208
8-1-78

Hr. 305

BNL 50817
UC-21
(Inertial Confinement Fusion -
TID-4500)

MASTER

CONCEPTUAL DESIGN OF A HEAVY ION FUSION ENERGY CENTER

A.W. Maschke
March 13, 1978

ACCELERATOR DEPARTMENT
BROOKHAVEN NATIONAL LABORATORY
ASSOCIATED UNIVERSITIES, INC.
UPTON, NEW YORK 11973

under contract No. EY-76-C-02-0016 with the
UNITED STATES DEPARTMENT OF ENERGY

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency Thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Synopsis

A Heavy Ion Accelerator system is described which is based upon existing technology, and which is capable of producing 150 MW of average beam power in 10 MJ, 200 TW bursts, 15 times per second. It consists of an rf linac which accelerates doubly ionized uranium ions to an energy of 20 GeV. Then by utilizing the well known procedure of multiturn injection, a 6.6 ms long burst of linac current is stored in 8 separate "accumulator" rings. At the conclusion of the filling process, a pulsed rf system bunches the beam in each of the 8 rings simultaneously. As the bunches decrease in length, they are then extracted from the rings and transported for about 1 km to one of 5 "boilers", in which the thermonuclear pellet has been placed. The 8 beams (2 opposing clusters of 4 beams each) are then focused simultaneously onto the pellet, resulting in a release of thermonuclear energy about 80 times larger than the input beam energy.

NOTICE
This report was prepared as an account of work sponsored by the United States Government. Neither the United States nor the United States Department of Energy, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness or usefulness of any information, apparatus, product or process disclosed, or represents that its use would not infringe privately owned rights.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Table of Contents

1. INTRODUCTION

This report describes the ignitor portion of a Heavy Ion Fusion energy complex. For the sake of this discussion, we assume that electricity is the principal output product. A plant size of 4 GW_e was chosen, consisting of 5 independent 1 GW_e boiler systems. The ignitor is expected to have an availability greater than 95%, whereas one of the 5 boilers will generally be shut down for maintenance. There are two classes of ignitor systems, intensive and extensive. E-beams and light ion beams are intensive, in that they are centered around a single boiler. Laser systems, and heavy ion accelerators are extensive, in that their beams can be transported substantial distances to a number of different boilers. In general, any extensive system will improve its economics by increasing the repetition rate until further increases either become technically impractical or no longer cost effective. In the case of an rf linear accelerator, the limit is around 10% duty cycle. The accelerator considered produces a beam of 20 GeV U^{2+} ions at a current of 160 mA. It pulses 15 times/sec and produces 10 MJ/pulse, and 150 MW average beam power.

1.1. Plant Size Assumptions

The general philosophy adapted for this design has been to take the most conservative approach possible and still arrive at a potentially viable energy option. Perhaps the most controversial assumption is related to the overall size of the plant. EPRI has expressed the desire of some utilities to build small plant additions. The average size of power plants has been doubling every 10 years for the past 40 years. As generator and boiler technology develops, and siting problems become more severe, the trend towards larger units, and larger stations (i.e. clusters of independent units) can be expected to continue.

In the case of a "high technology" power generation scenario, the cadres of skilled construction and operating personnel become an important consideration. This is one of the reasons why nuclear power stations tend to be built in clusters. For example, in Brazil, an 8 GW_e power station, consisting of 8 one GW_e nuclear reactors is under construction. Similar projects are being planned in Iran and Saudi Arabia.

For a heavy ion fusion power complex there must exist a skilled crew of operators to run and maintain the accelerator systems, and the tritium recovery systems, and the pellet fabrication facility, as well as the boiler system. All of these considerations tend to favor the large installation.

In addition, any fusion plant is going to have to contend with the tritium problem. This can be broken down into several categories.

1.1.1. Safeguards

Tritium is a weapons material, and as such is carefully controlled, as is plutonium. Whereas the situation is not entirely analogous to plutonium, there is nevertheless a "proliferation" aspect to the large scale production of tritium.

Two separate concerns exist. A non-nuclear, or non-thermonuclear nation, with an inertial confinement fusion facility could, conceivably, obtain information from the pellet explosions which would be of use in designing a thermonuclear weapon. Secondly, the large neutron fluxes available make it relatively easy to produce fissile material, either surreptitiously or on short notice. This might, in fact, be easier to do in a fusion plant than in a nuclear power reactor. Since tritium is very poisonous, and easily dispersed (which is not the case of Pu) into the biosphere, it presents a formidable threat in the hands of terriost groups.

1.1.2. Accidental Release

Because of the relatively large inventory of tritium in the boiler system, pellet fabrication facility and tritium recovery plant, there are numerous mechanisms by which accidental tritium releases could occur. At the present time the principal insurance against public damage in the event of an accident, is to make sure the site is large enough that at the site boundary the tritium will have dispersed sufficiently to insure that the relevant MPC's are not violated. This safety factor is another reason for preferring rather large systems, as opposed to many scattered small ones.

1.1.3. Tritium Recovery

The tritium produced in the thermonuclear burn, plus whatever is produced in the boiler, must be recovered, and returned to isotopic purity for use in the pellet factory. Detailed designs of such a facility have not yet been published. However, there is reason to believe that this recovery system will not be a negligible portion of the plant cost. This will also tend to push the economics towards larger sizes.

1.2. Input Energy Assumption

Current thinking on the part of LLL pellet designers places moderate confidence in 1 MJ being adequate energy to obtain an energy gain of $\sim$ 100. No doubt the confidence level increases as one goes to 10 MJ. Furthermore, the 10 MJ assumption allows one to pay 10 times more for a pellet which is probably easier to fabricate. Preliminary design of boilers to contain 10^9 joule explosions look reasonably tractable. Furthermore, the increased input energy comes at a rather modest cost. If 10 MJ was insuffficent, a 40 MJ machine could be built for about 50% more money. This is probably close to the credible limit for commercial HIF, i.e., if 40 MJ is not adequate, then the chances are the whole idea can be dropped.

1.3. Input Power Assumption

An input power of 200 TW was assumed. This left a factor of two "safety" factor over the pellet designers moderate confidence case. Also, it can be relatively easily achieved without resorting to space charge neutralization, a possibly over-conservative consideration. We note that the 40 MJ case, for 50% more money, could also produce an increase in power of a factor of 4 to 800 TW.

1.4. Pellet Gain Assumption

The pellet gain of 83 was selected to make the net power output of 4 GW_e. Not accidentally, of course, this is sort of the mid range of assumed pellet gains. At a gain of $\sim$ 40, the power plant output would drop to $\sim$ 1.8 GW_e. Since the ignition system cost is $\sim$ 900 M$, this is probably pretty close to the smallest gain which economics would allow. Lower gain scenario's can be made technically satisfactory by imploying fission blankets to increase the gain (perhaps by a factor of 5). It is most improbably that such a hybrid system could be made cost-competitive with a fast reactor, and has therefore been dropped from consideration here.

2. OVERVIEW OF THE PLANT OPERATION

There are two principal reasons for producing this design study. The first reason is to illustrate in some detail all of the steps which must be taken to obtain high energies and power, so as to make it apparent that in fact existing accelerator technology can produce the appropriate hardware. The second purpose is to produce a first pass at a cost estimate for such a machine. The bottom line for any inertial fusion driver is the cost per watt of average beam power. This particular scenario comes in with a capital investment of about $6 per watt of average beam power.

In arriving at the cost for the ignitor, we did not include the cost of the power plant (exclusive of ignitor) needed to run the ignition system. This is natural enough, because we have not, in this study, made an effort to cost the rest of the facility. In any complete analysis, this capital cost would be added on. Suppose we took 1 dollar/watt for the cost of the remainder of the plant. This is near the current power plant figure. The HIF accelerator requires 400 MW to operate it. (Efficiency of 37.5%). This would bring the "cost" of beam power from $6/watt to $8.66 watt. Note that if the ignitor was only 5% efficient, this power related cost alone would amount to $20/watt. This would make the low efficiency ignitor uneconomical even if its capital cost was negligible.

2.1. General Description

For the purpose of this study, it has been assumed that uranium ions are accelerated. In fact, one could use mercury, gold or bismuth without having a significant effect on the design. Prior to actually building such a plant, one would, no doubt, study the pro's and con's of different species quite carefully. The ions start out of a rather conventional source, and are accelerated first in a 500 kV high-gradient DC column. A beam of 50 mA of U^{1+} is produced at the end of the column. It would be preferrable to have a higher current, and a higher voltage. Both choices here are made because they are rather conservative. Performance at this level has not been attempted, to my knowledge, but does not represent a significant extrapolation from electromagnetic isotope separation experience. The design current for the linac is 160 mA. We obtain this by starting with 8 linacs of 20 mA each. As the ions gain energy, the bunches of beam are combined, until finally all the current is in a single structure. The accelerator starts out as a cascade of eight 2 MHz Wideroe linacs, injecting into four 4 MHz Wideroe linacs. At about 6 MeV we strip to U^{2+}, with about 50% efficiency, i.e., the current of U^{2+} is still the same as it was for the U^{1+}. At about 13 MeV we combine the beams into two 8 MHz Wideroe structures. At 30 MeV, the beams of 80 mA each are combined in a 48 MHz Alvarez linac. At 120 MeV this beam is matched into a 96 MHz Alvarez, and at 480 MeV into a 192 MHz Alvarez. The 192 MHz structure is continued until the final energy of 20 GeV is reached.

At this point the beam is injected into a very long (2π kilometer) "multiplier" ring. Ten turns are injected into this ring by means of multiturn injection into the horizontal phase space. At the completion of the 10 turn filling, the beam is extracted, the horizontal and vertical phase planes are exchanged via a series of skew quadrupoles, and the beam is multiturned into the aperture of a multiplier ring of 100 meter radius. We have now a current amplification factor of 100. This beam is now transferred to one of the 8 waiting accumulator rings. The linac beam could have been multiturned directly into the accumulator. However, because small losses occurring during multiturn could affect the vacuum, and since the vacuum requirement in the multiplier rings is at least an order of magnitude less than for the accumulators, it is safer to have separate rings for this purpose.

After all 8 accumulator rings are filled the beams are bunched with a low frequency (1st harmonic of the rotation frequency) pulsed rf system. This starts a longitudinal "implosion" of the bunch which carries the beam a factor of $\sim$ 10 over the "space charge limit" of the accumulator. This is possible because of the transient nature of the implosion. At this point the beams are extracted, and another factor of 5 increase in the current occurs in the $\sim$ 1 km drift from the accumulator to the boiler. At this point the 8 beams are simultaneously focused onto the pellet, with an instantaneous current of 2500 amperes in each of the beams.

The entire cycle takes $\sim$ 6.6 ms, and is repeated 15 times per second. The beam is transported alternatively from one boiler to another.

Each boiler operates 3.75 times/sec.

2.2. Summary Cost Estimate

It is not possible to make an accurate cost estimate without a detailed design. However, the machine components are familiar enough that costs for similar items of known cost can be applied without too much chance for error. An uncertainty factor for this estimate is probably no more than ± 25%. The estimate does not take into account cost savings that might be made by going to large scale production techniques. These techniques are generally not well known to accelerator builders, whose bread and butter has been one of a kind research machines, generally of much larger size. The philosophy has been adopted here that the cost savings obtained through large scale industrial production will be offset by reliability requirements, which for this system is much greater than for existing research accelerators.

2.2.1. Accelerator Systems	Millions of Dollars	
2 MHz Linac/Rf	4	
4 MHz	24	
8 MHz	8	
48 MHz	9	
96 MHz	26	
192 MHz (first stage)	23	
192 MHz (final sections)	258	
		352
2.2.2. Storage Rings/Transport		
8 Accumulators	80	
10 km Transport Lines	80	
Multiplier Rings	40	
		200
2.2.3. Conventional Construction		
Linac Housing	65	
Accumulator Tunnels	10	
Transport Tunnels	45	
Multiplier Enclosures	10	
Misc. Transport Enclosure	5	
		135
Total		687
EDIA (15%)		103
Contingency (10%)		79
Total Cost		869 M$

2.3. Overall Plant Economics

A recent EPRI study concluded that new electrical generating plants will have to produce electricity for between 3.5 to 4.5 cents/kW hr in order to compete with existing technologies. (Coal and nuclear). If this plant sold power for 4¢/kW hr, it would have a gross income of 1.23 billion dollars/year. This assumes 70% availability for the 5 stations. With these assumptions, one can proceed to assess various economic factors associated with the plant.

Many economic factors are not known at this time. The key unknowns are pellet costs and boiler costs. More "exotic" unknowns, such as interest rates and taxes, are not only unpredictable, but arbitrary parameters in the administrative space of Government.

The plant expenses can be divided between operating expenses and finance charges. The ignition system can be assumed to have an annual operating cost on the order of 10% of the entire capital investment ($\sim$ 87 M$/year). The rest of plant is probably 5% of capital cost/year, but this is just a guess based on existing plants.

If we assume values for amortization and interest on capital, say 12%, and assume 5% of gross income for taxes, then we can find a relation between pellet costs and plant costs/kW, exclusive of the ignitor.

If pellets were free, we could afford to pay 1100 $/kW for the boiler/generator system. On the other hand, with this set of cost assumptions, pellet costs of 1 dollar each would require us to build the boiler generator system for about 600 \$/kW. Since we have 5 GW_e installed capacity, the total plant cost would be 3.87 billion dollars. Since the plant uses 410 million pellets/year, it is likely that much of the cost would be associated with the capital investment in the pellet factory. Since the breakdown in pellet costs between capital and labor is not known at present, it is sufficient to merely specify a total cost.

It is worth noting that if the ignitor cost was twice as high, the cost of electricity would rise by about 12%. On the other hand, if pellets cost 2 dollars each, the price of electricity increases by 33%. The point is, the ultimate economic viability depends upon boiler and pellet costs, not on ignitor costs.

3. ACCELERATOR DESIGN

No substantive effort was made in this study to optimize the design in any real sense. Rather, the emphasis was put on exhibiting a design that requires the least departure from existing technology. It is important to emphasize that it is only within the past year that a development effort has been started to advance the state of the art in the area of heavy ion accelerators for inertial fusion. Therefore a similar study, started in a few years from now, could be expected to incorporate many new features which are at present only in the "concept" stage.

3.1. Preinjector and Ion Source

For purposes of this study a 500 kV Cockcroft-Walton accelerator is taken for the DC terminal. This is a conservative choice between a desire for high voltage to alleviate space charge problems versus a fear of breakdown and contamination damaging a higher voltage accelerating column. Extensive experience at GSI at 320 kV indicates that one could easily go somewhat higher in voltage. If 400 kV had be chosen for the terminal voltage, the overall design of the facility would not be altered appreciably.

Ion sources of a type suitable for injection into a preinjector acceleration column have been developed for protons with currents on the order of an ampere. Child's law for the extraction current density one might obtain requires that the current vary inversely with the square root of the atomic number. This means that we might expect currents about 15 times smaller for heavy ions. The performance assumptions made here require a beam of about 40 mA from each of 8 accelerating columns. Isotope separation sources, developed over the past 35 years, have routinely produced heavy ion beams of currents higher than required here. For purposes of this study, U^{1+} has been selected as the ion to be accelerated in the "pre-stripper" portion of the accelerator. The final choice would be decided on the basis of rather subtle differences between species. Isotopic purity, ion-ion cross-sections, stripping considerations, etc. will all play a role in the final choice. None of these considerations are expected to make a significant difference in either the design of the facility or its performance. If a particular species was found to have unusually small ion-ion charge-exchange cross-section one might chose to alter the scenario to take advantage of this larger accumulator lifetime. However, using a geometric cross-section as an upper bound, the system considered here loses only $\sim$ 1% of the beam.

3.2. Low β Linac Portion

A 500 keV heavy ion has a velocity of $\sim$ .002 c. This is about a factor of 2.5 times lower than any existing heavy ion accelerators.* Because the drift tubes become so small, it is necessary to go down

---

* A Model Heavy Ion Linac with $\beta \gtrsim .003$ has operated successfully at BNL.

in frequency as the velocity decreases. If one took the GSI Wideroe linac as an example, one might consider scaling that to 2.5 times lower frequency, i.e. ∿ 10 MHz. However, space charge forces are another factor which must be taken into account.

Longitudinal space charge forces become more severe as the bunches become shorter. Therefore, the maximum transportable current is inversely proportional to the frequency. If one assumes that some ratio of longitudinal space charge force applied to rf focusing constitutes a longitudinal current limit, then the following relationship follows.

$$i_{max} \propto \frac{E_p \bar{E}}{fA}$$

E_p is the kinetic energy of the ion, $\bar{E}$ is the average accelerating field, f is the frequency, and A the atomic number.

An estimate of what can be expected can be obtained empirically by examining the performance of an accelerator believed to be operating near its longitudinal space charge limit. The FNAL 200 MeV proton linac is probably the best example. Using the FNAL peak current figures of ∿ 300-400 mA, one obtains the following relation for heavy ions;

$$\frac{f}{\bar{E}} \approx 10 \qquad i_{max} \approx 20 \text{ mA}$$

For this study $\bar{E}$ was ∿ .2, and f = 2 MHz. This choice is not completely arbitrary. For instance, as one increases $\bar{E}$, and increase the frequency at the same time, the drift tubes become shorter, the aperture becomes smaller and the transverse focusing requirement becomes more severe. In the model described here, the longitudinal "synchrotron" oscillation frequency (ω_L) is below that of the transverse "betatron" oscillation frequency (ω_T). This situation remains throughout the linac. A choice of high gradient and high frequency could lead to a situation where $\omega_L > \omega_T$. As the particles gain energy we will eventually have to go to $\omega_T > \omega_L$. The coupling of transverse and longitudinal motion can give rise to emittance blow-up, and is to be avoided where possible. That is not to say that a higher frequency system may not be workable, but only to indicate a complication. In the spirit of this study, an approach was chosen that would exhibit the maximum likelihood of success with a minimum of computational effort. Obviously there is a lot of work to be done in the design of high current low-β accelerators. However, one should keep in mind that the low-β portion represents only about 5% of the entire ignitor system cost. If the number of systems were doubled, the cost difference would be scarcely noticeable.

So far the transverse space charge forces have not been discussed. The scaling here can be represented by the familiar expression:

$$i_{max} \propto \left(\frac{A}{Z}\right)^{1/3} \varepsilon_T^{2/3} (\beta\gamma)^{7/3} B_{p.T.}^{2/3}$$

where ε_T is the transverse emittance/π, and B_{pT} is the pole tip field of the quadrupole. One obtains a similar expression for the longitudinal current limit:

$$i_{max} \propto \left(\frac{A}{z}\right)^{1/3} \epsilon_L^{2/3/} \left(\frac{\bar{E}}{g_0}\right)^{2/3} (\beta\gamma)^{5/3}$$

where ϵ_L = longitudinal phase space/$\pi \equiv (\frac{\Delta p}{p} \cdot \ell)$.

ℓ = bunch half width $\frac{\Delta p}{p}$ = bunch momentum half width

$\bar{E}$ = average accelerating field

$g_0 \approx (1 + 2 \ \ell n \frac{b}{a})$ a = beam diameter

b = pipe diameter

Table I shows the sequence of events as one goes from the 2 MHz, 20 mA situation to the 160 mA, 192 MHz condition. We have taken the ratio of beam current to space charge limiting current to be unity for both transverse and longitudinal at the beginning of the system. Note that the most severe problem occurs at the injection into the 48 MHz Alvarez. If one was limited at injection into the 2 MHz structure, then an emittance blow-up of about $(1.34)^{3/2} = 1.55$ would be expected.

If we assume an adiabatic damping of the longitudinal phase space, and further assume that the phase length of the beam at the entrance to each new linac system is the same, we then find the following relation;

$$(z\bar{E})^{1/3} \frac{\beta}{f} = \text{const.}$$

This allows one to determine reasonable β values at which to jump the frequency without losing beam. For different frequency linacs

TABLE I

E_{in}	z	(β)	f	i	$\frac{i}{i_{TSC}}$	$\frac{i}{i_{LSC}}$
500 keV	1	.002116	2 MHz	20 mA	—1—	—1—
2.0 MeV	1	.004232	4 MHz	40	.63	1
6.4 MeV	2	.007570	4 MHz	40	.32	.704
12.8 MeV	2	.010706	8 MHz	80	.36	1
30 MeV	2	.016390	48 MHz	160	1.1	1.34
120 MeV	2	.032781	96 MHz	160	.69	1.34
480 MeV	2	.065562	192 MHz	160	.43	1.34

with the same average electric field, one sees that one must double β if one wants to double the frequency. For this example, of one did not want a bunch of greater phase length than that at the beginning of the 2 MHz structure, one obtains:

$$\beta \geq 6.2 \times 10^{-4} f \cdot (z\bar{E})^{1/3}$$

where f is in MHz, and $\bar{E}$ in MV/meter.

3.3. Low β Alvarez Portion

After a suitable length of 8 MHz Wideroe linac one can jump to an Alvarez structure at 48 MHz. The longitudinal acceptance is increased in the Alvarez because it has an average accelerating field in the neighborhood of 1 MV/meter. The first set of Alvarez tanks are ∿ 45 meters long. Then the frequency doubles, and after 135 meters of 96 MHz structure we go to 192 MHz. This is the frequency which will be kept for the remainder of the linac. We include about 80 meters of the 192 MHz linac in the low β portion because the electric field is still maintained at a relatively low level. Also, because the velocities are changing so rapidly in this portion, if a single rf system should fail, the entire beam would be down. In the later portion this will not be the case.

Alvarez linacs in both this frequency range and velocity range have been built previously and present absolutely no new scientific or technical problems. The beam current is another story. More than

50% of the rf power will be going into the beam. While the FNAL linac has done much higher currents, the pulse length was rather short, and depended upon energy stored in the cavities. The highest long pulse currents are about 100 mA at the BNL linac (pulse length ∿ 200 μsec). The 160 mA assumed for this design is 60% higher than that, but is not expected to present any serious problems.* The duty cycle assumed is 10% maximum. This is relatively modest compared to the 25-35% duty cycles used in existing heavy ion linacs.

3.4. Alvarez High β Section

The high β portion of the linac is not like any existing linac. Whereas the Alvarez structure obtains its best shunt impedance in the range between β's of .1 to .4, the existing proton linacs in this velocity range are of necessity quite different. For one thing, a 200 MeV linac has only a few tanks in this velocity range. Large scale production techniques which require extensive tooling were not a design option. Furthermore, while a proton will go through this velocity range in ∿ 50 meters, the heavy ion linac requires about 5 km. The change in structure from one tank to another is almost negligible. At the beginning of this section of linac, a synchrotron oscillation is about 82 meters long. Tanks 6 meters long would add 20 MeV to the beam. If a single 6 meter cavity was turned off, a 10% increase in the acceleration of 5 upstream and 5 downstream cavities could compensate for this loss. This fact makes the reliability of the high-β portion very much greater than it would be otherwise.

---

*Some design studies for high current linacs used to breed fissile material have considered currents as high as 300 mA.

3.4.1. Reliability Consideration

Because Heavy Ion Fusion power stations are likely to be larger than conventional single unit power sources, the reliability of the ignitor is of rather greater concern. At a conventional or nuclear energy center, single units are about 1 GW_e. In this scenario, the energy center produces 4 GW_e, and if the ignitor fails, there is a total interruption. It is clearly desirable to make such interruptions as infrequently as possible. An approach to this is to build in as much redundancy as possible. Existing large power stations have availabilities around 70%. An interconnected grid of 16 plants might be expected to have 4 GW "off-line" at any given time. In general, they would not go off-line simultaneously, and a typical abrupt change would involve only 1 GW_e. There would be plenty of time to bring up the power level in the other plants to make up the 4 GW. When the ignitor for a 4 GW system stops, there is the necessity to either abruptly shed load, or to increase the power level of other plants on the grid. This is a problem which must be studied for each specific application to determine the particular economic consequences of a 4 GW plant going down.

In this regard, it is worth noting that base-load electrical power generation is the most demanding use to which an energy center could be applied. Desalinization plants or irrigation pumping applications are not nearly as perturbed by a sudden loss of power. Another interesting characteristic of ignitor failure is the relatively short time for repair. An rf module can be replaced within a half-hour time span. Since two would have to fail in the same region (within 5 cavities) for a beam failure to occur, 15 minutes would represent an average interruption. If each of the 1000 rf modules has a 5000 hr mean time to failure, one unit would fail every 5 hrs. Given the half hour replacement time there is a 10% chance of another unit failing, and a .5% chance that it is close enough to interrupt the beam. Therefore, a 15 minute disruption would be expected every 10,000 hrs. or about once/year.

The vacuum system is easily capable of providing redundant pumping speed to maintain the required vacuum. Experience with large storage ring vacuum systems have shown that extremely reliable and leak-free systems can be built on the scale required here. Similar considerations apply to the focusing magnets.

An important factor for maintaining a high degree of reliability is preventative maintenance. Experience with existing machines would indicate that shutdown on an order of 8 hrs every two weeks would suffice for all routine maintenance procedures, and would even allow time to replace a linac tank if required. Whether or not extended shutdowns (a few weeks, say) would be required on an annual basis remains to be studied. All existing research accelerators have such shutdowns because of new developments being performed, or for fiscal reasons.

3.4.2. Design Considerations

The cavity design for the approximately 1000 cavities required in the high-β section is especially simple. Because the ions have an energy of about 1 GeV at the input end, the ratio of voltage gain/gap to total kinetic energy is very small (∿ 200 times smaller than for protons). What this means is that the effect of gap-defocusing is very small, and can be ignored without effecting anything. Also, the transverse emittance is about 10 times smaller than for a similar proton linac. The consequence of this is that the drift tubes do not require magnetic lenses placed in them. These two factors allow one to make substantial design simplifications. The focusing elements can be inserted in the inter-tank regions, where their outer diameter is not constrained, and they can be simply maintained. The drift tubes themselves need never be aligned, because there are no lenses in them. The tank becomes a simple welded steel structure, the inside of which is then copper plated.

The principal cost item for the high-β section is the rf system. It consists of ∿ 1000 2.5-3.0 MW rf drive systems capable of operating with a 10% duty cycle, and 6.6 ms long pulses. A number of options are available for the rf and a more detailed design is required to chose between them. 200 MHz kylstrons are just beginning to enter the market and appear like an attractive solution. Prices are still relatively uncertain. Triodes for this purpose are routine, but they may present more reliability problems. The rule of thumb, 10¢/watt for peak, $1/watt for average does not appear to be far off the mark.

There are definite cost savings associated with buying 1000 units compared with the 10-20 unit sizes that have been customary for existing ion linacs.

An interesting consequence of the small longitudinal phase advance in each cavity is that it is not necessary to provide amplitude modulation (i.e., feedback) on all of the rf systems. Roughly speaking, it is sufficient for only one cavity in 10 to control its amplitude during the pulse to adjust for time dependent beam-current fluctuations or drive fluctuations in the different rf systems.

3.5. Multiplier Rings

These rings are a novel part of this linac/accumulator scenario. The current multiplication in the accumulator is a factor of 100 over the linac current. This is obtained by stacking in the horizontal and vertical phase space. In principal, this could be done by putting 100 turns into a ring in one single operation. In practice, no one does it, and a detailed proposal would require a lot of work. Therefore, in the spirit of this design study, it was decided to make the most conservative technical assumption possible.

There are two multiplier rings. One with a circumference of 2π km and another one with a 100 meter radius. The large ring is a race track shaped and encloses the entire rf linac. The linac injects 10 turns into the horizontal phase space of the long multiplier ring. This multiturn injection process is straight-forward. The first multiturn injection into a strong focusing synchrotron was done at the BNL AGS, and since has been in use at many accelerator laboratories. The next step is to extract the beam from the long multiplier ring, and rotate the beam by 90^0, i.e. exchange horizontal for vertical phase space. This can be done either with a solenoid or with a series of skew quadrupoles (quadrupoles rotated 45^0 from their normal configuration). This beam is now multi-turned into the small multiplier, where once again 10 turn multiturn injection is performed. Upon completion of this multiturn process the resulting beam is adiabatically "bunched" by a small rf system on the 1st harmonic of the revolution frequency. Then the beam is extracted, without loss, and transferred to one of the 8 accumulator rings.

There are several advantages to this process over that of injecting directly into all the accumulators, even if one knew how. First, we require only the two sets of simple multiturn hardware. Since the beam remains in the multiplier only about 1/8 as long as in an accumulator, the vacuum requirement is much less severe. Therefore the chance that beam losses may effect the vacuum are reduced. Furthermore, since multiturn injection is the only function done in the multiplier rings, it is possible to take special precautions with regards to beam losses which would be awkward in the accumulators.

3.6. Accumulator Rings

There are two clusters of accumulator rings. A group of 4 are located in the same tunnel, and placed one on top of another. This technique was used, on a smaller scale, for the booster synchrotron for the CERN Proton Synchrotron. That system consists of 4 25 meter radius synchrotrons stacked a top one another. The accumulator rings are somewhat larger in aperture, and have a radius of 100 meters. There are no especially novel features of such a system. The fields in the magnets ($\sim$ 20 kg) could be produced by conventional magnets. A considerable power savings is produced by using warm iron magnets with superconducting coils. Table II gives a parameter list for the rings.

3.6.1. Transverse Phase Space Considerations

If we take an initial phase space area/π of 2×10^{-4} at the 500 keV input to the linac, and assume adiabatic damping throughout, then we wind up with 20 GeV U^{2+} ions with an ε of 1×10^{-6} meter radians. If we did "perfect" multiturn injection (i.e. no dilution), the area would be increased to 1×10^{-5} m-rad. The beam in the accumulator is assumed to have a transverse ε of 6×10^{-5}. This gives a "safety factor" of six. We expect a factor of two for the multiturn injection and another factor of two dilution in the low energy portion of the linac. These factors are based upon current experience, and conceivably could be improved upon. In any event we are left with a residual safety factor of 1.5. This is not very large, and illustrates

TABLE II

Accumulator Parameters

Radius	100 meters
Average Magnetic Field	1.6 Tesla
Beam Emittance/π	6×10^{-5} m-rad
Revolution Period	5.275 μsec
Average Circulating Current	16 amperes
Storage Time	≤ 6 ms
Vacuum	$\sim 10^{-10}$ torr
Betatron Oscillations/Revolution	~ 10
Vertical Semi-Aperture	5 cm
Horizontal Semi-Aperture	6 cm

the importance of determining the performance of the low energy sections before designing the final portions.

3.6.2. Longitudinal Considerations

Each accumulator acquires the longitudinal phase space area of a total of about 84 of the 2 MHz bunches. Each 2 MHz bunch has an area/π of about .008 volt-seconds. Therefore the entire accumulator has a phase space area/π corresponding to 1.3 volt-seconds, assuming a factor of two dilution in the linac. Now chromatic aberrations in the final focus restrict the momentum spread in the beam. Given the requirement to bunch the beam, in order to obtain the requisite "peak currents", this translates into a limitation on the longitudinal phase space. Taking a ½% value for $\Delta p/p$, and 20 ns for the half width of the bunch we obtain a requirement of 4 eV-seconds. As with the transverse case, we have a factor of 3 safety. A detailed design using chromatic corrections to allow a larger momentum spread could give us another factor of two. It is worth noting that the small safety factors for both the transverse and longitudinal phase spaces could be increased by adding more accumulator rings. If the number of accumulator rings were doubled it would only increase the igniter cost by 20%.

A novel feature of the accumulator rings is the rf system to compress the beam longitudinally. Experiments at BNL have demonstrated that a rapid bunching of the beam can produce beams of higher currents than the space charge limit would imply because of the transient nature of the bunching. Each of these accumulator rings contains 100 small, low impedance cavities. These cavities are driven by a spark-gap switched resonant circuit. A voltage of approximately 10 MV/turn is applied for 20-40 turns at the frequency of the first harmonic, i.e., around 200 kHz. Because systems to do this have not yet been built, and design work is just beginning, it represents the greatest cost uncertainty. The system is clearly buildable. Engineering is necessary to pin down the costs and produce optimized designs. A maximum upper bound on this system might add another 40 M$ to the overall ignitor cost of 870 M$.

3.6.3. Beam Lifetime Considerations

The storage times in the accumulator rings vary from 6 ms for the first one filled, to only a few hundred microseconds for the last one. Assuming that hydrogen is the principal background gas, then a vacuum of 10^{-10} torr will result in a lifetime for stripping, i.e. $U^{2+} \rightarrow U^{3+}$, etc. on the order of 400 ms. Therefore, on average, less than 1% of the beam will be lost on this account. Nevertheless, this represents about 1 MW of average beam energy lost in the accumulator. Special precautions will have to be taken to collect these particles on specially designed aperture stops. There are two reasons for this. One is that careless handling of these lost particles could cause physical damage to the vacuum chambers and/or deterioration of the vacuum. The other consideration is that of activation of the machine components. It is desirable to keep the machinery as free of residual radioactivity as possible. Fortunately, heavy nuclei with energies of 85 MeV/nucleon tend to stop before having a nuclear interaction. An appropriate choice of material can further minimize the amount of residual activity produced.

It is a design option, of course, to improve the vacuum to bring it into the 10^{-11} torr range. However, another effect becomes important. This is the collision of particles within the beam colliding among themselves.

There is no direct experimental measurement of the charge-changing cross sections, say $U^{2+} + U^{2+} \rightarrow U^{3+}$ + etc. Plausible estimates of the cross sections put the lifetime in these accumulators at about 1 second. This lifetime is design dependent. That is, a larger radius accumulator will have a longer lifetime. Lifewise, a lower betatron frequency would also increase the lifetime. Since the lifetime increases as $R^{5/2}$, the loss rate could be halved by increasing the radius from 100 meters to 132 meters.

3.7. Transport and Final Focus

Before the longitudinal bunching process has terminated in the accumulator rings, the beams are extracted and transported for a distance of about 1 kilometer to the boiler. During this time the beam continues to shorten, until at the end of the transport the instantaneous current in each bunch has risen to 2500 amperes, or a peak power of 25 TW/beam. The eight beams result in a total of 200 TW for about 50 ns, corresponding to the total input energy of 10 MJ. Pulse shaping can easily be done by shaping and timing of the eight separate bunches.

3.7.1. The Transport System

The beams are transported in tunnels containing 4 beams. The beams consist of series of quadrupoles whose strength increases somewhat as the beam gets closer to the boiler. The transport consists of 10 cm diameter iron quadrupoles, employing superconducting coils as an energy conserving measure. Since this is a "once-through" system, the vacuum over most of the 1 kilometer can be in the 10^{-8} torr range. In the last 100 meters, these will be a transition to a higher pressure, perhaps to 10^{-3} at the boiler directly. The principal thing which limits the peak current in these transport lines is the beams' own space charge, which tends to defocus the beam. The currents assumed in this study are able to be transported without difficulty. It is possible that methods currently under study to neutralize the space charge would allow one to obtain higher currents, or reduce the cost of the focusing system. Nevertheless, for purposes of this study, it was deemed appropriate to make the more conservative assumption and forego neutralization.

3.7.2. Final Focus and Boiler

For the last few meters of beam transport, the beam is within the boiler, and must be focused to a suitable spot. In this case, a spot size of about .5 cm is required. The magnetic quadrupoles would probably be 5-10 meters from the final focus. Because of the close integration of the final focusing elements with the boiler, no effort was made in this study to attempt a design or cost estimate for these lenses. These costs would be included as part of the boiler cost.

Each of the quadrupole focusing elements subtends a solid angle of about .015 steradians. All 8 of them then intercept about 1% of the pellet energy. One important consequence of this is that one can quite readily afford to take special precautions to protect the front surface of these lenses, which one might not wish to consider for the entire boiler. Therefore, the radius of the boiler and the focal length of the lenses do not have to coincide. It is quite plausible to have the lenses "protrude" into the chamber.

There have been suggestions made that an intense beam of heavy ions might propagate in a self-focusing mode through the chamber if the pressure was in the 1 torr range. If this turns out to be the case it would be a considerable simplification. Lacking experimental confirmation, it seems prudent to assure that one can obtain satisfactory performance in a vacuum. The "gas" expected in the boiler is in fact predominantly metal vapor. This would come from either the pellet, the walls or a liquid heat transfer medium in the boiler. These metal vapors are easily condensed out by a spray of colder material. This same spray of cold liquid metal will also extinguish any plasma in the chamber, which may have been residue from the previous shot. The metal spray is the logical equivalent of the exhaust stroke of an internal combustion engine.

NOTICE

This report was prepared as an account of work sponsored by the United States Government. Neither the U.S. nor the U.S. Department of Energy, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness or usefulness of any information, apparatus, product or process disclosed, or represents that its use would not infringe privately owned rights.

PRINTED IN THE UNITED STATES OF AMERICA

Available from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, VA 22161

Price: Domestic $ 4.50

Microfiche $ 3.00

Printed: April 1978

740 Copies