meters from a Nuclear Power Reactor at Bugey," *Nuclear Physics* B434 (1995): 503–534.

13. Y. Abe et al., "Indication for the Disappearance of Reactor Electron Antineutrinos in the Double Chooz Experiment," *Physical Review Letters* 108 (2012): 131801; An et al., "Observation of Electron-antineutrino Disappearance at Daya Bay"; J. K. Ahn et al., "Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment," *Physical Review Letters* 108 (2012): 191802.

14. Efficiency is defined as the ratio of neutrino interactions identified as such, relative to the total number of antineutrino interactions happening in the detector.

15. Y. Abe et al., "Indication for the Disappearance of Reactor Electron Antineutrinos in the Double Chooz Experiment"; An et al., "Observation of Electron-antineutrino Disappearance at Daya Bay"; Ahn et al., "Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment."

16. P. Vogel and J. F. Beacom, "Angular Distribution of Neutron Inverse Beta Decay," *Physical Review* D60 (1999): 053003.

17. This value is taken from A. Serebrov et al., "Measurement of the Neutron Lifetime Using a Gravitational Trap and a Low-temperature Fomblin Coating," *Physics Letters* B605 (2005): 72–78.

18. S. Oguri et al., "Reactor Antineutrino Monitoring with a Plastic Scintillator Array as a New Safeguards Method," *Nuclear Instruments and Methods* A757 (2014): 33–39.

19. V. Kopeikin, L. Mikaelyan, and V. Sinev, "Reactor as a Source of Antineutrinos: Thermal Fission Energy," *Physics of Atomic Nuclei* 67 (2004): 1892–1899.

20. T. A. Mueller et al., "Improved Predictions of Reactor Antineutrino Spectra," *Physical Review* C83 (2011): 054615; M. Fallot et al., "New Antineutrino Energy Spectra Predictions from the Summation of Beta Decay Branches of the Fission Products," *Physical Review Letters* 109 (2012): 202504.

21. F. Feilitzsch, A. A. Hahn, and K. Schreckenbach, "Experimental Beta Spectra from Pu-239 and U-235 Thermal Neutron Fission Products and Their Correlated Anti-neutrino Spectra," *Physics Letters* B118 (1982): 162–166; K. Schreckenbach et al., "Determination of the Anti-neutrino Spectrum from U-235 Thermal Neutron Fission Products up to 9.5 MeV," *Physics Letters* B160 (1985): 325–330; A. A. Hahn et al., "Anti-neutrino Spectra from Pu-241 and Pu-239 Thermal Neutron Fission Products," *Physics Letters* B218 (1989): 365–368; N. Haag et al., "Experimental Determination of the Antineutrino Spectrum of the Fission Products of ^{238}U," *Physical Review Letters* 112 (2014): 122501.

22. P. Huber, "On the Determination of Anti-neutrino Spectra from Nuclear Reactors," *Physical Review* C84 (2011): 024617.

23. P. Huber, "Anti-neutrino Spectra"; A. C. Hayes et al., "Reanalysis of the Reactor Neutrino Anomaly," *Physical Review Letters* 112 (2014): 202501.

24. Fallot et al., "New Antineutrino Energy Spectra Predictions from the Summation of Beta Decay Branches of the Fission products."

25. P. Huber, "Anti-neutrino Spectra."

26. The problem of antineutrino yields has recently received significant scrutiny. Until the 2011 work by a group from S. Mueller et al., "Improved Predictions," the results by Schreckenbach et al. Feilitzsch, Hahn, and Schreckenbach, "Experimental Beta Spectra from Pu-239 and U-235 Thermal Neutron Fission Products and Their Correlated Anti-neutrino Spectra"; Schreckenbach et al., "Determination of the Anti-neutrino Spectrum from U-235 Thermal Neutron Fission Products up to 9.5 MeV"; Hahn et al.,

"Anti-neutrino Spectra from Pu-241 and Pu-239 Thermal Neutron Fission Products," obtained in the 1980s at the Institut Laue-Langevin in Grenoble were considered the gold standard. The Saclay group revisited the previous results in an attempt to reduce the uncertainties. Instead, they found an upward shift of the central value of the average yield by about 3 percent while the error budget remained largely unchanged. Together with other factors, the previous experiments appear to observe a deficit in neutrino count rate of about 6 percent; this is called the *reactor antineutrino anomaly* and was first discussed by G. Mention et al., "The Reactor Antineutrino Anomaly," *Physical Review* D83 (2011): 073006. The result on increased flux was independently confirmed by one of the authors, Huber, "Anti-neutrino Spectra." A plausible explanation could come in the form of a new particle, a sterile antineutrino, which is not predicted by the Standard Model of particle physics. Given the far-flung consequences of the existence of this sterile antineutrino a considerable level of research activity ensued. For a recent review, see K. N. Abazajian et al., "Light Sterile Neutrinos: A White Paper" (2012), http://arxiv.org/abs/1204.5379.

27. P. Huber, "Spectral Antineutrino Signatures and Plutonium Content of Reactors."

28. F. P. An et al., "A Side-by-Side Comparison of Daya Bay Antineutrino Detectors," *Nuclear Instruments and Methods* A685 (2012): 78–97.

29. SCALE is a modeling and simulation suite for nuclear safety analysis and design developed and maintained by Oak Ridge National Laboratory under contract with the U.S. Nuclear Regulatory Commission, U.S. Department of Energy, and the National Nuclear Security Administration. It is available at http://www.ornl.gov/sci/scale.

30. MTU stands for metric tonne of uranium or heavy metal, i.e., all actinides.

31. D. Albright and C. Walrond, *Update on the Arak Reactor*, Technical Report (Institute for Science and International Security (ISIS), July 2013).

32. International Atomic Energy Agency, *IAEA Safeguards Glossary* (International Atomic Energy Agency, 2002).

33. D. Oberdorfer, *The Two Koreas: A Contemporary History* (Basic Books, 2001), 269.

34. D. Albright, "How Much Plutonium Does North Korea Have?," *Bulletin of the Atomic Scientists* 50, 5 (1994): 46.

35. Ibid.

36. O. Heinonen, Interview by PH, 16 April 2013.

37. Albright, "How Much Plutonium Does North Korea Have?"

38. Wit, Poneman, and Gallucci, *Going Critical*.

39. S. Hecker and W. Liou, "Dangerous Dealings: North Korea's Nuclear Capabilities and the Threat of Export to Iran," *Arms Control Today* 37 (2007): 6–8; S. Hecker, "Lessons Learned From the North Korean Nuclear Crises," *Dædalus Winter* (2010): 44–56.

40. S. Hecker, "Lessons learned."

41. D. Albright and K. O'Neill, eds., *Solving the North Korean Nuclear Puzzle* (Washington, DC: ISIS Press, 2000), 148.

42. In principle, Magnox reactors can be refueled under load, but the DPRK seems not to have mastered this technology at that time.

43. O. Heinonen, Interview by PH, 16 April 2013, Heinonen confirms, looking at Figure VI.2 of *Nuclear Puzzle* that this is an accurate description of the burn-up.

44. Albright and O'Neill, *Nuclear Puzzle*.

45. O. Heinonen, Interview by PH, 16 April 2013, Heinonen estimates that the upper limit is 0.5–1 kg, based on the detailed fuel burn-up data the IAEA obtained as part of its safeguards agreement for the IRT.

46. In principle, the amounts of the long-lived isotopes strontium-90, ruthenium-106, and cerium-144 will be different between the two irradiation histories which leads to differences in the low energy antineutrino spectrum below 3.6 MeV. However, extensive calculations show that the resulting event rate differences in 1992 are too small to be reliably detected.

47. Albright and O'Neill, *Nuclear Puzzle*, 148–149.

48. Albright and O'Neill, *Nuclear Puzzle*; O. Heinonen, Interview by PH, 16 April 2013, Heinonen does not believe that the liquid, high-level waste was transferred to building 500.

49. Albright and O'Neill, *Nuclear Puzzle*, 149–154.

50. D Albright and C. Walrond, *North Korea's Estimated Stocks of Plutonium and Weapon-Grade Uranium*, technical report (Washington, DC: ISIS, August 2012).

51. Y. Abe et al., "Direct Measurement of Backgrounds using Reactor-off Data in Double Chooz," *Physical Review* D87 (2013): 011102.

52. Overburden is commonly quoted in units of meter water equivalent (mwe), typically 1 m of rock/soil corresponds to about 2–3 mwe.

53. Ibid.

54. Abazajian et al., "Light Sterile Neutrinos: A White Paper."

55. Z. Djurcic et al., "PROSPECT—A Precision Reactor Neutrino and Oscillation Spectrum Experiment at Very Short Baselines" (2013), http://arxiv.org/abs/1309.7647; I. Alekseev et al., "DANSSino: A Pilot Version of the DANSS Neutrino Detector," *Physics of Particles and Nuclei Letters* 11 (2014): 473–482; A. P. Serebrov et al., "On Possibility of Realization NEUTRINO-4 Experiment on Search for Oscillations of the Reactor Antineutrino into a Sterile State" (2013), http://arxiv.org/abs/1310.5521.

56. O. Heinonen, Interview by PH, 16 April 2013; B. Reid and C. Gesh, phone interview by PH, 15 May 2013.

57. H. Menlove, personal communication.

58. S. Fetter, "Nuclear Archeolgy: Verifying Declarations of Fissile-Material Production," *Science and Global Security* 3 (1993): 237–259.

59. P. Heasler et al., "Estimation Procedures and Error Analysis for Inferring the Total Plutonium (Pu) Produced by a Graphite-Moderated Reactor," *Reliability Engineering and System Safety* 91 (2006): 1406–1413.

60. Albright and Walrond, *North Korea's Estimated Stocks of Plutonium and Weapon-Grade Uranium*.

61. O. Heinonen, Interview by PH, 16 April 2013.

62. Albright and O'Neill, *Nuclear Puzzle*; O. Heinonen, Interview by PH, 16 April 2013.

63. T. B. Cochran and C. E. Paine, *The Amount of Plutonium and Highly-enriched Uranium Needed for Pure Fission Nuclear Weapons*, Technical Report (Washington, DC: Natural Resources Defense Council, 1995).

64. O. Heinonen, Interview by PH, 16 April 2013; S. Hecker, *Report of Visit to the Democratic People's Republic of North Korea (DPRK) Pyongyang and the Nuclear Center at Yongbyon, Feb. 12–16, 2008*, Technical Report (Center for International Security and Cooperation, March 2008).

65. S. Kutchesfahani and M. Lombardi, "South Africa," in *Nuclear Safeguards, Security and Nonproliferation: Achieving Security with Technology and Policy*, ed. James Doyle (Boston: Butterworth-Heinemann, 2008).