

ThorCon™

the Do-able Molten Salt Reactor

Executive Summary

Jack Devanney

With major contributions from Lars Jorgensen, Jim Livingston,
Ralph Moir, A.C. Rodenburg, and Chris Uhlik.

Martingale, Inc

Sisyphus Beach

Tavernier, Florida

Version: Version: 1.08

Date: 2014-10-24

Copyright © 2011, 2012, 2013, 2014 Martingale Inc

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Contents

1 ThorCon Design Philosophy	5
1.1 The Goal: clean, dependable energy cheaper than coal	5
1.2 Producability	5
1.3 Reliability	6
1.4 Availability	6
1.5 Three Examples	9
1.5.1 Wigner and Hanford	9
1.5.2 Rickover and the Nautilus	10
1.5.3 Camp Century	10
1.6 The Big Caveat	11
2 The Pros and Cons of Liquid Fuel	12
2.1 History	12
2.2 The Molten Salt Reactor Experiment	13
2.3 The Pros	14
2.4 The Cons	16
3 The Technology	18
3.1 Preamble	18
3.2 The Can	18
3.3 The Hammock Suspension System	18
3.4 The Silo	19
3.5 The Power Module	19
3.6 The Silo Hall	19
3.7 The Recycling System	26
3.8 Passive Decay Heat Cooling	27
3.9 Passive Shutdown	33
3.10 The Silo Cavity	37
3.11 Steam Generating Cell	40
3.12 The Pot	41
3.13 Offgas system	49
3.14 Manufacturing and Erection	51
4 Plant Economics	54
4.1 Should-cost versus did-cost	54
4.2 Resource Requirements	55
4.3 Overnight Cost	59
4.4 Can Cost	61
4.5 Fuel Costs	63
4.6 Salt Cost	65
4.7 Staff cost	66
4.8 Total Cost	66
4.9 Market	68

Abstract

ThorCon is a simple molten salt reactor. Unlike all current reactors, the fuel is in liquid form. It can be moved around with a pump, and passively drained in the event of a casualty. Unlike nearly all current reactors, ThorCon operates at near-ambient pressure. ThorCon requires no new technology. A full-scale prototype can be operating within four years. The system, built entirely in blocks on a shipyard-like assembly line, can produce reliable, carbon free electricity at between 3 and 5 cents per kilowatt-hour depending on scale.

ThorCon employs a moderately high energy density resulting in a short (4 year) moderator life. A ThorCon plant is made up of one or more 250 MWe modules. Each module consists of two sealed *Cans*. Each Can houses a 250 MWe primary loop including a Pot (reactor), pump, and primary heat exchanger. The two Cans are duplexed. At any time, one Can is operating and the other is in cool-down or stand-by mode. The plant is designed so that the change out of a cooled-down Can is safe and quick.

The nuclear portion of the ThorCon is underground. The Cans are contained in silos fitted with water-cooled membrane walls. The Cans are passively cooled by radiation to the silo membrane wall and natural circulation. They operate at less than 300C. This allows a fuse valve rather than a freeze valve to be used to passively drain the primary loop in the event of core over-heating.

Each Can is fitted with its own fuelsalt drain tank also passively cooled by radiation to the silo membrane wall. There are no cooling penetrations into either the Can or the drain tank. The Cans hang from the top of the silo from which they are seismically isolated by elastomeric bearings.

The various components in the Can are hung from the Can lid via a hammock suspension system. This arrangement handles some otherwise very difficult thermal expansion problems within the tight confines of the Can, and allows the Can to be disassembled entirely from above.

ThorCon is designed so that once a Can has cooled down, it can simply be lifted from its silo, transferred to a special purpose vessel, and shipped to a specialized recycle facility. The difficult problems of disassembly and decontamination are shifted from the plant to this facility.

ThorCon uses an 8 year fuel-salt processing cycle. The used salt is allowed to cool-down in the non-operating Can for four years, eliminating the need for a separate, vulnerable spent fuel storage facility. The salt is then shipped to a specialized processing facility. All fuel waste handling problems are shifted from the plant to this facility. This processing will return almost all the uranium and eventually most of the plutonium in the salt back to the system, where most of these actinides will be fissioned.

ThorCon uses a simple sprayer system, already demonstrated in the MSRE, for removing volatile fission products to the offgas system. The ThorCon offgas system

1. recovers most of the energy in the offgas stream improving overall efficiency by about 1 percent,
2. reliably removes entrained salt,
3. confines most of the offgas radiation within the Can.

ThorCon requires no lithium-7. The baseline ThorCon uses NaF-BeF₂ (nabe) for both fuelsalt and secondary salt. Nabe is available and reasonably cheap. ThorCon is designed to accommodate the fact that nabe is less attractive neutronically than lithium based salts. Ground breaking neutronics calculations done by the Pacific Northwest National Laboratory indicate that ThorCon can operate with a U-235 enrichment less than 5%.

The baseline ThorCon employs a tertiary loop using solar salt to remove essentially all tritium. All the loops in the ThorCon are designed to produce natural circulation from decay heat in the event of loss of power. Both the secondary and tertiary loops are located in a steam generating cell which is designed to contain a massive rupture anywhere within these loops.

ThorCon is designed to be coupled to a super-critical steam cycle such as those being employed by

many coal-fired plants around the world. The first ThorCons can simply replace the coal fired boilers and their pollution in these plants.

ThorCon requires no new technology. There is no technical reason why a full-scale 250 MWe prototype cannot be operating within four years. The intention is to subject this prototype to all the casualties that the designers claim the plant can handle. This is the commercial aircraft model, not the Nuclear Regulatory Commission model. As soon as the prototype passes these tests, full-scale production can begin.

The entire ThorCon plant including the building is manufactured in blocks on a shipyard-like assembly line. These 150 to 500 ton, fully outfitted, pre-tested blocks are barged to the site. A 1 GWe ThorCon will require less than 170 blocks. Site work is limited to erecting the blocks. This produces order of magnitude improvements in productivity, quality control, and build time. ThorCon is much more than a power plant; it is a system for building power plants.

No complex repairs are attempted on site. Everything in the nuclear plant except the building itself is replaceable with little or no interruption in power output. Rather than attempt to build components that last 40 or more years in an extremely harsh environment with nil maintenance, ThorCon is designed to have key parts regularly replaced. Every four years the entire primary loop and its Can are changed out. The Cans are returned to a centralized recycling facility, decontaminated, disassembled, inspected, and refurbished. Incipient problems are caught before they can turn into casualties. Major upgrades can be introduced without significantly disrupting power generation. Such renewable plants can operate indefinitely; but, if a ThorCon is decommissioned, the process is little more than pulling out but not replacing all the replacable parts.

Assuming rational regulation, ThorCon can produce reliable, carbon free, electricity at between 3 and 5 cents per kWh depending on scale.

Acknowledgements

ThorCon builds on the immense amount of extremely high quality work that was done at Oak Ridge under the molten salt program. Without this effort, ThorCon would be unthinkable. ThorCon also would never have happened if Kirk Sorensen had not had the initiative to save, publicize this work, and make it readily available to all.

I have taken ideas from many other sources, some of which I am undoubtedly unaware of. But at least a couple of these borrowings must be credited.

1. The hexagonal core log configuration is copied from the Ebasco report. I would never have come up with this ingenious design in a thousand years.
2. I am indebted to many people at PNNL, especially Jim Livingston and Chad Painter, for the MCNP model and neutronics calculations. Jim's creation of an MCNP model of the rather complicated ThorCon core and silo in a very brief period is an amazing accomplishment. They are the people that proved that nabe could work.
3. The silo hall cross-section is inspired by a design by Ralph Moir and his mentor Edward Teller.
4. Chris Uhlik and Ralph Moir introduced the idea of off-site can recycling, not to mentioned many other improvements.
5. Chris also is the creator of the annular Serpent model of ThorCon, a clever idealization which allows us to do all sorts of analyses very efficiently.
6. The open standpipe pressure relief was suggested by A. C. Rodenburg. He is also the prime mover behind our choice of materials along with a host of other improvements,
7. Lars Jorgensen made substantial improvements to the fuelsalt processing system. ThorCon's offgas system is essentially a Jorgensen/Rodenburg product.
8. Per Peterson made a number of important suggestions including twisted tube heat exchangers and fluid diode pumps.
9. Dane Wilson has helped immensely on the materials side.
10. Dr. Ritsuo Yoshioka generously gave us his *ke5* code upon which the ThorCon point kinetics model is based.
11. Dick Engel, the major author of ORNL-TM-7207, the original inspiration for ThorCon, has made so many sage contributions, based in part on his actual experience with the MSRE that I've lost count.
12. For our burn up results, we are in deep debt to Jaakko Leppanen and Manuele Aufiero. Jaakko is the principle author of the remarkable Serpent package, which Manuele cleverly extended to liquid fuel reactors. Manuele took a lot of time from his day job to guide us through the idiosyncracies of Serpent.

I need to thank Ed Blandford, Lindsay Dempsey, Dave Devanney, Dick Engel, Robert Hargraves, Lars Jorgensen, Larry Kelleher, Joe Lassiter, David LeBlanc, Ralph Moir, Josh Nevin, Per Peterson, A. C. Rodenburg, Darryl Siemer, Art Stall, and Chris Uhlik for reviewing and commenting on ThorCon. Review does not constitute approval. None of the above have warranted this design nor have they been asked to. Errors of course remain my own.

Most importantly, I am grateful to Molly Devanney for putting up with a clearly deranged and obsessed husband.

1 ThorCon Design Philosophy

1.1 The Goal: clean, dependable energy cheaper than coal

Currently mankind consumes electricity at a rate of about 2,500 GWe. This will probably go to around 3,750 GWe by 2030. Over the next 20 years, the world will need roughly one hundred 1 GWe plants per year, about two plants per week, and this build rate needs to continue for the rest of the century. As things stand now, most of these plants will be coal fired. According to the MIT Technology Review, as of June, 2013, 1,199 coal plants are planned worldwide, with a nameplate capacity of 1,401 GWe.¹

Each one of these plants will require about 4 million tons of coal per year. Each one will produce between 400,000 and a million tons of ash per year. Each one will produce about 10 million tons of CO₂ per year. Each one will kill at least 9 miners per year (European numbers). Each one will shorten the lives of at least 300 people per year (European numbers) via pollution.² In aggregate, these 1200 new coal plants will require 6 billion tons of coal annually, produce 12 billion tons per year of CO₂, and kill or shorten the lives of at least 400,000 people per year.

Martingale has developed a molten salt reactor called ThorCon™ to prevent this from happening. But coal is more a symptom than a cause. ***If mankind is to prosper, then affordable, dependable power must be made available to all.*** If ThorCon can beat coal where coal is cheapest, it can provide cheap, clean, reliable electricity everywhere, including island states and other areas where coal is expensive or not available. If ThorCon can beat coal, ThorCon can beat energy poverty. But this will require:

Producability ThorCon must be cheaper than coal with zero CO₂ cost. The nuclear island must cost less than a coal fired boiler and its coal handling equipment.

Reliability ThorCon must be at least as reliable as coal. Even major failures must have modest impact on plant output.

Availability ThorCon must be available quickly and be capable of ramping up production very rapidly. The requirement is a ***full scale prototype operating within four years***

Each of these requirements has important implications for the ThorCon design.

1.2 Producability

As we have seen, we will need roughly one hundred new 1 GW plants per year for the foreseeable future. These are aircraft numbers. Boeing 747 production averaged 31 airplanes per year, 1966-2012. We need a production system, not individual fortresses. The system must encompass the entire plant, not just the reactor.

Such a system exists. Figure 1 is a picture of one the four Ultra Large Crude Carriers my company built in Korea in 2000 to 2002. These ships, Figure 1, are capable of carrying 440,000 tons of oil at 16 knots for 25,000 miles. They are the largest double hull tankers ever built. Each one of these ships required 67,000 tons of steel — much of it curved and some of it double curved — but only 700,000 man-hours of direct labor. Over half of this is in outfitting. A good yard will require about 5 man-hours to cut, weld, coat, and erect a ton of hull steel. The build time was less than a year. The cost was 89 million dollars.

¹Technology Review, Vol 116, No.3 June ,2013

² Markandya, A. and Wilkinson, P. Energy Generation and Health, Lancet, Vol 370, Sep 2007.

It is difficult to get a sense of scale from Figure 1. But as Figure 2 shows, the nuclear island of a 1 GWe ThorCon would fit inside three of the center tanks of this ship. The ThorCon requires one-fourth as much steel, and the structure is far simpler. The lesson is simple: we must bring shipyard-like productivity to nuclear power.

The shipyards achieve their remarkable productivity by block construction. Sub-assemblies are produced on a panel line, combined into assemblies, and then into fully coated blocks with piping, wiring, HVAC (and scaffolding if required) pre-installed. In the last step, the blocks, weighing as much as 600 tons, are dropped into place in a building dock.

Block construction not only creates order of magnitude improvements in productivity, but it also produces striking improvements in quality. Very tight dimensional control is automatically enforced. Extensive inspection and testing at the sub-assembly, assembly, and block levels is an essential part of the yard's productivity. Inspection at these levels is easy. Defects and faults are caught early and can be corrected far more easily than after erection. In most cases, they will have no impact on the overall project schedule. A decent tanker, operating in a very severe environment, will have an availability of over 95%.

The ThorCon uses exactly the same production process except the blocks are barged to the site and dropped into place.

1.3 Reliability

One of the least attractive features of current nuclear power plants is that they are extremely difficult and extremely expensive to repair. There are three reasons for this:

1. The difficulties of handling highly radioactive solid fuel elements.
2. The slow decay in reactivity after a shutdown. In a nuclear plant, if something breaks, you can't just go in and fix it.
3. Entombing the most critical components in a colossal, reinforced concrete mausoleum. This is required to handle a rupture in the very high pressure primary loop.

This has forced the designers and regulators to pretend things are going to last 30 or more years under some of the harshest conditions imaginable with nil maintenance. This in turn has created a stifling paperwork and certification system to attempt to meet this unrealistic goal. There is no limit to the amount of money you can spend attempting to do the impossible.

A successful design must avoid this dilemma. ***Everything but the building must be quickly replacable with modest impact on plant output.***

1.4 Availability

Many would say the most unrealistic of our three requirements is a full scale prototype operating within four years. Most would say "it can't be done." Indeed it can only be done if we adopt some stringent design constraints which can be summed up in our mantra: ***No New Technology***. ThorCon requires sound engineering, but nothing we don't already know how to do. ThorCon is a scale up of the remarkably successful Molten Salt Reactor Experiment (MSRE) described in Section 2.2. The MSRE is ThorCon's pilot plant. But in scaling up, we must be very disciplined in enforcing the No New Technology (NNT) rule.

In particular, the NNT rule implies

No Breeder There are two possible routes to a thorium breeder: two salt, and one salt. The two salt concept requires a barrier material that does not exist. The one salt design requires nearly



Length Overall(m)	380.00
Beam(m)	68.00
Depth(m)	34.00
Keel to Mast	73.96
Lightweight(mt)	67,591
Cargo Cubic(m3)	511,000
Ballast Cubic(m3)	149,500
Bunkers Cubic(m3)	12,900
Coated Area(m2)	350,000
Main Power(kW)	37,000
Prop diam(mm)	10,500
Steering gear(t-m)	870
Gen Power(kW)	3 x 1450
Boilers(kg/h)	2 x 45,000
Cargo pumps(m3/h)	3 x 5000
Ballst pumps(m3/h)	2 x 5000
IGS sys(m3/h)	18,750
Cranes(SWL)	2 x 20
Anchors(mt)	2 x 22
Windless	2 x 76
Winches	24 x 30
Firefight	CO2/foam
Accommodation	50
Lifeboats	2 x 50
Contract	1999-12-17
Keel-laying	2001-06-11
Delivery	2002-03-07
Essentially one off	
Cost	\$89,000,000

Figure 1: 440,000 ton tanker built in 1 year for \$89 million




Figure 2: 1 GWe ThorCon silo hall fits into 3 center tanks of ULCC

continuous, complex chemical processing of a very hot, extremely radioactive fuelsalt. This process has not yet been fully demonstrated even at laboratory scale. Both concepts need highly enriched lithium which doesn't exist in anything like the quantities required. The first generation ThorCon is a simple converter.

No flibe The salt that most molten salt designs are counting on is flibe: a mixture of very highly enriched ${}^7\text{Li}$, beryllium and fluorine. Unfortunately, flibe does not exist in anything like the quantities we need. The only proven process for creating highly enriched ${}^7\text{Li}$ required one-third of the USA's mercury stocks and was an environmental disaster. No one knows when flibe will be available in the quantities we need at a price we can afford. Until then, ThorCon must use a salt that is currently available and whose cost is known.

No Brayton cycle Many new reactor designs are based on using a closed Brayton (gas) cycle to turn heat into power. Unfortunately, closed loop Brayton turbines exist only at laboratory scale where they are encountering a series of vexing problems. Even open air Brayton turbines are developmental in the molten salt reactor context. ThorCon can feed its heat to the same standard super-critical steam cycle used by coal plants around the world for decades. The result is nil power loop development risk and ThorCons can simply replace current coal plant boilers and their pollution.³

No further scale up ThorCon is based on 250 MWe modules. The prototype will simply be one of these modules.

As soon as the prototype has proven itself, we can go straight to full scale production.

If we accept these constraints, we can be cutting steel for ThorCon in 2015.

1.5 Three Examples

There are those who think there is something fundamentally different about nuclear that mandates decade long project times. Here are three counter-examples, projects which faced far more difficult problems than ThorCon does.

1.5.1 Wigner and Hanford

Table 1 summarizes the chronology of the Hanford project, the plant that produced the plutonium for the first nuclear bomb. In April, 1942, Eugene Wigner arrived in Chicago and set out to design this plant. At the time, no one had even demonstrated that a chain reaction was possible. Little was known about nuclear cross sections or just about anything else. Wigner, on instinct, decided to use water as the coolant. He went straight to 250MWt when no zero MWt plant existed. In five months, his five man team using adding machines and slide rules completed the design. Only toward the end of that period was the first chain reaction demonstrated using a pile of graphite blocks in a squash court. In February 1943, Wigner convinced the Army to go with his water cooled plant and then scaled the design up another factor of two. About the same time Hanford, WA was chosen as the location, the locals evicted, and in August of 1943 construction started. A year later the plant went critical and in October, 1944 started producing plutonium. In 2.5 years, Wigner went from literally zero to 500 MWt. Wigner was furious that it took this long, blaming "too much money".

³ ThorCon can feed heat to a Brayton cycle more easily than to a steam cycle. As soon as the Brayton cycle becomes competitive with steam, ThorCon will switch over. But right now no one knows when that will be.

Table 1: Hanford chronology

1942-04	Wigner arrives in Chicago.
1942-07	Rough sketches of a water cooled pile. Told he's crazy. No way it can go critical. Helium is only way.
1942-09	Wigner group starts design of 250 MWt water-cooled pile.
1942-12	First sustained chain reaction ever.
1943-01	Wigner completes design. Group of 5 using adding machines and slide rules.
1943-01	Decision is made to locate at Hanford.
1943-02	Decision is finally made to go water. 500 MWt.
1943-08	Construction starts
1944-09	Pile goes critical. Wigner furious it took so long. Blames "too much money".
1944-11	First Plutonium.

1.5.2 Rickover and the Nautilus

The Nautilus chronology, Table 2, is more well known so I will just make a couple of comments.

1. Rickover was originally leaning toward a sodium cooled reactor. He was turned off by the high pressure and low efficiency of the pressurized water reactor (PWR). However, Alvin Weinberg, the inventor of the PWR, convinced him that the PWR could be shoe horned into the small space available on a submarine and that's all that mattered to the Navy. Efficiency was not a criteria. But Weinburg never thought the PWR was the right concept for civilian power. Later as head of the Oak Ridge National Laboratory, he shepherded the development of the MSR, oversaw the MSRE, and became the foremost proponent of molten salt. But Weinberg did not have Rickover's maniacal drive.
2. The decision to go pressurized water was not made until March, 1950. Shortly thereafter Rickover, against the advice of all, decided to go straight to a full scale prototype. At the time no such thing as a PWR existed at any scale. Rickover wasn't scaling up. He was going from nothing to full scale.

After the Nautilus success, Rickover became convinced that the pressurized water reactor was the only concept worth building. He wrote a mocking parody extolling the benefits of non-PWR concepts one of which was "unavailable". In fact, at the time, most of these concepts were far further along than the PWR was when he committed to a full-scale prototype. Rickover also put in place the extensive paperwork system that became the NRC regulatory process, guaranteeing that his feat would never be repeated in the United States. But a younger Rickover knew that paperwork was not the answer. In his words, "Good ideas are not adopted automatically. They must be driven into practice with courageous impatience."

1.5.3 Camp Century

An instructive exception to Rickover's control of American nuclear effort was the Army's successful small reactor program in the very late 1950's. Table 3 shows the chronology of one of these plants. Camp Century was located at 77N in one of the most inhospitable places on the planet, 6000 feet above sea

Table 2: Nautilus chronology

1949-02	Rickover given control of naval nuclear propulsion. Still leaning to sodium.
1950-03	Decision to go PWR. At the time, no PWR of any scale had ever been build. Just a Weinberg patent and some sketches. No one knew how to make control rods, cladding, bearings that could handle PWR conditions.
1950-??	Decision to go <i>straight to full scale prototype</i> , S1W, in Idaho. No pilot plant. Nil sub-system testing. Westinghouse, BuShips aghast.
1950-08	Construction of S1W starts. Delayed by bad winter.
1951-08	Electric Boat awarded Nautilus contract.
1953-03	S1W, the first PWR ever built, goes critical.
1954-01	Nautilus keel laid.
1954-09	Nautilus commissioned.
1955-01	“Underway under nuclear power”.

Table 3: Camp Century, PM-2A Chronology

1959-01-23	\$4,500,000 contract with American Locomotive Company signed 10 MWt, 2 MWe Plant Designed, Built, Tested.
1960-07-10	Plant arrives Thule on ship. 27 packages. Sledged inland. Erected in 78 days.
1960-10-03	First Criticality, all in cost \$5,700,000
1960-11-12	Plant operating on Greenland Icecap. Total time 22 months.

level on the Greenland Plateau, 800 miles from the North Pole. In January, 1959, the Army signed a 4.5 million dollar contract with the American Locomotive Company (ALC) for 10 MWt nuclear plant, dubbed PM-2A. ALC designed, built, and tested the plant in 16 months. The plant comprised 27 blocks. In mid-summer of 1960, the blocks were shipped to Thule, sledged 150 miles north, and erected in 78 days. In the summer of 1964, Camp Century was shut down. The PM-2A was disassembled and returned to the USA.

Points to ponder:

- Non-standard nuclear manufacturer.
- Plant built entirely on assembly line.
- Transported by ship in blocks to site.
- Erection time measured in weeks.
- Disassembled by reversing the process.

When you consider what these three projects accomplished, the ThorCon schedule not only becomes feasible, but appears downright dilatory. Eugene Wigner, for one, would not be impressed.

1.6 The Big Caveat

But ThorCon faces one problem that these pioneers did not. In the 1950's, nuclear power was regarded as the gateway to a whole new world. Now it is regarded as a faustian bargain capable of almost

unimaginable harm. This is a product of a long history of conflating the dangers of acute, high level radiation with the health implications of chronic, low level radiation

ThorCon requires a country that is willing to take a fresh look at the hazards associated with nuclear power, informed by what we now know about how the cell and higher level organisms respond to radioactive stress. We are confident that such an investigation will reveal that nuclear power is unequivocally safer than fossil fuels, especially coal. We are confident that such an examination will conclude that nuclear power can be regulated like any other hazardous activity. The best model would be commercial airline travel.

Such regulation should be based not on paperwork metrics, competition stifling certificates, and rerunning the same computer programs which were used to design the plant with the same positive results. Rather regulation should be based on envelope pushing, physical testing of a fullscale prototype including imposition of worst credible casualties. Once a design has passed all these tests, it must be replicated precisely and efficiently. The only feasible way of doing this is on assembly line. ThorCon requires a country that adopts this attitude.

2 The Pros and Cons of Liquid Fuel

2.1 History

ThorCon is a liquid fuel reactor. The nuclear fuel is dissolved in a molten fluoride salt. The reactor operates at high temperature (700C) and near ambient pressure.

The liquid fuel concept goes back to the earliest days of nuclear power. After World War II, there was a debate between the Argonne Lab led by Enrico Fermi and the Oak Ridge National Lab (ORNL) led by Alvin Weinberg on how best to use nuclear power in peace. Both men were thoughtful geniuses. At the time, both thought that the world's supply of uranium was very limited, and both knew that only 0.7% of this uranium, the isotope ^{235}U was fissile, that is, could be made to fission. This meant that for every ton of ^{235}U produced about 140 tons of uranium were required; and, in order to be useful in a reactor, that uranium had to be put through an expensive enrichment process in which the ^{235}U was separated from the much more common, but not fissile isotope ^{238}U .

In the early 1950's, the only non-experimental power reactor was the Navy's pressurized water reactor, which had been invented by Weinberg, and developed in a crash program for submarine propulsion. This reactor used highly enriched ^{235}U in solid form. The working fluid was water at very high pressure (160 bar) but at a rather low temperature (330C). Both Fermi and Weinberg believed that, if we tried to use this system for civilian power, we would very quickly run out of uranium.

Fermi argued for sticking with solid fuel based on ^{235}U , and the enrichment technology developed during the war. His solution to the fuel problem was to bombard non-fissile ^{238}U , which makes more than 99% of naturally occurring uranium, with high energy neutrons in the reactor. This converts some of the ^{238}U to plutonium which then can be fissioned. This became known as the fast breeder reactor. Most of the designs based on this concept use a liquid metal — often sodium — as the coolant.

Weinberg, following his mentor, the uber-genius Eugene Wigner, argued for a completely different approach. His idea was a liquid fuel reactor based on converting thorium to the fissile uranium isotope, ^{233}U . Thorium is 500 times more abundant than ^{235}U , much more easily mined, and requires no enrichment. The reactor is made up of a core of molten salt in which the ^{233}U is dissolved, surrounded by a blanket, also a molten salt in which the thorium is dissolved. The blanket gets bombarded with some of the core neutrons converting the thorium to ^{233}U . Both fluids are continuously circulated. The core

salt is run through a heat exchanger, and some processing to remove fission products and add some new ^{233}U from the blanket. The blanket fluid is processed to remove the ^{233}U which is sent to the core, and replaced with new thorium.

It turned out that there was far more uranium on the planet than Fermi or Weinberg thought. The industry fastened on the pressurized water reactor (PWR) that the Navy had developed. It was the quickest way to deploy civilian reactors. The PWR took advantage of the very expensive and difficult enrichment process developed for the bomb. The same companies who built the Navy's reactors could use nearly the same skills and knowledge to build the civilian reactors. And the PWR had an attractive business model. Once you sold a reactor, the customer had to come to you for the specialized fuel elements for the life of the reactor. In some cases, companies were willing to take a loss on the construction in order to obtain the fuel element cash flow.

Weinberg kept the liquid fuel concept alive at ORNL including building a 8 MW test reactor that ran successfully for four years, 1965 to 1969. But the program had almost no political support. The decision was made to focus all the nation's reactor research effort on the fast breeder. In 1976, the Nixon administration shut down the molten salt reactor program. A few years later the fast breeder program, which was experiencing skyrocketing costs was effectively shut down as well.

2.2 The Molten Salt Reactor Experiment

The liquid fuel effort at Oak Ridge began by attempting to solve the impossibly difficult problem of powering an airplane with a nuclear reactor. The engineers knew that they had to have both high temperature and low pressure. They tried many ideas before coming up with molten salt. This culminated in the Aircraft Reactor Experiment (ARE), a 2.5 MWt reactor which ran successfully for 1000 hours in 1954. The ARE operated at an outlet temperature of 860C. This required a large R and D effort which produced 100's of high quality reports on salt properties, corrosion, radiation response, etc.

The ARE was a remarkable achievement. But just about everybody involved knew the real target was civilian power. In 1956, Oak Ridge obtained modest funding for a civilian liquid fuel reactor. In 1959, 4 million dollars was allocated for a 8 MWt pilot plant, dubbed the Molten Salt Reactor Experiment (MSRE). Oak Ridge moved at a leisurely pace. It was not until 1962 that construction started. On June 1, 1965, the MSRE achieved first criticality. On May 23, 1966, the reactor was taken to full power. In December, 1966 they began extended runs at full power on ^{235}U . In January, 1969, the fuel was switched to ^{233}U . By end of 1969, Oak Ridge felt they had learned all they could from the extensively instrumented MSRE. This knowledge was carefully documented in scores of reports. In December, 1969, the MSRE was shut down after a total of 13,172 full power hours. In its last 15 months, the reactor had an 87% availability, an unprecedented number for a first of a kind pilot plant.

Overall, the MSRE was an extremely successful experiment answering many questions, and raising almost no new ones. It produced the knowledge base that ThorCon is built on. Without the MSRE, ThorCon's schedule would be unthinkable.

2.3 The Pros

The advantages of liquid fuel are manifold:

Efficiency ThorCon has a thermal efficiency of 44% similar to a modern coal plant. Light water reactors operate at about 33%. This difference translates into lower capital cost and a 40% reduction in cooling water requirement.⁴

Walkway Safety ThorCon like most liquid fuel reactors is designed so that, if the reactor overheats for whatever reason, the fuelsalt passively drains from the reactor to a tank where the chain reaction cannot happen. This requires no operator intervention. There is nothing the operators can do to prevent this drain.

ThorCon combines a negative temperature coefficient — reactivity decreases as the core heats up — with large temperature margins and sizable thermal inertia. Even if the fuelsalt somehow fails to drain and the control rods fail to scram, ThorCon will shut itself down. Section 3.8 offers an example.

Safety is further enhanced by the lack of pressure in the reactor. And there is no phase change in the event of a rupture. In fact, due to the very low vapor pressure of the salt, nil gas is released by a primary loop rupture. The difference between the ThorCon's operating temperature (700C) and the salt's boiling point (1400 C) produces a robust safety margin. Finally, most of the fission products, including the hazardous cesium and strontium, are salt seekers. They will stay dissolved in the fuelsalt in the event of a breach.

Low Part Count A solid fuel reactor requires thousands of highly engineered, highly stressed fuel pins. Some fast breeder designs require more than 50,000. Failure of a single fuel pin will shut the reactor down and force a difficult decontamination process. In the molten salt reactor's primary containment there are a few dozen components and only one major moving part, a pump impeller. And if something does break, you can remove the principle source of radioactivity, the fuel, in order to fix things.

Move fuel around with a pump Anyone visiting a solid fuel nuclear plant will be amazed and perhaps appalled by the refueling systems. Some such as Candu are remarkably ingenious. Some such as Westinghouse PWR are horrible kluges. Some employ complicated robotic fuel element shuffling devices that are supposed to operate for decades in a highly radioactive environment with nil maintenance. All are complex, failure prone headaches that drive the design.

All this disappears with liquid fuel. Moreover, we can adjust the fuel composition on the fly. There is no need for excess reactivity to account for the fact that in a solid fuel reactor the quality of the fuel deteriorates over time. Nor need you worry much about variations in the fuel makeup. Solid fuel reactors, especially fast breeders, require a very even isotopic fuel composition to avoid hot spots. In a liquid fuel reactor, any variation in the fuelsalt is quickly mixed away.⁵

⁴ ThorCon's operating temperature is limited only by material considerations. As materials improve, the temperature can be raised, and the thermal efficiency still further improved. At 850C, we can disassociate hydrogen from water efficiently and produce hydrogen based fuels.

⁵ Less important but still significant advantages of liquid fuel include:

1. Nearly automatic removal of the most troublesome fission product poison, xenon-135.
2. No need to fabricate fuel elements. This is particularly important when it come to burning spent fuel.

No mausoleum Pressurized Water Reactors operate at 160 bar pressure. High operating pressure means 9 inch thick reactor vessels and piping. Some of these forgings can only be done by a few specialized foundries. Worse, if we have a big piping failure, the pressurized water explodes into steam spraying radioactivity all over the place. We need a very large and strong containment structure. In the event of a loss of coolant, this gigantic structure must somehow be kept cool despite all the decay heat in the core.

Now the reactor, heat exchangers and all sorts of plumbing are entombed in this mausoleum where they are extremely difficult to repair or replace. Therefore, we pretend that they will need essentially no maintenance for the life of the plant.

But it gets still worse. These colossal structures must be made out of reinforced concrete. The rebar is as thick as a man's wrist and so dense that special concrete is formulated to allow it to flow into place. Reinforced concrete construction has four major characteristics:

1. It cannot be done on a panel line using block construction.
2. It is nearly impossible to automate.
3. It cannot be done in parallel. The containment dome becomes the critical path. Even the Chinese cannot build a PWR in less than 4 years.
4. It is difficult to inspect and extremely difficult to repair.

ThorCon's low operating pressure and lack of phase change allows us to use a simple steel containment structure that can be manufactured as blocks on a panel line using less than 2 man-hours per ton of steel and assembled in months if not weeks.

2.4 The Cons

Despite the manifest advantages of liquid fuel and the success of the MSRE, the molten salt program disappeared into history. The obvious question is: why?

Politics aside, after the success of the MSRE, ORNL made a critical mistake. They focused almost all their effort on a thorium breeder. Unfortunately, a molten salt, thorium breeder was not and still is not feasible. There are two routes to a molten salt thorium breeder: two salt, and one salt. The two salt concept requires a barrier material that does not exist. The one salt design requires nearly continuous, complex chemical processing of a very hot, extremely radioactive fuelsalt. This process has not yet been fully demonstrated even at laboratory scale. Both concepts need highly enriched ${}^7\text{Li}$ which doesn't exist in anything like the quantities required. Almost all the sporadic attempts to resurrect the liquid fuel reactor have foundered on these problems.

ThorCon is a thorium converter, not a breeder. ThorCon requires periodic additions of fissile fuel. And the first generation ThorCon is not a particularly efficient converter. Only about 25% of its power comes from converting thorium to ${}^{233}\text{U}$. ThorCon derives its ability to produce power cheaply not from its use of thorium, but from all the other advantages of liquid fuel.⁶

However, even if it is only a converter, a liquid fuel reactor must still solve a number of difficult technical problems: Here are the cons:

Labile radioactivity A liquid fuel reactor can move fuel around with a pump. A liquid fuel reactor can use gravity to remove fuel from the core. Don't have to worry about meltdown when the fuel is already molten. Liquid fuel is the key to simplicity, trivial refueling, low part count, and safety. But liquid fuel does mean that the entire primary loop is extremely radioactive. Alvin Weinberg, who had a way with words, put it best: "vast amounts of radioactivity in labile form". Super-hot fission products are moving around and end up where we don't want them: primary loop pump, primary heat exchanger, drain tank, or worse. ***Something fails, you can't just open it up and fix it.***

High temperature A molten salt reactor operates at low pressure and high temperature. Low pressure means safety and major savings in piping and containment. High temperature means high thermal efficiency, much less rejected heat, and the real possibility of producing hydrogen based fuels economically. But high temperature is a two edged sword. It creates difficult thermal expansion issues. Do you operate the cell at furnace temperatures? If so, all kinds of instrumentation, materials problems. If not, how do you cool the cell without freezing the salts on upsets?

Moderator Life ThorCon is a *thermal* or slow reactor. It requires neutrons that have been slowed down by a factor of 5000 to sustain a chain reaction. The material that does the slowing is called the *moderator*. ThorCon uses graphite as the moderator. Graphite is damaged and eventually destroyed by radiation. If the reactor uses a low power density to obtain long moderator life, the result is a massive core. We can forget about assembly line construction. We need 20 plus years to prove the system, and we still have a mess at the end of core life, assuming the core actually lasts that long. If the reactor uses a high power density, we end up with compact core, but only a 4 or

⁶ The decision to not attempt a breeder improves proliferation resistance. In thorium breeder designs, it is theoretically possible to extract weapons grade ${}^{233}\text{U}$, although the process is so difficult a strong argument can be made that any country with the capability of extracting and building a ${}^{233}\text{U}$ bomb would find it much easier to make a bomb conventionally.

In the ThorCon, the ${}^{233}\text{U}$ is always denatured, that is, diluted with lots of ${}^{238}\text{U}$. Even if a bombmaker extracted the uranium, he would still have to enrich it. And since the ${}^{233}\text{U}$ would be contaminated with highly radioactive ${}^{232}\text{U}$, he would have to run his centrifuges in a hot cell. No sentient bombmaker is going to go this route.

5 year moderator life, ***and the used-up moderator needs a multi-year cool down before it can be safely handled.*** Even then disassembly and decontamination pose major problems. And to make matters still worse, the design must handle significant dimensional changes in the graphite as it is irradiated.

Reactor vessel life uncertainties The MSRE revealed the reactor vessel will be subject to helium embrittlement and tellurium attack. ORNL developed solutions but they were never fully tested at the flux and fluence levels which most molten salt designs anticipate using.

The Offgas System In a molten salt reactor, the gaseous fission products are continuously removed from the reactor. This is good news from a fuel burn and radioactive waste point of view. And it means that there are far less volatile fission products trapped in the reactor which could be spread around by a containment breach as happened at Chernobyl and Fukushima. But it also means we must handle an extremely radioactive gas stream contaminated with a salt mist continuously and reliably.

Tritium A molten salt reactor will create up to 60 times as much tritium as LWR. The health hazards associated with tritium are often vastly over-stated. However, to be readily acceptable, the molten salt design should match or better the LWR in tritium emissions.

No Lithium 7 The salt that most molten salt designs are counting on simply doesn't exist in anything like the quantities we need. The only proven process for creating it, Colex, was an environmental disaster.

The ThorCon solves these problems in a robust, simple fashion. ***And it does so without requiring any new technology.*** ThorCon is the Do-able Molten Salt Reactor.

3 The Technology

3.1 Preamble

This document is a high level summary description of ThorCon. It assumes the reader has some familiarity with nuclear power and a modest acquaintance with liquid fuel reactors in particular. This summary does not include the 900 plus pages of back-up calculations and reports supporting this design. These are described in some detail in the 220 page ThorCon Design Control Document.

3.2 The Can

The heart of ThorCon is a sealed *Can*, Figure 3.⁷ The Can contains a 550 MWt reactor, which we call the *Pot*, a primary loop heat exchanger (PHX), and a primary loop pump (PLP). The pump takes liquid fuelsalt — a mixture of sodium, beryllium, uranium and thorium fluorides — from the Pot at 704C, and pushes the fuelsalt over to the PHX at a rate of just under 3000 kg/s. Flowing downward through the PHX, the fuelsalt transfers heat to a secondary salt, and is cooled to 564C in the process. The fuelsalt then flows over to the bottom of the Pot, and rises through the reactor core, which is mostly filled with graphite slabs, called moderator. This moderator slows the neutrons produced by the fissile uranium, allowing a portion of the uranium in the fuelsalt to fission as it rises through the Pot, heating the salt to 704C, and (indirectly) converting a portion of the thorium to fissile uranium. It's just that simple; and just that magical. Rock burning 1.01.

Figure 3 shows some of the Can numbers. The Pot pressure is less than 2 bar gage, but the outlet temperature of 704C results in an overall plant efficiency of about 44%, and a net electrical output per Can of 250 MW. After start up, to produce this electricity, the Can requires 880 kg of 20% enriched uranium per year.⁸ The Can is a cylinder 11.6 m high and 7.3 m in diameter. It weighs about 400 tons. The Can has only one major moving part, the pump rotor.

3.3 The Hammock Suspension System

One of the difficult problems facing a molten salt reactor is accommodating thermal expansion. The ThorCon handles this by hanging the Pot and the PHX from the Can lid by cables, allowing the Can at about 300C to expand independently from the primary loop structure which is at 700C at the top and 550 C at the bottom. More importantly, the cables allow the Pot and the PHX to be pushed apart as the primary loop heats up. The cables (actually TZM rods fitted with an eye at either end)) allow the PHX not only to move laterally but also to tilt and rotate. The hammock system allows all this to happen in the very tight confines of the Can.

Almost all the vertical expansion is downward. The drain line is hung from the PHX to Pot line and has no direct physical connection to the Can. So this vertical movement is unrestricted and the drain line at Can temperature is free to expand independently of the primary loop. The lack of any physical connection between the Can and the primary loop below the Can lid also allows us to drop the Can wall and bottom from the Can lid and primary loop during the disassembly process at the remote recycling facility, without breaking any connections inside the Can.

⁷ The drawings in this document tend to be cartoonish. This is due in part to the fact that they are created dynamically by the ThorCon DNA model. The DNA model is a totally parametric description of the plant. One can change any of the independent variables in the design, press a button, and (nearly) all the results, including costing, drawings, 3-D models tables, analyses and reports will change accordingly.

⁸ ThorCon can operate on a wide range of fuels including enrichments down to less than 4%.

3.4 The Silo

The Can is located in an underground Silo as shown in Figure 4. The top of the Silo is about 29 m underground. This sketch shows the secondary salt loop in green. The secondary salt is a mixture of sodium fluoride and beryllium fluoride containing no uranium or thorium. Hot secondary salt is pumped out of the the top of the Primary Heat Exchanger to a Secondary Heat Exchanger where it transfers its heat to a mixture of sodium nitrate and potassium nitrate commonly called solar salt from its use as an energy storage medium in solar plants. The solar salt, shown in purple in Figure 4, in turn transfers its heat to a steam loop, shown in red, creating supercritical steam, and also reheating that steam to increase the plant's efficiency. Table 4 shows the main parameters of the four loops.

	Mass Flow	Hot	Hot	Cold	Cold	Fluid
	kg/s	C	bar	C	Bar	numbers are mol fraction
Fuelsalt	2994	704	10.5	564	4.0	NaF-BeF ₂ -ThF ₄ -UF ₄ 76/12/9.5/2.5
Secondary Salt	1534	621	10.5	454	20.0	NaF-BeF ₂ 57/43
Tertiary Salt	1414	598	12.0	344	1.0	NaNO ₃ -KNO ₃ 55/45
Steam Main	225	538	248	288	260	
Steam Reheat	162	538	38	343	39	

Table 4: The Four Loops (flows are per module, fuelsalt is fresh)

Directly below the Can is the Fuelsalt Drain Tank (FDT). In the bottom of the Can is a *fuse valve* shown in blue in Figure 3. The fuse valve is merely a low point in a drain line. At normal operating temperatures, the fuelsalt in the fuse valve is frozen creating a plug. But if the Can heats up for any reason, the plug will thaw, and the fuel salt will drain to the FDT. Since the drain tank has no moderator, fission will stop almost immediately. This drain is totally passive. There is nothing an operator can do to prevent it.

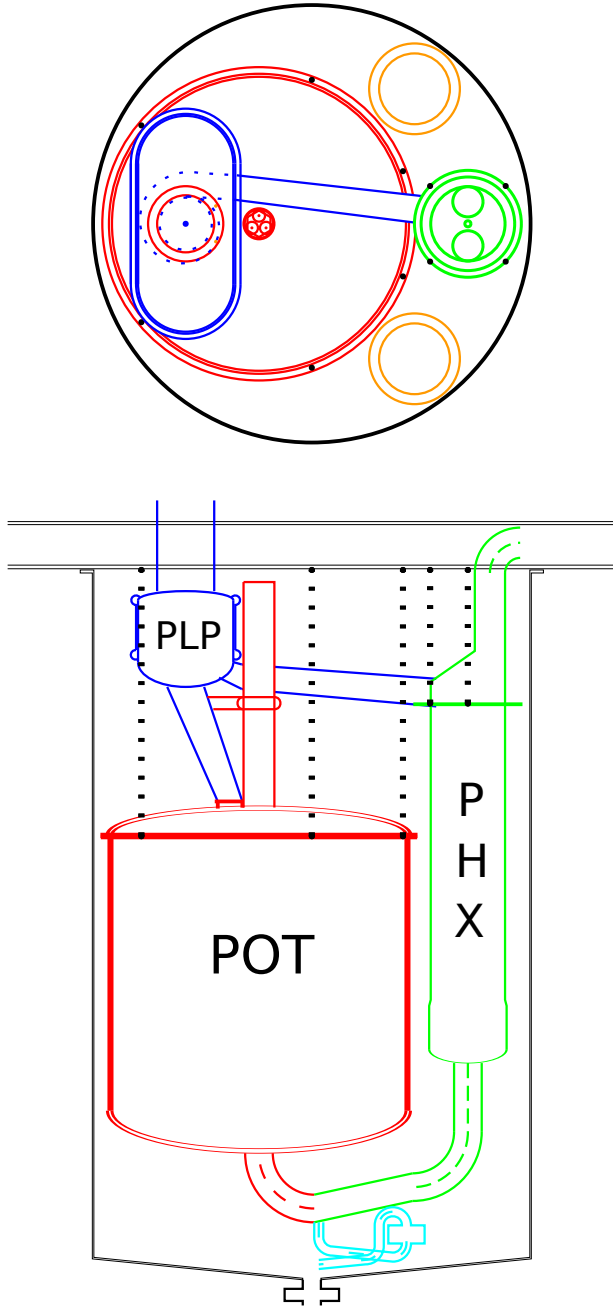
3.5 The Power Module

Achieving 250 MW's of electricity from such a compact reactor requires a moderately high power density of 16 MW per m³ of active core. This results in a short moderator life. We must be prepared to replace the core graphite every four years. Therefore, the Cans are arranged in pairs to make up a *Power Module*, as shown in Figure 5. At any one time, just one of the Cans of each module is producing power. The other is in cool down or stand-by mode. This is very much like a duplex filter. When one side needs to be taken out of service, you flip to the other side, transferring the fuelsalt from the old Can to the new. The old Can then remains in its silo for a little less than 4 years, at which point the Can decay heat will be down to less than 1 kW.

3.6 The Silo Hall

A ThorCon plant is made up of one or more power modules. Figure 6 is a sketch of a four module, 1 GWe plant. The solid red circles are the Can silos. The modules are arranged in a long *Silo Hall* served by a 500 ton bridge crane, as shown in Figure 4. This entire structure is underground. At one end of the hall is the control room. This end also houses the *Transfer Module* comprising salt storage tanks, offgas Xe and Kr storage, fuelsalt transfer casks, and two Can transfer pits.

Figure 3: The ThorCon Can: a Pot, a Pump, and a Still



Can electrical output(MW):	250.0
Can thermal output(MW):	550.0
Fuelsalt:	NaF-BeF2-ThF4-UF4
Mol percent:	76/12/9.5/2.5
Fuelsalt flow rate(kg/s):	2934
Pot temperature in(C):	564.0
Pot temperature out(C):	704.0
Loop transit time(seconds):	13.8
Pot inlet pressure(bar g):	1.65
Primary Loop Pump(kW):	1006
Annual 20% LEU fuel(kg):	880
Annual 20% LEU fuel(liters):	460.0
Overall plant efficiency(%):	45.0
Percent thorium fueled(%):	abt 25
Can OD(m):	7.259
Can Height(m):	11.738
Pot OD(m):	4.961
Pot Height(m):	5.717
Can weight (no salt) (kg):	381,767
Fuelsalt weight (kg):	40,515
2014-10-21T16:41:37Z	
Version:	1.08

Figure 4: Silo hall cross-section

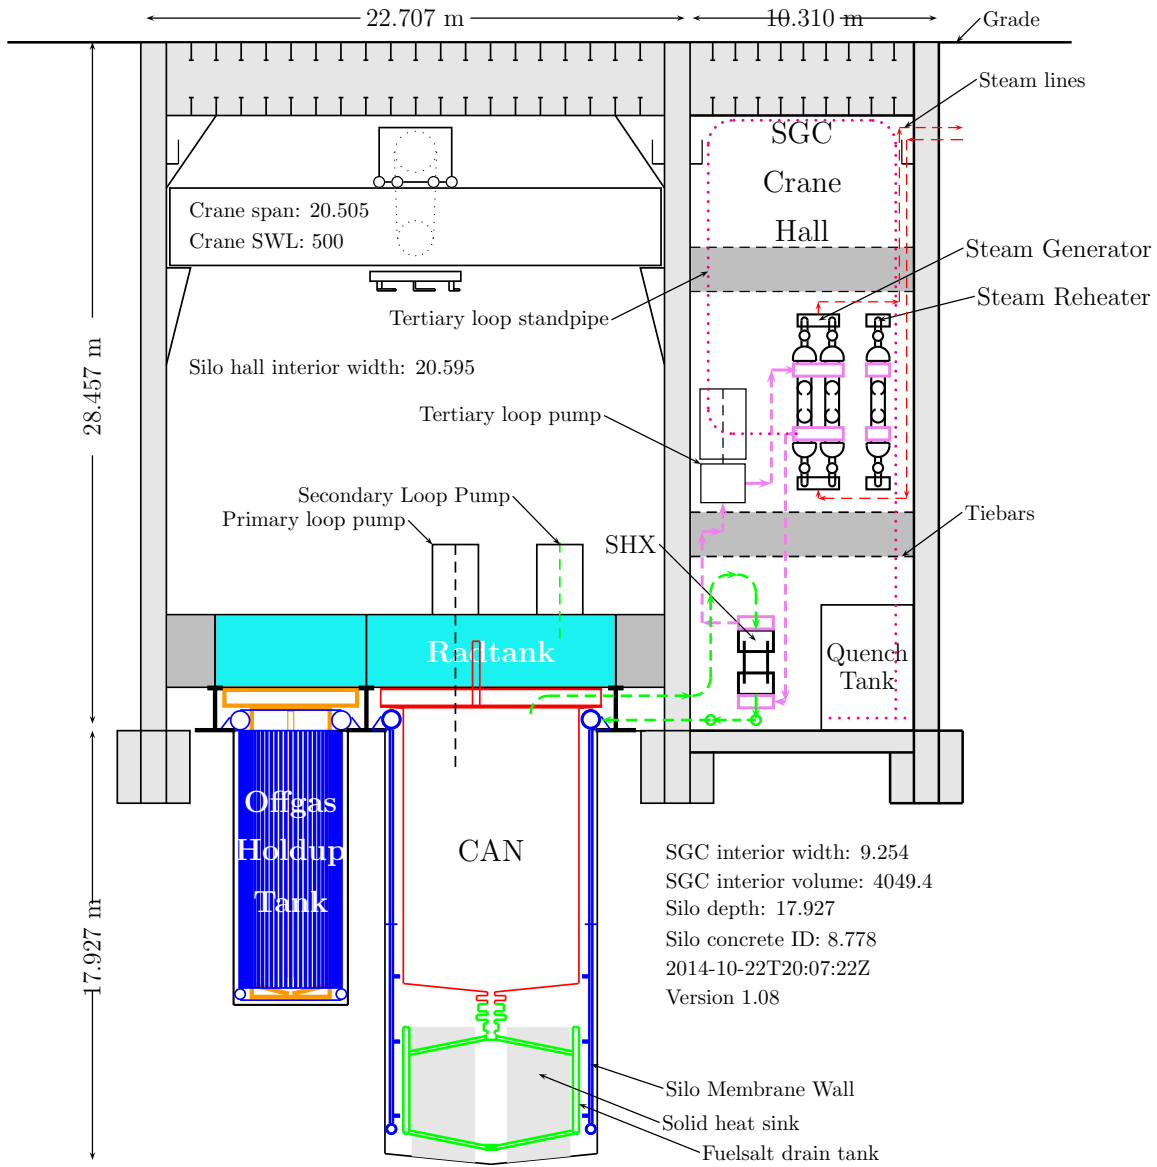


Figure 5: Plan view of module




Figure 6: Silo hall plan

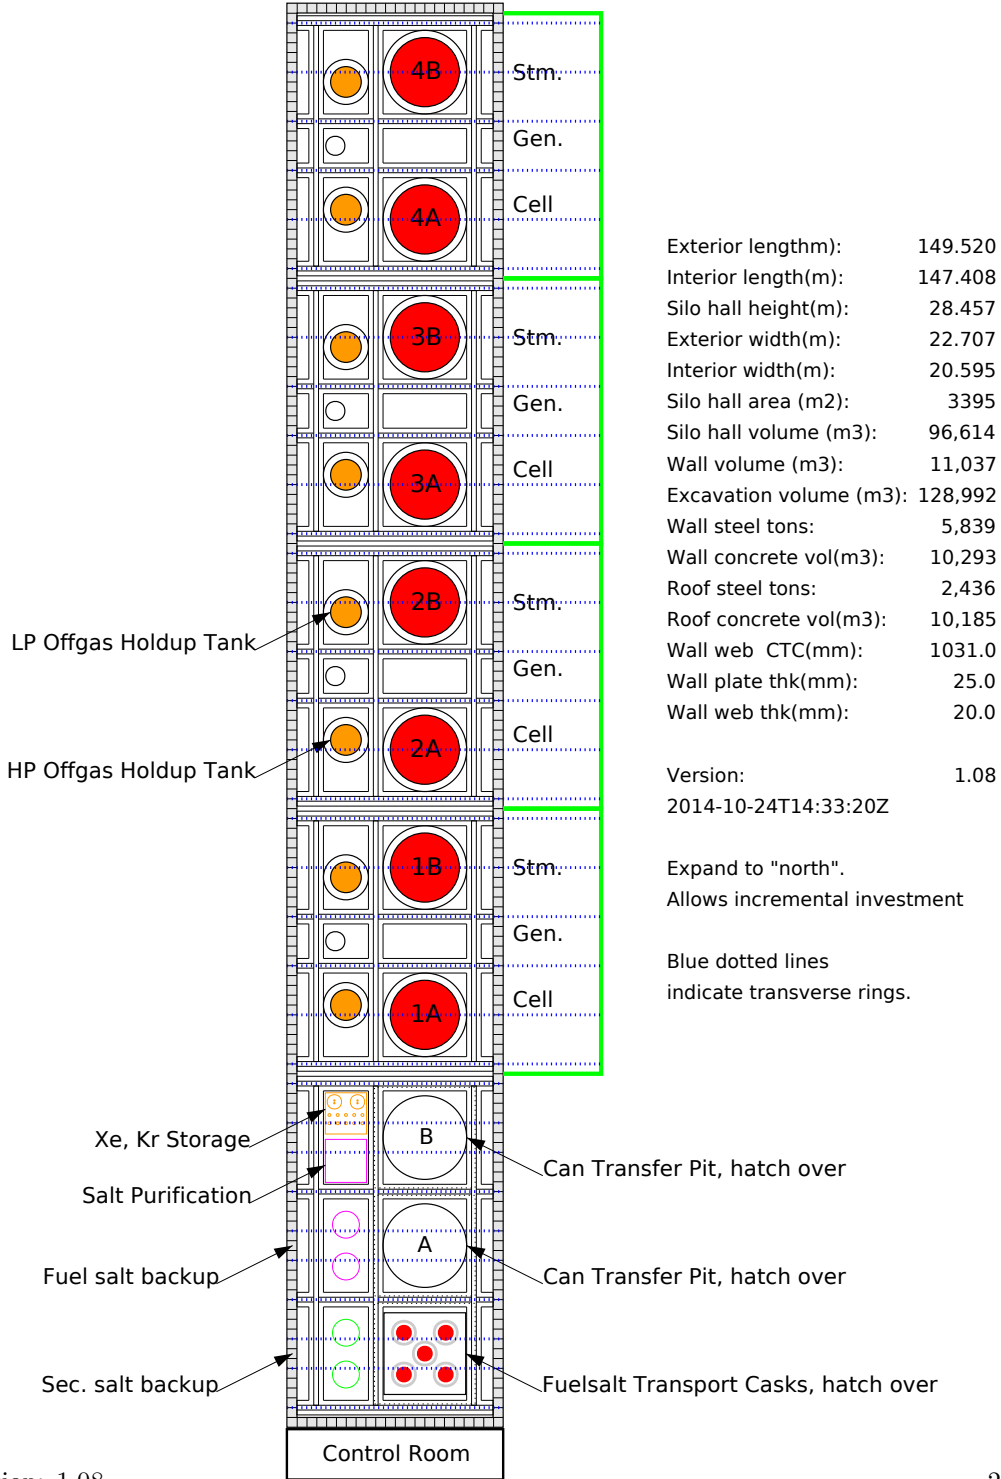


Figure 7: Site plan

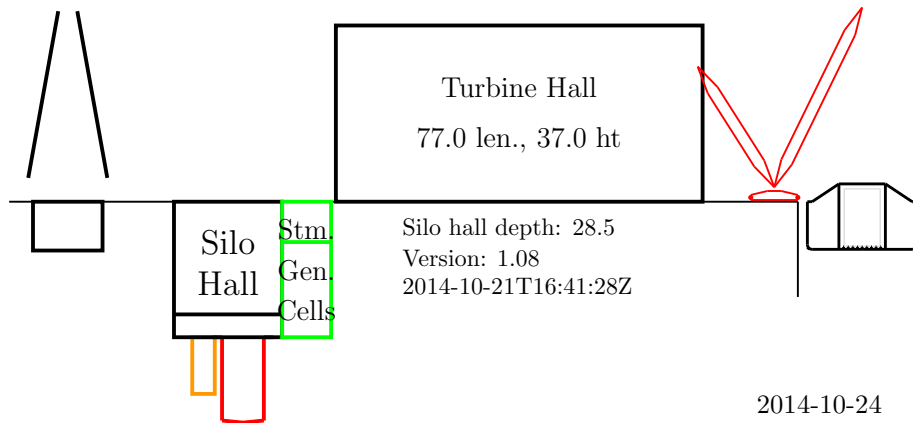

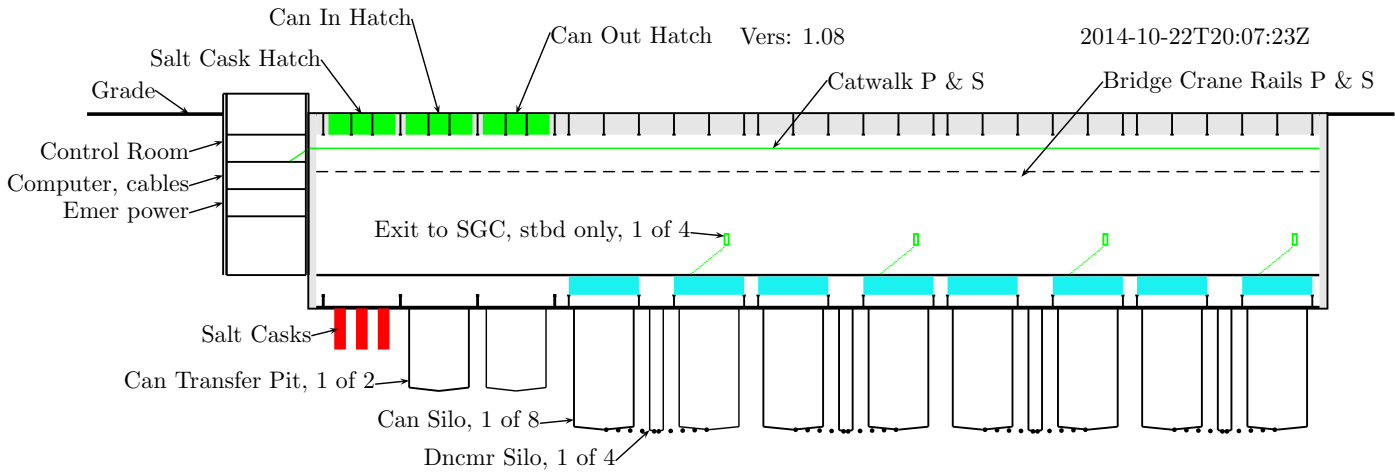


Figure 8: Silo hall profile looking west



3.7 The Recycling System

In the ThorCon system, no complex repairs nor fuelsalt processing are attempted at the plant. Each 50 ThorCon plants are supported by a centralized recycling facility (CRF). Normally, Cans are changed out every four years and the fuelsalt is reprocessed every eight years. When the Cans or fuelsalt need replacing, they are shipped to the CRF in a special purpose Canship. The difficult problems of disassembly, decontamination and waste handling are shifted from the plant to this facility.

After a used Can has cooled in its silo for nearly four years, the Can can be safely lifted out of its silo, transferred to special purpose *Canship*, and shipped to the CRF. Figures 7 and 8 indicate how this is done. The silo hall bridge crane moves the used Can to one of two Can Pits. There is a large hatch above each pit. These hatches are shown in green in Figure 8. A crawler crane transfers a replacement Can from the ship to one of these pits, pulls an old Can from the other, and loads it on the Canship.⁹ This leapfrog process continues until all the Cans are replaced.

When we need to change out the fuelsalt due to the build up of fission products — which will happen on an eight year cycle — the old salt will remain in its Can for close to four years. During this period, the old fuelsalt is as well protected as the salt in the operating Can. ***There is no need for a separate, vulnerable, spent fuel cooling and storage system.*** By the time, we pump the old fuelsalt to shipping casks, shown in red in Figure 8, its decay heat will be down to 80 kW, 0.25% of the original.

The fuelsalt going both ways will be unattractive weapons material. The uranium will be both fully denatured and, after the initial load, contain enough ^{232}U to further complicate a bombmaker's life, while at the same time allow tracking of any diversion. The plutonium will contain sufficient ^{238}Pu to make even a fizzle weapon infeasible for all but the most advanced weapons states. This material will be significantly more anti-proliferation resistant than the MOX fuel which is currently being transported.

This system of regular replacement of the most critical components means that major upgrades can be accomplished without significantly disrupting power generation. And since the returned Cans are disassembled and fully inspected, incipient problems will be caught before they can turn into casualties.

Fuelsalt processing will start out with separation of the uranium by fluorine volatility and then recovery of the salt by vacuum distillation. Cesium will boil off early in this step and be sequestered. Most of the uranium and salt will be returned to the fuel stream.¹⁰ Currently, it is not economically feasible to remove the plutonium and other transuranics from the fission products, so the *ash* from the distillation will be stored at the recycling facility. Each 1 GWe ThorCon will produce about 8 m3 of this ash every 8 years.¹¹

It is far more economic to upgrade a single centralized facility than 50 plants. Eventually, plutonium separation will become feasible, and most of the plutonium will be returned to the plants to be fissioned.

⁹ The Canship is 150 m long, 23 m beam, 14 m air draft, with a light draft of 2.5 m, allowing it access to most major rivers outside of Europe.

¹⁰ Burnup numbers indicate that the uranium can be recycled 3 or 4 times before ^{238}U build up will force re-enrichment or disposal.

¹¹ With nil krypton or cesium, the ash will have little gamma activity. Handling and storage will not be difficult.

3.8 Passive Decay Heat Cooling

A unique and critically important feature of ThorCon is the *membrane wall*. Each silo is fitted with a membrane tube wall, shown in blue in Figure 9. Figure 10 is a better sketch of the membrane wall. The membrane wall is made up of vertical steel tubes connected by strips of steel plate which are welded to the tubes. These tubes are filled with water and connected by circular headers, top and bottom. The top header is connected by a riser to a heat exchanger in the pond.¹² The outlet of the heat exchanger is connected to the wall bottom header by a downcomer. Figure 11 is an overall sketch of this arrangement. Figure 12 zooms in a bit on a module. In Figure 12, the membrane wall piping at the top of the silo is shown in blue. The downcomers enter from the left at the middle of the grid, disappear down the *downcomer silo*, and at the bottom of the silo fan out to the membrane walls. This is indicated by the light dotted bluelines. The risers are the blue pipes with arrows pointing to the left.

The Can is cooled by thermal radiation to the silo membrane wall. This heat converts a portion of the water in the wall tubes to steam. This steam/water mixture rises by natural circulation to the cooling pond, where the steam is condensed, and returned to the bottom of the can membrane wall.

The silo membrane wall also cools the Fuelsalt Drain Tank (FDT). The FDT is located directly below the can as shown in Figure 4. The drain tank is tall, thin rectangular trough that has been wrapped into a circle. This arrangement provides sufficient radiating area to keep the peak tank temperature after a drain within the limits of the tank material. This cooling process is totally passive, requiring no operator intervention nor any outside power.¹³

The silo membrane wall is what makes the ThorCon work.

1. The membrane wall allows us to keep the Can interior below 300C during normal operation. The fact that the membrane wall is always operating is an important safety feature. If a problem develops in the membrane wall loop, we will find out before a casualty occurs rather than during.
2. The wall allows us to capture any tritium permeating through the Can or drain tank in the inert gas in the annulus between the Can/FDT and the walls.
3. The wall cools more rapidly as the Can/FDT tank heats up, but more slowly as the Can/drain tank cool down, which is exactly what we want to handle both emergencies and avoid salt freeze ups.
4. The wall maintains a double barrier between the fuelsalt and the membrane wall water, even if the primary loop is breached.
5. And the wall does all this without any penetrations into the Can or the fuelsalt drain tank.
6. The membrane wall protects the silo's concrete lining.

After a full power drain, the power that must be transferred to the membrane wall peaks at 4.7 MW about 2 hours after the drain, Figure 13. At that point the FDT tank temperature is about 960C. Calculations show that the membrane wall can remove more than 20 MW of power, if it had to.

¹² The pond also serves as a reserve of cooling water in certain casualties. It can be operated once through, or fitted with open or closed loop cooling towers depending on the availability of water at the site.

¹³ Some recent light water reactor designs claim to be passive. But on closer inspection they require a set of valve operations to realign the system. Both an active response and power to implement that response are needed.

Figure 9: Cross-section of silo looking north




Figure 10: Silo Can Membrane Wall

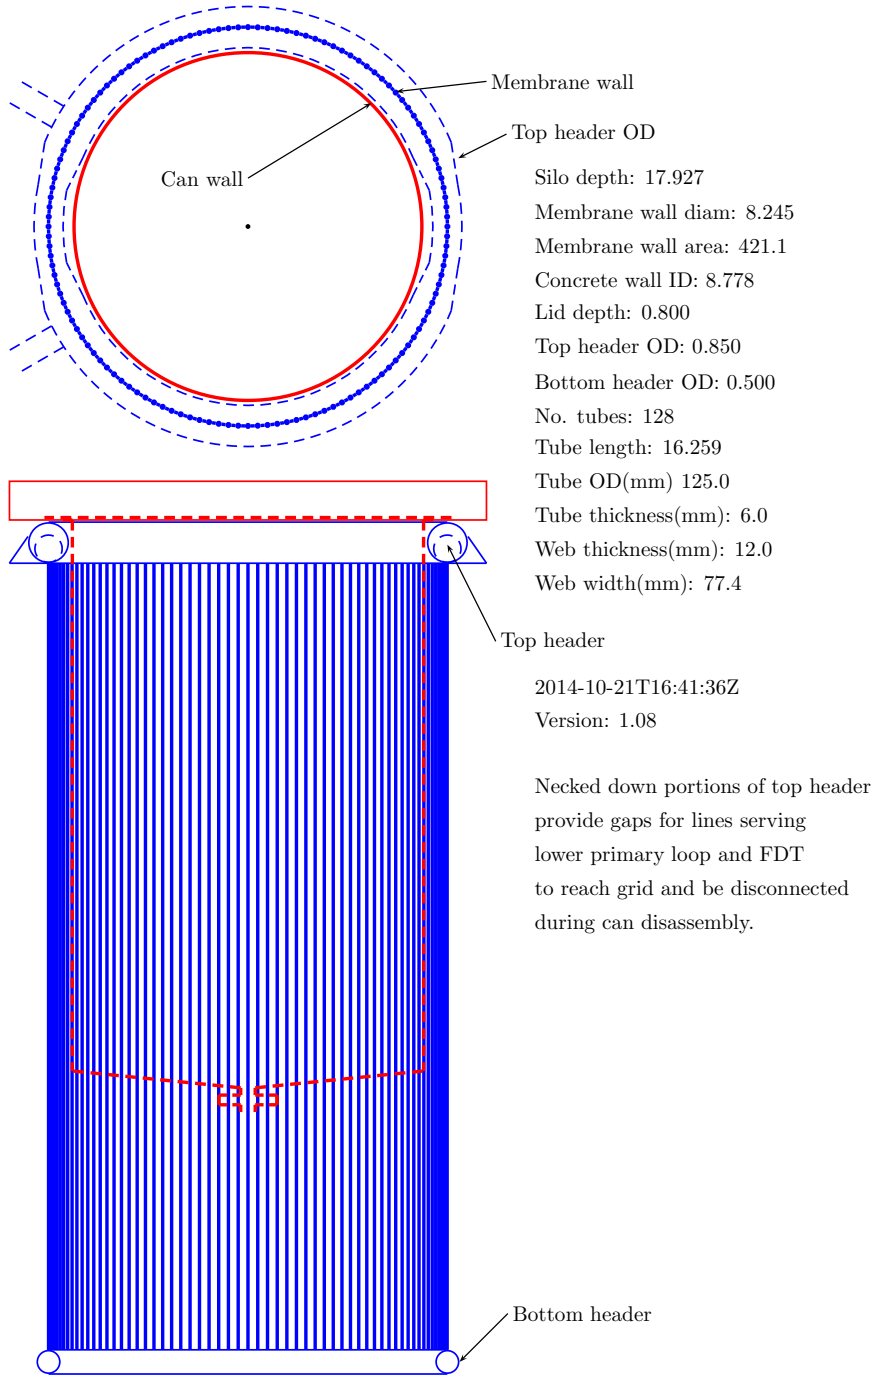


Figure 11: Silo Can Cooling System


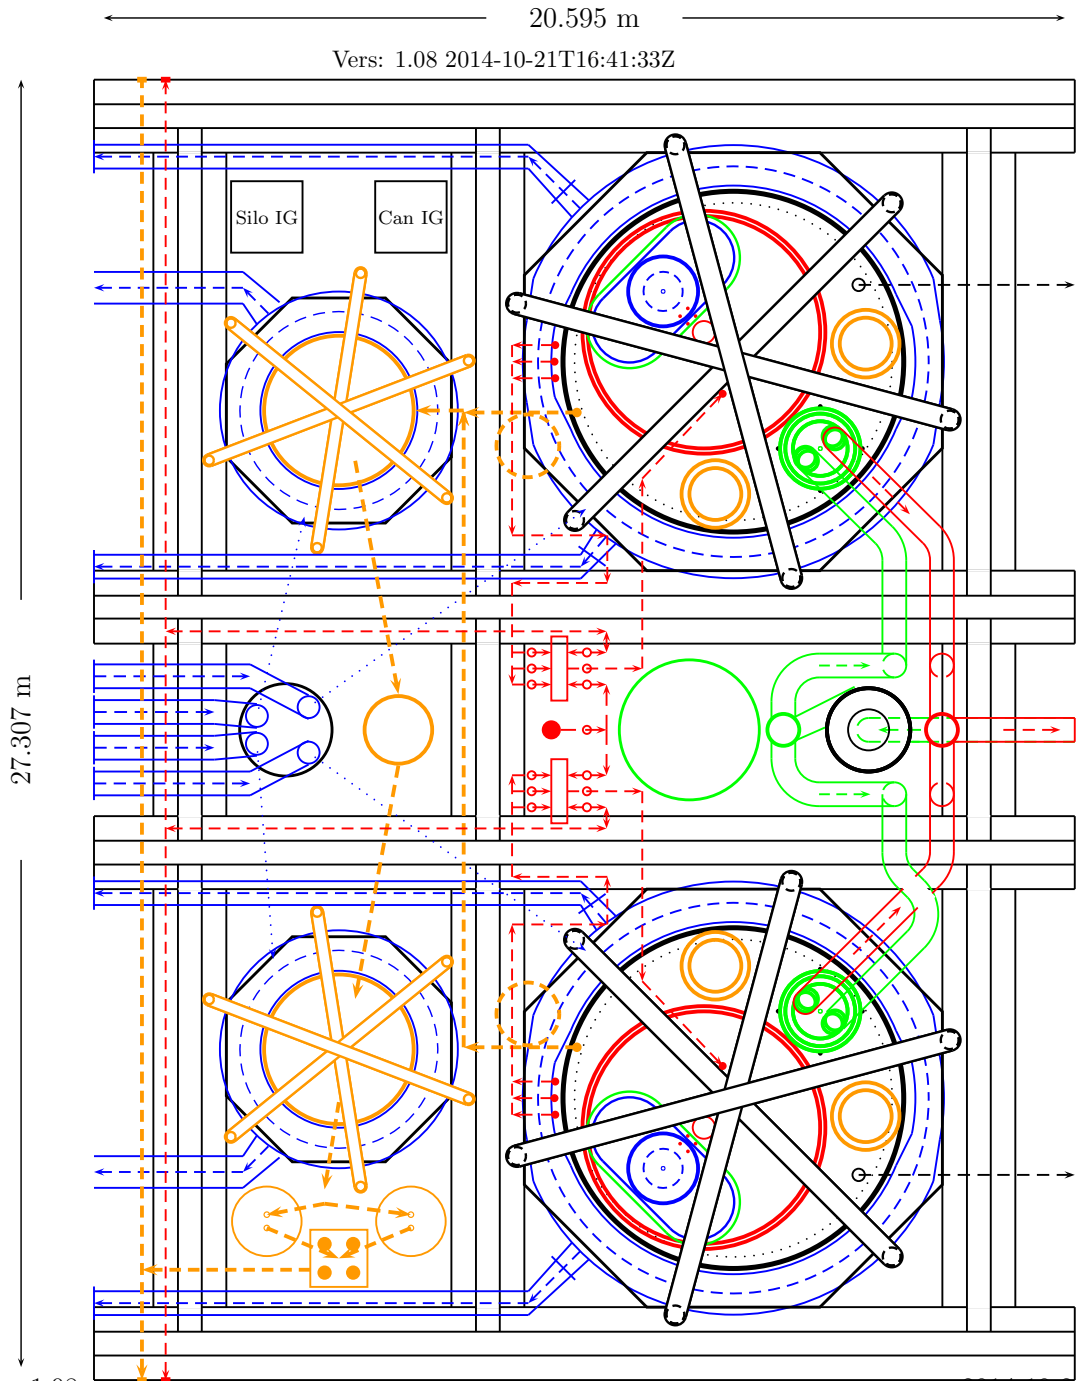


Figure 12: Plan view of grid block



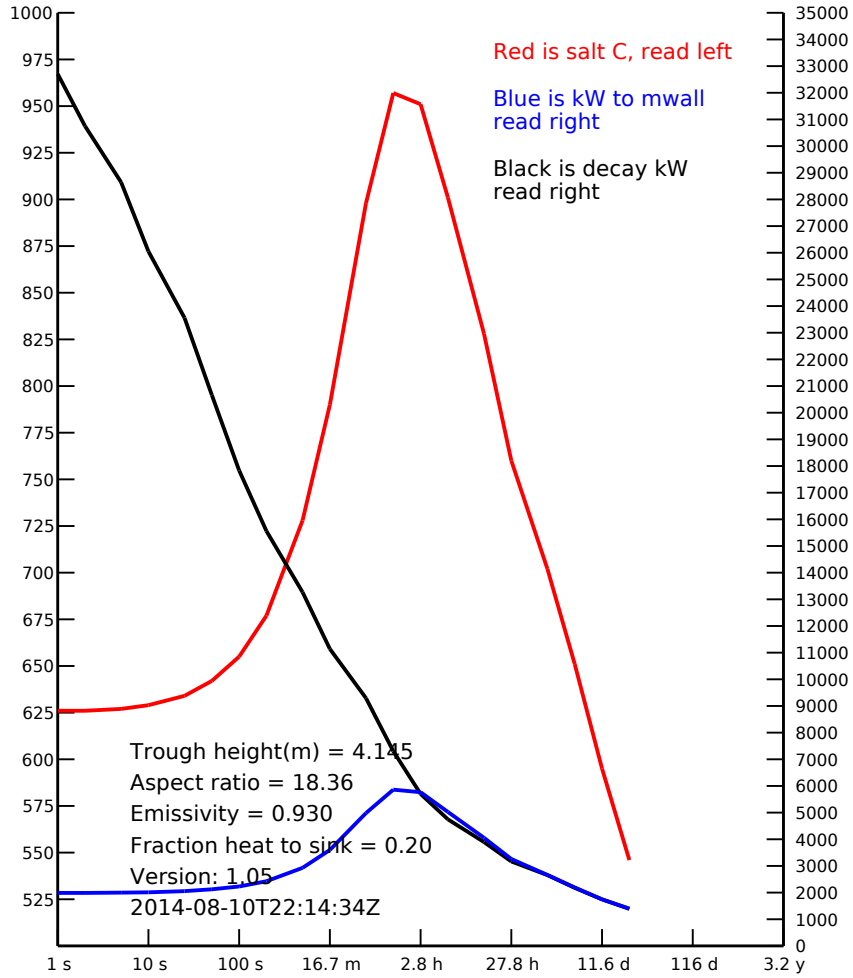


Figure 13: FDT salt temperature profile, full power drain

3.9 Passive Shutdown

Membrane wall cooling is remarkably robust against both modelling errors and casualties which overheat the Can or the FDT. There are two reasons for this:

1. Radiation to the wall goes as T^4 . Any increase in Can or drain tank temperature drastically increases the heat radiated to the membrane wall.
2. Increased temperature markedly improves the fuelsalt's heat transfer properties, mainly by decreasing viscosity.

The loss-of-heat-sink, fails-to-drain casualty is a striking example.

Figure 14 shows what happens to the fuelsalt temperature in this casualty, assuming a successful scram. We have about 13 hours before the primary loop reaches 950C. If we do nothing, the primary loop temperature will stay above 950C for about three days, peaking at 1010C. We expect the primary loop to fail about 10 hours into that period due to creep. When it does, the fuelsalt will drain to the refractory lined bottom of the Can and then to the drain tank.

Even if the control rods fail in a loss-of-heat-sink, fails-to-drain casualty, ***ThorCon will shut itself down***. ThorCon combines a negative temperature coefficient — reactivity drops as the core temperature increases — with large temperature margins and sizable thermal inertia. Figure 15 shows what happens if we lose the secondary loop heat sink, and the primary loop fails to drain, and *the scram fails*. The fuelsalt will heat up, but this increase in temperature automatically reduces the reactor power. If we do absolutely nothing, the fuelsalt temperature will rise to 756C in about 120 seconds at which point the fission power output of the reactor will be essentially zero. ThorCon will produce about 6600 MJ of energy during the shutdown, equivalent to about 20 minutes of decay heat.¹⁴

We can upper bound this no-scram casualty by assuming the fuelsalt starts off at 756C. In other words inject all the shutdown energy into the system at time zero. Figure 16 shows the results. ***The peak temperature is only about 25C higher than the same casualty with scram***. ThorCon safety is not dependent on the proper functioning of the control rods.

¹⁴ Our transient model assumes no heat loss and no decay heat. Heat loss slows down the shutdown, but the decay heat will speed it up. Early in the casualty the decay heat is larger than the heat loss, so the actual shutdown will be faster than Figure 15 claims.

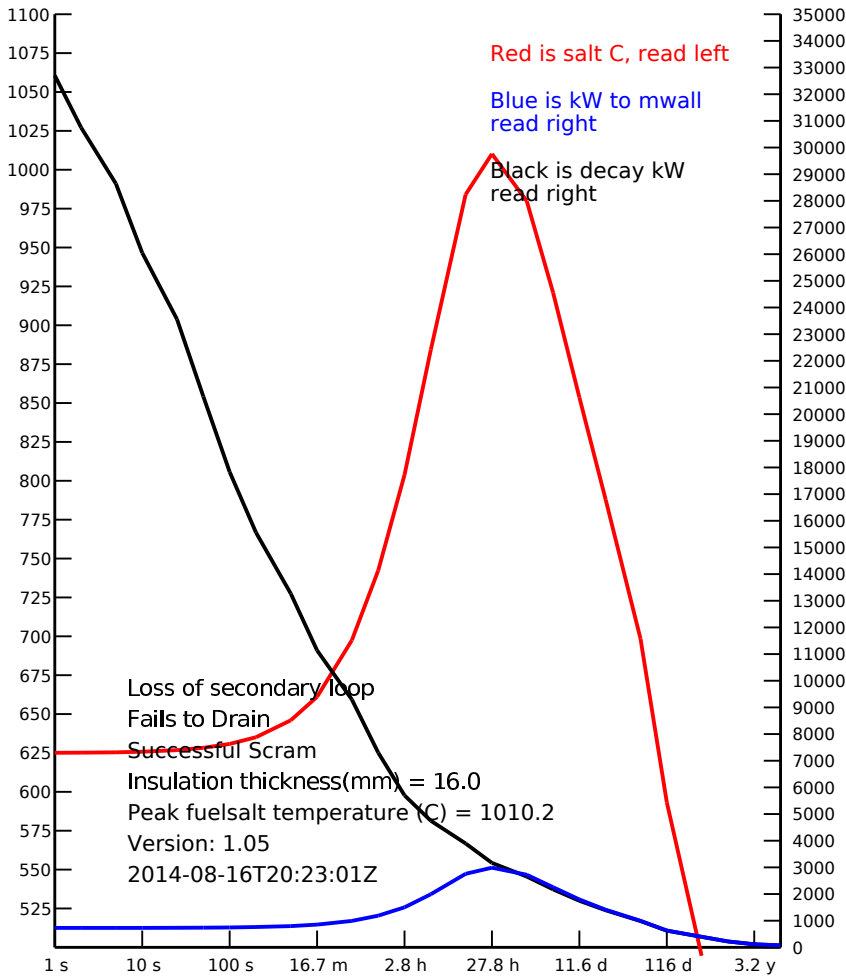


Figure 14: Response to Loss of Heat Sink, No Drain, with Scram

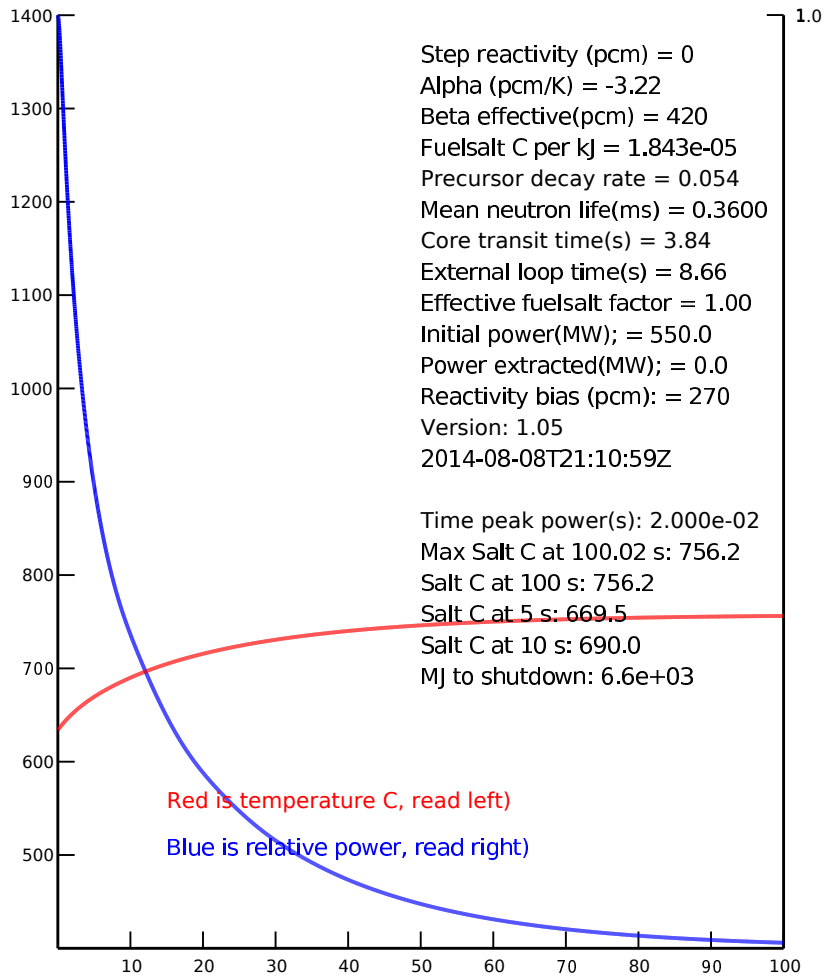


Figure 15: Shutdown after loss of heat sink, no drain, no scram

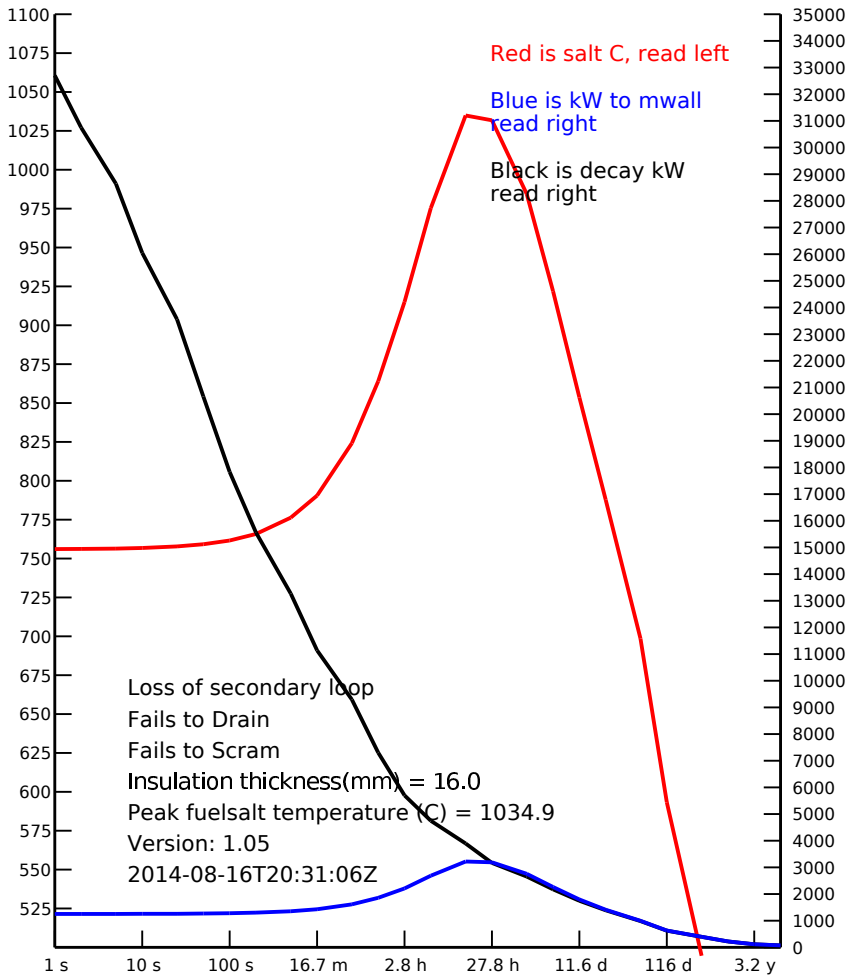


Figure 16: Response to Loss of Heat Sink, No Drain, NO Scram

3.10 The Silo Cavity

The Can/silo interface provides an obvious point at which to provide seismic isolation for the primary loop. Figure 9 indicates how both the Can and the Can membrane wall are hung from the module grid structure and are free to expand downward. The Can proper is flanged to the Can lid shown in plan view in Figure 17. The Can lid is a 50 mm thick disk whose diameter is slightly larger than the membrane wall top header. This disk serves as part of the gamma shield. Three beams forming six radial spokes are welded to the top of the Can lid. Each of the spoke tips rests on a lead core, elastomeric bearing.

The position of these bearings is indicated by the dashed circles in Figure 17. Figure 18 shows a detail of the can lid support. In the ThorCon, seismic isolation comes almost for free.

Figure 9 also shows the silo hall floor radiation barrier which is a 3.0m deep, square *radtank* filled with borated water above each Can. The radtank sits atop the module grid structure. Borated water is a more effective, more easily replaced, neutron shield than borated concrete and takes advantage of Admiral Rickover's observation that "water has no cracks". A 270 mm layer of lead in the radtank bottom is the primary gamma shield. The 50 mm thick radtank top, which is also the silo hall deck, serves as a secondary gamma shield. Personnel have full access to the silo hall, except during Can transfers.


The radtank serves as part of the *silo cavity*. The area beneath the radtank, within the grid beams, and between the Can and the silo wall, the silo cavity, is gas tight and inerted. Thus, ThorCon has a total of four barriers between the fuelsalt and the environment:

1. The primary loop piping.
2. The Can/FDT.
3. The Silo Cavity.
4. The Silo Hall.

The Fuelsalt Drain Tank sits on its own elastomeric bearings. The connection between the Can and the FDT from bottom to top consists of a bellows, a FDT shut off valve, a trunnion operated, clamp-style coupling, and a shut off valve attached to the Can. The Can and the FDT are nearly a single entity as far as an earthquake is concerned. The drain line rupture casualty is largely eliminated.


When it is time to change out a cooled Can, the radtank is drained and lifted away. The valve/coupling actuator lines are connected above the Can, the shut-off valves closed, and then the coupling opened. The Can is lifted off its seismic bearings, and moved to a transfer pit. When the Can is lifted out of the silo, both it and the FDT are sealed, nearly eliminating particle born contamination.

Figure 17: Can and Lid Structure Looking Down



Can OD(m):	7.259	Can lid disk OD(m):	8.415
Can wall thickness(mm):	25.00	Lid thickness(mm):	50.00
Can lid beam length(m):	9.988	Lid flange bolt PR(m):	3.730
Lid beam depth(mm):	800.00	Lid Bearing PR(m):	4.705
Beam flange width(mm):	450.00	Version:	1.08
Beam flange thick(mm):	50.00	2014-10-21T16:41:37Z	
Lid beam web thick(mm):	21.00		

Figure 18: Detail of Can Lid and Seismic Isolation



3.11 Steam Generating Cell

The ThorCon is a four loop design. Heat is transferred from the fuelsalt in the primary loop in the Can, to a secondary salt, and then to a tertiary salt, and finally to the steam cycle. The secondary and tertiary loops for each module are contained in the Steam Generating Cell (SGC) as shown in Figures 5, and 19. In Figure 4, the secondary loop is shown in green, the tertiary loop in magenta, and the super-critical steam loop in red.

The SGC is divided into thirds:

1. The Secondary Heat Exchanger (SHX) cell at the bottom.
2. The Steam cell in the the middle, which contains the steam generator and the steam reheater.
3. The SGC crane hall at the top.

The decks between the three cells serve both as structural tiebars and radiation barriers. As tiebars they cut the SGC wall spans by a factor of three relative to the silo hall. This means the SGC can withstand an internal over-pressure of 5 bar using essentially the same scantlings as the silo hall. This in turn implies that the SGC cell has the capability of containing a massive rupture in the 260 bar super-critical steam loop within the SGC.

ThorCon can run on a range of fuelsalts. The baseline ThorCon uses eutectic $\text{NaF} \cdot \text{BeF}_2 \cdot \text{ThF}_4 \cdot \text{UF}_4$ for the fuelsalt. $\text{NaF} \cdot \text{BeF}_2$ is called *nabe*. The baseline ThorCon fuelsalt is nabe spiked with a 82%/R18Barren nabe is used in the secondary loop. The ThorCon can tolerate minor leaks between the primary and secondary loops.¹⁵

Nabe is not the best salt neutronically. It will absorb about 6% of our neutrons as opposed to 1.5% for salts based on highly enriched ^7Li . But nabe is available, reasonably cheap, and, according to our neutronics calculations, plenty good enough for a thorium converter. Nabe increases our fuel costs but that increase is much less than the cost of a lithium based salt if it were available, which it is not.

The tertiary loop salt is 55/45 $\text{NaNO}_3 \cdot \text{KNO}_3$, commonly known as solar salt from its use as an energy storage medium in solar plants. The primary function of this loop is tritium capture. The ThorCon is designed to capture most of the tritium in the Can itself. But at this point, we can't be sure how effective this will be. Almost all the tritium which escapes the Can tritium removal system will end up in the tertiary loop where it will be oxidized to tritiated water.

The baseline ThorCon will release less tritium than a PWR. Also the low freezing point of solar salt (220C) allows us to use existing super-critical steam cycles without the need for any special feedwater pre-heating. Thanks to the tertiary loop, a steam leak in the steam generator will not produce any toxic, corrosive chemical reactions. The tertiary loop costs us about 2.5% in efficiency, which increases the capital cost of the plant by about 4% for the same electrical output.

Finally, the tertiary loop allows us to place a totally passive barrier between the supercritical steam and the secondary salt. Perhaps the worst ThorCon casualty would be a massive rupture in the steam generator which over-pressurizes the tertiary loop which in turn creates a tube rupture in the SHX which over-pressurizes and ruptures the PHX tubes. This scenario would create a toxic, radioactive mixture of high pressure steam and fluoride fuelsalt. To prevent this, the tertiary loop is fitted with an open standpipe shown dotted in Figure 4. A standpipe is the simplest, surest form of a pressure relief device. In the event of a rupture in the steam generator, the steam /solar salt mixture will simply exhaust out the standpipe and into a quench tank at the bottom of the SGC.

All four loops are arranged so that the level at which they receive heat is well below the level at which they reject heat. This is the reason that the PHX is raised above the Pot in the Can. In the event of

¹⁵ The secondary loop is at a higher pressure than the primary loop, so any leaks will be into the primary. This is the opposite of PWR.

loss of power, there will be sufficient natural circulation through these loops to remove the decay heat without dumping the fuelsalt to the integral drain tank. ***The ThorCon has two nearly independent, totally passive paths for removing decay heat.***

All four loops including their heat exchangers and pumps are constructed of SUS316Ti, a standard stainless steel. This largely eliminates galvanic corrosion. Most molten salt designs envision using a low chromium, high nickel specialty steel called Alloy N for the surfaces contacting the fluoride salts. But after the MSRE, ORNL did a series of experiments with stainless steel and fluorides salts. Some of these tests ran for 45,000 hours (5 plus years). The tests showed that, provided the salt was free of impurities and maintained in a reducing condition, the effective SUS316 corrosion rate was about 0.025 mm/year.¹⁶ In other words, ***an extra 1 mm thickness is worth about 40 years.***

SUS316 is much more radiation resistant than high nickel alloys. Alloy N requires specialized fabrication techniques and frequent re-annealing since it work hardens quickly. SUS316 is more easily worked and the skills required are widespread. Alloy N is a specialty steel with only a handful of suppliers, and possibly long lead times. SUS316 is a standard steel with many suppliers and stockists. SUS316 is far cheaper than Alloy N and the cost is much more predictable.

But at the end of the day, the reason why ThorCon is able to use SUS316 in this very demanding environment is that ThorCon has the ability to replace everything easily. In fact, in our costing we are assuming we replace the entire Can including the primary loop every four operating years. We expect to recycle the Can and the primary loop four or more times; but we only need these components to last four years.

3.12 The Pot

Figures 20 and 21 show the reactor proper. It is a simple cylindrical vessel with ellipsoidal heads, The reactor core comprises 84 vertical, hexagonal graphite *logs*, surrounding a central hexagon which contains the control rods. The logs are the moderator which must be replaced every four years.

The core is encircled by the radial reflector. The radial reflector consists of 36 graphite wedges. Each pair of wedges is connected on their outer periphery by a *shield* segment. The shield segments, outlined in blue in Figure 20, contain pockets filled with boron carbide powder. These pockets are shown in light blue. The boron in these pockets reduces the neutron flux at the wall by about a factor of 100. This means we can be confident not only that the walls will last four years but also that we will be able to recycle the Pot vessel many times. However, in our costing, we assume the entire Can must be replaced every four years.

The peripheral shield segments are interlocked with the reflector via *tabs*, jigsaw puzzle-like T-shaped knobs which fit into matching *slots* in the reflector. Thus the shield ties all the graphite together. Between the shield and the Pot wall is an annulus through which fuelsalt flows to cool the Pot wall. Our neutronics calculations indicate that the reflector graphite will be able to be reused for at least three 4 year cycles, but in our costing we are assuming it is replaced every four years.

Notches on the outside ends of each shield segment fit into axial SUS316 ribs welded to the Pot wall. They prevent the graphite from rotating while allowing the Pot wall to expand both radially and axially relative to the graphite.

¹⁶ Effective corrosion refers to the depth of voidage due to loss of chromium from the surface. The weight loss corrosion rate is much smaller.

Figure 19: Steam generating cell looking west

Version 1.08 at 2014-10-21T16:41:48Z




	Figure 20: Plan view of Pot at Mid-Height	
Number of fuelsalt logs:	84	Hot Pot vessel ID(mm): 4860.96
Salt Volume in logs(m3):	4.495	Cold salt annulus width(mm): 5.00
Moderator kg:	66089	Hot salt annulus width(mm): 25.24
Side reflector kg:	42197	Salt volume in annulus(m3): 1.75
Shield kg:	10402	Shield thickness(mm): 100.00
B4C Fraction:	0.100	




Figure 22 is an isometric view of the log. The design is based on vertical graphite *slabs*. These slabs are fitted with vertical ridges or *nubs* separating them from neighboring slabs and providing the salt channels or *slots*. The slabs are arranged around a Y-shaped yoke to form a hexagonal *log* about 230 mm on a side. Figure 23 shows a plan view of the arrangement. The dark gray is the yoke; the light gray are the slabs; the light gray protrusions are the nubs. Figure 24 shows the core log in profile.

The slabs are attached to the central yoke by top and bottom collars and dowels. Guide rods protrude from the both ends of the log. These rods fit into sockets in the top and bottom reflectors. A hexagonal *hat* of solid graphite is fixed to each upper guide rod. These hats form a shield for the top reflector.

The Pot design must not only accommodate differential thermal expansion ***but also the contraction and subsequent expansion of the graphite due to radiation damage which will be different in different parts of the Pot.*** The segments of the bottom reflector are tied together so that the entire bottom reflector expands and contracts radially as if it were a single piece. The bottom reflector rests on 36 radial ribs welded to the Pot bottom. It is anchored radially at the center of the Pot by a vertical protrusion from these ribs. But the bottom reflector is not connected to the Pot bottom. Above the bottom reflector is an array of hexagonal graphite *plugs*, one for each log. These plugs have a central socket into which the bottom log guide rod fits. The plugs are attached to the bottom reflector from the bottom at the plug/log centerline. Thus each plug is free to contract toward its log centerline. The lower log guide rods fit into sockets in the plugs.

The top reflector is similar to the bottom reflector. It too expands and contracts as a single piece. It is anchored at the Pot lid centerline by a column around the control log, but it is not attached to the Pot lid. It has a series of sockets into which the upper log guide rods and lifting lugs fit. The hats on the upper guide rods form an inner reflector as do the plugs at the bottom. The whole assembly is free to move vertically.

When the Pot is empty, the bottom reflector with its plugs rests on the Pot bottom ribs, the logs drop down and rest on the plugs, and the top reflector rests on the top of the logs. When the Pot is full, the top reflector floats up against the Pot lid ribs, the logs float up in the top reflector sockets and the bottom reflector floats up against the logs.

This arrangement:

1. Allows the graphite and the Pot vessel to expand/contract totally independently.
2. Allows each log to expand/contract axially independently.
3. Shields the bottom and top reflector from radiation contraction. Thanks to the log hats and the plugs, the neutron flux at the top and bottom reflectors is more than 200 times smaller than the flux in the core. This means we can ignore contraction in the bottom and top reflectors and that means we can use these reflectors to keep the logs in position in plan view. If these reflectors were exposed to significant fluence, then they would be contracted when the logs returned to their original dimensions at the end of moderator life. The reflectors would attempt to crush the logs together snapping off the guide rods.
4. Each hat and plug is free to contract towards its log centerline avoiding any build up of stresses due to horizontal flux gradients. Their freedom to expand/contract vertically handles the very strong vertical flux gradient without imposing a stress.
5. The contraction shows up in a gap between neighboring logs/plugs/hats which is spread evenly throughout the core. The contraction will increase the salt/moderator ratio; but this slow change can be handled by adjusting the fuel composition. Fortunately, when the log at mid height is contracted 2%, the ends of the log will contract only about 0.5%.¹⁷ When the ends are contracted

¹⁷ This will produce flux flattening at mid-height which is what we want.




Figure 21: Pot Profile: Full Hot Condition

Figure 22: Log Isometric





Figure 23: Core Log Looking Down. This is NOT a 3-D drawing.



Hexagon Side(mm): 220.03

Log OD(mm): 440.06

Log Apothem (mm): 190.55

Log ID/width (mm): 381.10

Slab Thick (mm): 39.84

Nub height (mm): 3.80

Nub width (mm): 6.00

Fraction salt: 0.1112

Log moderator kg: 778

Log salt volume(m3): 0.0529

Slot hydraulic diameter(mm) 7.45

ThorCon Version: 1.01

2014-05-25T19:12:31Z

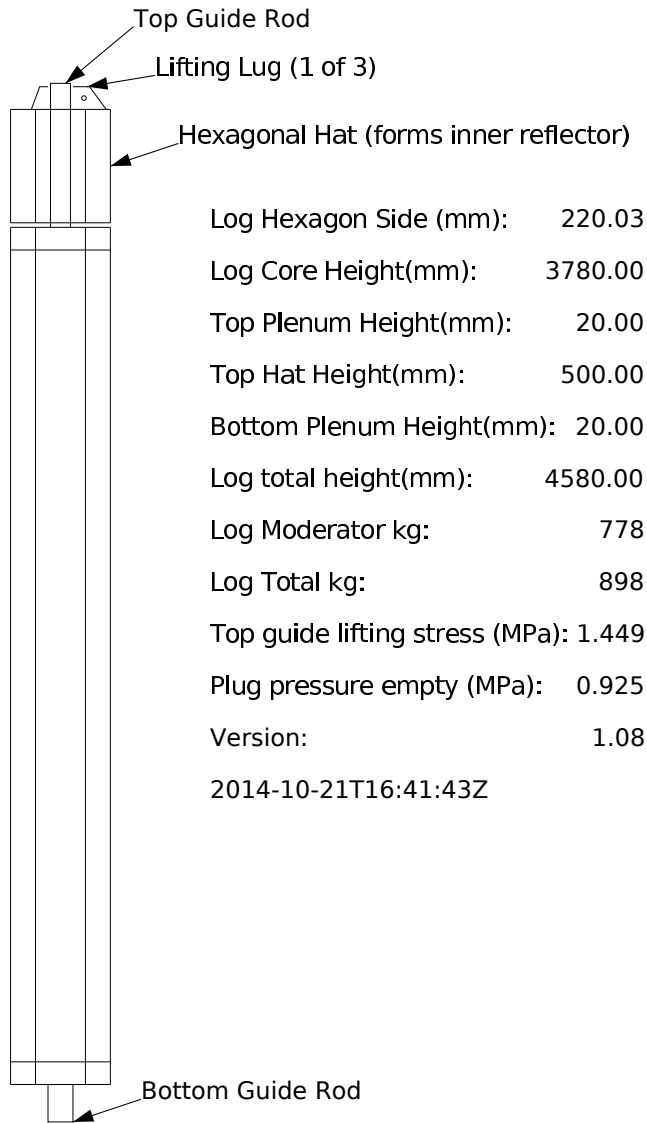


Figure 24: Core Log Profile

about 2%, the mid-height will be nearly back to original. So little orificing is required.

6. Makes assembly and disassembly straightforward. Assembly consists of dropping the bottom reflector in place, dropping the logs and radial reflector in place and locking them together with the shield segments, bolting the shield ring in place, dropping the top reflector in, and dropping the lid on top. Reverse to disassemble. Except for 36 vertical shield ring bolts, nothing inside the Pot has to be connected or disconnected.

Both MNCP and Serpent models of ThorCon have been created to determine the neutron flux distribution through time. Both are full 3-D models of the entire Can silo including the radtank. These models have been used to determine the design's initial fuel requirements, stability, control rod worth, moderator and reflector life, fuel addition schedule, build up of fission products and production of U-233 through time, and the decay heat curves. The results are summarized in the Design Control Document. These results have been incorporated into the plant costing. According to these models, the baseline design meets all our requirements. See the Design Control Document for a description of these models and a summary of the results.

3.13 Offgas system

A critically important advantage of a liquid fuel reactor is that the volatile fission products, xenon and krypton, are continuously removed rather than building up in the fuel elements destroying cladding and consuming neutrons. This is essential to the ThorCon's ability to burn up transuranics, to convert thorium to ^{233}U , to go for long periods without reprocessing the salt, and to dramatically reduce the amount of radioactive gases released in a containment breach. But it also means we must continuously process this small but extremely radioactive gas stream. This is the job of the offgas system.

Each ThorCon module will produce about 0.12 kg/day of Xe and Kr, generating a remarkable 1600 kW of decay heat. The ThorCon offgas system comprises six stages:

1. Extraction and 25 min holdup in the primary loop pump header tank. This step uses the same sprayer system successfully used by the MSRE to extract noble gases from the fuelsalt. An important luxury of a converter over a breeder is that a converter is not dependent on rapid removal of ^{135}Xe . A converter not only is not forced to a complicated xenon sparging system but is far more robust against unpleasant surprises. Calculations show that, even if the xenon removal rate is an order of magnitude less than we expect, the result is a modest increase in fuel costs. 530 kW is recovered this stage.
2. One hour hold up in the in-Can Offgas Recuperators. The offgas vented from the top of the sprayer system is directed to the two OffGas Recuperators (OGR's). The OGR's are in-can, offgas holdup tanks/heat exchangers which transfer the early stage offgas decay heat to the secondary salt. They also separate the entrained salt from the offgas and return this salt to the primary loop. A nice bonus is that we recover the bulk of the offgas waste heat. 410 kW is removed at this stage.

More importantly, most of the offgas radiation is confined within the Can, where there are four barriers between the offgas and the biosphere. Virtually all the ^{137}Xe and ^{90}Kr will have decayed to ^{137}Cs and ^{90}Sr before the offgas leaves the OGR. These two particularly troublesome fission products will not show up in the offgas stream. Rather they will oxidize to fluorides and dissolve into the fuelsalt. In the event of a breach of all barriers, they will stay in the salt.

The OGR's will also keep almost all the iodine in the fuelsalt. The concern here is ^{131}I . Iodine, like fluorine, is a halogen. At the redox levels at which ThorCon operates, the ratio of iodine gas to iodine dissolved is less than $1.0\text{e-}8$. The only way for ^{131}I to escape the Can is via its precursor ^{131}Te . ^{131}Te has a half-life of 25 minutes. Even if all the ^{131}Te made it into the offgas which seems

quite unlikely given a boiling point of 988C and the froth, 7/8ths would decay in the PLP header tank and OGR's. And we can expect plate out on the OGR packing as well.

3. 12 hour hold up in the Low Pressure Holdup tank. Each module is fitted with two offgas hold up tanks, shown in orange in Figure 5, These tanks are packed with high void metal packing. Both tanks are contained in membrane wall silos and are cooled by radiation to the membrane wall and natural circulation in the same manner as the Cans. The two tanks operate in series. The first tank operates at near ambient pressure and has a residence time of 12 hour. It extracts 90% of the remaining decay heat and cools the gas to less than 200C. 600 kW is removed at this stage. Both holdup tanks are fitted with 1 m3 sump tanks. They will be changed out on a 32 year cycle. Over that period about 260 kg of gas daughter fission products will build up in the LP HUP sump tank. The sump tanks contain CuF2 granules which will sequester both cesium via CsF and iodine via CuI. Among the nuclides which will be so sequestered is ^{134}Cs and ^{131}I . Preventing the release of these nuclides will drastically reduce the impact of a complete containment breach.
4. Compression and 115 hour hold up in the High Pressure Holdup tank. 65 kW is removed at this stage. This tank is essentially identical to the LP tank but operates at a higher pressure and a lower temperature. Almost all the remaining decay heat is removed during this period. This makes cryogenic cooling at the next stage possible and allows almost all the neutron absorbing ^{135}Xe to decay to the essentially stable ^{135}Cs . The gas leaving the HP HUP is passed through a uranium or titanium getter bed which captures practically all the tritium in the valuable tritide form.
5. Helium separation is accomplished in two stages. The offgas leaving each module's HP HUP is piped to the module's Cryogenic Separation Unit (CSU) shown in orange in the SW corner of Figure 12. The CSU consists of a vertical U-loop with a finned nitrogen pipe in the center of the loop. Liquid nitrogen is fed through this pipe and the offgas is fed through the finned annulus. The Xe and Kr freeze on the pipe and fins and build up. The helium passes through the loop and is returned to the module's working Can via a compressor and filter. Each module has two CSU loops. Once every 24 hours the offgas flow is switched to the other CSU loop, and the first loop is allowed to warm up, evaporating the Xe and Kr. This gas is compressed to 100 bar and fed to one of four HP storage bottles submerged in a tank between the loops. The helium that is in the CSU loop at the time of the warm up also ends up in these bottles. The time to fill a bottle is 41 days. The minimum holdup time in the module storage bottles is 130 days, This allows essentially all the ^{133}Xe (half-life 5.25 days) to decay to stable ^{133}Cs .
6. Second stage helium separation and Xe and Kr storage in high pressure bottles. The first stage separation reduces the helium mass flow by about a factor of 400. But helium is so much lighter than Xe and Kr the residual helium still takes up most of the volume in the Stage 1 bottles. This is why we need two stages of separation. After 130 days plus of holdup, the Xe and Kr and residual helium from each module is fed to a single CSU unit in the Transfer Module at the control room end of the silo hall. There the process is repeated to separate out almost all the remaining helium. The second stage cooler reduces the helium content by another factor of 200. This allows four years of Xe and Kr from a 1 GWe ThorCon to fit into ten 100 liter bottles. At this point, the only radionuclide remaining is ^{85}Kr with a 10.76 year half life. The decay heat is down to less than 2 kW and is spread fairly evenly over the bottles. These bottles can be cooled by naturally convected air. The current market price for Xe is \$1200 per kg and that for ^{85}Kr , which is used in crack detection, is \$330 per kg. Every four years the Canship returns these ten bottles to the Centralized Recycling


Facility where the xenon and krypton can be separated and sold.¹⁸

In addition, the system will have to handle decay heat from the fission product metals that do not form fluorides, and which are carried over into the offgas system. Most of this material will plate out in the Offgas recuperators and the HP Holdup tank. We are estimating this will add 200 kW to the heat load of both of these components. However, there is large amount of uncertainty here so the system is designed to handle treble this number.

3.14 Manufacturing and Erection

The ThorCon is designed so that essentially *everything* can be manufactured on a shipyard-like assembly line. The yard can either be an actual shipyard or a special purpose reactor yard. Figure 25 is a very preliminary sketch of what a starter reactor yard might look like. At full scale production, we will need one or two yards with five to 10 times this capacity.

Figure 25: Starter 10 GWe/y Yard block diagram, 200,000 tons steel per year



A good shipyard will require about 5 man-hours to cut, weld, coat, and erect a ton of hull steel. The yards achieve this productivity by block construction. Sub-assemblies are produced on a panel line, and combined into fully coated blocks with piping, wiring, HVAC (and scaffolding if required) pre-installed. In the last step, the blocks, weighing as much as 600 tons, are dropped into place in an immense building dock. **The ThorCon uses exactly the same production process except the blocks are barged to the site and dropped into place.** The essential difference between shipyards and most other assembly lines, such as aircraft manufacturing, is that shipyards build blocks on the assembly line, not

¹⁸ The supply of xenon and krypton from ThorCons will undoubtedly push the price down; but, from the point of view of society as a whole, this is a plus.

the final product. The final product is put together elsewhere. ***Thinking in terms of blocks rather than final product is a key element in the ThorCon philosophy.***

But ThorCon's structure is far simpler and much more repetitive than a ship's. The silo hall, SGC's, and SGC crane hall will use concrete-filled, steel plate, sandwich walls as shown in Figure 6. This results in a strong, air-tight, ductile wall. The wall cells are 25 mm thick steel plate about 1 m on a side. Each wall block consists of 8 to 15 cells. Each power module requires 7 such wall blocks as shown in Figure 26.¹⁹ The walls will be put together on a panel line in the yard as approximately 150 ton blocks. A similar structure is used for the roof. A 1GWe ThorCon will require about 17,000 tons of steel for silo hall, SGC crane hall, and SGC cells, all simple flat plate. A properly implemented panel line will be able to produce wall blocks using less than 2 man-hours per ton of steel.

Upon arrival at the site, the wall blocks will be dropped into place and welded together using the automatic hull welding machines the yards have developed for this purpose. The wall cells will then be filled with concrete. Nil form work is required.

Similarly the silo membrane walls, rad tanks, etc will be manufactured on an assembly line. Each module will require a total of 31 blocks.


The grid block, Figure 12, is the single most important block. Almost all the plant's piping and much of the wiring is in the grid. All this will be pre-installed and tested in the yard in assembly line mode. Extensive inspection and testing at the sub-assembly and block levels is an essential part of the yard's productivity. Inspection at the block level is easy. Defects and faults are caught early and can be corrected far more easily than after erection. In most cases, they will have no impact on the overall project schedule.

The grid block is about 21 m wide, 28 m long and max 4.8 m high. This block is barge transportable well up just about every major river in the world.²⁰ All the other blocks are narrower. However, the baseline ThorCon will move Cans around in the vertical position which will require an air draft of about 15 m. This exceeds the height limit on most European rivers.

¹⁹ The end module needs 2 extra blocks. The Transfer Module requires 8 wall blocks. A 1GWe ThorCon will be made up of 38 wall blocks.

²⁰ ThorCon will use ocean going barges such as the Marmac 20 class (76.2 m length x 22.95 m beam) which can carry 5000 tons on 3.9 m draft.

Figure 26: PMOD wall blocks



4 Plant Economics

4.1 Should-cost versus did-cost

It is received wisdom that nuclear energy is costly. This was not always the case. In the 1960's and early 1970's, most USA nuclear plants were built at an overnight cost of less than one thousand 2013 dollars per kW.²¹ Most of these plants were about 500 MWe and the total construction time was 4 to 5 years. But in the mid-70's real cost started rising, and in the early-mid 80's after Three Mile Island they exploded. At the same time, construction periods stretched out to 8 or 10 or more years. Orders stopped after Three Mile Island in 1979, but the last plants to be completed in the late 1980's had capital costs in excess of ten thousand 2013 dollars per kW. Vogtle 3 and 4 currently under construction will probably exceed this number. Obviously, nuclear is now expensive.

But should it be? Economists tell us that in reasonably competitive market: multiple providers, nil price power, no big secrets, no major barriers to entry, no big externalities, market cost measures the value of the resources consumed by an activity. This is the *should-cost*. In situations where these conditions do not apply, there can be a titanic discrepancy between the should-cost and the dollars actually expended, the *did-cost*.

Consider the difference between commercial ships and naval vessels. Table 5 compares a 360,000 ton displacement Very Large Crude Carrier(VLCC) with the US Navy LPD class. The VLCC can carry 320,000 tons of crude oil. The far smaller LPD is a 25,000 ton transport designed to carry 700 marines and their landing craft (two air cushion vehicles) and aircraft (4 helicopters or 2 Ospreys). The LPD has one 30 mm gun, four 50-cal machine guns, and two compact RAM close-in missile launchers for armament.²²

	VLCC	LPD
Length Overall(m)	333.0	208.5
Beam(m)	60.0	31.9
Full Load Draft(m)	22.0	7.0
Displacement(mt)	360,000	25,300
Accommodations	40	1002
Power	1 x 35MW	2 x 15MW
Speed	16kt	(flank) 22kt
Cargo capacity	350,000m3	2229m2+2190m3
Ballast capacity	150,000m3	abt 5000m3
Construction time	1yr	3 to 8 yrs
Cost	\$80,000,000	\$1,700,000,000

Table 5: Comparison of VLCC and LPD

The VLCC is 14 times larger and 20 times cheaper.²³ VLCC contracts are fixed price usually with

²¹ To convert from dollars at time of build to 2013, I have used the Consumer Price Index, which probably overstates the cost of the early plants since inflation in most material costs has been less than the CPI inflation.

²² Each RAM launcher weighs about 6000 kg and costs \$440,000 exclusive of pre-launch target detection.

²³ The price of a VLCC varies with the market. During a tanker market boom, the price can rise to 120 million or more. During a slump, it will drop to about 60 million which is about the yard's marginal cost of building the ship. A good yard can very profitably build a VLCC for 80 million dollars.

stiff penalties if the ship is not delivered within a few weeks of the target date.

Of course, the VLCC was not built with the same stringent quality control backed up by extensive paperwork as the naval ship. As a result, on average a VLCC will experience involuntary offhire time of about 15 days per year. This includes a two week dry docking every 5 years. Most ships do better than 15 days, but some VLCC's don't live up to this standard. A VLCC that has more than 30 days offhire per year in the first 15 years of her life is regarded to be a lemon. She will probably cost the yard a customer.

In contrast, LPD availability reflects the kind of standards that can be expected when obscene amounts of taxpayer money are applied to the problem. Table 6 shows a bit of the history of the lead ship, the San Antonio, LPD-17. The performance of the eight sister ships has not been much better. They were all delivered late and have experienced essentially the same set of problems. Availability, generously defined, has been in the 50% to 60% range. The initial cost per ship has remained at over 1.5 billion (Navy numbers), despite the fact that multi-ship contracts were supposed to reap economics of scale. Both admirals who were responsible for the LPD program were promoted.

If the job of building a 22 knot, 25,000 ton ship capable of carrying 700 marines, a couple of helicopters and a couple of air cushion vehicles were put out for competitive bid to the the world shipyards, I am quite confident the price would come in under 50 million dollars, quite possibly well under. Here we have an example of a difference of a factor of 30 between should-cost and did-cost. And the ships would perform per spec.

4.2 Resource Requirements

The whole point of the longish introduction to costing is that in situations where there may be a divergence between should-cost and did-cost, we must look at the resources required: the steel, the nickel, the concrete, the graphite, etc. One of the functions of the ThorCon DNA model is to accumulate material requirements by sub-system. This is done recursively in a manner which allows the designer to drill down as far as desired. Table 7 shows only the top level, nuclear island only, for a four module, 1GWe plant. This table is at least 80% complete.

These are not big numbers. The ship in Figure 1 required 67,000 tons of steel. She took less than a year to build and cost 89 million dollars. As we have seen, a 1 GWe ThorCon silo hall would fit inside her center tanks. Clearly, if we are going to make a comparison, we need a smaller ship. Perhaps a 150,000 deadweight ton, Suezmax tanker will do. The Suexmax requires more steel (23,000 vs 15,000 tons) and is larger overall (270 m by 50 m by 23 m versus 150 x 30 x 34). The ship's structure is far more complex and subject to tougher loads. The Suezmax has far more coated surface. The Suexmax has a 17 MW low-speed diesel, three 1 MW diesel generators, five 1 MW pumps, two 20 ton cranes, etc. **A good shipyard can profitably build a Suexmax tanker for sixty million dollars.** The ThorCon has sixteen 1 to 2 MW pumps. Other than the two cranes, these are the ThorCon's only large moving parts. The Suezmax has hundreds, if not thousands, of large moving parts.

Of course, the ThorCon has a bunch of stuff the Suexmax does not. Table 8 is a list of the most important. It appears that a generous first guess — **assuming shipyard-like productivity** — at what a 1GWe ThorCon should cost is about 200 million. This is nuclear side only, and excludes the Can Recycling Plant and the Can Ship(s).

Table 9 compares the ThorCon's steel and concrete requirements versus other sources of electrical power per average MW produced. The ThorCon non-nuclear numbers are from ORNL-TM-4515 which are for a 1000 MWe PWR low temperature steam plant. The 40% more efficient ThorCon plant will be smaller.

- 1996-12** Contract awarded. The budgetted cost of the ship is \$617 million.
- 2000-08** Construction started. Supposed to be commissioned 2002-07. Navy admits cost is now up to \$861 million. CBO estimates cost at 1.3 billion.
- 2003-07** San Antonio launched.
- 2004-12** Towed from Avondale to Pascagoula. Could not move under own power despite being christened in 2003.
- 2005-??** Attempted sea trials. Navy came up with 15,000 deficiencies. Some of these were major enough to compromise watertight integrity.
- 2006-01** Inexplicably Navy accepts ship waiving the unresolved issues. She is commissioned, but still can't deploy. Northrop-Grumman gets extra money "for post-shakedown availability".
- 2007-03** Failed to finish sea trials, complete failure of one steering system, major defects found in 3 of 17 sub-systems. Ship is now 840 million dollars over budget.
- 2007-06** SecNav Winter writes builder "23 months after commissioning of LPD 17, the Navy still does not have a mission capable ship".
- 2008-08** After a further series of problems and legal wrangling between Navy and builder, San Antonio finally deployed on first mission in late August, 2008. Most sources put the total taxpayer cost at 1.5 billion or higher. Some say 1.7 billion, one says 1.8 billion. Navy itself says cost may go to 1.85 billion. Stern gate failure delays departure 2 days.
- 2008-10** Got as far as Bahrain in October. Extensive oil leaks. 30 welders and fitters flown out from USA for at least two weeks of repairs.
- 2008-11** All four main engines out of commission.
- 2009-02** During transit of Suez, one screw suddenly went into reverse, sending the ship out of control and aground.
- 2009-??** Ship's XO Sean Kearns refuses Captain's mast, is court-martialed, and then acquitted after testifying that ship officers had been pressured to declare the ship was ready to deploy when she wasn't. Defense provided copious evidence supporting claim.
- 2009-07** Inspections reveal that 300 m of piping must be replaced. Reduction gear shavings found in main engines.
- 2010-03** San Antonio to Norfolk for 4-5 month overhaul costing 5 million. But inspectors finds bolts in the main engine foundation improperly installed, extensive bearing damage. Problems include bent crankshaft. Repairs now expected to take about 11 months and cost at least \$30 million.
- 2011-04** San Antonio still in repair. Navy starts an investigation into "issues with the San Antonio". Maintenance firm Earl Industries fired. Earl had won the 75 million dollar contract despite not being low bidder on the basis of "exceptional" performance on past contracts. Earl still has USN carrier maintenance contracts.
- 2011-05** San Antonio leaves yard, and after trials declared ready for duty.
- 2011-07** Unable to maintain full power. Returns to yard for repairs.
- 2012-03** San Antonio given the Navy's Battle Effectiveness Award, beating out four of her sisterships. Gets to paints a big E on super-structure.

Table 6: Abbreviated History of LPD Lead Ship, San Antonio

Table 7: 4 Module ThorCon Top-Level Resource Requirements, Nuclear Side Only

ThorCon Version 1.08	2014-10-21T12:41:25Z
Plant lead(mt)	2474
Plant tzm(mt)	4
Plant graphite(mt)	1300
Plant hay230(mt)	188
Plant imtp(mt)	127
Plant graphite_rings(mt)	57
Plant sus316(mt)	1427
Plant steel(mt)	16771
Plant c_c(mt)	2
Plant Ni(mt)	77
Plant sus304(mt)	758
Plant conc_hitemp(m3)	2332
Plant excavation(m3)	215084
Plant concrete(m3)	42131
Plant iron(mt)	12778

Overall on the nuclear side, the ThorCon has about a 2:1 advantage in steel and at least a 5:1 advantage in concrete over its nuclear competitors, despite its use of the structurally inefficient sandwich walls to avoid the crippling problems associated with reinforced concrete.

ThorCon and coal use the same steam side. But a 1000 MWe coal plant has to handle and burn 10,000 tons of coal per day. It has to attempt to handle at least 1000 tons of ash per day. Even if we are unconcerned about the 26,000 tons per day of CO₂, the sheer bulk of the fuel requires a coal plant to use about twice as much steel and concrete as ThorCon. When it comes to initial resource consumption, only a gas plant beats ThorCon.

Table 8: Adjustments

1500 tons synthetic graphite	20 usd/kg	30,000,000
500 ton silo hall crane		20,000,000
1300 tons of SUS316	6 usd/kg	7,800,000
2,500 tons of lead	3 usd/kg	7,500,000
40,000 m3 of poured concrete	160 usd/m3	6,400,000
12,900 tons of pig iron	0.5 usd/kg	6,500,000
220 tons of Haynes 230	25 usd/kg	5,500,000
500 ton crawler crane		5,000,000
200,000 m3 of excavation	25 usd/m3	5,000,000
800 tons of SUS304	4 usd/kg	3,200,000
5 tons of carbon-carbon	360 usd/kg	1,800,000
80 tons of Nickel	20 usd/kg	1,600,000
2,300 m3 of erected concrete	300 usd/m3	700,000
4 tons of TZM	50 usd/kg	200,000
9 tons of boron carbide	20 usd/kg	200,000
Total		101,400,000

Table 9: Comparison with other power sources

	Steel MT/MW	Concrete M3/MW
ThorCon (90% CF)	15 + 22	42 + 43
1970 PWR (UCB)	40	190
Current PWR (NEI, MacKay)	67	520
AP1000 (Westinghouse)	42 + 22	
Wind (UCB 25%CF)	460	870
Coal (UCB)	98	160
Gas (UCB)	3.3	27

4.3 Overnight Cost

Table 11 summarizes the results of a slightly more formal CAPEX costing exercise for a four module, 1 GWe plant. For ThorCon, the Can and its contents are treated as a consumable (see Section 4.4) and is not included in Table 11. While the costing is not yet complete, we should be able to bring this plant in at under \$800 per kW, overnight excluding Cans, fuel and salt. Our target is 600 million dollars. However, these are not first-of-a-kind numbers. These cost estimates are based on having a ThorCon assembly line up and running.

A key assumption underlying these cost figures is rational regulation. There is no limit to how expensive regulatory paperwork and delays can make a project. Table 11 assumes a regulatory system similar to that in place for airplane manufacture, with the focus on exhaustive testing of prototypes, and thereafter, quality control, and product liability. The ThorCon design is conservative, robust, and far safer than the already safe LWR. These bold claims should be thoroughly tested by imposing a full range of actual casualties on the prototype, not by rerunning the same computer programs by which the plant was designed. Assuming the ThorCon passes these tests, then it's a matter of quality control, enforced primarily by manufacturer liability.

The other key assumption is assembly line production. A shipyard can build a 320,000 ton capacity tanker (VLCC) with a steel weight of over 40,000 tons for less than 500,000 man-hours in direct labor. But we are not talking about just the labor for the steel; but all direct labor including kilometers of piping (much of it large diameter), wiring, installation of 35,000 kW machinery, 250,000 m² of high tech, thick coatings, accommodations, everything. The ThorCon is designed to bring shipyard-like productivity to nuclear plant construction,

The critical path in building a greenfield ThorCon plant will be the steam side. The nuclear portion consists almost entirely of modular components built on assembly lines. A very large, complex ship can be built in less than a year. The same thing is true of a ThorCon. There is no technical reason why a ThorCon cannot be built in less than two years, once the manufacturing system is up and running. However, for costing purposes we have assumed a four year construction period. Table 10 converts overnight cost into levelized cost per kWh as a function of the initial cost and the discount rate for a 1 GWe ThorCon assuming a 90% capacity factor.

Table 10: CAPEX USD per kWh as a function of initial cost and discount rate

Plant MWe:	1000.00
Plant life:	32
Construction Period:	4.00
Capacity factor:	0.90

Initial Cost MM USD	Discount Rate		
	0.05	0.10	0.15
800	0.00692	0.01236	0.01922
1000	0.00865	0.01545	0.02403
1200	0.01038	0.01854	0.02883

Table 11: Summary of Overnight Cost, 4 Module, 1GWe Plant (Cans excluded)

ThorCon Version 1.08	Estimate	Margin	Budgetary	Number	Estimate	Budgetary
Pond	1.569	2.000	3.137	1	1.569	3.137
Pond condensers	0.233	1.300	0.303	8	1.863	2.421
Silo Hall empty	25.303	1.200	30.364	1	25.303	30.364
Silo Hall crane	14.961	1.200	17.954	1	14.961	17.954
Module grid	1.714	2.000	3.428	4	6.856	13.712
Can Silo Radtank	0.998	1.100	1.098	8	7.983	8.781
Can Silo and Membrane Wall	0.718	1.300	0.934	8	5.747	7.472
Fuelsalt Drain Tank, Heat sink	2.114	1.500	3.171	8	16.909	25.364
Offgas Holdup tanks and silos	0.429	1.300	0.557	8	3.428	4.457
PLP motor/impeller	0.300	2.000	0.600	5	1.500	3.000
Secondary Loop Pump	0.595	1.300	0.774	5	2.975	3.868
Steam Generating Cell, empty	3.035	1.300	3.945	4	12.139	15.781
Secondary Heat Exchangers	0.221	1.500	0.332	4	0.885	1.327
Tertiary Heat Exchangers	0.267	1.500	0.401	4	1.070	1.604
Tertiary Loop Pump	0.443	1.300	0.575	5	2.213	2.877
Coolant Salt Drain tank						
Non-Module grid	1.714	2.000	3.428	2	3.428	6.856
Transfer Pits						
Off-gas clean-up system						
Salt processing system						
Crawler Crane	5.000	1.200	6.000	1	5.000	6.000
Control Room						
Total nuclear plant					113.830	154.975
Land						
Steam Plant Buildings	23.140	1.025	23.718	2	46.280	47.437
Feedwater System	37.866	1.025	38.813	2	75.732	77.625
Steam Piping	24.432	1.025	25.043	2	48.864	50.086
Turbine incl foundation	59.010	1.025	60.485	2	118.020	120.970
Condenser	8.609	1.025	8.824	2	17.218	17.648
Cooling towers	5.685	1.025	5.828	4	22.742	23.311
Cooling other	18.612	1.025	19.077	2	37.224	38.155
Electrical	38.555	1.025	39.519	2	77.110	79.038
Instrumentation	16.548	1.025	16.962	2	33.096	33.923
Total steam and electrical					476.286	488.193
Total plant					590.116	643.168

4.4 Can Cost

For ThorCon, the cost of the Can is a consumable. Table 12 displays our estimate of the material cost of a Can.

ThorCon Version 0.99		2014-04-29T11:14:44Z	
Can Cost		Code: 22223000	
	Amount	Unit Cost	Cost
Can c_c(kg)	598	500.00	299,152
Can tzm(kg)	519	50.00	25,972
Can graphite(kg)	184689	20.00	3,693,782
Can sus304(kg)	94928	4.00	379,713
Can graphite_rings(kg)	3550	9.00	31,946
Can sus316(kg)	86261	6.00	517,563
PLP Pump(kW)	2213	75.00	165,961
Heating Tape(m2)	168	1600.00	40,000
Aerogel Insulation(m2)	168	63.00	10,553
Account Can total			5,164,647

Table 12: Can Material Cost

The weight of the empty Can is about 400 tons. 162 tons of this is graphite blocks/logs that just have to be dropped into place. If we assumed shipyard productivity, a Can would require about 3000 man-hours. The Can is simpler and far more repetitive than a ship. But this is nuclear; so we dumbly multiply this by a factor of 20. 60,000 man-hours at a generous 50 USD/mh fully built up is 3 million dollars.

A reasonable initial goal here would be 6,000 man-hours with a target after learning of beating the shipyard number. The Can and its contents will require more thorough inspection and testing than a ship. But at the sub-assembly, assembly, and block level such testing and inspection is not difficult and in many cases can be automated or semi-automated. In a rational world, a factor of two over shipyard labor requirements would be a gross over-estimate.

We have to transport the Can to the plant, and 8 years later transport it to the Can Recycling Plant. The levelized cost of this function is about 250,000 dollars per round trip.

Finally, we recycle what we can, sell what we can as scrap, and dispose of the rest.²⁴ Looking at Table 12, a reasonable goal here would be a net salvage value of 5 million dollars per Can. If so a recycle plant handling one Can per week, would generate 250 million dollars of value per year. Even a quite expensive facility would net a large portion of this. But for now we assume recycling operates at a loss of 5 million dollars per Can.

We then round up the overall figure to a total “overnight” cost of 15 million dollars ignoring the fact that half of the transport costs and all the recycling cost occurs 7 or more years after the Can is delivered to the plant. A properly functioning system, after settling in, should be able to cut this number in half.

Anyway assuming a net overnight cost of 15 million USD, a four module plant with one initial spare, a four full power year Can life, a 90% capacity factor, 10% discount rate, a 32 year plant life, and we start off by providing each module with two Cans and everytime we pull out a Can we provide the plant with a new Can, the levelized Can cost is about 0.34 cents/kWh. Table 13 shows an idealized schedule. We ignore the (probably imaginary) wind down period at the end, providing four Cans in the middle of year 31 even though we are assuming nil remaining plant life. Table 13 makes the point that Can cost is not a really strong function of discount rate.

²⁴ Essentially all the graphite will be recycled, grinding it to powder and starting from scratch where necessary.

Number of modules:					4
Plant full power MW:					1000.0
Can overnight USD:					15,000,000
Can full power years:					4.00
Capacity factor:					0.90
Plant economic life:					32.0

Month	Outlay 1.0e6 USD	PV COST 0% 1.0e-6 USD	PV Cost 5% 1.0e6 USD	PV Cost 10% 1.0e6 USD	PV Cost 15% 1.0e6 USD
0	135.0	135.0	135.0	135.0	135.0
53	60.0	195.0	183.4	174.4	167.4
107	60.0	255.0	222.2	200.0	184.6
161	60.0	315.0	253.4	216.7	193.8
215	60.0	375.0	278.4	227.6	198.7
269	60.0	435.0	298.5	234.7	201.3
323	60.0	495.0	314.6	239.3	202.7
377	60.0	555.0	327.6	242.3	203.5

Levelized USD/kWh		0.00244	0.00285	0.00340	0.00403
-------------------	--	---------	---------	---------	---------

Table 13: Idealized Can Replacement Schedule and Cash Flows

4.5 Fuel Costs

Essentially all the cost of the fuel is in the ^{235}U . The ^{235}U requirements can be broken down into the end period inventory (loop inventory plus FDT inventory), cool down inventory, and the net additions. The second column of Table 14 shows the mass fraction of ^{235}U required at the beginning of each 8 year period per module.²⁵ This costing is doubly conservative in that it assumes a new fuel charge every 8 years. In reality, the uranium for the second, third and fourth “new” charge will be processed old fuel.²⁶ But the average annual fuel additions are based on reuse which means they get larger in each succeeding period. Table 14 also assumes that none of the plutonium or other transuranics is returned to the plants. Eventually, we expect to reinject and burn most of the TRU.

At the end of each period we will have collected an inventory of ^{235}U in the FDT from the need to make room for the additions. The third column shows the amount of ^{235}U pushed out of the primary loop annually. Roughly speaking this is about one-third the net additions.

On average every plant will need another one-third of this inventory to account for the cool down period between years 8 and 12 in each cycle. This assumes we have enough plants staggered so that this inventory can be shared.

Finally, we have the net additions. These have been taken directly from the Serpent burn up runs. Under these assumptions, the unit fuel cost is about 0.61 cents per kWh. It is a weak function of the discount rate.

Under these assumptions, ThorCon is not all that ^{235}U efficient. Over 32 years, we feed about 12,000 kg of ^{235}U to a module and end up with a 5,000 kg of unused ^{235}U enriched to 9%. Currently for costing purpose we are putting zero value on this inventory. 9% is highly enriched by normal standards. And there is still some thorium left. We can get at least another 8 year cycle out of this salt.

More basically, we need to burn this ^{235}U . There are a number of options, but ideally we’d re-enrich. Separating the uranium from everything else is straightforward and the result is UF6 which is the form the enrichers require. The ^{235}U content could be boosted back to 20% at about 48 SWU per kg ^{235}U . Currently, an SWU costs \$90 and this is likely to drop to less than \$50 in the near future. Obtaining ^{235}U for \$2500 to \$5000 per kg is remarkably cheap.

The problem is this uranium will contain 10 to 20 ppm U-232. But down the road if there is enough of this material, a means will be found to recover the ^{235}U . Laser enrichment is particularly attractive in this regard because it will avoid enriching the U-232.

²⁵ Table 14 is based on full power years (capacity factor = 1.0) to match the Serpent numbers. Lower capacity factors will stretch this out with nil change in the unit fuel costs.

²⁶ This assumes rapid expansion in the ThorCon fleet since the inventory at the end of each 8 year period is close to double the amount required at period start.

ThorCon Version:: 1.08		2014-10-19T20:05:52Z			
Discount rate (pct):	5.00	Capacity factor:	1.00		
Fuel U235 percent:	20.000	Feed U235 percent:	0.720		
Tails U235 percent:	0.20	U3O8 USD per kg:	66.139		
USD per SWU:	90.00	U3O8 to UF6 per kg:	7.50		
U3O8 USD per kg of U235:			14849.12		
Conversion USD per kg of U235:			1427.88		
SWU USD per kg of U235:			20417.37		
U235 USD per kg:			36694.38		
Thorium USD per kg:			50.00		
Year	New Loop Inventory U-235 kg	New FDT Inventory U-235 kg	New Cool. Inventory U-235 kg	Net Additions U-235 kg	New Thorium Th-232 kg
0	636	41	0	153	11964
1	0	41	0	153	0
2	0	41	0	153	0
3	0	41	0	153	0
4	0	41	0	153	0
5	0	41	0	153	0
6	0	41	0	153	0
7	0	41	0	153	0
8	636	52	320	174	11964
9	0	52	0	174	0
10	0	52	0	174	0
11	0	52	0	174	0
12	0	52	0	174	0
13	0	52	0	174	0
14	0	52	0	174	0
15	0	52	0	174	0
16	786	73	350	185	11964
17	0	73	0	185	0
18	0	73	0	185	0
19	0	73	0	185	0
20	0	73	0	185	0
21	0	73	0	185	0
22	0	73	0	185	0
23	0	73	0	185	0
24	895	89	456	190	11964
25	0	89	0	190	0
26	0	89	0	190	0
27	0	89	0	190	0
28	0	89	0	190	0
29	0	89	0	190	0
30	0	89	0	190	0
31	0	89	0	190	0
U235 PV million USD:					218.776
Thorium PV million USD:					1.463
Levelized USD per kWh:					0.00606
Gross value of final 4985 kgs of U235 million USD:					182.93

Table 14: Per module fuelsalt ^{235}U requirements with no reuse

Table 15: Levelized Cost of Salt excluding fuel

Version: 0.85		2013-07-29T15:04:35Z		
Full power output (kW):		250000.0		
Capacity factor:		0.90		
Fuel salt kg:	48434	\$/kg: 35.00	Yrs: 8	
Secondary salt kg:	65034	\$/kg: 35.00	Yrs: 32	
Tertiary salt kg:	50000	\$/kg: 0.50	Yrs: 32	
Discount	Fuel salt	Secondary salt	Tertiary salt	Total
Rate pct	cents/kWh	cents/kWh	cents/kWh	cents/kWh
0.0	0.01075	0.00361	0.00004	0.01440
5.0	0.01267	0.00696	0.00008	0.01971
10.0	0.01466	0.01102	0.00012	0.02580
15.0	0.01667	0.01524	0.00017	0.03207

4.6 Salt Cost

A 1 GWe ThorCon will require about 450 tons of the fluoride salt, nabe, which is NaF-BeF₂ (57/43 mol pct). Just about all the cost of nabe is in the beryllium. Nabe is 46% BeF₂ by weight. A conservative estimate of the cost of BeF₂ is \$75 per kg or \$375 per kg of Be.²⁷ You can buy toothpaste quality sodium fluoride off the web for less than a dollar per kg. Our current estimate of the cost of nabe is \$35.00 per kg. Each module will require about \$4,200,000 worth of nabe, about 40% of which will be fuelsalt,

ThorCon operates on a fuel cycle of 8 years at which point fission product and trifluoride build up will require us to replace the fuelsalt. The U, Np, and eventually Pu will be extracted and fed back into the reactor. We should be able to recover this salt but for costing we are assuming we will have to buy new fuelsalt every 8 years.

We will also need about 50 tons of *solar salt*, 55:45 NaNO₂-KNO₂ for the tertiary loop. This salt, whose name comes from its use in solar energy storage tanks, costs about 50 cents per kg.

Table 15 converts the salt expense to cost per kilowatt-hour as a function of discount rate. This table does NOT include the uranium and thorium in the fuelsalt. Unit salt cost will be less than 0.03 cents per kWh. This is almost in the noise. The ThorCon's use of nabe means that salt costs are a non-factor. If we screw up and have to dispose of some salt, we can afford to do so. This is most certainly not the case for flibe, even if it were available.

There is little direct incentive to recover the nabe in the fuelsalt. However, we do it anyway to reduce the volume of radioactive waste, that has to be stored until the fission products have decayed to background levels.

²⁷ Beryllium in the form of BeF₂ is cheaper than beryllium metal. BeF₂ is a precursor of Be metal avoiding the expensive reduction step.

4.7 Staff cost

Table 16 shows a preliminary estimate of the staffing for a four module ThorCon. The main difference between this table and staff numbers for other advanced reactor concepts is that ThorCon has far fewer people in what is variously called Engineering, Configuration Control, and Work Control. These people design and implement corrections and improvements to the plant, manage refueling outages, and see that all this work is properly documented. For a standard US nuclear plant, this can be more than 300 people.

Administration and Training	25
Operations	42
Maintenance	30
Janitorial, grounds, cafeteria	10
Security, Loss Prevention	72
Engineering, work control	10
Trainees	20
Total	209

Table 16: Four Module On-site staff

ThorCon has no refueling outages.²⁸ All the plant does is load and empty the fuelsalt casks on a 8 year interval. Almost all the work associated with changing out a Can is offloaded to the Can Recycling Plant. The on-site work is limited to taking the new Cans off the Canship and loading the old Cans on. More importantly, any plant modifications are designed, implemented, overseen, and documented by the central ThorCon design office. Once again the airplane model, rather than the USA NPP model. The job of the small on-site engineering staff is to report problems, suggest improvements to the central office, and monitor any changes after they have been implemented. The bottom line is that a four module ThorCon can easily be staffed with about 200 people.

The annual payroll will depend critically on where the plant is located. For now we conservatively assume, an average fully built up cost of \$150,000 per man-year. For a four module plant and a capacity factor of 0.90, the corresponding levelized cost is about 0.5 cents per kWh.

4.8 Total Cost

Putting this all together, and assuming a 10% real discount rate we end with Table 17 which displays the cost of electricity in dollars per kWh as a function of the overnight cost.

Even at the higher initial costs and a generous cost of capital — at least once the technology is proven — the ThorCon will be competitive with everything with the possible exception of gas at below \$3/MMBtu and zero CO2 price. In areas where either very cheap gas is not available or carbon capture is required, the ThorCon reigns supreme. If otherwise unusable super-critical steam plants are available, the cost of ThorCon power drops to about 2 cents per kWh. The ThorCon is cheap.

Since the ThorCon is modular, it is subject to weaker economies of scale than a standard monolithic reactor. Our costing indicates that a one module, 250 MWe plant will have an overnight cost that is 35% that of the large plant. This will add less than a cent per kWh to the levelized costs. There will be some diseconomies in plant staffing but they will add less than 0.5 cents to the levelized cost. A 250 MWe ThorCon can produce power at less than 5 cents per kWh. This broadens the market for the ThorCon

²⁸ The turbine(s) must be taken off-line for 10 to 14 days every four years for maintenance and overhaul. This outage will be largely managed by the central office.

Version: 1.04	2014-07-20T14:51:02Z		
Plant MWe:	1000.00		
Plant life:	32		
Construction Period:	4.00		
Capacity factor:	0.90		
Cost of capital:	0.10		
Overnight Cost	\$800 MM	\$1000 MM	\$1200 MM
Unit Capex	0.01236	0.01545	0.01854
Unit Can Cost	0.00340	0.00340	0.00340
Unit Fuel Cost	0.00599	0.00599	0.00599
Unit Salt Cost	0.00020	0.00020	0.00020
Unit staff	0.00493	0.00493	0.00493
Unit waste	0.00100	0.00100	0.00100
Total \$/kWh	0.02788	0.03097	0.03406

Table 17: Total Levelized Cost USD per kWh, 1 GWe ThorCon

to small island nations that are now paying 25 to 40 cents per kWh for their electricity. In this regard, it is important to note that this is **dispatchable** power. ThorCon by design is a load follower. ThorCon can function both as base load and, within limits, as a peaking plant.²⁹

More generally, in looking at levelized cost numbers such as Table 17, it is extremely important to compare like with like. Not all forms of energy are created equal as far as the grid is concerned. At the top end are highly reliable and highly dispatchable sources of energy, power that can both be counted on and adjusted quickly to meet changes in demand. Hydropower and aircraft derivative gas turbines are examples. At the other extreme are energy sources that are neither reliable nor dispatchable. Since such sources require full back up, to have a net positive economic effect, the cost of such electricity must be less than the value of the fuel they save, including any requirement for additional, fuel inefficient, spinning reserve.³⁰

ThorCon is extremely reliable and reasonably dispatchable. It sits near the top end of this spectrum.

²⁹ Of course, as the capacity factor drops the LCOE will go up. Since the Can is a consumable, ThorCon has an unusually high marginal to fixed cost ratio for a nuclear power plant. About 1/3 of ThorCon's unit cost is marginal, which ameliorates a drop in capacity factor.

In certain unusual situations, it could make sense to set up ThorCon with a dedicated peaking capability. One way would be to use solar salt from the tertiary loop to store energy. The simplest solution would be a two tank system. During low demand periods, a portion of the solar salt from the SHX would be directed to the hot tank being replaced by salt from the cold tank. In high demand periods, the hot tank salt would be pumped through a steam generator and that steam fed to a peaking turbine, the salt ending up in the cold tank.

³⁰ This assumes the grid is not mandated to take power from such sources when it is uneconomic to do so. If this is not the case, such sources exacerbate the demand variations imposed on the rest of the grid sometimes forcing the grid to replace less expensive, less CO2 intensive sources with more expensive, more intensive.

Unreliable, non-dispatchable power can be transformed to reliable, dispatchable by providing enough storage. But if so, the cost of that storage must be included in the levelized cost.

4.9 Market

The market for reliable, dispatchable electricity at these prices is immense. This reflects the enormity of the energy problem facing the planet. The USA EIA predicts that world net electricity consumption will go from a rate of 2500 GW in 2013 to 4500 GW in 2040.³¹ This is roughly in line with the Hargraves argument. Hargraves points out that, in order to induce world population to level off at about 9 billion in 2050, we need to lift per capita income to about \$20,000 yr.³² This corresponds to an average per capita energy consumption which is about half current USA consumption. Even if the US and other high consumers cut their consumption to this level — an unlikely scenario — the growth required to bring the rest of the world up to that level will result in more than doubling energy consumption by 2050.

According to the above projections, over the next 25 years, the world will need something like one hundred 1 GWe plants per year, about two plants per week just for new electrical generation. As things stand now, most of these plants will be coal fired. Indeed, as of June, 2013, some 1,199 coal plants are planned worldwide, with a nameplate capacity of 1,401 GWe.³³ Flipping half these plants to ThorCon would require five reactor yards like Figure 25, and represent a 50 billion dollar per year business, just for the yards.³⁴ Since ThorCon is designed to be coupled to a standard super-critical steam plant, the option of replacing the boilers in the coal plants that are being built today becomes viable further down the road.

At 3 cents/kWh, replacing fossile fuels in transportation via electrification or synthetic fuels becomes a real possibility. In this case, as Shu points out, the requirement becomes more than one 1GWe plant per day.³⁵

³¹ EIA, International Oil Outlook, 2013.

³² Hargraves, R. Thorium: Energy cheaper than coal, p 150-151.

³³ Technology Review, Vol 116, No 3. as of June, 2013,

³⁴ This will require an expansion in enrichment capability. But if the demand is there so will the supply. Urenco has indicated they can have a 20% LEU enrichment plant operating four years after the decision to make the investment is made.

³⁵ Shu, F., Global Change and the Energy Crisis, UCSD, 2013. www.physics.ucsd.edu/~tmurphy/phys239/shu.energy.pdf.