Thoughts on Fusion Priorities

David E. Baldwin, General Atomics
Presented at the 1999 Fusion Summer Study
Snowmass, CO, July 12-24, 1999

When this meeting was initially conceived, we all assumed this would be the first of a few such meetings addressing fusion opportunities and priorities for the next decade. Events have passed us by, however, and we now have three review/advisory groups looking to the output of this first meeting to provide input for their recommendations. Given this turn of events, we should make every effort to grasp this as an opening to give these groups a message we really want them to hear. Otherwise, we will miss a opportunity not to be repeated.

If we expect to do this in the period of two weeks, we need to focus our attention on a limited subset of the wider range of issues we might otherwise discuss; we cannot afford the luxury of reviewing and discussing the entire program. In the few minutes I have here, let me outline what I see as major topics for our further discussion, focusing principally on MFE owing to the major changes which have occurred in that program over the last several years. The order of topics I have selected is not one of priority, but one which seeks to provide some sense of logical progression.

Consensus and Governance

It is widely recognized that we must develop a shared vision for the program and a consensus that has widespread community buy-in. Webster's defines consensus as "a judgment arrived at by most of those concerned," or a "group solidarity in sentiment and belief." It need not, and probably cannot, be unanimous; but it must be wide and deep. It is not enough simply to agree on fusion energy as an ultimate goal; our consensus must be at a level useful for defining strategy and priorities, and it must not represent just the lowest common denominator.

In a democratic society, consensus and governance are opposite sides of the same coin. Consensus cannot be implemented without a means for decision making and execution, and governance cannot succeed without the "consent of the governed." The program failures of our recent past, i.e., our failure in our last three major fusion initiatives, BPX, TPX and ITER, are ample evidence that something isn't working. There has been an ongoing disconnect between what we strive to do as a program and what we are willing to support as a community. Admittedly, we are not solely in control of our destiny, there are always external factors at work, and a stronger union of consensus and governance is a necessary-but-not-sufficient condition for success; but we could do much more to make our influence felt. In the past, it was often sufficient for a single lab to build an internal consensus behind its director; consent of the community was more a matter of securing the required program budget increase. In these times of "national" projects and constrained budgets, multiple institutions share a direct interest in any new initiative, and all share a budget impact. We need a better process of first evolving what we want to do as a community and then supporting the effort get it accomplished.

The issue of governance becomes most sharply drawn in the duality of science and energy. In a science program a strong, centralized agency control is neither appropriate nor the most effective. However, for the energy mission, we do need a management or governing mechanism for decisions which move us toward our goal, a mechanism which thinks and acts for the program as a whole. The OFES is engaging the FESAC in a stronger role to this end, but taken alone this doesn't provide an adequate solution. Also, we can't convene a Snowmass-like meeting like this every time we need to make a decision. We need to develop a process of more representative governance, combined with peer review, which we can respect and support even on the occasion when we don't agree with its priorities or decisions. This topic is not within the formal charge to this workshop; but it is so important that we should not, yet again, let pass the opportunity to begin a process of convergence.

Science *and* Energy

Although fusion research has now been designated a "science program," we must not lose sight of the fact that we are funded because of our energy goal. For now, we are told to focus on the "R" of R&D, i.e., to focus on developing the scientific building blocks for successful fusion energy. Our record of scientific progress over the last decade has been one in which we all can justly feel pride. In effect, we have been successful at just the activity on which we are now being told to focus. The job isn't finished, of course, and there remain important things to do; but we must also look forward to when and how we can again move on toward our energy goal.

The subject matter of our science in pursuit of the energy goal should be quite broad. Not only should it encompass the traditional topics of high-temperature and burning-plasma physics, but it should include the science of materials, energy conversion, magnets and heating systems for MFE, drivers and chambers for IFE, etc. To reach our challenging goal we will require improvements in such areas as these, accomplished by understanding and exploiting the underlying science.

Executing a science mission requires a variety of scales of facilities, a variety which is determined by the range and character of science to be explored. Where scientific questions can be addressed in small, specialized facilities, they certainly should be. However, exploring plasmas under fusion conditions which integrate many complex phenomena generally entails large facilities, sophisticated diagnostics, and substantial modeling and simulation, particularly as one begins to ask more and more detailed scientific questions. We must recognize the altered role of our large facilities. Whereas once they were associated with "parameter pushing" and were therefore deemed unscientific (which, in itself, was an inaccurate and unfair characterization), they now have become national scientific centers equipped to carry out in parallel multiple detailed explorations. In fact, it is in our ability to explore and understand the details of a significant level of integrated performance which has brought about many of our recent scientific successes. As we move forward, and particularly as we broaden the program to explore new concepts, we must also be careful to preserve this essential core capability.

Breadth *vs.* Depth

The recent restructuring of the fusion program is both valuable and timely in providing an avenue for continual injection of new ideas – any healthy R&D program must keep open such an avenue. Having come as far as we have in developing plasma science using the tokamak, there are important questions to answer for a power-plant application: "was one or another of the alternatives abandoned prematurely or for the wrong reason, i.e., for a reason which we now see a way to get around?"; or "could research into an alternative teach us how to further improve the tokamak either in a relatively pure form or in some hybrid configuration." In reexamining other fusion alternatives, however, we should be seeking prompt, definitive answers to these questions, not be embarking on a multitude of open ended and weakly related research programs. (The anticipation that most Concept Exploration experiments will not graduate to Proof of Principle is a recognition of this point.) There were valid reasons why the tokamak raced ahead of its competitors – recall that it proved to have stable operating regimes with good heat containment. What we need to determine, in the light of today's increased understanding, is whether these advantages were peculiar to the tokamak or whether they could be made to apply as well to another configuration having more desirable power-plant attributes.

The question of breadth versus depth is an important one for this meeting to discuss, for there is real risk of spreading our resources and efforts too thinly The fusion program risks adopting a strategy within which we divide our efforts to the point that we can't enter into anything very deeply. Our research has matured to the point that quality, frontline research requires both the heating systems to reach new operating regimes, like the recently discovered internal-barrier modes, and a substantial diagnostic complement. High-power NBI has proved a potent experimental tool, as have RF and the newly emerging ECH. Today's theory and modeling require in-depth levels of information for testing and validation of our understanding. To meet these needs, for example, the DIII-D facility is currently equipped with 3 multi-MW heating systems and has over 60 separate instruments providing space and time resolved measurements of most quantities needed – and the request list keeps growing! Investigation into any other magnetic fusion concept which intended to provide power and to delve scientifically to the same depth would require a similar installation. Recall, I made the point earlier that frontline fusion science generally requires frontline facilities.

There is a second aspect of depth, that of moving on to the next step which in itself incorporates a sense of progress to a greater depth. For any program, having an identified next step provides a focus and standard by which priorities within the prevailing program can be judged. For MFE, of course, the next-step means exploring self-heated plasmas. I'll return to more on that subject later on.

Physics Program Structure

Another topic deals with how we think about and how we represent our research program to the outside world. Today, we are all too prone to describe the program in terms of differences, rather than similarities. We talk about "tokamaks" *vs.* "non-tokamaks" or about six or seven "alternates," as if each was entirely unique. As a result, to outsiders we have come to appear to be subject to the adage "if you don't know where you are going, one path is as good as another." I submit that we do have a pretty good idea of the most probable features of a reactor and that our best strategy is to pursue that vision with vigor, but always to keep an open mind for valuable new ideas, insights or possibilities. To do this, our magnetic fusion program structure can much better be described as seeking to answer two fundamental questions:

What is the optimal configuration having a strong toroidal field, $B_T > B_P$?
The tokamak, ST and stellarator all rely on strong B_T and on a very similar body of underlying physics. So far, high-B_T configurations have provided our most successful experimental approaches for achieving reactor-like conditions. Projected to a reactor level, strong B_T appears to provide the required stability and confinement, but it raises issues associated with either superconducting magnets or high recirculating power fraction. The operative question is how best to capitalize on the virtues of strong B_T by choice of aspect ratio, toroidal symmetry, means for steady-state operation, etc., to optimize the advantages strong B_T introduces. In the end it will be a question whether the benefits thus optimized outweigh the costs.

What is the optimum configuration having weak toroidal field, $B_T < B_P$? Are the beta limit and confinement high enough? Can B_T be created by plasma currents alone? Can B_T nearly vanish?
The RFP, spheromak, FRC and dipole are all addressing answers to these questions. If the required stability and confinement can be achieved with low B_T – the question of confinement dependence on size will remain a crucial one – the costs and complexities associated with high B_T in a reactor could be reduced. Many other features would remain little affected, however, e.g., the constraints introduced by interaction with material walls. The idea of a relaxed Taylor state for the plasma is appealing, but the cost in transport and/or resistive losses near a cold surface is high – "self-organization" sounds great until you ask what it costs to keep the current going! Nonetheless, the payoff from positive answers to these questions could be great, so they are worth exploring; but the risk of negative answers must be recognized as being high.

If we were to structure our program around these two questions, it would be seen outside as having much greater coherence. The strong-B_T grouping, forming the mainline of the program, aims at reactor having many common features and similar cost-of-electricity (COE); the issue is the best way to do it. Research into the possibility of $B_T < B_P$ is pursued for its potential high payoff.

Ultimately, of course, the two questions are addressing two aspects of the broader question of the optimum toroidal confinement configuration. Our much increased understanding of the underlying physics and our expanding computer simulation and modeling capabilities provide the means for linking them. There is some advantage, however, in first breaking the broader question into two component parts and addressing them more-or-less separately. The situation is somewhat analogous to the direct- and indirect-drive approaches to IFE.

If we accept the validity of structuring the MFE program around these two fundamental strategic themes, then we have a structure for many of the discussions to go on here at Snowmass and for planning and discussing our research investigations. The questions then become "what are the essential issues to be addressed to answer each question?", "which can be addressed in existing facilities?", "what kind of new facilities are required?", "what kinds of increased computational and modeling capabilities are needed?", etc. These questions are very similar to those set out in the charge to this meeting, but by focussing them within this binary structure we can give more coherence to, and achieve a broader consensus for, our final conclusions.

Power Plant Optimization

In describing fusion as a science program having an energy mission, it is common to describe this duality as seeking to exploit the underlying science to "optimize" our power-plant concept. However, for this to be meaningful, we must agree in some sense upon what we are aiming for, or what constitutes a direction of improvement. Unless we can agree upon that, we run the serious risk of talking past one another.

My expectation that the first-generation fusion power plant will be a superconducting tokamak having some form of "advance tokamak" operating mode. The ARIES and ITER studies have provided a range of projections which vary considerably, depending on a number of critical plasma performance assumptions. To many in this audience, the resulting pictures are "too complex," or "too costly," or in some other way not what they had hoped for fusion. We all entered fusion research to pursue a dream, but we also should not be misled by unworkable fantasies. The reality for fusion is that we will have to deal with all the implications of the D-T fuel cycle, and it will come in ~1000 MWe unit sizes – implying some 7 or so meters for the chamber wall determined by allowable wall-loads of several MW/m^2. This reality is borne out by the ±15% consistency of projected COE for all MFE reactor designs, with the balance-of-plant accounting for ~50%. (IFE designs <u>target</u> somewhat lower costs, but this goal then rolls back <u>very</u> serious cost and other requirements on all system components.)

The consistency of the conclusions of the MFE reactor studies carries with it another important message and warning. We should not be too cavalier in criticizing the tokamak as a reactor, because in all likelihood its characteristics are close to those of any MFE reactor. A fusion reactor design is a multi-variant optimization within physics, technology and engineering constraints – confinement, beta, stability, heat and neutron wall loads, steady-state, simplicity, cost of sub-elements, etc. One misses the interaction of these multiple constraints in asserting that a given alternate is superior to the tokamak in some particular way and, therefore, automatically should make a better reactor. My two earlier strategic questions address this question of optimization in a way which keeps the greater picture in mind.

Burning Plasmas and Program Balance

Burning plasma (BP) research has in many quarters been called the next frontier of fusion research for both the science and energy components of the mission. The BP concept has strong emotive appeal to a community which has seen ignition as its Holy Grail for many years. The debate has centered around questions like "where do BP issues fit in the spectrum of other program priorities?", "what are the science benefits, for both the tokamak and alternatives, to expected from a BP experiment?", "how could such a project be funded?", and "is international collaboration essential?". I have felt all along that this Snowmass meeting should be used as a opportunity to discuss these and related questions and, hopefully, thereby to arrive at a sufficient consensus that we as a community could speak with something approaching a single voice.

Let me say at the outset that I believe that any new initiative on the scale of a BP experiment will, to a great extent, be required to bring in its own funding. The current base program (~$200^{+}M/yr. for MFE) is barely adequate to support the broader scientific program necessary either to support a BP initiative or to otherwise move forward toward fusion energy. It was this reality which forced us to think in terms of an international collaboration when we were targeting a facility the size of ITER, or even of ITER-RC. This is still the reality for a lesser facility. This nation could certainly afford to mount a $1B-class facility, should it choose to do so. However, whether it has will to do so is more problematic, for as we all know, energy is not on the national agenda. At best, any decision on such an initiative will have await the outcome of an ITER construction decision by Japan and Europe, so a ~5-year planning horizon would not be not unreasonable.

Since last year's Madison meeting there has been within a segment of the U. S. fusion community a strong sentiment for a so-called "multiple-machine strategy." Within this view, it would be necessary to split the BP and steady-state advanced-tokamak components of the ITER technical objectives into two different facilities, even though splitting them would necessarily mean being unable to investigate overlap issues – leaving that to a future step.

The current NSO studies, which have produced the pre-conceptual FIRE design, suggest that it might be possible to address burning-plasma and some steady-state issues (for ~20 seconds) in a single Cu-coil facility. The true feasibility of this dual goal is still open to question, although it offers a most interesting prospect. The FIRE design relies on AT operating modes at high field, as have been developed in DIII-D and C-mod, in more aggressive way than would have been prudent at the time of the BPX design. It will be tough challenge for the NSO study to identify a facility having such capabilities for the ~$1B cost target. However, if FIRE's costs and AT capabilities stand up to further scrutiny, the community would have something on the table for consideration as a next step, possibly as a domestic project but more likely to have some amount of international involvement. The programmatic benefit would be to provide the community a vehicle for moving on to this new regime using a means which also tests the best shot of the tokamak as a power-plant candidate. Certainly, in addition, the excitement engendered by a BP initiative would prove a powerful attraction for new, younger workers in the field.

The BP-physics questions which should be discussed at this meeting are numerous: the dimensionless parameters to be accessible and their ranges, the range of BP issues which can be addressed, the degree of simultaneous BP and AT operation, the transferability of results from tokamak to other alternates, etc. (This transferability has been very successful in other physics areas, such as transport, stability, divertor behavior, etc., and so should it be for BP issues.). We now have studies ranging from the full ITER down to Ignitor, with ITER-RC, PCAST, BPX and FIRE in between. Depending on what is demanded of a BP facility, I would expect there could be a maximum in the benefit/cost ratio at a point other than an endpoint.

Having expressed support for the kind of study the NSO is currently undertaking, we also should pay close attention to the emerging ST operating experience base, both by giving facilities like NSTX high priority and by folding their results into our planning. The point is that, were the ST to live up the projections of its advocates – and if we must wait ~5 years there is time to test this point – a next step based on the ST could provide the program with an exciting option. At a projected cost comparable to facility like FIRE, an ST conceivably could offer burning-plasma physics, steady-state operation and an eventual role as a facility for blanket development. Because the ST is so closely related to the advanced tokamak, we already have been able to skip the CE stage, and there are good reasons to be optimistic that its POP performance projections will be met. However, even then, the step from NSTX to a burning-plasma machine would be large in the performance needed. Some kind of phased operation would be called for, with successive stages contingent on continuing good performance.

Materials Testing

Finally, let me turn to the question of materials development. The issue of low activation materials is widely accepted as a critical element in the ultimate attractiveness of fusion. It is also the view that, before deployment in a test-reactor situation, specialized facilities will be required to test fusion materials and components made from them. Commonly, these consideration lead to advocacy of a beam-driven point neutron source (PNS) for small-sample irradiation and testing and a fusion-driven volume neutron source (VNS) for component testing. However, the PNS suffers from small irradiation volume and only approximately correct neutron spectrum, and the VNS suffers from the logical weakness of building a facility to test the very components needed to build it in the first place. Much has been said on this subject because of its importance to the attractiveness of fusion, and I have only two points to add.

First, testing is necessary, but it is not enough, however. At an time when elements of the materials and computational communities are talking about "designer materials", we should be also be thinking about how to approach the fusion materials question from a much more fundamental level where large-scale modeling would play an important role.

Second, the U.S. also should pay closer attention to the work going on in Novosibirsk on the Gas Dynamic Trap (GDP). A neutron source based on the principles of the GDT could produce ~1 MW of D-T neutrons at ~1 MW/m^2 for a cost considerably less than a tokamak-based VNS and requiring advanced components only over a limited region. It would have more test area than a PNS and could meet many of the objectives of the VNS. The Russians are carrying out experiments on the concept now, with very encouraging initial results, and are proposing a full hydrogen prototype.

An objection raised to the GDT is that it is not on a path to a power plant. I believe this to be a specious argument – after all, would a PNS lie on such a path? The right question is "what technique for generating the required neutron spectrum would provide the needed information at the lowest cost?"

Summary and Closing Remarks

To recapitulate what I have said so far, we need to sharpen the strategic focus within which we cast our science program. The strong-B_T path, and most likely the superconducting advanced tokamak, represents the most probable path to an MFE reactor. We should explore improvements both in that basic approach and more broadly, although we would be remiss not give priority to the advanced version of the tokamak approach which has proved so successful over the years. The MFE program is scientifically and technically ready to take the BP step, but we must find a way that allows that step to go forward without sacrificing other activities important for fusion. Finally, we should keep both our options and our minds open for technical approaches which might offer greater benefit and/or lower cost. An ST burning plasma facility or a GDT neutron source provide excellent examples.

The Community Plan for fusion at the recommended level of $300M/yr. contains four essential elements -- it would make good scientific use of our existing facilities, it would broaden and fill gaps in our MFE portfolio, it would embrace IFE as an important fusion energy option, and would it target 2004 to assess our readiness for one or more major initiatives in fusion research. The priorities I have outlined above are consistent with that Plan, both in pursuit of our ongoing fusion science efforts and in preparing us for the important next step.

Starting tomorrow, we will split into Working Groups. As an organizing deliverable for each Group, I suggest the question "What do you want the summary report, which will be carefully read be the external reviews and more widely, to say about your piece of the fusion subject?" In your discussions, I urge you to keep your eye on the big picture, to be mindful of the "fishbowl" in which we are operating, and to remember that consensus is essential if we wish our voice to be heard constructively outside our community.