The Fukushima Accident

Peter Bernard Ladkin
University of Bielefeld CITEC and Causalis Limited

On March 11, 2011 an unusually strong earthquake, now called the Tohoku earthquake, occurred off the north-eastern coast of Japan's main island, Honshu. While the earthquake caused some destruction, the ensuing tsunami about three quarters of an hour later largely destroyed numbers of towns and devasted cities such as Sendai.
http://earthquake.usgs.gov/earthquakes/eqarchives/poster/2011/20110311.php
http://en.wikipedia.org/wiki/2011_Tōhoku_earthquake_and_tsunami

While there is some evidence to think that the earthquake itself caused some damage to critical systems at the Fukushima Daiichi ("Fukushima Number One") nuclear plant, located at the edge of the ocean in Fukushima province, some few hundred kilometers north of Tokyo, and indeed cut the supply of outside electricity, the tsunami flooded the plant, including the basements in which the back-up power generators were situated. There are videos of this happening posted on the WWW. An event in which local power from the turbines as well as both primary external and secondary emergency power is lost is known as a "station blackout":
http://mrzine.monthlyreview.org/2011/lochbaum240311.html There are batteries which supply power for a few hours in the event of a station blackout.

What's The Problem?

There are six reactors at Fukushima Daiichi. At the time of the earthquake, Reactors 1-3 were operating, Reactor 4 was defueled – its fuel was stored in the Spent Fuel Pool inside the Reactor Building along with spent fuel from previous operation – and Reactors 5 and 6 were in "cold shutdown" (see below for terms). Reactors 1-3 shut down automatically immediately upon detecting the strong earthquake.

The reactors here are a GE design, known as a Boiling Water Reactor, BWR, with Mark I containment in Reactors 1-5 and Mark II containment in Reactor 6. In a shut-down of this kind of reactor, the chain reaction which sustains the usual power generation is halted by the insertion of "control rods", made from neutron-absorbing material such as boron or cadmium, directly into the reactor core. The rods are inserted automatically from below. However, current reactor designs require continued and continual cooling as the radioactivity in the core remains. Although it decreases in the usual exponential manner over time, a matter of days for the by-product iodine-131 but many months for the other major by-product caesium-137, the reactor must be actively cooled for years until the radioactivity in the core decreases to a point at which passive cooling (that is, just letting it siphon off its heat through static, passive heat sinks) suffices.

The core, plus the cooling/heat transfer water in a BWR is contained first in a Reactor Pressure vessel (RPV) and associated piping taking the superheated steam into turbines, as shown in Figure 1. (This figure is also used on p12 of the UK Office of Nuclear Regulation Interim Report http://www.hse.gov.uk/nuclear/fukushima/interim-report.pdf .) The primary prophylaxis against a release of radioactive substances into the environment in case of a physical failure is physical containment. The RPV is thus enclosed in a Primary Containment Vessel (PCV); the Reactor Building itself is designed to be the secondary containment structure.

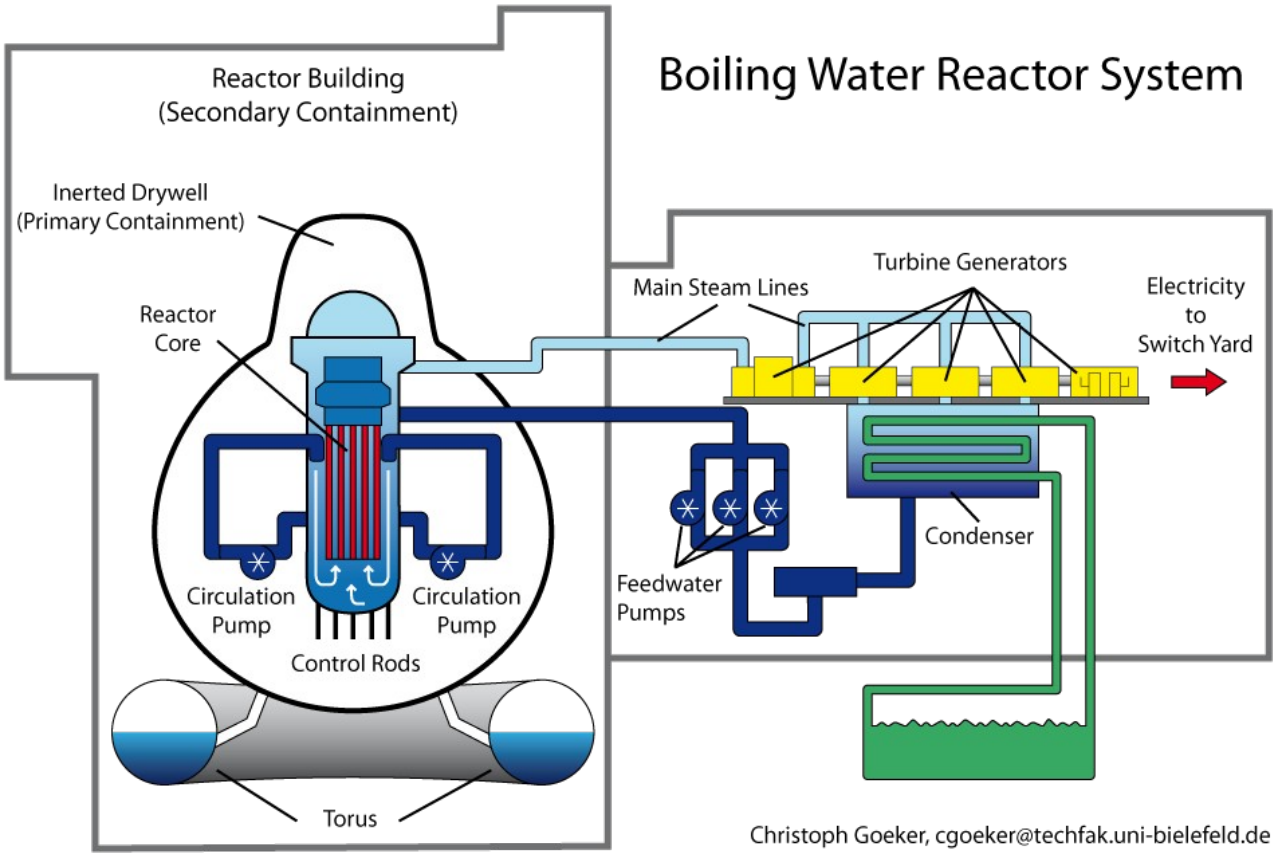

Figure 1: A Schematic Diagram of a BWR with Mark I containment (graphic: C. Goeker)

The water in the RPV is (super)heated by the heat of the core chain reaction, and converts to steam, which is led to pass through the turbines, which generate electricity. After passing through the turbines, the steam is condensed and passed, cooler, back into the RPV. The RPV and some cooling/generation system piping is containing within the concrete-and-steel Primary Containment Vessel (PCV), itself enclosed in the Reactor Building (secondary containment), which also contains the Spent Fuel Pool (see below), as shown in Figure 2 below, also Figure 4, p14 *op. cit.*

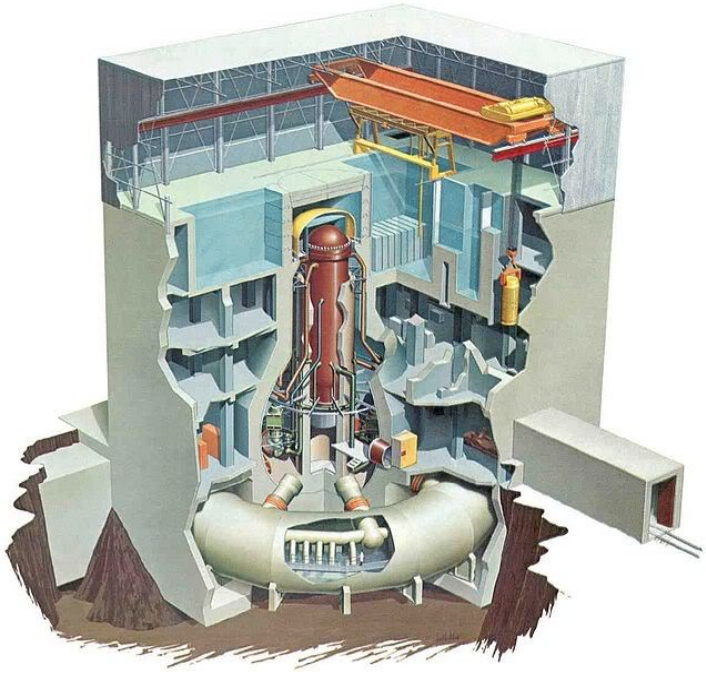

Figure 2: The BWR/Mark I Reactor Building (by kind permission of GE)

The Turbine Building is a separate structure, but the system of RPV, generation/cooling circuit, and turbines is part of one enclosed pressure system in the form of a loop. For more technical details of BWR systems and their various cooling systems, see the US Nuclear Regulatory Commission's Reactor Concepts Manual http://www.nrc.gov/reading-rm/basic-ref/teachers/03.pdf

The phenomenon of persistent active cooling is basic to contemporary reactor design. From being on-line, after a shutdown, the reactor **must** - and this is a "must" without exceptions - be actively cooled for years. This is and has been well-known for half a century. The first reactor at Fukushima Daiichi went on-line forty years before the accident, in 1971, was operating at the time of the tsunami, and had been recently granted a licence to continue operating until 2021. http://en.wikipedia.org/wiki/Fukushima_Daiichi_Nuclear_Power_Plant

The reactor itself is not the only installation in a BWR with Mark I containment which requires cooling. The Spent Fuel Pool (SFP) is located in the upper stories of the reactor building, as shown in Figure 2, under the orange-and-yellow crane used for transferring fuel rods out of the top of the reactor into the pool, and vice versa if refueling. Spent fuel must be actively cooled until the residual radioactivity is low enough that passive cooling – that is, allowing the heat to distribute itself through the immediate surroundings, which then transfer the heat to the atmosphere – suffices. This process takes years. The SFP at Reactor 4 was also being used to store the current fuel; the reactor had been defueled. The fuel in the SFPs requires significant active cooling; more so of course when the current fuel is sitting there as well.

A reactor is said to be in "cold shutdown" when the temperature of the cooling matter is below its atmospheric boiling point, in the case of water 100°C. In cold shutdown, the water in the core cooling is no longer superheated and the cooling system may be opened to the atmosphere, say for some maintenance tasks, without a high-pressure release of suddenly-generated steam. In normal operations, the cooling water is only lightly contaminated, but in the case of, say, a meltdown (see below) it might well be highly contaminated, as indeed it is in some of the units at Fukushima. Reaching cold shutdown normally takes days after shutdown, but the reactor must still be actively cooled, indeed for years after shutdown. The Fukushima Daiichi Reactors 1-3 may achieve cold shutdown by the end of 2011, that is, somewhat over nine months since shutdown http://www.dowjones.de/site/2011/09/japan-trade-min-possible-to-achieve-cold-shutdown-of-fukushima-daiichi-by-year-end.html

All the active systems, not only for active cooling, but also for the control of semi-passive cooling systems such as those driven by steam turbines operating from residual core heat, require a supply of electricity to operate. When there is no local power (the reactors are shut down), no external power (the grid is not supplying), no backup power (the generators are out of commission), one is reliant solely upon the electricity stored in the batteries, which lasts at most hours, whereas cooling must continue for years.

"It is assumed" that a station blackout can be somehow "fixed" within the few hours which the batteries last. This one obviously wasn't and couldn't be.

The result of all this is that the reactor cores in the then-operating Reactors 1-3 have melted down (so far as is known six months after the accident – nobody and nothing is able to get close enough to see). That is, the heat from the continuing radioactivity was enough to convert the cores into liquid, which becomes even hotter because the same amount of radioactivity is now concentrated in a smaller volume, and likely melts – has melted – partially or completely through the concrete-and-steel floor of the Reactor Pressure Vessel and maybe the Primary Containment Vessel.

Measuring Bad and Worse

It used to be assumed by nuclear engineers as well as Hollywood that a “meltdown” was the ultimate catastrophe with nuclear-reactor operations. It was assumed that if a meltdown occurred, **all** the radioactive substance in the core would somehow make it into the environment, rather rapidly. Although we cannot yet say how much will make it into the environment, it is clear the release has not been as rapid as Hollywood imagined and the nuclear-engineering industry feared. The country operating the reactor is required by treaty to tell the International Atomic Energy Agency how much it thinks **has already** made it into the environment. The Japanese regulator NISA said one month after the accident: about a tenth as much as Chernobyl. The estimate has subsequently been doubled, and we are not yet done counting, as recent ocean research by Woods Hole Oceanographic Institute has shown http://green.blogs.nytimes.com/2011/09/28/fukushimas-contamination-produces-some-surprises-at-sea/

In fact, the official estimate of radioactivity released on 12 April 2011 was largely based upon atmospheric releases, due in part to three oxyhydrogen explosions, two of which were captured on video posted on news WWW sites and WWW video sites within hours. There appears to be no such record of the third, in the defueled Reactor Number 4, with which the main concern was keeping the Spent Fuel Pool, containing also the core fuel of the defueled reactor, cooled. Atmospheric release occurred also in escaping steam, still a problem at time of writing, as well as release into quantities of water which are open to the atmosphere, such as that which accumulated in the Turbine and Reactor Buildings of various units, and in trenches surrounding some of those buildings, and which the operator has been attempting to pump into containment and processing structures. Some of this, according to Woods Hole, must still be escaping, because the concentration of radioactivity they have measured in specific areas of ocean water, about 10,000 Becquerels per cubic meter (a Becquerel, Bq, is a count of the number of decay events occurring per second) remained constant from May to July, whereas if leaks had ceased, dispersal would have caused this to reduce considerably over this time (in April, the count was 100,000 Bq per cubic meter. “Background”, from natural causes as well as the result of atmospheric nuclear-bomb tests, is about 1.5 Bq per cubic meter as measured in 2010). Rather than overt releases of water from accumulations, it could be that this radioactivity is being carried through groundwater and thereby into the ocean, say from the melted cores.

The International Atomic Energy Agency (IAEA) has a scale, called the INES, which measures accidents on a scale of 0 to 7. The Handbook is quite detailed, at over 200pp: http://www-ns.iaea.org/tech-areas/emergency/ines.asp . However, roughly speaking the types of accident severity can be classified into loss of redundancy (called “defence in depth” in traditional nuclear-power thinking) which alone may take one only up to INES Level 2; serious mechanical-engineering failure, which may take one up to INES Level 5; and significant environmental release of radiation, which is solely responsible for Levels 6 and 7.

The IAEA says that the responsible government determines the INES Level. The Japanese government, after the oxyhydrogen explosions in Reactor Buildings 1, 3 and 4, and the suspected events in Reactor 2, initially classified each of Reactors 1-4 as INES Level 5, and Reactors 5 and 6 as Level 0. The French disagreed, saying it was at least Level 6. But the protocol is not open to debate: if the Japanese said Level 5, then Level 5 it was.

Presumably, until the estimate of environmentally-released radioactivity was completed (in April), the main proof of accident consisted in the physically-destroyed plant, and this in four cases (Units 1-4) was given the most severe INES classification for such events, which is Level 5. After the estimate was publically issued, the events in Units 1-3 was jointly classified as a single Level-7 event (presumably since contamination was supposed to have issued from all three), Unit 4 remained a Level-5 event and Units 5 and 6 Level 0.

In fact, there is clear evidence of loss of redundancy in cooling Units 5 and 6. The data clearly show alternating cooling and heating of Reactor Pressure Vessels and Spent Fuel Pools, which means that of the four required cooling systems, fewer than four were working (indeed, in early days only one) up until about September 15: an obvious loss of defence in depth. This observation and analysis by Bernd Sieker was presented at the Eleventh Bieleschweig Workshop: http://www.rvs.uni-bielefeld.de/Bieleschweig/eleventh/ ; Sieker's data plots are available at http://nuxi.homeunix.org/Fukushima/Fukushima-Plots.pdf .

This all shows that the protocol used for classification is not ideal. The French were right: indeed, they were optimistic. As was the Japanese organisation that disagreed with NISA's estimates on the original Level-7 classification move. Sieker is right. The people performing the classification and reporting it to IAEA were wrong as they did so; sometimes one suspects deliberately wrong (Units 5 and 6) and sometimes just making mistakes in judgement where others knew better. Whatever the purpose of classification, objectivity is better: being right is even better. The classification should surely aim first to be right, and that entails serious engagement with dissenting technical views, for which there is no place in the current classification protocol.

Just Bad? Or Worse? What Went Right

It has been claimed by industry insiders as well as knowledgeable observers for some time that a meltdown was not necessarily the worst consequence of a lack of cooling. Meltdown allows radioactivity to contaminate ground and groundwater, which is relatively immobile compared with, say, air. A Spent Fuel Pool which overheated, boiled away its cooling water, exposing fuel rods directly to the air, would release radioactivity directly into the air, where it would be tranported much more quickly, and less predictably, allowing far less time to evacuate people in its path, and in principle contaminating far greater areas of planet. Even worse, active fuel, as was stored in the Spent Fuel Pool of Reactor 4 (SFP 4), which remained uncovered and uncooled could actually restart a chain reaction, in which case most of the material – the entire contents of a core - would burn into the atmosphere very quickly – a completely uncontained and uncontrolled reaction of rather greater intensity than the reactor itself when operating. Not good.

This hasn't happened. Due to emergency measures, namely shipping Putzmeister mobile concrete pumps with 62-meter arms , some even larger, and 20+wheels each, from Germany and the US http://www.pmw.de/cps/rde/xchg/pm_online/hs.xsl/9419_ENU_HTML.htm , and pumping 100 tonnes of water (about equal to 100,000 liters) per day into SFP 4 in the first few months, any releases from SFP 4 have been kept to what most observers agree is a minumum.

So this accident is a catastrophe. But it is by no means worst-case.

The Putzmeister process for cooling the SFPs was by any measure a success. Another success was the emergency cooling of the reactors by seawater.

The salts in seawater, when used to cool the RPV of a BWR, are extremely corrosive and can coagulate in and slowly destroy the piping. So this is an emergency measure of limited temporal effectiveness. Not only that, but if indeed subsystem failure is induced through using seawater, there is no apparent way to get in immediately and fix it: the internals of the cooling system are open to the core; the cooling water is highly radioactive if the core has been compromised, and when a leak ensues through mechanical failure of some sort – which becomes almost inevitable if you were to use seawater for too long and the reactor has not reached “cold shut-down” at which it is cooled to below the atmospheric boiling point of the coolant, all that substance converts into steam and is uncontained, and so contaminated that no human can get anywhere near it to fix it. Cold shut-down

takes some days to occur in normal operations, and may be attainable only with considerable difficulty, if at all, in case of a meltdown (as we are seeing at time of writing with Reactors 1-3). This is all well known. During the first few days of the accident, the manager of the plant decided on seawater cooling of the RPVs, and reports say he refused to stop doing so even when commanded by senior authorities (the government, deriving, some say, from the Prime Minister himself) to do so.

Had the seawater cooling ceased within the first days of the accident, there is little doubt that the meltdown would have been far worse than it is, maybe even reaching Hollywood standards. Seawater cooling was finally substituted with freshwater cooling, and nothing had catastrophically broken. By all measures, a success.

Should We Have Known? Should We Have Guessed? Did We Blow It? Or Is It Bad Luck?

The design basis at Fukushima Daiichi for flood resistance is reported to have been a little under a 6-meter wave. The bluff on which the plant stands used to be higher, but was reduced to make construction and cooling-water management easier http://search.japantimes.co.jp/cgi-bin/nn20110712x2.html .

Was the design basis appropriate? Obviously not. Was it reviewed? Yes, as far as we can tell. But when and how and what the conclusions were, the "safety case" if you like, has not been made public, despite the very public consequences of this mistake.

"Mistake", I say? Was it? This was initially reported as a "1,000 year tsunami", with recent geological research having uncovered evidence of a similar event 1,100 years ago or so. So let's see: a thousand years is 10^7 operating hours (ophours); that means based on this reasoning that the expectation of flooding at this plant is 1 in 10^7 ophours. And what is the expectation that flooding leads to a severe accident? Here is industrial and organisational sociologist Charles Perrow in his 2007 book The Next Catastrophe (Princeton University Press):

> [p134] *Emergency power at nuclear plants is provided by diesel generators (which have a long history of failing to start and other problems). Clearly visible in some places, these generators could be taken out with grenades. Or, a hurricane could do the work of the terrorists' dynamite and take out the power, and the storm could easily render the emergency generators inoperative as well.*

He then recounts an occasion, in 1996 at Nine Mile Point, when power and emergency power were out for twenty minutes. And again,

> [p173] *[Nuclear power plants] are vulnerable to natural disasters. There have been emergency shutdowns in the face of hurricanes, for example, though no storms or floods have as yet disabled a plant's external power supply and its backup power generators.*

These plants were known to be vulnerable to flooding taking out backup power. US nuclear power plants, except for Diablo Canyon and San Onofre in California, are not located on the ocean, so the word "tsunami" is not to be expected here, but "hurricane" is, along with "tornado" or "torrential rain". Such events would occur to Perrow's readers. (Lee Clarke has noted that such US plants as Salem and Oyster Creek might also be vulnerable to flooding from oceanic events.) Indeed, the nuclear engineer, trainer, and failure analyst David Lochbaum, the leader of the Nuclear Project at the Union of Concerned Scientists, has written of station blackouts caused by flooding, and the risks, in his weekly "Fission Stories" series at http://allthingsnuclear.org/ . I understand Lochbaum first drew attention to the inadequacy of station-blackout provisions in 1992.

There is a good engineering case to be made, indeed I would say it is pretty clear, that a "1,000-year tsunami" of this magnitude is almost certainly going to cause a station blackout for longer than the batteries will hold, in other words the chances of an severe accident due to loss of cooling are better than even, and in any case certainly way more than one in ten. The conclusion follows that the hazard that a "1,000-year tsunami" event occurs and causes an accident has a likelihood of the order of magnitude of 1 in 10^7 per ophour.

Observe that this is two orders of magnitude less than what is acceptable for certifying a commercial airliner (certification guideline says it is acceptable that a "catastrophic" event from a single cause be shown to occur not more frequently than 1 in 10^9 ophours. A "catastrophic" event is one that "*prevent[s] the continued safe flight and landing of the aircraft*"; see E. Lloyd and W. Tye, Systematic Safety, CAA Publications 1982, p37). That is two orders of magnitude less probable.

And what about severity? For a commercial airplane, loss of everyone on board plus the airplane plus maybe some people on the ground (usually few, even in the case of the airplane landing in a city, for example AA587 in Queens in November 2001, although, as we saw, up to about 3,000 if it flies directly into a very large skyscraper, as on September 11, 2001, which is unlikely unless deliberate). In contrast, estimates of deaths due to the Chernobyl accident range from a few hundred (comparable with a commercial aircraft accident) to tens of thousands and more (see below), to which we could apply discount factors, as economists and epidemiologists do, for not being immediate.

So it looks as if, even with the 1,000-year event, the risk known to have been taken with this power plant is two orders of magnitude more likely that that known to be taken with a commercial airliner. What in safety-engineering reasoning accounts for this enormous difference in conventionally-"acceptable" risk?

But in fact a "1,000-year" tsunami was not needed to cause this accident in this way. See Figure 6, p22 of the Interim Report of 19 May 2011 of the UK Office for Nuclear Regulation on the implications for the UK nuclear industry: http://www.hse.gov.uk/nuclear/fukushima/interim-report.pdf . There are four tsunamis in the last 120 years on the east coast of Japan whose magnitude exceeded the design basis for Fukushima Daiichi, and two of those waves were over 20 meters.

Indeed, the tsunami data was well known. The Washington Post reported an astonishing instance of seismologist Yukinobu Okimura being put down by a TEPCO official for raising the issue of tsunamis at a meeting at the industry regulator in 2009: http://www.washingtonpost.com/world/japanese-nuclear-plants-evaluators-cast-aside-threat-of-tsunami/2011/03/22/AB7Rf2KB_story.html . At the Eleventh Bieleschweig Workshop, Robin Bloomfield presented some evidence that the tsunami threat was being taken seriously by the operator and regulators, and old assumptions and behavior revised: http://www.rvs.uni-bielefeld.de/Bieleschweig/eleventh/BloomfieldB11Slidesv01a.pdf

So, I judge, yes, "we" should have known, and indeed were on our way. Should we have "guessed"? Most certainly, and that back in 1971 at the latest. Did "we" blow it? I would say obviously, yes. Or was it just bad luck? Well, one could say it was "luck", indeed of the bad sort, that led the tsunami to occur when it did, rather than, say X years in the future after we had potentially sorted out the tsunami risk, starting first with the acknowledgement of its possibility. On the other hand, random-seeming events are like that: they just occur at some point. Engineers of safety-critical systems are supposed to take such phenomena into account, and it looks very much, also from newspaper

reports, as though this was not adequately done (see, for example, the translation of an article in the Süddeutsche Zeitung on p43 of my A Fukushima Diary http://www.rvs.uni-bielefeld.de/Bieleschweig/eleventh/LadkinFukushimaDiary.pdf). So, no, it wasn't *just* bad luck. There was some amount of not paying attention.

So what does "paying attention" consist in here?

How One Engineers Safety-Critical Systems

The first step is common to all engineered systems, indeed to many human artefacts.

Step 1. Say what the system is supposed to do, along with maybe some amount of how you think it is going to do it.

A nuclear power plant is intended to generate electricity through using the heat of an atomic chain reaction to move fluids or gases through a turbine, the energy of whose subsequent motion is converted into electricity.

One designates a system as *safety-critical* if one thinks that the system is capable, either in normal operation or in failure behavior, of some behavior which causes harm (harming people or distressing the environment). *Safety* is in some sense the absence of harm; ensuring safety is ensuring the absence of harm. Let us assume there is a general social mandate to ensure safety as far as reasonably possible (this assumption can be justified, but this lies beyond the scope of this note).

Functional safety concerns those aspects of safety that are associated with the intended function, or associated failure behavior, of the system. For example, the strip of metal which fell off an airliner on take-off in Paris, lay on the runway, and according to the investigators punctured the tire of a Concorde on take-off on July 25th, 2000 is a matter for functional safety: the strip had a function, was attached so to perform this function, and the attachment failed. However, if a part is coated with a substance whose application could harm the workers applying it, this is not a matter of functional safety concerning the part as a subsystem of the finished artefact. It nevertheless might well be a functional-safety concern for the factory which applies this coating, for applying the coating is a function of the coating factory. Functional safety concerns the operation of a system in its context, and it is this aspect which I am addressing here.

Step 2. Decide whether the system is capable of harmful behavior.

This may seem obvious. However, the constant recall of children's toys should suggest to us that it is not always obvious if an artefact is capable of harmful behavior. It is clear that a nuclear power plant is capable of harmful behavior.

Step 3. Since the system is capable of some harmful behavior, or capable of inducing, or even not hindering some harmful behavior, list all the harmful behavior of which you believe the system to be capable/to be capable of inducing/not to hinder, taking into account the behavior of its environment which enables this harmful behavior.

The trick here is to be complete.

In the case of nuclear power, one can well believe the list to be more or less complete. Harm consists of :

- (a) the usual kinds of local physical harm connected with the malfunction of large industrial

plant which operates under extremes of temperature and pressure, and subject to chemical reactions amongst the substances used. We know about these mostly from long experience with all kinds of plant
- (b) the harm associated with release of radioactive substances into the environment

(Notice that I have already applied here an intellectual grouping. Instead of listing all the things that can go physically wrong with plant, I use the phrase I wrote in (a). And instead of listing all the things that can go wrong when the environment contains substantial quantities of radioactive substances, I designate rather the initiating event in (b). It is possible to give a complete listing by using such grouping, but it should also be clear that completeness attained in this conceptual manner contrasts with the detail and precision needed for engineering concrete avoidance and mitigation measures. Indeed, I hold this to be one of the great challenges of safety engineering: to attain completeness while enabling the necessary detail and precision.)

As we have seen, the INES from the IAEA recognises (b) as the major factor in severity, the other factors (loss of "defence in depth" and physical malfunction or destruction of plant) come under (a). So it seems as if Step 3 is pretty well covered also in the case of nuclear power plants. The influence of the environment is taken into account in so far as the distribution of any released radioactive material over the earth and into the atmosphere is taken into account.

Step 4. Associate with the harmful behavior from Step 3 all the consequences of that harmful behavior.

We have harm from Step 2, and we have harmful behavior, from Step 3. This step associates the specific harm with the specific harmful behavior.

For nuclear power at this level of generality, the harm resulting from (a) is harm to workers, which involves not only the usual industrial trauma but also potential exposure to radiation. This is traditionally controlled through the use of crude measurement devices (dosimeters) and limits on the maximum exposure to which workers may be subjected over a given period of time. These limits are apparently flexible, as the Fukushima accident showed: they were raised by the regulator for continuing workers at the plant to the level previously reserved for once-only exposure in an emergency situation. It should be clear that they are controversial, as indeed is the harm generated by (b). "Plausible" estimates of deaths due to the Chernobyl accident differ by orders of magnitude. The United Nations Scientific Committee on the Effects of Atomic Radiation (UNSCEAR) reported in 2005 57 direct deaths in the accident and 4,000 additional cancer deaths among the 600,000 people regarded as significantly exposed (from the Wikipedia article **http://en.wikipedia.org/wiki/Chernobyl_disaster#Assessing_the_disaster.27s_effects_on_human_health** which has links to archived original documentation). However, Greenpeace reported in 2006 an estimate of 250,000 cancer cases, of which 100,000 would be fatal **http://www.greenpeace.org/international/en/publications/reports/chernobylhealthreport/** This does not take into account other human health effects, which people believe have been seen as far away as Berlin (see articles by Watts, Pflugbeil and Sperling, as well as effects on other living organisms by Achazi, in *Atomkraft als Risiko*, ed. Lutz Metz, Lars Gerhold, Gerhard de Haan, Peter Lang Verlag, 2010). There are serious estimates of harm even higher than this: see Alexey V. Yablokov; Vassily B. Nesterenko; Alexey V. Nesterenko, Chernobyl: Consequences of the Catastrophe for People and the Environment (Annals of the New York Academy of Sciences, Wiley-Blackwell, 2009.

With an order of magnitude difference in serious estimates of deaths, let alone other illnesses and environmental damage, we may conclude that a consensus on Step 4 is lacking for nuclear accidents. It's not like an airplane or train accident, or apparently even an oil spill.

Step 5. (Hazard Analysis, HazAn): For all the behavior you have identified in Step 3, identify precursor behavior, or states of the system+environment, at which point there is still room for intervention before the harmful behavior inevitably ensues.

If there is a severe thunderstorm in the flight path of your aircraft, you can deviate. Severe thunderstorms often contain weather sufficient to induce loss of control in any aircraft. Once you're in it, luck (mostly) lets you out. But you want to identify the hazard and deviate before that. When the pressure is rising in your pressure vessel, you want to try to contain that rise before it rises to the point at which the vessel ruptures. So uncontrolled undesired rising pressure is a hazard. If two commercial high-performance aircraft come within a few hundred feet of each other, they may not reliably be able to avoid collision, so coming that close, an "airprox", is a hazard. Since a BWR must be continually cooled, even after shutdown, at a known rate, in order to prevent core meltdown, overpressure in heat-transfer mechanisms and breach of containment, station blackout is a hazard.

Step 6. For each hazard, designate the consequences (namely the harmful behavior of which the hazard behavior, or entry into the hazard state, is a precursor).

This makes clear that HazAn is a technical grouping, for the purposes of mitigation and avoidance, of the harmful behavior identified in Step 3. With a hazard is associated a *severity*, which is variously the harm which may result from the hazard.

There is some controversy about how to assess severity. Traditionally, engineering takes the worst-case outcome. Loss of power on all engines in a commercial airliner has "catastrophic" severity, even though British Airways Flight 9 in 1982 landed safely at Jakarta after having had the engines flame out due to a volcanic ash encounter; British Airways Flight 038 crashed at London Heathrow airport in January 2008 after losing significant engine power on short final, with one broken leg and two other minor injuries, but a very broken aircraft; and Air Transat Flight 236 glided to a successful emergency landing in the Azores islands after losing engine power through fuel exhaustion. The classification is appropriate: many or all the people on board might well have died in any of these cases, and their survival has little to do with the event which caused the emergency. But there are also situations in which worst-case classification might not be quite as appropriate. If I fall off a mountain bike on a hill trail, the worst case is that I am run over by the off-road vehicle following right behind. But this is so unlikely, and so easily mitigated (let him by!), that mountain bikers, with reason, usually do not think this way.

Step 7. Design into the system avoidance mechanisms for that behavior. If you can't find any, then design-in some mechanisms which mitigate the consequences identified in Step 4.

Step 8. Demonstrate that the mechanisms you have devised in Step 7 are adequate, to some measure of "adequate".

Step 7 is self-explanatory. Step 8 is usually undertaken by judging "risk", usually taken as some combination of the likelihood of a hazard associated with its (modified) severity, further combined over all hazards. The original definition of risk, and that still used in finance, is due to Abraham de Moivre, in his paper De Mensura Sortis in the Proceedings of the Royal Society in 1711 (see Hald, de Moivre, McClintock, International Statistical Revue 52(3), December 1984, available through JSTOR for a commented reprint): the expected value of loss. However, engineering definitions differ from this (see Ladkin, Causal System Analysis, Chapter 5, 2001 **http://www.rvs.uni-bielefeld.de/publications/books/CausalSystemAnalysis/Chapter_5_Problems_calculating_risk_via_hazard.pdf** , a modified version of Ladkin, Hazards, Risk and Incoherence, RVS-Occ-98-01,

Faculty of Technology, University of Bielefeld, 1998 **http://www.rvs.uni-bielefeld.de/publications/Reports/risk.html**).

This is an idealised process. In fact, all of these steps are iterated many times during the life of a complex project. Steps 2 through 8 are applied recursively to subsystems as they are designed.

Problems With This Process

For a complex system, the 8 steps of this idealised process must be performed by a variegated human organisation, indeed a collection of partly cooperating, partly competing formal organisations, which within themselves are rarely if ever perfect. Very often, parts of these organisations will want to keep certain relevant matters secret, or at least undisclosed to certain other groups, for a variety of reasons.

One could imagine a regulatory regime for certain safety-critical systems defining specific documentary products to be generated for each step of this idealised process, which products would enable a rational reconstruction of an argument that a system is adequately safe (whatever "adequately" is taken to mean).

A question then arises as to who shall see and analyse the rational reconstruction, and who shall assess that it is correct. Generally speaking, *quis custodiet ipsos custodes*?

Application to Fukushima

Although there is considerable ongoing discussion about what we have seen during the unfolding of the Fukushima accident, there are some general, generally agreed, conclusions one can draw.

First, quite apart from the builder in (originally) 1971, it seems that the operator TEPCO did not accurately estimate the risk from tsunamis to its Fukushima Daiichi plant and the regulator, NISA did not correct this misestimate, for whatever reason. That seems to be a failure of Steps 6, 7, 8: the consequences (station blackout and thereby loss of cooling after batteries were exhausted) do not seem to have been appropriately assessed (Step 6); the avoidance (a seawall capable of withstanding a tsunami of, I understand, less than 6 meters, less than some which have occurred on the Japanese east coast in the last 120 years) was obviously inadequate; and the oversight (discussion of the reasoning behind any claim that the tsunami risk had been adequately considered) also inadequate.

There was and is clear room for engineering improvement.

The first question, then, is whether other crucial safety aspects of the design and operation of these plants are in a similar condition of having been inadequately assessed, handled and shown to be so.

The second question is whether a different form of human organisation around nuclear power will circumvent these inadequacies. This is in general the *quis custodiet* problem.

And the third question is whether such a human organisation, if it exists, is implementable by the polity.

General Application

This isn't just about Japan. These three questions arise for all societies and political cultures.

Japan is one of the most advanced societies in the world in terms of engineering and its applied

benefits to society, including arguably the most well-functioning public transportation systems in the world as well as some of the most sophisticated city living (which depends heavily on engineering). And, until recently, one of the apparently best-functioning nuclear power-supply systems for electricity.

Some engineers and politicians have been able to argue, after Chernobyl, that "something like that couldn't happen here" because of the obvious lack of control, both intellectual and physical, over the artificial situation induced by the reactor operators. There is no such question arising with Fukushima. Robin Bloomfield suggested a more appropriate general stance in his presentation at the 11th Bieleschweig Workshop: "We don't have tsunamis like that, but what is it that *we* have missed in our operations?" http://www.rvs.uni-bielefeld.de/Bieleschweig/eleventh/BloomfieldB11Slidesv01a.pdf . The main lesson to be learned by many from the Fukushima accident is that "something similar could happen here!"

In the aftermath of the accident, I argued that a public safety case for all safety-critical engineered systems would be a positive step (equivalent in the terms of this note to making Step 8 public): http://www.abnormaldistribution.org/2011/03/27/fukushima-the-tsunami-hazard-and-engineering-practice/ . Martyn Thomas and Nancy Leveson pointed out privately that there are significant issues around retention of intellectual property by the companies involved: a full safety case involves details of the design, and these designs are how companies succeed in being better than others and they – legitimately – are reluctant to reveal them. Besides, arguments can be blocked socially or by the polity for other reasons, as I considered in http://www.abnormaldistribution.org/2011/04/14/the-epidemiology-of-memes-and-its-effect-upon-safety/ As I said there, I don't know how to ensure that such arguments cannot be blocked in general.

The safety case for other safety-critical systems is not generally publically-available, for example for commercial aircraft or for oil-drilling operations. Why should it be for nuclear power? The answer, for me, is that total costs of nuclear power accidents are an order of magnitude more expensive than the total costs of accidents in these other sectors (see Slide 10 of my Bieleschweig talk http://www.rvs.uni-bielefeld.de/Bieleschweig/eleventh/Ladkin11Bieleschweig.pdf for some crude estimates), even without taking into account the current externalities (Slide 11 of the same set). The severity of significant nuclear accidents is greater than in most if not all other engineering endeavors. (Here one might think of the Bhopal accident, and comparisons may be instructive. More people were killed directly that in either the Chernobyl or Fukushima accidents. It seems also that dangerous substances involved are persisting in the environment – see https://secure.wikimedia.org/wikipedia/en/wiki/Bhopal_disaster#Ongoing_contamination . But there are likely different factors at play here than purely technical ones. My thanks here to Martyn Thomas for the reference.)

Also, certain fundamental issues with nuclear power plants have been solved by no society yet: what to do with the waste. There are ways to handle existing waste: the question is whether they, given the vagaries of human operation, suffice. Germany may be thought to have a relatively good regime, but even here there have been questionable practices and decisions concerning some repositories, such as at Asse: http://www.sueddeutsche.de/politik/schwere-vorwuerfe-es-gab-nie-ein-sicheres-endlager-asse-1.691321 (report in German). The question here is whether anything can reasonably be left to sit for 10,000 years if there are ongoing maintenance problems on a time scale of decades, three orders of magnitude smaller. The United States has publically suffered a lack of what is termed a "long-term solution". But even the German regime only works (if one believes it does work) because the amount of waste is limited: the issue is "solved" only because of Germany's planned exit from nuclear power generation. The "solutions" issue is in many societies overtly political, and seems likely to remain so. The issue of waste disposal needs to be part of the public discussion, and a public "safety case" would be one way of ensuring it is so.

The *quis custodiet* question of the adequacy of the safety assessment and reasoning must somehow be resolved. It seems to me likely, given the aftermath of significant nuclear accidents, that, as with the issue of nuclear waste products, the polity will somehow be involved. How? How does technical engineering and the polity interact to form policy? A fundamental political problem here is that safety involves the *absence* of something, namely harm. In a politically relevant sense, everything was somehow OK with Japanese nuclear power production up until 10 March 2011, and then it suddenly, on March 11 and subsequently, was not. The polity does not seem to me to cope with such point-in-time phenomena in the way in which safety engineers might wish it would.

For another example of point phenomena and how they are dealt with in society, consider driving behavior and severe road-traffic accidents. A deadly road-traffic accident might be and is often socially localised to those involved and their friends and relations, whose lives have been irreversibly changed. An accident such as Fukushima has such consequences for proportionately very many more people.

The social phenomenon associated with road-accidents, many small groups of victims and associated sufferers, each group largely socially isolated from each other group, offers no guide to dealing with a point-of-time event with potentially hundreds to thousands to tens of thousands of victims and associated sufferers.

(For a similar contrast in social sensitivity, consider STEC illnesses, deriving from some strains of the common animal-gut bacterium E. coli. Individual occurrences happen all the time, in their low-thousands in the EU and a hundred thousand in the US each year, and we hear little of them. However, cases which appear to be causally related to each other engender a very different set of social reactions. See my note http://www.cs.york.ac.uk/hise/safety-critical-archive/2011/0631.html for a little more detail.)

Finally, and maybe most importantly, who believes that Fukushima is the last significant nuclear accident that the world will ever see? After four thousand years of engineering, ships still sink. After a hundred and seventy years of engineering, trains still crash into each other. After a hundred years of engineering, commercial airplanes still crash into land and sea and, after seventy-five years of ATC, occasionally into each other. So where is the next big one going to happen? How? How many people, how many even of us, would have been that interested had I asked this question on March 10, 2011?

Coda

Charles Perrow pointed out that there are some significant issues not addressed above. Indeed not. Here are some.

- There appears to be some evidence that at least one cooling system lost its integrity as a direct consequence of the quake, before the tsunami hit. For example, http://www.independent.co.uk/news/world/asia/the-explosive-truth-behind-fukushimas-meltdown-2338819.html . I am not aware of any reliable assessment of the anticipated severity of the accident even if the flooding had not taken out secondary power.

- Perrow has drawn attention to "normal accidents", accidents that happen to tightly-coupled, interactively complex systems during normal operations, and has developed significant theory about them (Perrow, Normal Accidents, Basic Books 1984, reprinted and extended Princeton University Press, 1999). Perrow observes "*[t]his failure .... was not a "normal accident," [in Perrow's sense] the unexpected interaction of multiple failures; it was a*

design accident, the failure to design in safeguards for extreme events that are rare but possible." (private communication, 14 March 2011). Issues of normal accidents in nuclear power plants are very important, and I haven't addressed any of them here – as Perrow says, Fukushima was not one.

- I haven't discussed claims that nuclear power is "clean", and how to adjust those claims to account for the accidents, which tend to be "dirty". But unless one believes there will never be another nuclear accident, ever, such claims must be adjusted.

- I haven't discussed the trope of "*We're OK, we do nuclear power properly, but those people over the border, they don't pay so much attention, they don't do it right, they are the ones to watch.*" This trope is moderately pervasive. If normal accident theory is true, as it may well be, it is disingenuous to use this trope to deflect attention from one's own safety efforts, for no one is immune from accidents if normal accident theory is true.

- I have argued that an appropriate engagement of the polity with engineered-system safety involves a certain transparency – of information, about safety measures, their strengths, weaknesses, their failings; as well as of intent and planning. Transparency is also an issue in the response to accidents and emergency planning for their mitigation. It has been questioned by some as to whether the necessary transparency was evident in the operator's and regulator's response to the Fukushima accident. This particular issue, while important, I take to be outside the scope of this note.

Acknowledgements

I acknowledge the substantial contribution of the presenters at the 11th Bieleschweig Workshop on the Fukushima accident, in particular Robin Bloomfield, Lee Clarke, John Downer, John Knight, Nancy Leveson, Charles Perrow, Martyn Thomas, Axel Schneider and Bernd Sieker, and the distinguished discussants, to the views discussed here. The workshop arose through the ProcEng mailing list we set up in Bielefeld for discussion of this accident; thanks to Jan Sanders for installation and maintenance. I have benefitted from insightful comments of Martyn Thomas, Charles Perrow, Bernd Sieker, Lee Clarke and Jan Sanders on previous versions of this note.

Disclosure

I am committed to open provision of information, discussion, and scientific results and reasoning concerning the safety of engineered systems. A version of this article will appear in the Proceedings of the 20th Safety-Critical Systems Symposium, Springer-Verlag, London, to be published on February 8, 2012, and will be available at www.springerlink.com . The publisher has kindly agreed for me to publish this version, with appropriate acknowledgement, on my professional WWW sites. I am very grateful for the publisher's support of my commitment to open provision.