Antineutrino Monitoring for Heavy Water Reactors

Eric Christensen, Patrick Huber,[*] and Patrick Jaffke
Center for Neutrino Physics, Virginia Tech, Blacksburg, Virginia 24061, USA

Thomas E. Shea
TomSheaNuclear Consulting Services, Görgengasse 10/25, 1190 Vienna, Austria

(Received 10 April 2014; published 25 July 2014)

In this Letter we discuss the potential application of antineutrino monitoring to the Iranian heavy water reactor at Arak, the IR-40, as a nonproliferation measure. An above ground detector positioned right outside the IR-40 reactor building could meet IAEA verification goals for reactor plutonium inventories. While detectors with the needed spectral sensitivity have been demonstrated below ground, additional research and development is needed to demonstrate an above-ground detector with this same level of sensitivity. In addition to monitoring the reactor during operation, observing antineutrino emissions from long-lived fission products could also allow monitoring the reactor when it is shut down, provided very low detector backgrounds can be achieved. Antineutrino monitoring could also be used to distinguish different levels of fuel enrichment. Most importantly, these capabilities would not require a complete reactor operational history and could provide a means to reestablish continuity of knowledge in safeguards conclusions should this become necessary.

DOI: 10.1103/PhysRevLett.113.042503 PACS numbers: 28.50.Dr, 01.75.+m, 14.60.Lm, 89.20.Dd

The IR-40 reactor in Iran is of particular concern, since the design thermal power of 40 $\mathrm{MW_{th}}$ combined with the choice of moderator, heavy water, makes this reactor ideal for plutonium production for nuclear weapons [1]. Iran states that this reactor will be used for the peaceful purposes of isotope production for medical uses and scientific research. It remains to be seen whether Iran will operate the reactor at all and, if the IR-40 becomes operational, whether it will operate as designed or with some modifications that make it less amenable to weapon plutonium production [1], or whether an extra-territorial siting arrangement might allay proliferation concerns [2].

If the IR-40 goes into operation, the International Atomic Energy Agency (IAEA) will need to confirm that its operations are as declared, using a combination of methods that are reliable and cost effective. Antineutrino monitoring could complement other methods and provide important additional assurance to the international community that Iran continues to honor its commitments. Existing safeguards methods are ill suited to deal with possible breakout scenarios or situations when inspector access is intermittent. The historic example of the Democratic People's Republic of Korea (the DPRK) and its interactions with the IAEA and the international community from 1992–1994 included both intermittent denials of inspector access and the DPRK's eventual breakout from the Treaty on the Non-proliferation of Nuclear Weapons. As a result, the question of plutonium production in the DPRK prior to 1994 is still unresolved; see, for instance, Ref. [3].

Antineutrino monitoring was first proposed more than 30 years ago [4] and is based on the fact that the number of antineutrinos produced and their energy spectrum depends in a well-defined manner on the reactor power and on the relative contribution to fission from the various fissile isotopes: uranium-235, plutonium-239, uranium-238, and plutonium-241. In a recent analysis [5] we were able to show that the application of antineutrino monitoring would have been able to provide timely information about plutonium production in the DPRK—even given the actual, constrained and intermittent access by IAEA inspectors. We have applied the techniques developed in Ref. [5] to the specific case of the Arak IR-40 reactor in Iran to show that antineutrino detectors could provide the IAEA with a resilient high-level monitoring capability not offered by any other known technique.

The IR-40 is capable of producing 10 kg of weapon-usable plutonium per year. A safeguards regime for the IR-40 must be able to verify that the actual plutonium production agrees with the declarations made by Iran, and that the plutonium produced remains accounted for. Obtaining plutonium from most reactors and, in particular, from the IR-40, requires the reactor to be shut down for the irradiated fuel to be removed. To quantitatively address the diversion problem involving plutonium from a known reactor, two questions have to be distinguished: the total amount of plutonium produced in the reactor and the amount of plutonium actually residing in the reactor core. The former can be inferred from the complete power history of the reactor, whereas the latter requires additional detailed information on the fueling history of the reactor or a method to directly assess the core state in terms of average fuel burnup. It is the agreement or disagreement of these two quantities, the total produced and actual core plutonium, which may indicate whether or not a plutonium *diversion* has taken place.

The power history of a reactor (total integrated reactor power can also be used to estimate the production of tritium in the heavy water inventory) can be inferred by measuring the primary coolant flow rate and temperature drop using a thermohydraulic monitoring system, a method the IAEA already employs in some research reactors [6]. The core burnup is not usually measured directly but is inferred from knowing the type of fuel that goes into the reactor core and on a burnup calculation based on the power history of the reactor. For discharged fuel typically only the fact that individual fuel elements emit intense ionizing radiation is verified using Cherenkov light. The key to the relatively high reliability of this chain of inferences is to maintain continuity of knowledge (COK) by employing containment seals and surveillance measures.

Once COK is lost, recovery is difficult and may be limited. More sensitive monitoring methods are available to detect complex removal scenarios, although these methods are seldom used because they require isolating individual fuel elements, require lengthy measurement periods, and are expensive to employ. Antineutrino monitoring could provide a robust and nonintrusive alternative method to recover from a loss of the COK.

Consider a hypothetical IR-40 example inspired by the historic DPRK record: assume that there has been full safeguards access for $N - 1$ months but, in the Nth month, COK is lost. Assume further that the reactor is shut down at the beginning of the Nth month. There could be many reasons for such events to happen, spanning the gamut from legitimate operational reasons to a mere technical glitch over a diplomatic standoff, to an attempt at proliferation with a wide range of measures taken to delay detection and reprisal. (For an extended period without inspector access, secondary means of monitoring reactor operation, e.g., infrared satellite imaging, could detect reactor operation and provide a rough estimate of reactor power.) In the Nth month of our hypothetical IR-40 scenario, the power history is interrupted, but for a sufficiently short time such that the extra burnup that could be achieved is very limited and therefore, does not play a major role. But did a refueling take place? The basic task is to reestablish verifiable knowledge of the core state without being able to rely on a power record or uninterrupted containment and surveillance.

In Ref. [5] we showed that measuring the composite energy spectrum of antineutrinos emitted from a reactor could allow the burnup and, thus, the plutonium content to be estimated accurately and in a timely manner. We make the same assumptions here about the detection system as in Ref. [5], i.e., 4.3×10^{29} target protons at a hypothetical efficiency of 100%; Table I shows how this requirement translates into actual detector mass. We envisage a system where the whole detector with supporting electronics fits inside a standard 20 ft shipping container, i.e., a mass of less than 21.6 ton. Smaller detectors would also work but the times required to achieve the performance we cite would be correspondingly longer. Furthermore, we assume a neutrino energy threshold of 1.8 MeV (a threshold of 4 MeV would reduce sensitivity by about 60% for reactor-on measurements) and sufficient background rejection capabilities to allow for surface deployment, for a recent proof of principle demonstration see Ref. [7]. Additional research and development into surface-level background rejection is needed to demonstrate a system with the capabilities outlined here. Even moderate overburden would greatly reduce the need for such research and development, and likely would allow existing detector designs to meet our stated goals. We estimate the diameter of the IR-40 reactor containment building to be approximately 34 m and, therefore, with the shipping container positioned right against the exterior of the reactor containment building and the reactor core center being 5 m above grade level, the antineutrino detector would be located 19 m from the center of the reactor core. (more precise distances could be obtained during design information verification activities at the IR-40.)

TABLE I. Actual detector mass in ton as a function of efficiency for a mineral oil based liquid scintillator (EJ-321 L) with 8.6×10^{22} protons per gram and a polyvinyltuloene based solid scintillator (EJ-200) with 5.1×10^{22} protons per gram.

Efficiency [%]	25	40	60	80
Liquid scintillator	20.1	12.5	8.4	6.3
Solid scintillator	34.0	21.3	14.2	10.6

Assuming the reactor is running at full power when inspector access is resumed, following the methods given in Ref. [5], the antineutrino emissions could be used to determine the core plutonium content and, thus, to also determine whether or not the reactor had been refueled during the period when the inspectors were not allowed access. This burnup based analysis relies on standard reactor physics calculations made using commercially available software (we carried out a reactor simulation of the IR-40 using the two dimensional neutron transport analysis code NEWT and the depletion code Origen. Both codes are from the SCALE software suite [8]). It provides a means to correlate the fission rates of the various fissile isotopes in the reactor core. For our hypothetical IR-40 example, we assumed that the core in its original configuration contained 10 ton of natural uranium dioxide, and that the reactor ran at its design power of 40 $\mathrm{MW_{th}}$. Our model was derived from a full three dimensional analysis developed by Willig, Futsaether, and Kippe [9]. Our results in terms of isotopic composition for the major fissile isotopes and all of the main plutonium isotopes agree to within 1%–2% with the corresponding values reported by Willig, Futsaether, and Kippe.

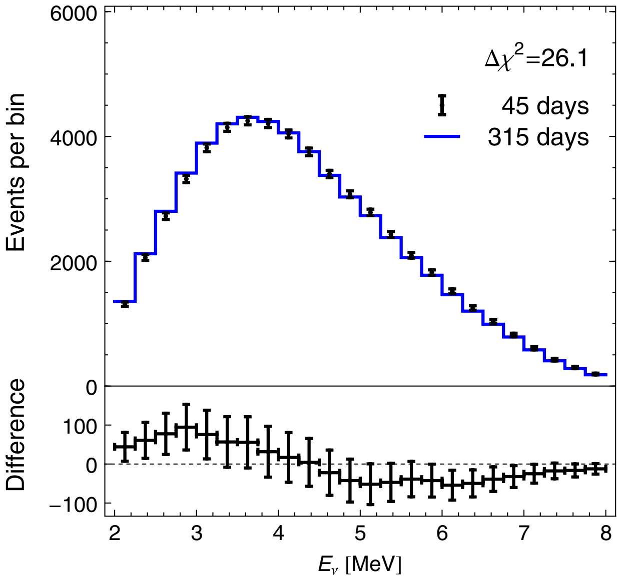

FIG. 1 (color online). In the upper panel, data points show the event rate spectrum obtained in a 90 d data taking period for a core of average age of 45 d. The error bars indicate the statistical error in each bin. The blue line indicates the corresponding expected event rate spectrum for a core of average age of 315 d. The lower panel shows the difference in event rates between the 45 d core and the 315 d core and the corresponding statistical error bars.

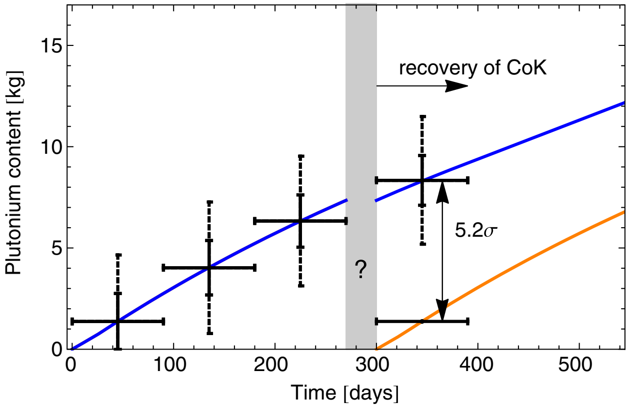

FIG. 2 (color online). Shown is the 1σ accuracy for the determination of the plutonium content of the reactor as a function of time in the reactor cycle. The data taking period is 90 d each. Dashed error bars indicate the accuracy from a fit to the plutonium fission rate f_{Pu}, whereas the solid error bars show the result of a fit constrained by a burnup model. The blue/dark line indicates operation without refueling and the orange/light line indicates operation with a refueling after 270 d.

In Fig. 1 we show the resulting event rate spectrum for a core of 45 d average age (data points with statistical error bars) and for comparison the expected event rates for a core of 315 d of age (blue line). Clearly, the older core has a much softer antineutrino spectrum, which is because of the much higher plutonium content as fission of plutonium produces a softer antineutrino spectrum. The difference in χ^2 between the two cores is 26.1 units corresponding to about 7 kg difference in plutonium content. The visibility of this effect does not rely on extremely good energy resolution since the spectral feature is essential bimodal: below about 4 MeV the rate goes up and above it goes down. As far as systematics is concerned, the main impact will come from the uncertainty on the energy scale and assuming a reference detector at a reactor of known core state is available, the relative energy scale uncertainty between detectors will be appropriate; using the Daya Bay measured value of 0.2% [10] the plutonium accuracy decreases by about 60%.

The quantitative results of our IR-40 analysis in terms of plutonium content are shown in Fig. 2, where the vertical axis shows the amount of plutonium in the reactor core as a function of time. The blue curve shows the evolution of plutonium content assuming that no undeclared refueling has taken place, whereas the orange curve assumes that the previously irradiated core, containing 8 kg of plutonium, was replaced with a fresh core after 270 d of irradiation. Here, 270 d was chosen since according to Willig, Futsaether, and Kippe the content of plutonium-239 drops to 93% after 270 d and thus 270 d represents the longest operational period that still yields weapon-grade plutonium (even lower grade plutonium can be (and has been) used to make nuclear explosives and 93% does not constitute a sharp boundary). Within the first 90 d after the putative IR-40 shutdown (shown as gray vertical band) the two cases would be distinguished unequivocally by analyzing the antineutrino monitoring data. Even partial core refuelings corresponding to as little as 2 kg of removed plutonium could be detected at 90% confidence level. Alternatively, a full core refueling would be detected within about 9 d at 90% confidence level.

If the IR-40 remains shut down after the loss of COK, the antineutrino detector still offers a method to assess the core state by measuring the antineutrino emissions from the long-lived fission fragment isotopes: strontium-90 with a half-life of 28.9 y, ruthenium-106 with a half-life of 372 d, and cerium-144 with a half-life of 285 d. In the decay chains of these three isotopes, antineutrinos are emitted with sufficient energy to be detected by a standard antineutrino detector using inverse beta decay. These long-lived fission fragment isotopes have direct fission yields in the percent range and thus their abundance is large and directly proportional to the burnup of the fuel. By measuring these antineutrino emissions it could be possible to assess the approximate fuel burnup and plutonium content, and to determine whether a major removal of spent fuel had taken place.

The measured antineutrino rates from these fission products would be much smaller than the antineutrino measurement rates during reactor operation. Overall, this measurement puts very stringent requirements on detector

TABLE II. Long-lived isotope measurement: The event rates are the total rates integrated over energy between 1.8 and 3.6 MeV reconstructed neutrino energy and the time to detection (TTD). Results are shown for different times since shutdown (TSS), background (BKG) scaling, and confidence level (CL).

		270 d reactor runtime					
BKG	TSS	90% CL			99% CL		
factor	[d]	Signal	BKG	TTD [d]	Signal	BKG	TTD [d]
2	30	81	3923	56	226	12471	178
	0	37	806	23	93	2102	60
1	30	39	911	26	101	2487	71
	90	46	1261	36	120	3503	100
0.5	30	20	228	13	49	578	33
		180 d reactor runtime					
1	30	55	1857	53	152	5780	165

performance and to demonstrate those capabilities will require significant research and development. In Ref. [5] we estimated (based on data from [11]) that there will be about 43 background events per day per ton of detector from beta-delayed neutron emission from cosmogenically produced lithium-9 and about 1 background event per day per ton from fast neutrons; this neglects any contributions from the hadronic component of backgrounds which are to be expected at the surface. Dealing with the fast neutron background at the surface will require excellent particle identification, which needs to be demonstrated. In Table II we show the time required to achieve a 90%/90% confidence level detection of removal of all the spent fuel contained in the reactor core as a function of the time since shut down, when the core removal occurs and for different background levels. Note, that this result is based on an analysis binned in energy which provides a gain in sensitivity of about 30% compared to a pure rate-based analysis since the signal and background shapes are very different. Also, above 3.6 MeV there is only background and no signal, and thus this sideband can be used to effectively fix the magnitude of the background. As previously stated, the size of the signal is proportional to the burnup of the spent fuel; hence, the longer the reactor has been running the easier this measurement becomes. Even in the low burnup case, which would be characteristic for the production of weapon-grade plutonium, this measurement could be performed with a background rates as estimated: if data taking starts within a month after shutdown, a 90% confidence level confirmation of the presence of the core can be achieved within 30 days or less.

Given the proliferation concerns regarding the IR-40, it has been suggested that the reactor could be modified to make it less suitable for the production of weapon-grade plutonium. One possibility would be to modify the reactor to use low-enriched uranium (LEU) instead of natural uranium (NU) as a fuel, and changing the moderator from heavy to light water [1]. A detailed neutron transport reactor physics calculation has been reported by Willig, Futsaether, and Kippe [9]. They concluded that changing the moderator from heavy to light water could be detrimental to reactor safety. Instead, it has been proposed to use a heavy water moderator together with fuel enriched to 3%, providing a use for the existing Iranian stock of LEU. This LEU configuration could reduce the annual plutonium production from 10 to 3.9 kg with a slightly smaller fraction of plutonium-239.

If LEU fuel were introduced into the IR-40, antineutrino emissions could also be used to distinguish a natural uranium fuel core from a low-enriched uranium configuration by tracking the rate of change in the plutonium fission fractions in the reactor, a technique we term differential burnup analysis (DBA). The basic observation behind DBA is that both configurations follow the same overall burnup pattern: specifically, for the uranium-235 fission fraction F_{U235}, and the plutonium-239 fission fraction F_{Pu239}. Being on the same overall path implies that looking at a single snapshot in time t_1, the resulting single pair of values of $F_{U235}(t_1)$ and $F_{Pu239}(t_1)$ could not be used to distinguish the two configurations. This is illustrated in Fig. 3, where the time evolution of the fission fractions in uranium-235 and plutonium-239 is shown for

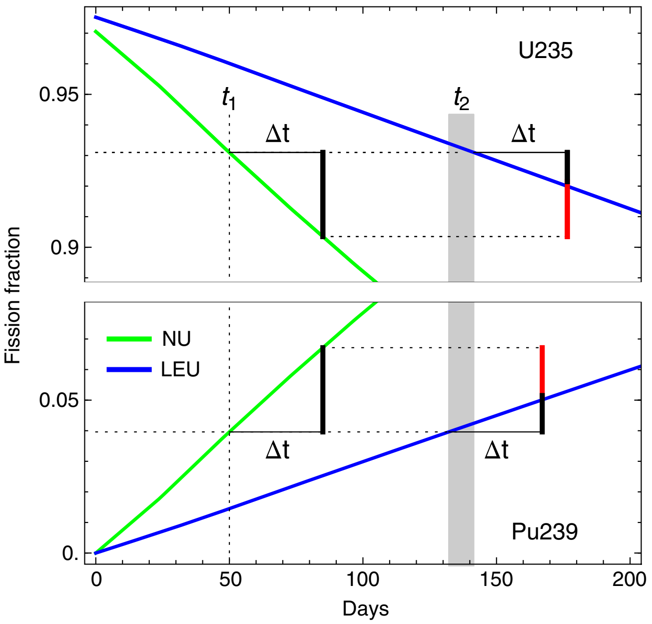

FIG. 3 (color online). Shown are the fission fractions in uranium-235, upper panel, and plutonium-239, lower panel, for a natural uranium fueled core (NU) in green and for a 3% enriched uranium fueled core (LEU) in blue as a function of time elapsed in the reactor cycle. The fission fractions in both isotopes at time t_1 for the NU core match those at a later time t_2 for the LEU core as indicated by the horizontal dashed lines. The change of fission fractions after a fixed time interval Δt, the so-called differential burnup, is indicated by the thick vertical black lines. There is a distinct difference in differential burnup between the LEU and NU cores for both isotopes as indicated by the thick red lines.

both the NU and LEU cores. The pair of fission fractions $F_{U235}(t_1)$ and $F_{Pu239}(t_1)$ for the NU core is nearly identical to the pair $F_{U235}(t_2)$ and $F_{Pu239}(t_2)$ for the LEU core. This identity is approximate since it would require two slightly different values of t_2 for uranium-235 and plutonium-239 to achieve exact identity, as indicated by the width of the gray vertical band. This effect is, however, too small to distinguish the two configurations. The speed at which both configurations move along this path is significantly different, therefore comparing the differential burnup $F(t+\Delta t) - F(t)$, shown as thick vertical lines, for both configurations gives rise to a measurable difference between the NU and LEU cores, shown as thick red lines. Note, that the uranium-238 fission fraction does not contribute to this distinction since it stays constant for both core configurations and the plutonium-241 fission fraction is present only at a very small, basically unmeasurable level. This measurement, due to its time differential nature, is quite robust against systematics, provided the detector is stable in time, as demonstrated by Daya Bay [10]. Applying DBA to the case at hand we find a 90% confidence level distinction between the two configurations solely based on antineutrino measurements within about 167 d.

In summary, we have shown that if antineutrino monitoring of the Iranian IR-40 reactor were instituted, it could provide a complete assessment of the reactor core in terms of burnup and plutonium content with a sensitivity exceeding standard IAEA verification requirements while meeting the timeliness criterion of 90 d. This information could be available in a timely manner and could be obtained by placing a detector outside the reactor building. This technique does not rely on a declaration of reactor power since the power could be inferred from the antineutrino signal simultaneously with the core state. In case the reactor is shut down for extended periods, monitoring antineutrino emissions from long-lived fission products could make it possible to verify the presence of the spent fuel inside the reactor core for up to several hundred days after the shut down. In combination, these techniques could allow a graceful and timely recovery from a loss of the COK. Furthermore, differential burnup analysis could provide a means to distinguish different fuel enrichment levels. Other safeguard methods alone could not achieve this performance, and are likely to be more intrusive and labor intensive.

Antineutrino monitoring would work as well for *any* reactor from a few megawatts thermal power to small modular reactors to large scale commercial nuclear power reactors. Also, it can and should be combined with existing monitoring techniques to enhance effectiveness against a host of future possible developments.

While the results of theoretical analyses are promising, antineutrino reactor monitoring still needs crucial research and development in terms of background rejection and rugged detection systems as well as a precise calibration of reactor antineutrino fluxes. Looking ahead, and noting Iran's willingness to extend IAEA access into aspects of its nuclear program that are not available in other states, Iran may itself wish for the IAEA to include antineutrino monitoring in the safeguards approach for the IR-40, providing a real-world opportunity for a full scale demonstration to enhance the credibility of the global nonproliferation system.

We would like to thank L. N. Kalousis for useful clarifications regarding Ref. [11], and H. Kippe for providing the detailed reactor model of the IR-40 as the starting point for our calculation. This work was supported by the U.S. Department of Energy under Contract No. DE-SC0003915 and by a Global Issues Initiative Grant by the Institute for Society, Culture and Environment at Virginia Tech.

---

*pahuber@vt.edu

[1] O. Heinonen, Foreign Policy (2011), http://www.foreignpolicy.com/articles/2011/01/27/can_the_nuclear_talks_with_iran_be_saved.
[2] T. E. Shea and F. A. Morris, in *Proceedings of the 50th Annual Meeting of the Institute for Nuclear Materials Management (INMM)* (Institute of Nuclear Materials Management, Deerfield, Illinois, 2009).
[3] J. S. Wit, D. Poneman, and R. L. Gallucci, *Going Critical: The First North Korean Nuclear Crisis* (Brookings Institution Press, Washington, DC, 2007).
[4] A. A. Borovoi and L. A. Mikaelyan, Sov. At. Energy **44**, 589 (1978).
[5] E. Christensen, P. Huber, and P. Jaffke, arXiv:1312.1959.
[6] Tech. Report, IAEA (2002), http://www.iaea.org/Publications/Booklets/TeamingInspectors.
[7] S. Oguri, Y. Kuroda, Y. Kato, R. Nakata, Y. Inoue, C. Ito, and M. Minowa, Nucl. Instrum. Methods Phys. Res., Sect. A **757**, 33 (2014).
[8] http://www.ornl.gov/sci/scale.
[9] T. M. Willig, C. Futsaether, and H. Kippe, Science and Global Security **20**, 97 (2012).
[10] F. An *et al.* (Daya Bay Collaboration), Nucl. Instrum. Methods Phys. Res., Sect. A **685**, 78 (2012).
[11] Y. Abe *et al.* (Double Chooz Collaboration), Phys. Rev. D **87**, 011102 (2013).