Optimal Performance for a Fusion-Neutron-Driven Commercial Waste Transmutation Facility

Charles D. Bowman, ADNA Corporation, Los Alamos, NM 87544
cbowman@cybermesa.com

Abstract

A thermal-spectrum graphite-moderated molten-salt subcritical facility for transmutation of nuclear waste from commercial power is described. The facility is designed to accommodate a horizontal cylindrical or line source of neutrons produced either by proton spallation or by fusion. The major objectives of this design are: (1) economic practicality achieved by cheaper structural materials and the minimization or elimination of reprocessing, (2) delay of the need for geologic storage of high-level commercial waste by several hundred years, and (3) lower radioactive toxicity for the eventual stored waste remnant. The spallation neutron drive consumes about 9 % of the electric power generated from fission in the system and the cost per neutron is 74 MeV of electric energy. If the power consumption for a fusion neutron drive is limited to 9 % also, the corresponding cost of d-t fusion neutrons is shown to be 185 MeV. This is 250 times more than the energy that may be expended per d-t fusion for a pure fusion electric power plant and therefore sets the bar for economically useful fusion lower by 250 for a fusion-driven transmuter than that for a pure fusion power plant. In addition to transmutation, the paper describes potential advantages of fusion neutrons for generating fission electric power from thorium and for combined thorium power generation and LWR waste burning

Table of Contents

Introduction

The interest in transmutation of waste is now more than ten years old and the process has been studied at a significant level in almost every nation with a commercial nuclear power program. The greatest focus has been to try to adapt fast reactor technology, which has been ready for commercial deployment for a decade. In spite of its technical maturity, fast reactor technology has not been commercially deployed anywhere owing to economic impracticality. Frequent and expensive reprocessing and expensive materials of the reactor set costs for this technology well above commercial light water reactors. In spite of much international interest in such a system for accelerator-driven transmutation, the economic prospects are even more discouraging since an expensive accelerator must be added to an already too expensive technology for transmutation. In fact it is probably safe to say that no accelerator-driven technology will be economically practical with the present conception of subcritical reactor, accelerator, and reprocessing plant. One of these components must go if economic practicality is to be achieved. The ADNA Corporation believes that the external neutron source can be used, in addition to reducing the waste toxicity and other undesirable properties, to essentially eliminate the reprocessing component from transmutation technology. In effect the externally produced neutrons displace the need for recycling and reprocessing. A thermal neutron spectrum and liquid fuel enable the most complete and lowest cost means for maximizing the benefits of external neutrons[1].

While existing spallation technology using GeV protons on a liquid lead target appears to provide an economic and practical solution to the external neutron source component, the possibility for implementing fusion neutrons as the external source should be examined. Perhaps the most remarkable finding of this report, to the author, is that the energy cost of neutrons from a fusion-driven transmutation system may be higher by a factor of 250 than the allowable cost for a pure fusion electric power plant. Perhaps an existing approach to fusion technology is already capable of driving an economically practical transmuter. Furthermore, the spallation process is well understood and can be improved only modestly from the present status, while fusion technology has built-in “stretch.” If a fusion-driven system is practical with the energy cost of fusion neutrons a factor of 250 higher than that for a pure fusion system, then clearly greatly enhanced performance is possible each time one reduces the fusion neutron cost by a factor of two or three. “Stretch” is discussed at some length in this report with respect to several options for fusion-produced neutrons including the transmutation of commercial waste, the generation of virtually unlimited energy from thorium without reprocessing, and the production of energy from thorium with concurrent commercial waste burning.

This paper is a beginning effort in integrating a fusion neutron source into a subcritical transmuter. The prospects for doing this successfully depend on achieving advantageous coupling of the fusion neutrons into the transmuter, completing the full scale neutronic analysis of the several modes of operation, assuring that the mechanical engineering, and power production features of the design are practical, and designing in the required safety and non-proliferation features of the system. For example, most transmuter designs receive a beam moving vertically downward, which is a difficult geometry for a fusion device and also for a spallation source to some degree. This paper describes a design in which the cylindrical transmuter has been turned on its side so that the beam enters horizontally and perhaps from both ends on the axis. The spallation and fusion neutron approaches are synergistic in that the spallation approach for a thermal spectrum molten salt system can be tested in a demonstration facility in the near future using the Los Alamos LANSCE accelerator or other accelerator. Meanwhile a practical fusion neutron source can be developed concurrently with the transmutation demonstration and the two technologies integrated together as the next step along the path to a practical fusion-driven transmutation and fission energy production system.

Horizontal transmuter design

The plan view for a 750-MWt subcritical system is shown in Fig. 1. The system is driven by two accelerators at each end of the cylindrical device each operating at an energy of 1 GeV and a current of five megawatts beam power. Sufficient neutrons are produced to achieve the design fission power with k_{eff} = 0.96.

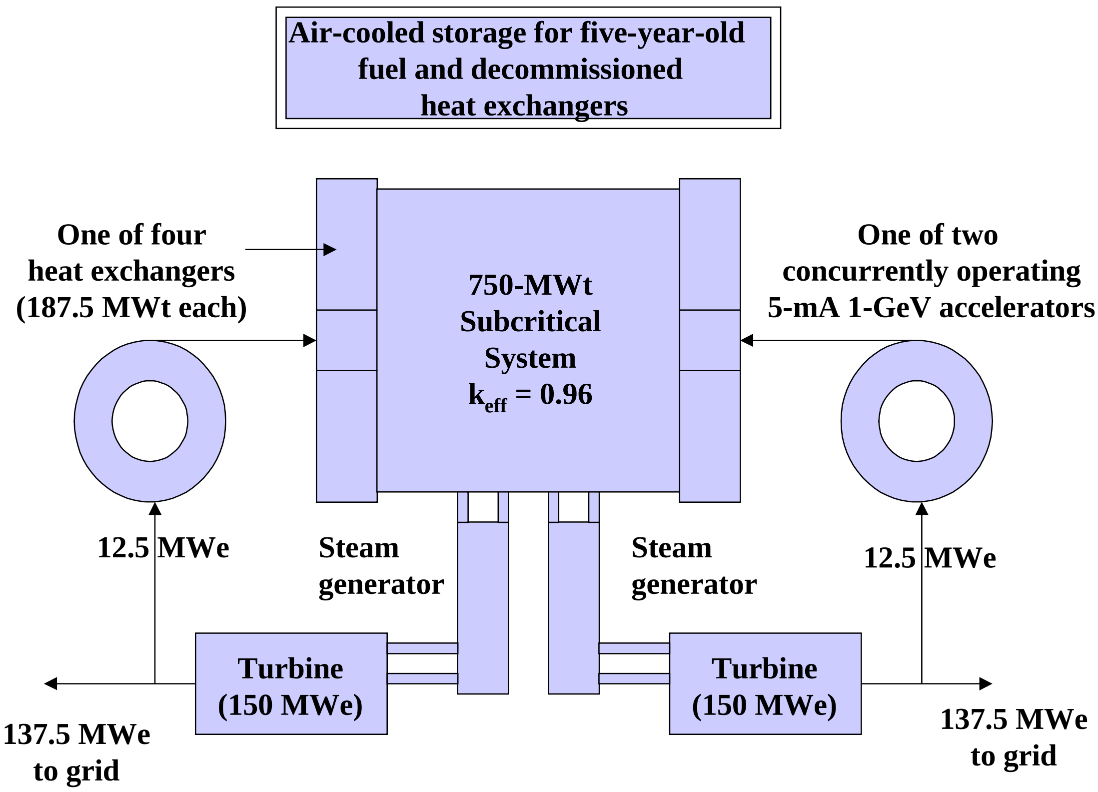

Fig. 1. A plan view of the subcritical system shown with accelerator beam entering the transmuter on the horizontal from both directions. This geometry should better enable incorporation of a fusion device for the spallation source since essentially all fusion systems exhibit horizontal design.

One of the principle problems of an accelerator-driven system is that the present rate of unanticipated beam interruptions is too high for a practical power source[2]. Since the two independent accelerators will exhibit uncorrelated beam interruptions, the power should very seldom drop by more than a factor of two. Also since most beam interruptions are short, the design includes a large inventory of secondary salt feeding the steam generators so that full electric power output can be maintained for beam interruptions lasting a few tens of seconds. It is important to point out that the beam interruption problem is much less severe for molten salt fuel, which is not damaged by rapid temperature changes in the salt, than for solid fuel. Solid fuel is susceptible to severe damage by sudden and frequent changes in fuel temperature even by only a factor of two and no solid reactor fuel has been tested under the frequent temperature cycling expected even with advanced higher reliability accelerators. The design shows two steam generators and turbines being driven from one tank of secondary coolant salt (not shown in Fig. 1) also to improve power reliability. A natural-convection air-cooled storage facility is shown for five-year-old canister-stored used fuel, for decommissioned heat exchangers, and other radioactive components.

More details of the transmuter are shown in Fig. 2. The neutron source region is on the central axis. Neutrons produced there moderate and diffuse into a surrounding

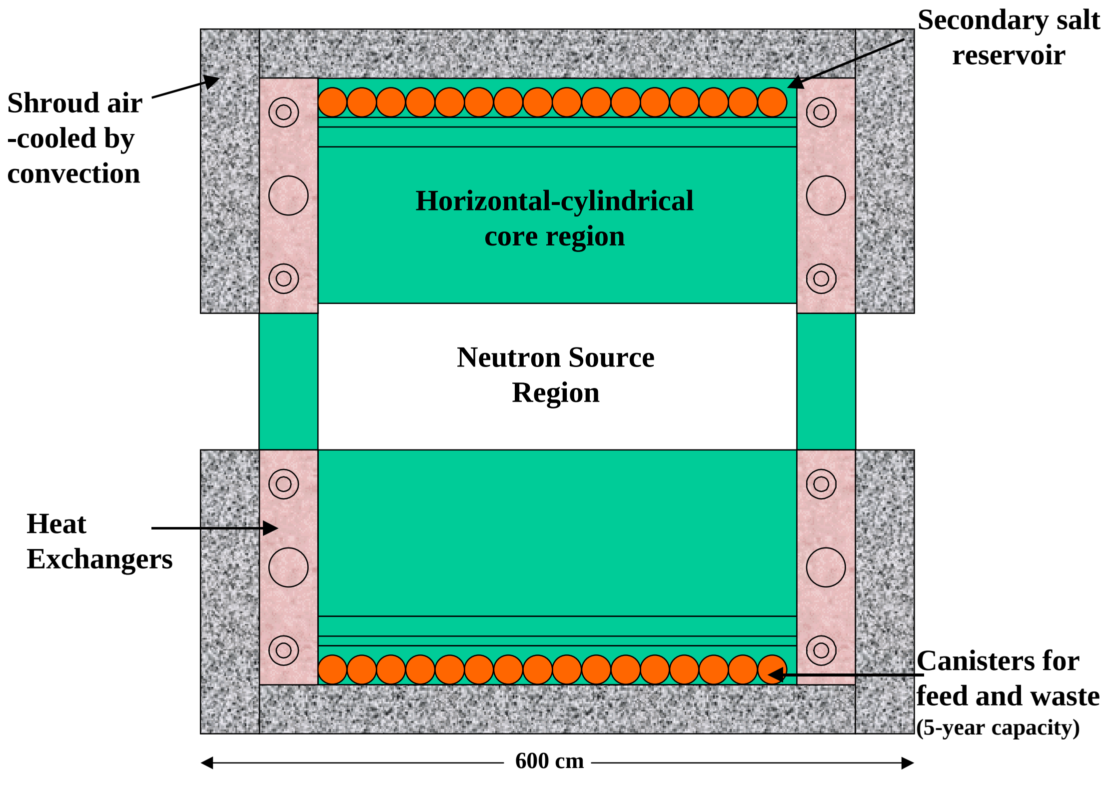

Figure 2. Plan view of horizontal transmuter including neutron source, core region, heat exchangers, secondary salt reservoir, canisters for input and output salt, and the air-cooled shroud.

graphite-loaded core region with a 14 % molten salt volume fraction. The carrier salt is a $^{7}LiF\text{-}BeF_2$ mixture in the approximate mole ratio of 67/33 with a few percent of actinide and fission product for waste burning and a melting point of about 500 C. Alternatively for energy production from thorium, the mixture is ^7LiF -ThF_4 in the mole ratio of 72/28 and a melting point of 570 C. The heat is transferred to a non-radioactive secondary salt by four internal heat exchangers using the secondary salt NaF-$NaBF_4$ in the mole ratio of 8/92 with a melting point of 385 C and an operating temperature in the range 650 to 550 C. The reservoir for the secondary salt surrounds the core with a volume sufficient to allow a few tens of seconds of electric power production should power generation be lost by beam interruption.

Canisters for input and output salt storage are immersed in the secondary salt as shown in Fig. 3. They store as much input salt and output salt as required for five years of operation using the same canisters for both input and output storage. The system is surrounded by a shroud for convection air-cooling sufficient for entirely passive fission after-heat removal.

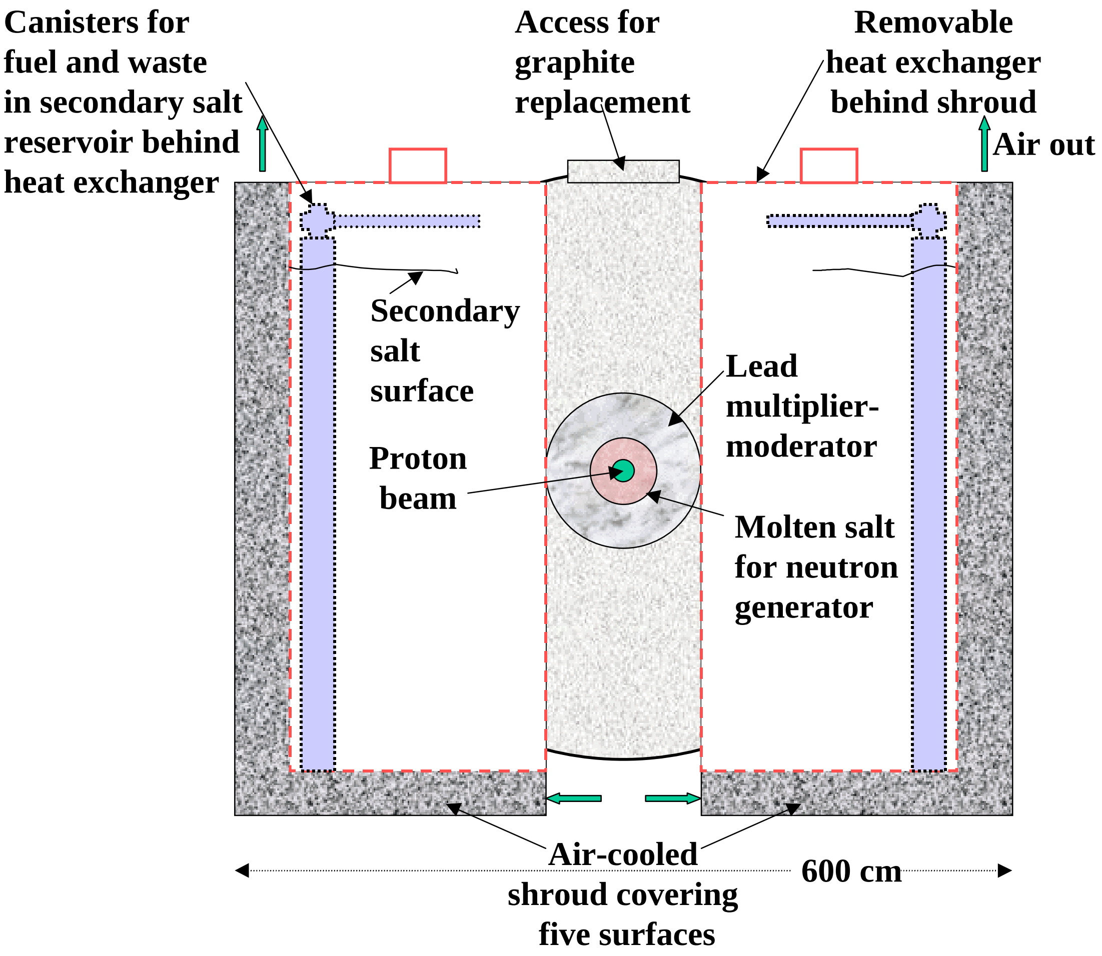

Fig. 3. End view of the horizontal transmuter showing input and output salt canisters stored in the secondary salt and the shroud. The shroud surrounds the transmuter on its diameter and the heat exchangers at the ends and it is filled with stone. Air enters through the bottom center and moves by convection upward and out.

The 50-cm thick shroud is filled with stone to enhance heat storage capacity and heat removal. The graphite would float in the surrounding salt if not constrained and it is therefore possible to remove the graphite by picking it up from the surface of the molten salt. The salt must be transferred to the storage canisters before new graphite can be installed. The graphite exchange process might be as follows:

1. Transfer the radioactive salt to the storage canisters.
2. Replace the salt in the core with fresh non-radioactive salt.
3. Remove the graphite constraint allowing the graphite to float to the surface.
4. Remove the floating graphite from the non-radioactive salt adding salt as graphite is removed to enable the graphite to float high enough for removal.
5. When graphite removal is complete, pump the non-radioactive fresh salt back into storage canisters so as to empty the core.
6. Replace the removed graphite with the core empty of salt.
7. Constrain the graphite from floating, fill with the originally removed radioactive salt, and resume equilibrium operation.

The graphite lifetime is estimated to be about 40 years. The power density averaged over the core for an average power of 750 MW thermal is 16.6 watts/cm^3. For moderation in graphite (and also salt which has about the same moderating power) the corresponding flux above 50 keV is 6.7 x 10^{13} n/cm^2-s. The better graphites[3] can withstand a fluence of 3 x 10^{22} n/cm^2 above 50 keV before the linear swelling of the graphite reaches 1 %. The corresponding graphite lifetime in this flux would be 14 years. However the above estimate assumes that all the neutron slowing down takes place in graphite, which is the case for most graphite reactor designs such as the HTGR or the pebble bed. However, in this system, which is heterogeneous with salt flowing through 14-cm diameter channels in the graphite, things are somewhat different. All fission neutrons are born in the molten salt, which cannot be structurally damaged by the neutrons. The neutrons lose about 2/3 of their energy in the salt so that only 1/3 is left for the graphite and the graphite lifetime is therefore enhanced by a factor of 3 to about 40 years. Of course a prudent design should allow simple replacement of the graphite as described above as actual operating experience might show the need for more frequent replacement.

The proton beam enters horizontally from both ends through a Leiss[4] window as shown in Fig. 4 where it is dispersed before passing through a conical surface into the molten salt where the beam stops after producing neutrons. The neutron production per proton in the light-element salt is much lower than in lead but the neutron energy is much larger. These higher energy spallation neutrons move into the surrounding molten lead and multiply there with the end result being the same neutron production per proton as if the beam had struck the lead directly (about 30 neutrons per 1-GeV proton). The advantage of this concept is that most of the proton beam power is deposited into the salt. The smaller energy that the lead receives can be transferred to the salt through conduction eliminating the need for a separate cooling system for lead. The lead therefore only moves by passive internal thermal convection. Since the lead is not directly struck by the beam, there is very little spallation or fission of the lead as (n,xn) reactions dominate so the lead does not build up a high radioactivity or a large inventory of other elements. It is expected that the lead will last for the lifetime of the transmuter, but it is designed so that it can be readily removed if necessary. For operation with a fusion source, access from both ends is a clear advantage and obviously the spallation target can be removed and replaced by the fusion source.

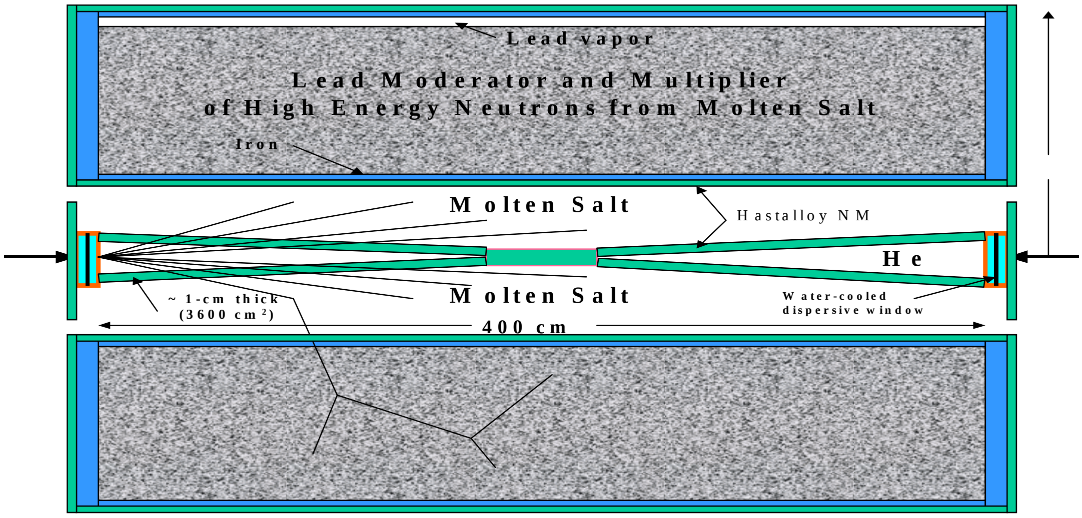

Fig. 4. The salt-lead spallation target. Beam enters from both ends of the target through a dispersive window and passes through a conical surface into the molten salt. High energy neutrons produced by the beam in the light-element salt are multiplied as shown in the bottom left by (n,xn) reactions in lead. For operation with a fusion source, access from both ends is a clear advantage and obviously the spallation target can be removed and replaced by the fusion source.

The Leiss window is shown schematically in Fig. 5. The window is a four-foil system with accelerator vacuum on one side and the target helium gas on the other side at about one atmosphere. The four foils divide the window into three chambers with water cooling flowing through the outer chambers and high pressure helium in the inner chamber to stabilize the windows against flutter. The mass of the water and the four foils is substantial so that the beam emittance is greatly broadened after exiting the window. This target design, which has been used for high power beam for many years, seems to be well suited to this particular application. The water cooling could of course be replaced by flowing lead or salt if safety required it. The lifetime of this window is estimated to be about one year for the 5 mA 1-GeV proton beam and so the window must be designed for frequent replacement.

When using liquid fuel, it is highly desirable to maintain three barriers between the fuel and the outside world. This window in combination with the conical section provides two of the barriers. The beam loading on the conical section is low from the perspective of heat load and radiation damage since the dispersed beam is spread out over the full 3600 cm^2 conical surface. The structural metal identified for the molten salt is the well studied Hastalloy N modified, which is a high nickel alloy, and therefore susceptible to damage by thermal neutrons through the (n,p and n,a) reactions. Since this surface is exposed mostly to unmoderated neutrons, the neutron damage rate to the Hastalloy NM should be acceptable from an engineering perspective, but provision should be made for replacement of this element during the life of the transmuter. Probably this can be made practical through a process similar to that described above for replacement of the graphite. For operation with a d-t fusion source, the lead would provide an estimated multiplication of about 2.5 and it should therefore be considered a vital part of the neutron source geometry. However for the lower energy d-d neutrons, the lead would provide no multiplication and would only provide neutron moderation. Therefore for the d-d source, the lead probably is unnecessary.

Leiss Window for High Power Beam Transmittal and Dispersion

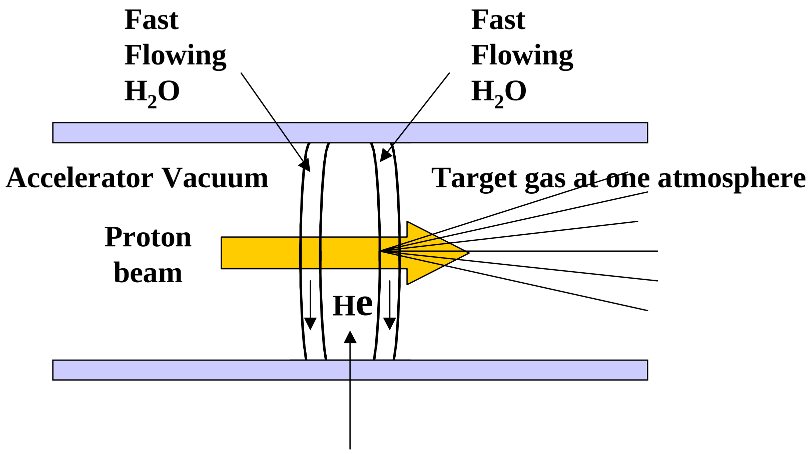

Fig. 5. The Leiss window is a four-foil system defining three volumes. Water flows through the outer two volumes providing cooling for all four foils. Static high pressure in the inner volume prevents water-induced fluttering in the foils. The substantial material in the beam with this window introduces dispersion into the beam sufficient for the beam to illuminate the full conical surface that the beam must traverse to reach the salt.

Commercial WasteTransmutation

While transmutation of commercial reactor waste is the prime objective of this workshop and this has been the prime focus of the accelerator-driven studies, there are other applications that might be more desirable for taking the fullest advantage of the nuclear energy resource. The implementation of the accelerator-driven system for three applications is first discussed and then it is shown how successful introduction of a fusion source might enhance the performance beyond that possible with the accelerator.

Fig. 6 illustrates the key[1] to maximum advantage for a liquid fuel system for transmutation and energy production. The upper portion of the figure showing the conventional approach includes an accelerator driving a transmuter with reprocessing on the back end to remove fission product and return unburned actinide. All other transmutation systems other than the one described here work this way. The lower portion shows the new approach that includes the accelerator but eliminates the expensive and troublesome reprocessing. There is a continuous flow of fresh salt bearing Pu and minor actinide to be

CONVENTIONAL APPROACH

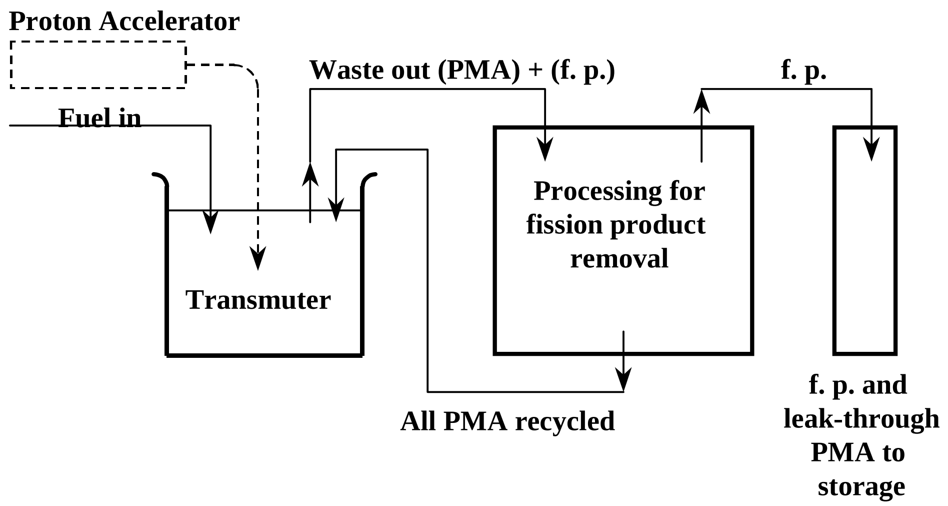

NEW APPROACH

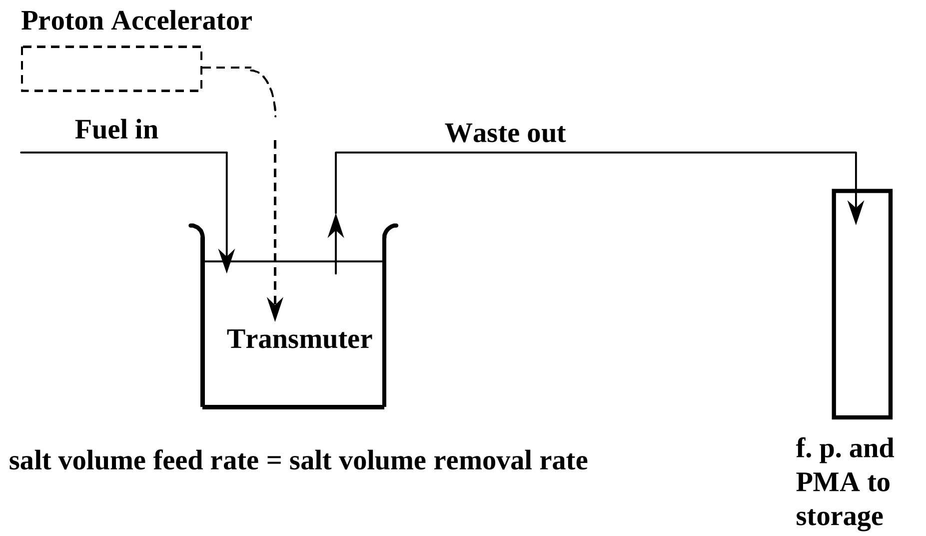

Fig. 6. The conventional and new approach to transmutation and energy production

burned and a continuous removal of salt including fission products and unburned actinide. The chemical and isotopic composition of the waste going in is very different from that coming out as shown in Fig. 7 where the back shows the feed and the middle labeled Tier-1 shows the output. (The front row labeled Tier 2 showing the results[5] of burning of the Tier 1 output in a second conventional cycle using the same transmuter is not discussed further in this report.) The objective is to accomplish enough transmutation in a single pass to make the back-end reprocessing unnecessary.

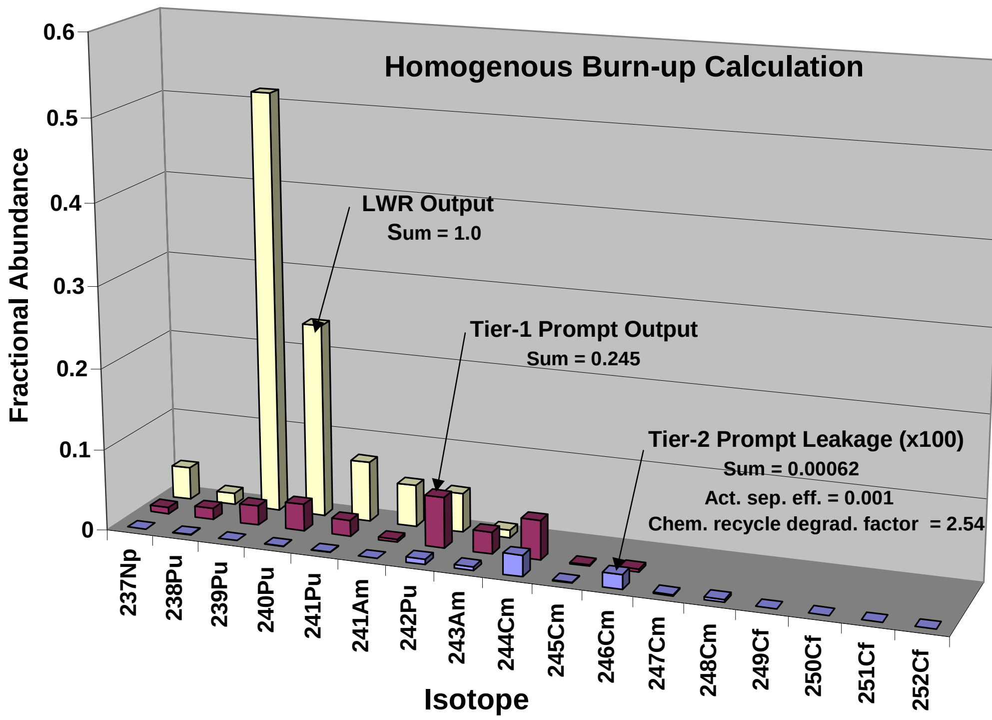

Fig. 7. Homogenous burn-up results. The data at the back shows the input isotopic distribution from LWR spent fuel. The sum of the bars adds to unity. The middle row shows the output from a single pass through the transmuter as in the bottom of Fig. 6 when the structure of the transmuter is homogenized. The bars add to 0.245 for a single pass, which far exceeds any burn-up possible for a single pass through a fast spectrum transmuter. Improved performance down to a remnant of 0.07 for the thermal spectrum has been found for a single pass from a more realistic heterogeneous calculation[6], but has not been confiormed. This degree of burn-up in a single pass might be sufficient as described in the text. If one wishes to go further in burn-up, the front row shows the results if one takes the output of the middle row and passed it through again with back-end separations.

The ATW Roadmap Report[2] emphasizes that a reduction to 10 % of the fed actinide is sufficient for the planned implementation of Yucca Mountain because for reduction below 10 % defense waste begins to dominate the leakage from the facility. The back row of Fig. 7 shows the isotopic composition of the feed actinide and the bars sum to unity. The middle row shows the output from a single pass through the transmuter as in the bottom of Fig. 6 when the structure of the transmuter is homogenized. The bars add to 0.245 for a single pass, which far exceeds any burn-up possible for a single pass through a fast spectrum transmuter. Improved performance down to a remnant of 0.07 for the thermal spectrum has been found for a single pass from a more realistic but unconfirmed heterogeneous calculation[6]. So it seems possible perhaps to reach the necessary degree of burn-up *in a single pass* with the approach shown at the bottom of Fig. 6 and a highly thermal spectrum. With 90 % of the input actinide destroyed and the remnant plutonium of no use for weapons, a repository receiving this waste need no longer be referred to as a plutonium mine that must be guarded into perpetuity.

The new approach must be brought to the equilibrium condition in which the input and output rate and the isotopic compositions are stabilized, and it requires a managed protocol to reach this equilibrium. However once in equilibrium the system exhibits two interesting features. First, the isotopic composition of the output and the mixture inside is obviously the same. Therefore the output from one system can be used to start a second transmuter that is *initially in equilibrium*. The first transmuters on line become the mothers of many daughters with the output from the mothers becoming the starting inventory for the daughters.

The second interesting feature is that the input isotopic composition is instantaneously and irreversibly transformed to the final composition as the input is fed and mixed. This is a non-proliferation advantage of this system over all others since in a treaty abrogation scenario partially burned fuel can be removed at any point and the Pu recovered for weapons use. By contrast the equilibrium Pu inventory demanded for operation of a fast spectrum transmutation system for the U. S. far exceeds the Pu presently existing in the U. S. LWR spent fuel stockpile. Every nation, even one as small as the Czech Republic, with LWRs and a fast spectrum transmutation system would own enough Pu for an inventory of weapons as large as that of the U. S. today and this situation would persist with fast spectrum transmutation for as long as nuclear power exists in its present-day LWR form. With the new approach this concern disappears.

The front row of Fig. 7 illustrates the ultimate burn-up possible[5] if one feeds the output shown in the middle row to the same system except with back-end processing to achieve the lowest possible remnant. The front row bars have been multiplied by 100 to make them readable on the same scale as the other distributions.

Improving transmutation with fusion neutrons

For the purposes of this paper, the important question is how the new approach might be improved with fusion neutrons. Obviously, if the fusion source turned out to be cheaper than an accelerator either in capital or operating costs, transmutation costs would be reduced perhaps making a deciding difference on the cost practicality of transmutation. However, the other point relates to performance enhancement. The degree of burn-up of Fig. 7 is constrained by the need to maintain k_{eff} near 0.95 so that the neutron multiplication is about 20. A significantly smaller remnant could be achieved if k_{eff} were reduced to 0.90, but the accelerator capital and operating costs doubles and operation under this condition is probably not economically practical. While one might look forward to accelerator improvements allowing practical operation at lower k_{eff}, spallation is not a new technology and there is limited room for improvement. This is not the case for the fusion source as illustrated in Table 8 comparing the implementation of spallation and fusion neutrons. The comparisons are normalized to production of 2 X 10^5 Mev of heat either by fission or by fusion. As given in the second column for fission, the number of fissions required is 1000. Assuming a thermal to electric conversion efficiency of 44 %, the electric energy generated from the 2 X 10^5 MeV of heat is 0.88 x 10^5 MeV.

Table 1. A comparison of spallation and fusion neutron source effectiveness.

	Spallation (1 GeV on Pb)	d-t Fusion Alone	d-d Fusion Alone	d-t Driven Subcritical System	d-d Driven Subcritical System
Reactions	1000 fissions	11,400 fusions	50,000 fusions	1000 fissions	1000 fissions
Heat energy generated	2 X 10^5 MeV	2 x 10^5 Mev	2 x 10^5 Mev	2 x 10^5 Mev	2 x 10^5 Mev
Electric energy generated	0.88 X 10^5 MeV	0.88 X 10^5 MeV	0.88 X 10^5 MeV	0.88 X 10^5 MeV	0.88 X 10^5 MeV
Input neutrons for 1000 fissions	104	________	_______	41.6	104
Electric energy consumed	7700 MeV	7700 MeV	7700 MeV	7700 MeV	7700 MeV
Percent of generated electric power	8.76 %	8.76 %	8.76 %	8.76 %	8.76 %
Electric energy per source neutron	74 MeV	0.69 MeV	0.15 MeV	185 MeV	74 MeV

The number of spallation neutrons required to produce the 1000 fissions must take into account the fact that only about $k_{eff}S/\nu$ of a source of S neutrons start fission chains. The multiplication factor for these fissions is $1/(1-k_{eff})$ so that $S = 1000\nu(1-k_{eff})/k_{eff}$. For $k_{eff} = 0.96$ and $\nu = 2.5$ the result is 104 neutrons to induce the 1000 fission events. Each 1-GeV proton produced by a 45 % efficient accelerator produces 30 neutrons with a cost per neutron of 73 MeV and a total electric energy cost for the 104 neutrons of 7700 MeV. This amounts to 8.76 % of the electric power generated.

The third column of Table 1 lists a similar set of parameters for pure fusion energy production via the d-t reaction. To generate the 2 X 10^5 MeV of heat energy from a d-t reaction releasing 17.6 MeV per fission requires 11, 400 fusion reactions. For this comparison, we require that the electric energy generated from this heat, the electric energy consumed in generating the fusion, and the percent of the electric energy generated be the same as for the fission process. In that case we find at the bottom of the column that one can spend no more than 0.69 MeV per fusion event. If one goes through the same process for d-d fusion, the electric energy that may be spent per fusion is only 0.15 MeV. Both are highly challenging goals that must be met for an economically deployable system.

The most interesting part of the table is in columns 5 and 6. The requirements for a subcritical system driven by d-t fusion neutrons begins with the reminder that essentially all of the energy comes from fission and that 1000 fissions are required to produce the 2 X 10^5 MeV of heat energy. With the lead multiplier providing 2.5 neutrons per 14-MeV neutron, the total number of neutrons required to drive the 1000 fissions is 41.6 and the energy that may be spent to generate each neutron under the same conditions as the accelerator-driven system is 185 MeV. This is *250 times larger* than for a practical pure d-t fusion power reactor. In terms of a practical application for d-t fusion technology, the bar is lowered by a factor of 250 if the neutrons are to drive a subcritical fission system. Put another way, if the temperature of a plasma is 50 keV, only about 0.1 % of the ions must undergo neutron-producing reactions to have a practical driver for a subcritical system. The corresponding energy expenditure for d-d fusion is 74 MeV, and the bar is lower by a factor of 500 for d-d fusion. Of course the capital cost of the fusion source also must be compared with the capital cost of the spallation source and appropriate adjustments to the above comparison made if they are not equivalent.

The above comparison is the basis for the “stretch” advantage for the fusion neutron source. At best the spallation source can be improved by perhaps a factor of two by increasing both the thermal electric conversion efficiency and the accelerator efficiency from about 0.45 to 0.60. However, we have seen above that the energy expenditure per fusion for a pure fusion device must be 250 times lower than that required for an effective subcritical system driver. If the pure fusion goal is realistic as has been assumed for the lifetime of the fusion program, then fusion neutrons can offer great advantages for fission energy as the fusion programs make progress toward the goal of pure fusion with a factor of about 250 of “stretch” built in compared to a factor of two for the accelerator. For example, if the energy requirement per fusion were reduced by technical advances from 175 MeV to 17.5 MeV, only about 1 % of the electric power would be consumed for producing neutrons. Alternatively by spending the same percent of the electric power on neutron production, ten times as many neutrons could be produced with the same overall economics as with the accelerator and k_{eff} could be reduced to 0.705. For the waste transmuter, this would mean a single-pass burn-up far better than the factor of ten reduction likely with the accelerator. This is yet another example of trading off external neutrons against back-end reprocessing chemistry. It is clear therefore that *fusion neutrons for enhancing subcritical fission system performance provide a sustained motivation for progress toward a pure fusion power plant.*

Thorium burning in subcritical systems

There are many advantages in burning thorium over uranium for nuclear energy generation. Since there is four times as much Th as U in the world, the fission energy resource can be extended by a corresponding factor four. Second, no weapons-useful material can be recovered since the burning of a small fraction of uranium along with the thorium results in denaturing the ^{233}U isotopic enrichment to levels of no use for weapons. Third weapons-related isotopic enrichment technology could be outlawed since there is no need for this technology for the Th cycle. Fourth, the weapons-related reprocessing technology can be deferred essentially indefinitely as we will see below. Fifth, the need for geologic storage may be deferred indefinitely or perhaps obviated as will also be shown. Lastly, criticality is not only unnecessary; nuclear power and criticality are fully decoupled since criticality is impossible under any circumstances..

The subcritical thorium burner offers striking advantages over a thorium reactor. To implement the thorium reactor as a breeder or even a converter requires on-line back-end reprocessing to remove the fission product poisons almost as fast as they are produced[7]. This high dependence on chemistry has been a major deterrent to a practical role for thorium burning. However in a subcritical system, this serious chemistry constraint can be eliminated as shown in Fig. 8. The external neutron source enables the system to

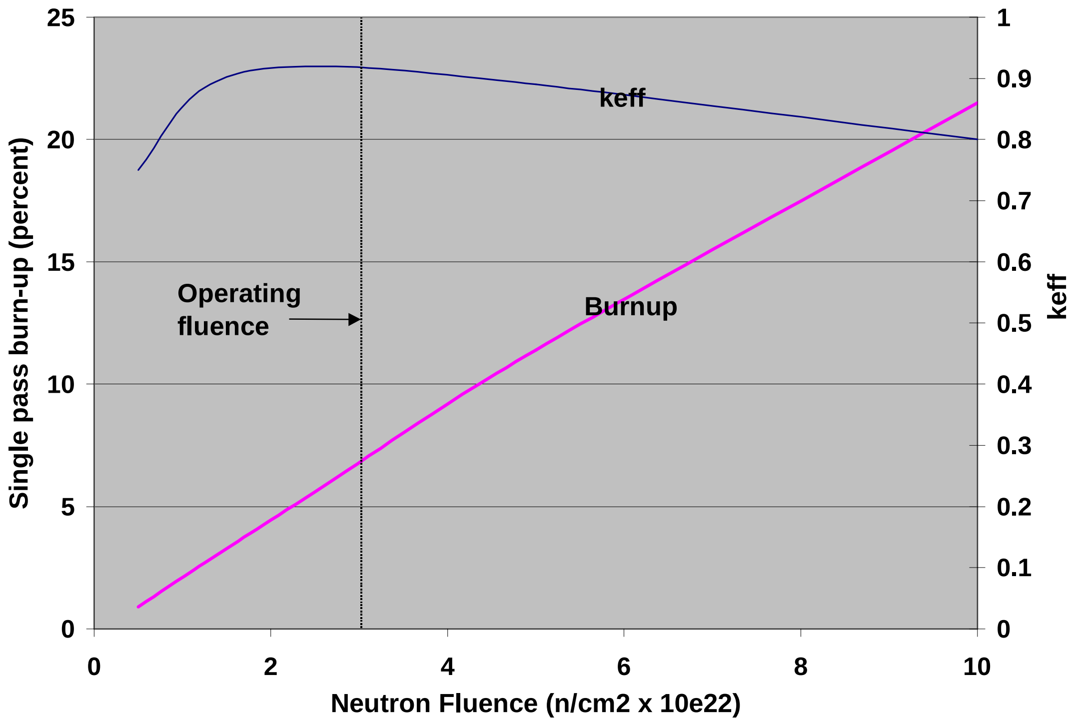

Fig. 8. Single-pass thorium burning. For the continuous feed and removal of liquid fuel, an equilibrium can be established with the fuel flow rate establishing an average fuel dwell time and an associated average neutron fluence. The value of k_{eff} is small for small dwell time or neutron fluence. A higher fluence is accompanied by a higher ^{233}U concentration and higher k_{eff}.

However as the fluence increases, the fission product poisoning effect also increases and this overpowers the effect of growth of ^{233}U. The k_{eff} curve reaches a maximum near 3 x 10^{22} n/cm^2. At this fluence the degree of single-pass burn-up of Th is 7 %. This is ten times higher burn-up than is achieved in LWRs in a single pass without reprocessing. Since there is four times as much Th as U in the world, this technology extends the nuclear energy resource by a factor of 40 beyond that available without reprocessing.

operate with k_{eff} less than unity providing an interesting opportunity. For the continuous feed and removal of liquid fuel, an equilibrium can be established allowing the fuel to flow through at a rate such that the fuel is exposed for an average dwell time corresponding to an average neutron fluence associated with a particular value of k_{eff}. The value of k_{eff} rises from zero with zero fluence and a higher fluence is accompanied by a higher ^{233}U concentration and higher k_{eff}. However as the fluence increases, the fission product poisoning effect also increases and this overpowers the effect of growth of ^{233}U. The k_{eff} curve reaches a maximum near 0.92 and a fluence 3 x 10^{22} n/cm^2. At this fluence the degree of single-pass burn-up of Th is 7 %. It should be noted here that Fig. 8 is the result of homogenous calculations and that significantly better performance should result from heterogenous calculations presently underway with k_{eff} = 0.95 for the same fluence.

For comparison the LWR burns only 0.6% of the mined uranium. Therefore this system enables ten times higher burn-up of mined thorium than achieved in an LW for burn-up of mined uranium. Since the world's thorium inventory is four times as large as that for uranium, this technology extends the nuclear energy resource by a factor of 40 beyond that available with LWR technology.

The main drive from the beginning behind fast reactor technology was the extension of the uranium resource by breeding and the implementation of reprocessing. Essentially the same objective is achieved with this technology for thorium with the accelerator extending the resource and eliminating the need for the reprocessing. For a fast breeder reactor, the plutonium doubling time is an important parameter. For a breeding ratio of 12 % per year, about six years is required for producing sufficient Pu for starting a daughter reactor. Thus over a 40-year lifetime one fast breeder reactor can become the mother of about seven daughter reactors. The equilibrium Th burner is capable of the same. Since the output salt is the same as the equilibrium salt inside, the output salt can be used to start in equilibrium a daughter subcritical thorium burner *without reprocessing*. The residence time for the salt inside a thorium burner operating with an equilibrium fluence of 3 x 10^{22} n/cm^2 and a thermal flux of 2 x 10^{14} n/cm^2-s is about five years. Every five years sufficient output salt is accumulated for starting in equilibrium a daughter subcritical thorium-burning system. So one mother thorium subcritical burner after about five years to reach equilibrium can initiate over its 40-year lifetime seven daughter subcritical thorium-burning systems in automatic initial equilibrium. Each of these daughters can initiate eight more units over their lifetime. The parallel with the fast breeder with a breeding ratio of 12 % is almost exact except that the subcritical thorium system does not require reprocessing and fuel refabrication.

Passing the output fuel to launch daughter systems not only eliminates the need for purchasing a new beginning fuel loading, it also eliminates the period for approach to equilibrium when operation must be more closely managed. Therefore until the deployment of these types of systems is saturated, the output fuel from mother units can be sold at a substantial price thus helping in overall cost reduction in the early period of deployment.

The external source also eliminates the need for enrichment, fuel fabrication, fuel refabrication, storage, an extensive nuclear transport infrastructure, and a large inventory of weapons-useful plutonium. By feeding a small amount of uranium with the thorium, the ^{233}U is denatured to uninteresting weapons material with little overall degradation in performance. This new technology therefore enables an essentially complete divorce of nuclear energy and nuclear weapons technology. Finally, all these benefits are received in a system for which a criticality accident is not only unlikely but is physically impossible. No combination of feed of natural fuels or removal of salt enables criticality. There is also no accidental reconfiguration of the fuel that would result in criticality since dispersal of the salt in graphite enables the closest approach to criticality and dispersal in natural material or concentration reconfiguration in any geometry will result in lower criticality.

Fusion neutron impact on thorium burning

As in the case of waste transmutation, the fusion neutron source can have an enormous impact beyond that of the accelerator. Remembering that the d-t fusion equivalent energy expenditure of 185 MeV would match spallation as a external neutron source, if the cost of d-t fusion neutrons were lowered by a factor of about three, the fluence could be raised by a factor of three and the percent burn-up also increased by a factor of three to about 20 % in a single pass as seen from Fig. 8. Alternately after a first cycle of burning 7 % of the thorium in a single pass using accelerator-produced neutrons, the advent of fusion-produced neutrons at 1/3 the energy cost would allow this salt to be recycled for burning 14 % more of the thorium at $k_{eff} = 0.80$ without reprocessing. As the energy efficiency of fusion-neutron generation is increased, this can continue on to even higher burn-up without reprocessing. With a potential “stretch” of 250 for fusion, the number of recycles is almost indefinite, the need for beginning again with a pure thorium feed is postponed indefinitely, and the need for reprocessing may be postponed indefinitely. With the potential always to reuse the thorium output and the indefinite postponement of reprocessing that would give rise to a high level waste stream requiring geologic storage, the need for geologic storage of *high level* waste could be postponed indefinitely. Since the present LWRs consume only 0.6% of the uranium resource and the fusion technology integrated with equilibrium thorium burning would enable use of at least half of the thorium without reprocessing, the nuclear energy resource would be enhanced by the factor of 4 x 0.5/ 0.006 = 333. If LWRs provide 20 % of the world’s electric power needs for 50 years, this new technology might carry the world at substantially increased energy use rates perhaps for 1000 years.

Closure on this scenario comes when the pure fusion goal is reached producing neutrons at the lowest cost or the neutrons become so cheap on the way to the pure fusion goal that destroying all of the remnant long-lived fission product species such as ^{99}Tc and ^{129}I and any remaining minor actinide species is economically practical. It should be noted here that separation of long-lived fission product species is not essential. The fission products ^{129}I and ^{99}Tc are burned down by a factor of two in a thermal neutron fluence of 3 x 10^{22} n/cm^2. If these isotopes are recycled without separation several times, the over all reduction may reach a factor of ten or more, which may be sufficient depending on nuclide transport in a repository. In that case, the actinide and long-lived fission product remnant might be small enough to fall into the class of low level waste. Since criteria for low-level waste storage can be met in almost any state in the U. S., the need for a central place for storage of high level commercial nuclear could disappear. The above scenario would require interim and accessible storage of non-weapons-useful material.

In the event that a better energy source becomes available before the thorium resource is exhausted, one may always use inexpensive fusion neutrons to close out thorium-burning without leaving a long-term high-level waste stream.

Burning LWR waste with thorium

The introduction of the above thorium-burning technology would leave the LWR waste stream for resolution. This could be handled by the LWR waste transmutation approach outlined earlier, or this waste could be burned along with thorium without front-end reprocessing. It should be remembered that the graphite-moderated thermal spectrum offers better neutron economics that an LWR and therefore exhausted LWR fuel can be further burned in the graphite-moderated reactor, and even further burned in a graphite-moderated subcritical system. This is true even without any separations. That is, all of the fuel can be fed into the graphite subcritical system including the ^{238}U and the fission product. The only requirement therefore is transformation from oxide or metal to fluoride; no separations are necessary although some might naturally occur in the process of fluorination.

Since the spent LWR fuel feed with no separations is more reactive than the thorium feed, it is clear that the mixture of thorium and LWR waste will be more reactive than thorium alone. A figure similar to that of Fig. 8 is being prepared showing a maximum k_{eff} for a particular fluence and an associated burn-up percent of the thorium and the LWR fuel. A potentially higher value for k_{eff} may be used instead for burning with a higher fluence and a resulting greater burn-up per single pass. As in the previous cases, as neutrons are made cheaper by advances in fusion neutron generation technology, the fuel can be rerun through the system without back-end reprocessing. The detailed performance of this system remains to be calculated, but the performance of the similar thorium-only systems suggests that the feed fraction of commercial waste should be about 20 mole percent of the thorium feed. The presence of the waste will enable enhanced burning on the first cycle because of the plutonium and ^{235}U present, but in subsequent recycles the system will perform essentially as the thorium burner.

Next steps

It seems to be a reasonable assumption that this meeting will uncover a promising substantial role for fusion neutrons in LWR waste burning and energy production from thorium. The prospects discussed in this paper would be considerably strengthened by the following studies.

- Evaluation of the array of fusion technologies for most effective and lowest cost integration into a subcritical system
- More quantitative evaluation/comparison of spallation and fusion neutron drive
- Neutronic optimization of performance of all three systems
- Confirmation of the design of the heat exchangers, and the other components of the power-production system that feed back to neutronic design.
- More complete understanding of the disposition of fission products including the categories of noble gases, near noble metals, the alkali metals, and the lanthanides derived from the operation of a molten salt loop with non-radioactive surrogates for fission products.
- Procedures for feed and removal of salt and exchange of graphite.
- Practical integration of a fusion neutron source into a subcritcal system
- Demonstration that three-barrier containment of high-level radioactivity is practical
- Conceptual design for a deployable unit capable of operating in all three modes
- Cost analysis for the conceptual design of a deployable system
- Conceptual design of a lost-cost accelerator-driven demonstration system with power in the 25-50 MWt range

Summary

The impact of fusion neutrons on waste transmutation and energy production from thorium is summarized in Fig. 9. The primary focus of this meeting is on LWR waste transmutation, but this paper draws attention to the possibility that waste transmutation and millennium-scale energy production might be linked advantageously. It also shows that externally neutron driven systems probably can eliminate the need for reprocessing and might eliminate the need for long-term high-level waste storage. The elimination of reprocessing is a vitally important advantage. The generation of nuclear energy with reprocessing using MOX fuel in today's LWRs is only marginally economic and would not be economic without a government subsidy for the capital cost of the reprocessing plant. All of the present concepts for waste transmutation in subcritical systems except that proposed here depend heavily on reprocessing and an expensive accelerator as well. Economics might allow the addition of either an accelerator or reprocessing to nuclear energy technology, but not both.

The transmutation of LWR spent fuel referred to[1] in other ADNA publications as Tier 1 is shown in the upper part of the Fig. 9. Plutonium and minor actinide are removed from the LWR spent fuel and fed to an accelerator-driven once-through transmuter operating at k_{eff} of about 0.96 where 90 % of the feed actinide is burned away and the weapons-usefulness of the material is eliminated. This degree of transmutation might be sufficient for waste storage at Yucca Mountain since the radioactivity of the defense waste begins to dominate at the 10 % remnant level. Cheaper neutrons from an improved accelerator

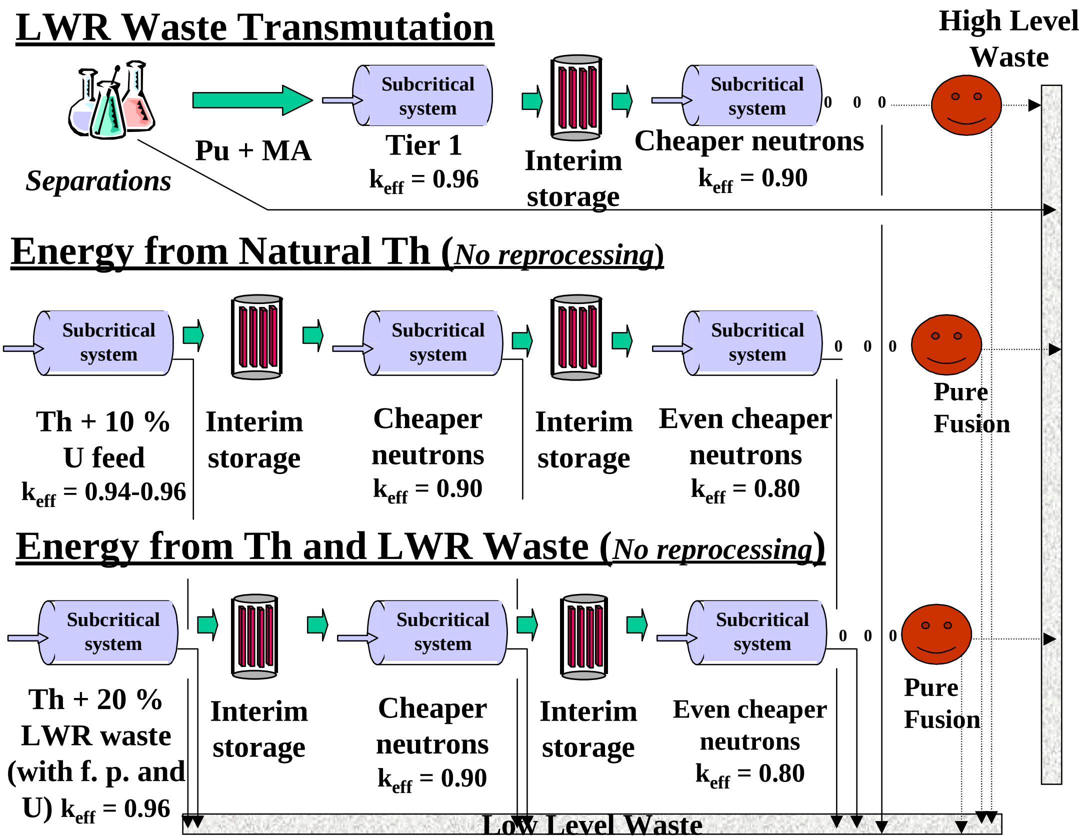

Fig. 9. Comparison of subcritical thermal-spectrum molten-salt systems for LWR waste transmutation, energy production from thorium, and the combined burning of both accompanied by the elimination of reprocessing and possible elimination of high-level geologic waste storage depending on societal demands and definitions.

or a fusion neutron source would allow the 10 % remnant to be returned to the system for equilibrium burning with $k_{eff} = 0.90$ and a further reduction of the remnant to perhaps 5 % or better. The "stretch" inherent in the fusion neutron source development could lead to further recycles without reprocessing continuing perhaps until the remnant does not meet high level waste definitions. However, the initial front-end separation generates a high-level waste stream and therefore the continued recycling with cheaper neutrons cannot eliminate high-level waste completely. Therefore while waste burning alone appears definitely to provide significant benefits probably at acceptable cost, it is lacking in closure.

The middle portion of the figure shows the thorium burner with 10 % uranium fed to denature the ^{233}U to non-weapons-useful material. With an accelerator it appears likely that this system can operate at $k_{eff} = 0.94\text{-}0.96$ while consuming in a single cycle 7 % of the Th-U fed. This is ten times the single pass uranium consumption achieved in the LWRs and therefore would extend the nuclear energy resource by a factor of 40 since there is four times as much thorium in the world as uranium. It has been shown above that the system also caries an inherent doubling time of about five years that can be implemented without reprocessing. It probably could provide the world's energy needs for several hundred years. If neutrons could be produced more cheaply by a factor of two, this energy output could be about tripled by reburning with $k_{eff} = 0.90$. As the fusion neutron source is improved this could go on and on until the thorium (and uranium) resource is depleted perhaps near the end of the millennium. Along the way low level waste is generated but no high-level waste need be released until the end. The end would not come until the energy source had been displaced by a new source. In the happy event that fission is succeeded by pure fusion and the abundant neutrons that can come with it, there would be sufficient neutrons to transmute any remaining fission product or actinide waste which fell into the category of high-level waste.

The thorium burning possibility in the end might not give rise to high-level waste, but it would leave the LWR waste behind. The bottom of the figure shows the combined burning of thorium and LWR waste. All of the LWR waste could be fed into the system without front-end separation and still exhibit reactivity greater than that of pure thorium burning. Successive reburning with improved neutron sources could over time burn almost completely the thorium, uranium, and long-lived fission product without back-end recycling, with the pure fusion system supplying as many neutrons as needed for clean-up at the end.

The combined burning of thorium and LWR waste is ultimately the most desirable technology since it would satisfy the world's energy need for a long time and go far toward resolving the nuclear high-level waste problem without front-end or back-end reprocessing. Of the three, the highest impact and lowest cost technology to develop and deploy is thorium burning, so this technology might deserve the most attention now. Although the burning of LWR waste alone currently has the most political drive behind it, it seems unlikely that any waste-burning-only solution will be soon driven out of the market place by combined thorium and LWR waste burning. In closing it must be mentioned that these projections require confirmation with more detailed neutronic studies, simple experiments, and engineering design.

Acknowledgements

The author is much indebted to Drs Edward T. Cheng and Curtis W. Bolton III for encouraging this evaluation of fusion-driven subcritical thermal-spectrum liquid-fuel systems. The help of Dr. Andrei Tchistiakov with the Monte Carlo calculations for evaluating the neutronics and the ultimate performance of the described systems is much appreciated. The analysis and evaluation of the accelerator-driven concepts behind this paper by Dr. Igor Slesserev were exceedingly helpful.

References

1. Charles D. Bowman, "Once-Through Thermal-Spectrum Accelerator-Driven LWR Waste Destruction Without Reprocessing," Nuclear Technology 132, 66 (2000)
2. DOE Report to Congress, "A Roadmap for Developing Accelerator Transmutation of Waste (ATW) Technology, DOE/RW-0519, Office of Civilian Reactor Waste Management (1999)

3. J. R. Engle, H. F. Bauman, J. F. Dearing, W. R. Grimes, E. H. McCoy, and W. A. Rhodes, “Conceptual Design Characteristics of a Denatured Molten-Salt Reactor With Once-Through Fueling,” ORNL/TM-7207 (1980)
4. This window was designed by J. E. Leiss at NIST for use with high power electron beams.
5. Charles D. Bowman, “Once-Through Thermal-Spectrum Molten-Salt Accelerator-Driven Waste Transmutation: Performance Calculations,” Report ADNA/01-03 (2001)
6. Olivier Meplane, “Technical Report on the Bowman Reactor,” University Joseph Fourier (Grenoble), Institute des Sciences Nucleaires (1998)
7. Alvin M. Weinberg, “Molten Salt Reactors,” Nuclear Applications and Technology 8, 102 (1970)