COMPLETELY AUTOMATED NUCLEAR REACTORS FOR LONG-TERM OPERATION II:

Toward A Concept-Level Point-Design Of A High-Temperature, Gas-Cooled Central Power Station System

Edward Teller†, Muriel Ishikawa† and Lowell Wood†
and
Roderick Hyde§ and John Nuckolls§

†Stanford University, Stanford CA 94305-6010

§University of California Lawrence Livermore National Laboratory
Livermore CA 94551-0808

Table of Contents

ABSTRACT

We discuss a new type of nuclear fission power reactor optimized for the generation of heat for use in obviously *safe, economic, and long-duration electricity production in large central power stations. These reactors are fundamentally different in design, implementation and operation from conventional light-water-cooled and -moderated reactors (LWRs) currently in widespread use. They feature a low-average-enrichment initial fuel loading which lasts the entire 30-year, full-power design life of the power-plant, and which is intended never to be removed from the reactor. The reactor contains a cylindrical core comprised of a nuclear ignitor and a much larger nuclear [breeding+burning] wave-propagating region containing natural thorium or uranium fuel, a surrounding neutron reflector and radiation shield, distributed means for implementing a thermostating function on the reactivity and local power density, a redundant pressurized-gas coolant transport system, and automatic-and-redundant heat-dumping means to obviate concerns regarding all classes of loss-of-coolant accidents during the plant's operational and post-operational life.These reactors are proposed to be situated in suitable environments at ≥ 100 meter depths underground. After optimization and learning exercises with prototypes, their operation will be completely automatic, with no powered mechanisms, no operator controls and no provision for human access during or after their operational lifetime, in order to avoid both error and misuse. The power plant's heat engine and electrical generator sub-systems are located above-ground, and are connected to the nuclear heat-source only with readily sealed coolant conduits. This paper outlines a concept-level point-design of a 1 GWe member of this*

type of reactor, one oriented to production of high-temperature, high-pressure coolant-gas and directed toward 60% efficiency, combined-cycle electricity generation.

INTRODUCTION AND OVERVIEW

A fundamentally new design of nuclear power reactors for 21st century electricity generation applications is needed for at least four fundamental reasons: concerns about nuclear reactor accidents, the economically non-competitive posture of modern nuclear electricity relative to the best fossil-fired generation technology, worries about nuclear weaponry proliferation (via diversion of reactor products to make nuclear weapons), and the potential long-term shortage of nuclear fuel for high-intensity world-wide applications.

An entirely fresh look at nuclear-energized large-scale power generation is called for, particularly in the light of rapidly increasing demand for electricity in the capital-poor Third World and growing concerns as to the environmental consequences of meeting this demand with fossil fuel-fired systems. It isn't clear that basic parameters of currently available types of nuclear power reactors will permit them to meet such demands.

This paper sketches a concept-level point-design of a full-scale member of a representative of the general class of power reactors which we believe may be suitable for satisfying world-wide needs for electricity in the 21st century.

We are considering a reactor serving as a $\leq$ 2 GW heat source for electric power production at the 1 GW level, emplaced $\geq$ 100 meters underground and operating for 30 years without human access after the start of power production, the purpose of such isolation being to avoid both error and misuse. The power plant's heat engine and electrical generators are located above-ground, and are connected with the nuclear heat source by coolant conduits carrying high-temperature helium gas rather than water. Heat-to-electricity conversion efficiencies as high as 60% may be achieved by using such a reactor to drive a combined-cycle system centered on aeroderivative turbines (technologically descended from jet engines) similar to those already widely used in the most efficient natural gas-fueled generating stations (which are also the economically pace-setting ones).

Our current reference-design reactor contains a cylindrical core comprised of a small nuclear ignitor and a much larger nuclear burnwave-propagating region. The latter contains natural thorium or (possibly depleted) uranium fuel, and functions on the general principle of fast breeding. The entire core is surrounded by a neutron reflector and a radiation shield. Uniform temperature throughout the core is maintained by a large multiplicity of thermostating modules which, through the action of simple automatic controls transporting isotopically-enriched lithium when the local material temperature rises into the regime corresponding to a coolant-gas-temperature design-value of ~1000 K, depress the local neutron flux and thereby reduce the local power production. Triply-redundant primary means of transporting heat up to the generating station are provided, and entirely independent, triply-redundant energy-dumping means are included in this design to passively transport afterheat out of the core in the event of a loss-of-coolant accident or after the end-of-operational-life.

When the core's initial loading of nuclear fuel is exhausted after approximately thirty years, the reactor is permanently shut down by the addition of a neutron absorber to the core. The core's spent fuel is allowed to beta-decay in the reactor. The surrounding hot dry sand of the engineered heat-dump provides long-term protection against environmental conditions which might otherwise eventually induce leakage of residual radioactivity into the biosphere.

MOTIVATIONS OF THE PRESENT WORK

Nuclear reactors for central-station electricity generation have not been widely accepted by the public in most technically advanced nations, even though they are relatively safe. What is urgently needed in present circumstances is that they be seen to be *obviously* safe.

Compared to other large-scale electrical energy sources, modern nuclear reactors <u>are</u> safe. Three Mile Island, the location of one of the two big malfunctions during the past two decades, saw an accident expensive in dollars and disastrous for the utility owner, but injurious to no one. Chernobyl was far more serious and did kill dozens of people outright (and perhaps hundreds of others, on a statistical basis). Nonetheless, the human-life-loss per kilowatt-hour of nuclear electricity generated over the past several decades is still below the loss-of-life per kilowatt-hour arising from fossil-fired electricity generation over the same interval.

The remarkable fact about these two large-scale accidents is that both of them crucially involved human error in power plant operation, a contingency only little considered in the design of these systems. (In addition, the Chernobyl design was burdened with positive temperature- and void-coefficients of neutronic reactivity.) It is our opinion that such accidents can be ruled out satisfactorily if we can separate a functioning power reactor and its resulting radioactivity from <u>any</u> human interference (following an extended interval of design-debugging and -optimization).

The first new design measure to achieve *obvious* safety is to construct power reactors underground. This has also been advocated vigorously by

Andrei Sakharov, following similar reasoning. Logically, the reactor's mass-transporting connections to the surface also must be automatically and robustly shut down, in event of a major malfunction.

A second new design measure is to make the function of the reactor fully automatic and *obviously* safe from all types of human error. The exclusion of human access - and thus all types of close-up inadvertence - requires that the reactor fuel elements should not be replaced, that the reactor should function without human maintenance and that, after the conclusion of power production, the reactor and its fission products should remain in place indefinitely. Indeed, present plans to collect reactor products at a safe central location has the intrinsic difficulty of hazards arising from accidents during the initial removal and storage and the subsequent transportation of the spent fuel assemblies.

We are therefore studying reactors which deliver high-grade (i.e., high-temperature) heat-on-demand. A reactor of interest must *automatically* deliver more thermal power when more electricity is demanded from the associated central station - and it must not overheat when the power demand is reduced.

A fundamental requirement for any practical power reactor is that it should be economically competitive. One aspect of this requirement is met by providing a high degree of intrinsic and obvious safety at low cost, so that expensive safety mechanisms, personnel and regulations become unnecessary. Another is to minimize operations-and-maintenance costs and a third is to achieve exceptionally high thermal-to-electric conversion efficiencies.

A critical consideration to be kept in mind is the recent rather drastic decrease in the cost of fossil-fired electricity generation due to technological advance. During the past decade, electrical plants have been put into operation which work with 60% thermal efficiency and produce electricity with a total unit energy cost (TUEC) of $0.04 per kilowatt-hour, moreover in quite modest minimum sizes (e.g., a few tens of megawatts) and with environmental impacts (e.g., atmospheric and thermal pollution) which have already been reduced to remarkably low levels. This TUEC for such "combined-cycle" electricity is roughly half of that available from modern American nuclear power plants, which also produce twice as much waste heat per kilowatt-hour. Advanced nuclear fission power reactors - and, indeed, all emerging nuclear energy sources - must be economically competitive with the best modern conventional alternatives. Moreover, the costs of these alternatives must be expected to decrease, as the new combined-cycle generation technology matures.

A second fundamental requirement is connected with the long-term availability of nuclear fuel. At present, 3000-4500 tonnes of uranium metal are required to fuel a single 1 GWe plant for a 30 year interval (depending on whether or not reprocessing is done). Sources of only 3 million tonnes of uranium metal are identified in the United States, unless the production cost is increased significantly over $100 per pound. Thus, fuel for 1000 GWe of nuclear power - about the present-day level of non-nuclear generating capacity in the U.S. - can be guaranteed only for less than a half-century, unless breeder reactors are developed and deployed. Considering world-wide requirements and assuming an average per capita consumption of only one-third that of the U.S., a total electrical generating capacity of 10,000 GWe would be required. The known world-wide reserves of uranium ore accessible at $100/pound of uranium metal produced would suffice for LWR-based electric power production only for a modest fraction of a single century, accentuating the requirement for much more efficient use of nuclear fuel.

GENERAL FEATURES OF THE NEW TYPE OF REACTOR

Our work through the present time focuses on the conceptual-level design of a new type of nuclear power reactor aimed at simultaneously providing all these basic features and requirements which appear to be prerequisites for widespread nuclear electricity generation in the first half of the 21st century. Naturally, specifics must be considered if substantial progress is to be made. We are currently analyzing a baseline point-design of a representative member of this new type of power reactor from the standpoints of materials-and-structures, heat transfer and nucleonics, and have simulated its structural, thermal and nuclear performances with digital computer-based models. We present some interim results of these ongoing analyses. (The models themselves are discussed in Appendix A.)

We consider primarily thorium-fueled reactors, for thorium is widely distributed geographically in high-grade ores, and is cheaply available. Such a reactor must be a breeder, for reasons of efficient nuclear fuel utilization and of minimization of requirements for isotopic enrichment. It must be a fast breeder because the high absorption cross-section of fission products for thermal neutrons does not permit the utilization of more than about 1% of thorium (or of the more abundant uranium isotope, U^{238}, in uranium-fueled versions), without removal of fission products. The pertinent neutron cross-sections vs. neutron energy are shown in Figures 1 and 2.

The conventional approach is to remove fission products by reprocessing or by actively transporting the fission products away from the nuclear reaction zone. This requires human intervention, which we wish to exclude entirely or, alternatively, relatively complicated arrangements of uncertain reliability or safety. Diversion of fissile materials to military purposes is a risk inescapably associated with either the intermittent or continuous reprocessing of nuclear fuel. (Also, closing nuclear fuel cycles via reprocessing and *obviously* safe long-term storage of fission products has yet to enjoy universal satisfaction, even in most of the technologically advanced countries.)

We are considering designs in which the nuclear reactions are driven primarily by fast neutrons, most of which have energies within a factor-of-ten of the MeV-scale temperatures at which they are evaporated from nascent fission fragments. Several substantial advantages are thereby obtained. First, the pertinent ratios of cross-sections are much more favorable to breeding, and thus to a potentially high fuel burn-up. Second, the preferred structural materials, such as Nb and Ta, have relatively small absorptions at these high neutron energies, so that comparatively large amounts of them may be employed without imposing significant damage on the reactor core's neutron economy. Third, the net neutron multiplicity of fission is significantly greater than the breeding threshold-value of 2.0 at fission neutron energies than it is at thermal neutron energies, and notably so for U^{233}. Fourth, neutron-engendering processes peculiar to fast neutrons, such as fast fission of Th^{232} or U^{238} and (n,Xn) reactions on all actinides, while each relatively unimportant, collectively contribute non-negligibly to the overall neutron economy.

We are therefore now investigating the extent to which this general type of design permits the simultaneous attainment of one-time fueling with low average enrichment materials for three decades of full-power production and of *obvious* great safety margins.

FUNDAMENTAL FEATURES OF THE REACTOR CORE

Our analyses have led us to provisionally conclude that this new type of power reactor in a 1 GWe-scale format will be a right circular cylinder of approximately 3 meters diameter and 10 meters length. This core basically consists of a small nuclear ignition region containing fuel enriched in U^{235} (albeit not to an extent supportive of diversion to assembly of a nuclear weapon), embedded in a much larger breeding+burning section containing U^{238} (natural or depleted uranium) or, preferably, Th^{232}. Core-averaged fissile isotopic content (e.g., ~1% by mass) is comparable to that of natural uranium. This, together with the feasibility of very high average fuel burn-up in the designs which we consider (e.g., 50%), indicates a total requirement for perhaps 100 tonnes (rather than 3000-4500 tonnes) of as-mined fuel for a 1 GWe reactor's entire three-decade full-power operational life. Core nucleonics are discussed in Appendix B.

As core coolant, we plan to employ pressurized helium, rather than water. This permits the utilization of thermal energy at substantially higher temperature, avoids all hazards arising from water reactions with high-temperature materials, and provides favorable independence of the core's neutronic reactivity on coolant-voiding. Key features of the core's coolant system are shown in Figure 3. Core thermal transport issues are also discussed in Appendix C.

Reactor core construction materials, e.g., Nb, Ta, W and Re, are widely available. They are chosen primarily for their superior long-term, high-temperature mechanical and chemical properties, which are retained to adequate extents under large fluences of high-energy neutrons. Core structural design issues are treated in Appendix C.

Management of nuclear power production in the reactor's core is fully automatic in all respects, over the entire three-decade interval between nuclear ignition and final core shutdown at the power plant's end-of-operational-life. Participation of human operators is needed only at the commencement of the core's ignition and at the final, irreversible negation of its reactivity.

Fully automatic regulation of nuclear power production is performed by uniformly distributed, functionally redundant thermostating modules. Each of these acts to absorb strongly the local neutron flux when the local material temperature exceeds the design-value, thereby quenching local nuclear power production and assuring thermal homeostasis of every portion of the core, over wide ranges of coolant flow and temperature, fuel composition, neutron spectrum and neutron flux. Each thermostating module acts by reversibly inserting neutron-avid liquid Li^6 into a small (<100 cm^3) compartment located in the coolant-flow from a source outside the neutron reflector, under the drive action of a thermostating bulb filled with neutron-indifferent liquid Li^7 which is emplaced in adjacent, substantially hotter nuclear fuel. A 3-D lattice of such thermostating modules, each functionally-independent from all of its fellows and emplaced during manufacture throughout the core, serves to regulate the matter temperatures everywhere, at all times. See Figure 4.

REACTOR CORE OPERATION UNDER NORMAL CONDITIONS

At the commencement of the reactor's operational life, the centrally-positioned nuclear ignitor module is driven critical by one-time removal of neutronic poison and, through concurrent nuclear fission and high-gain breeding actions, commences to launch a nuclear deflagration wave into the adjacent unenriched fuel. This wave first diverges radially from the centrally positioned, on-axis nuclear ignitor until portions of it reach the outer edge of the cylindrical fuel mass, where it is resolved into two oppositely-directed, axially-propagating waves. One such wave moves toward each of the two ends of the cylindrical core at a (exceedingly low peak) speed determined at all times by the instantaneous thermal power demand on the reactor (and upper-bounded by the leisurely beta-decay of Pa^{233}, the rate-limiting step in Th^{232}-U^{233} breeding). When only very low power is demanded, the neutronic reactivity is driven to zero by action of the thermostatic controls, and the deflagration-wave stalls; when heat is removed from the core at a greater rate, the thermostating controls cool and thus raise the neutronic reactivity to slightly positive levels, and the [burning+breeding] wave re-commences its advance.

Fuel moderately enriched in fissile material is generated behind each of the two wave-fronts. These two increasing masses of enriched fuel then continue to burn, until fission product accumulation and fertile isotopic depletion (at a 50% core-averaged fuel burn-up) finally drives the core's neutronic reactivity negative. Figure 5 illustrates typical conditions ahead of, within, and behind this pair of nuclear deflagration waves, and Appendix B discusses core nucleonics in more detail.

Heat from the burning fuel is removed from the reactor and transported to the surface by a coolant-flow consisting of pressurized helium-gas. The helium is circulated axially through the core of the reactor by a 2-D hexagonal lattice of pipes, as depicted in Figure 3. The nuclear fuel is emplaced between the pipes and transfers heat to them by thermal conduction, whereupon it heats the circulating helium. The coolant loops are designed with complete three-fold redundancy and, under normal operating conditions, are actively powered by pumps at the surface. The primary in-core structure of the reactor is the array of coolant pipes; these are currently planned to be constituted of a tantalum alloy in order to achieve the required long-term, high temperature creep resistance. The entire reactor is emplaced within a containment vessel which, because it is located outside the hot, neutron-rich core, can be implemented with a conventional high-quality steel alloy. Details are discussed in Appendix C, and key features are shown in Figures 3 and 6.

From an operational standpoint, the reactor, which is emplaced $\gtrsim$ 100 meters underground to provide a substantial additional measure of biospheric safety, is fully automated. As already noted, no operator controls of any kind are provided, other than the one to start it up initially and another one to finally shut it down after three decades of (nominally, full-)power production. Fuel is loaded into the reactor only when its core is built, and spent fuel is never removed. The reactor is designed to be maintenance-free, and is never to be accessed after it first commences operation. Requirements for specially-trained, nuclear-cognizant personnel during plant operation are thus strictly minimized.

REACTOR CORE BEHAVIOR UNDER ABNORMAL CONDITIONS

The ability of *each* of the set-of-three independent primary coolant loops to remove the thermal power from the reactor's core naturally invites partitioning of the coolant conduits and the heat engines on the surface into three distinct modules, each capable of entirely adequate core cooling. Doing so drastically reduces the likelihood of a common-mode loss-of-coolant accident, relative to the likelihood of such events in typical reactors.

Independent of this triply-redundant primary cooling, a further large safety margin is provided by fully automatic, entirely passive rejection of nuclear afterheat into the underground environment surrounding the reactor, via *another* triply-redundant set of coolant loops within which the coolant is convected entirely passively. The engineered heat-dump into which the heat is transported is design-rated to accept the full-power afterheat of the reactor for an indefinitely long duration, and is described in Appendix D.

Thus we design to preclude core damage due to excessive temperature, both in the course of arbitrarily severe loss-of-coolant accidents during full-power operation and in a zero-maintenance scenario following final reactor shutdown, in an extremely reliable, six-fold redundant manner. See Figure 6.

Nuclear afterheat will be removed during loss-of-coolant accidents and after end-of-operational-life using the same basic method (convection by gaseous helium) as during normal full-power operation. We constrain our design quite conservatively by the assumption that none of the three independent primary coolant loops nor their circulating pumps will be available in these two circumstances. To perform appropriately in such situations, we provide a second coolant system, as indicated in Figures 3 and 6. This secondary, entirely independent coolant system consists, as did the primary system, of a triply-redundant array of cooling loops. Helium coolant-gas is circulated at sufficient rates through these pipes by a passive thermosyphon, whose driving pressure head is simply due to the fact that the reactor is emplaced below the engineered heat dump to which the afterheat is rejected for indefinitely long periods. Details are provided in Appendix D.

Redundant, engineered, one-time-operation coolant pipe closures act to robustly seal off the reactor from the biosphere, either under operator direction at end-of-operational-life or automatically in the event of large-scale fission product entrainment in the coolant-gas stream, functionally backing-up automatically actuated mechanical valves. These are described in Appendix E.

REACTOR DECOMMISSIONING CONSIDERATIONS

The reactor vessel, positioned in the engineered heat-dump, serves as the spent-fuel burial cask. The core's post-end-of-life afterheat is passively transported through the triply-redundant, passively-convected coolant loops and coupled into a highly redundant network of heat pipes threading throughout the surrounding engineered heat-dump. Indeed, the operation of this engineered heat-dump after final shutdown serves to generate a ~100 meter-diameter "bubble" of hot, dry sand around the reactor vessel. This spatially-extensive engineered heat-dump, described in Appendix D, also serves to decouple the reactor vessel from the ambient groundwater environment, thereby ensuring the vessel's long-term mechanical integrity and thus its suitability as the burial cask for the spent fuel.

Active measures to decommission the reactor and to remove spent fuel from its core (including several tonnes of residual fissile material) are thereby completely obviated. Burial cask integrity is ensured for multi-millennia intervals, sufficient to see total beta activity in the core decay by a million-fold, to residual levels less than 10 kilocuries.

DESIGN VARIANTS

This new type of power reactor has a great deal of inherent variability implicit in its design. We have emphasized a particular design variant with great inherent safety realized by automatic power control and afterheat removal. Moreover, we have focused attention on one which generates high-temperature heat, believing that a unit of 1000 K heat indeed has an economic worth twice that of a thermal unit at the 600 K temperature typically delivered by LWRs, and that the likely-somewhat-greater costs of a high-temperature heat-supply will be significantly exceeded by the assuredly-greater economic benefits of combined-cycle generation.

It must be noted that major variations on the fundamental theme sketched above may be feasible. For instance, we have evaluated a 600 K water-heating design (of appropriately minimized intra-core coolant volume) which functions in nucleonic, structural and most heat transfer senses quite similarly to the helium-heating one just discussed. We have also examined uranium-fueled cores, and have verified that they function basically the same as thorium-fueled ones in nucleonic senses (though the latter have superior high-temperature mechanical properties, avoiding Pu and its low melting-point).

A major design variable is the peak specific fuel power. For a typical thorium-fueled core design with our baseline peak fuel specific power of 200 W/gm, we pre-expand the metallic fuel by 3- to 4-fold above its full-density specific volume, i.e., so that the initial mean fuel density is 3-4 gm/cm^3. (Doing so lowers the intra-fuel peak pressure occurring at maximum burnup, albeit at some loss in the fuel's thermal transport qualities, and lowers worst-case structural requirements.) Operating at lower peak specific power in the fuel or with significantly different initial fuel mean density offers the prospect of reactor cores - and thus reactors - of quite different configurations.

In summary, this new type of nuclear power reactor is remarkably rich in design possibilities.

SOME NATIONAL ENERGY STRATEGY CONSIDERATIONS

From a national energy systems standpoint, reactors of this new type act to supply high-temperature heat-on-demand throughout their entire operational lives, and thus constitute a base for reliable, large-scale electrical energy supply at a predictable, capped cost over multi-decade intervals. Since they have no day-to-day operator controls and are never re-fueled or serviced, they generate no requirements for highly-trained personnel or special materials (e.g., enriched fuel) after they first commence operation. Their exceptionally efficient use of unenriched fissionable material - thorium or natural (or depleted) uranium - obviates all long-term fuel supply issues. Each such reactor uses (in its ignitor) only ~0.001 of the fissile material - enriched uranium - already in international commerce.

Since they are never re-fueled, these new types of reactors present no spent-fuel handling, transport, reprocessing, or disposal issues, and they greatly reduce military diversion concerns. This last point may be of the greatest importance.

The high-grade, high-pressure heat produced-on-demand by this new type of reactor is intended to support combined-cycle electricity generation, with its especially favorable economic and environmental impact indices, e.g., half of the TUEC for electricity and half of the waste heat rejected to the environment of modern LWRs.

Because of its extraordinarily great, multi-layered set of safety features, this new type of reactor may ultimately be suitable for underground siting in urbanized areas. Such siting would also substantially reduce capital charges for transmission systems, as well as make reasonably high-grade waste heat readily available, e.g., for space heating and cooling applications.

MATERIALS AVAILABILITY CONSIDERATIONS

High quality manufacture of this new type of power reactor, either for indigenous use or for export, is within the capabilities of any technologically developed country. Fuel cladding and intra-core coolant piping materials used in ~25 tonne total quantities per GWe core may have to be imported, but these are available from multiple commercial sources.

Moderately enriched uranium for use in the ignitor modules of reactor cores can be competitively procured in any needed quantity from Russian, European or American sources by any NPT-member nation. Enriched lithium isotopes for use in filling the thermostating modules are similarly available in pertinent quantities.

SOME ECONOMIC CONSIDERATIONS

At the present time, the economic pace-setting technology for large-scale central station electricity production is natural gas-fired combined-cycle generation. The typical total unit energy cost (TUEC) of such conventionally-generated electricity is ~$0.04 per kilowatt-hour. Of this amount, approximately $0.010 represents capital costs, $0.004 is attributed to operations and maintenance, and ~$0.025 is expended to buy the natural gas fuel.

According to our present estimates, the additional capital cost-equivalent of the new type of nuclear heat source (~$0.010 per kW-hr) will be comparable to the $0.010 per kW-hr capital cost-equivalent of a gas-fired, combined-cycle turboalternator system. Thus, the total capital charges of a combined-cycle central station energized by the new type of reactor will contribute a total of ~$0.020 per kW-hr to the TUEC; the already-minor cost contribution due to operations and maintenance ($0.004 per kW-hr) is likely to be reduced. The principal avoided-cost of the new type nuclear reactor-energized variant relative to the gas-fired variant of combined-cycle generation would be the dominating one of the natural gas fuel-cost; in contrast, initial nuclear fuel costs are capitalized and nuclear refueling costs would be zero, as refueling is never done. (The capital cost of the system to reliably dump afterheat is a small fraction of total plant capital costs.) Thus, the present pace-setting TUEC of $0.04 per kW-hr for combined-cycle gas-fired generation may be decreased to less than $0.03 per kW-hr for this nuclear-energized combined-cycle electricity.

However, careful engineering and economic analysis of a complete design will be necessary to develop significant confidence in these estimates. Working with additional collaborators, we hope soon to have a reasonably detailed design available for consideration by the entire nuclear engineering community.

CONCLUSIONS

Large-scale, central-station electricity production using nuclear heat sources has become an unattractive prospect in a number of technically advanced countries. This situation has persisted for two decades in some cases. It is perhaps not a severe exaggeration to suggest that the crisis facing nuclear power is a fundamental, structural one, compounded of substantial technical and economic components. Even more important is the need to deal with public perception.

If popular opinion is to accept - let alone invite - re-introduction of nuclear power generation in the majority of the OECD nations, the several existing and largely independent arguments against its usage must be effectively answered. Minor variations on widely deployed nuclear power technologies may not suffice.

In token of the seriousness with which we view the current situation, we are attempting to develop one such fundamentally new approach, which we have sketched above. In a number of salient respects, it substantially simplifies the *obviously* safe winning of high-temperature heat from nuclear power reactors for use in modern thermal engines. See Figure 7.

This new approach offers the prospect of assured proper reactor utilization not just when the best-trained and most highly-motivated technicians are its operators, but also when the least-trained and most careless operators may be in charge. It permits nuclear power generation capabilities to be made available with high confidence regarding materials diversion to countries which may not have highly stable political arrangements. It fully addresses nuclear fuel supply issues, even when intensive, world-wide nuclear electrification is considered. It potentially doubles the economic value of a unit of nuclear-derived heat, by delivering it at substantially higher temperature for conversion to electricity with combined-cycle technology. We expect that this set-of-features may significantly lower both the economic and the non-economic costs of long-term, large-scale nuclear electricity production, reducing such costs to highly competitive levels.

Consequently, we believe that this new approach, or one similar to it, may satisfy via nuclear power much of humanity's requirements for electricity in the 21st century. Indeed, some such novel approach may be *necessary*, if nuclear power is to fulfill its potential.

APPENDIX A: MODELING

In order to investigate and quantify physical-technological feasibility, we have employed digital computer-based performance simulation of several different types of models of the new class of reactors which we discuss. These models support detailed studies of the **nucleonics**, the **structural** and the **heat-transport aspects** of reactor functioning.

NUCLEONICS

For maximum design flexibility and fidelity-of-modeling, we have examined the neutron transport and nuclear reactions in our model reactor designs with Monte Carlo-based means. To perform this function, we have employed the general-purpose TART95 neutron and gamma-ray three-dimensional transport- and reaction-modeling code-set developed and distributed by the Lawrence Livermore National Laboratory (LLNL). This software package represents a development effort whose scale is of the order of 10^2 man-years and an associated code-validation effort of the order of 103 man-years in size. TART95 and its ancestors have very frequently been employed for calculation of the reactivity of critical assemblies. However, they lack time-dependence, in that they generate snapshots-in-time of the transport phenomena which they model.

We have used the latest-released version of the LLNL ENDL (Evaluated Nuclear Data Library) as the physical data source for this code, which we have employed exclusively in the 175 neutron energy-group mode, with TART95's most recently upgraded thermal scattering and resonance cross-section multi-band-averaging features.

Model reactor designs in our studies are resolved into several hundred spatial zones, usually possessing axial symmetry, and a few hundred different materials. Sixteen isotopes are usually carried in each zone, representing both fertile and fissile isotopic components of nuclear fuel, in addition to reflector and coolant elements, structural materials, and various neutronic poisons (including fission products, carried as an ENDL-standard mix). Thus we ensure proper reactivity dependence on temperature and accurate representation of the course of long-term, possibly high fuel-burnup reactor operation. Our models often employ homogenized materials, whenever the physical scale-lengths of different materials are less than or equal to neutronic mean-free-paths for neutron energies of interest. Indeed, we constrain zone dimensions to be less than a neutronic mean-free-path for any principal reaction (e.g., radiative capture or fission), over the entire neutron spectrum of interest. We employ different spatial zones, many carrying a unique material composition, in our models to account for substantially different material (or isotopic) compositions throughout the reactor.

The behavior of the isotopic fractions in each zone of a model problem are integrated in time with a fourth-order Runge-Kutta integration scheme, which couples the standard fissile and fertile isotopes of the actinide elements to each other and to fission products, using the reaction rates just calculated for the conditions in each of the spatial zones of the problem by TART95. (We typically follow ~155 neutron-driven reactions in each zone, including all fission, radiative capture, elastic and inelastic scattering, (n,Xn), (n,p), (n,d), (n,t), and (n,a) reactions, each in a properly energy-dependent manner, whenever the corresponding cross-sections are at-or-above the 0.01 millibarn level.) Neutron absorption on all non-actinide isotopes is implicitly accounted for. The newly-updated isotopic abundances in each zone are then inputted as a newly-reformulated "problem" to the TART95 code for another cycle of neutron transport and reaction calculations, thereby completing the basic set of operations of a single integration time-step.

The magnitude of the time-step of the kinetics integration is determined by the maximum permitted fractional change (usually 5-10%) in any of the major isotopic concentrations in any zone of the problem. (As would be expected, this 'critical value' is typically the fissile isotopic or the fission product concentration in the leading edge of the nuclear deflagration-wave propagating into the unenriched fuel-charge.) Between 150 and 500 time-steps suffice for an integration simulating 3 decades of reactor operation, depending on choices made regarding initial conditions and time-step controls.

The top level of our nucleonics modeling program-set, which we call BURNBRED (for 'burn' and 'breed'), combines the TART95 neutron transport-and-reaction package with the isotope kinetics integration package, and provides input, control and editing functions. For the present study, it has been hosted on an IBM-type personal computer (IBM PC) system. A model run for GW-scale reactors of ~250 materials, ~400 zones and 16 isotopes per zone operating over a simulated 3-decades requires 30-100 hours of computing time. In one simulation, several trillion floating-point arithmetic operations are performed and several billion bytes of intermediate neutronic reaction data-sets are processed using the computing system's hard-disk memory.

Such calculations often include use of a heat-transport feature in the BURNBRED package which permits the modeling of thermostatic module control of the time-dependent neutronic reactivity variation of the reactor's core, relative to *a priori* specifications of coolant flow through the core.

While time-dependent reactor modeling by our Monte Carlo-based approach would certainly be an extravagant expenditure of computing resources by standards of even a decade ago, the total [capital+operating (electricity)] costs of a single few-day-long calculation on our modern PC is of the order of $3. The human time-to-assimilate the results of such a problem-run and specify the design of the next problem to be modeled is usually not much smaller than the duration of the run itself; thus a far faster computer could not be effectively employed. Indeed, the computing system which we currently use (centered on a Pentium 166 MHz chip interfaced to 512 Kbytes of pipeline-burst cache memory) has been benchmarked to be roughly twice as fast as a single CPU of the fastest supercomputer generally available, the CRAY-YMP, when executing the extremely memory-reference-rich and highly scalar (e.g., branch-intensive) instruction-sequences characteristic of our modeling tool-set.

STRUCTURES AND HEAT TRANSFER

The present baseline-design core configurations, discussed in Appendix C, have been largely designed and analyzed with semi-analytic methods organized into spreadsheets.

These analyses determine the coolant pumping requirements by balancing the fluid drag in the coolant-tubes threading the core's fuel-charge (which tubes dominate total loop losses) against the available pressure-head. The fluid drag is modeled using turbulent pipe friction formulae, while the pressure-heads are either specified (for the pumped primary coolant flow) or calculated (for the thermosyphoned secondary coolant flow) from the thermal-gradient-derived density profiles.

The heat transfer from the fuel into the coolant pipes is studied with finite element codes (discussed below). These detailed results are then scaled with size and power levels for use in the design-&-analysis spreadsheets. The transfer of heat from a pipe wall into the coolant-stream is calculated using turbulent boundary-layer heat-transfer coefficients.

The stress levels encountered in the coolant tubes are readily found, once the pressures and wall thicknesses are known. The former are determined by the pumping-power analysis just sketched, while the amount of tube-wall material is limited to levels that can be tolerated within the core's neutron economy.

These relations are then combined to develop and analyze specific reactor designs. The coolant mass-flows required are set by the inlet and outlet temperatures, while the size-scale of the coolant tubes is determined by the desired temperature drop across the fuel (i.e., from its hottest point to the nearest coolant-tube wall). The aggregate area of coolant tubes and the fuel's specific volume are set by desired values of the overall neutronic thickness of the fuel-charge (i.e., the core fuel-charge's density-radius product) and the peak temperatures of the coolant-pipe material. These factors then determine the inlet pressures needed in both the primary and secondary loops in order to circulate the coolant. These pressures and the quantities of coolant-pipe material available dictate the stress levels in the pipe walls.

The transfer of heat from the porous, burning fuel-charge into the coolant-pipes threading through it is studied with detailed finite-element thermal codes. These analyses employ volumetric heat production in the fuel, thermally nonlinear conduction through the fuel and the walls of the coolant tubes, and boundary-layer heat-transfer coefficients to finally deliver the heat into the bulk of the coolant-gas. These detailed analyses are particularly useful when considering the performance of the 2-D lattice of coolant-tubes during situations when some of the loops comprising the lattice are inoperative.

The structural behavior of the present array of fuel and coolant-gas tubes is straightforward; the coolant tubes are cylindrical and their stresses can be found analytically. Other, earlier designs which we have considered involved coolant tubes and fuel configurations with much more complex cross-sections; these systems have been studied with coupled thermal and mechanical finite-element codes. These codes remain available for more detailed analyses of future reactor designs, whether these are based upon the present simple configurations or are ones utilizing more complex layouts.

APPENDIX B: NUCLEONICS

The first phase of the nucleonics of the new type of reactor likely is familiar to students of fast breeder reactors. A centrally-positioned nuclear ignitor moderately enriched in U^{235} has a neutron-absorbing material (e.g., a borohydride) removed from it by operator-commanded electrical heating, and becomes neutronically critical. Local fuel temperature rises to the design set-point and is regulated thereafter by the local thermostating modules. Neutrons from the fast fission of U^{235} are mostly captured at first on local U^{238} and Th^{232}.

(Uranium enrichment of the ignitor may be reduced to levels not much greater than that of LWR fuel by introduction into the ignitor and the fuel region immediately surrounding it of a radial density gradient of a refractory moderator such as graphite; high moderator density enables low-enrichment fuel to burn satisfactorily, while decreasing moderator density permits efficient breeding to occur. The optimum ignitor design involves trade-offs between proliferation robustness and the minimum latency from initial criticality to the availability of full-rated-power from the fully-ignited fuel-charge of the core; lower ignitor enrichments require more breeding generations and thus impose longer latencies. We consider use in the ignitor of uranium enriched to $\leq$ 20% U^{235} to be highly proliferation-resistant. Even though such material could be diverted to use as relatively high-performance feedstock for isotopic enrichment to weapon-useful material, the complexity of doing so is not qualitatively smaller than an *ab initio* effort using natural uranium.)

The core's maximum (unregulated) reactivity slowly decreases in the first phase of the ignition process because, although the total fissile isotope inventory is increasing monotonically, this total inventory is becoming more spatially dispersed. By proper choice of initial fuel geometry, fuel enrichment vs. position, and fuel density, it may be arranged for the maximum reactivity to still be slightly positive at the time-point at which its minimum value is attained. Soon thereafter, the maximum reactivity begins to increase rapidly toward its greatest value, corresponding to the fissile isotope inventory in the region of breeding substantially exceeding that remaining in the ignitor. A quasi-spherical annular shell then provides maximum specific power production. At this point, we refer to the core's fuel-charge as "ignited." Similarities to "standard" types of fast breeder reactors now become more distant.

The spherically-diverging shell of maximum specific nuclear power production continues to advance radially from the ignitor toward the outer surface of the fuel cylinder. When it reaches this surface, it naturally breaks into two spherical zonal surfaces, with one surface propagating in each of the two opposite directions along the axis of the cylinder. At this time-point, the full thermal power production potential of the core has been developed, and we characterize this epoch as that of the launching of the two axially-propagating nuclear deflagration waves. (We choose to ignite at the center of the core's fuel-charge and thus to generate two oppositely-propagating waves simply in order to double the mass and volume of the core in which power production occurs at any given time, and thus to decrease by two-fold the core's peak specific power generation, thereby quantitatively minimizing thermal transport challenges.)

From this time forward through the break-out of the two waves when they reach the two opposite cylinder-ends, the physics of nuclear power generation is effectively time-stationary in the frame of either wave, as is suggested by the snapshots of Figure 5. The speed of wave advance through the fuel is proportional to the local neutron flux, which in turn is linearly dependent on the thermal power demanded from the reactor core (via the collective action of the thermostating modules).

When more power is demanded from the reactor via lower-temperature coolant-gas flowing into the core, the temperature of the two ends of the cylindrical core (which are closest to the coolant-gas inlets) decreases slightly below the thermostating modules' design set-point, Li^6 is thereby withdrawn from the corresponding sub-population of the core's constellation of modules, and the local neutron flux is permitted to increase to bring the local thermal power production to the level which drives the local temperature up to the set-point of the local thermostating modules.

However, this process is not effective in heating the coolant-gas significantly until its two divided flows move into the two nuclear burn-fronts. These two portions of the core's fuel-charge - which are capable of producing significant levels of nuclear power when not suppressed by Li^6-loading - then act to heat the coolant-gas to the temperature specified by the design set-point of their modules - provided that the nuclear fuel temperature doesn't become excessive (and regardless of the temperature at which the coolant-gas arrived in the core). The two coolant flows then move through the two sections of already-burned fuel centerward of the two burn-fronts, removing residual nuclear fission and afterheat thermal power from them, both exiting the fuel-charge at its center. This arrangement clearly encourages the propagation of the two burn-fronts toward the two ends of the fuel-charge, by "trimming" excess neutrons primarily from the trailing edge of each front. See Figure 5.

This description implies that the core's neutronics are essentially self-regulated. As long as the fuel density-radius product of the cylindrical core is $\geq$ 200 gm/cm^2 (i.e., 1-2 mean free paths for neutron-induced fission in a core of typical composition, for a reasonably fast neutron spectrum), this is basically the case. The primary function of the neutron reflector in such core designs is to drastically reduce the fast neutron fluence seen in the outer portions of the reactor, such as its radiation shield, structural supports, thermostating modules and outermost shell. Its incidental influence on the performance of the core is to improve the breeding efficiency and the specific power in the outermost portions of the fuel, though the value of this is primarily an enhancement of the reactor's economic efficiency: outlying portions of the fuel-charge are not used at low overall energetic efficiency, but have burnup levels comparable to those at the center of the fuel-charge.

Final, irreversible negation of the core's neutronic reactivity may be performed at any time by injection of neutronic poison into the coolant-gas stream, likely the primary loops which extend to the surface but possibly also the afterheat-dumping loops connecting the reactor to the engineered heat-dump. At the present time, it appears that lightly loading the coolant-gas stream with a material such as BF3, possibly accompanied by a volatile reducing agent such as H_2, will serve satisfactorily to deposit metallic boron rather uniformly all over the inner walls of the coolant-tubes threading through the reactor's core, via exponential acceleration of the otherwise slow chemical reaction $2BF_3 + 3H_2 \rightarrow 2B + 6HF$ by the high temperatures invariably found there. Boron, in turn, is a highly refractory metalloid, and will not migrate from its site of deposition. Its more-or-less uniform presence in the core in <100 kg quantities will suffice to negate the core's neutronic reactivity for indefinitely prolonged intervals without involving the use of powered mechanisms in the vicinity of the reactor.

During its operational lifetime, the reactor core generates several dozen kilomoles of neutrons in excess to the minimal needs of its neutron economy. These are available to make up losses due to parasitic absorption by structural materials, to leakage, etc. In the design which we overview in this paper, we have conservatively allocated half of this excess to make up such losses, and half to be absorbed in the thermostating modules by Li^6. In a more optimized design, the total losses due to structural materials and leakage will be increased to ~80% by introduction of more such material in order to realize even lower peak stresses and higher long-term reliability of proper system performance, and perhaps for other functions such as internal support to the nuclear fuel against gravitational loads. (After all, the Li^6 merely functions as a "neutron dump," in analogy to the "engineered heat-dump" reviewed in Appendix D; structure performs functions of positive utility.)

APPENDIX C: THERMAL & MECHANICAL DESIGN

The reactor's basic function is to serve as a high temperature heat source; thermal energy is constantly removed from it and used to generate electricity. In order to maximize both the safety and the economic return of this reactor, we require it to function in a hands-off fashion over a long interval. Accordingly, the reactor must safely and reliably operate, without maintenance or repairs, at high temperature for many years. These requirements, together with particular features associated with the reactor's nuclear operation, dictate fundamental features of the thermal and mechanical design of the reactor.

The reactor generates huge amounts of thermal energy, which must be continually removed in the form of ~2 GW of thermal power. From the global viewpoint, this removal is done to generate electricity but, from the reactor's view, its purpose is simply cooling of a magnitude always sufficient to maintain core temperature compatible with structural integrity. The reactor is designed with two sets of coolant loops: a large primary set for use during normal operations, and a smaller secondary loop-set for removal of afterheat during loss-of-coolant accidents and for multi-century intervals after the reactor's shutdown.

The coolant loops contain both in-core and out-of-core portions. The in-core parts are particularly challenging because they function at high temperature in a severe environment and cannot be accessed or maintained. By contrast, the out-of-core portions of the primary loop-set are cool, readily accessible, and can be serviced normally. The out-of-core portion of the secondary loop-set must be designed for hands-off, no maintenance operation, but is relatively cool and located in a much more benign environment than the core of the reactor. Thus, in the present discussion, we'll concentrate on the design of the heat transfer system internal to the reactor.

The reactor is a high aspect ratio cylinder; its fuel-charge is ignited in its middle and burns axially in both directions; at a given time, most power production occurs in two fairly narrow disks of the fuel-charge. The coolant loops direct coolant-gas flow axially along the length of the reactor in many small pipes, passing through and cooling the two burnfronts. The fuel is placed between the pipes, transferring heat to them conductively. The entire fuel-charge is contained within a pressure vessel which provides the global confinement of the fuel and the reaction products thereof.

There are many possible choices for the coolant and structural materials of the reactor; we will discuss only our baseline-design set. The reactor coolant is chosen to be pressurized helium gas. (Helium has good heat-carrying capability and is inert from both chemical and nuclear perspectives, so its use will not constrain the reactor's performance.) The in-core structural material, which is also used for the coolant pipes, is tantalum. (Tantalum is chosen because of the high temperature regime in which we require the reactor's core to operate and the need to maintain excellent creep resistance over a nominal 30-year full-power operational life, as well as its generally good mechanical workability and chemical corrosion resistance. From a nuclear standpoint, it is not an ideal material, acting as a significant neutron absorber, even when using a fast neutron spectrum; therefore, strict limitations are imposed on the fraction of tantalum that can be used.) The pure element is not employed, as much stronger Ta-based engineering alloys exist; our baseline-design utilizes properties typified by alloys ASTAR-811C and NAS-36. Little is known about how the properties of these alloys change in the course of high-fluence fast-neutron irradiation, although similar materials typically gain strength but lose ductility. Since we will only operate at low strain and stress levels, these alloys should remain serviceable. Another issue centers on the effects of partial transmutation during their in-core service life; some of the tantalum will be changed into tungsten, and some of the tungsten into rhenium. These elements are initially present (W at the 10% level and Re at the 1% level, respectively) in the alloys, and are themselves strong, refractory materials, so this gradual change of the material elemental composition during service-life may be quite benign, although this has not yet been demonstrated. Material selection for the primary fuel-containment vessel is much less challenging, since this pressure-shell will be outside the reactor's core, and thus will operate in a less stringent thermal and neutron environment, as well as not posing a neutronic threat to the core's neutron economy; a high-strength steel alloy will likely be utilized.

Helium gas is used to transport heat out of the core both during normal operations and for afterheat removal; the principal distinctions between these two functions arise in how the gas is pumped and how the heat is ultimately removed, outside of the core. The primary coolant loop transports more heat, but does so under normal operational conditions. Pumps are used to circulate the gas flow between the underground reactor and the surface powerplant. The secondary loop transports less heat, but under more stringent conditions, either during loss-of-coolant accidents or entirely unattended, over very long intervals of post-operational life. We cannot depend on pumps for powering gas circulation in the latter case, but instead will use the natural thermosyphon action present from having a heat source located below a heat sink. (Obviously, heatpipes may also prove to be an attractive option for this passive heat removal function.)

We have focused on two different configurations, indicated in Figure 3, for the placement of coolant pipes within the core; many other possibilities exist. The simpler of the two geometries is an array of three separate-and-distinct periodic manifolds of parallel pipes interpenetrating to form a 2-D hexagonal grid. The fuel is positioned, initially as an open-celled metallic thorium foam, in the space between the pipes. Each piece of fuel is in (physical, and thus thermal) proximity to pipes belonging to each of the three separate coolant manifolds, so there will still be coolant flow proximate to each and every fuel-parcel, despite failure of one or even two of the coolant manifolds. In this particular design-variant, there is not a separate secondary set of coolant loops to remove afterheat; instead, the same coolant pipes and gas which perform the primary cooling also serve to remove the afterheat. While this configuration is effective under normal operations as well as most loss-of-coolant scenarios, the use of common-function coolant-loops does pose some issues.

For this reason, a more complex, but more robust, six-pipe layout has been more seriously considered. The pipes are once again laid out in a 2-D hexagonal close-packed grid, but now consist of two different species: three coolant pipe-sets are integrated into three separate-and-distinct primary coolant manifolds, while the other three separate-and-distinct coolant manifolds comprise an independent secondary cooling system. Any given parcel of fuel is thermally well-coupled to each of the six different coolant-pipes, so we can maintain local cooling for both normal operations and loss-of-coolant accidents, if one or even two of each of the primary and secondary coolant system were to fail.

The actual design of a reactor coolant system must simultaneously satisfy several conditions. Clearly, we must have sufficient coolant flow to remove the reactor's heat at acceptable peak temperatures, and enough pressure-head must be available to push this flow through the reactor's core. The coolant pressure is selected by balancing its flow-enhancing benefits against its structural penalties. Large-diameter coolant pipes and a correspondingly large aggregate flow area ease the pumping task, but increase the internal temperature drops within the fuel and across the pipe-to-gas-flow boundary layer. One way to reduce these temperature drops is to reduce the packing density of the fuel, but doing so increases the core's volume and its fuel inventory. There are clearly tradeoffs to be made during the detailed engineering design of a reactor; we present one particular set of preliminary choices below.

The physical dimensions of a six-pipe reactor core of 2 GWt rating are given in Table 1. The primary and secondary coolant-pipes have the same diameters (although not the same wall thickness) and are closely packed (i.e., are nearly touching).

Table 1.

Length	**Diameter**	**Depth**	**Fuel Mass**	**Total Pipe Area**	**Pipe Diameters**
Meters	*Meters*	*Meters*	*Metric tons*	*Meters*^{2}	*Centimeters*
10.0	2.4	100	75	1.3	0.45

The basic power and thermal parameters of this reactor are given in Table 2.

Table 2.

Peak Specific Power	T_{in} - Inlets	T_{out} - Primary	T_{out} - Secondary	ΔT - Into flow	ΔT - In fuel
Watt/gm	*Deg K*	*Deg K*	*Deg K*	*Deg K*	*Deg K*
200	500	1000	1200	150	250

We require circulation of 770 kg/sec of helium in order to extract the full 2 GWt power from the reactor and 33 kg/sec in the secondary loop to remove peak afterheat. The in-core pipes dominate the total loop flow resistance; to reduce this as much as possible the coolant flow is split into two streams, each traveling half the length of the reactor, i.e. between an outlet at the middle and two inlets at each of the two ends. When calculating the pumping work needed to circulate the coolant, we've assumed that only two of the three pipe manifolds are operational; if all three are functional less work is required, while if only one is available the reactor will automatically operate at less than full power.

In Table 3, we list the inlet pressures needed to insure flow circulation under these conditions, as well as the amount of pipe material used and the resultant peak stresses encountered. During normal operation of the reactor, the pipe walls in both the primary and secondary coolant loops reach peak temperatures of about 900^{o}C and are subject to creep. The reactor core will normally be pressurized (thereby shifting the task of pressure containment away from the high temperature, neutronically irradiated pipe walls and onto the relatively cool, unirradiated material of the reactor vessel); however, the pipes are designed for adequate creep resistance anyway. During a loss-of-coolant accident, the secondary-loop coolant-pipes reach peak temperatures of 1100^{o}C, but this lasts only a comparatively very brief interval, rapidly dropping as the core's afterheat power declines; these higher temperature loads thus cause prompt, but not creep, stress.

Table 3.

P_{in} - Primary	P_{in} - Secondary	Wall - Fuel Ratio	Creep Stress	Prompt Stress
Bars	*Bars*	*Atom-fraction*	*Kpsi*	*Kpsi*
53	84	0.33	35	43

We note in passing that we also provide a modest-scale heatpiped-based heat-transport system coupling into the sand immediately surrounding the reactor in order to thermally ballast the reactor's pressure-shell during loss-of-coolant accidents and after end-of-operational-life. Its mechanical integrity relative to hot creep is thereby ensured.

APPENDIX D: ENGINEERED HEAT-DUMP

All nuclear reactors continue to generate thermal power via nuclear afterheat long after their reactivity has been negated. While this power is at most 6-7% of the reactor's output power prior to shutdown, it is still far greater (by as much as a factor of 10^4) than can be tolerated within the reactor core or its immediate surroundings; the afterheat energy must be extracted from the core and disposed of. There are two issues associated with such heat rejection: the long time-interval over which the rejection must occur, and the requirement to reject this heat at peak power levels during loss-of-coolant accidents, when all or part of the normal heat-transport system of the reactor may be unavailable. In order to aptly address these issues, we have designed the new type of reactor with a passively operating coolant-transport system, discussed in Appendix C, which is capable of removing afterheat at the required power levels from the reactor itself. In this Appendix, we briefly discuss the other necessary aspect of afterheat removal: where the afterheat energy goes after it is removed from the reactor core.

We design to deposit the entire afterheat energy into a simple thermal sink, utilizing the heat capacity of loosely-packed sand under which the reactor is emplaced. The primary attraction of such a choice is that this overburden will always be present, while convective sinks might not be, given the long time-intervals and severe loss-of-coolant scenarios for which we believe any power reactor system must be designed. Another significant advantage is that this heated overburden will, during the long interval of post-operational life, act as something of a shroud, thermally inhibiting the permeation of groundwater into the immediate vicinity of the reactor , particularly if it is utilized with clay barriers deployed as in modern engineered sanitary landfills. (Water may thereby be confidently excluded as a possible corrosion agent which might impair the reactor's role as the spent-fuel's burial-cask.)

The total afterheat energy which the engineered heat-dump must absorb is, in the worst-case of abrupt end-of-life shutdown of the reactor from full-power-production without benefit of any post-end-of-operational-life cooling, about 1.5 x 10^{15} J (~0.35 MT). Since the material removed to emplace the reactor is unlikely to be an optimal heat-sink material (primarily due to its thermal characteristics), we specify that it be replaced with sand. This material can absorb ~ 1 kJ/cc, starting from room temperature and going up to near-melting, implying typical heat-dump dimensions of a rectangular prism about 130 meters on a side and 100 meters in depth. This volume will be positioned above the reactor, and won't extend significantly underneath it (except for a modest distance, for groundwater-sealant and seismic decoupling purposes). There are two principal reasons for this choice: avoidance of more excavation to greater depth than really necessary, and exploitation of the gravitational head made available by positioning the heat-sink above the heat-source to power a natural thermosyphon which will serve to reliably and entirely passively circulate heat-transfer fluids between the -source and -sink.

We have already described, in Appendix C, the helium gas-loaded thermosyphon which removes afterheat from the reactor and transports it up through the subterranean heat-dump. In addition to this system, we also require a transport system to distribute the afterheat energy in an adequately uniform manner throughout the heat-sink, since the thermal diffusivity of sand is insufficient for this function on the requisite time-scales.

We specify a two-component heat-pipe network to serve as this transport system. The afterheat energy is initially lifted vertically from the reactor and towards the surface by the action of the thermosyphon. At periodic intervals along this thermal trunk-line, heat is coupled into a first set of "arm" heat-pipes, which transport it quasi-horizontally (with a slight upward tilt in order to gravitationally assist heat-pipe functioning) across the width of the heat-dump. A second set of "finger" heat-pipes, thermally tapping into each "arm" at periodic intervals, carries heat sideways (in the direction orthogonal to the "arm"), thereby distributing it throughout the entire quasi-horizontal layer of sand. This arrangement of many parallel heat-pipes, coupled thermally but not fluidically, provides a highly redundant - and thus functionally robust and systemically reliable -- heat distribution system threading the entire engineered heat-dump.

Our present, incompletely optimized design employs a set of stainless-steel heat-pipes whose working-fluid is mercury, and is comprised of approximately 50 "arms" and 2000 "fingers." The fingers are spaced approximately 4 meters apart. This arrangement is sufficient to sink the peak afterheat power of $\lesssim$ 150 MW from our baseline-design reactor, and to adequately dispose of the monotonically declining afterheat power at all times thereafter.

During the initial phase of heat-dump operation, the afterheat energy heats only the cylinders of sand proximate to the fingers. After approximately a month of operation, these cylinders begin to merge into each other and the entire heat-dump is thereafter heated semi-uniformly. The array of finger heat-pipes directly contacts sufficient sand to readily disperse the earliest phases of afterheat, at times when the afterheat power level is highest and thermal diffusion has not spread this heat-pulse into the sand very far from the finger pipes.

As noted in Appendix C and depicted in Figure 6, we also employ thermostatically-actuated thermal switches (with a contrast-ratio of >100:1). These serve to connect the heat-dump with the ascending coolant-pipes only when the reactor core's temperature exceeds the design set-point, thereby adequately preserving the thermal capacity of the heat-dump for use when needed.

APPENDIX E: COOLANT-PIPE CLOSURES

In order to provide extremely reliable, very swift closing of coolant pipes under emergency conditions (e.g., when significant fission-product entrainment in coolant-gas is sensed), we provide one-time-operation closures which are known to offer the required features. Such fast closure systems would be used in addition to a redundant system of automatically-actuated conventional valves.

In underground testing of nuclear explosives and their military effects, it is often required to close steel pipes of 1-3 meters diameter enclosing gas or vacuum interiors on time-scales of milliseconds, with extremely high reliability of complete closure and maintenance of pre-existing internal conditions. Highly engineered, high-explosive-energized means for plastically flowing the relatively thick-walled pipes of carefully chosen composition into a cold-welded-closed configuration have been demonstrated repeatedly to be eminently suitable for such functions.

Use of such pipe closures precludes contamination of high-value experiments by explosive debris transiting the pipe's interior at far higher speeds than would possibly occur under even emergency circumstances in the situations of present interest. These closures are of such quality that they maintain high-vacuum conditions downstream of the closure-point after the closure-event (in spite of the appearance of occasionally very high pressures on the upstream side).

ACKNOWLEDGMENTS

We gratefully acknowledge helpful discussions with many colleagues on nuclear engineering issues. Tom Reed provided a great deal of useful technical and politicoeconomic feedback, as well as overall impetus. George Zimmerman offered characteristically insightful guidance on computer-based modeling issues, and Dermott Cullen has aided in many ways, most particularly in the peculiar use and with respect to *ad hoc* performance enhancements of the TART95 code package, as well as in creation of illuminating graphics. Gordon Wenneker contributed in a most appreciated manner in the creation of presentation aids. Victoria Wood graciously and patiently donated use of her personal computer.

UCRL-JC-122708 Pt 2

Last Changed Date: 02-11-97-- By: kab

URL: http://www-phys.llnl.gov/adv_energy_src/ICENES96.html