Neutrinos and Arms Control: Thinking Big about Detection of Neutrinos from Reactors at Long Distances

John G. Learned

Department of Physics and Astronomy, University of Hawaii, Manoa

First Draft late October 2003, revision of February 4, 2004.

ABSTRACT

We discuss a somewhat futuristic plan for a world network of enormous neutrino detectors, which may be employed for monitoring the activity of all reactors on earth. Three (or more) cubic-kilometer size instruments with sensitivity down to about 1 MeV, placed in the deep oceans (or possibly lakes) can record the electron anti-neutrino fluxes from reactors, no matter where they reside. Using known power levels from the roughly 440 operating reactors, one can detect and monitor any new reactor, in particular one which may be producing illicit nuclear weapons material. Such a signal cannot be hidden or jammed, and via tomography may be located to a precision of order 20 kilometers over a one year timescale.

The suggested array would have many ancillary applications, ranging from detection of nuclear bomb tests, to studies of neutrinos from supernovae from throughout our supercluster and seeking the decay of protons to significant levels. First estimates are that such an array of detectors, after industrial development particularly in the area of photodetectors, could cost in the range of a new particle physics accelerator or an aircraft carrier.

I. INTRODUCTION: NEED FOR INTERNATIONAL REACTOR MONITORING

In the present world political environment when nuclear weapons capability is proliferating to more countries, there appears to be a need for the international community to keep track of reactor activities around the world, with the specific focus of monitoring nuclear weapons fuel production. We do not propose herein to justify the need for or cost of such monitoring, but simply note that it is a subject of significant concern. Our goal in this paper is to explore the question of whether this is possible and what would be required to remotely monitor the output of all reactors on earth. The particular question we attempt to address is whether we might detect, localize and measure the output of new, unannounced reactor facilities.

What we shall show in the following is that indeed with some, perhaps substantial, investment in technology development, it may be practical to carry out such a project. Not only do we conclude that such a project is likely to be feasible, but that it may be affordable. It is very important to recognize that such a project is not a single application device. Such monitors would have tremendously important additional dividends, from the detection of clandestine nuclear testing to elementary particle physics, astrophysics, and geophysics. Hence a point in favor of such a plan is that it would attract and engage a large scientific community in realizing and operating the facility.... a win-win engagement between the defense and science communities.

II. GOAL: DETECT REACTORS FROM THE WHOLE EARTH

Without further justification we assert that what is needed is a series of detectors of scale size of about one kilometer or of effective mass of about a gigaton (1 km^3 of water = 10^9 tons), and with a threshold sensitivity in the range of 1 MeV of energy deposition from neutrino interactions. More precisely, what is needed is the ability to record 2-6 MeV electron anti-neutrinos via the inverse beta decay process,

$$\bar{\nu}_{\rm e} + p \rightarrow e^+ + n.$$

The positron annihilates promptly, producing a flash of light proportional to the neutrino energy above the reaction threshold of 1.8 MeV. The neutron then wanders randomly about until it is captured on a proton in the medium, after some hundred microseconds and a distance of order of a meter. This further releases the characteristic 2.2 MeV binding energy of deuterium. The signature then consists of two flashes of light, which must be close in space, time, and in intensity, with the second being close to 2.2 MeV equivalent. This provides a clean signature, almost devoid of background. In fact this is the process that Reines and Cowan used to make the first detections of neutrinos in 1953-6 [1].

Reines and Cowan detected neutrinos at just a few meters distance from the reactors. Experiments over the years have moved to ever-greater distances seeking evidence for neutrino oscillations. In January 2003 the KamLAND experiment reported the detection of neutrinos from reactors all around Japan, with a mean distance of 180 km [2]. Indeed they have made the first credible claim for observation of electron anti-neutrino oscillations on earth (which results are in concert with the suspected electron neutrino oscillations of solar neutrinos, and lay to rest the "solar neutrino problem".) The KamLAND detector is illustrated in the following Figure 1. The sensitive volume consists of a 1000 ton balloon of liquid scintillating material, surrounded by 1879 photomultiplier tubes (PMTs), and achieving a response of about 250 photoelectrons (PE)/MeV of deposited energy (perhaps more sensitivity than needed for the reactors, but the experiment is designed to move on to study low energy solar neutrinos).

Figure 2 shows the spectrum of reactor neutrinos, which falls steeply with energy. On the other hand the cross section rises swiftly, with energy above threshold squared, as shown. Thus the net detected rate has a maximum around 4 MeV, extending from about 2 to 8 MeV in neutrino energy.

The initial report from KamLAND is summarized in Figures 3 and 4, comparing rates as a function of distance with previous experiments, and showing the observed spectrum of events. This result represents the culmination of almost 50 years of efforts in the detection of neutrinos from reactors, with instruments growing larger and more sensitive with time.

Figure 4 shows the spectrum of events as observed, as expected without oscillations and with the best-fit oscillations parameters. Note that at energies below 2.6 MeV there is significant background due to neutrinos from radioactive decays in the earth, mostly Uranium and Thorium.

The effective energy range for reactor detection is thus about 1.0 to 6.0 MeV in the prompt energy (positron energy).

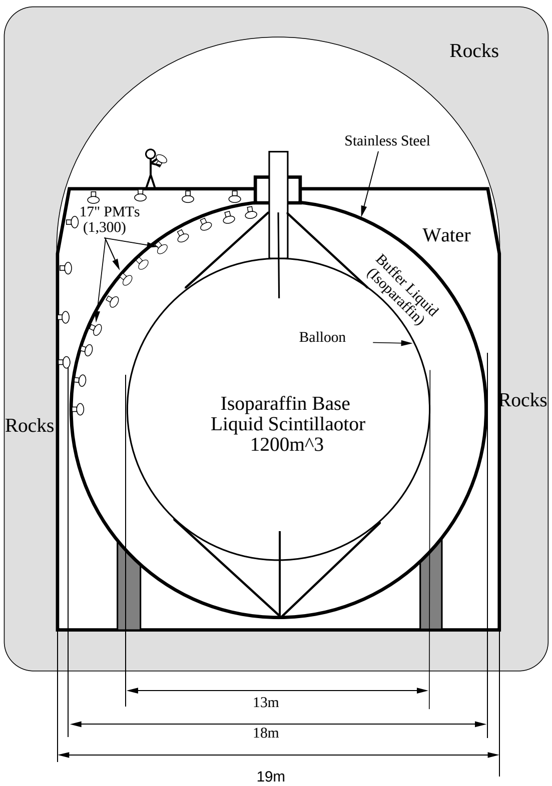

FIG. 1. Schematic cross section of the 1000 ton KamLAND anti-neutrino detector in Japan, which sense reactors hundreds of kilometer distant[3].

One can scale the reactor detection rates from these KamLAND the measurements. The rate is

$$R \cong 832/day \times \left(\frac{P}{1\ GWt}\right) \times \left(\frac{1000\ km}{D}\right)^2 \times$$

$$\left(\frac{M}{1\ GTon}\right) \times F\left(E, D/E\right)$$

where P is the nominal reactor thermal power (about 3 times the electrical power if a power reactor), D is the distance (km), M is the detector effective mass (Gigatons, or $1 km^3$ of pure water), and F(E,D/E) is a term taking oscillations into account. The latter term will be 0.5 to 1.0 depending upon energy and distance (about 0.6 at KamLAND for the reactors around Japan).

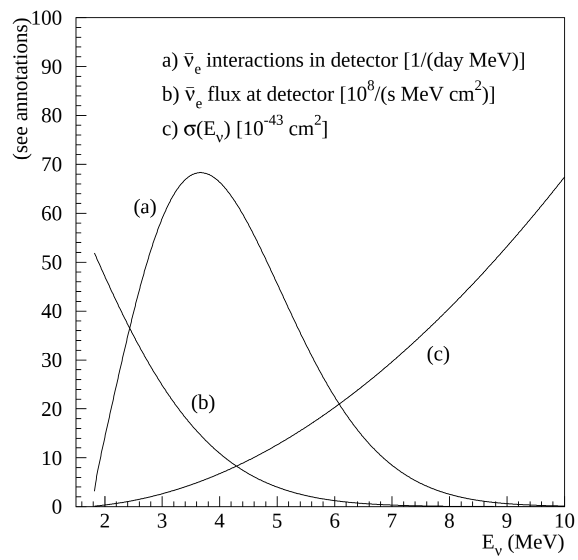

FIG. 2. The spectrum of neutrinos from a typical reactor (b), cross section for inverse beta (c) and net event spectrum, above the 1.8 MeV threshold energy (a)[3].

If the detector location is deep underground or under water (roughly > 4 km water equivalent depth), there will be little background. At lesser depths there will be backgrounds due to down-going cosmic-ray muons generating neutrons which can make false coincidences with other background counts to fake inverse beta decay events. Such are not fatal, but need to be accounted for if a more shallow depth is under consideration. Herein we shall take the simplifying assumption of adequate depth to escape this concern (which might be relaxed in later considerations).

A. Sum of All Reactor Powers

As of 2002 there were a total of 440 power reactors in the world. The total energy produced was 2574 TWe-Hrs in 2002, or the equivalent of 881 GWt average power. The rate in a nominal 1 km^3 detector that is at an average of 6000km distance from these reactors will be about 17,000 nuebar counts/day. Thus from Poisson statistics the rate will be determined to $1\ \sigma = 130$ nuebar counts/day in one day, or one may say that the total rate will be measured with a precision of $0.77\%/\sqrt{(\text{days})}$. Thus a typical 2 GWt reactor at a range of 1000

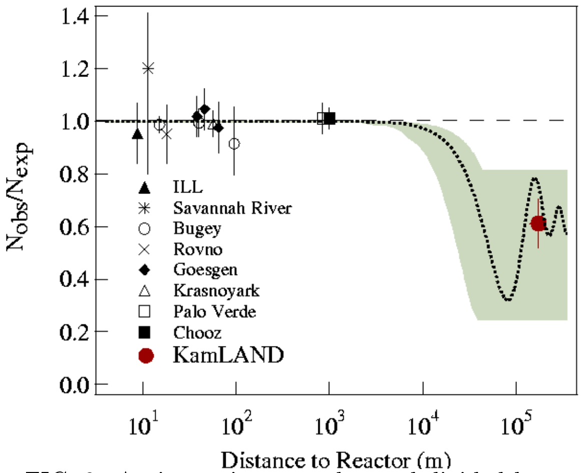

FIG. 3. Anti-neutrino rate observed divided by expected without oscillations as a function of distance showing KamLAND results with effects of depletion due to oscillations.[2]

km producing 1543 nuebar counts/day will produce an excess over the total reactor sum rate of 12 σ each day!

This assumes we know the rest of the world's contribution and that it is stable. Of course the rest of the world's reactors will be turning on and off (though they operate at typically 95% on-time, and when on produce full power), but this information is known or at least knowable. The information about many reactors is already available on the web, day by day, and presumably it could be made available for all reactors by the IAEA, at least to monitoring authorities. The knowledge of the reactor output by cooperating entities needs study, but is at least good to about 1% and there are claims that it may be generally better than that [4]. Moreover the prediction of today's rate based upon yesterday's rate should be good to the order of 1% times something less than 5% (perhaps 2%) (as pointed out to me by Giorgio Gratta). A new clandestine reactor operating at substantial distance would stand out in the received $\bar{\nu}_e$ rate alone. Now we need to discuss whether several detectors can localize the new reactor.

III. DIRECTIONALITY

Inverse beta decay is not inherently directional. Actually there is a slight backwards directionality to the positron production $(1 - 0.102 cos\theta_{e+})$

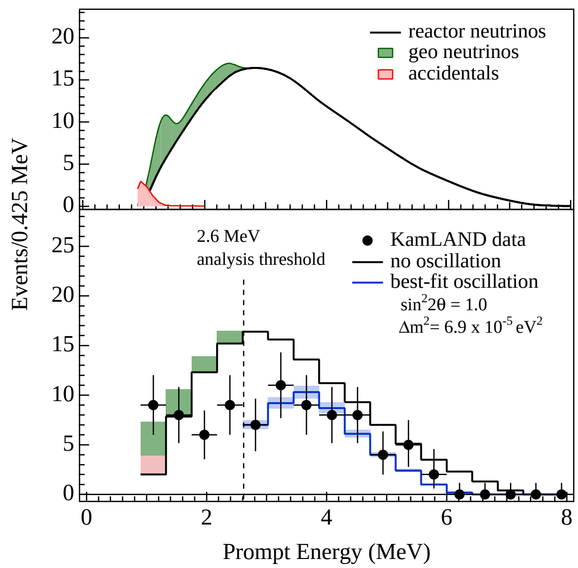

FIG. 4. The KamLAND event spectrum (positron energy) observed, and the expected without oscillations in the smooth curve. The contribution of U and Th decay neutrinos from the earth is indicated in color below 2.6 MeV.[2]

and the neutron acquires some kinetic energy, up to 100 keV. In consequence, there is a weak correlation between neutrino direction and direction from positron annihilation point to neutron capture location. This was observed in the CHOOZ experiment [5]. However, the effect is so weak (18° with 2700 events) that it is probably not useful in locating new reactors at long distances.

There are interactions (neutral current) between the anti-neutrinos and the electrons in the medium, which scatter the electrons in the direction of the neutrino motion. The cross section is down by about a factor of five at relevant energies, and one does not have the nice inverse beta signature. Nonetheless, this reaction needs further investigation to see if information can be extracted. It is not easy to estimate the utility of this at present, since it depends upon the background rate for single counting events from which reactor events must be discriminated (and of course scattering from solar neutrinos, which we can treat as a well known background now!). It has been estimated that in the Super-Kamiokande experiment this reactor signal amounts to a few percent of the total background rate in the 5 MeV range. We will defer but not dismiss further consideration of the potential to employ electron events.

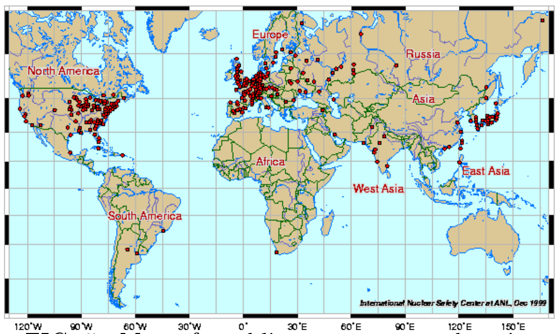

FIG. 5. Map of world's power reactor locations in 2002.

The main method for locating reactors we shall now discuss is tomography, using known counting rates and at least three stations around the world.

A. Tomography

Since most power reactors are well known in output as function of time (1-2%, maybe better), day by day at least (to the IAEA), any new unknown reactor (with say 2GW thermal power) will contribute an average of 43 counts/day at the great range of 6000 km. We can use the relative increases in rates in two well-separated detectors to figure out the locus of possible locations of the new reactor. These loci turn out to be circles on the earth. (The geometry is similar to that of the equipotentials from two point charges, which are spheres. The intersection of two spheres is a circle). If one further knows the reactor power output, then the location is determined. There will be an ambiguity of two solutions, but one is likely to be ruled out by geographical considerations in the real world.

Clearly, if we have three detectors around the world, then we can solve for the unique location, and the power output as well. And we can monitor that power output as a function of time. Further detectors will constrain solutions, eliminate imprecision in cooperating reactor power outputs, and make monitoring of multiple locations possible. It is hard to generalize much at this stage since the distribution of present reactors around the world is very non-uniform. We need numerical studies optimizing the locations of such detectors, consistent with available locations for deployment as well as regions to be specially watched.

To get some idea of the sensitivity of this position detection we have made a simplified numerical calculation. We take two detectors separated by 40 degrees on earth (for simplicity separated in longitude by 40 degrees at the equator), a distance of about 4500 km. We assume a reactor of 2GWt added to the uniform counting rate of 17,000 nuebar events per day at each detector due to all the rest of the world's reactors. In Figure 6 we display the asymmetry $[(n_1\text{-}n_2)/(n_1+n_2)]$ between the counting rate excesses (n_1,n_2), which demonstrates the circularity of the locus of solutions for constant asymmetry (perhaps not so obvious on this Mercator projection). In Figure 7 we show the resolution as a function of location, after one year of counting for a new reactor lying along the equator. Of course in this situation there is no position resolution in latitude. One sees that the resolution is about 20 km in longitude over a substantial region between the two detectors, better if closer to one extreme, and getting worse if on the far side of either detector. We assumed a reactor power as 2 GWt, but if it is 200MWt, ten times smaller, the resolution will be worse by about a factor of about three. The IAEA goal of detecting production of 8 kg of Pu from a 3 GWt reactor in 30 days would appear to be achievable [4].

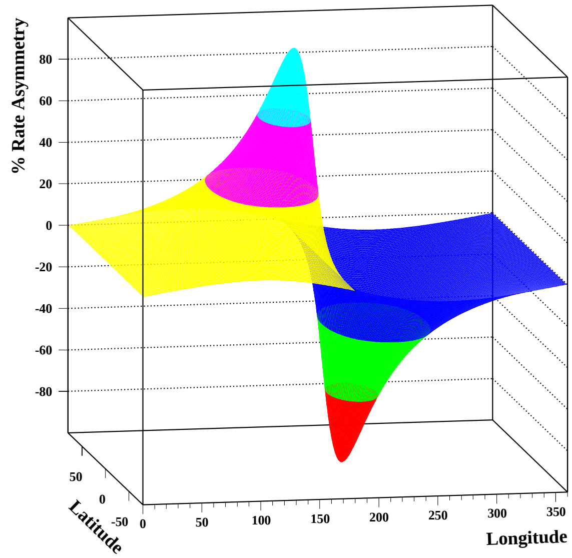

FIG. 6. Counting rate asymmetry for an unknown reactor versus longitude and latitude for two detector arrays located on the equator with longitudes separated by 40 degrees,

Obviously more sophisticated calculations are needed for the three (or more) detector case, and we should include realistic backgrounds in real world geometry of reactor locations, as well as accounting for the effects of oscillations. But the point remains that one can aim for something like 20 km resolution with such a system. Integrating over a year or so of observations, one can measure the power, and one can track the time history of the operation on a day by day basis.

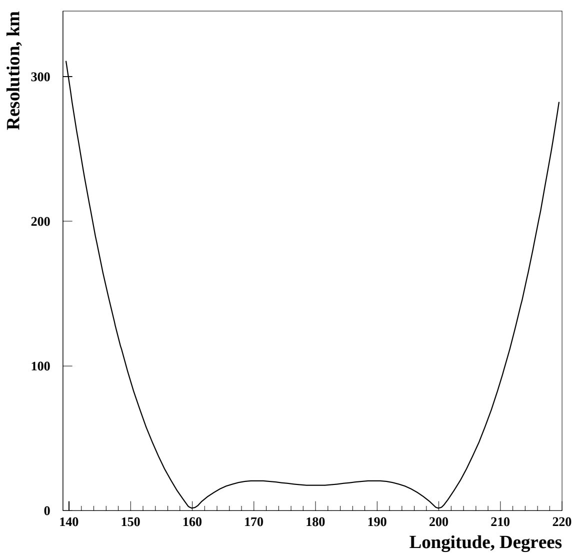

FIG. 7. Numerical example of unknown reactor longitudinal resolution after one year. Two km^3 detector arrays are assumed, at the equator and longitudes separated by 40 degrees. Unknown reactor is taken as along equator.

B. Spectrum With Oscillations Depends upon Distance

Aside from the other possibilities for directional measurement, there is another handle one may employ in detecting the range of a clandestine reactor. This is the fact that the detectable spectrum of the reactor changes with distance due to neutrino oscillations. One can see this in Figure 8, but it is somewhat better than this since this figure applies to the sum of all the Japanese reactors detected in KamLAND. One can certainly, in principle, employ spectral information along with count rate to infer distances at ranges out to around 1000 km. However, it is a somewhat complicated game, and needs simulations to explore the resolution, work beyond the scope of the present conceptual paper. I believe it has some promise when used in conjunction with the other directional information as already discussed.

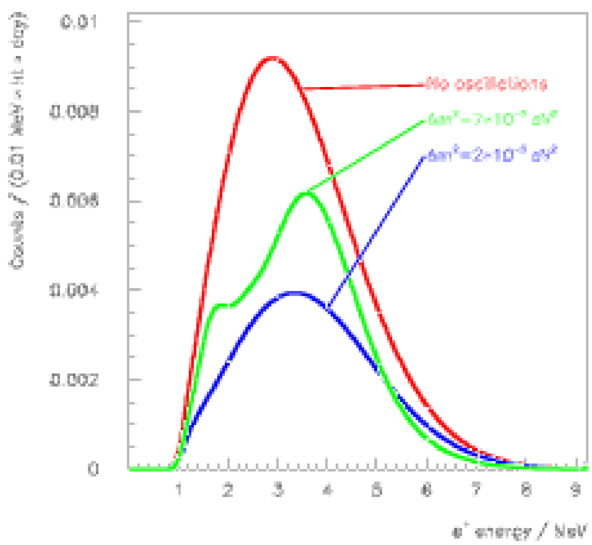

FIG. 8. Positron spectrum at KamLAND for no oscillations and several values of Δm^2. One may also read this also as a measure of the spectral distortion of events from differing ranges: close, roughly 180 km, and 1000 km[3].

In this context, detection of spectral differences due to reactor load (Uranium versus Plutonium) will be very difficult. For ranges of a few hundred kilometers however, with the presently considered km^3 detector, spectral studies may yield useful information.

IV. DESIGN CONCEPT

Using neutrinos to monitor reactors or detect nuclear explosions is not a new idea. What really is new herein is the consideration of the possibility that we can achieve a detector of gigaton mass at a cost which may be affordable in the not too distant future. The keys, as we discuss in the following are the use of water as the medium, with some added material for light-output enhancement, plus, most importantly, progress in light detection technology.

One need not do complex calculations to be convinced that nothing else but water is affordable: a gigaton, or 2 x 10^{12} pounds, of material costing \$1/pound would cost \$2 trillion! However we can certainly employ pure water. Employing seawater would bring unacceptably high rates of radioactive backgrounds from ^{40}K, U and Th. Possibly it would be desirable to load the water with low-^{40}K salt. Salt is not expensive, helps neutron detection, and would help counteract the 3% buoyancy of seawater. As we discuss below, a further additive to increase light output is probably needed.

Similarly one has not much choice about location for such a detector. Deep mines are simply not affordable for such volumes, at least in useable geometries. One might think about, say, lining tunnels of 100 m^2 cross section, but then one would need 10,000 km of tunnel! If one excavated a tunnel of the largest cross section possible from the standpoint of rock stability at reasonable depth, one would have a tunnel of perhaps 50m x 100m x 200 km. Even if this were practical, the cost of excavation would still be in the range of a hundred billion dollars.

The only remaining possibilities are in the deep ocean or deep in an ice sheet as at the South Pole (like the ICECUBE experiment). The use of deep ice is worthy of some more consideration, but it is not obvious how one would afford to install enough light collection and detection to get the low energy threshold resolution needed. Since it is known from the AMANDA studies that the ice has significant scattering (order of 25 m or less effective scattering length), one would have to have light collectors placed densely enough to localize events on that distance scale. One cannot contemplate melting the entire mass of ice to dope and refreeze it for greater light output, nor could one wait for material to diffuse through the volume. In sum, I do not see that it is likely that one can use the deep ice for this application.

Instead, it seems the obvious venue is in the deep ocean for such a huge detector. A fringe benefit of the ocean is that such a detector is potentially mobile and reconfigurable (and repairable, as instruments frozen in the ice are not). Hence from here on I will assume the location to be the deep oceans of the world, more than about 4 km. Certainly studies should be made as to whether lessor depths such as are available in Lake Baikal ($\tilde{1}$.4 km) would suffice. Gratta points out that Lakes Baikal and Victoria ($\tilde{1}$.1 km) have the virtues of being distant from nuclear ships and reactors, and offering geopolitically interesting locations.

A. Photodetector

The photodetector for this project is the single biggest technological/cost hurdle, one that will require significant research and industrial effort. Photomultiplier technology in some ways has not evolved much in the last half century, since invention. It still employs large glass envelopes and requires excellent vacuum. The 50 cm diameter photomultipliers used in Super-Kamiokande and KamLAND are beautiful devices, but are difficult to handle and would need pressure housings for the deep ocean. Large photomultipliers cost about \$1/cm^2 in photocathode area. One can see right away that if we are to cover an area of order 10^7 m^2 with detectors, the cost will be in the neighborhood of \$100 billion. Hence a goal of research leading to such a detector would be to develop new photodetection capability with cost reduction of a factor of 100, or at the very least a factor of ten. One may imagine a factor of ten being realized from automation of the manufacture of present style photomultipliers (perhaps with new hybrid multiplier structures). The present Hamamatsu large PMTs use hand blown glass envelopes and much hand labor in construction. In contrast, large cathode ray tubes for television, manufactured in great numbers, can be bought for prices on the order of a factor of ten less than photomultipliers. So, one may hope for this level of cost savings by creating enough market demand. Still such tubes need to be made to be pressure resistant for the deep ocean, and there are non-trivial associated mechanical costs (support, shielding magnetic field, connectors, etc.).

What we need is a new technology that makes a leap forward from the beloved old glass envelope designs. For example, recent work has been conducted with making electron detector of flexible sheets (micromegas technique [7]). Also, there have been some explorations of organic detection materials [8]. As Gratta has pointed out, an even better solution might be to seek a mechanism for significant light amplification at production, and thus enable the detection to involve imaging optics (think CCD cameras).

For the present exercise (again, admittedly optimistic), I will imagine the possibility for photodetector "wallpaper", material which is photosensitive, pressure tolerant, flexible, and produced in mass quantities for a price of order \$100/m^2. Obviously, significant investment is needed in R&D to realize this leap into twenty first century photodetection, but it seems not impossible. One should note that this technological leap would generate significant repercussions in other areas of science and commerce, perhaps in itself spawning new industries.

Our fallback would be next-generation mass-produced PMTs, but which would then dominate (and perhaps double) the entire project cost.

B. 10 Megaton Modules

Without much justification I will assume that the gigaton detector consists of 100 modules at 10^7 ton each. One large detector would be almost impossible to handle, vulnerable to single point failure, and would in any case require subdividing internally. As usual in this sort of optimization game, costs increase with the number of subdivisions, so we ought to minimize the number of modules, consistent with practicality. There have been many discussions of megaton scale detectors, a jump in volume of a factor of 20 from Super-Kamiokande (to MegaK or HyperK in Japan, UNO in the US, for example). A spherical balloon of 134 m radius will have a volume of 10 megatons of water and would represent a factor of ten jump from proposals under present discussion. There is precedent in oil tankers of this size (lengths to around 350 m, with mass up to a megaton) which negotiate the world's oceans. For present purposes I will assume a flexible bag with pressure tolerant photodetectors and electronics on inner wall. Giant bags have been discussed in the past for oil and water storage and transport.

A metal (steel) tank would be conceivable, built afloat. However, costs become rather large: if we imagine a 5 cm thick tank hull and a total of 100,000 ton structure, the cost could be of order \$100 M. If a flexible bag (incorporating the detector and electronic sheets) is practical, then the bag structure costs may be nearer to \$100/m^2 (industrial conveyor belt material) or a total of around \$20-30M (plus detector and electronics costs).

Anchoring forces do not seem to be a major concern, being less than 30 tons in the typical ocean bottom currents of less than 10 cm/sec. The 3% buoyancy of the ocean is not a trivial problem, since it amounts to 300,000 tons in a 10-megaton module! This might be one motivation to employ a steel tank, if affordable. Otherwise one probably needs to consider adding low potassium (low ^{40}K) salt to the water. In any event the detector is likely to be buoyant, and will need to be hauled down to the ocean bottom, from where it can be brought back for service or redeployment elsewhere.

One needs to fill the detector with pure water with minimal radioactive contamination (compared to seawater), and with good optical properties. One might consider taking the already purified water from the Antarctic glaciers. But, reverse osmosis filtering is now sufficiently developed that one can afford to use filtered fresh water (at a cost of order of a few dollars per m^3).

C. Sensitivity

We need a sensitivity to observe reactions down to about 1 MeV, and particularly good enough to resolve the 2.2 MeV signal of neutron capture on hydrogen. In Super-Kamiokande, utilizing only Cherenkov light with 40% PMT wall coverage, we have a response of about 10 photoelectrons (PE) per MeV. If we want to maximize the surface to volume ratio of the detector we should make it as large as possible. If we take the suggested 10-megaton spherical modules, the diametrical distance of 268 m is somewhat more than twice the absorption length probably achievable for 400 nm light in pure water. Most light rays will, of course, travel less than this distance. We can compensate for this somewhat by covering a large fraction of the wall surface. Let us assume 100% wall coverage with new flexible photodetector material. If this cannot be achieved then we can back off to a smaller volume with more total surface area per array, but less photodetection coverage fraction overall, somewhat compensating.

At the Super-Kamiokande sensitivity level this amounts to a Cherenkov signal of about 22 PE from the neutron capture. Without detailed study, taking into account the noise rates discussed below, it is not clear if this is an adequate signal. It is possible that we can add some material to raise the sensitivity without costing the directionality inherent in the Cherenkov radiation (as was done by the LSND group, using oil with slight doping). An old problem has been that there are no available water-soluble scintillating or readily useful wavelength shifting materials (that I know of).

There seems to be no reason for this situation to persist, as I have learned from talking to several chemists. A gain of a factor of 3 in sensitivity seems possible from wavelngth shifters tapping the very blue light which does not travel far. It may be difficult to do this while reserving water transparency. Scintillating material would have the added virtue of giving some sensitivity to particles below Cherenkov threshold (e.g. kaons and recoil protons). The latter is useful in rejecting backgrounds and in proton decay studies.

Another means of increasing sensitivity would be to add something to shorten the neutron capture time and increase the resulting light output. Adding $GdCl_3$ in solution is an interesting candidate [11]. It has a huge neutron cross section, and would give about 8 MeV instead of 2.2 MeV. Doping at the level of 0.1% at a cost of $3/kg would amount to $30M in a 10 Megaton module. However the total needed is beyond the present world reserves, so this seems not practical. Is anything other than $NaCl$ affordable?

D. Noise Rates

Noise rates in good quality large area PMTs are on the order of 1 count/cm^2/sec. For a 10-megaton balloon this would amount to 2.2 counts/nanosecond. If we have nanosecond resolution then this is equivalent to 220 keV every nanosecond. With a threefold addition in sensitivity as suggested above, this would mean that a neutron capture signal of 88 PE would include about 3% background counts.

This noise rate amounts to about 1/40 the noise rate in the open ocean. So, if we should add salt the net ^{40}K would have to be roughly less than 2.5% of that in typical sweater.

For present purposes, I will also assume satisfactory levels of U, Th, Radon in the filtered water. While this looks to be a manageable issue, it needs study. In any case, if the optical backgrounds should prove to present a problem, then we can make the modules smaller or subdivide them.

E. Module Size Optimization

As one sees in reading the foregoing, we need study of costs as functions of sensitivity, noise rates, resolution, etc. More modules would yield better sensitivity, would have lower noise rates, would present easier mechanical handling, are closer in size to present experience, would presumably offer higher reliability for the functioning of the entire array, but would imply higher total array costs. The present working number (100) is just a guess based upon having gone through this sort of exercise in the past: I am guessing that 10 megatons is about as big can be handled, and that the economics will drive us to as large a module as can be managed.

V. COST SCALE

In the following Table I we present a first guess at costs for a ten-megaton module adequate to detect electron anti-neutrinos from reactors. Most crucially we have assumed significant progress in photodetectors achieving an area cost reduced by a factor of 100. One can see that if that major advance is not achieved, and we fall back to a reasonably justifiable factor of ten economy-of-scale gain in more traditional glass PMTs then the photodetectors would approximately double the module (and array) costs.

TABLE I. 10 Megaton Module Cost Scale

Photodetectors: if 100x improvement achieved	$22M
Bag: $100/m^2	$22M
Electronics: same as PMT	$22M
Water: $0.01/gal	$25M
Anchoring, mechanical	$10M
Calib, control, comm	$10M
Salt, low ^{40}K $0.01/#, 2.5%	$5M
Doping, increase light	$5M
TOTAL	**$121M**

A. Gigaton Array Cost

The cost of the entire array of 100 ten-megaton modules will thus be about $12 billion for 1 km^3 of sensitive detector. One would send power and commands out and return data on an electro-optic cable, similar to those in use by the telecom industry around the world. There are some engineering challenges in installing and cabling the array, but connection and service with robots and submarines will be practical. The cost of the cable (order of $10/m) and shore station will be less than the cost of one module and is ignored for now.

One question for study is whether the array needs to be this large. For example if the subject for monitoring is not at great range (as, say, in monitoring North Korea from the Japan Sea), then smaller arrays may suffice. However, we need at least 3 such arrays around world for tomography to locate new surreptitious site anywhere.

I realize that setting down explicit and not well documented cost figures opens this paper to criticism, even dismissal. But because there is no question that an array of this mass could be constructed, I think that it is necessary to make the argument that it *may* be affordable. I would ask critics to respond with their most cost effective alternative.

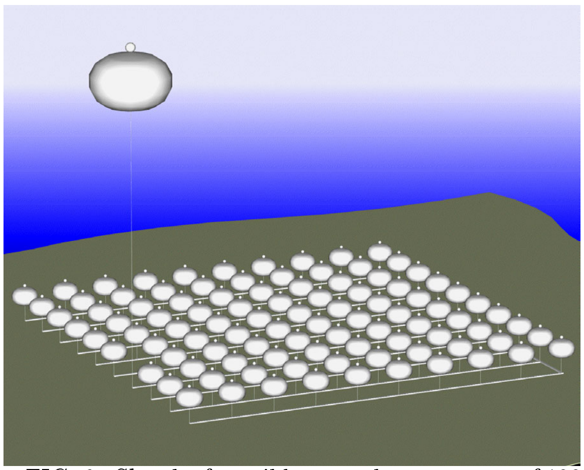

FIG. 9. Sketch of possible ocean bottom array of 100 ten-megaton modules anchored in a square array on the deep ocean bottom, typically 4-5km depths. One module is shown in retrieval/repair configuration. Array would be cabled to shore for power, command and control and data transmission.

B. Array Layout

In the sketch in Figure 9 we indicate an ocean bottom array of dimensions about 4 km by 4 km with 100 of the ten-megaton modules. This allows some space between modules for ease of mooring, but convenient connection into the same cable harness to shore. Probably there would be at least two redundant cables to shore for reliability.

The physics spin-off of such an array is apparent, as such would constitute a tremendous neutrino telescope useful for many studies. Configuring the modules spaced as closely as practical in a horizontal plane would permit use of the array to search for high-energy neutrinos, detecting near horizontal high energy muons crossing the entire array. For low energy neutrino studies the relative locations of the modules makes little difference. In general, for neutrino astronomy purposes, a near equatorial location is preferable.

Alternative configurations may certainly be considered. For example a distributed array (say over a scale of 100 km) would permit measurement of the gradient across the array, at least for strong signals. One might consider a linear geometry or perhaps a triangular geometry. Another approach might be to distribute all the modules throughout the world's oceans. If we had 300 modules total, they might be spread out to gain more position resolution in a rather more complicated tomographic reconstruction. Such a dispersal would enhance the geological study capabilities. Study is needed.

VI. DETECTION OF NUCLEAR EXPLOSIONS

For fifty years people have thought about detecting atomic bombs via neutrinos, and in fact Fred Reines had first been urged by Enrico Fermi to use this means for the initial discovery of neutrinos' existence. The difficulty is, of course, a weak and fleeting signal. The beauty of this method for test-ban-treaty monitoring is that neutrinos cannot be faked, jammed or shielded. And detection measures weapon yield. With the presently discussed km^3 detector we could certainly detect a 100-kiloton (TNT equivalent) device out to a distance of 1000 km with about 2000 counts in a few seconds [9]. Such a signal would give a precise measurement of detonation time, and would confirm and add to seismic measurements, permitting calculation of location and yield with only one neutrino detector. We will not pursue this issue here, as we need studies to see what can be done employing information from three arrays, with small explosions and with detection information combined from other sources.

VII. DETECT NUCLEAR POWERED VESSELS

The array discussed here would certainly also be able to detect nuclear submarines and ships out to substantial range, if they are running at reasonable power levels (when in port they are usually nearly shut down) [10]. A 100 MWt marine reactor would contribute about one sigma to the world total nuebar count rate at 1000 km range in two days, and would thus be marginally detectable. At 100 km range however it would be easily detectable, and if the array is distributed over some distance on the ocean bottom, the signals might be used to roughly track the submarine, but not with great precision, nor in real time. The tracking accuracy should not enough to cause worry to military planners concerned with destabilizing exposure of submarines to attack. On the other hand, such nuclear vessels traveling known routes with known power levels, would provide an excellent calibration signal for the array, and be useful in some precision studies of neutrino oscillations.

VIII. LARGE PROGRAM OF UNPRECEDENTED OTHER PHYSICS STUDIES

Such an array of instruments as we are considering herein would provide a cornucopia of scientific studies. This scientific bounty would have the benefit of attracting a large and active scientific community to become engaged in the design, construction and operation of such a huge and complex instrument. The science from such facilities would occupy a large group of physicists, geologists and astronomers for decades. Some of the obvious topics are:

- Proton decay search to $>10^{36}$ yr, tests all SUSY models.
- Solar neutrino temporal variation to 0.13%$\sqrt{}$(days).
- $\sim$ 1.5 Type II Supernova/month ($\sim$100 counts/few sec, E=10-50 MeV), clear signal from all of Virgo Cluster (no confusion with bomb signals of lower energy) [13].
- Measurement of "relic neutrinos" from all past supernovae.
- Neutrino point source astronomy.... MeV to PeV
- Far detector for neutrino-factory physics, measuring neutrino properties.
- Cosmic ray studies of origins and composition.
- Search for neutrinos from dark matter annihilations (from earth, sun and galaxy).
- Geophysics in study of earth density including the core.
- Search for natural geo-reactors anywhere in the earth.
- Detailed study of earth radioactivity and heat flow.

For a review of many aspects of neutrino astronomy see [12]. Note that the science which can be obtained with this array speaks directly to at least six of the eleven questions in the recent (Turner, 2003) review of the National Academy of Sciences ("Connecting Quarks with the Cosmos"/citeTurner2003).

IX. SUMMARY: A NEW OPTION FOR THREAT MONITORING

We have proposed the study of a huge array of anti-neutrino detectors, an array designed primarily to monitor nuclear reactors anywhere on earth. A 1-km^3 instrument with 1 MeV sensitivity located in the deep ocean is certainly scientifically and technically possible: there are no in-principle problems. Practicality is a question of technology scaling and economics. A first and admittedly optimistic attempt to estimate the costs for such a device indicates that it may be built on scale of \$10-20 billion. Significant optical detector development is needed, plus other technological and engineering studies.

Three such instruments placed around the world would allow monitoring of *ALL* the world's reactors on a daily basis, plus detection and location of new reactors to few tens of kilometers. Such arrays would have some ability also to detect nuclear explosions.

Moreover the construction of such a system would provide a huge pure scientific program enlisting strong scientific community involvement, with high spin-off in science and technology, and the almost certain significant discovery of unexpected phenomena due to exploring so far into new scientific territory.

What next? First I would hope that this paper can serve as a stimulus to others to jump beyond our current thinking about future neutrino detectors. Perhaps more detailed studies will follow, along with serious technological exploration and development. The next detector stage could be in the range of a 10 megaton unit, and then onwards toward a full gigaton neutrino detector.

X. ACKNOWLEDGEMENTS

I want to acknowledge many useful comments from Giorgio Gratta, Gene Guillian and Sandip

Pakvasa about this manuscript. Also Bob Gauldie for asking the right question. But if this is all nonsense the blame is mine alone.

XI. REFERENCES

[1] F. Reines and C.L. Cowan Jr., 1953, Phys. Rev. **92**, 830. F. Reines and C.L. Cowan Jr., 1959, Phys. Rev. **113**, 273.

[2] K. Eguchi, *et al*, KamLAND Collaboration, Phys. Rev. Lett. **90** 021802 (2003); hep-ex/0212021.

[3] C. Bemporad, G. Gratta, and P.Vogel, Rev. Mod. Phys. **74**, 297 (2002); hep-ph/0107277.

[4] A. Bernstein, A. Wang, G. Gratta, T. West, J. Appl. Phys. **91**, 4762 (2002); nucl-ex/0108001.

[5] M. Apollonio, *et al*, CHOOZ Collaboration, Phys. Rev. D**61**, 012001 (2000); hep-ex/9906011.

[6] J.F. Beacom and P. Vogel, Phys. Rev. D**60**, 033007 (1999); astro-ph/9811350.

[7] P.S. Barbeau, *et al*, subm. Nuc. Instr. Meth. (2003); hep-ex/0304013.

[8] P. Peumans, V. Bulovic, S.R. Forrest, Appl. Phys. Lett. **76**, 3855 (2000).

[9] A. Bernstein, T. West, V. Gupta, Science and Security **9**, 235 (2001).

[10] J. Detwiler, G. Gratta, N. Tolich, Y. Uchida, Phys. Rev. Lett. **89**, 191802 (2002); hep-ex/0207001.

[11] J. Beacom and M. Vagins, subm. Phys. Rev. Lett. (2003); hep-ph/0309300.

[12] J.G. Learned, K. Mannheim, Ann. Rev. Nucl. Part. Sci. **50**, 679 (2000).

[13] G.A. Tammann, Proceedings of the 1976 DUMAND Workshop, A. Roberts ed., U. Hawaii, p.137, 1976.

[14] M. Turner, /it et al, "Connecting Quarks with the Cosmos: Eleven Science Questions for the New Century (2003)", http://www.nap.edu/books/0309074061/html.