The Politics of Accelerated Stress Testing

Edmond L. Kyser, Cisco Systems, Inc., San Jose, California
Eugene R. Hnatek, Xilinx, Inc., San Jose, California
Mark H. Roettgering, eLoyalty Corp., Sacramento, California

The technical literature and various technical conferences delve into the myriad details of the Environmental Stress Screen (ESS) process, the ESS profiles to be used for testing, the required equipment characteristics, etc. Almost everything that can be written about the virtues of ESS and the inherent technical details has been written. We contend that it is not the technical aspects of ESS that dominate decision-making. The real issues, for most companies, are of a political nature. ESS implementations become political when the functional organizations that bear the short-term costs of ESS do not get credit for the long-term benefits. Various factions within most large corporations rise to the surface to question processes such as ESS from a self-serving viewpoint. Justifying the need for continuing with ESS eats up a lot of time in meetings, evaluating databases and developing position presentations. In this article we discuss these commonly encountered political issues, provide a process for their resolution and conclude with recommendations for corporate ESS management.

Today's fast time to market and concern with low price may be taking our focus off quality and reliability. Frank Burge of *Electronic Engineering Times* in his September 27, 1999 editorial put it this way. "In a world where price is king, are we painting ourselves into a corner – a corner where design quality gives way to price or time, eliminating steps in the design verification/test process or choosing suppliers strictly on price? Are we back to making the numbers at any cost?"

The decision whether or not to perform ESS on a specific product is a typical example of the quality vs. cost problem with which many companies struggle, including our own. One of the problems in being able to make a decision based on data is the fact that very little real data (whether from current or equivalent products) is available to determine the value of ESS. Typically at stake are millions of dollars in investment capital, thousands of square feet of manufacturing floor space, tens of person-years, and the reputation for quality and possibly the profitability of the corporation. A typical product flow diagram for Accelerated Stress Test (AST) processes is shown in Figure 1. Many separate stakeholders of the corporation are involved in this complex process, the core of which is manufacturing ESS: Product Development, Manufacturing, Field Service, Engineering Services, Sustaining Engineering and Information Services. Figure 2 illustrates a common hierarchy of these groups, each of which typically has its own agenda and point of view. Traditional guidelines, established product requirements documents and standard procedures may not be sufficient or appropriate. Benchmarking is difficult. Evangelizers for specific approaches to increasing reliability are quick to offer their services and opinions, often at loggerheads with one another. Industry standards are rare and often ambiguous. Perhaps most significantly, the benefits (and the associated costs) realized from the program do not accrue proportionately to the functional units that bear the costs.

Acronyms

AST	Accelerated Stress Test
CPU	Central Processing Unit
ESS	Environmental Stress Screen
FCS	First Customer Ship
IBP	Indifferent Buying Price
MTBF	Mean Time Before Failure
NPS	Net Present Saving
NPV	Net Present Value
NTF	No Trouble Found
PWA	Printed Wire Assembly
WACC	Weighted Average Cost of Capital
WIP	Work in Progress
WPC	Whole Product Cost

The goal of an ongoing AST program, such as implementation of manufacturing ESS, is to make cost effective improvements in the field reliability of the hardware being tested. Figure 3 shows a normalized field failure distribution for five recent Tandem products, all of which were subjected to 100% manufacturing ESS.

All of the products represented in Figure 3 show the same pattern of a high initial return rate that decreases more or less asymptotically to a stable return rate in about two years. This is a classic characteristic of products that are most likely to benefit from an ESS program.

Survey of Attitudes on ESS

During a recent IEEE workshop on Accelerated Stress Testing, we conducted a survey of attitudes towards AST to determine if there was a common experience among industry practitioners that could be leveraged as the science evolves. The issue was defined as: "Within your company, where do you see support for or opposition to ESS, and why?" The organizational results are summarized in Figure 4. The respondents' reasons behind the support and opposition are shown in Tables 1a and 1b.

As the comments in the tables indicate, many of the reasons given are similar and can therefore be combined. The resulting grouped categories of opposition and support are shown in Tables 2a and 2b. In cases where the reasons appeared to be ambiguous, require other processes to be considered, or deal with educational or organizational issues, the category *out of scope* was used. Out of scope does not imply that the reasons are invalid, just that they will not be addressed in detail here.

As Tables 2a and 2b indicate, we are left with two potential sources of benefit and a large bucket containing several cost factors: additional time, reduced manufacturing yields, test costs and repair costs. The fear of product damage will be handled explicitly as part of the question of improved reliability.

The survey results we have been discussing represent the opinions of 32 individuals from 22 corporations active in ESS. One of the most striking results is that the same issues or organizations appear in both the positive and the negative columns. Obviously, there are strong differences of opinion and a lack of mutually acceptable (accurate and meaningful) data upon

---

The work from which this article was written was performed while the authors were employed by Tandem Computers Division of Compaq Computer Corp., Cupertino, CA. This article is based on their paper which was presented at ESTECH 2000, the Institute of Environmental Sciences and Technology 46th Annual Technical Meeting, Providence, RI, April 2000.

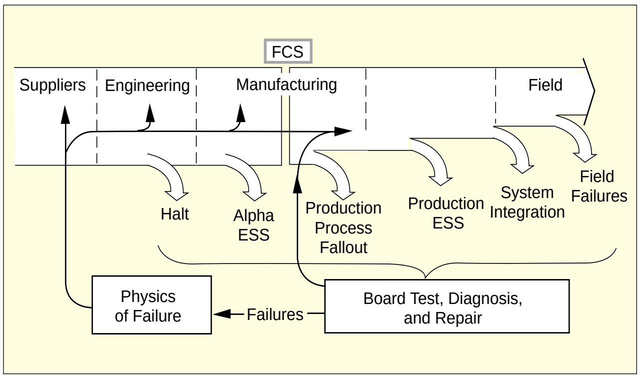

Figure 1. Product flow for AST programs.

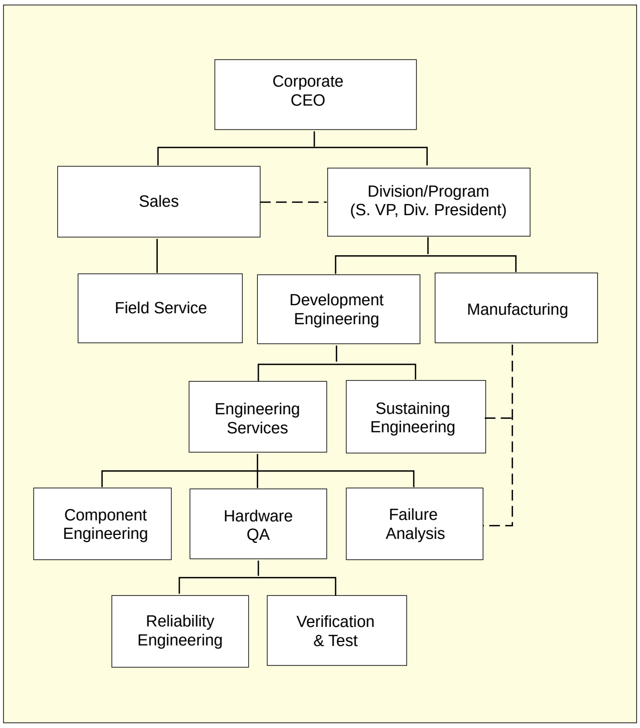

Figure 2. Generic corporate reporting structure.

which to base decisions. This is equivalent to stating that there is a high degree of uncertainty about many important aspects of a manufacturing ESS program. Without a structured methodology in place to address this uncertainty, a common ground within the corporation may never be found.

We maintain that what is needed is a common metric of success that accommodates all of the above reasons – since all are valid in the opinion holder's frame of reference. How is one going to net out all the above positives and negatives? The problem can be formulated as follows:

$$\begin{bmatrix}\text{Positives}\\ \text{Benefits}\\ \text{Support}\end{bmatrix} - \begin{bmatrix}\text{Negatives}\\ \text{Costs}\\ \text{Opposition}\end{bmatrix} = \begin{bmatrix}\;?\;\end{bmatrix}$$

We propose that the *metric of success* is the dollar and the method of *netting out* the positives and negatives is to discount all cash flows to net present value and calculate a net present cost. Instead of taking a best guess at exact amounts of the costs and benefits, all uncertainty should be explicitly stated so that conflicting opinions about possible outcomes can be addressed simultaneously. This process is detailed in the following section.

Decision Model

A review of the pluses and minuses of ESS raised by the practicing community quickly reveals the major source of organizational problems that arise in an ESS implementation. The majority of the costs are easily identified and can be quantified with a high degree of accuracy. The manufacturing organization bears essentially all costs – using many common manufacturing metrics (end-to-end yield, inventory turns, WIP days, etc.) and ESS is a negative. On the other hand, the benefits, while identifiable, possess the following characteristics. They are highly uncertain, difficult to quantify with any degree of accuracy, difficult to measure, require an explicit value statement by management and are not immediately realized. The benefits are realized by the corporation as a whole, essentially through downstream cost-avoidance (lower field service and warranty costs) and through increased sales (product reputation).

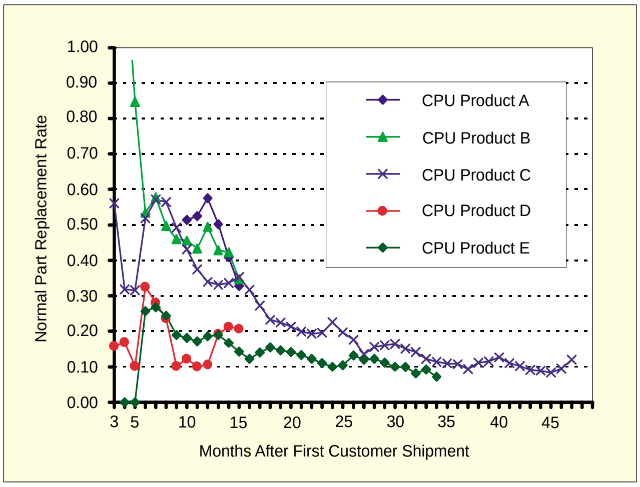

Figure 3. Field data – part replacement rate.

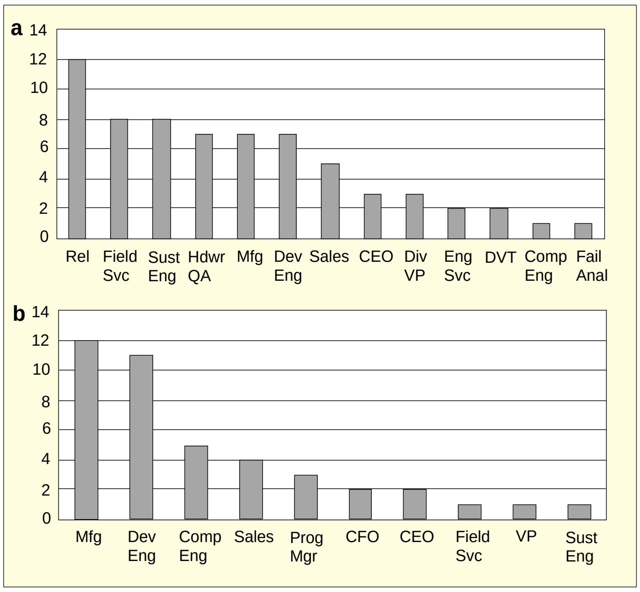

Figure 4. a) Organizations favoring ESS. b) Organizations opposing ESS.

It can be said that the problem with ESS acceptance is that it is high in both organizational and technical complexity. Technical complexity arises from the large number of strategic and operational decisions and processes that need to be in place for an ESS program to function in an efficient manner. Organizational complexity is inherent when:

- Costs and benefits are realized by different groups.
- Uncertainty allows a variety of advocates and opponents to champion opinions without fear of refutation by data.
- There is a lack of strong cross-functional leadership from management.

Unfortunately, management attempts to solve problems of this nature by attacking the "people problem" first through team-building, facilitation, consensus-building, etc. Despite these well-intentioned tactics, the underlying technical complexity invariably remains and with it, the conflict. What is needed is a framework in which to solve the technical complexity *first*. Through creating a technically accurate and compelling business model, organizational disagreements can be addressed in a methodical and rigorous manner. Arguments such as: “Doesn’t believe the benefits” can be addressed by explicitly addressing which parts of the model are inconsistent with the beliefs of the opponent. Consequently, if the model is accepted and the inputs are agreed to, the resulting netted-out cost or benefit of ESS should stand on its own, leaving nebulous and ambiguous arguments without legs.

We propose a normative decision model as the best method for solving the technical complexities of ESS. In this framework, we must first clearly identify exactly what we are modeling. Stated here:

“What is the net present value of all future product costs for a unit which is to undergo ESS subtracted from the net present value of all future product costs for a unit which will not undergo ESS?” We call this quantity Net Present Savings or NPS.

The NPS we compute is a marginal savings on a per-unit basis. This eliminates the requirement to consider facility and capacity issues. We also assume that all other manufacturing processes remain the same. We do not explicitly consider the potential benefit of reduced run-in times here, although the framework allows for it. One last assumption is that we are discussing a particular ESS screen for a particular product – the selection or modification of screen parameters to maximize NPS is not performed here, although we have used the methodology to do parameter optimization at Tandem/Compaq. The model and theoretical results discussed in the following analysis were built using Analytica® analysis software from Lumina Decision Systems.

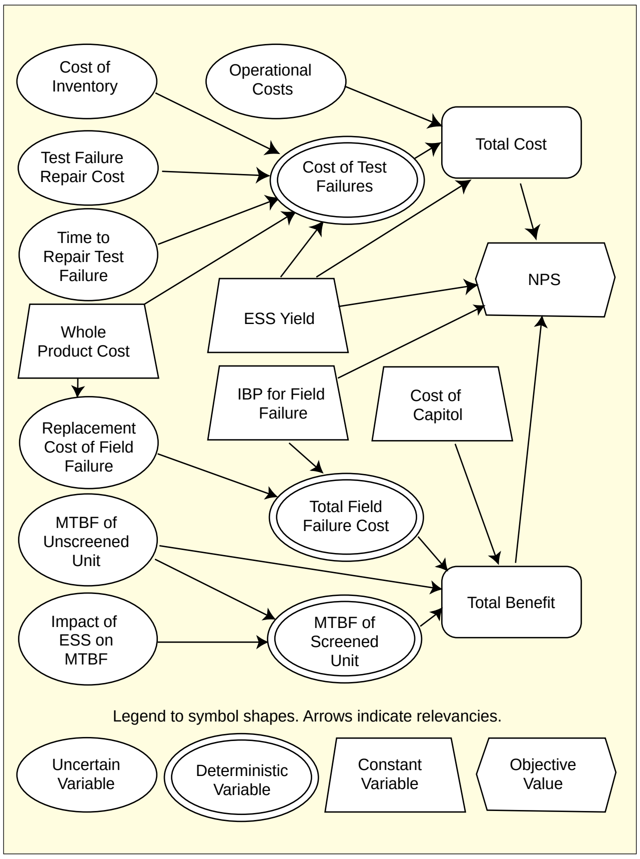

Figure 5. Influence diagram.

The influence diagram of Figure 5 illustrates the factors that have been included in our model. Based on the factors identified in the ESS survey, we will model seven uncertain or random variables (single ovals) and four constant variables (trapezoids). Double ovals indicate deterministic variables (those variables known exactly once the inputs are chosen). A summary of the model variables is given in Table 3. Values are representative of our experience with a broad range of CPU products.

We use the lognormal distribution to express the uncertainty in almost all random variables included in this model. The lognormal has a sharp lower bound of zero and is positively skewed. For most cost and time parameters, these characteristics are highly desirable.

Table 1a. Reasons for supporting ESS.

Key	**#**	**Stated Reason**	**Comments**
a	11	Increased reliability/quality	Hard to measure; hard to quantify benefits; compare to *n*
b	9	Sales advantage/ customer satisfaction	Same as *a*, but more difficult to quantify
c	6	Reduce field service costs	Equivalent to *a*
d	4	Reduce DOA/early life fails	Equivalent to *a*
e	3	Identify failure modes in house	Benefits seen only by redesigning to avoid failure modes
f	2	Better product	Equivalent to *a*
g	2	Reduced field returns	Equivalent to *a*
h	2	More efficient than run-in	Weibull analysis can help determine this
i	1	Identify process failures	Equivalent to *a*+*e*
j	1	Improve yields	Equivalent to *e*

Table 1b. Reasons for opposing ESS.

Key	**#**	**Stated Reason**	**Comments**
k	11	Additional cost	Virtually all opposition is cost based. Easier to measure than benefits
l	9	Outside of component specs, design limits	Equivalent to *n*
m	6	Additional WIP Time	Additional step assumes all else equal – part of *k*
n	4	Decreases manufacturing yields	Easy to measure, easy to quantify. Compare to *a*
o	3	Afraid of damaging good product	See comments on *a* – effect on reliability is uncertain
p	2	Seen as critical of known good process	'Known good' implies improved reliability is of no benefit or screen is no good
q	2	Doesn't understand process	Education issue
r	2	Difficult test to run/ diagnose failures	Part of *k*+*t*
s	1	Additional handling problem	Equivalent to *k*+*m*
t	1	Repair costs	Part of *k*,*s*
u	1	Run-in more efficient	See *h*
v	1	Doesn't believe in benefits	See *a*

Table 2a. Revised reasons for ESS support (benefits).

Key	**#**	**Stated Reason**	**Comments**
a	25	Increased reliability/quality	Hard to measure; hard to quantify benefits; compare to *n*
b	9	Sales advantage/ customer satisfaction	Same as *a*, but more difficult to quantify
*	7	Out of scope	

Table 2b. Revised reasons for ESS opposition (costs).

Key	**#**	**Stated Reason**	**Comments**
k	19	Additional cost	Virtually all opposition is cost based. Easier to measure than benefits
n	15	Decreases manufacturing yields	Easy to measure, easy to quantify. Compare to *a*
m	5	Additional time	Additional step assumes all else equal – part of *k*
o	5	Afraid of product damage	See comments on *a* – effect on reliability is uncertain
*	4	Out of scope	

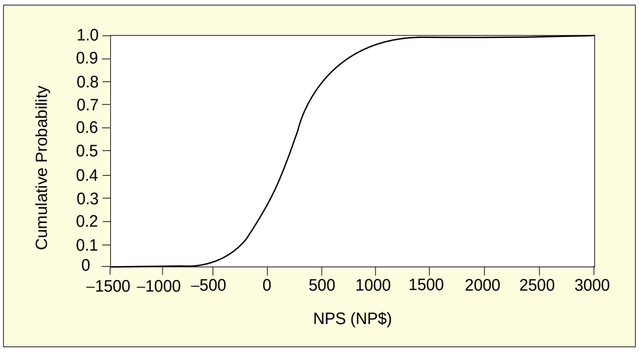

Figure 6. Probabilistic model output.

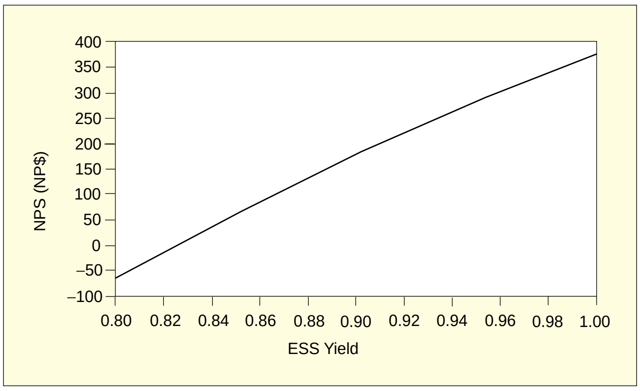

Figure 7. Sensitivity of NPS to variations in screen yield.

Field Failure Cost is one of the more difficult parameters in the model for most corporations to assess. We have broken it into two parts based on the results of the conference survey discussed above: replacement cost or warranty cost; and reputation cost. Replacement cost can be assessed directly through careful consideration of all contributing costs, but reputation cost (re-buy, word-of-mouth, etc.) may best be derived by discussing the Indifferent Buying Price (IBP) of a field failure.

Suppose there was a wizard who was able to perform the following feat: moments before a field failure occurs, the wizard calls the CEO of your company and offers to allow you to secretly swap out the failing unit before the failure takes place, for a price. The CEO’s IBP for the field failure is the price up to which he/she is willing to pay the wizard for this service – the CEO would pay any lower price (in addition to the replacement cost) but would refuse to pay any more. Although IBP for field failures could be different for the same product depending on customer and application differences, a well thought out value for the IBP will be equivalent to the reputation cost of a failure. Both replacement and reputation costs are valued at the time in the future when the failure takes place. With this groundwork in place, our model simply computes the Total Benefit per unit for performing ESS as the difference between the present value of the total failure cost of a screened unit vs. an unscreened one.

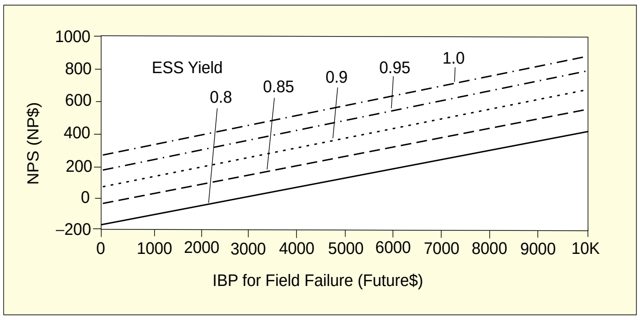

Figure 8. Sensitivity of NPS to yield and IBP.

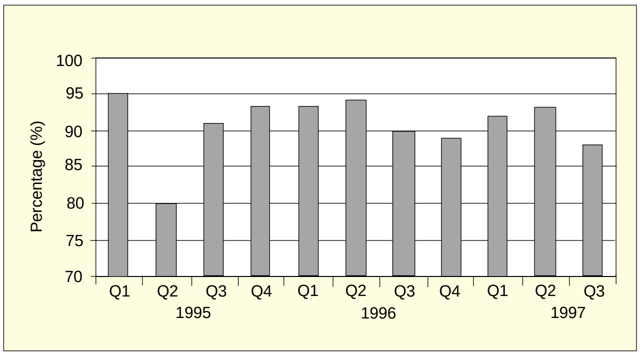

Figure 9. PWA ESS yields.

Figure 6 displays the results of the model discussed above. The expected value of NPS is $180/unit, a good return on a $50 test. The cumulative distribution of NPS contains much more information, however. Indeed there is a 30% chance that this generic ESS program will lose money on a per unit basis. On the other hand, there is just as likely a chance that a net benefit of more than $325 per unit will be realized.

Any dispute with the conclusion that the hypothetical ESS program represented by this model and corresponding parameters is a *good bet*, should be stated in the context of the model or its parameters rather than with more abstract terms. A good bet is a deal with an uncertain but positive expected outcome.

Table 3. Model variables and their descriptions.

Key	Variable Name	Units	Comments	Value
A	Cost of Inventory	%/week	Includes depreciation and liquidity effects	Lognormal(0.5, 1.5)
B	Test Failure Repair Cost	NP$	Material and labor costs for debug and repair	Lognormal(1500, 1.5)
C	Time to Repair Test Failure	Weeks	WIP time	Lognormal(6, 1.5)
D	Replacement Cost of Field Failure	Future$	Material and labor (warranty costs)	Lognormal(H/2, 1.25)
E	MTBF of Unscreened Unit	Years	Mean Time Before Failure	Lognormal(5, 1.5)
F	Impact of ESS on MTBF	%	Factor by which ESS improves unit MTBF	Normal(20%, 15%)
G	Operational Costs	NP$	Variable cost only (no fixed costs)	Lognormal(50, 1.5)
H	Whole Product Cost	NP$	Used to calculate inventory, depreciation and replacement costs	5000
J	ESS Yield	–	Probability of passing ESS screen	90%
K	IBP for Field Failure	Future$	–	2000
M	Cost of Capital	%/year	Time vs. money discount rate	15%
N	Cost of Test Failures	NP$	Total cost of fail, debug, repair cycle	$((1/J)-1(B+H(((1+A)^{C})-1))$
P	Total Field Failure Cost	Future$	Includes direct and indirect costs	D+K
R	MTBF of Screened Unit	Years	See E	E(1+F)
T	Total Cost	NP$	Total additional cost of ESS	N+G(1/J)
W	Total Benefit	NP$	Total downstream benefit per unit derived from ESS	$(P/1+M)^{R})-(P/(1+M)^{E}$
X	NPS	NP$	Per unit net present savings	W-T

NP$ is Net Present Dollars.
Future$ is dollars not discounted to present value.

Lognormal(*x*,*y*) is a distribution with mean *x* and geometric standard deviation *y*.
The range [*x*/*y*, *x***y*] contains about 68% of the probability mass.

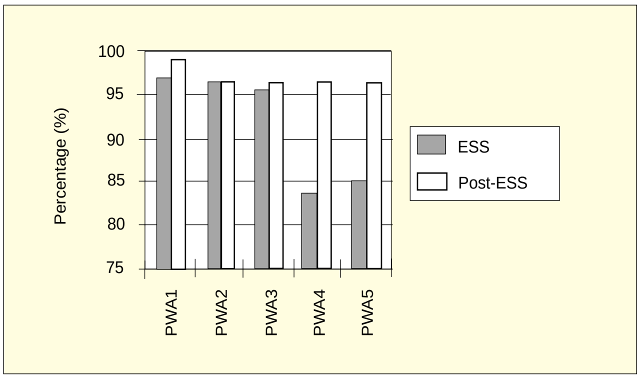

Figure 10. Manufacturing yield by PWA type for 3Q97.

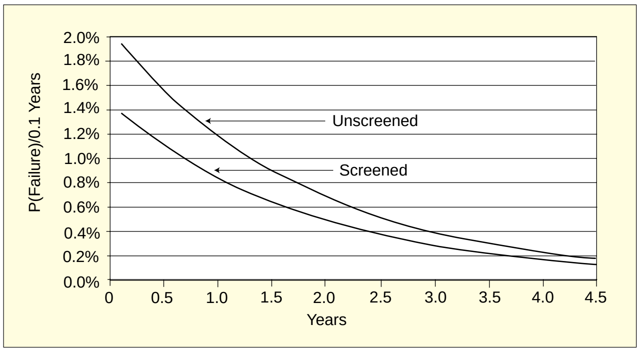

Figure 11. Failure rate vs. time for split population comparison of screened and unscreened field performance.

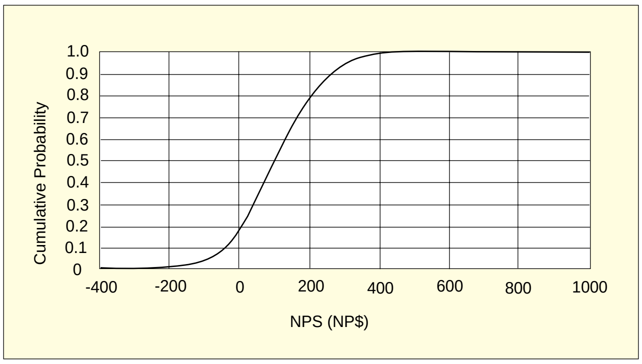

Figure 12. Calculated NPS using value model and field data.

By encoding differing points of view in the form of parametric uncertainty and incorporating all stakeholder's concerns into the model structure, discussions are moved from the political realm into the technical one.

Although this is a generic example, it is useful to demonstrate how insights may be gained through further analysis. One such analysis may be to address the following concern: "What if the screen yield required to achieve a 20% improvement in MTBF is either higher or lower than 90%?"

Figure 7 shows the mean (average or expected value) NPS as a function of screen yield. As can be seen, any screen parameter set with a yield higher than 83% would be considered valuable. Similarly, the effect of IBP on the value of our generic ESS program can be investigated graphically in Figure 8. For a yield of 90%, the program would still have a mean value of $60 per unit even if the reputation cost (IBP) of a failure were valued at $0.

Table 4. Input variable importance.

Variable	Importance
Impact of ESS on MTBF	0.871
Replacement cost of field failure	0.337
Test failure repair cost	0.211
MTBF of unscreened unit	0.086
Time to repair test failure	0.040
Operational costs	0.013
Cost of inventory	0.001

Finally, it is enlightening to examine the degree to which the uncertainty in the input variables contributes to the variation in the output variable, NPS. Table 4 lists the absolute rank-order correlation between NPS and the listed uncertain inputs. This analysis indicates that (as expected) the greatest opportunity to reduce uncertainty in the value of this hypothetical ESS program is to refine the impact of ESS on MTBF estimate. Conversely, expenditures of effort on refining any of the bottom four variables in the table will do little to reduce the uncertainty in the estimate of per unit NPS.

Data

When an ESS program is initiated, there must be decisions made in the face of many uncertainties. As the program progresses, real data become available and the initial probability estimates can be replaced by real numbers. In this section, we show some of the manufacturing and field data collected at Tandem division of Compaq relating to our ESS program and answer some of the issues raised earlier.

One roadblock to ESS is usually stated as follows: "We can't afford the yield loss in manufacturing caused by ESS." Translation: "We need to ship every unit we build in order to make our revenue target. We can't worry about reliability at this point. ESS yields cost us money, both in shippable units (lost revenue) and reworked or scrapped PWAs." Let's look at two examples of the yield of CPU PWAs subjected to manufacturing ESS. Figure 9 is a composite bar chart showing combined manufacturing ESS yields for five different CPU products. Note that the ESS yield remains essentially constant. As process and component problems were solved, new problems emerged and were addressed. In this case, given the complexity of the products, 100% ESS was required for the entire life of each product. Figure 10 shows a breakout detail of the products included in the last bar of Figure 9 and adds post-ESS yields for each of the five. This chart shows the value of conducting ESS in production and the potential impact of loss in system test or the field if ESS was not conducted. Notice the high ESS yield of mature PWAs (PWAs #1-#3) but the low ESS yield of new boards (PWAs #4 and #5). The benefit of ESS for new products is evident here. Note particularly that the Post-ESS yields for both mature and immature products are equivalent, indicating that ESS is finding the latent defects. Nonetheless, the value of ESS must be constantly evaluated. At some point in time when yield is stable and high, it may make sense to discontinue its use for that PWA/product.

The ideal data set would allow the creation of a failure rate vs. time plot or hazard function for a split population of screened and unscreened product. Unfortunately, this type of data is never available until several months or even years have elapsed. Figure 11 displays one real-world example. It took careful data mining from over five years of run-time and of more than 2000 total field installations to produce this information for a Compaq CPU server product. If this data were known ahead of time, ESS implementation decisions would have been simple – screen yields would be known (92%) and so would the effect of the screen on the MTBF of shipped product (14%). Assuming all other model parameters discussed earlier apply to the product producing this field data, the cumulative distribution in Figure 12 reflects the value of its ESS program. The expected value is $110 with only a 20% chance of being negative.

The amazing thing about these field data when applied back into the model is that our screening decision and our expected value remain virtually the same. The main benefit we gained through the gathering of field data is that our uncertainty about the true value of NPS has been reduced. The 10%-90% range has shrunk from [–260, 650] to [–20, 260]. Considering that five years had elapsed after data recording began, the product is long past its end-of-life. The reason a structured framework for dealing with this uncertainty about the future is so valuable is that it allows corporations to make the best decision possible – at the time the decision needs to be made.

Conclusions

Decisions relating to Environmental Stress Screening involve political aspects within a corporation to far greater a degree than almost any other manufacturing process. A correct decision on whether or not to perform ESS on a particular product requires a measure of success (metric) that is acceptable to all of the different corporate stakeholders. We suggest that Net Present Value of all associated costs and benefits is the metric of choice.

In addition to a robust cost-benefit model, the nature of ESS requires that there be a strong company champion for ESS at a high level who can adjudicate the inevitable disagreements, focus on corporate goals and provide direction for the discipline. However, the champion needs to remember that the goal of ESS is not *reliability at any price*, but rather *reliability at the right price*.

Finally, since ESS decisions assisted by a model of NPS will be based on probabilities estimated before actual data exist. Yield and failure data must be obtained to verify the initial probabilistic estimates. The resulting data should be used to improve screening decisions for future products and to modify current practices. The politics of fear and regret should not be allowed to interfere with a sound decision making process.

References

1. Henrion, M., Arnold, B., et al, *Analytica® Users Guide*, Lumina Decision Systems, Los Gatos, CA, 1999.
2. Roettgering, M., Kyser, E., “A Decision Process for Accelerated Stress Tests,” EEP Vol. 26-2, ASME, 1999.

The authors can be contacted at: eugene.hnatek@xilinx.com.