SLAC–PUB–5875
August 1992
(A/E)

DESIGNING RELIABILITY INTO ACCELERATORS *

ANDREW HUTTON
Stanford Linear Accelerator Center, Stanford University, CA 94309, USA

1 Introduction

For the next generation of high performance, high average luminosity colliders, the "factories," reliability engineering must be introduced right at the inception of the project and maintained as a central theme throughout the project. There are several aspects which will be addressed separately:

Concept
Design
Motivation
Management Techniques
Fault Diagnosis

2 Concept

"Quality is not defined by zero defects. That's only a partial way of looking at quality, and it shows a misunderstanding of what the job is....The customer is interested in performance...whatever is part of performance or style, unless those elements improve, the company fails." (W. Edwards Deming).[1]

In building accelerators, we are always on the technology frontier and must take risks to provide the performance that is our sole justification for being here.

2.1 Concept of single risk

We are obliged by the performance requirements to extrapolate the state-of-the-art. The concept of single risk means that not more than one parameter in any system should be extrapolated beyond existing technology. Usually, technology or performance improvements will be required in many different areas of the machine design. The important point is that these improvements should be kept as separate as possible to avoid unexpected and unpredictable interactions, which almost always have a negative impact on the performance.

---

* Work supported by Department of Energy contract DE-AC03-76SF00515
Presented at the Advanced Photon Source Reliability Workshop, Argonne, IL, January 29–31, 1992.

This design concept is the single most important factor in reducing the risk in building a new machine. If there are several extrapolations of parameters, possibly interacting to enhance each other, the machine will be extremely difficult to commission, and operation will also be difficult to predict. It is much safer to estimate the effect of a single major step in an otherwise well-understood set of parameters.

3 Design

Having adopted a parameter set, based on the concept of "single risk," each subsystem should be designed to maximize reliability. In each subsystem, the single risk must be identified and theoretically evaluated to the very best of our ability, inviting experts from all over the world to help (it is no use being proud!). Prototypes of each single risk item should then be built and evaluated.

It is important to evaluate every aspect of the design theoretically, particularly the new risks. But it is even more important to build prototypes to check the theory. It is only with this two-pronged approach that it is possible to gain confidence in the design and estimate the safety margin.

It is important to examine every major component of the machine at this stage to evaluate how to:

Minimize the likelihood of failure
Diagnose a failure (control system)
Repair a failure (modular design)
Minimize consequences of failure (partial redundancy)
Minimize return to operation (common spares)

At the earliest stages of the design, the engineering choices should be made to maximize the reliability of the whole machine and to reduce the time needed for diagnosis and repair of faults.

It should be noted that the greatest return on investment occurs during the concept and engineering phases of the project. Attempting to "add" quality at a later stage is extremely inefficient (see Figure 1).[2]

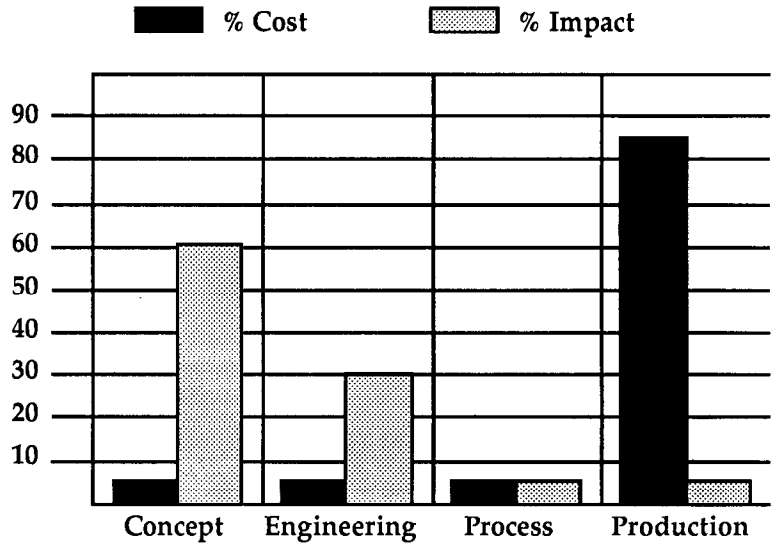

Figure 1. Comparison of Cost and Impact of Quality.

4 Motivation

The only way that reliability can be assured is by making every member of the project team feel personally responsible for the quality of his, or her, subsystem and a vital part of the success of the whole project. It is the role of Project Management to instill a common vision in each and every person working on the project and to provide the tools necessary for each engineer to evaluate his, or her, success in reaching the goals. Quality cannot be *imposed* from above nor ensured by stringent inspection. It takes a dedicated and motivated project team who understand that quality is of paramount importance. It is the function of the Project Leader to promote this common vision of the importance of quality and motivate the project team. The best definition of a leader that I have seen is due to Laotzu (600 BC)[3]

A leader is best
When people barely know that he exists,
Not so good when people obey and acclaim him,
Worst when they despise him.
"Fail to honor people, They fail to honor you;"
But of a good leader, who talks little,
When his work is done, his aim fulfilled,
They will all say, "We did this ourselves."

The emotional commitment, "buy-in," of the project team is vital to the success of the project. This should be accomplished by delegating to each physicist or engineer the responsibility of designing and building their system as they think best within the context of the global machine optimization and by giving them the authority over the resources necessary to carry out the task. This is not sufficient, however. They must all work together as part of a team for the benefit of the whole project.

> "A system consists of components. Any company, any industry, consists of components that are different activities. All the components of the system must contribute to the system, not exist for their individual gains.
>
> For example, the travel department in a company is not there to save money on travel, but to serve the whole company, so the job of the travel department should be to put a passenger down at his destination, physically fit for the job, even if the travel department has to pay a premium rate to put him there. Every component must serve the whole system." (W. Edwards Deming).[1]

The management principles of W. Edwards Deming[4] were designed to foster cooperation among the workers in an industrial context. However, the basic principle is extremely well-suited to an accelerator project (or indeed to any high-tech project). The intention is that everyone working on the project should be trying to build the best possible accelerator with the resources available, not trying to build the best subsystem to the possible detriment of the whole project.

Deming uses the word sub-optimize to describe the process where individual groups optimize their subsystem rather than the whole project. It is easy to show that this always leads to an inferior overall result, but not so easy to see how to avoid it. The difficulty is to balance the desire of each team member to perform well as an individual (this is an especially strong motivating force in our field) with the necessity that the group performs well as a team. This is the heart of the motivational problem that must be resolved if the accelerator is to be reliable.

Deming proposed fourteen principles which should be adhered to if an organization is to build a quality product. They are given in the Appendix for completeness. Not all of these principles can be applied to a government-funded laboratory but the direction of the recommendations is extremely well-adapted to scientific projects.

4.1 Teamwork

Probably the most important factor that can improve the quality and reliability of the machine *and simultaneously reduce the cost* is the creation of a team or teams for the design and implemention. "I can't stress strongly enough that when people work together as a team, the job can be done with fewer people," (Tom Stallenkamp, General Manager of Large-Car Operations, Chrysler Corporation, explaining how the new LH sedans had been developed in three-and-a-quarter years instead of the usual four-and-a-half to five years and with fewer people).[5] The function of the Project Leader is then to coordinate the activities of the different teams and to ensure that people are working together simultaneously rather than sequentially.

Each team should consist of a small group of people with different skills and backgrounds who, together, search for the best solution. Coordination between the teams is vital to ensure that the "best solution" is also the best for the whole project, as discussed above.

The team concept is shown in the Joiner Triangle[6] (Figure 2), developed by Joiner Associates Consulting Group, who specialize in helping companies learn how to use quality management principles. It is intended to show the importance of the scientific approach (a given in our field) and the team approach to create quality. The creation of strong cohesive teams is one of the best guarantees of a reliable accelerator project.

Figure 2 The Joiner Triangle

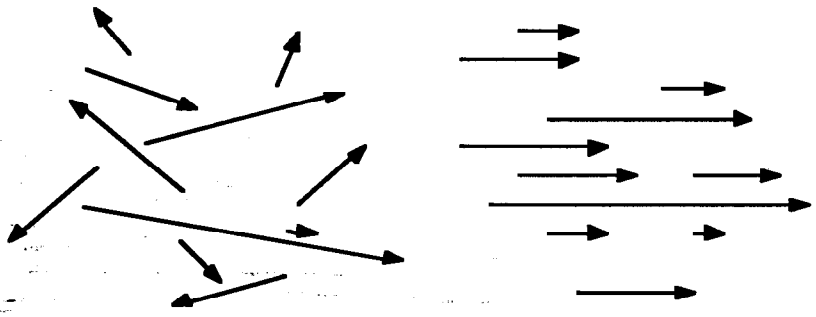

Everyone "Just Doing HisBest" Aligning the Arrow

Figure 3. Aligning the Arrows.

The team must have a common goal. General Motors calls this "aligning the arrows"—see Figure 3.[2] This is important in GM as "our strength has always been in our individual people and sometimes, they are not all aimed in the same direction." These words will resonate with many accelerator physicists and engineers. In our chosen field we are privileged to work with a large number of extremely talented individuals - but they are not always working towards a common goal. [The quotations are from W. Scherkenbach, who implemented the Deming philosophy at Ford Motor Company in the early 1980's ("Quality is Job 1"), and is now Group Director for Statistical Control and Process Improvement Methods at the Buick-Oldsmobile-Cadillac Group of General Motors. The adoption of the Deming management approach at GM is considered the major factor in the recent upsurge in quality in the Buick Division and the award of the Malcolm Baldridge National Quality Award to Cadillac in 1990.]

5 Management Tools

In the same way that there are now tools to help in establishing and track costs and schedules, there are also tools to establish a reliability budget and then to monitor progress during the project construction. These tools enable a reliability goal to be assigned to each sub-system, and to derive the expected reliability of the total project. The reliability budget provides a way to make each engineer feel part of the reliability of the whole project and helps to create a team.

5.1 Reliability Engineering

The goal of reliability engineering is as follows:

1. Provide a global reliability analysis of the entire project to assign a reliability "budget" for each system.
2. Provide reliability analysis of each system to assign a reliability budget to each subsystem.
3. Construct a probabilistic model of the project to derive the optimum decisions.

These steps provide the framework for assigning a reliability goal to each group.

Figure 3 shows some of the inputs and outputs into a probabilistic model (this representation is due to Dr. Jerrel Thomas of Failure Analysis Associates. He uses these techniques to help companies make rational decisions to improve the reliability of their products or plants).[7]

To this date, no accelerator project has succeeded in setting up a probabilistic model which covers the whole machine. The problem is the lack of input data, but some preliminary steps in this direction have been taken as a result of the Advanced Photon Reliability.

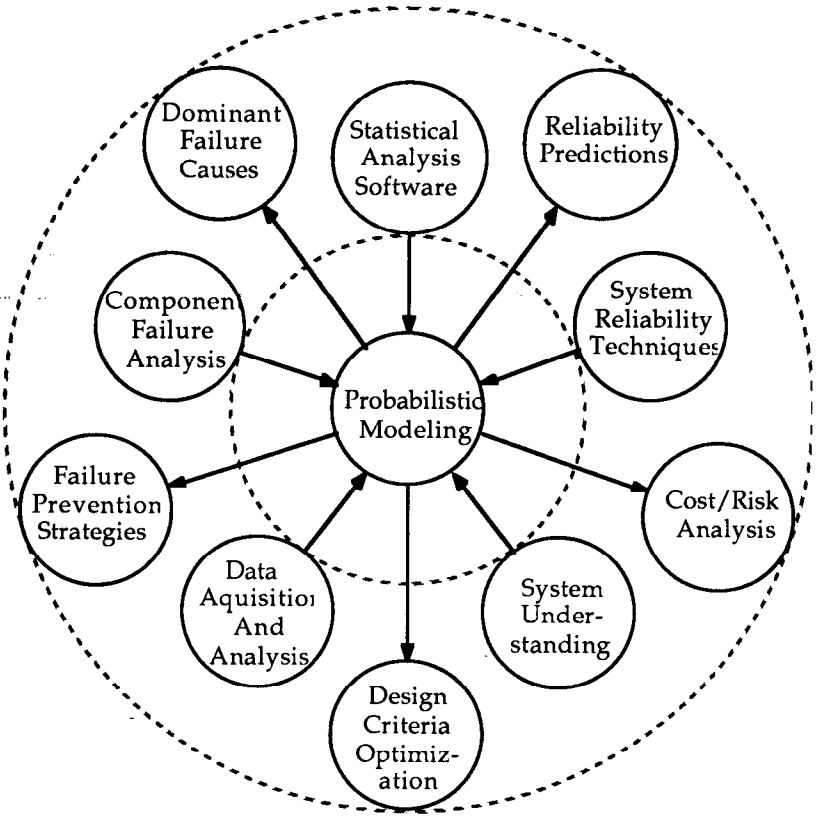

Figure 3. Inputs and Outputs in Probabilistic Modeling Workshop.[8]

The output of the probabilistic model should be a reliability goal for each group. This quantifies the balance between the desire of an individual to do the best possible job on the subsystem for which he or she is responsible and the need to ensure a proper balance in the project to avoid suboptimization. Until we learn how to build the model properly, we will be forced to assign quantitative reliability goals derived by extrapolation from existing machines.[9] Table 1 shows how data from the Tevatron and its injectors has been used to obtain a set of goals for the SSC Injector chain. This is an excellent first step. The next step would be to examine each category and estimate the slope of the reliability versus cost curve to determine which systems could be improved with little investment and which systems would suffer little from a reduced budget. The data from Fermilab constitute a valuable contribution to the reliability studies of other machines. Data from other operating machines would permit quantitative correlation of effects.

Assigning a reliability goal is an extremely important step in creating a cohesive team and helping all of the members of the project team to get a global view of reliability. It also ensures that everyone understands that reliability is vital to the project. "I have come to the conclusion that our western culture is one which 'manages by scorecard.' If quality is important, put it on the scorecard."[2]

Table 1 Fermilab Downtime/Availability and SSC Availability Comparison

FermiLab					SSC		
	Events	Total Hours	Mean Down Time	Avail-ability	Avail-ability	MTBF Hrs	SUB-SYSTEM
INJECTOR TOTAL	1607	881	0.55	0.894			
Linac	581	110	0.19	0.987	**0.983**	11.4	Linac
Booster	322	66	0.2	0/992	**0.988**	16.5	LEB
Main Ring	454	448	0.99	0.946	**0.935**	15.2	MEB
Utilities	47	118	2.52	0.986	0.992	315.1	All Utilities
Controls	158	75	0.47	0.991	0.995	291.7	Gas
Misc	45	64	1.42	0.992	0.996	329.1	Safety Intrlock
TEVATRON TOTAL	517	1121	2.17	0.865	**0.875**	17.3	HEB
Correctors	46	62	1.35	0.993	0.996	560.8	Corrector PS
Cryogenics	47	75	1.59	0.989	0.994	438.6	Cryogenics
Injection	10	16	1.61	0.998	0.999	2397.8	Inject/Abort
Magnets	8	333	41.63	0.960	0.876	2020.4	Magnets
Misc	84	99	1.18	0.988	0.993	326.3	Misc
Power Supply	74	134	1.81	0.984	0.991	310.7	Ring PS & Reg
QPM	43	60	1.40	0.993	0.996	590.2	Quench Protect
Quenches	113	195	1.73	0.977	0.996	1346.2	Mag.Quenches
RF	47	61	1.31	0.993	0.996	556.1	RF System
Vacuum	9	30	3.31	0.996	0.998	2143.1	Vac System
Cen He Liq	4	13	3.19	0.998			
Controls	32	43	1,35	0.995			
TEV+INJ TOTAL	2124	2002	0.94	**0.759**	0.987	69.8	ALL INJECTORS
					0.823		COLLIDER
					0.800		SSC TOTAL

5.2 *Reliability and Availability*

It is helpful to clarify some basic definitions at this point:

Reliability: The probability that a component does not break down in a given time.

Availability: The fraction of scheduled time that the machine is performing properly. It includes the time needed to fix a failed component.

In synchrotron light sources there is an extremely high premium on *Reliability* which tends to be less important in high energy physics. In High Energy Physics machines, *Availability* is the most important. This puts additional requirements on reducing the time for diagnosis and repair, as well as the time needed for returning the accelerator conditions to those prior to the failure.

5.3 Project Control

The reliability budget allows the performance of each subsystem to be quantified before the whole complex is running. Corrective action can be taken by evaluating the reliability of prototypes and early production units. This is usually done anyway in a qualitative fashion, but reliability analysis can do this quantitatively and cost-effectively.

Wherever possible, the use of stress tests should be employed in evaluating prototypes. For example, electronics can be tested at higher temperatures, equipment can be tested in a high vibration environment, magnets and power supplies can be tested at higher currents, water circuits can be over-pressured, etc. This type of prototype testing can identify weaknesses which can be removed by design changes or help to define the test program for the production units to ensure that the weakest link still meets specification.

Typically, the early testing of prototypes will help ensure that the reliability goals have been met. Note that it is bad for a subsystem to exceed the reliability goal by a large factor if this has been done by over-investment of scarce resources. The obvious result is that some other subsystem will be under-funded and the reliability of the whole project will suffer (suboptimization).

5.4 Contract Specification

It is usual in a contract to specify the performance, delivery, and payment schedule. Usually the contract will have performance bonus and penalty clauses. In addition, each component should be specified in terms of its reliability, and a reliability performance bonus and penalty clause included in the contract. This requires a lot of care in writing the specification, but the advantage of specifying reliability in the contract is that it can help ensure that the successful bidder has the expertise to successfully complete the contract.

Inspection of the completed product is not enough to ensure quality, it takes a collaborative effort and there has to be a financial incentive for the vendor to devote the effort needed.

The Japanese business style, which is now starting to be adopted in America,, is to develop a long-term relationship with the suppliers of vital components. The supplier then becomes part of the "team." In an accelerator project, it would be preferable to be able to develop components in collaboration with the vendor, but US government contract regulations make this extremely difficult, if not impossible, except in a few special cases. The recent DOE directive encouraging joint development contracts between government laboratories and industries (CRETA) is a welcome step in the right direction.

6 Fault Diagnosis

When the accelerator is completed, there will be an extremely high premium on rapid fault diagnosis and repair. An efficient control system is needed that can monitor every aspect of the machine operation and provide quick diagnosis of problems or equipment failures. If a fault is diagnosed in a piece of equipment, all the information that exists about the piece should be available in the control room. This will involve everything from the construction information, initial calibration, installation coordinates, maintenance history, drawings, etc. This goes beyond what has been provided by the control system in previous generations of machines.

6.1 Data Management

The entire data management system should be based on a commercial relational database such as ORACLE. This means that initial specifications should be recorded, procurements tracked, all the acceptance test data logged, all calibration data and alignment references stored, information about installation recorded and, later, the maintenance history kept up-to-date. All equipment should be tracked in the database: mechanical, electrical, services, instrumentation and controls. This is a large investment of effort that will pay off in the long run.

Since all of the new accelerator projects will need essentially similar software, this is an area where everyone stands to gain by collaborating on a common system. CERN has been using ORACLE since 1981 and it is now used for a very wide range of applications.[10] Exploratory talks are now taking place between SLAC, CERN and several other laboratories on development of a joint system.

Table 1. LEP Data Base functions

Inventory Management
- Laboratory and Office Equipment.
- Storage handling of mechanical and electrical equipment with information concerning volume, weight and location.

Project Management Facilities
- Budget estimation and control.
- Query facilities for supplier information, stores catalog items, current expenditure and contract follow-up data which is loaded into an ORACLE database from the ADP databases.

Documentation Catalog

Office Tools for Managing
- Key Distribution.
- Office and laboratory space allocation.
- Personnel phone numbers.
- Work requests.
- Overtime.
- Plastic card allocation for purchasing stores items, site access, etc.

Machine Construction Applications
- Planning facility using critical path analysis.
- Machine installation logistics.
- Drawings catalog and approval facility.
- Cables installation planning and management.
- Personnel protection system.
- Electronic circuit components database.
- Machine equipment testing and management.
- Transport and contract work management.
- Survey metrology database.
- Accelerator alarm system.
- Extended CAD/CAM database (EUCLID).

6.2 CAD System

It is advantageous to integrate the CAD system into the database to combine the advantages of both systems. Specifically, the database will be used for organizing, controlling, and accessing all documentation. The more common integration of CAD with CAE and CAM will be less important than the integration of the drawings with other data. One major problem that has still to be resolved is the choice between solid modeling and two-dimensional drawings. While everyone agrees that solid modeling is, in principle, better because it is more complete, there is no general agreement as to what constitutes an acceptable additional overhead to have it. Those laboratories that have invested in solid modeling too early have been left with an extremely negative impression. The format that will be used for storing drawings is also not settled. For two-dimensional drawings, standard protocols already exist for document exchange (IGES, DXF), but there is no generally accepted standard for three-dimensional solids.

6.3 Accessibility of Information

Failures can well occur on weekends or in the middle of the night (in fact this always seems to be the case). This means that the cognizant engineer or physicist would not be available immediately, so the system must work without him (or her). This also means that the database experts and the CAD experts would also not be available, so the system must work without them either. The system must therefore provide easy access to information for the operations staff. Updating or modifying the data should not be too easy to prevent errors creeping in.

6.4 Troubleshooting Example

Let us examine how the database would be used in the case of a focusing error being detected by the control system (software has been developed at SLAC for this type of error-checking).[11] The database should be interrogated for additional information on the power supply (what the voltage should be for a given current, what the response should be to a given control system instruction, etc.), so that the operators can perform a standardized test remotely. If the power supply is faulty, the database should be interrogated for location, model type, recent maintenance history, and location of spare parts. The circuit diagrams and board layouts are also needed from the CAD system. This information should be available to the operators in the control room as well as the maintenance crews.

If the power supply checks out, but a change in the resistivity of the magnet is detected, the ring must be accessed to perform tests on the magnet. This is a major source of lost time and the intervention must be kept to a minimum. All of the relevant calibration data on the magnet, including detailed drawings from the CAD system, should be available to the maintenance crew before making the access.

6.5 Information Management

An ideal system would provide:

a) An isometric view of the magnet, including the position of test points and the expected measurement values at these points.

b) A check out scheme involving a series of pre-established steps. Ideally, these steps would be performed using equipment attached directly to a portable computer terminal, linked to the machine control room and the database. (Portable here means that it can be mounted permanently on a golf cart).

c) As diagnosis proceeds, the technicians working at the remote location may well require detail and assembly drawings. This information should be available on the portable computer terminal and, ideally, on a portable hardcopy machine.

d) Finally, after the problem has been diagnosed, the defective part will be repaired or replaced. This change must be recorded, including the serial number of the new part (preferably barcoded) and the database should be updated automatically.

e) If a mechanical or electrical modification is made to the part, this must also be recorded and, at a later date, the drawings updated to reflect the change.

f) Following the repair, a test procedure should be run on the equipment to ensure that it performs properly. This should also be computer controlled and the new calibration numbers stored in the database.

Every element of this system is currently in use in industrial or scientific applications. This kind of integrated data system will be necessary to maintain high availability of the accelerator.

6.6 Requirements for the 'Ideal' Information Management System

a) The system must be comprehensive.
The system must be so easy to use that no-one will be tempted to bypass the system and keep information only in notebooks.

b) The information must be up-to-date.
Updating information must be easy so that the latest version is correct. This is particularly important for maintenance history. Older versions should be archived but available.

c) Information must be filed coherently.
The interface used to access information must be easy and self-evident. It must be easier to get the latest version from the system than to find a hardcopy of an old version in your office.

d) The information must be available.
While it is desirable that only an authorized person enter information into the system, it is vital that anyone be able to access the information.

7 Summary

Any new accelerator or collider should be designed from the outset to be reliable by adopting the single-risk philosophy. The most cost-effective way to make a reliable machine is to invest heavily in the concept and engineering design phases, later efforts to "add on" quality are bandaids; they are expensive and ineffective. The consequences of a failure, and its impact on the system availability, should be minimized by careful design. Downtime should be minimized by developing an integrated information management system to assist in diagnosis and repair.

It has been pointed out by many authors[12] that the cost of designing and building a reliable machine is not more than the cost of an unreliable machine. Designing in quality from the start and instilling a "quality attitude" reduces the amount of rework, simplifies installment and alignment, and improves efficiency.

The most important single element in building-in reliability is the experience and attitude of the Project Team. The experience and competence of today's accelerator builders is well appreciated, the equally important aspect of morale and attitude must be addressed if a real "factory" is to be built. It is the function of the Project Leader to create the appropriate environment.

APPENDIX

Deming's 14 Points

1. Create constancy of purpose for improvement of product and service.
   "A company's role is to stay in business and provide jobs through innovation, research, constant improvement, and maintenance."

2. Adopt the new philosophy.
   "Quality must become the new religion. We can no longer afford to live with mistakes, defects, poor workmanship, bad materials etc."

3. Cease dependence on mass inspection.
   "Inspection with the aim of finding the bad ones and throwing them out is too late, ineffective and costly. Quality comes not from inspection but from improvement of the process."

4. End the practice of awarding business on price tag alone.
   "Purchasing departments customarily seek the lowest-priced vendor. Frequently, this leads to supplies of low quality."

5. Improve constantly and forever the system of production and service.
   "Improvement is not a one-time effort. Management is obligated to continually look for ways to reduce waste and improve quality."

6. Institute training.
   "Too often, workers have learned their job from another worker who was never trained properly. They can't do their job because no one tells them how."

7. Institute leadership.
   "The job of a supervisor is not to tell people what to do or to punish them but to lead. Leading consists of helping people do a better job and of learning by objective methods who is in need of individual help."

8. Drive out fear.
   "Many employees are afraid to ask questions or take a position. The economic loss from fear is appalling. It is necessary for better quality and productivity that people feel secure."

9. Break down barriers between staff areas.
   "Often staff areas—departments, units, whatever—are competing with each other or have goals that conflict. They do not work as a team so they can solve or foresee problems."

10. Eliminate slogans, exhortations, and targets for the workforce.
    "These never helped anybody do a good job. They generate frustration and resentment."

11. Eliminate numerical quotas.
    "Quotas take account only of numbers, not quality or methods. They are usually a guarantee of inefficiency and high cost."

12. Remove barriers to pride of workmanship.
    "People are eager to do a good job and distressed when they can't. Too often, misguided supervisors, faulty equipment, and defective materials stand in the way. These barriers must be removed."

13. Institute a vigorous program of education and retraining.
    "Both management and the workforce will have to be educated in the new methods, including teamwork and statistical techniques."

14. Take action to accomplish the transformation.
    "It will take a special top management team with a plan of action to carry out the quality mission. Workers can't do it on their own, nor can managers. A critical mass of people in the company must be committed.'"

The fourteen points are taken from *Out of the Crisis,* by W. Edwards Deming.[4] The quotations in italics are from *The Deming Management Method,* by Mary Walton[13]. Not all of these points can be applied to an accelerator project, but the more of them that are applied, the higher the quality of the accelerator will be.

References

1 W. Edwards Demming. Interview in *Automobile Magazine*, Ann Arbor, Michigan, June 1991.

2 William W. Scherkenback. *Deming's Road to Continual Improvement*, SPC Press, Inc., Knoxville, Tennessee, 1991.

3 Laotzu. *The Way of Life*, trans. Witter Brynner, Perigee Books, New York, 1944, Cited in Reference 2.

4 W. Edwards Deming. *Out of the Crisis*, MIT Center for Advanced Engineering Study, Cambridge, 1986.

5 Tom Stallenkamp. Interview in *Automobile Magazine*, Ann Arbor, Michigan, June 1992.

6 Peter R. Scholtes, et al. *The Team Handbook*, Joiner Associates, Madison, Wisconsin, 1988.

7 Jerrel M. Thomas and Caleb S. Davis. *A Synthesis of Failure Prevention and Reliability Methods*, FAA-M-81-2-9, Failure Analysis Associates, Menlo Park, California, 1981 (unpublished).

8 Advanced Photon Source Reliability Workshop, Argonne National Laboratory, Argonne, Illinois, January 29—31, 1992 (unpublished).

9 K. Dixon. Presentation to the LEB Preliminary Design Requirements Review, SSC , Waxahatchie, Texas, 22 July 1990 (unpublished).

10 Josi Schinzel. LEP-DI/JS (89-134), 1989 and private communication.

11 S. Kleban, M. Lee and Y. Zambre. *GENI: A Graphical Environment for Model-Based Control*, N.I.M. A293, 1990.

12 Philip B. Crosby. *Quality is Free*, McGraw-Hill, 1979.

13 Mary Walton. *The Deming Management Method*. Perigee Books, New York, 1986.