NATIONAL AERONAUTICS AND SPACE ADMINISTRATION
Technical Report

Case Study of 'Engineering Peer Meetings' in JPL's ST-6 Project

November 20, 2003

PREPARED BY:
Lawrence P. CHAO
Design Division
Stanford University
Department of Mechanical Engineering
Terman 551
Stanford, CA, 94305

MENTOR:
Dr. Irem Y. TUMER
Computational Sciences Division
NASA
Ames Research Center
MS 269-3
Moffett Field, CA, 94035

EXECUTIVE SUMMARY

This design process error-proofing case study describes a design review practice implemented by a project manager at NASA Jet Propulsion Laboratory. There are many types of reviews at NASA: required and not, formalized and informal, programmatic and technical. Standing project formal reviews such as the Preliminary Design Review (PDR) and Critical Design Review (CDR) are a required part of every project and mission development. However, the engineering peer reviews that support teams' technical work on such projects are often informal, ad hoc, and inconsistent across the organization. This case study discusses issues and innovations identified by a project manager at JPL and implemented in "engineering peer meetings" for his group.

Table of Contents

1. BACKGROUND

1.1. Design Process Error-Proofing

From the lessons of the Mars Climate Orbiter and Mars Polar Lander to the Intel processor and Firestone ATX tires, the consequences of design failures are well-known, but errors to some degree have still been accepted as part of design. By learning from past design errors, the key factors can be identified. The design process includes knowledge, analysis, communication, execution, change, and organization at different levels. Design errors usually occur in one of these areas of the design task, depicted in Figure 1. (Chao 2003) In any task, the agents involved must perform an *analysis* of the situation to determine what must be done. Design tasks require *knowledge* by the agents to perform the task. The agents must *communicate* the requirements to begin the task and the completed work once they *execute* the task. However, at all times, the agents and information are subject to *change* from other areas in the organizations, as well as noises or uncertainties.

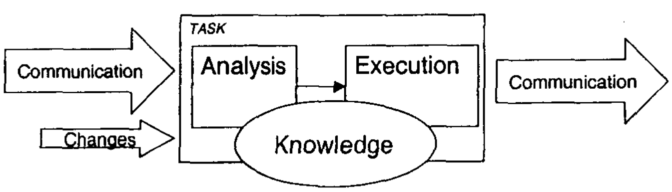

Figure 1: KEY FACTORS IMPACTING ERRORS MAPPED ON TO A TASK

While specifications are set by manufacturing, the challenge for design is that information and even requirements are still very dynamic and fluid. From complex engineering calculations to forward-looking business decisions, there are a range of issues to be addressed. **Design process error-proofing** (Chao 2001) attempts to predict, detect, and prevent problems that occur during product development that affect product quality, cost, and time-to-market.

Leading organizations currently have limited methods in predicting and preventing errors from occurring during design. Many tools which reduce product development errors are currently available - traditionally design reviews, but also structured methods. Project management and analysis through design structure matrix (DSM), design process failure modes and effects analysis (dpFMEA), and project quality function deployment (projQFD) are also important. (Chao 2003) By framing, applying, and integrating such with information systems can aid an organization towards identifying and solving design errors and error-proofing the system.

If possible, error-proofing measures should be proactive and preventive. The application of these error-proofs to the problem areas should be fairly straightforward. The error-proofs found or developed should apply to specific areas already. In addition to identifying the proper error-proofs, the team should develop a plan for implementation, including identifying responsible parties to oversee implementation and setting dates to begin and complete the work. In their implementation, we have found and developed three dimensions. The **error factors** refer to the nature of the error, including *knowledge, analysis, communication, execution, change,* and *organization* components, that the corrective action is designed to mitigate. **Solution level** refers to the level at which the error-proof tries to fix: either at the *problem*, the *process*, or the *system.* The **robustness level** refers to the method by which the error-proof mitigates future errors, ranging from *prevention, detection, inspection*, to *process improvement and aid.*

Higher-level error-proofs are more robust because the prevention of an error can be absolute. It does not rely on someone to act upon a cue. Even if an error is detected, it may not be mitigated. Design reviews can miss errors. They are weak processes extremely reliant on the skills and knowledge of the reviewer. This is not to say that higher level error-proofs are always "better." The lowest level error-proofs can refer to design tools, like design for manufacturability methodologies. Though they may not prevent a specific error, they help with creating a better design and prevent failure to satisfy the customer, for example, whereas the higher levels are more geared towards preventing human errors, such as errors in communication.

Ideally, these design process devices should concentrate on active prevention rather than detection since, during the design process, the desired outcome or right answer is unknown. Instead though, engineers can rely on their intuition, experience, or even the initial specifications. In the end, design process error-proofs should share attributes of good manufacturing poka-yoke. They should be simple and be designed for a specific problem or attribute. Engineers and managers should be able to implement them several times and integrate them into the design process rather than patch them in and add complexity to the system.

By understanding the nature of design errors and the strengths and weaknesses of the tools used to combat them, organizations can work towards higher levels of error-proofing. Design process error-proofing is not a universally understood, much less universally implemented, approach. As important as any method is the understanding that there are system-level issues related to design errors that can be improved.

1.2. Case studies

This case study is one of several benchmarking development process error-proofing efforts worldwide. The goal of these studies was to share and educate on product development errors in the hopes of better understanding them to better implement error-proofing methods.

This effort collected error-proofing examples for product development through real examples. Design process error-proofing is a new field with a lack of resources and understanding, but many organizations already have projects which undertake similar challenges. This is an opportunity for an organization to document and explore an existing problem and solution.
Other cases include:

- *Lean Engineering and Probabilistic Design* (GE Aircraft Engines)
- *Data-Driven CAD Robustness* (GE Aircraft Engines)
- *The Gate Model and Project Platform Management* (ABB)
- *Engineering Process Templates* (GE Aircraft Engines)
- *System Engineering Model* (Hewlett-Packard)
- *Use Case Library* (Hewlett-Packard)
- *Knowledge Engineering* (General Motors)

Additional information is available at the URL:

http://mml.stanford.edu/Research/DPEP/ep.casestudy.html

This report describes lessons and experiences in engineering peer reviews in design at NASA from a project manager (**PM**) of the New Millennium Program's Space Technology 6.

1.3. NASA Reviews

For decades, the National Aeronautics and Space Administration has applied effective design principles with appropriate peer reviews and periodic systems design reviews to result in high reliability aerospace design. NASA has a well-defined life cycle which consists of several phases. Like many organizations, NASA uses these phases as a means to organize decision points. Requirements definition begins in Phase A, with refinements and baseline occurring in phase B. Lower level requirements are derived between phases B and C, and major requirement definition is completed for all levels by phase C. Design reviews are at key transition points along this life cycle, outlined in Figure 2

All NASA missions and spacecraft are subject to a technical design review process. The primary objective of this program is to enhance the probability of success by identifying potential or actual design problems in a timely manner. There are a number of system reviews which are performed throughout the lifecycle.

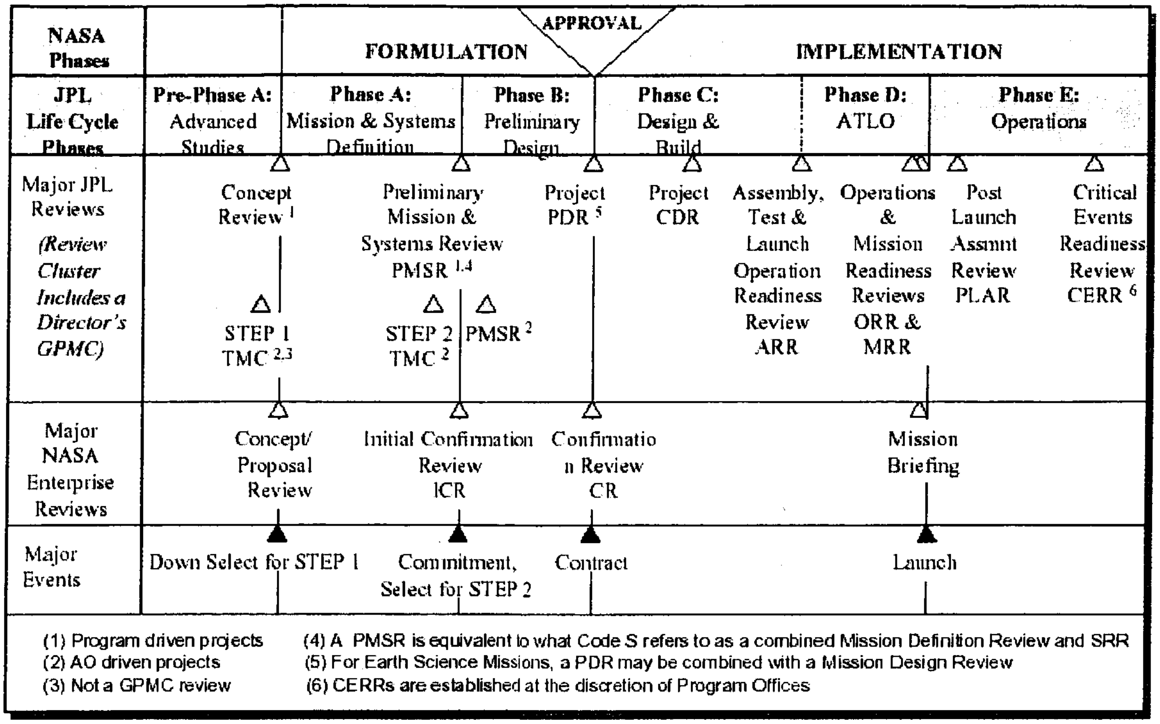

Figure 2: JPL LIFE CYCLE INCLUDING MAJOR REVIEWS

In the NASA life cycle, two key reviews are the PDR and CDR. The **Preliminary Design Review (PDR)** is the first major review of the detailed design and is normally held prior to the preparation of formal design drawings. PDR's are conducted to confirm that the approach for the system's design is ready to proceed into the detailed design phase. A PDR is held when the design is advanced sufficiently to begin some testing and fabrication of design models. Detail designs are not expected at this time, but system engineering, resource allocations and design analyses are required to demonstrate compliance with requirements.

The **Critical Design Review (CDR)** is held near the completion of an engineering model, if applicable, or the end of the breadboard development stage. This should be prior to any design freeze and before any significant fabrication activity begins. The CDR should represent a complete and comprehensive presentation of the entire design. CDR's are conducted to demonstrate that the detailed design is complete and ready to proceed with coding, fabrication, assembly and integration efforts.

While standing reviews like the PDR and CDR are highly structured and formalized, the technical peer reviews that are an important pre-review for these programmatic reviews are not. It is in these informal reviews that the engineers and managers must work out the details that can be missed in formal reviews. The key is success in these peer reviews is a balance between discipline and freedom.

1.4. History

Virtually every organization has had examples of design reviews which failed to detect errors in design. As a manager at GE Aircraft Engines once said, "design reviews are like 100% inspection with an imperfect gauge. One NASA study found that about 80% of post-launch problems and failures possibly could have been identified in design reviews. (Quinn 1994) This study recommended that a series of informal, but structured, detailed peer reviews be conducted prior to the Preliminary Design Review and, especially, the Critical Design Review.

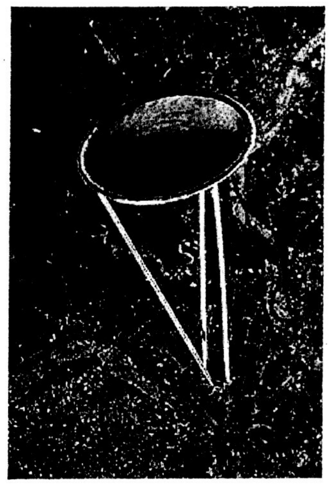

Figure 3: THE INFLATABLE ANTENNA

As one example of how design review procedures can allow errors, a project manager at JPL shared with us the lessons from a NASA JPL project which experimented with a new technology of an "inflatable" antenna. This inflatable antenna was 14 meters wide and mounted on three 28 meter struts. It would lay the groundwork for future technology development in space structures, which have the potential to be 10 to 100 times less expensive than conventional structures.

The antenna experienced unexpected dynamics during the initial ejection and inflation of the structure, but the correct final shape was attained. After full inflation of the antenna structure, the spacecraft began rotating unexpectedly. Because the technology used a new mylar, most engineers wanted to spend the time reviewing it, and only used what time was left to review the rest of the system. The box which would store the antenna only had one engineer working on it, who was not able find all the errors in it. Because the system lacked discipline in scheduling and allocating resources for a proper system review, the team was not able to identify the error.

2. PEER MEETINGS

2.1. Introduction

Before becoming a project manager, the PM of ST-6 never liked the review process and always felt he would conduct them differently. While the reviews did have the proper elements in terms of agenda and expertise, he didn't like the psychology of the reviews. He felt it was conducted like a trial with opposing prosecution and defense looking out for their interests and not the necessarily working together looking for weaknesses in the design. One side is trying to look smart by defending their work, while the other side is trying to look smart by attacking it. It's not truly a team that is assessing the situation.

In NASA, the engineers and managers are asked to have reviews very early and very often. The purpose of these reviews is largely to ensure that the team does not go in the wrong direction. However, these early reviews often occur before the team is really ready to lock in on ideas or numbers, but they feel the need to produce viewgraphs with their preliminary thoughts and sell it as a more mature idea. And once something is in writing and on a viewgraph, it gets notoriety. These estimates can be based on completely unrealistic assumptions. Costs are particularly dangerous as it is very hard to increase that amount, even if the initial guess is much too low.

Formal reviews can even act as a barricade. The key to doing the reviews early is to keep them informal and use peer reviews rather than the formal, mandated reviews. These "**peer meetings**" need to be conducted in the right way. Both types of reviews are necessary. It is important though to structure them so they complement each other.

2.2. Background

The New Millennium Program (NMP) is a technology program at NASA Jet Propulsion Laboratory in Pasadena, California jointly funded by the NASA Office of Space Science and the Office of Earth Science. It was established in 1995 with the goal of speeding space exploration through the development of highly advanced technologies. The purpose of the program is to conduct testing of breakthrough technologies in the space environment to minimize risk of first use. The Space Technology 6 (ST-6) Project is developing technologies which will improve a spacecraft's ability to autonomously make intelligent decisions on what information to gather and send back to the ground as well as determine its attitude and adjust its aim. The goal of the program is to develop unmanned missions that no longer require (delayed) control by ground crews.

Figure 4: ARTIST'S CONCEPTION OF THE INERTIAL STELLAR COMPASS AND THE AUTONOMOUS SCIENCECRAFT EXPERIMENT

In 2004 and 2006, the NMP's ST-6 will try two experimental technologies. The Autonomous Sciencecraft Experiment (Sciencecraft) and the Inertial Stellar Compass (Compass). Sciencecraft will enable a spacecraft to decide what science observations to make, then process, and return data on its own. Compass will enable a spacecraft to continuously sense its position and recover after temporary malfunctions or power loss.

In the "world series" of reviews done so far for the $25 million New Millennium ST-6 program, there have been:

- 8.2 kg of documents
- 117 reviewers
- 14 reviews
- cost of reviewers was $208K to NASA
- 522 items of ISO compliance
- 321 design principles
- 21 versions of requirements documents based on different viewers
- 1325 charts produced

In this program, the PM has instituted his program of "**engineering peer meetings.**"

2.3. Method

It is important to place a priority on peer reviews and plan for them from the very beginning. Though ideally, peer reviews would be done on every subsystem and component in the greatest of detail with all the foremost experts, that is not always possible. Likely, the team needs to do some sort of review on every subsystem but spend different amounts of time and resources on varying subsystems. For example, some programs may be more hardware or software oriented. It is up to the project manager and the subsystem leader to look at areas where there are problems. Software can often have overruns, and in addition, it is often difficult to do a formal review on it. As such, peer reviews are great for software. Other times when peer reviews are essential are when the team lacks experience with the type of project or technology.

Before the peer review begins, it is important to identify what is needed out of the review. Is it a "rubber stamp" process? Or should the nature be just to gather as many smart people together to share general comments? The review can try to understand the process identify things that can go wrong (*insight*) or just be a check to make sure the process has been done (*oversight*).

2.3.1. Scheduling

As pre-reviews, a good rule of thumb is to conduct the peer reviews about a week to 10 days before the formal review. This gives enough time to react to suggestions and criticisms. To do a review too early, even as much as a month, likely means the review of something that is not a true representation of the project by the time the formal review comes.

However, the planning of these peer discussions must start to take place well in advance. It can take up to six months to get a good handle of a medium size mission. It takes a month to figure out good questions to ask, a month to collect questions about similar subsystems in the past, a month to find out who to talk to, two months to conduct the interviews, and another month to put the data into something that's reasonable and use probabilistic risk tools.

Nonetheless, the review process is a continuous process. Each review likely uncovers new problems and continually changes the risk posture of the project and future reviews may be needed.

2.3.2. Personnel

Peer reviews are extremely reliant on the skills of both the presenters and the reviewers. Not only must they be technically sound to understand the issues, but they also need the verbal and communicative skills to explain the background and analysis.

It is up to the program manager to decide who to include in reviews. Currently there is no organized system or list of reviewers in place for managers to refer to.

The process is quite informal but intuitive in many ways. The main constraint is to choose people who are not working on the same project. For example, if the subsystem is in electronics, reviewers should be people who are in electronics, like former cognizant engineers. When experts are needed from other areas, often the best place to start is a manager of that section or line organization.

Peer review discussions are usually driven by the design team presenting areas they are not comfortable or confident with. Whenever possible, it is important for the reviewees to recognize the issues and invite reviewers who are familiar and experienced in this area. For example, if the issue is in a field programmable gate array (FPGA), it would be ideal to find someone who has been working with it for many years now. Though it sounds obvious, this requires thought and research by the design team to identify experienced personnel.

In addition, it is better if the project manager is not present. This allows a more relaxed atmosphere. The reviewers can simply chat with the project manager informally to update on the process. The value of peer reviews is not passing or failing them. It is to prepare for the formal review. The PM has said that some of his best formal reviews have come after failing a several peer reviews.

2.3.3. Format

It is important to change the way people come into these peer reviews. The format is important, and should emphasize that these aren't official, formal peer reviews. It is better to just have them in an engineer's office. If it's in a conference room, not only does it make it difficult to find and look up material, but it also affects the psychology and the formality of the exchange.

Even the sitting arrangements have psychological impact in peer reviews. By putting the review in a round table format and not having people stand-up and present the material, it prevents the exchange from being too formal or adversarial with a prosecution-defense mentality. The PM recommends having only three, perhaps four, members at a time in peer meetings. From his experience, the best arrangements include one or two team members with two peer reviewers. If there are 3 or more reviewers, then not only does it limit the dialogue, but it is necessary to make copies of the papers and makes the process more formal. With only a few people, the team can just look over the same sheet of paper and have a more intimate dialogue.

One of the things that makes peer reviews under this PM's projects unique is the **brainstorming session** that starts each peer meeting. He aims to make the experience different in many psychological aspects. The peer reviews begin with the participants in the room just chatting about the problem and the issues for a while without writing anything down. This allows the conversation to be more free-flowing and also lets the participants to be more relaxed about talking about issues without worrying about details that they will be held accountable for later. The PM has actually tested this technique on different subsystems and found that teams often stay with numbers even if they are not necessarily accurate. When the reviews required the teams to write down numbers, the group often "forced" the numbers and they evolved only 30%, while with teams that weren't allowed to write things down, their numbers evolved by as much as 400%. When talking about costs, it is important to ask open instead of leading questions. The initial cost conservations should be one-on-one, and should start with discussions on the maximum cost figure before the nominal or minimum.

2.3.4. Risks and Worry Generators

The key to the PM's peer reviews are a list of "**worry generators**." These are templates of types of questions that reviewers should be concerned about, and are generated from historical and lessons learned sites. Most reviewers don't prepare anything before their reviews. Because of the formality of the mandated activities for reviews in addition to the stigma of failure at any organization, it is often difficult to find documentation which talks about failures. Euphemistic language that doesn't help discuss how to catch errors is often prevalent in reports.

This PM's worry generators were inspired by his participation in a number of design reviews. He found that certain experienced reviewers consistently asked questions in certain areas. By putting them down and formalizing them, and then talking to the review office and managers at different departments, he was able to collect a list of general questions to start any review. There are specific individuals in the organization which can help with certain aspects of the worry generators, such as cost reviewers and ISO reviewers. There are worry generators for each subsystem, with separate ones for software and hardware. However, before each review, it is up to the project manager or subsection leader to review relevant lessons learned and past failure history documents to find specific relevant areas that should be discussed in the peer reviews. For example, power generator issues are likely common for different missions.

Worry generators also include topics that engineers don't like to talk about, such as costs and non-technical issues. Engineers often feel above programmatic issues and wouldn't talk about such topics in their presentation. Because issues aren't technical doesn't mean they aren't relevant. It is necessary to worry about even environmental issues, like the allocations needed in reserve if a contractor on the east coast is snowed out. The worry generators are to be used to carry forward discussion into the group brainstorming. Examples of worry generators include:

- *What can go wrong? What would be the source(s) of the problem? What are the chances of that happening? How certain am I of that?*
- *Do I really know what stresses this process and how this process will respond over its full range of loading?*

- *What are the triggers that might break the process? What signals that it is broke, when, and to whom?*
- *If I were trying to make this process/system/step NOT work, what things could I do?*
- *What things in the environment outside of my area of focus do I not know about fully and which of them might interact with my area? What would be the results? How would I know?*
- *How ready is the technology to be used? How do I know that?*
- *For how long will the technology be usable? Why? What limits the time? How many times will I have to turn the technology over in a 10-15 year period?*
- *Are there any dicey design issues I can foresee? Any test and integration challenges?*
- *How much is the human part of the process affected? How much training, re-staffing, overcoming of resistance will there be: Why? How certain am I?*
- *Can every thing needed be manufactured, delivered, tested, and made operational in time? What happens if not? How likely is that?*
- *Are there any things that could be unsafe or hazardous? How? Under what conditions?*
- *What could one do about each of these things to:*
  - *avoid the risk (e.g., change requirements, other options)?*
  - *transfer the risk (sell tasks, insure)?*
  - *abate the consequences (reduce the impact; provide back up; insulate/isolate)?*
  - *reduce the likelihood of occurrence?*
  - *keep track of the situation, watching for triggers, leading indicators?*
- *Is this big enough to do something about, keep track of and worry about, or accept and ignore?*
- *How could we get more information to narrow the uncertainty of what we know about these potential risks?*

The PM also has a well-defined and established process for dealing with risks. Often engineers find the whole process a waste of time. When it comes to risk management and cost estimation, he hires a person full-time or half-time to perform collect data and enter it into "risks lists" for the whole project. This individual goes from office to office prodding people for data. He or she polls engineers to see what the maximum and nominal costs for each risk are and assigns probabilities. Using historical data and systems like MAIS (Mishaps and Anomaly Information System), the analysis is done not only for the likelihood of the risks occurring but also the money required to mitigate them.

2.3.5. System Review

Peer reviews often don't discuss system interactions as the system review takes place in the formal reviews. Usually these formal reviews have the leads of the peer reviews come and report. The PM tends not to discuss system interface issues in the peer reviews because the problem becomes much too big and costly for a peer review to handle. The topic can dominate the conversation and prevent the group from talking about the subsystem's issues. He does advocate having a separate system peer review where members from the subsystem peer reviews attend. The participants report on the problems of current subsystems and discuss the system implications. There are also interface review documents which can aid this meeting. In addition, system issues are discussed at the weekly meetings with all representatives from the subsystems. Because the project the PM is working on is a new technology project, he does not have the requirements in configuration management that some larger projects do or in using the past failure report databases. However, these are other resources which can aid the system review.

3. CONCLUSIONS

3.1. Case observations

Different projects require different review practices. Larger projects have some advantages in terms of having more experienced people, as well as mentors available for the inexperienced. Smaller projects do not have this luxury – the projects are the training grounds for new engineers. In addition, larger projects are by definition required to have a certain rigor, like with configuration control. As a result, there is also less flexibility in larger projects. Because the large projects have more resources, the quality of both the engineers and reviewers are also usually higher on larger projects. Because the formal reviews are more rigorous, these projects can benefit from peer reviews as a good pre-work and a pre-review.

The design review process is a weak process which depends on the individuals involved, both the manager and engineer, reviewer and reviewee. Still, many at NASA do not recommend more requirements and formalism necessarily as a way to improve the process. Because design processes can vary so greatly in terms of not only size and technology, but even in domain and technology, it is not feasible to create a universal checklist for all reviews. There is an inherent tradeoff between the specificity of items in the list with the size and burden of performing such a review. The key to improving the design review process is to treat it as an activity of insight and not just oversight.

This PM believes the best way to improve the peer review process is to target project managers and create a guide on how to conduct reviews, beginning with a background on the psychology of reviews. That can be followed with a training class, probably no longer than 2 hours. The first half-hour session can be on the current state of design reviews, a second on peer review techniques, and a third would be a retroactive example. General guides like these would likely be accepted at both the upper management and project management levels.

3.2. Error-proofing context

The PM believes design reviews are a dangerous tool. They can give the users a false sense of security. For most organizations, design reviews are the only line of defense against errors in the design process. Even if they are applied universally, they are still an imperfect gauge susceptible to human errors and will always miss some problems. Nonetheless, they will always be an essential part of any organization's efforts to error-proofing the development process. The keys to using design review as a part of the design process error-proofing toolkit are to recognize their inherent weaknesses.

By considering the solution and robustness level of design reviews, shown in Figure 6, one can recognize what they can and can not do. Like the quote said, design reviews are an "inspection" activity of level III robustness. Even if they identify an error, there are detection and rework delays involved. Ideally, the error should be caught as soon as possible if not prevented altogether, and not merely at the end of the phase. At best, design reviews are a process level (level 2) fix, but often times they aim to identify and fix specific problems reactively. However, improvements to design review organization can impact the system and not just a particular problem or process.

Level 3: *Fix the system*
Level 2: *Fix the process*
Level 1: *Fix the specific problem, very reactive*
Level 0: *Denial phase, rationalize without fixing*

V. Prevention: *Eliminate the possibility of performing an erring action*
IV. Detection: *Detect the error immediately after being made*
III. Inspection: *Source inspection of completed parts*
II. Improvement: *Improvement to simplify or guide the process*
I. Tool: *General tool to aid analysis and design*

Figure 5: SOLUTION AND ROBUSTNESS LEVELS TO CONSIDER FOR ERROR-PROOFING

Unlike other error-proofing solutions, design reviews are a system that is, for the most part, already in place at most organizations. Improving the process does not require large capital investments in technology or even a change in the process necessarily. Design reviews can be impacted immediately. Design reviews can be focused on the key error factors through guidelines, checklists, and training. The key is to make these reviews as robust and consistent as possible and supplement them with additional efforts. Our research has found that the project manager is the key to implementing a strong design review and error-proofing process at the project level. That can only be done with good pre-work, gathering both strong reviewers and reviewees, and being committed to reviewing consistently.

ACKNOWLEDGEMENTS

We sincerely appreciate the assistance of Artur Chmielewski with ST-6 at the Jet Propulsion Laboratory for his time and sharing his knowledge and experience of engineering peer meetings.

Thanks also to Chester Borden for arranging this case study, David Bell for his assistance in this interview, and Peter Putz for his comments on the report. Special thanks to Irem Tumer for her patient organization of this research and insightful comments about the whole project.

REFERENCES

Chao, L., and Ishii, K. 2003. "Design Process Error-Proofing: Failure Modes and Effects Analysis of the Design Process," *Proceedings of the ASME DETC: DFM,* Chicago, IL.

Chao, L., and Ishii, K. 2003. "Design Process Error-Proofing: Development of Automated Error-Proofing Information Systems," *Proceedings of the ASME DETC: DAC*, Chicago, IL.

Chao, L., Beiter, K., and Ishii, K., 2001, "Design Process Error-Proofing: International Industry Survey and Research Roadmap," *Proceedings of the ASME DETC: DFM*, Pittsburgh, PA.

Quinn, James. 1994. "Flight P/FR's and the Design Review Process." *JPL D-11381.*

WEB REFERENCES

Error-Proofing Case Studies
http://mml.stanford.edu/Research/DPEP/ep.casestudy.html

Goddard System Review Office - Design Review Guidelines
http://arioch.gsfc.nasa.gov/301/html/design.html

Jet Propulsion Laboratory
http://www.jpl.nasa.gov/

New Millennium Project
http://nmp.jpl.nasa.gov/program/program.html

New Millennium ST-6
http://nmp.jpl.nasa.gov/st6/ABOUT/index.html