Technical Issues in Keeping the Nuclear Stockpile Safe, Secure, and Reliable

Marvin L. Adams
Texas A&M University

Sidney D. Drell
Stanford University

ABSTRACT

The United States has maintained a safe, secure, and reliable nuclear stockpile for 16 years without relying on underground nuclear explosive tests. We argue that a key ingredient of this success so far has been the expertise of the scientists and engineers at the national labs with the responsibility for the nuclear arsenal and infrastructure. Furthermore for the foreseeable future this cadre will remain the critical asset in sustaining the nuclear enterprise on which the nation will rely, independent of the size of the stockpile or the details of its purpose, and to adapt to new requirements generated by new findings and changes in policies. Expert personnel are the foundation of a lasting and responsive deterrent.

Thus far, the United States has maintained a safe, secure, and reliable nuclear stockpile by adhering closely to the designs of existing "legacy" weapons. It remains to be determined whether this is the best course to continue, as opposed to introducing new designs (as envisaged in the original RRW program), re-using previously designed and tested components, or maintaining an evolving combination of new, re-use, and legacy designs. We argue the need for a stewardship program that explores a broad spectrum of options, includes strong peer review in its assessment of options, and provides the necessary flexibility to adjust to different force postures depending on evolving strategic developments.

U.S. decisions and actions about nuclear weapons can be expected to affect the nuclear policy choices of other nations – non-nuclear as well as nuclear – on whose cooperation we rely in efforts to reduce the global nuclear danger. This emphasizes the importance of the United States structuring its nuclear programs with the maximum transparency, making clear their unambiguous purpose and our intention to reduce reliance on nuclear weapons.

I. INTRODUCTION

I.A. Motivation

This paper was stimulated by a workshop[1] that was organized to identify and discuss technical issues involved with keeping the U.S. nuclear stockpile safe, secure, and reliable in the coming decades. We were specifically asked to address the following questions.

---

[1] The Role Of Nuclear Weapons In National Security Policy: Technical Issues. Workshop sponsored by the American Physical Society, the American Association for the Advancement of Science (AAAS), and the Center for Strategic and International Studies. Held April 24, 2008, at AAAS Headquarters, Washington DC.

1. How can the nuclear stockpile be kept safe, secure, and reliable in the face of aging, the end of nuclear testing, and calls in some quarters for drastic downsizing of the nuclear enterprise? What are the critical elements of the program that must be supported?

2. Will the legacy science-based stockpile stewardship program, Life Extension Programs (LEPs), and enhanced surveillance of the existing stockpile be able to do the job, or will substantial modernization be necessary?

3. If required, is modernization best done through concepts such as the Reliable Replacement Warhead (RRW), through a new but undefined concept that Congress has called Advanced Certification, or through an evolving combination of RRW and legacy stewardship?

I.B. Stockpile Stewardship to Date

Stockpile Stewardship includes increased surveillance of stockpile weapons along with programs designed to improve:

- our understanding of weapons-relevant physical phenomena;
- our ability to perform experiments, without nuclear yield, that improve our understanding of such phenomena;
- our ability to computationally simulate weapon explosions and provide realistic assessments of uncertainties in simulated quantities;
- most importantly, our ability to integrate the results from research, surveillance and re-manufacturing programs into a capability to assess stockpile issues, develop corrective actions as needed, and assess the impacts of those actions.

To date, stockpile stewardship has been successful by at least two fundamental measures. First, it has discovered causes for serious concerns in the stockpile, which is what it should do if such causes exist. Second, it has successfully addressed these concerns. To address some concerns, some components of some weapons have been replaced, sometimes with components that are not identical to the original designs. Such component replacement is one method by which "Life Extension Programs" (LEPs), which aim to extend the life of weapons in the current stockpile, can achieve their goals.

Although stockpile stewardship has been successful to date, continued success is not assured. A key goal of this paper is to describe elements of an ongoing stewardship program that we believe are necessary for continued success.

As time marches on we must be prepared to discover further issues that could lead to serious stockpile concerns that will have to be successfully addressed. These may arise from birth defects, design errors, or aging effects such as corrosion. The current approach of the LEPs, which have replaced individual weapon components to extend the life of weapons in the legacy stockpile, may not be the best, or even an adequate, procedure for resolving all future concerns found by future surveillance and analysis of the stockpile. One concern with the current LEP approach is that multiple small changes to the original design of a given weapon could conceivably interact to produce an unforeseen change in performance. This concern is the context for the technical questions that this paper is asked to address. For example, questions 2 and 3 concern the LEP and RRW programs. The latter program (no longer active) was devised to explore whether a new design – aiming for high confidence and long shelf life, perhaps at the expense of yield/mass ratio – could have advantages over existing tested, but aging and altered, designs.

Some preliminary remarks are in order. A “new” design would not be subjected to a nuclear-yield test and thus would not benefit from such a straightforward confidence-builder. However, any “new” design proposed in the foreseeable future would be based upon previously tested designs, and the alterations from tested designs would be intended to increase confidence in performance. Traditional LEPs also incorporate components with changes in their manufacturing process and design. Thus it is possible to weigh confidence in a LEP verses RRW design only in the context of specific individual designs, and reasonable experts may differ in their conclusions.

I.C. Context and Background

The questions above cannot be answered without also asking “What are these weapons for? What nuclear arsenal and infrastructure does the United States need to meet our international security goals in today’s world and for the foreseeable future?”[2]

Following a summit meeting in Moscow, November 13, 2001, Presidents Bush and Putin bluntly stated: “The United States and Russia have overcome the legacy of the Cold War. Neither country regards the other as an enemy or threat.” They called for “the creation of a new strategic framework to ensure the mutual security of the United States and Russia, and the world community”, and asserted “that members of NATO and Russia are increasingly allied against terrorism, regional instability and other contemporary threats.” These statements seem to mark the formal end of the era of nuclear deterrence between Russia and the United States. However, the nuclear policies and postures of the two nations have not reflected such a major change. The Nuclear Posture Review released selectively by the Pentagon in late 2001 recognizes that the U.S. must be prepared to respond to changes in our relations with Russia and other potential threats but that these may be done by drawing on what it called a responsive force, essentially a reserve force in addition to operationally deployed forces which could be available “in weeks, months, or even years,” as required to meet a developing crisis.

The clarity of the bipolar U.S.-Soviet world has now given way to the ambiguities and uncertainties of a world in which international security is threatened by transnational terrorists, unstable and failed states, and regimes that scorn a world order based on broadly accepted principles. The dangers inherent in such a stew are magnified by easier access to nuclear technology, inadequately protected stockpiles of plutonium and highly enriched uranium, the growing availability of missiles worldwide, black-market nuclear supply networks, and a trend toward acquisition of “latent” nuclear weapons capabilities through the possession of the entire nuclear fuel cycle.

These developments have led to increasing concerns that nuclear deterrence as implemented during the Cold War is becoming decreasingly effective and increasingly hazardous, if not obsolete. Various initiatives have been proposed to adapt today’s nuclear postures to address these concerns and to reduce nuclear dangers that we face in today’s changing world. Such actions would necessarily begin with the United States and Russia, who collectively possess about 90% of the nuclear warheads in the world, continuing to reduce their nuclear arsenals. Additional steps could be to cease the production of fissile materials and strengthen the NonProliferation Treaty (NPT) with measures to ensure effective monitoring of compliance. Such actions require close cooperation with many nuclear and non-nuclear nations. We can expect that the nuclear policy choices of non-nuclear states – for example, whether they cooperate and whether it is in their security interests to join the nuclear club – will be affected by

---

[2] An elaboration of this discussion can be found in a Report published by the Arms Control Association in October 2007: “What Are Nuclear Weapons For? Recommendations for Restructuring U.S. Strategic Nuclear Forces,” by Sidney D. Drell and James E. Goodby (http://www.armscontrol.org/pdf/20071104_drell_goodby_07_new.pdf)

our decisions about our nuclear forces. This is evident from their frequently expressed concerns about our commitment to Article VI of the NPT, in which the nuclear states agreed to negotiate toward cessation of the nuclear arms race and toward nuclear disarmament.

The effect of our choices on other nations must be weighed, along with other costs and benefits for U.S. national security, when we set requirements and initiate new programs for our nuclear arsenal and infrastructure. In weighing all factors, it is important to distinguish actions and changes that are *required* for the U.S. to maintain a safe and reliable nuclear deterrent from those that merely reduce the cost or difficulty of achieving this. Neither difficulty nor high cost should remove an option from consideration if it carries important benefits. We have already recognized such a distinction in honoring a moratorium on nuclear-yield underground tests (UGTs) for the past 16 years together with pursuing an ambitious, multifaceted science-based stockpile stewardship program.

I.D. Remainder of Paper

In the following section we address the first question listed above, which concerns the critical elements of a program to keep the stockpile safe, secure, and reliable. The subsequent section addresses the second and third questions, which concern the spectrum of choices for maintaining the stockpile. At one end of the spectrum is strict-as-possible adherence to existing stockpile designs; at the other is replacement using new designs that are based upon, but include changes from, tested designs and are not subjected to a confidence-building nuclear-yield test. We follow these technical sections with a brief discussion of policy considerations. We offer concluding remarks in the final section.

II. FIRST QUESTION

The first question has two parts: How can the nuclear stockpile be kept safe, secure and reliable in the face of aging, the end of nuclear testing, and calls in some quarters for drastic downsizing of the nuclear enterprise? What are the critical elements of the program that must be supported? Our answer also has two parts. The first describes the expertise that will continue to be essential for maintaining a safe, secure and reliable stockpile, independent of changes in the strategic environment, policy, or specific program elements. In accord with testimony and public statements by national laboratory directors and others responsible for stewardship of the nuclear stockpile, we submit that the foundation of enduring, responsive deterrence is not weapons or facilities, but rather the expert personnel in the complex. The second part of our answer describes the elements that we believe must be part of any program that will successfully maintain a safe, secure and reliable stockpile, including tools and facilities that the complex's expert personnel must have in order to achieve success. Again, we argue that these elements are independent of external environment or specifics of policy, including the size of the arsenal, within the realm of foreseeable scenarios.

II.A. Expertise

The most important component of a responsive infrastructure that can maintain a safe, secure and reliable nuclear weapons enterprise is personnel with the requisite expertise. This includes experts in surveillance, dismantlement, manufacturing, design, assessment, basic science, experimentation, computational simulation, etc. Success of the program, whose elements we discuss below, depends entirely on the expertise of the people who execute it.

The assessment of nuclear weapons is not a simple matter and will not in the foreseeable future lend itself to simple formulaic approaches. This is especially true given untested conditions or perturbations to tested designs, both of which are inevitable. Because of the extreme conditions generated during a nuclear explosion and the myriad phenomena that simultaneously occur, the expertise required to answer stockpile questions is broad, deep, and often unique to nuclear weapons. This expertise is best obtained by dedicated personnel who devote the bulk of their careers to nuclear weapons.

Expert personnel constitute more of a deterrent to evolving threats than do facilities or even existing weapons. Given sufficient resources, people with the appropriate expertise can respond quickly to unanticipated problems or changes in requirements and can provide confidence in the solutions they produce. Without such people, no amount of resources will yield timely solutions, in which confidence is justified, to new problems.

Expert personnel have always been important to the complex. This was obvious in the earlier years when new and improved weapons were being designed and introduced into the stockpile to meet new military requirements. Although we no longer call for new designs for new military missions, we still rely on the expertise of designers to assess and solve potential problems as they are identified over time. They must devise and assess possible solutions that can be developed and employed with confidence without relying on nuclear explosive testing. This challenge calls for both innovation and adherence to change discipline so as not to introduce more unknowns that could result in lower confidence in the redesigned weapon. The record of success to date testifies to the combined triumph of the stockpile stewardship program and the outstanding people who execute it.

Our comments about expertise apply beyond the designers of nuclear explosive packages. They are equally valid for personnel responsible for the non-nuclear components of the warhead. In addition, they apply to personnel who advance weapons-relevant science, enable weapons-relevant experiments, create and evaluate simulation codes, design and execute surveillance of stockpile weapons, and devise and execute the manufacturing and assembly of components.

If the number of nuclear weapons in the stockpile shrinks dramatically, one might be tempted to believe that the number of expert personnel in the complex should shrink as well. However, a smaller stockpile does not imply simpler or fewer problems that must be solved by expert personnel. It is possible that a smaller stockpile could require more effort to maintain: with fewer weapons there is less room for error in assessment of performance, and there is potentially an expanded set of environmental variables to consider for a given weapon design.

As we mentioned in the introduction, maintenance of existing designs in the stockpile has already led to the replacement of original parts with parts that were not identical to the original designs. We are referring not to limited-life components that were designed to be replaced, but to parts that were not designed for replacement. Assessing the impact of such a change requires a great deal of nuclear-weapons expertise, especially when the altered design cannot be tested. Changes that appear to be small may not have small impact on performance. We do *not* foresee a future in which simulation codes are so accurate and fool-proof that personnel with reduced expertise can make assessments in which confidence is justified. This would be true even under a policy in which the complex is directed to adhere as closely as possible to existing, tested designs. Some changes to the designs of some components will be necessary, and the assessment of any change should be performed (and reviewed) by personnel with broad and deep weapons expertise.

An increasingly broad spectrum of national-security problems requires nuclear-weapons expertise[3]. An example is the increasing need to assess proliferation risks under a variety of

---

[3] This has been emphasized by NNSA Administrator Thomas D'Agostino, in testimony to the Strategic Forces Subcommittee of the House Armed Services Committee on February 27, 2008, as follows: "In addition, our 21st

scenarios. Another is the desire for "nuclear forensics" to help identify the origins of interdicted nuclear material, radiological dispersion devices (exploded or interdicted), and nuclear explosive devices (exploded or interdicted). Yet another is the need for a capability to disarm and disable interdicted devices. There is also is the ongoing – perhaps growing – need to verify treaties and monitor nuclear weapons-related technologies: expertise will be needed to train the relevant inspectors. These additional requirements underscore the need for expert personnel as the foundation of a responsive nuclear-weapons complex.

We have argued that important and challenging U.S. national-security problems will require substantial nuclear-weapons expertise in the foreseeable future. However, importance and challenge alone may not be sufficient to attract the outstanding people that are needed. One of us (MLA) has regular contact with many graduate students and has advised ten Ph.D. graduates who have worked at nuclear weapons laboratories. Outstanding graduates have many opportunities, both in the private sector and in laboratories outside the weapons program, and they seek positions that are not only important and challenging but also interesting and perceived to have long-term stable futures. If there is a perception that the nuclear weapons program is declining, it may be difficult to attract the best and brightest. Because the problems are so difficult and so important, however, it is important to attract the best and brightest. Any strategy for maintaining the stockpile must address this issue. Maintaining strong scientific programs at the national weapons laboratories in scientific fields related to, but broader than, the disciplines directly applicable to understanding the warheads themselves can provide important inducements for attracting top-flight scientific and engineering personnel.

II.B. Other Elements

Here we describe elements that the nuclear weapons program must include in order to maintain a safe, secure, and reliable nuclear weapons stockpile. We submit that these elements are necessary independent of the size of the stockpile or the details of its purpose, even in the limit of the smallest size conceivable. As we have asserted above, a flexible and responsive deterrent rests on a foundation of expert personnel, but these personnel must be embedded in a program that allows them to retain and hone their expertise while providing the tools necessary for appropriate responses to surprises and policy changes.

II.B.1. Search and Discovery

One cornerstone of the nuclear weapons program must be vigilance in the search for and discovery of problems in the stockpile, arising for example from design errors, aging, or birth defects in the original manufacturing process. This must continue to include rigorous and aggressive surveillance and forensics designed to detect problems early, with high confidence that problems are not missed. The record on this to date appears to be very good, with the surveillance program having supported twelve years of successful annual assessments. In addition to surveillance, search and discovery must include modern assessment of legacy designs. This can be viewed as time-delayed peer review that brings to bear tools and knowledge developed since the designs were originally certified.

Once a problem is discovered, personnel in the complex must develop an understanding of it, devise solutions for it, and assess the proposed solutions. Understanding and assessment rely on

---

century enterprise will continue to leverage the scientific underpinnings of the historic nuclear weapons mission to respond to a full range of national security challenges that we have, and beyond nuclear weapons sustainment but shift those more towards nuclear counterterrorism and nuclear nonproliferation activities. And as an example, we provide technical support to the Defense Department and the FBI and emergency render-safe and post-event nuclear technical forensics activities. And a lot more needs to be done in that area and we're going to be looking to shift more towards that area."

theory, experiment (including past nuclear tests), computation, and peer review, as described below. Solutions depend on the expertise and imagination of the responsible personnel.

II.B.2. Computational simulation

A second essential cornerstone must be a capability for high-fidelity computational simulations of phenomena relevant to nuclear weapons. This capability provides tools that are needed for discovering and understanding problems in the stockpile, devising solutions to problems, and assessing impacts of proposed solutions. Such a capability relies upon several key elements, including computational hardware, model development and validation, software implementation of models and verification thereof, quantification of uncertainties and assessment of predictive capability, and ability to adapt simulation capability in response to new findings and/or changes in the spectrum of problems that need to be addressed.

High-fidelity simulations of some important phenomena – phenomena that in many cases dominate the uncertainty in our simulated results – remain well beyond the capability of today's computers. While the next generations of computers (Roadrunner and then Sequoia) will lessen this problem, they will not eliminate it. Pursuit of advanced computers, as in the Sequoia acquisition, should continue, but simultaneously there should be continued cost/benefit evaluation of how much reduction in uncertainty can be gained by further increases in computing power.

The development of scientific understanding of physical phenomena requires theoretical and experimental investigations. Experiments inform theorists and test their theories. Theories are incorporated into mathematical models that can be solved (approximately) with computers. One might imagine that all needed models are well in hand by now, but this is not the case, especially for "multi-scale" problems that are ubiquitous in the simulation of weapons phenomena. In these problems we may have accurate models for behavior at the atom-by-atom scale, for example, but not at the much larger scales on which we must perform simulations. Practical models are currently under development for phenomena that in many cases dominate our uncertainties, and this must continue in the near future. This is an important component of the "Predictive Capability Framework" that guides efforts in the NNSA Advanced Scientific Computing program. The experiments that inform and test theories relevant to weapons phenomena include past nuclear tests as well as ongoing tests with no nuclear yield. Simulation is also playing an increasing role in informing theories and in testing some models for consistency with more fundamental models.

Model "validation" requires implementation into software and testing against reality, where "reality" ultimately must come from experimental measurements. "Validation" is not a yes/no proposition, but rather should be a quantitative assessment of model error. It is not possible to quantify predictive capability without quantifying model error, and it is not possible to quantify model error without measured data. This is part of what drives the need for an experimental program as one of the cornerstones of the nuclear weapons program, as we shall describe below.

Model implementation into software is more challenging today than in past decades because of the complexity of the computing hardware on which the software must operate. It is difficult to make efficient simultaneous use of thousands of processors, without processors constantly waiting for results from their neighbors. This difficulty increases with processor count and with the heterogeneity of the computing platform. Roadrunner is highly heterogeneous, and one Sequoia platform is expected to have well over one million processing cores. For such platforms, successful implementations of mathematical models of physical phenomena require integrated efforts by teams with expertise in theory, discretization techniques, iterative methods, and computer science. A key part of implementing a model is "verification," in which the team attempts to ensure that the implementation is correct (bug-free) and that the errors made in discretization and iteration are quantifiable and controllable.

Given a set of models, the numerical methods used to approximate the models, and a software implementation, simulations can be performed. However, simulation results are not very useful unless their accuracy is known. Quantification of uncertainties in simulated quantities of interest is extremely challenging, especially when measured data are sparse and uncertain. This task requires rigorous analysis that includes comparison of simulation and measurement, as well as quantification of measurement errors, numerical errors, model sensitivity to uncertainties in input quantities, etc. The challenge is compounded if the problem in question is not "close" to events that have been experimentally measured, as can be the case in problems that arise under stockpile stewardship. The computational science community worldwide is attempting to address the general problem of quantitative assessment of predictive capability, but this is a relatively new exploration and methods for this assessment remain under active development.

Our final remark about the simulation cornerstone is that the required simulation capability is not static. We cannot foresee every problem that will need simulation; thus, the complex must retain the ability to adapt its simulation capability to new requirements generated by new findings, changes in policy, and additions to the slate of national-security problems that require weapons expertise. The ability to adapt rests upon people who are actively engaged in the elements described above, including model development and validation, software implementation and verification, and assessment of predictive capability.

II.B.3. Robust Experimental Program

A third essential cornerstone must be a robust and advanced experimental capability. The integral role of experimental science is evident in the preceding discussion concerning the cornerstone capability in high-fidelity simulation, in which experiments were mentioned many times as essential ingredients in the development of modeling and simulation capabilities. Below we describe some of the many roles played by an experimental program and stress the need to maintain an experimental capability.

Experiments are essential in discovery science for phenomena that are not sufficiently understood, as in the ongoing efforts in the National Boost Initiative and in efforts to resolve "energy-balance" questions. In this role they inform model development and help to assess model validity. Data archived from past nuclear tests play an essential role, but they are not sufficient to answer all of the questions that arise in the course of stockpile stewardship. Experiments can often be devised to mock up problems uncovered through stockpile stewardship, thereby helping assess the effects of those problems and also the accuracy of simulations applied to them. Similar experiments can help with assessment of the impact of proposed solutions.

It is important to recognize that the ability to respond to problems requires an ongoing capability to design and field experiments. It is not possible to anticipate all future problems and perform the needed experiments in advance. .

Experimental measurements are an essential ingredient in the validation of mathematical models of physical phenomena, the quantification of uncertainties in simulations, and the assessment of predictive capability. This is detailed above in the discussion of simulation capability. We remark that ultimately our predictive uncertainty can be no smaller than the uncertainty in experimental measurements upon which our simulation capability is founded. Thus, in some cases there may be strong motivation to develop improved experimental diagnostics with smaller and better characterized errors.

An example of improved diagnostics with smaller and better-characterized errors is the Dual-Axis Radiographic Hydrodynamic Test (DARHT) facility, recently brought into operation at Los Alamos with multiple pulses and two axes. DARHT will provide significantly more precise measurements of experiments of a kind that have been done for many years. A different example is the National Ignition Facility (NIF), currently nearing completion at Lawrence Livermore, which will enable new kinds of experiments by creating physical conditions that were previously unattainable in the laboratory. If NIF performs as predicted it will enable new "discovery" and "model-development" experiments for phenomena that are not sufficiently understood today, and will enable validation experiments for models and codes under physical conditions that are significantly closer to those created by nuclear explosions. As yet another example, sub-critical experiments at the Nevada Test Site permit us, with advanced instrumentation, to explore properties of key weapons materials and allow us to test, for example, whether different manufacturing processes have a significant effect on the behavior of newly built components. A number of additional examples exist of important small- and large-scale experimental facilities, but hopefully these three illustrate the importance of such facilities.

II.B.4. Peer Review

In the absence of a confirmatory nuclear test, it is appropriate to take great pains to carefully assess any modification to or replacement of an existing tested design. It is equally important to subject the assessment to careful scrutiny to either discover any weaknesses in the assessment or to build high confidence in its validity and rigor. An essential part of such scrutiny is detailed independent peer review. This requires an independent set of weapons experts using an independent set of analysis tools (such as different simulation codes). We amplify on this in the section that addresses the second and third questions.

II.B.5. Research and Development

The program elements described above cannot succeed without research and development (R&D). At the risk of belaboring the obvious, we shall list a few examples (certainly not an exhaustive list) of how some of these elements depend on R&D.

The development of understanding of weapons-relevant physical phenomena and the incorporation into mathematical models are R&D efforts. In many cases the experiments that address questions related to the phenomena require R&D before they can produce the needed data with the needed precision. Statistically rigorous quantitative assessment of model error (how closely model output matches measured data) and predictive capability (how closely a model predicts what has not been or cannot be measured) remains an area of active R&D.

Discretization methods for mathematical models of physical phenomena are an active area of research. Accuracy often depends on tailoring the discretization method to the physics described by the model as opposed to blindly implementing an existing method. Further, existing discretization methods have known flaws that are continuously being mitigated and removed by ongoing R&D in applied mathematics. R&D is also needed to develop algorithms that can efficiently solve the discrete equations that model coupled physics on computers that contain 10^6 or more processing elements. It has been difficult to achieve efficiency on 10^4 processing elements; 10^6 is uncharted territory.

Quantification of uncertainty and assessment of predictive capability are tremendously important tasks in an environment without testing. Methodology for these difficult tasks is in a relatively early stage of development in the scientific computing community, and techniques for assessing the accuracy and reliability of this methodology are essentially nonexistent. R&D is underway in the community to improve this situation.

In addition to enabling the program elements listed above, R&D supports the exploration of options across a wide spectrum of choices for addressing stockpile problems that will surely arise as stewardship continues. We discuss this further in the section that addresses the second and third questions.

Finally, we emphasize that a vibrant R&D program will be very helpful, and probably essential, for attracting outstanding people into the complex. A nuclear weapons program without active R&D will likely be viewed as dead or dying by people who are making decisions about where they want to invest decades of professional effort.

II.C. Summary

Any program that is expected to maintain a safe, secure, and reliable nuclear weapons stockpile should be tailored to support a nuclear-weapons policy that must be formed by the next and subsequent administrations. While specific programs and emphases should be policy-dependent, the ingredients we have listed will be required under any foreseeable policy. These ingredients form the basis of a nuclear weapons program that can be flexible enough to support U.S. policy as it inevitably evolves in response to a changing world.

III. SECOND AND THIRD QUESTIONS

III.A. Introduction

The second question is whether "legacy" stockpile stewardship, LEPs, and enhanced surveillance of the existing stockpile can do the job (i.e., maintain a safe, secure, and reliable stockpile) or whether substantial modernization will be necessary. The third question is whether modernization, if necessary, is best done through concepts such as RRW, through Advanced Certification, or through a combination of RRW and "legacy" stewardship. The wording of these questions suggests either/or choices, whereas we see a rich continuum of options. We combine and re-state these questions as follows:

> Will the U.S. be able to maintain a safe, secure, and reliable stockpile by adhering closely to original designs of today's stockpile weapons, will we need to introduce new designs, or will we need an evolving combination of both? Where does Advanced Certification fit in?

Because Advanced Certification has not yet been defined, we cannot discuss it with confidence. Suppose we take it to be a procedure whereby any weapon design is subjected to rigorous assessment, with strong peer review of the assessment, before it can be certified. Then Advanced Certification should be applied to any altered weapon, regardless of where the alteration falls on the spectrum described above. We take this to be independent of the question of where the weapons program should fall on that spectrum as the future unfolds.

As mentioned in the introduction, there is a spectrum of modernization options ranging between two extremes. At one extreme we contemplate making heroic efforts to replicate weapons as they were introduced into the stockpile, using hazardous materials and outdated manufacturing processes, regardless of costs or modern assessments of the designs. At the other extreme we contemplate replacement of aging weapons using new designs that have not been subjected to nuclear tests, trusting our ability to predict their performance. To date the LEPs have operated toward the replication end of the spectrum (but see discussion in next paragraph), while the RRW program was intended to be a venture toward the other end. A "hybrid" strategy in between could involve repackaging "old" components, tested as part of old designs, into new designs.

LEPs have not stayed completely at the replication end of the spectrum. Outside the Nuclear Explosive Package (NEP) of a weapon there are components that can be tested to a significant extent. Some of these have been replaced by newly designed components, which is an example of the new-design end of the spectrum. Some components inside the NEPs have been replaced by components that differed from the originals to some extent in design and/or processing. For example, the replacement W88 pit employed new personnel using new manufacturing processes and different facilities. In certifying this pit we have declared success in venturing, to some degree, away from identical replicates, even inside the NEP. Our confidence in the modified warhead is based on careful experiments and analyses performed as part of the Stockpile Stewardship Program. It is reasonable to assume that in future LEPs we will continue to see replacement components that are not exact replicates of the originals.

III.B. Point on the Spectrum

We believe that the best option along the spectrum will depend on the issue being addressed and the requirements that are imposed. That is, our answer to the main question of this section is: "It depends on the problem." If we encounter a design in which there is high confidence, with a component that is not difficult to replicate almost identically, replication may be the optimal solution. But given a hypothetical design for which modern assessments produce low confidence and in which there is hazardous material and parts that are not easy to replicate, a new design with high margin and not-too-high uncertainty could be the optimal path. This new design could include all new components or, conceivably, some components from previous tested designs. If policy were to require that stockpile weapons contain certain advanced features, then new designs could be the most reasonable option.

Because we do not know what problems will arise in the future or what requirements will be imposed upon the stockpile, and because the best choice among the spectrum of options is problem- and requirements-dependent, a flexible response capability must include the ability to apply and assess a range of options. It would not be wise for the complex to focus only on replication or only on new designs, for example, for then it may not be prepared to apply the best solution to the next problem that arises.

III.B.1. How do we choose?

We should continue stockpile stewardship and continue to evaluate what it tells us about the merits of various options. As part of or in conjunction with this, we should prepare to respond to evolving requirements and potential surprises by thoroughly assessing multiple options (i.e., various points on the spectrum). Consider the "new-design" end of the spectrum, which the RRW program was slated to explore. As originally conceived, RRW-1 was intended as a replacement for Navy delivery platforms with design features that would enhance the safety of the weapon in abnormal environments such as fires and incorporate new features to preclude unauthorized use. In addition the warhead was required to meet narrowly specified military characteristics and to incorporate design features to ensure that it would age gracefully, without any required refurbishments for several decades and – crucially – without triggering a need to resume nuclear-explosion testing.

Before it was cancelled, RRW-1 had not yet reached a point where one could say with confidence whether it could achieve its goal of developing a warhead, with improved safety, security, and reliability, that could be deployed without underground tests and with higher confidence than the legacy warhead it was designed to replace. Moreover the program had not yet instituted the desired peer review process – either by LANL of the LLNL design that had been selected, or by LLNL within its own house – largely due to lack of time and funding. It was also unclear whether the RRW program would be financially sustainable, for detailed cost analyses had yet to be performed.

The RRW Program asked important questions that should be answered. The three weapons laboratories did excellent work, for which they should be highly commended, in attempting to answer whether the RRW-1 design could successfully achieve its ambitious goals within the restraints that were imposed: no new designs for new military missions and no underground explosive tests. These restraints were imposed on the basis of a judgment by the government that they are important in order not to harm prospects for achieving U.S. strategic/political goals of reducing the nuclear danger and of strengthening the nonproliferation regime, which rely on broad international cooperation.

A detailed technical review of the RRW program that had been mandated by Congress was completed by JASON in the summer of 2007. In its abstract, as presented to NNSA, the review summarizes its findings as follows:

NNSA tasked JASON to conduct a technical review of the Reliable Replacement Warhead (RRW), with a focus on the LLNL/Sandia design, now call WR1. This report summarizes our findings and recommendations. The design of a new warhead, without new nuclear explosive tests, relies on the scientific connections and traceability of that design to (1) the legacy nuclear explosive test data, (2) established physics, and (3) new and ongoing experiments. The WR1 design is pursued with these principles in mind, but certification is not yet assured. The certification plan presented needs further development. For example, additional experiments and analyses are needed that explore failure modes, and assess the impact on performance of new manufacturing processes. Substantial work remains on the physical understanding of the surety mechanisms that are of high priority to the RRW program. Establishing that the case for confidence in any RRW has been satisfactorily made will require a new peer review process. In addition to certification issues, it is too early to assess how the WR1 will impact the modernization and streamlining of NNSA's production complex.

Based on what we have learned so far from the LEP and RRW experiences, we suggest the following: Given the difficulty of the challenge to create a certifiable new design, a program for maintaining a safe, secure, and reliable stockpile should explore a wider segment of the spectrum of options. We still have far to go before answering whether new designs can be created that incorporate all the desired attributes, can be fielded without UGTs, and provide confidence as high as or higher than we have currently in the legacy weapons. Further, there is always the possibility that a new design could contain unrecognized flaws that could lead to unanticipated problems, perhaps generating more significant findings than today's designs. In making design changes it is important to keep this danger in mind and exercise careful change discipline. As we attempt to resolve these issues, perhaps it would be prudent to also address them for a point in the middle of the spectrum, such as new designs that repackage components from previously tested designs.

Exploration of multiple options along the spectrum would create greater flexibility to respond appropriately to changing requirements and to potential surprises (such as those that may arise from continued search and discovery under stockpile stewardship). Each point in the spectrum represents an approach that has strengths and weaknesses. Until an approach is explored in some detail we will not know these strengths or weaknesses in sufficient detail to support informed decisions about how best to address a given stockpile issue. Investigation of multiple options thus supports the flexibility and responsiveness that is desired in the complex. It also meshes well with the R&D program described in our answer to the first question, provides a means by which personnel (including designers) can hone their expertise, and could help attract outstanding minds into the weapons program.

Because of the success of stockpile stewardship to date, there appears to be time to carefully assess multiple options along the spectrum: While tomorrow could bring surprises, today we know of no burning issues in our stockpile that appear to be beyond the capability of our expert personnel to solve, given reasonable resources and implementation of program elements such as those we described in the previous section. Also, before committing to a point on the spectrum, we need guidance from the policy-makers as to the missions, composition and the size of our future arsenal.

It appears that NNSA will have to maintain an infrastructure capable of providing the necessary support, including life extension programs, for elements of the legacy stockpile for the time being. Our suggestion to explore multiple options could conceivably lengthen this time period to a marginal extent, compared to a scenario in which a single option is chosen immediately, but we view the exploration as a necessary pre-requisite to decisions that could eventually eliminate the need for the legacy infrastructure. The explorations should be more generic until there is higher confidence, supported by a strong consensus emerging from peer review, of what improvements can be accomplished and deployed without UGTs. This would enable the government to implement changes if it judges them to be merited on balance, based on their supporting our strategic/political goals of reducing nuclear danger and strengthening the NPT, and their contribution to improved efficiency and economy, in addition to their technical merit. Such actions may well be short of deploying a new RRW equipped with a full suite of desired features. On the other hand, if policy assigns heavy weight to the full suite of desired features, a new-design new-component option could be judged most meritorious on balance.

We also want to emphasize the importance of a strong and credible peer review process as an essential component of any certification program. To this end we quote the finding and recommendation in the JASON 2007 Review of the RRW Program, which pointed out that it "must play a larger role than provided for by current NNSA guidelines." It also recommended "that NNSA establish a RRW peer review mechanism with the following elements:

- the process must be visible, funded, and administered to assure the nation that all expertise available has been applied to a rigorous evaluation of the new design;
- it is imperative that is effectiveness be examined periodically by an independent organization;
- the peer review team should be broadly constituted and have authority to pose formal tests of a computational or experimental nature to the design team;
- issues identified through peer review must be documented, tracked and follow a formal process of closure with participation by the peer review team;
- responsibility for conducting peer review should be assigned to the weapons design laboratory not leading the design effort."

IV. POLICY CONSIDERATIONS

While our discussion of the three questions posed for this workshop has focused on the technical issues, it will be essential to keep in mind the strategic/political implications of actions initiated in maintaining and modernizing U.S. nuclear forces. Not only what we do but how we do it will be important. As we discussed in the Introduction, U.S. decisions and actions about nuclear weapons can be expected to affect the nuclear policy choices of other nations – non-nuclear as well as nuclear – on whose cooperation we must rely in efforts to reduce the global nuclear danger. Relevant factors include structuring new programs with the maximum transparency, and making clear their unambiguous purpose and our intention to reduce reliance on nuclear weapons. In addition as we continue to reduce the size of the arsenal it will need to retain a flexibility to adjust to different force postures depending upon evolving strategic developments. For example, we may want larger fractions of the force in a reserve posture that can be reconstituted in a timely manner depending upon potentially emerging threats. The need for flexibility will also impact how we organize R&D and Advanced Certification programs. This point was addressed in a broader context for the US defense establishment by Richard Wagner and Ted Gold in 1990 in their paper "Long Shadows and Virtual Swords" (published by the American Association for the Advancement of Science, in a volume entitled "Science and International Security" edited by Eric Arnett). They emphasized that, in circumstances such as we are facing today with an uncertain and changing strategic security landscape, maintaining a strong technical base remains of utmost importance. Rather than measuring the success of R&D primarily in terms of its leading to new deployments through a direct acquisition process, we should view its success in terms of preparing potential capabilities that will be available if and when needed at a later date. In this paper we have emphasized the importance of such a program, which we hope will emerge from the government's updated nuclear posture review.

There are a number of specific issues on which an updated nuclear policy review can have a major impact on how to optimize the balance among Advanced Certification, RRW, and maintaining legacy elements in our modern program. An obvious one is the overall size and production capacity of the modernized infrastructure. This will depend on the total number of warheads required in the arsenal – including both the actively deployed and the responsive, or reserve, forces, as well as the potential surge demands we need to be capable of meeting in the event of unexpected strategic developments or potential technical surprises.

An additional way in which an updated nuclear posture review would have large leverage on the U.S. weapons program is in the determination of the targets we intend to hold at risk. The hardness of targets of interest, together with recent and continuing improvements in accuracy of delivery systems, will determine the appropriate mix of weapons yields. The yield required to generate a given effect is sensitive to the accuracy of the delivery systems.. With the improved accuracy of modern weapons, it is conceivable that lower yields, possibly achievable via tested single-stage weapons in which we have high confidence, could meet broad targeting requirements.

A third factor is the nature of anticipated hazardous environments. With recent improvements in accuracy, not to mention reduced stockpile numbers, fratricide has receded as a concern. Similarly the world today is moving ahead with limited ABM systems, presumably cooperative between the U.S. and Russia, with hit-to-kill or proximity-fused conventional warheads. There is little apparent interest in nuclear-tipped ABM systems that raise concerns about neutron-induced pre-initiation. Meanwhile one must question whether the impact of such hazardous environments should continue to raise concerns about weapon reliability due to the reduced performance margins they cause for some of our weapons – concerns that impose burdens on our weapons program. The Cold War era stockpile-to-target sequence requirements need a formal high-level review in view of the considerable burden they put on the current SSP. Must they be retained or, based on an assessment of their costs vs. benefits, should these requirements be relaxed in the new post-Cold War strategic environment, especially in light of other technology advances?

V. SUMMARY

Stockpile stewardship is working. It has uncovered causes for serious concerns in the nuclear weapons stockpile and has successfully addressed these concerns. Addressing stockpile concerns has involved replacing components, sometimes with parts that are not identical to the original designs. To address a concern, the weapons program faces a spectrum of options ranging from replication of original parts at one extreme to fielding a new design at the other. An interesting option between these extremes is re-use/re-packaging of components from previously tested designs. We argue that the best option along this spectrum depends upon the issue being addressed and the requirements that are imposed. These cannot all be known in advance. We conclude that it is wise to develop a responsive capability by exploring several points along the spectrum through an active research and development program.

Continued success of stockpile stewardship is not a foregone conclusion. A responsive infrastructure that can continue to maintain a safe, secure and reliable nuclear weapons enterprise must include several key components. We argue that the most important component is top-notch expert personnel, without whom confidence in the U.S. nuclear deterrent will erode along with U.S. ability to respond to changing threats and other national-security requirements. These experts must be engaged in a stable program that includes several cornerstone elements: 1) Vigilance in the search for and discovery of problems in the stockpile; 2) High-fidelity computational simulations; 3) A robust experimental program; 4) Strong peer review; 5) An active research and development program exploring a range of stockpile options.

U.S. decisions and actions about nuclear weapons can be expected to affect the nuclear policy choices of other nations – non-nuclear as well as nuclear – on whose cooperation we rely in efforts to reduce the global nuclear danger. This emphasizes the importance of the United States structuring its nuclear programs with the maximum transparency and of making clear their unambiguous purpose and our intention to reduce reliance on nuclear weapons.

ACKNOWLEDGEMENTS

We thank the American Physical Society, the American Association for the Advancement of Science, and the Center for Strategic and International Studies for organizing the workshop that led to this paper. We thank Lars Bildsten and Raymond Jeanloz for helpful comments and suggestions.