TRANSITIONING TECHNOLOGY FROM R&D TO PRODUCTION

by

Seward W. Pulitzer III

S.M. Mechanical Engineering, Massachusetts Institute of Technology, 1998
B.S. Mechanical Engineering, Rensselaer Polytechnic Institute, 1996

Submitted to the MIT Sloan School of Management and the Department of Electrical Engineering and Computer Science in partial fulfillment of the requirements for the degrees of

Master of Business Administration
and
Master of Science in Electrical Engineering and Computer Science

In conjunction with the Leaders for Manufacturing Program at the Massachusetts Institute of Technology, June 2008

Signature of Author _____
MIT Sloan School of Management
Department of Electrical Engineering and Computer Science
May 9, 2008

Certified by _____
Duane Boning, Thesis Advisor
Professor and Associate Department Head
Department of Electrical Engineering and Computer Science

Certified by _____
Roy Welsch, Thesis Advisor
Professor of Statistics and Management Science
Center for Computational Research in Economics and Mgmt Science

Accepted by _____
Terry P. Orlando
Chair, Department Committee on Graduate Students
Department of Electrical Engineering and Computer Science

Accepted by _____
Debbie Berechman
Executive Director of Masters Program
MIT Sloan School of Management

TRANSITIONING TECHNOLOGY FROM R&D TO PRODUCTION

by

Seward W. Pulitzer III

Submitted to the MIT Sloan School of Management and the Department of Electrical Engineering and Computer Science on May 9, 2008 in partial fulfillment of the requirements for the degrees of Master of Business Administration and Master of Science in Electrical Engineering and Computer Science.

Abstract

Corporate research and development (R&D) drives progress in the high-tech industries. Companies that advance the state-of-the-art in product performance enjoy significant advantages over the competition. However, although technical achievement may be required for competitiveness, it is far from sufficient. The successful creation of a prototype is in no way an indication that large quantities of identical units can be economically or reliably produced. Indeed, transitioning a new technology from the laboratory in which it was created to a production environment can be as challenging as the actual development.

Many of the obstacles on the path to production stem from process variability. The randomness inherent in every manufacturing operation introduces risk into the transition process. As products increasingly become more technologically complex, the need for close coupling between R&D and production groups also grows. However, current trends toward distributed development and outsourced manufacturing work to segregate these groups, increasing the chances that risk factors will be overlooked.

This thesis examines the process by which new technologies are transitioned from the laboratory to the shop floor, and addresses some of the risks related to process-uncertainty that arise. In particular, it focuses on those challenges that are difficult to recognize. Research for the project was conducted during a six month internship at the Raytheon Company. The transition-to-production process was segmented into three phases, and three related products were used as cases with which each phase could be studied. The problems that appeared in these cases were addressed, and the "lessons learned" were then generalized into a set of guidelines applicable to a broader range of situations.

Thesis Supervisor: Duane Boning
Title: Professor and Associate Department Head
Department of Electrical Engineering and Computer Science

Thesis Supervisor: Roy Welsch
Title: Professor of Statistics and Management Science
Center for Computational Research in Economics and Mgmt Science

Acknowledgments

I would like to thank the Leaders for Manufacturing Program and its director, Don Rosenfield. I would also like to thank my thesis advisors, Prof. Roy Welsch from the Sloan School of Management, and Prof. Duane Boning from the Department of Electrical Engineering and Computer Science, for their time, assistance, and patience during the work leading to this thesis.

Additionally, I would like to thank Roger Hinman, John Day, and Dave Devietro for their guidance, support, and friendship during my internship at Raytheon. Their belief in the benefits of collaboration between academia and industry, and their tireless support of the LFM program in particular, is a model for all such relationships.

Finally, I want to thank my mother, my father, and my sisters, Brooke and Morgan. It is no exaggeration to say that, had it not been for their love, strength, and patience, this thesis would not exist. I owe each of them more than I can say. That said, if they ever let me go back to school again, I will find each of them, sit them down, and make them read every page of this document.

Table of Contents

List of Figures

List of Tables

Acronyms

ANOVA	Analysis of Variance
AOI	Automated Optical Inspection
ASIC	Application Specific Integrated Circuit
BGA	Ball-Grid-Array
BJT	Bipolar Junction Transistor
C&E	Corporate and Eliminations
C3I	Command, Control, Communication, and Intelligence
CBT	Cross Business Team
CCA	Circuit Card Assembly
CL	Center Line
CMMI®	Capability Maturity Model[1] Integration
CMOS	Complimentary Metal-Oxide Semiconductor
COTS	Commercial Off-The-Shelf
DDS*x*	Direct Digital Synthesizer (generation *x*)
DEI	Data Entry Interface
DFM	Design for Manufacturability
DoD	U.S. Department of Defense
DOE	Design Of Experiments
DREX	Digital Receiver/Exciter
DTICS	Defense Trusted Integrated Circuits Strategy
EEPROM	Electrically Erasable/Programmable Read-Only Memory
ESR	Electronically Steered Radar
EAV	Entity-Attribute-Value model
FCS	Frequency Control Solutions
FNC	Future Naval Capability
GUI	Graphical User Interface
HALT	Highly Accelerated Life Test
HBT	Heterojunction Bipolar Transistor
HiTCE	High Thermal Coefficient of Expansion
HTOL	High Temperature Operating Life
IBT	Integrated Business Team
IAD	Integrated Air Defense
IADC	Integrated Air Defense Center
IC	Integrated Circuit
IDS	Integrated Defense System
IIS	Intelligence and Information Systems
IO	International Operations
IPD	Integrated Product Development
IPDS	Integrated Product Development System
IRAD	Internal Research and Development
IT	Information Technology

[1] *CMMI* and *Capability Maturity Model* are registered trademarks of Carnegie Mellon University.

JBT	Joint Battlefield Integration
LCL	Lower Control Limit
LSL	Lower Specification Limit
MCM	Multi-Chip Module
MD	Missile Defense
MMT	Maritime Mission Systems
MPW	Multi-Project Wafers
MS	Missile Systems
MSL	Moisture Sensitivity Level
NCS	Network Centric Systems
NTSP	National & Theater Security Program
O&S	Operations and Support
PACI	Process Analysis and Control Interface
PAVE	Precision Acquisition Vehicle Entry
PAWS	Phased-Array Warning System
PVD	Physical Vapor Deposition
REQM	Requirements Management (CMMI Acronym)
RF	Radio Frequency
SBIR	Small Business Innovation Research
SiGe	Silicon-Germanium
SAN	System Assigned Number
SAS	Space and Airborne Systems
SFDM	Shop Floor Data Management [system]
SMT	Surface Mount Technology
SQL	Structured Query Language
SSC	Surveillance and Sensors Center
SVTAD	System Validation, Test, and Analysis Directorate (a.k.a. SVT&AD)
TAPO	Trusted Access Program Office
TFA	Trusted Foundry Access
TRS	Test Requirement Specification
TS	Technical Services
UCL	Upper Control Limit
USL	Upper Specification Limit
VBS	Virtual Business System
V&V	Verification and Validation
WIP	Work in process
WYSIWYG	What You See Is What You Get

1 | *Introduction*

Note: Much of the data in this thesis has been disguised to protect Raytheon confidentiality.

In most high-tech industries it is research and development (R&D) that takes center-stage. Exciting developments in the corporate research labs of innovative companies, both large and small, stimulate investment in new technologies and drive the state-of-the-art forward. Organizations, companies, and even governments formulate plans based not only on current technical capabilities, but on predictions of where the level of technology will be years in the future. With so much focus on R&D, laboratory achievements are often viewed as immediate corporate windfalls.

But technological achievement by no means ensures commercial success. As corporate R&D becomes increasingly sophisticated, the path from success in the lab to success in the marketplace becomes more difficult, complex, and uncertain. Many companies have been forced to learn the harsh lesson that transitioning a new technology from R&D to production can be equally as difficult as developing the technology in the first place.

The difficulty of this transition-to-production is typically compounded by the fact that the people, equipment, and organizational structures that favor research are often diametrically opposed to those required for effective production. Whereas R&D is typically associated with free-thinking academics, research equipment, and small tightly-knit groups, large-scale production tends to require a disciplined workforce, manufacturing equipment, and a great deal of communication across organizational boundaries.

In practice, the incompatible needs of R&D and production are reconciled through some sort of transition. In startup companies, particularly those centered on a single product, the opposing factors often cause the firm to transform itself once the technology is mature and ready for production. Typically accompanied by large infusions of capital, this metamorphosis from an R&D to a production company is a tumultuous event in the corporate lifecycle. For larger established companies, the reconciliation is effected by divorcing the R&D and production segments, then performing the transition as a handoff from one to the other. In this mode, research centers are tasked with bringing advanced technologies to some state of maturity, usually embodied as a functional design. Once this threshold is reached, the design is handed to a manufacturing group for production. This mode has the benefit of allowing the development and manufacturing centers to adopt whatever form works best for them, but is suffers from several drawbacks as well. First, in an era of complex and highly integrated systems, it is difficult, if not impossible, to decouple the functionality of the technology from the fabrication techniques with which it will be produced. And second, the handoff of information from R&D to manufacturing can be complicated and error-prone. The latter difficulty is often exacerbated by cultural differences between research and production groups and the differing beliefs and assumptions that go along with them.

1.1 Problem statement

One of the most important factors for success when transitioning a product from the laboratory to production is the recognition and mitigation of risk arising from process variability. Unfortunately, such risks are not always easy to spot when dealing with complex technical products. Too often, risk recognition and mitigation efforts are

narrowly focused on purely technical issues, such as the range in which a particular parameter is expected to shift during a manufacturing operation, and neglect other "soft" factors, such as the organizational structures in which the risk mitigation takes place or the nature of the interface between the R&D and manufacturing groups.

Process risk management in today's manufacturing environment is more difficult than ever. As technology advances and products become more complex, the coupling between R&D and production strengthens. But as the manufacturing environment is increasingly characterized by geographically-distributed development and outsourced manufacturing, it becomes harder to maintain an effective relationship.

The challenge undertaken in this thesis is to examine the process whereby new technologies are transitioned from the lab to production, identify some of the risks related to process uncertainty, and formulate solutions to mitigate these risks. Specifically, the goal is to uncover *hidden risks*, those factors that impact the success of the transition-to-production process but that are difficult to spot.

1.2 Approach

Research for the work was performed, over the course of a six-month internship, at the Raytheon Company, an established corporation with distinct engineering and manufacturing groups.

To characterize the transition-to-production process, three phases of evolution were defined:

Phase I. A new technology is emerging from an R&D center and has been tapped for use in a product. The design is in the final development stages but questions may still exist regarding its achievable performance. The technology may be in the process of being qualified for use in a particular application. The emphasis is still on development.

Phase II. The functionality of the technology has been demonstrated, and the design may be undergoing modifications to make it more practical and/or

economical to manufacture. The emphasis is shifting from development to production.

Phase III. The technology is in some level of production. R&D work is considered to be complete, and the emphasis is entirely on production.

It is important to recognize that the definitions of these phases are somewhat arbitrary. Rarely is the path from development to commercial production a simple one, and multiple technological, manufacturing, and business activities often overlap. Nevertheless, they provide a useful framework for the investigation of the transition process.

Since the development cycles of advanced technological products can span years, yet the duration of the internship was only six months, an approach is adopted wherein three products, each at a different stage in their transition, are chosen to serve as case studies. By studying these cases, the challenges and risks present in the product's phase of transition, specifically those dealing with the identification and control of variability, can be examined. Each of the subsystems chosen is part of a single functional assembly: a circuit board used for signal amplification in radar systems. This helps to maintain consistency between the cases with regard to the underlying technology and the environment in which they were developed. The three components, each of which is discussed in detail in subsequent chapters, are introduced in Table 1-1 below, and their functional relationships are illustrated in Figure 1-1. The DDSx, a chip used for synthesizing analog electrical signals, is a Phase I component. The assembled circuit board, the CCA, is used for the study of Phase II. Finally, the SAW device, an oscillator employed to generate a high-accuracy clock signal, is the Phase III focus.

Table 1-1. Subsystems for case-studies

Device	Function	Transition Phase
DDSx	Signal Synthesizer	I
SAW	Oscillator	II
CCA	Circuit Board	III

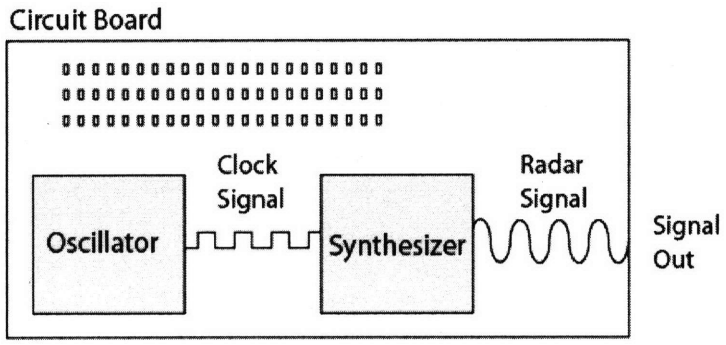

Figure 1-1. Conceptual relationship among case-study subsystems[1]

Risk related to variability exists at all three stages of the transition process, but the focus of the efforts to mitigate it changes as the process matures. The focus of each stage is summarized in Table 1-2. In Phase I, when the technology itself is still in a state of flux, potential sources of variability may not known. Risk mitigation activities in this phase therefore focus primarily on identifying potential sources of randomness that may exist. In Phase II, as the design becomes more finalized, efforts shift to find ways to eliminate those factors, or else arrive at a design that is robust to them. Finally, in Phase III, after all possible steps have been taken to minimize sources of variability at the design level, the emphasis shifts to process control, i.e. minimizing randomness in the fabrication processes used to produce the technology. Although each phase is most strongly associated with a single focus area, there can be considerable overlap; a typical pattern is illustrated in Figure 1-2.

[1] Note that this is figure is a *conceptual simplification* of the actual assembly, presented here to illustrate the general roles of each of the components. It should not be interpreted as representative of the physical system.

Table 1-2. Focus areas for risk mitigation

Phase	Focus Area
I	Variation awareness
II	Design for manufacturability
III	Process control

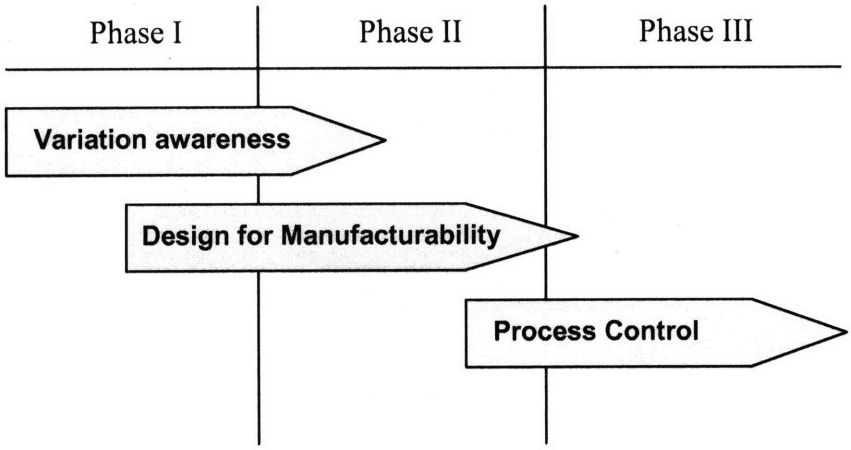

Figure 1-2. Variation control focus throughout the transition phases

1.3 Thesis layout

In Chapter 1, a high level overview of the objective of the internship and the approach taken is described. It defines the transition-to-production process in terms of three phases, and describes the changing areas of focus for managing variability throughout the transition.

In Chapter 2, business perspectives on the defense industry are presented to provide context for the work performed. Factors making the sector unique, such as the nature of government contracts, barriers to entry, and national security requirements, are discussed so the reader may understand the environment in which the work presented takes place.

In Chapter 3, a brief history of the Raytheon Company is given, along with an analysis of its organization and culture. An introduction to phased-array radar technology is provided, as is a description of Raytheon's product development process.

In Chapter 4, the role of variation control in Phase I of the transition, *variation awareness*, is examined using the selection, development, and qualification of a digital signal synthesizer chip, the DDSx, as a case study. The focus of the study is on how not only engineering factors, but also those stemming from organizational structure and communication, can contribute to variation risk.

In Chapter 5, the Phase II *design for manufacturability* approach is explored. During the internship, a statistical experiment was designed and executed using DOE[1] techniques to characterize a circuit board soldering process, identify design factors suspected of influencing the output (or screen out those that did not), and predict yields. The subsystem used for this case is the soldering of components on the CCA (circuit-card assembly).

In Chapter 6, the Phase III processes that are suited to prototyping or low-volume production must be replaced or adapted for use in higher volume manufacturing. This chapter discusses the creation of a statistical process control system for the purpose of collecting data on high variability processes and bringing them under control. The subject of the study is the manufacturing of a SAW oscillator device. The thrust of the chapter is on how an effective yet flexible system can be implemented.

In Chapter 7, the conclusions reached in this work are summarized.

Finally, the Appendices include additional information pertaining to the work discussed. Appendix A describes the Test PCB that was used for the DOE experiment in Chapter 5, Appendix B contains the data resulting from experiment, and Appendix C includes the results of the statistical analyses.

---

[1] DOE: Design of Experiments

2 | *Business Perspective*

The manner in which new technologies are transitioned to production can strongly depend on the characteristics and requirements of the industry. This chapter briefly describes the attributes of the defense industry in which Raytheon operates.

2.1 Defense industry overview

The dynamics of the American defense industry are different from those of other commercial sectors. The most salient difference is the dominance of U.S. government contracts in the marketplace. Restrictions arising from national security concerns, profit structure, and contract awarding criteria are all unique to the industry.

Since the ability of a company to be successful in the defense industry largely depends on its ability to win government contracts, it is critical for contractors to be able to demonstrate their engineering and manufacturing capabilities. Contracts are typically awarded to those companies that have proven themselves as "knowledge leaders" [Padgalskas 2007]. While past performance is an obvious method for demonstrating competency, another is compliance with government recommended practices and submission to third-party evaluations.

One such evaluation is the Shingo Prize for Operation Excellence. Established in 1988, the purpose of the award is to promote awareness of lean manufacturing concepts.

Each year, examiners from the Shingo Prize Academy conduct on-site evaluations of business facilities. Companies are evaluated based on the following criteria[1]:

- Leadership Culture and Infrastructure
- Manufacturing Strategies and System Integration
- Non-Manufacturing Support Functions
- Quality, Cost and Delivery
- Customer Satisfaction and Profitability

Although the Shingo Prize is not unique to defense, high scores constitute a valuable credential when vying for government contracts and are therefore highly sought after in the industry.

2.1.1 Contracts

The nature of the government contracts existing in the American defense industry is well described in Raytheon's 2006 annual report:

> U.S. government contracts include both cost reimbursement and fixed price contracts. Cost reimbursement contracts provide for the reimbursement of allowable costs plus the payment of a fee. These contracts fall into three basic types: (i) cost plus fixed fee contracts which provide for the payment of a fixed fee irrespective of the final cost of performance, (ii) cost plus incentive fee contracts which provide for increases or decreases in the fee, within specified limits, based upon actual results as compared to contractual targets relating to such factors as cost, performance and delivery schedule, and (iii) cost plus award fee contracts which provide for the payment of an award fee determined at the discretion of the customer based upon the performance of the contractor against pre-established criteria. Under cost reimbursement type contracts, the contractor is reimbursed periodically for allowable costs and is paid a portion of the fee based on contract progress. Some costs incident to performing contracts have been made partially or wholly unallowable for reimbursement by statute, FAR or other regulation. Examples of such costs include charitable contributions, certain merger and acquisition costs, lobbying costs and certain litigation defense costs.

---

[1] From the Shingo Academy: http://www.shingoprize.org (accessed 7 December 2007).

Fixed-price contracts are either firm fixed-price contracts or fixed-price incentive contracts. Under firm fixed-price contracts, the contractor agrees to perform a specific scope of work for a fixed price and as a result, benefits from cost savings and carries the burden of cost overruns. Under fixed-price incentive contracts, the contractor shares with the government savings accrued from contracts performed for less than target costs and costs incurred in excess of targets up to a negotiated ceiling price (which is higher than the target cost) and carries the entire burden of costs exceeding the negotiated ceiling price. Accordingly, under such incentive contracts, the contractor's profit may also be adjusted up or down depending upon whether specified performance objectives are met. Under firm fixed-price and fixed-price incentive type contracts, the contractor usually receives either milestone payments equaling up to 90% of the contract price or monthly progress payments from the government generally in amounts equaling 80% of costs incurred under government contracts. The remaining amount, including profits or incentive fees, is billed upon delivery and acceptance of end items under the contract.

2.1.2 Barriers to Entry

Although many small companies do business with the Department of Defense through SBIR[1] programs and others, the barriers to entry for firms desiring to compete with the top contractors are significant. The products are usually technologically sophisticated and highly complex, and require massive capital resources and logistical support. In addition, a long track record of success is typically needed to demonstrate capability in the contract proposal process. Exceptions do exist, however. Founded in 1997, L-3 Communications experienced a rapid rise to the number eight spot[2] in just nine years.

2.2 The Market

The term *defense industry* is typically used to refer to all manufacturers that produce equipment and supplies for the militaries of the world. Although the companies comprising the industry are usually commercial (i.e. not government agencies), the vast majority of revenue comes from government spending. As a result, the fortunes of

---

[1] SBIR: Small Business Innovation Research

[2] Based on annual revenues from defense contracts.

defense firms depend not only conventional market factors but also on the global geopolitical landscape.

In 2007, the United States was responsible for nearly half of worldwide military spending. Not surprisingly, the size of the global arms market is closely tied to the Pentagon's budget which, in fiscal 2007, was almost $450 billion (see Table 2-1). At the 2006 Vision Conference of the Government Electronics and Information Technology Association (GEIA), the organization revealed that although it expects U.S. defense spending to grow to $609.4 billion over the next decade, the real buying power of the budget is expected to fall 16.3% due to inflation [Keller 2006].

Table 2-1. U.S. Dept. of Defense Budget[1]

Department	FY 2005	FY 2006	FY 2007
Dept. of the Army	167,261	132,019	111,712
Dept. of the Navy	133,560	132,492	127,322
Dept. of the Air Force	131,673	128,895	130,386
Defense-wide	69,981	76,837	70,114
Total	502,476	470,242	439,534

2.3 The Competition

Raytheon competes with the world's largest defense contractors; Table 2-2 lists the top ten, by sales, in 2006. "Technical superiority, reputation, price, past performance, delivery schedules, financing, and reliability are among the principal competitive factors considered by customers in these markets" [Raytheon 2006 Annual Report]. The firms are large, usually domestic, and most focus heavily on defense. Of the five largest competitors, with sales ranging from $18 to $36 billion, only Boeing generated less than 75% of its revenue from military sales.

Although they are constantly vying for market share, the magnitudes of some government contracts often cause them to partner. For instance, on 29 April 2002, the

[1] U.S. Department of Defense Budget for FY 2007.

Navy awarded the $2.9 billion, four-year DD(X)[1] contract to a team fielded by Northrop Grumman and Raytheon; the losing team of General Dynamics and Lockheed Martin was then brought onboard as subcontractors. Such cooperation in the industry is necessary and common.

Table 2-2. The top 10 defense contractors worldwide in 2006[2].

Rank	Name	Country	2005 Rank	Defense revenue	% of total revenue
1	Lockheed Martin	U.S.	1	36,465	98.0
2	Boeing	U.S.	2	30,791	56.1
3	Northrop Grumman	U.S.	3	23,332	76.0
4	BAE Systems	U.K.	4	27,500[3]	79.0
5	Raytheon	U.S.	5	18,200	83.1
6	General Dynamics	U.S.	6	16,570	78.0
7	EADS	Germany/France[4]	7	9,120	22.5
8	L-3 Communications	U.S.	13	8,549	90.5
9	Thales	France	9	8,523	70.0
10	Halliburton	U.S.	10	7,552	36.0

[1] The contract, known as DD(X) design agent, is for development of a new family of ships that can operate with smaller crews and will use technologies to avoid radar detection, share information and communicate more effectively.

[2] [*Military Information Technology* 2007]

[3] Estimated US dollar amount; actual was £13,765.

[4] Incorporated in the Netherlands.

3 | *Raytheon Company*

This chapter introduces the Raytheon Company and describes its business performance, organizational structure, and system for developing technological products. It also introduces the reader to phased-array radar systems, the components of which are discussed in subsequent chapters.

3.1 Overview

Raytheon Company (NYSE: RTN) is a defense contractor headquartered in Waltham, MA. With net sales of over $20 billion in 2006, Raytheon is one of the largest such businesses in the world [Raytheon 2006 Annual Report]. From page 31 of the 2006 annual report:

> [Raytheon] develops and provides technologically advanced, integrated products, services and solutions in [its] core defense markets: Sensing, including radars and radio-frequency systems and infrared and electro-optical sensors and systems, Command, Control, Communications, and Intelligence (C3I), including tactical communication, command and control and intelligence systems; Effects, including missiles, precision weapons and information operations; and Mission Support, including full life-cycle services and training.

Raytheon has approximately 80,000 employees worldwide, of which roughly 15% are unionized. Sales to customers outside the U.S. accounted for $3.7 billion, or 18% of 2006 revenue. Of this, $1.3 billion derived from foreign military sales[1] through the U.S. government.

The company is organized into six principal *business segments*:

1. Integrated Defense Systems (IDS)
2. Intelligence and Information Systems (IIS)
3. Missile Systems (MS)
4. Network Centric Systems (NCS)
5. Space and Airborne Systems (SAS)
6. Technical Services (TS)

Revenues are generated fairly equally across these segments (see Figure 3-1). The research discussed in this thesis was undertaken within IDS, which is the largest of the segments. The business provides complex system integration services and products for military and civilian markets, including long-range ballistic missile defense radars. In 2006, it had 13,400 employees and $4.2 billion in revenue (21% of the year's $20.3 billion in sales).

Drilling further down, IDS is subdivided into five *business areas*[2]:

1. Future Naval Capability (FNC)
2. International Operations (IO)
3. Joint Battlespace Integration (JBI)
4. Maritime Mission Systems (MMS)
5. National & Theater Security Programs (NTSP)

---

[1] The U.S. Department of Defense's Foreign Military Sales program facilitates the sale of defense equipment and military training to foreign governments. Rather than dealing directly with the defense contractor, the customer interfaces with the Defense Security Cooperation Agency.

[2] NTSP consolidated the prior business segments of Integrated Air Defense (IAD) and Missile Defense (MD).

The company's stock history is fairly stable. While increases in its share price have, for the most part, lagged those of its competitors (Figure 3-2), it has closely tracked the Dow Jones Defense Industry Index (Figure 3-3). Although it is subject to sometimes rapid changes in government defense spending, stable revenue streams from workhorse products such as the Patriot Missile System, plus the smoothing effects of long-term and large-scale contracts, help to buoy the firm through choppy periods.

Breakdown by business segment		
Integrated Defense Systems	IDS	4,220
Intelligence and Information Systems	IIS	2,560
Missile Systems	MS	4,503
Network Centric Systems	NCS	3,561
Space and Airborne Systems	SAS	4,319
Technical Services	TS	2,049
Other	Other	828
Corporate and Eliminations	C&E	(1,749)
Total		**20,291**

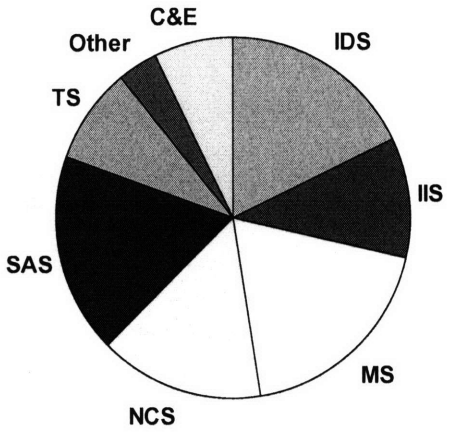

Figure 3-1. Raytheon 2006 revenue ($ millions)

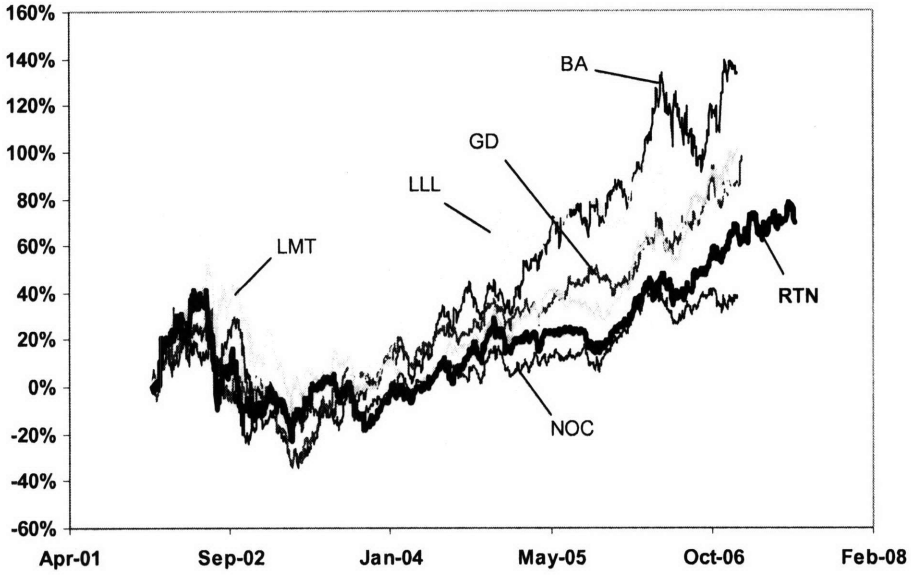

Figure 3-2. Raytheon stock price versus competitors[1]

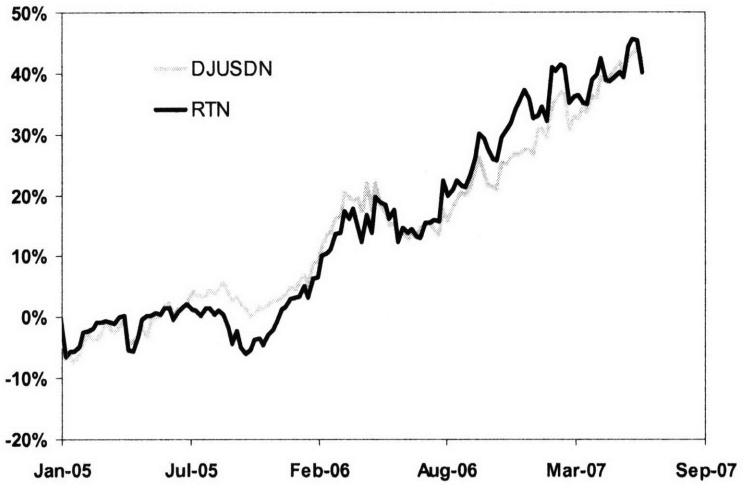

Figure 3-3. Raytheon stock price versus the Dow Jones Defense Industry Index

[1] LMT = Lockheed Martin; LLL = L-3 Communications; GD = General Dynamics; BA = BAE Systems; NOC = Northrop Grumman; RTN = Raytheon Company.

3.2 Organizational analysis

The organizational structure of the company was examined using the popular *Three Lens* management model developed at MIT. This model provides a framework in which subject organizations can be fully characterized in terms of their strategic, cultural, and political attributes.

3.2.1 Strategic organization

Raytheon is structured as a matrix organization. Within the IDS business areas are five Integrated Business Teams (IBTs). Each is aligned with a particular defense initiative, and is responsible for pursuing, winning, and executing contracts. Strategic, operational, and financial support is provided by 12 Cross-Business Teams (CBTs). The hierarchical organization of IDS is shown in Figure 3-4, and the relationships of the IBTs to the CBTs are shown in Figure 3-5.

Given the nature of the company's contracts and the large number of highly specialized technologies it provides, its organizational structure is appropriate. Although matrix organizations are often viewed as flawed, it is difficult to conceive of any viable alternatives, an argument bolstered by the fact that most of the major defense contractors utilize such a scheme to some degree. However, Raytheon does take matrix organization to an extreme. In some cases, cross-functional and cross-business responsibilities are defined down to the workgroup level, and no one person is accountable for a particular area. In these cases, it can be difficult to initiate change because authority is fragmented.

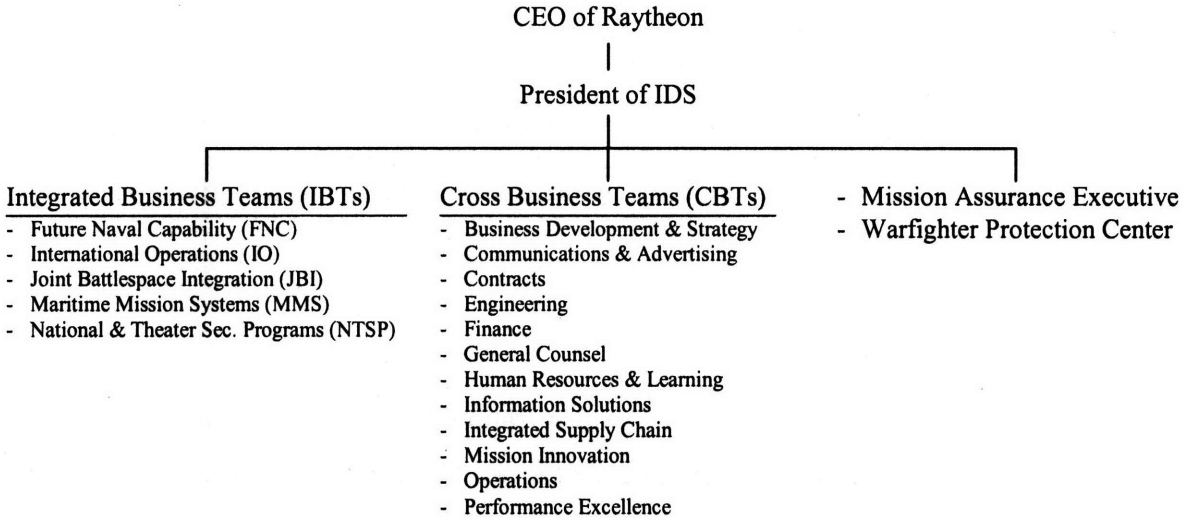

Figure 3-4. IDS organizational hierarchy.

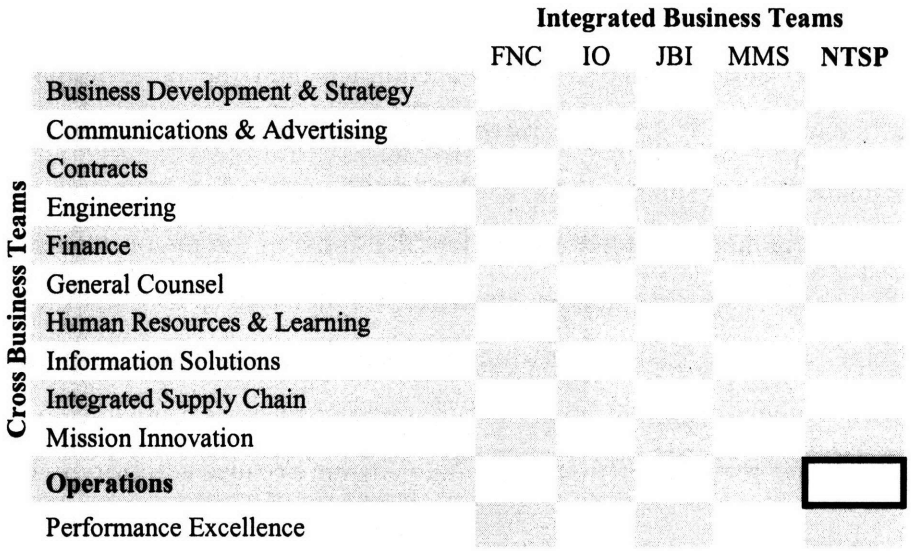

Figure 3-5. IDS matrix organization.
(The highlighted area shows the region in which the research was conducted.)

3.2.2 Company Culture

Workforce

One of the most salient characteristics of the Raytheon culture is the fact that many of its employees have been with the company for 20 or more years, a trait that is increasingly rare among American organizations, particularly those dealing with technology. And when speaking with younger members of the Raytheon workforce, it is

not unusual to find that at least one of the employee's parents work at the company as well. As a result, there is an unusually strong sense of shared identity between the company and its workforce, as well as a shared pride. Company sponsored functions such as "family fun days" in the summer are anticipated and well attended.

Although the makeup of the company's workforce is one of its main strengths, it is also one of its biggest problems. 75% of the Raytheon's workforce is 43 or older, and there is a large age gap between current employees and new hires. As more and more of the existing workforce retires, it will become increasingly difficult to fill their seats.

Table 3-1. Workforce demographics

Birth Year	Percent
1900-1945	10
1946-1964	65
1965-1980	20
1981-1999	5

Anonymous management

Also unusual in a company of this size is the lack of high-profile executives. Although many employees know who the CEO is (William Swanson in 2007), he does not possess the celebrity of more well-known corporate leaders such as Bill Gates or Jack Welch. Executives below the CEO remain relatively anonymous to all except those who report directly to them. Although this might be partially attributable to the nature of work in the defense industry, it seems more strongly linked to the company's culture. The blue-collar atmosphere of the shop floor extends, to an unusual degree, upward into middle- and even upper-management. A majority of managers were raised and educated locally. A common background and common experiences help to unify the workforce and cement the corporate identity.

3.2.3 Internal Politics

In general, given the size of the organization, politics plays a minimal role. Most of the political maneuvering within IDS revolves around program affiliation. Due to the

program-oriented focus of most teams and workgroups, decisions are sometimes made that benefit the program with little regard to how they impact the company as a whole. Not surprisingly, the political power bases that exist stem from these programs, with program funding serving as the source of that power. Large-budget programs have the most clout, and their leadership teams will usually win out when vying for resources. Although there is some logic behind this inherent prioritization system, it largely disregards other important factors such as profitability, scheduling, importance to future business, and others.

The main disadvantage of this system is that there is no motivation for program personnel to engage in long-term or company-wide improvement projects. As a program evolves, various technical, organizational, and logistical challenges are encountered. Sometimes, opportunities present themselves whereby a broad solution could be put into place that would ease the burden on future projects. Unfortunately, solutions such as these lie outside the scope of most projects, and few managers are willing to accept the financial responsibility. Long-term improvements that could benefit the company are, therefore, often passed over in favor of stop-gap solutions. This cycle perpetuates, and subsequent programs repeatedly face the same challenges.

3.3 Integrated Product Development System

Integrated product development (IPD) refers to the simultaneous development of a product, the preparation for its manufacture, and the performance of necessary marketing activities. The goal of IPD is to reduce the time-to-market, thereby reducing capital outlays and gaining an advantage over slower competitors. In an era of intense competition and rapid technological development, the importance of being first-to-market is well known. From a McKinsey report, "...products that come on to the market that have kept to budget, but are 6 months late, may lose 33% of profits during a 5-year period. Products which are 50% over budget but, on the other hand, have kept to schedule may lose 5% of profits over the same 5-year period" [Ottosson 1996]. Although these figures may vary considerably across companies and industries, the importance of rapid development and production is clear.

While there is general agreement in industry on the importance of IPD, the strategies employed to achieve it differ from company to company. Not surprisingly, whereas some companies have successfully incorporated IPD into their organizational framework, others have failed. Many hurdles exist for a firm attempting the transition to IPD, including finding an approach appropriate to the company and industry, modifying the organizational structure to align with IPD objectives, and maintaining clear communication across functional boundaries. One of the most common points of failure is a poor linkage between the R&D and business strategies. Kaminski contends that much of the problem stems from the vertical structure many companies utilize wherein groups are organized according to function (e.g. finance, engineering, marketing, etc.) [Kaminski 1994]. These vertical groups, often referred to as "silos," strive to maximize their performance as measured by group-specific metrics. The engineering group might, for instance, seek to increase production efficiency by designing and building automation equipment. Unfortunately, the drive toward functional excellence often comes at the expense of cross-functional performance. Extending the prior example, if it takes the engineering group an extra three months to automate the process, and that shifts the time-to-market by three months, then the loss in sales could potentially offset the efficiency gains. Such tradeoffs are the motivators behind IPD.

Raytheon's system for new product or process introduction and management is IPDS, the Integrated Product / Process Development System (conceptually illustrated in Figure 3-6). Quoting from the company's internal documentation, "IPDS is the Raytheon enterprise system that defines the standard organization processes used by all businesses to ensure program success." Raytheon describes IPDS as being:

- *A Raytheon company standard*

  IPDS serves as a common process baseline compatible with other process standards such as the Department of Defense's CMMI (Capability Maturity Model Integration). It is also the foundation for process improvement.

- *A knowledge sharing framework*

  The system integrates the contributions of all engineering and business disciplines

into a single lifecycle process. It collects common and local best practices, and provides the "enablers" (methods, tools, and training) needed by the programs.

- *A common program planning tool*

  IPDS is the key to the alignment of the process with the organization.

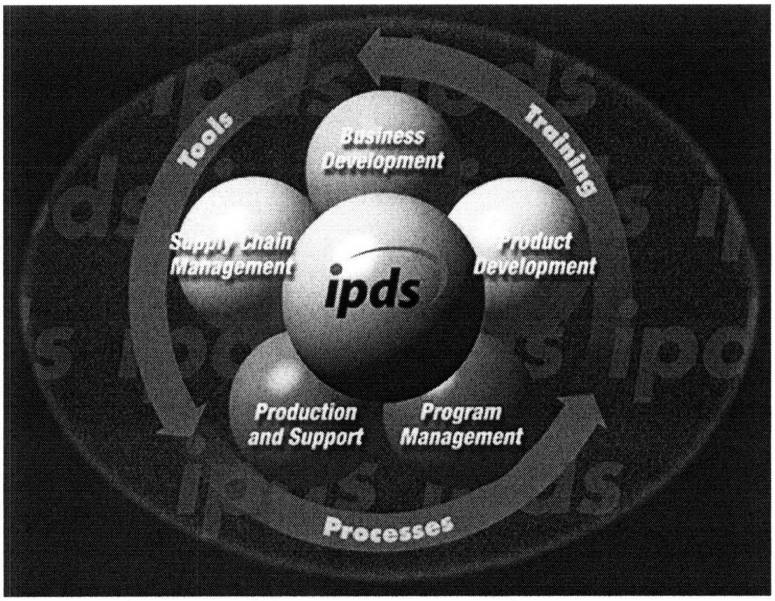

Figure 3-6. Raytheon IPDS integration

IPDS defines a lifecycle for new product development in terms of a series of seven stages, each comprising a set of tasks. The seven stages from business planning to final deployment and support, and the interrelations among them, are illustrated in Figure 3-7; the full list of supporting tasks follows.

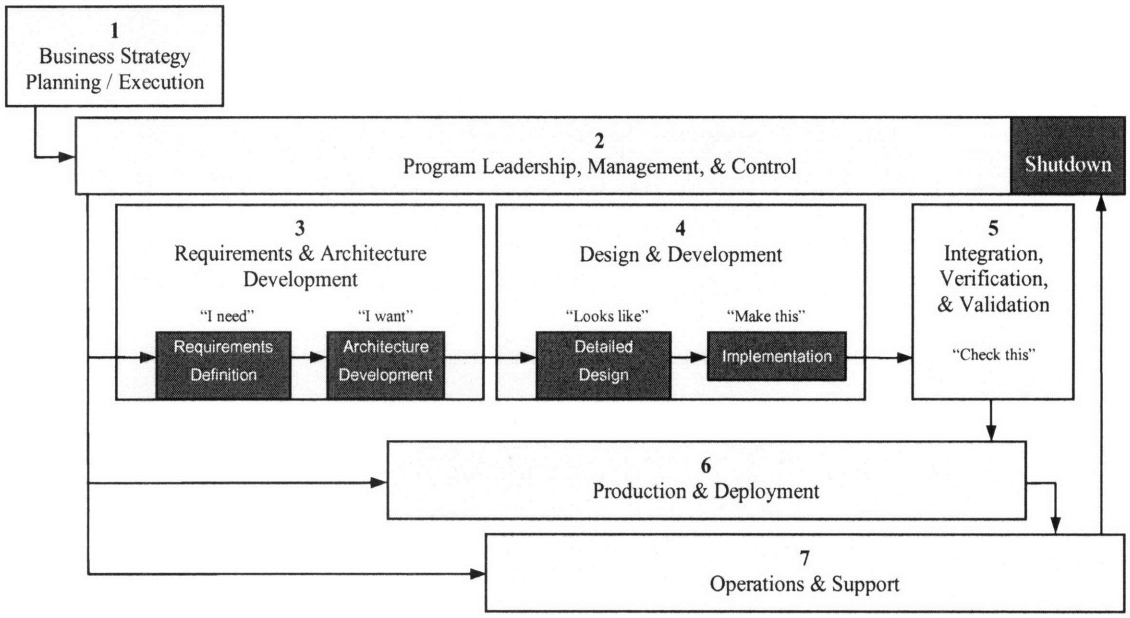

Figure 3-7. The IPDS Lifecycle

IPDS Stages and Tasks

- Stage 1 – Business Strategy Planning/Execution
  - Strategic Planning; Capture / Proposal Process
- Stage 2 – Program Leadership, Management and Control
  - Leadership and Control
  - Pre-Contract Activities; Start-Up Activities
  - Execution Management; Execution Support
- Stage 3 – Requirements and Architecture Development
  - System Requirements; System Architecture
  - Product Requirements; Product Architecture
  - Component Requirements
  - Requirements Management
- Stage 4 – Design and Development
  - Preliminary Design; Detail Design
  - Build Preparation, Component Build and Verification
  - Supply Chain Support
  - Test Equipment and Site/Facility Design and Development
  - Supportability Preparation
- Stage 5 – Integration, Verification, and Validation (V&V)
  - Product Integration; Product V&V
  - System Integration; System V&V
  - Production Needs Assessment
- Stage 6 – Production and Deployment
  - Production Planning
  - Production Material; Production Site/Facility Equipment Duplication and Improvement
  - Production Mfg. & Acceptance; Production Delivery
- Stage 7 – Operations and Support (O&S)
  - O&S Management
  - Mission Support and Services
  - O&S Continuous Improvement

3.4 Phased-Array Radars

Although Raytheon manufactures a wide range of products, the ones with which this work is most concerned are phased-array radar systems. It is therefore useful to briefly discuss the product and the technology on which it is based.

In contrast to articulated radars that rotate to sweep a single beam throughout the scan path, phased-array radars[1] employ a large number of individual antenna elements and exploit the principle of electromagnetic wave interference. Brookner provides a good summary of their principle of operation [Brookner 1985]:

> A phased array is typically made up of a flat, regular arrangement of radiating elements; each element is fed a microwave signal of equal amplitude and matched phase. A central oscillator generates the signal; transistors or specialized microwave tubes such as traveling-wave tubes amplify it. If the signals all leave the array in phase, they will add up in phase at any point along a line perpendicular to the face of the array. Consequently the signal will be strong, capable of producing detectible echoes from objects that lie in its path, along the array's boresight, or perpendicular axis, and within a small angle to each side.
>
> At greater angles to the boresight individual signals from different radiating elements must travel different distances to reach a target. As a result their relative phases are altered and they interfere destructively, weakening or eliminating the beam. Thus outside the narrow cone, centered on the array's boresight, in which constructive interference takes place, objects produce no detectable echoes. Because of the characteristics of interference patterns, the width of that cone is directly proportional to the size of the array. With every element radiating in phase, the beam is in effect steered straight ahead.
>
> Now suppose the signals from each of the radiating elements are delayed electronically by amounts that increase steadily across the face of the array. Each delay causes a signal to lag a fraction of a wavelength behind the signal from the adjacent element. A change in the relative phases of the signals is the result. Now the zone in which the individual signals add up in phase to produce a strong sum signal, capable of detecting targets, lies not straight ahead, down the boresight of the antenna, but off to the side in the direction of increasing phase delay. The angle of the beam reflects the magnitude of the phase shift, the size of the array and wavelengths and the wavelength of the signals. Again the beam takes the form of a slender cone surrounded by regions of destructive interference. *The radar beam has been deflected without a physical movement of the antenna (author's italics).*

---

[1] Also called Electronically-Steered Radars (ESRs).

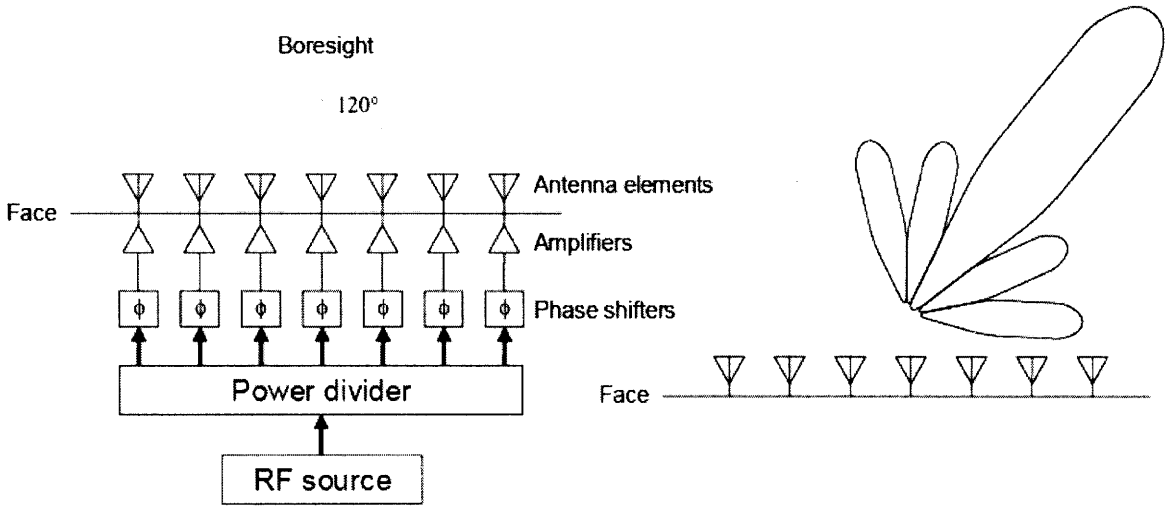

Figure 3-8. Phased-array radar schematic
[Adapted from Fenn et al. 2000]

In other words, the steering of the beam in a phased-array radar is accomplished electronically by adjusting the phases of the signals fed to the individual antenna elements in the array; there is no mechanical movement. This technique, shown schematically in Figure 3-8, typically affords steering capability 60 degrees off the perpendicular axis, which means a single flat array of elements can scan a total of 120 degrees. A phased-array is therefore sometimes constructed as a three-face pyramid; each face scans 120 degrees, and the whole assembly is capable of a 360 degree view. An example of this construction is shown in Figure 3-9.

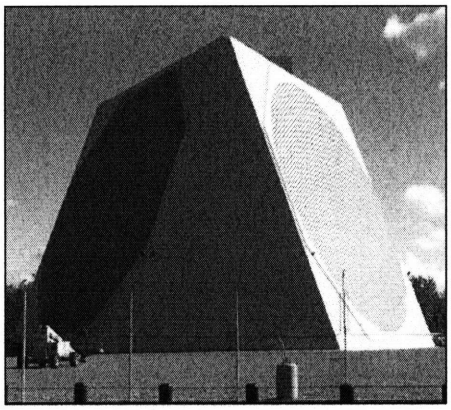

Figure 3-9. Raytheon PAVE-PAWS phased-array radar[1]

[1] PAVE: a U.S. Air Force code-word with murky origins and even murkier meaning; the most commonly used meaning is **P**recision **A**cquisition **V**ehicle **E**ntry. PAWS: **P**hased-**A**rray **W**arning **S**ystem. The image shown is the 90 foot diameter array at the Clear Air Force Station in Alaska.

4 | *Phase I: Variation Awareness*

In the early stages of research and development, potential sources of variability are often hidden. In many cases, the underlying physics of a new technology may not be fully understood. In other cases, although engineers may have a complete picture of the technology's workings, the complexity of the system and its sensitivity to environmental factors may cloud predictions of the technology's manufacturability. In this first stage, effectively addressing variability often requires a combination of uncovering and dealing with known sources, while simultaneously making provisions for yet-to-be-discovered sources. Not surprisingly, successfully performing this balancing act can be difficult.

While known sources of variability can be addressed using a range of engineering approaches, some of which are discussed in Chapters 5 and 6, strategies to guard against future unknown sources may lie as much in the domains of organizational structure and communication as in science. To use a simple analogy, if one's car has a malfunctioning engine, the steps required to fix it can be mapped out if one has the necessary technical knowledge and tools. If, on the other hand, one is worried about unknown problems developing at unpredictable times in the future, a different approach is called for. Since it is not practical to carry around every tool for every conceivable situation, it is more effective to have an organizational structure and communication channels in place that

can field the challenges as they arise (in this case, a towing company, a repair shop, and a phone with which to call them).

This chapter presents a case study of a technology in Phase I. It describes how a decision was made between selecting an internally-developed component and a commercially available one for use in a program, how the development team diligently worked to minimize risk throughout every stage, and how, despite their best efforts, hidden risks stemming from manufacturing variability nevertheless appeared. *The discussion in this chapter draws no conclusions as to whether the decision was right or wrong. Indeed, it will be shown that the chosen path was eventually successful. Rather, this case study demonstrates the strong coupling between R&D and production, as well as the importance of the organizational framework in which the link occurs.* Finally, the lessons that can be learned from this case are discussed.

The selection occurred in 2007; the discussion below presents the case as a sequence of contemporaneous decisions and processes.

4.1 Background: Choosing a DREX synthesizer

At the heart of a phased-array radar system lies the electronics that generate the high-frequency signals used to transmit the scanning beams and receive the returned signals. It is one of the great paradoxes of the industry that the performance of these radars, which are very large and highly complex systems, is wholly dependent on thumb-size electronic chips.

Program A is a large-scale phased-array radar project. The program's radar design utilizes an array of Digital Receiver/Exciter (DREX) modules (circuit board assemblies) to generate the required high-frequency signals. The core of each board is comprised of two digital signal synthesizers, one to generate the transmit signals, and one to generate the signal needed by the receiver. The radar requires 100 DREX modules for each of the three antenna faces, so a total of 600 signal synthesizers are needed. Additionally, although Program A is expected to be the main consumer of these synthesizers, identical chips are also required by Programs B, C, and D. Needs in these projects are 5, 10, and 15 chips per program, respectively.

A critical part of the DREX development effort is the selection of a suitable digital synthesizer. Although many COTS[1] synthesizers are available in the market, the program's performance requirements eliminate nearly all commercial options. Program A needs a chip capable of generating a high quality signal, characterized by low levels of phase noise and sufficient rejection of spurious harmonics, at high frequencies.

4.1.1 Option 1: Off-the-shelf

At the time of selection, there is only one COTS component considered suitable for the program, the DX synthesizer. The DX had been extensively researched by the DREX development team. Designed in Taiwan and packaged by a contract manufacturer in Malaysia, the latest version of the chip is a third generation design. This version is intended to operate at relatively high frequencies whereas prior generations have operated at lower speeds. The power consumption specified by the manufacturer of the chip is considered to be very good. Testing indicates that the component performs well with respect to phase noise, but that the presence of spurious harmonics could be a problem.

4.1.2 Option 2: Internal Development

The main competitor to the DX for the Program A signal synthesizer is Raytheon's home-grown DDSx (generation *x* Direct Digital Synthesizer). DDSx development had begun as an independent IRAD[2] program at Raytheon-A[3] in 2002, completely independent of Program A. The chip is considered to be an evolutionary upgrade to the prior "best in class" DDS chip. Funded by Raytheon business units A and B, a proprietary design is being produced for Raytheon by its spin-off company, Acme Engineering[4]. Initial wafer fabrication took place sometime in 2005.

The module was originally comprised of an ASIC, an EEPROM, and two chip-resistors. However, since Program A's system requirements were such that the EEPROM

---

[1] COTS: Commercial Off-The-Shelf

[2] IRAD: Internal Research and Development

[3] Business unit disguised

[4] The names of some contractors have been changed to protect Raytheon confidentiality.

and resistors are unnecessary, the components have been removed from the design and the module has been redefined as a packaged ASIC.

The core of the module is the ASIC die. The die is fabricated using a proprietary BiCMOS process at the foundry, Progressive Semiconductor[1]. Each wafer contains 50 die. It has five metal layers and is "bumped" using Progressive's own process. The dies are encapsulated into a BGA[2] package with a ceramic substrate.

Silicon-Germanium (SiGe) BiCMOS semiconductor technology, on which the DDSx is based, appeared roughly 20 years ago but reached maturity only recently. Commonly used in ICs to create heterojunction bipolar transistors[3], SiGe is presently the technology of choice in high-frequency T/R modules because it "offers the potential of reasonable power levels and performance ... together with yields unreachable using [other] semiconductors, the ability to integrate RF and logic, and prices approaching those of conventional CMOS" [Schiff 2004]. The upshot for circuit designers is the ability to integrate high-frequency RF, analog, and digital functionality on the same die. Although compound semiconductors using SiGe offer performance advantages over their monolithic counterparts, the technology does have drawbacks. One of these is the sensitivity of HBT performance to wafer sheet-resistance[4].

4.2 Problem statement

In the spring of 2006 the DREX development team has to make a decision between the commercially available DX and the internally developed DDSx. In reaching this decision they must weigh a number of technical and economic factors. Ultimately, the challenge is to find the solution that economically meets program requirements while minimizing risk.

---

[1] The names of some contractors have been changed to protect Raytheon confidentiality.

[2] BGA: Ball Grid Array

[3] Heterojunction bipolar transistors (HBTs) differ from their homojunction cousins (BJTs) in that they use differing materials for the base and emitter regions, a feature that facilitates high frequency operation.

[4] Expressed as Ω/□ (read as *ohms per square*), sheet-resistance is the average resistivity of a thin film multiplied by the film depth, typically the transistor junction depth.

The goal of the discussion here, however, is not to re-examine the decision itself, but rather to analyze how it is made. The DREX synthesizer selection process is used as a case with which to study the issue of variation awareness during Phase I of the transition process. Knowing at the beginning of the study that potential problems are eventually averted, the main objective is to learn how these risks creep into the process despite diligent work on the part of the development team to ensure a secure path. Ultimately, these findings generalize into lessons applicable to the broader world of Phase I transitions.

4.3 Approach

The strategy adopted by the decision makers is to first understand the context in which the decision must take place. This necessitates an understanding of the organizations involved and the manner in which they interact with each other. Once the participants are established, the planned decision making process is analyzed.

4.3.1 Organizations involved in DDSx development

To comprehend the involvement of the various organizations in the DDSx development effort and their relationships to each other, it is first necessary to understand the Trusted Foundry Access (TFA) program of the National Security Administration (NSA).

Microelectronics fabrication facilities are rapidly migrating out of the United States. Companies with the capability of producing advanced semiconductor devices are either going out of business, being acquired by foreign firms, or moving their facilities off-shore [Carlson 2005, Streit 2005]. This trend jeopardizes the ability of American defense contractors to fabricate the ASICs vital to their systems in a secure environment. In response to the growing threat, the U.S. Deputy Secretary of Defense, in October of 2003, launched the Defense Trusted Integrated Circuits Strategy (DTICS) out of which grew the NSA's Trusted Access Programs Office (TAPO) and the TFA program. The objective of TAPO is to guarantee secure access to microelectronics fabrication services to U.S. defense firms. Some of the key attributes of the TFA program are [*Military Information Technology* 2007]:

- Guaranteed access to a trusted U.S.-operated foundry supplier for mission-critical applications.
- Fabrication at various levels of design classification.
- Access for low-volume requirements of less than 100 parts, targeting leading-edge technologies that would otherwise be unobtainable.
- Quick turnaround cycle times to meet schedule requirements.
- Application engineering support, providing technology selection and implementation assistance.

DTICS defines three "trust categories" that range from Category I (vital to mission effectiveness) to Category III (needed for day-to-day business). Systems such as the DDSx that fall into Category I *must* use trusted foundry services for their ASICs.

To realize TAPO's objective, NSA entered into a contract with Progressive Semiconductor that provides trusted access to an array of technologies at the company's multiple facilities. Under the contract, Progressive is only responsible for basic wafer fabrication using a specified process; the customer is responsible for the layout and physical design, as well as the packaging and testing of the fabricated dies[1]. Among the process technologies available to customers under the contract is Process X, Progressive's BiCMOS technology, which is the process required by the DDSx.

Typically, once a contractor has been granted access to the TFA program, the firm communicates directly with Progressive. However, the use of Multi-Project Wafers (MPWs) introduces an intermediary. MPWs are wafers on which dies for unrelated circuits, often from unrelated companies, are fabricated. They are typically used for research and low-volume production runs where the cost of creating a full mask for a single design would be prohibitive. The bundling of the designs onto an MPW mask, and the shielding of proprietary information amongst the companies that share the mask, is

[1] However, the TFA contract does provide the ability for the customer to gain access to Progressive's front-end design services.

performed by a third party. For the TFA program, TAPO hired the National Nuclear Security Administration's (NNSA's) Kansas City Plant (KCP) for MPW management. Although the use of MPWs can greatly reduce fabrication costs for a contractor, it introduces logistical complexity; since KCP plays the role of Progressive's customer, the contractor whose design is being fabricated cannot communicate directly with Progressive. The resolution of technical process-related issues must therefore be indirect, and can be very difficult.

In this case study, Raytheon is developing the DDSx within the TFA framework. With this in mind, the main players in the effort are summarized below (with some information repeated for clarity). Note that the names of some contractors have been changed to protect Raytheon confidentiality.

- *NNSA* (Washington, D.C.): A semi-autonomous government agency within the Department of Energy, NNSA owns KCP and provides MPW management services to TAPO.

- *KCP* (Kansas City, MO): The organization is part of the NNSA's nuclear weapons complex. It assembles and manufactures components for national defense systems, although it also provides services to university and corporate customers. The facility is managed by *Honeywell Federal Manufacturing & Technologies* under contract with NNSA. KCP's role in the DDSx program is to provide MPW services to Raytheon.

- *Progressive* (Au Sable Forks, NY): Under contract with the NSA, Progressive provides the SiGe BiCMOS fabrication services required for the chip. Based on MPW mask-sets provided by KCP, Progressive produces and ships unpackaged, untested dies.

- *Reynolds Electronics* (Japan): The facility produces the DDSx packages and ships them to Reynolds-USA.

- *Reynolds Electronics* (Wilmington, CA): The site assembles the die and the package, and ships a finished but untested part.

- *Acme Engineering* (Harkness, MD): An ASIC design house, Acme Engineering produced the DDSx chip design under contract with Raytheon. Although approximately 30 designers once worked at the company, during DDSx development that number dwindled to a single person. Acme Engineering spun off from Raytheon in January 2002. Daniel P. Burnham, then Chairman and CEO of Raytheon, said, "Formation of this startup represents a disciplined, market-driven approach to penetrating high-growth markets without increased financial risk to Raytheon" [Neely 2007].

- *Raytheon Site 1* (Fern Lake, MA): Manufacturing and operations.

- *Raytheon Site 2* (Jay, NM): Engineering.

- *Raytheon Site 3* (Keene, VT): Initial design and testing work.

The communication channels utilized by these groups are somewhat complex. Figure 4-1 is a representation of the networks that exist prior to, and during, the transition period. The most salient features of both are the numbers of players involved, and the inability of Raytheon to communicate directly with Progressive. Note that once the transition process began, lines of communication were opened between Raytheon Site 1, the other Raytheon sites, and KCP.

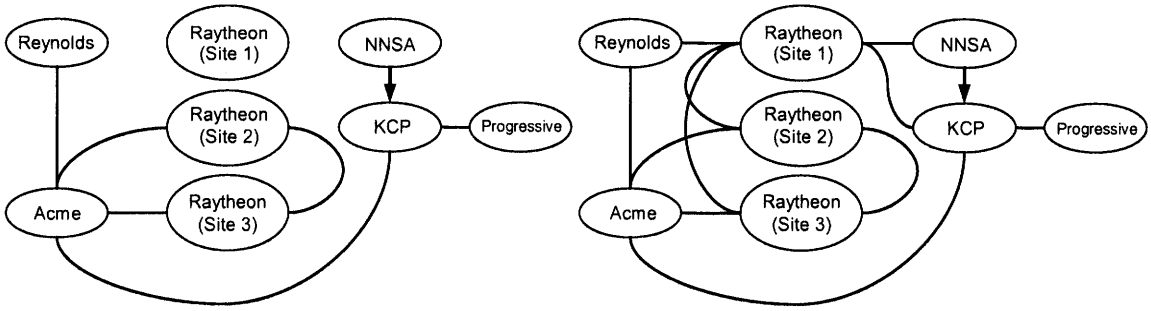

Figure 4-1. DDSx communication map: pre-transition (left) and during transition.[1]

4.3.2 The decision

In April 2006, the pros and cons of the DX and DDSx are compiled. A trade-off matrix is created to aid in the decision making process. A greatly simplified version of this matrix is shown in Table 4-1. Weights are associated with the categories to calculate an overall desirability index; performance is heavily weighted, risk and cost are moderately weighted, and schedule is lightly weighted. The resulting scores are shown in Table 4-2.

[1] Lines without arrows indicate bidirectional communication.

Table 4-1. Decision matrix

Factor	Parameter	DDSx	DX
Technical Performance	Phase noise	○	○
	Spurious harmonics	○	●
	Speed	○	○
	Power	●	○
Schedule	PCB	○	◊
	Design impact	○	●
	Yield	◊	◊
	Lead time	◊	◊
	Test & characterization	○	○
Risk	Cost	◊	○
	Obsolescence	○	◊
	Technical specifications	○	○
	Schedule	○	○
	Company	○	◊
Cost	Recurring	●	◊
	Reliability	●	○

○ = good, ◊ = satisfactory, ● = poor

Table 4-2. DREX option desirability-index scores

Option	Score
DDSx	4.28
DX	3.28

In the final analysis, although there are cost and reliability issues associated with the DDSx, it is found to be the winner by virtue of superior performance. The decision is therefore made to commit to the DDSx for use in the DREX.

4.4 Outcome / Results

In late 2006, eight MPWs are fabricated by Progressive under the trusted foundry program. A handful of die are selected for packaging (by Reynolds) and testing. The initial test results indicate a problem: although the units function well at lower clock speeds, some fail to function properly at the target speed. Simulations performed by the

Site-3 engineers indicate the most likely cause to be variation in the sheet-resistance of the wafers, a process-dependent property. Significant variation in sheet-resistivity exists across multiple wafers and within single wafers. Progressive guarantees a sheet-resistivity of ±25% from target but, due to the secrecy surrounding the process, will not divulge the true control capability. Investigations are further hampered by the inability of Raytheon personnel to speak directly to Progressive due to the Raytheon-KCP-Progressive relationship.

None of the options for addressing the problem is appealing. The stop-gap solution involves performing additional wafer runs wherein different runs will target different sheet-resistivity values in the hopes that some will emerge with an acceptable value. Not only will this be expensive, but an acceptable outcome is by no means assured. On the other end of the spectrum is a redesign that will be more tolerant of variable sheet-resistivities.

While these and other strategies are being formulated to deal with the sheet-resistivity issue, the project management team learns that, of the four programs slated to incorporate the DDSx, only Program B requires operation at the target clock speed; the remaining programs can accept lower speeds. Since many of the units that fail to operate at the high clock speed are capable of functioning at lower speeds, this information significantly boosts the wafer yield. The team decides to earmark those chips capable of high-speed operation for Program B (which only needs five chips), and relegate the remainder to the other programs.

Ultimately, what could be an extremely costly problem for the company is luckily avoided due to a timely relaxation of the requirements. When one considers the case in retrospect, among the questions that arise are these:

1. How did the team find itself in the dangerous position of having committed to a solution that, by the original requirements, would not work, despite diligent efforts to ensure the opposite?

2. Why were the requirements for the other programs unnecessarily demanding at the outset?

Variability awareness

Although it could be argued that a variety of factors helped to shape the chosen solution, one issue stands out more than any other, *the risk associated with manufacturing variability*, which was largely neglected in the decision making process. While significant attention was paid to technology readiness (i.e. the maturity of the synthesizer design), as well as to the cost and scheduling issues associated with the fabrication process, the two issues were treated as uncoupled, and little heed was paid to their interaction. In other words, the impact of process variability on design performance was mostly ignored[1].

Of course, the recognition of process risk is not enough; the magnitude of the risk must also be assessed, a task that is often far more difficult.

> Within the arena of circuit and system design, variation awareness is becoming crucial. Design for manufacturability encompasses a wide array of tools and techniques to improve the performance and yield of circuits in the face of process and environmental variation. These include analytic and simulation methods to understand the impact of variation on performance or yield; variation reduction approaches, particularly increased regularity in design; and robust design and active compensation approaches to ensure operation given expected variations [Boning et al. 2007].

The challenges in developing cutting-edge technologies can be enormous, and it is natural that the nominal performance of the design would receive most of the attention. But, if the technology is to be successfully transitioned to production, *variation awareness* is critical. Without knowing the sensitivity of the design to process-dependent

---

[1] It should be noted that this conclusion is the result of the observations of, and documents available to, the author. It is entirely possible that discussions along these lines did in fact occur in the period leading up to DDSx selection. Indeed, the inclusion of *yield* as a category in the decision matrix is evidence of the recognition of process variability as a factor in selection. However, irrespective of whether or not the issue arose, it does not appear to have been given much consideration in the final decision. To support this position, although *yield* was included in the matrix, it was considered to be a scheduling factor, rather than a risk factor, and was accordingly assigned the minimal weighting.

parameters, and whether or not the allowable bounds for those parameters fall within the process capability, proceeding with production is a risky endeavor.

Of course, uncovering which parameters are susceptible to process variation can be extremely difficult, particularly when dealing with new technologies. Parameter sensitivity studies, even when performed through simulation, are often time-consuming and expensive. Nevertheless, it is crucial to identify potential trouble spots. Once the process-sensitive parameters have been identified, their influence on production can be minimized by working on two fronts: (1) modification of the design, and (2) control of the fabrication process [Boning et al. 2007].

The first of these options, commonly referred to as *design for manufacturability (DFM)*, accepts the inherent existence of process variation, and seeks to produce a design that is robust. In other words, the goal of DFM is a design that is tolerant of process variation, irrespective of whether that variation is due to systematic or random process deviations. Such a design is not always possible, especially when the envelope of technological capability is being pushed. Even when it is possible to produce process-tolerant designs, the effort required to do so, as well as the additional part count that often results, may be costly. For these reasons, designers are judicious in their selection of parameters to design against. In the case of the DDSx, the recognition of sheet-resistivity as a process-dependent parameter to which the design was sensitive did not occur until after the initial fabrication run, well after the ASIC mask-set had been created. Although the value of a design that was robust with respect to sheet-resistivity became immediately apparent, the extremely high cost required to generate a new mask-set, as well as the time needed for the redesign, effectively eliminated the option. Given the demanding pace to which the DDSx design team was subjected, and the high performance targets for which they were aiming, it is unfair to say that the sensitivity to sheet-resistivity should have been recognized. It does, however, highlight the importance of including DFM concepts earlier, rather than later, in the design cycle.

Process control across organizational boundaries

The second option for combating variation, control of the fabrication process, concerns a reduction of randomness at the source. Even though some level of randomness

is present in every process, improved equipment, tighter control of environmental variables, and other factors can oftentimes reduce variability. A problem surfaces, however, when this approach is applied to the case of the DDSx, namely, that Raytheon was not responsible for fabrication and therefore had no direct control over the process. In fact, once the development team traced the performance problems to variations in sheet-resistivity, they sought to obtain detailed information from Progressive concerning process capability with respect to this parameter, and an indication of whether or not anything could be done to tighten control. But the inability of Raytheon to communication directly with Progressive (at least in a technical capacity) by virtue of the TAPO arrangement, combined with Progressive's need to safeguard what it viewed as the confidential details of its proprietary fabrication process, effectively eliminated the second option of improved process control as well. Although the question of whether or not any useful modification could have been made in such a sophisticated process is debatable, the fact remains that the avenue was closed to exploration because of communication limits between the involved organizations.

The interface between program requirements and R&D

The remaining issue, overstated performance requirements, was a function of the interface between the customer programs and those involved in the development effort. In this case, it appears that there was little structure to the interface, and that assumptions were made on both sides of the boundary. The specific instance of required clock speed is a good example. Once it became known that Program A did not, in fact, require operation at the higher speed, and that only one of the programs did require such operation, the question of how the erroneous specification found its way into the development requirements was raised by the development team. To the author's knowledge, no conclusion was ever reached. It is possible the requirement began with Program B, which was the only one that required high-speed operation. If such was the case, one can imagine a scenario where, if the high clock speed was assumed to be a simple thing to attain, it could have been applied to all programs. It is equally plausible that the specification was an artifact of the initial stages of the DDSx development effort which began before the program requirements were known. Irrespective of the cause, a

mechanism to map the program needs to the development effort, and to monitor the connection throughout development, was lacking.

Since Raytheon's Integrated Product Development System is designed to be CMMI compliant, it is useful to examine the issue through a CMMI lens:

> The purpose of Requirements Management (REQM) is to manage the requirements of the project's products and product components and to identify inconsistencies between those requirements and the project's plans and work products...
>
> Part of the management of requirements is to document requirements changes and rationale and to maintain bidirectional traceability between source requirements and all product and product component requirements [*CMMI® for Development, Version 1.2*: 420].

Clearly, at the time the inconsistency between the source (the program) requirements and the product (the DDSx synthesizer) requirements was recognized, bi-directional traceability did not exist. It is tempting to dismiss the problem as a documentation glitch, but this seems inadequate. After all, Raytheon produces a wide range of highly complex systems, and consistently demonstrates that it has the capability to management program requirements. An alternative explanation is that the DDSx development program "fell through the cracks" with respect to requirements management in particular, and CMMI compliance in general, because of its roots as an independent project. Whereas development efforts conceived in support of a specific program are automatically subject to Raytheon's CMMI processes, orphan projects may not be. In many ways this is desirable, since the conventional wisdom holds that highly structured or regulated environments tend to stifle technological achievement. But it does highlight the need for an independent project to be subjected to CMMI processes *once tapped for use in a program and placed on its critical path*. It is at this juncture that the formal requirements mapping could, and probably should, take place.

4.5 Generalization

The study teaches several important lessons that have applicability outside the specific circumstances of this case.

First, the case study illustrates the importance of being aware of potential sources of manufacturing variability, and their impact on functionality, as early as possible in the development process.

Second, the case study highlights the dangers present when the organizational context in which process variation is being addressed is not factored into the risk equation. Although a particular problem may be considered tractable from a technological standpoint, constraints on communication between the groups tasked with solving the problem and/or misaligned interests can be equally challenging. This lesson is particularly relevant in light of current trends toward distributed development and outsourced manufacturing.

And third, the case study underscores the need for well-defined interfaces between development groups (technology makers) and programs (technology users). Specifically, it demonstrates the need for a mechanism to incorporate independently-begun projects into a program's critical path by establishing the interface and formally mapping program requirements to developmental targets.

5 | *Phase II: Design for Manufacturability*

"Design for manufacturability (DFM) seeks methods to improve performance and yield given process and environmental variation, through robust design, increased regularity, and other approaches" [Boning et al. 2007]. Without question, a strong focus on DFM is a critical part of commercial technology development. The point at which DFM receives that focus, however, often depends on the technology readiness level of the design. Although the theoretically optimum approach is to tailor a particular technology for production from the outset, in practice this rarely happens. The challenges inherent in bringing a new technology to fruition are often so daunting that the issue of how to produce it in quantity is flagged as "future work." Eventually, however, the emphasis must shift.

The shift to a DFM focus as a technology passes into Phase II can be difficult for a variety of reasons, not least of which is the cultural difference between R&D and production groups. Things that are obvious points of concern for one group may not be obvious to the other. For example, an R&D engineer may struggle for years to create a specialized circuit. He builds his prototype boards by hand, and carefully solders all his components. To him, with his intimate knowledge of such circuits, it is obvious that any change in the value of a particular critical resistor would impact performance. Yet to the manufacturing engineer who receives the design, and who probably has little knowledge

of how the circuit actually works, the critical resistor looks like all the others. He therefore allows the standard tolerance on its value, and a problem has been created. Although this example is trivial, it should be apparent how the problem can compound in complex systems with thousands of parts, and with multi-person teams on either side of the handoff.

This chapter presents a DFM case in the context of circuit board manufacturing. The performance of the board was critical, and it was known that it could be impacted by variations in the production process. The challenge was to identify those design parameters that influenced the outcome, quantify their effects, and, if possible, determine optimal values that could yield a robust design. Finally, it was necessary to predict process yields for the purpose of deciding whether design modifications were warranted.

The design of the board occurred in 2007; the discussion below presents the case as a sequence of contemporaneous decisions and processes.

5.1 Background: The amplifier CCA

Among the many components that comprise the phased-array radar are a variety of circuit card assemblies (CCAs) on which the electronics that drive the system are implemented. The CCAs required by the Program A radar include 5,568 amplifier boards. Because they process radar excitation signals, these boards are designed for high frequency operation. For efficient operation, the dynamic components of the CCA are chosen so as to tune the board to a specific frequency. Furthermore, since the distances between the components on the board affect its operation due to factors such as the self-inductance of the traces (which come into play at high frequency operation), it is not only the selection of the components that is critical but also the physical positions of the components on the board.

As the amplifier CCA is being designed, the engineering team recognizes that variations in the values and positions of board components can be problematic. The team performs simulation-based sensitivity studies to uncover the most vulnerable components and model the impact that perturbations in their parameters will have on system performance. Their studies show that the location of five key SMT capacitors on the CCA is critical. Specifically, the results suggest that if any one of these capacitors is,

through inaccurate placement, offset by more than ±10 mils in the widthwise direction, the resulting change in impedance of the current path will shift the frequency of the board outside the specification limits, and the CCA will require manual rework to correct it. The recognition of capacitor placement as a manufacturing risk leads to two questions: (1) what is the accuracy and repeatability of the CCA manufacturing process, and (2) is it possible to optimize the placement operation by changing any design parameters?

The CCAs are to be populated with SMT components on an automated production line using a standard solder reflow process. It is this process that determines how accurately the critical capacitors are placed. The primary operations involved in the process are as follows:

1. *Solder paste application*: The CCAs are fed, via a conveyor belt, into a machine that applies solder paste to the pads using a metal stencil.

2. *Component placement.* The boards proceed into a "pick-and-place" machine that utilizes one or more actuated mechanisms to rapidly pick up SMT components from feeder reels and place them at a programmed location on the board. The components adhere to the boards because of the tackiness of the solder paste.

3. *Solder reflow.* After placement, the boards are transported on the conveyor into a reflow oven where they pass through several heating zones. A controlled reflow profile is achieved by controlling the temperatures of the zones and the speed with which the CCA traverses them. As the boards pass through, the solder paste melts and extraneous substances such as flux burn off. The boards cool as they emerge and the liquid solder solidifies into conductive joints.

4. *Optical inspection.* After reflow, the boards are inspected by Automated Optical Inspection (AOI) equipment to verify the existence and placement of all components. These machines use a camera to visually inspect each component site. If a component is found to be skewed or missing, it is flagged and the board

is transferred to an operator for rework which typically involves manual resoldering of the component.

Although it is possible for components to jiggle during conveyance through the line such occurrences are rare, and the placement accuracy is thought to depend primarily on two steps: (1) the accuracy of placement on the solder paste by the pick-and-place machine, and (2) the degree to which the components shift during solder reflow as the solder paste melts. Of these, the reflow operation is considered to be the most critical. Since it is well known that, due to the surface tension of molten solder, components tend to "self-center" during reflow [Aravamudhan], there is a suspicion that the accuracy with which the components are initially placed by the pick-and-place machine is irrelevant.

5.2 Problem statement

The challenge is to identify design and process parameters affecting the accuracy of capacitor placement in the manufacturing line, quantify their effects, and determine values to optimize yield. Ultimately, the question of whether a design modification is warranted must be addressed.

5.3 Approach

Capacitor placement is defined by three degrees of freedom: translations x and y, and rotation θ (see Figure 5-1). The values are defined as offsets from nominal values, so for a perfectly placed capacitor, $x = y = \theta = 0$. The electrical characteristics of the CCAs are such that the most critical parameter is a shift in the x direction. According to the sensitivity studies, this offset must be held to less than ± 10 mils from center.

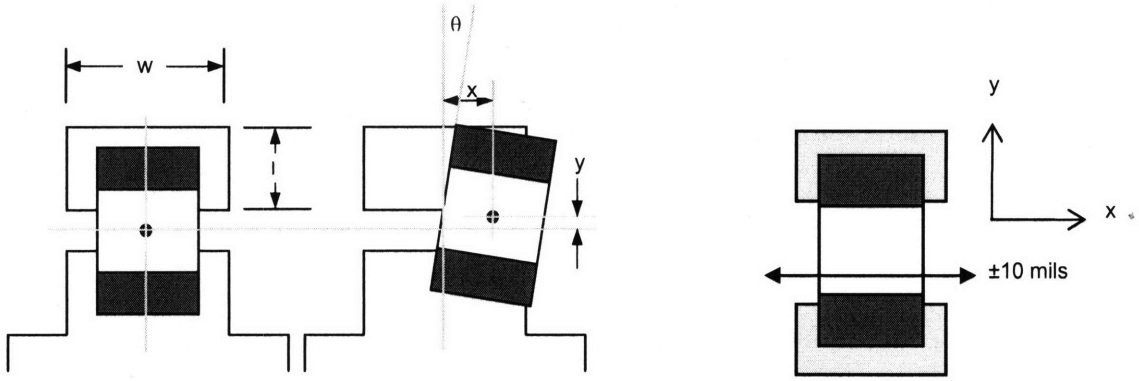

Figure 5-1. Conventions (left) and the specification limits

Many design and process parameters are thought to influence the final location of the capacitors. Discussions are held with the design and manufacturing teams to narrow the list to seven that are considered the most important and controllable. These parameters are:

1. *Solder pad width, w* (i.e. dimension in the *x* direction).

2. *Solder pad length, L* (i.e. dimension in the *y* direction).

3. *Area ratio, AR.* In the literature, one of the main suspects for the skewing and/or tombstoning[1] of components during reflow is thermally imbalanced pads [Straus 1998]. If one of the solder pads is attached to a large section of copper and the other is not, then the one attached to the copper will take longer to heat in the reflow oven. The solder paste on the hotter pad will melt first, and the capillary forces and surface tension of the molten solder can pull and twist the capacitor. Since all copper regions on the CCA have the same thickness, the ratio of one area to another is used to define the parameter.

4. *Full mask v. solder dam.* Rather than using a conventional solder-mask design whereby all regions of the board are masked except the pads, the CCA design

[1] *Tombstoning* refers to a phenomenon whereby SMT components are tipped on end by forces that develop in the melting solder. The term derives from the morbid yet entertaining observation that the vertical components resemble miniature tombstones on the circuit board.

employs several "solder dams." Under this scheme, rectangular solder "pads" are defined on large regions of exposed copper using thin lines of mask material. Capacitors footprints defined using complete solder masks rather than outlines are referred to as "full mask."

5. *Double.* There is concern that having two capacitors in close proximity to one another could exacerbate thermal imbalance problems. *Double* means the capacitor has another capacitor adjacent to it.

6. *Via.* Some of the solder pads on the CCA have vias to the ground plane. In addition to electrical grounding, it is hypothesized that these vias could act as thermal shorts to ground, thus keeping the pad cool in the oven and causing a thermal imbalance with respect to the other pad.

7. *Offset in x.* Although it is well known that the surface tension in molten solder will cause SMT components to "self-center" during reflow, it is not known (1) whether this effect will still work in the case of a large offset in initial capacitor placement, and, if it does, (2) whether those components that are initially offset will center with same accuracy as those that are initially centered.

Other parameters such as the oven reflow profile and solder paste thickness are known to influence the outcome, but they are not included because they are considered to be fixed properties of the process. Of course, it must be noted that the degree to which these factors are truly fixed is unknown since no means exist in the process to actively monitor and control them. Nevertheless, Raytheon's manufacturing experts are confident that the factors are repeatable enough to be excluded from the experiment.

5.3.1 Experiment design

To identify the factors that will influence the process, and quantify their effects, an experiment is set up using DOE techniques. In order to economically generate enough data, dedicated test PCBs are used instead of the actual article. This permits the inclusion of hundreds of capacitors (and hence data points) on each board, rather than the five afforded by the product boards. This also means that the test articles will not be exactly

identical to the production boards, and introduces new potential variation; specifically, it means that the capacitors will be located in different positions in the oven. However, the past experience of the process experts suggests that oven positioning does not significantly affect the results, and that the additional data offered by the test boards warrants the change.

Since the capacitor piece-cost is small, a full-factorial design is defined in the seven variables. Three levels are chosen for the continuous variables w, l and AR to capture any nonlinear responses that might be present. Each of the binary variables *fullMask*, *double*, and *via* have two levels. Although the *offset* variable is continuous it is decided to use just two levels corresponding to "no-offset" and "offset." The value for the offset is a roughly 25% widthwise overhang on the pad; this will cause a capacitor to almost touch the pads of the adjacent capacitor, and therefore represents the maximum recoverable offset that could occur. The levels for l, w, and AR are selected based on the engineering experience of the manufacturing personnel and other factors governing the board layout.

With three 3-level and four 2-level variables, the resulting full-factorial test matrix requires 432 capacitors. According to the design, each board requires 216 capacitors, with two boards, designated I and II, constituting one "experiment." The doubling of capacitors for the *double* variable leads to the inclusion of some duplicate capacitors (i.e. some of the capacitors on a board share the same set of factor values), and the final design consists of 228 capacitors per board. All boards are identical in design; the difference between boards I and II in each experiment is that all components on board I are centered, while all components on board II are offset. A conceptual layout of the board is shown in Appendix A, and a photograph of one of the test boards is shown in Figure 5-2. Note that the outer shape of the board is the same as that of the real article. For practical reasons such as the number of boards per PCB panel, it is decided to run five replicates, resulting in a total of ten boards.

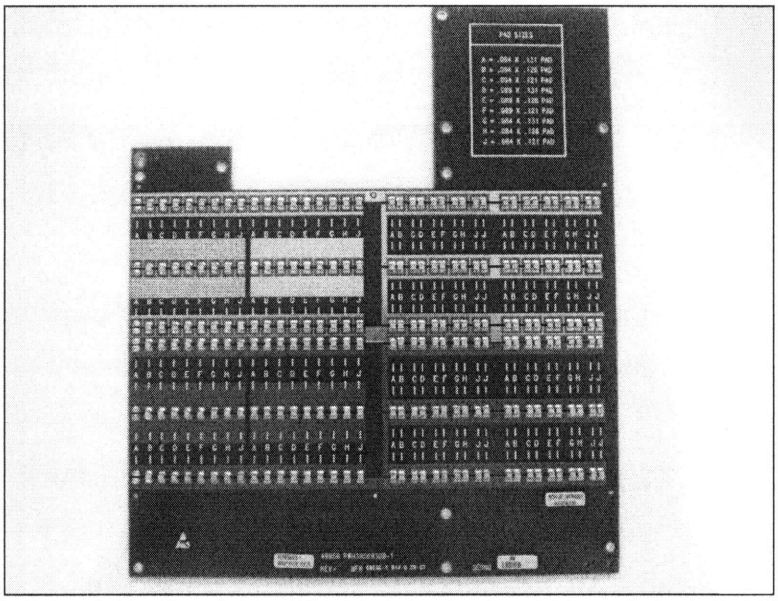

Figure 5-2. Test circuit board used in DOE

The following coded levels are used for the parameters:

Table 5-1. Levels for the parameters

		Level	
3-level	**1**	**2**	**3**
Solder pad width, *w* (mils)	121	126	131
Solder pad length, *l* (mils)	84	89	94
Area ratio, *AR* (approx.)	1	2	3
2-level	**0**	**1**	
Mask type, *fullMask*	dam	full	
Single or double, *double*	single	double	
Vias present, *via*	no	yes	

7 variables:

On each board:

w (3):	solder dam widths: *w1, w2, w3*
L (3):	solder dam lengths: *L1, L2, L3*
AR (3):	area rations: *AR1, AR2, AR3*
fullMask (2):	solder mask dam v. full
double (2):	isolated component v. adjacent
via (2):	via under pad v. no via

Across boards:

x offset (2):	placement offset, 0 and +33 mils

216 + 12 duplicates = 228 capacitors per board

2 boards per experiment x 5 replicates = 10 boards

Figure 5-3. DOE Summary

5.3.2 Execution

The experiment was run on an automated SMT line. The pick-and-place machine is programmed to place the capacitors with zero offset (i.e. centered). Five boards are then fed into line. The major process steps are:

1. Application of solder paste.
2. Placement of capacitors by the pick-and-place.
3. Optical inspection to record component x, y, and θ prior to reflow.
4. Solder reflow.
5. Optical inspection to record component x, y, and θ after reflow.

After the five zero-offset boards are completed, the pick-and-place machine is reprogrammed with the specified offset. The five offset boards are then run. It should be noted that by feeding the boards in this way, as opposed to randomly sequencing the zero- and nonzero-offset boards, the potential for block effects is introduced. However,

since the time required to reprogram the pick-and-place machine is small (less than ten minutes), and since the environment of the assembly area is carefully controlled, it is unlikely that any factors act to create different results on the first and second blocks.

As the experiment is set up on the shop floor, it is discovered that the six capacitor footprints along the left side of the board were placed too close to the edge, and if capacitors are placed there they will interfere with the conveyor gripping mechanism. They are therefore omitted, and each board is populated with 228 – 6 = 222 capacitors. The common parameters for the lost capacitors are (in coded units): $w = 3$, $L = 3$, *double* = 0, *via* = 1; they vary with respect to *AR*, *fullMask*, and *offset*. It is not believed the loss of these capacitors will affect the results.

Placement data, namely the measured x, y, and θ offsets of the capacitors, is recorded by the AOI machine on the line. The same machine is used for measurement before and after reflow.

5.4 Outcome / Results

The data from the experiment was analyzed using statistical methods to characterize the properties of the placement and soldering operation, and conclusions were drawn regarding the behavior and yield of the process.

5.4.1 Characterization

The data from the experiment supports several valuable conclusions. Of particular interest are those design parameters found to have no significant effect on the outcome. While they do not draw immediate attention to themselves, the fact that the process is relatively robust to them answers some long-held questions and could have an impact on future designs.

Several groups, listed in Table 5-2, were defined to categorize the distributions; data for these groups is included in Appendix B, and the analysis of variance (ANOVA) results are presented in Appendix C. The groups were selected so as to highlight the effects of the variables. *Design* refers to the group of capacitors having exactly the same parameters as the current CCA design (the single exception is area ratio – since the capacitors on the current CCA have a range of area ratios, none of them correspond

exactly, and so all were included in this group). Because this group has a relatively small number of samples (15), another group, *Baseline*, was defined that includes all the variables that do not have a significant impact on location (*L, AR, fullMask,* and *double*). Since this group has a much larger sample size (180), its statistics are probably more representative of what would actually occur during production.

Table 5-2. Defined groups

Groups (samples)	*w*	*L*	*AR*	*fullMask*	*double*	*via*	*offset*
No offset (1,110)	all	all	all	all	all	all	0
Offset (1,110)	all	all	all	all	all	all	1
Design (15)	2	2	all	0	0	0	0
Baseline (180)	**2**	**all**	**all**	**all**	**all**	**0**	**0**
With initial offset of 33 mils (180)	2	all	all	all	all	0	1
With vias (180)	2	all	all	all	all	1	0
With narrower solder pads (210)	1	all	all	all	all	0	0
With wider solder pads (180)	3	all	all	all	all	0	0

The main conclusions from the experiment are:

1. **Process characterization**. With respect to the placement and soldering of CCA capacitors on the SMT production line, the process (placement accuracy in the *x*-direction using current baseline values) is characterized by:
    - Specification: ±10 mil
    - Process mean: 0.487 mil
    - Process standard deviation: 2.552 mil
    - Process capability: $C_p = 1.306$, $C_{pk} = 1.243$
    - Reject rate: 0.01% per capacitor

2. **Factor screening**. Table 5-3 summarizes the parameters found to have a statistically significant effect on final capacitor location within the ranges used in this experiment[1].

Table 5-3. Factors affecting placement after reflow

	x	y	θ
Solder pad width	x		
Solder pad length		x	
Area ratio			x
Full mask v. solder dam			
Double v. single		x	
Via v. no via	x	x	
Offset v. centered	x		x

3. **Self-centering capability**. If a capacitor is inaccurately placed (specifically, offset in the x-direction) by the pick-and-place machine:
   - The capacitor *will* substantially re-center itself during reflow.
   - The capacitor *will probably not* be centered as accurately as one placed correctly. As might be expected, the distribution is skewed toward the side to which the capacitors were offset.

4. **Effects of vias**. The presence of vias on the solder pads has some impact on location accuracy. In the experiment where a large number of vias was used, the standard deviation for capacitors on vias was 2.776 mils versus 2.552 with no vias. It is believed that this is due to interactions between the surface tension of the solder and the via holes, or leakage of the solder through the via holes, or both. Since the current CCA design utilizes solder pads with single vias, it is not thought that they will have a significant impact on production. *However, the observed effects were unexpected, and are highlighted for consideration in future CCA designs.*

[1] A p-value of 0.01 was used as the threshold.

5. **Effects of solder pad widths.** There is a correlation between the width of the solder pad and the amount to which the capacitor is likely to be off-center, as summarized in Table 5-4. The most obvious conclusion is that wider pads permit a greater range of motion for the capacitors.

Table 5-4. Dependence of variability on solder pad width

w (mils)	Std. Dev.
121	2.270
126	2.552
131	3.307

6. **Bimodal *y*-direction distributions.** With respect to the final location in the *y*-direction after reflow, it appears that the capacitors tend to "snap" into one of two stable positions, leading to a bimodal distribution. The affect is more pronounced when vias are present. Although this behavior occurs in *y*, a non-critical direction, it is nonetheless an unexpected and interesting result.

7. **AOI measurement limitations.** Odd results were obtained when the pre-reflow positions of the offset capacitors were measured: the measurements show extremely high consistency (standard deviation = 0.894), and the mean (approx. 22 mils) was less than the programmed offset of 33 mils.

8. **Rejection rates.** Rejection rates for various groups are shown below in Table 5-5. A *reject* is a capacitor that falls outside the ±10 mil specification limits. Board rejections rates are calculated assuming five critical capacitors per CCA. Two sets of rates are shown: those using the sample mean, and those assuming a centered, or zero, mean. Although the fictional centered rates differ slightly from the actual, the conceptual simplicity of dealing with a centered process when communicating with design and manufacturing teams makes them the more useful numbers.

Table 5-5. Predicted reject rates[1]

		With sample mean			Centered		
Groups	Samples	Probability outside ±10 mils	Probability of board rejection	Rejects per 1000	Probability outside ±10 mils	Probability of board rejection	Rejects per 1000
Design	15	0.000084	0.000420	0.42	0.000042	0.000210	0.21
Baseline	**180**	**0.000117**	**0.000585**	**0.58**	**0.000089**	**0.000445**	**0.44**
With initial offset	180	0.061802	0.273104	273.10	n/a	n/a	n/a
With vias	180	0.000502	0.002507	2.51	0.000315	0.001574	1.57
With narrower pads	210	0.000011	0.000055	0.05	0.000011	0.000055	0.05
With wider pads	180	0.003405	0.016909	16.91	0.002497	0.012423	12.42

5.4.2 Optimization

The experiment identified three factors that influenced capacitor placement in the critical dimension: (1) solder pad width, (2) the presence of vias, and (3) centering of the component on the solder paste. Of these, the first two are adjustable design parameters (the third was studied purely to understand the self-centering characteristics of the process, and the results confirm the relatively obvious conclusion that capacitors placed in a centered position are more likely to end up centered than those that are offset). This information was then used to address the main question facing the design team: is the current design adequate, or are modifications required?

An inherent trade-off existed with regard to solder pad width. Although the results confirmed the intuitive belief that narrower pads would yield more accurate positioning, this was not necessarily the best overall solution. Other factors unrelated to component placement were known to impact the frequency of the completed circuit boards. Since the easiest method for tuning an off-frequency board is to adjust the position of a capacitor, tightening all the solder pad widths would remove this capability. From this standpoint, the best compromise would be to narrow all but one of the pads, leaving the remaining one as the designated "tuning" capacitor. Although the gains from this change were not

[1] Assuming a normal distribution.

known, the data does provide an upper bound on the improvement: if *all* of the pads were narrowed, the reject rate could be expected to decrease from 0.58 to 0.05 per 1000 boards, a gain of 0.53 per 1000.

The difference in rejection rates between the boards with and without vias – 2.51 and 0.58 per 1000, respectively – translates to 1.93 boards per 1000. Since the experiment used multiple vias on all pads, whereas the CCA has only single vias on some pads, this difference is an extreme upper bound. Although some gains might be achieved by moving vias off of the pads, the impact on performance due to impedance changes would mean that the redesign effort and cost would be significant.

Ultimately, the main question was whether, in light of the results, a redesign was warranted. It was estimated that an inaccurately placed capacitor would require 8.9 minutes to rework at a labor rate of $23.35 per hour. Given the Program A need of 5,568 boards, the upper bound on possible gains was 0.53 + 1.93 = 2.51 per 1000, or 14 boards, which represents a total rework cost over the run of $50. Although the cost of a redesign effort was not known, it was certainly several orders of magnitude larger. *Therefore, although the results indicated that some optimization was possible, the potential cost savings did not warrant the effort. From a business standpoint, the current design is adequate.*

5.5 Generalization

Although the experiment was conducted on a specific circuit board design, the results obtained are general and relevant to any board subject to soldering on the SMT line. The characterization of component placement accuracy, and what design variables affect it, will be useful in the creation of future high frequency circuit layouts.

6 | *Phase III: Process Control*

In an ideal world, every new product entering Phase III (emphasis on process control) would have a proven production line with specialized high-volume manufacturing equipment waiting for it. While this does occasionally happen, it is also true that some technologies are produced using the same machines and techniques that were used for prototyping. Production methods that have evolved this way typically have at least one of the following characteristics:

- WIP is moved by hand
- Workflow is erratic
- Many of the operations are performed by hand
- Machines with different capabilities are used for the same operation.
- Little or no process data (e.g. cycle times, yields, etc.) is recorded

The last point, lack of process data, is especially paralyzing, for without it the potential benefits from correcting any other factor cannot be measured. It is for this reason that the capturing and analysis of process data is a necessary first step when adapting developmental fabrication techniques for higher-volume manufacturing.

This chapter describes a sophisticated device that evolved in this manner and the problems that plagued its production. It further discusses how the need for a statistical process control system was identified, the approach taken to implement such a system, and the software tools that resulted. Within the context of this specific case, it is shown that the usefulness of such a system derives not just from its functionality, but from its flexibility and the ease with which it can be brought to bear on new processes.

The design of the system occurred in 2007; the discussion below presents the case as a sequence of contemporaneous decisions and processes.

6.1 Background: SAW devices

When a compressive or tensile stress is suddenly induced on the surface of a substrate, the resulting pressure wave will propagate throughout the substrate at the speed of sound. When the direction of propagation is along the surface of the substrate, the amplitude of the wave will typically decay exponentially with the depth of the substrate, and the wave is referred to as a Surface Acoustic Wave (SAW). SAW devices comprise a class of electromechanical components that utilize this effect to realize electrical oscillators, transducers, and filters.

A common embodiment for a SAW device consists of a piezoelectric substrate, usually quartz, and two sets of interdigitated electrodes. As shown in Figure 6-1, when a voltage is applied to the input electrode, it induces a longitudinal stress in the piezoelectric substrate. The wave then propagates along the surface until it reaches the output electrode where the piezoelectric effect translates the stress in the material into a voltage. When used as an oscillator, the received signal is amplified and fed back to the receiving electrode which causes a corresponding wave to travel in the opposite direction. Since the sonic speed of the substrate is known, it is possible to tailor the transit-time of the wave by adjusting the distance between the electrodes. If the geometry and material condition are carefully controlled, SAW devices such as these can function as highly accurate frequency sources.

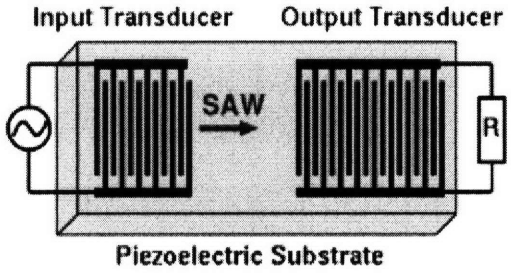

Figure 6-1. A typical SAW device

The FCS group at Raytheon manufactures a variety of SAW devices for different applications. The devices are grouped into four product families: bandpass filters, oscillators, dispersive delay-lines, and non-dispersive delay lines. Bandpass filters with center frequencies from 25 MHz to 2 GHz are produced, as are oscillators with output frequencies of 480 MHz. The oscillator frequencies are trimmed to within ±2 ppm during production, and are designed to drift less than 1 ppm per year[1]. Not surprisingly, such a degree of precision places enormous demands on the manufacturing process. To achieve the required level of geometric control, all fabrication steps are performed in a cleanroom environment.

6.2 Problem statement

The Frequency Control Solutions (FCS) group has its roots in R&D, and continues to function as a research center for the company, developing technologies and exploring new applications. Yet it is also responsible for the reliable and profitable production of certain components, including SAW devices.

The evolution of the SAW production system followed a common path. Like other processes that originate in the development phase, the current system is essentially an expansion of what was originally a prototyping process. But there are good reasons for keeping it: production takes place in a 10k cleanroom[2], many of the operations require batching, and although the output may be high by R&D standards, it is not high enough

---

[1] From Raytheon sales literature.

[2] A class 10,000 cleanroom, characterized by fewer than 10,000 0.5μm or smaller particles per cubic foot, and fewer than 70 5μm or smaller particles per cubic foot.

to warrant elaborate cleanroom automation equipment. In most respects, the production system in use is well suited to its purpose.

However, the sensitivity of the process causes it to suffer from low and widely varying yields. For instance, for all SAW devices manufactured between January 1st and October 15th in 2007[1], 49% of the units were scrapped. Interviews with area personnel revealed that, historically, scrap rates range between 30 and 70%. COGS[2] for a typical SAW device is, on average, $31.42 (exact cost depend on device type and frequency), and production levels are roughly 1,000 per year. Although some of the parts are scrapped before completion thereby reducing the invested cost, the inability to measure certain parameters during processing means that many parts do reach completion and are scrapped during final testing. If all the units scrapped over the given time period had been scrapped in the final process step, the estimated loss would have been over $9M. Although this is an upper bound, there is little doubt that motivation to improve the production yield exists.

A simplified schematic of the SAW production process, with confidential details omitted, is shown in Figure 6-3. The processing of subassemblies A and B are typically carried out in parallel and then the two are joined for final processing. Subassembly A begins as a wafer of piezoelectric material such as quartz. Each wafer yields about nine units, although this depends on the type of device. Subassembly B is the cover assembly, and there are commonly four covers per wafer. The subassemblies are diced independently and then fused together to form the final assembly. Included among the final assembly steps is a process to trim the device frequency to the desired target. The characteristics of the trim process are such that it can only decrease the frequency of a device, and has a range of roughly 5 ppm. Devices having frequencies outside of this window must be scrapped. The cumulative value added in each step of the fabrication process is shown in Figure 6-2.

---

[1] Manufacturing start dates.

[2] COGS: Cost Of Goods Sold

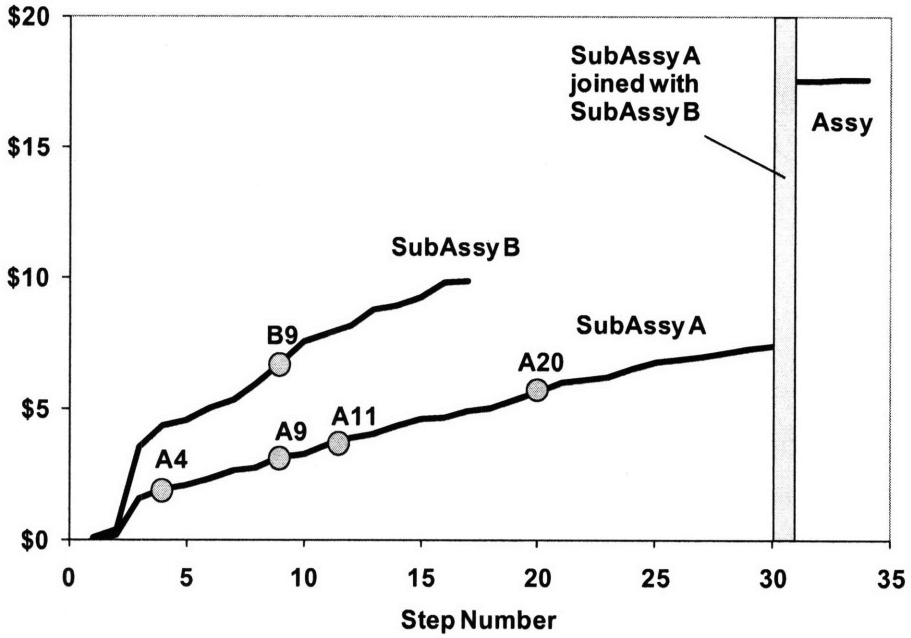

Figure 6-2. Scrap costs (disguised) for SAW process steps. The critical steps are labeled.

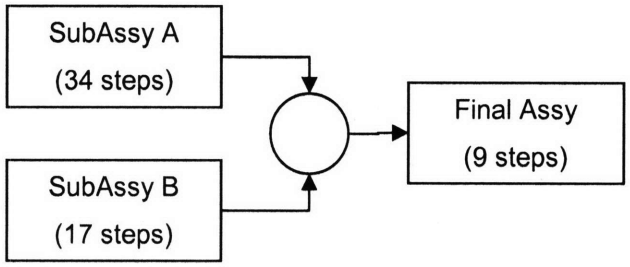

Figure 6-3. SAW device manufacturing process

At several points in the process, thin-film physical vapor deposition (PVD) is used to create metal layers on the device. It is well known among the engineers and operators familiar with SAW device production that this operation is responsible for much of the performance variability they observed in the final devices. The nature of the variability is unknown, however, as is the correlation between deposition errors and final frequencies.

PVD is a fabrication process whereby evaporated metal is deposited onto a substrate in a vacuum chamber to form an extremely thin film. Several variants of the technique exist; the one employed in SAW device production utilizes an electron beam evaporator whereby a high-energy beam from an electron gun is used to boil the material

to be deposited. Metal vapor condenses and accumulates on the substrate at a rate of approximately 10 nm/s. There is enough inherent variability in the process that open-loop, or timed, deposition is not adequate. Therefore, a monitor crystal is used to indirectly observe the thickness of the deposited layer. This crystal, a small piezoelectric resonator, has a known electromechanical resonance. As metal is deposited onto its surface in the evaporator, along with the target wafer, the added mass lowers its resonance. Circuitry in the evaporator controller electrically measures this frequency shift and calculates the thickness of the deposited layer. A single crystal can be used for multiple depositions until the metal thickness reaches a point where the crystal must be discarded.

Although the process is controlled using the monitor crystal, there are a number of factors that act to reduce the overall accuracy of the operation. These include:

- The metal thickness on the monitor crystal is not directly observed, but is extrapolated from its frequency shift.
- The monitor crystal is located in a different part of the vacuum chamber than the target wafer.
- Even within the target wafer the deposited thickness can be non-uniform.

Five process steps suspected of being problematic, four for subassembly A and one for subassembly B, are identified by manufacturing personnel; they are listed in Table 6-1. All but one, *Mill 1*, involves thin-film metal deposition. Figure 6-2 shows the cumulative added value, or scrap costs, for the process steps, with the five critical steps labeled. Process rejects occurring further along in the process are, of course, more expensive because they have accumulated more work. Note that there are raw material costs associated with both subassemblies, but they are relatively small and therefore neglected.

Table 6-1. Critical SAW operations

Operation	SubAssy	Step No.
PVD-1	A	A4
Mill-1	A	A9
PVD-2	A	A11
PVD-3	A	A20
PVD-4	B	B9

The primary problem is that process variations occurring in some or all of the steps are thought to contribute to the high scrap rate. To begin to address the issue, it is necessary to have data from each of the steps. Currently, a large amount of process data, roughly five year's worth, has already been collected using the existing legacy system. However, due to the limitations of the system, including the lack of input validation[1], the existing data is extremely difficult to use. Although this situation has existed for years, operators have continued to collect the data.

The main challenge lay, however, not just in the creation of a system for process data collection and monitoring, but in making it flexible, extensible, and capable of rapid deployment to new areas and processes. The reasons why a more capable process control solution had not already been implemented has little to do with technology, for a wide range of data systems exist in the marketplace, but rather with the cost required to implement a new system and integrate it with the existing databases.

The problem, then, is to create and implement a system to enable the company to take a set of fabrication operations that had originally been geared toward R&D, such as those in the SAW area, and quickly establish processes control.

---

[1] *Input validation* refers to the practice of ensuring, through software, that only data in the proper format is entered. It prevents, for example, a user from entering text where a number is expected, or from inputting values that are clearly outside a sensible range. The legacy system in the SAW cleanroom had no such validation. As a result, in one field where a nominal value of "100" was expected (representing 100 Angstroms), a range of variations was present, including "100," "100A," "100Ang," "0.100," and others. Although a person familiar with the process can discern the meanings, data analysis software cannot. Therefore, years of legacy data was, for all practical purposes, useless.

6.3 Approach

Although the SAW production area is the main focal point, the adopted strategy is to create a system that can be seamlessly deployed to any process in the facility[1]. To meet this objective, the solution needs to either build upon, or interface with, Raytheon's existing data management and monitoring system, VBS (Virtual Business System). VBS consists of a graphical front-end for information display and manipulation, and a relational database for storage. The solution also needs to be flexible enough to meet a wide variety of manufacturing needs.

The technical approach focuses on two areas: (1) database design (the "back-end") and (2) user interface design (the "front-end"). Since part of the overall strategy is to leverage the existing VBS database, design in this area deals with choosing data structures that are efficient yet generalized enough to accommodate virtually any process. Likewise, user interface design focuses on functionality and ease-of-use, but also has to satisfy the generality requirement.

6.3.1 Database design

Once the process measurements and SPC parameters have been identified, the next step is setting up tables to store the data using the existing VBS IT infrastructure.

VBS data is stored in a relational database. As stated by E. F. Codd in his seminal paper [Codd 1970], "the relational (or model) view of data ... provides a means of describing data with its natural structure only – this is, without superimposing any additional structure for machine representation purposes." The structure of a relational database provides for, among other things, what Codd referred to as "data independence," that is, a decoupling of the format in which the data is stored from the way in which it is used by an application. Such independence permits one to modify and grow the database without impacting the client applications.

---

[1] In theory, the system could be deployed to any area in the global enterprise. As of this writing, VBS had seen use in several other North American facilities and was expanding; any tool built on its framework was instantly available to users company-wide.

The VBS database uses the common SQL (Structured Query Language) protocol. SQL is a standard, and by far the most popular, language for interacting with relational databases.

Database design almost always involves a tradeoff between generalization and performance. Due to the ways in which most database queries are executed, those schemes that enable the best performance (measured both in terms of speed and in the simplicity of the required programming) often afford little or no flexibility with regard to what data is stored. This is not a problem in those cases where the information to be stored is fully defined at the outset, but in those scenarios where growth is anticipated, committing to a fixed structure is often impossible. An alternative to the rigid approach is a general scheme commonly known as the Entity-Attribute-Value (EAV) model, in which the names of attributes are explicitly identified for entities within the data structure along with their corresponding values. Using this scheme, information can be stored in a generalized way, and few limitations exist on future expansion.

The difference between the two approaches may best be illustrated by way of an example. Suppose there exist two processes for which we wish to record data. Process A has three parameters: *temperature*, *pressure*, and *processing time*, while Process B has two: *length* and *width*. The different tables that would be used to store this information are shown in Table 6-2. Note that, in these tables, the SAN (System Assigned Number) is an index used to maintain the uniqueness of the records. In the pre-defined approach, a separate table would be used for each process, and the columns of each would be specific to the process. Conversely, using the EAV approach a single table would be used to store all of the data (note that, in this example, *entity* corresponds to *process* and *attribute* corresponds to *param*). It must be borne in mind that the simplicity of the single EAV scheme is not without cost. The table will obviously be larger, and, in general, more computational effort will have to be expended to extract and sort data. In this application, however, the benefits of a flexible system outweighed the drawbacks, and the EAV model was chosen.

Table 6-2. Example: Comparison of pre-defined and EAV data storage approaches.

Pre-defined

Process A

SAN	Temp	Pres	Time
1	69	234	132
2	68	234	145

Process B

SAN	Length	Width
3	30	5
4	40	6

EAV

SAN	Process	Param	Value
1	A	Temp	69
1	A	Pres	234
1	A	Time	132
2	A	Temp	68
2	A	Pres	234
2	A	Time	145
3	B	Length	30
3	B	Width	5
4	B	Length	40
4	B	Width	6

Data structures are then designed to enable the storage of process information and the definition of SPC parameters. Since the data concerning all current and future processes has to comply with these structures, it is necessary to ensure that they possessed adequate flexibility. One of the more mundane yet challenging aspects of this requirement is the ability to cope with arbitrary process batching. For instance, suppose there exists a wafer-level operation that one wishes to control. And further suppose that, once diced, the wafer will yield ten units, each with a unique serial number. Conceivably, one might want to correlate the wafer-level process data with the performance of the finished unit, which implies a recording of the process data on a unit-by-unit basis. However, since it is a batch process, the units in the batch must be correlated with a single data point. For a single operation, this correlation is trivial; but when batching is generalized, as it must be, to arbitrary groupings within a process flow, and amongst entirely different processes, the problem becomes more difficult. The chosen solution is to allow for completely arbitrary groupings, correlated by a single SAN. Relationships amongst the groupings can then be reconstructed later according to the specifics of the process batching and the data correlations to be analyzed.

The *Data* and *Batch* tables, the fields of which are shown in Table 6-3 and Table 6-4, respectively, store the process data and serial number grouping information. The *SPC Parameter* table, shown in Table 6-5, stores the information that defines the control

parameters for each process, and warrants further explanation. In this table, a process is constituted by a unique *Mfg_Area-Operation-Resource* combination[1]. *Param* is the parameter for which SPC control is being applied, for example, deposited metal thickness. *Target* is the value that was intended for that process. For example, an operator might have intended a deposition of 30Å (the *target* value) of metal on a wafer, but the actual deposition might have been 33Å (the *param* value). Inclusion of a target value in the table is necessary for those operations that perform variable functions, often because they act on different part types. Extending the metal deposition example, if type A parts require a metal thickness of 30Å, but type B parts need only 25Å, then is makes little sense to control the absolute thicknesses. There are multiple approaches for normalizing the values, and they vary by operation and machine type. For the sake of simplicity, the scheme chosen for this system was to apply control to the error between the realized and intended values, i.e. *param – target*. *UCL*, *CL*, and *LCL* are the SPC control limits[2] that, when combined with a chosen set of rules, will define a process fault. *LSL* and *USL* are the part specification limits[3]. Strictly speaking, the specification limits are not a property of the process and, in fact, a process might be used on different types of parts with wildly different specifications. However, many single-part processes do exist, and the users wanted the fields to be included as *optional* parameters to enable the prediction of process rejection rates.

---

[1] Actually, since the *operation* names are unique, the inclusion of the *manufacturing area* is not required to define a process; however, it does help to clarify the concept. The field is included in the table primarily as a means to filter relevant operations in the interface. Without it, the user would be presented with a daunting list of all operations that exist in the company.

[2] UCL, CL, and LCL are acronyms for the Upper Control Limit, Center Line, and Lower Control Limit, respectively.

[3] LSL and USL are acronyms for the Lower and Upper Specification Limits, respectively.

Table 6-3. Fields of the *Data* table

Field	Description
SAN	Unique record identifier
OPERATION	The operation for which data is being recorded
PARAM	The name of the operation parameter
VAL	The recorded value for the operation parameter

Table 6-4. Fields of the *Batch* table

Field	Description
SAN	Unique record identifier
SFC_NUMBER	Unit serial number
SHOPORDER	This field was used to interface with the data in Raytheon's shop floor data management system
ITEM	Part number
DTTM	Date/time stamp of the data record (system assigned)
TEST_DATE	Date/time of the operation (user entered/editable)
LOGON_ID	Username of the operator who entered the data

Table 6-5. Fields of the *SPC Parameter* table

Field	Description
SAN	Unique record identifier
MFG_AREA	Manufacturing area
OPERATION	The operation to be controlled
RESOURCE	The resource (machine, typically) to be controlled
PARAM	The parameter to be controlled
TARGET	The parameter target (i.e. the intended value)
LCL	Lower control limit
CL	Centerline
UCL	Upper control limit
LSL	Lower specification limit
USL	Upper specification limit

6.3.2 Interface design

With the data structures thusly defined, the next step is the design of the interface. Most of VBS is built using the LabVIEW language. Developed by National Instruments initially as a language to interact with the data acquisition hardware the company manufactures, LabVIEW is a graphical programming language. As stated by National Instruments [1]:

> The ... LabVIEW graphical development environment helps create flexible and scalable design, control, and test applications ... The NI LabVIEW graphical dataflow language and block diagram approach naturally represent the flow of data and intuitively map user interface controls to data, so programmers can easily view and modify data or control inputs.

[1] http://www.ni.com, accessed 24 October 2007.

Written in LabVIEW, the various VBS interfaces are compiled and then distributed via a locally installed "Launch Pad" application. Although the SQL database is centralized, all computations executed by an interface are performed locally by the client computer on which it is running. The benefit of this approach is that the computation effort is distributed, and does not tie up a central server. The drawback is that the results of every database query must be sent from the server to the client, resulting in heavy network traffic.

Although an entire suite of tools can be developed to interact with the information in the SPC database, the system, at a minimum, requires two main interfaces: (1) one for manual data capture, and (2) one for analyzing process data and defining control parameters.

The Data Entry Interface

The need for new LabVIEW interfaces to be compiled and distributed conflicts with the goal of creating an SPC system that permits new operations to be added quickly and easily. Using the conventional approach, if SPC was to be applied to a new operation, a member of the VBS team would have to create an interface that includes the parameters unique to that operation. Although the programmer could do this by modifying an existing interface, the process would still be time-consuming and error-prone. Furthermore, the interface would need to be compiled and deployed to the operator stations for use, something that could be problematic and tedious if more than one iteration was required before a suitable design was reached.

To address this problem, a scheme is adopted whereby a generalized data entry interface (DEI) that includes features common to all operations is designed, and the variables specific to the operation are defined parametrically. Using this approach, a single interface can be compiled and distributed *once*. The addition or modification of operations will then require only the manipulation of parametric data, and no software changes will be needed.

Hard-coded into the DEI are controls to select the manufacturing area and operation for which data is being collected. Once these are chosen by the user, the system queries the Shop Floor Data Management (SFDM) database for the shop orders

associated subject to that operation. The user can then add serial numbers from those shop orders to define the process batch. In addition to the saving of data, the DEI also enables the user to modify or delete past record sets. A screenshot of the interface is shown in Figure 6-4.

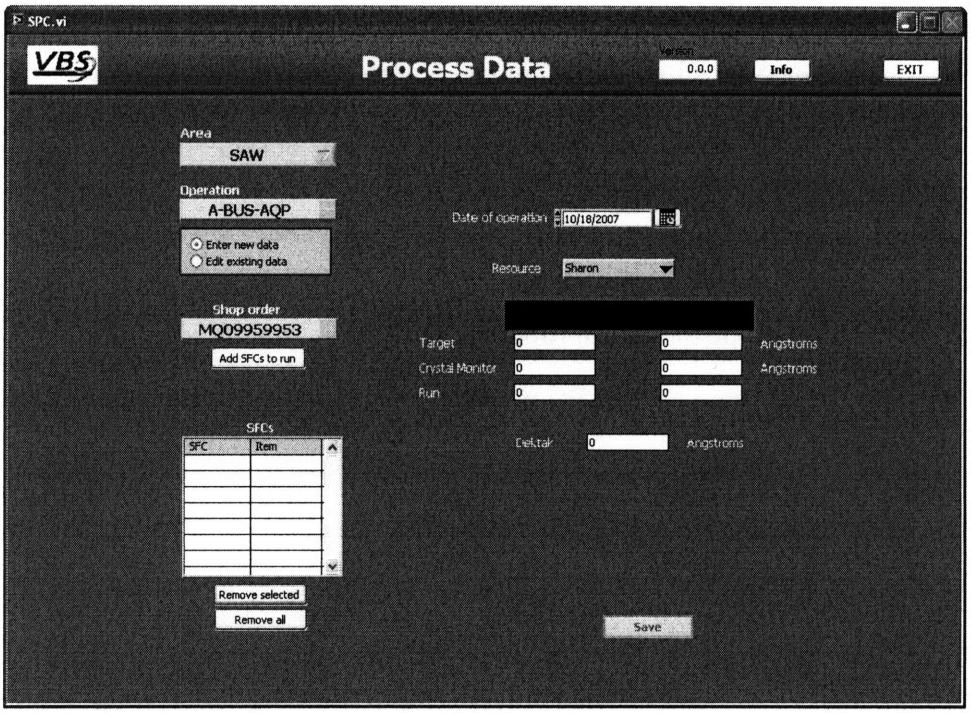

Figure 6-4. The data entry interface (black areas redacted)

The flexibility of the DEI derives from the way in which its controls[1] are defined. At the core of the interface are two tables that exist in the same database as the other SPC tables: the DEI controls table and the DEI pull-down values table, the fields of which are shown in Table 6-6 and Table 6-7, respectively. The bulk of the screen layout is defined in the controls table. The number of controls, their types, physical layout, and the SPC parameters to which their values are to be mapped are stored here. An auxiliary table, the pull-down values table, is needed to store the values for pull-down selection controls which have discrete and pre-defined values (an example is a pull-down from which the user would select which resource/machine was being used).

---

[1] The term *controls* is used here in the programming sense where a control can be any of the familiar graphical objects used for user input. Examples include text boxes, pull-down list, and check-boxes.

Table 6-6. Fields of the *DEI Controls* table

Field	Description
MFG_AREA	Manufacturing area
OPERATION	Operation
POS_X	The x-position of the control on the screen (in pixels)
POS_Y	The y-position of the control on the screen (in pixels)
WIDTH	The width of the control (in pixels)
HEIGHT	The height of the control (in pixels)
TYPE	The type of control (options: *numeric*, *string*, *label*, *pull-down*)
PARAM	The name of the SPC parameter to be mapped to this control value
F1	Wildcard field for type-specific attributes
F2	Wildcard field for type-specific attributes
F3	Wildcard field for type-specific attributes
F4	Wildcard field for type-specific attributes

Table 6-7. Fields of the *DEI Pull-Down Values* table

Field	Description
OPERATION	Operation
PARAM	The name of the pull-down control that will contain this value
VALUE	The value to add to the pull-down control

The DEI is a *self-generating application*. When the user selects an operation, the interface queries the tables defined above for the controls with which it is to be associated. It then dynamically generates the required controls, lays them out on the screen, and maps their values to the proper SPC parameter names. When the user clicks *save*, corresponding entries are recorded in the SPC data and batch tables. Similarly, when prior record sets are viewed or edited, the data is read from the database and

dynamically mapped to the corresponding control. Although this approach results in a visual layout that is not as crisp as what could be achieved if a custom interface were designed for each operation (since the layout must accommodate those operations that have a large number of controls, operations with fewer controls tend to look somewhat sparse), it offers the distinct advantage of being able to add and modify operations without any software changes. Instead, changes need only be made to the data in the controls tables. Thus, new operations can be brought into the SPC system in a matter of minutes and with little cost.

The Process Analysis and Control Interface

The DEI enables the collection of SPC-relevant data, but it does nothing with regard to process analysis or control. For this, another application is created, the SPC Process Analysis and Control Interface (PACI), shown below in Figure 6-5. The PACI enables users to observe the distribution of process data and its relevant statistics, and then use that information to set up control parameters. To set up or modify control values, the user first uses the data filter to select a dataset representative of a time when the process was thought to be under statistical control. He then enters the sample size for the operation, n. The PACI then calculates and displays the estimated process mean (μ), standard deviation (σ), centerline (CL), and natural upper and lower control limits (UCL and LCL, respectively) [Montgomery 2001: 208]:

$$\textit{Natural "3}\sigma\textit{" control limits:} \qquad \begin{aligned} UCL &= \mu + 3\frac{\sigma}{\sqrt{n}} \\ CL &= \mu \\ LCL &= \mu - 3\frac{\sigma}{\sqrt{n}} \end{aligned} \tag{6.1}$$

The user then enters the desired control limits as part of the SPC parameters. Although the user will in most cases want to use the suggested natural limits, the interface permits the use of arbitrary bounds. The additional SPC parameters are the control parameter, target (if there is none, the user can select *NA*), and part specification limits. As previously discussed, the specification limits are not part a property of the

process; but they are included as optional values so that, if used, the PACI can display process capability (C_{pk}) and predicted reject rates.

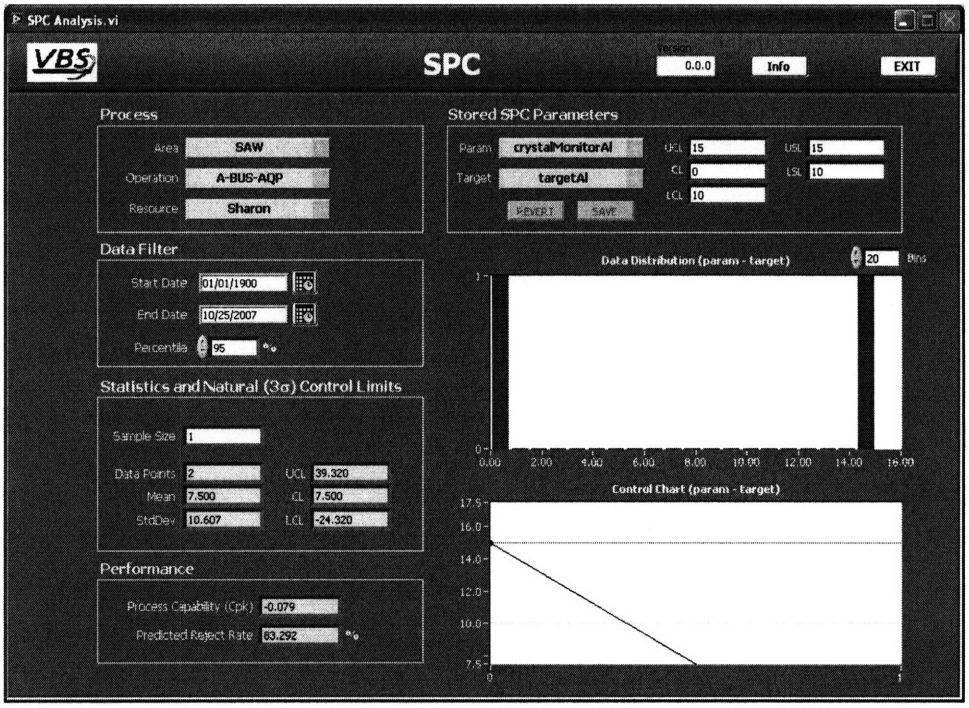

Figure 6-5. The process analysis and control interface

6.4 Outcome / Results

In this section, the deployment of the SPC tools is described, as is the resulting outcome of the project. The lessons from the case are then generalized.

6.4.1 Deployment

The DEI was deployed in the SAW cleanroom on 9 October 2007 for beta testing[1]. To ensure that no data was lost, a phased approach was adopted wherein data was entered into the legacy SFDM system as well as the new VBS-based DEI; the

[1] *Beta testing* refers to a standard phase of software development wherein the product is released to a limited audience for use in the field.

complete switchover to VBS would then occur once a suitable level of confidence in the new tool was reached.

Once the DEI was deployed, development efforts shifted to the PACI. On completion of the tool, legacy SAW production data from SFDM was ported to the VBS database. Due to the unstructured nature of the data (the same lack of structure that instigated the development of the DEI in the first place), the transformation was extremely tedious and time consuming. In some cases, reasonable interpretation of the data was impossible, and some entries were omitted. Eventually, however, a large percentage of the dataset was salvaged and adapted to the new format, enabling its insertion into the VBS database. It was thus possible to evaluate the functionality of the PACI tool before a statistically significant dataset was available from the shop floor.

The extracted data contained process information for 39 part numbers in eight product families (types of SAW devices)[1] spanning five years; approximately 8,500 parameter entries for two operations were salvaged. Figure 6-6 below is a screenshot of the legacy data being analyzed by the PACI. Although the dataset was not complete it did provide a starting point with respect to process statistics and control limits.[2]

---

[1] SAW devices such as oscillators have "dash numbers" which differentiate a single device family into parts tuned to different frequencies. For instance, part "x-y" would be a part of family x (e.g. an oscillator) and frequency y (e.g. 1 GHz).

[2] Note the large gaps in the control chart in Figure 6-6. In examining the legacy data it was discovered that such gaps were common, and typically arose from the ad-hoc nature of the manufacturing process. In this case, the gaps occurred because an alternative to the *Sharon* machine was used for the process for a time, either because *Sharon* was broken or it was being used for another product.

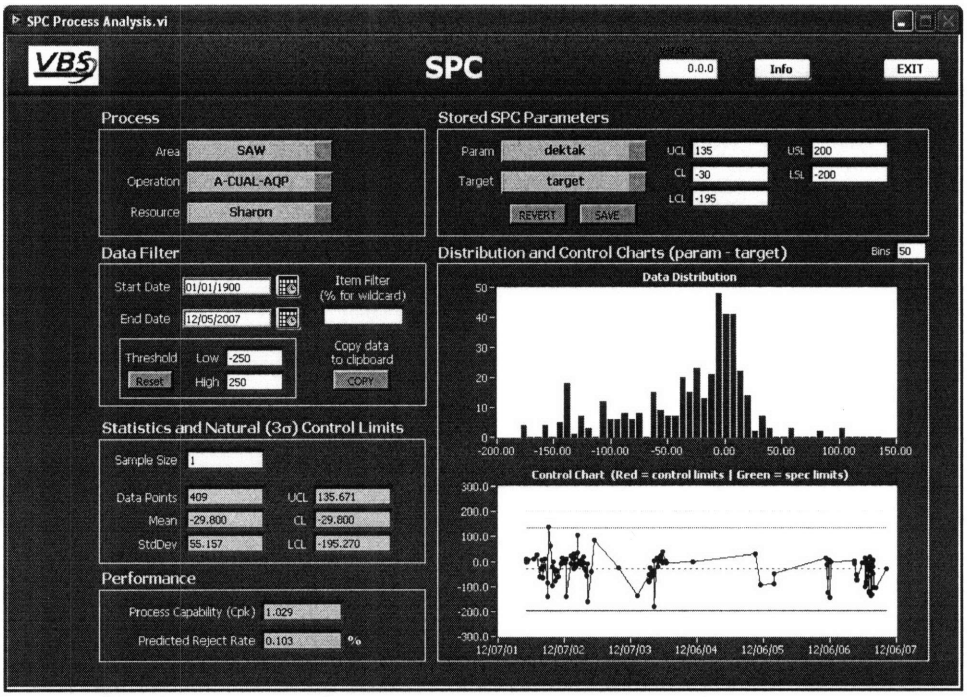

Figure 6-6. Screenshot of PACI used on legacy data

6.4.2 Status

Although the SPC software had been deployed and was in use in the SAW cleanroom, competing projects and time constraints prevented it from being used enough to generate a useful dataset prior to the conclusion of the internship. Even though legacy data does exist, corrupt values and missing entries preclude its use here to demonstrate the software's use.

While the capabilities of the software would ideally have been demonstrated using actual data, it is nevertheless possible to illustrate its potential value using a hypothetical example. Suppose the metal deposition machine being used for step B9 (operation PVD-4) periodically has a problem: as the monitor crystal accumulates metal and nears the end of its useful life, its signal begins to drift. Further suppose that B9 is not currently subject to statistical process control. Once the need for close monitoring is realized, the flexibility of the software will enable personnel to quickly configure a data entry screen for B9. When a statistically significant dataset has been generated using the DEI, the PACI can then be used to calculate and implement control parameters.

Thereafter, any drift of the monitor crystal will immediately show up on the PACI as a violation of the control rules, and production can be halted until the monitor crystal is replaced. Since the small drift in resulting metal thickness is likely to be detectable only through statistical means, if process control is not implemented on the operation it is conceivable that production will continue for several days; only when the resonant frequencies of the final units are tested will the fault be discovered. Since the scrap cost for operation B9 is $6.74 per unit, and production rates are approximately 1,000 per year, the cost of, say, a two day delay in fault detection would be $36.93. (Note that this value is disguised to protect Raytheon confidentiality; the actual value is significantly higher.)

Is the new software directly responsible for the savings? Yes and no. After all, existing commercial SPC tools could be used instead and achieve the same result. *The intrinsic value of the software derives instead from the ability to expand it to other manufacturing areas, and bring additional processes under statistical control, quickly and with minimal cost.* Therefore, although the above example illustrates a direct cost savings of $36.93, the true benefit of the approach is avoiding the additional investment that would otherwise have been required to bring B9, and all other operations requiring monitoring, under control.

6.4.3 Future Work

GUI layout tool

With the DEI framework in place, it is possible for tools to be created to further simplify and accelerate the creation of data capture screens for new operations. Currently, setting up the controls and parameters for a new operation is simple and straightforward, but it requires a system "administrator," i.e. someone with intimate knowledge of, and direct write access to, the database. It also requires the control information to be entered in a tabular format, something that many people are uncomfortable with. To expand this capability to the common user, it is possible to construct a tool that enables controls to be visually created and positioned on the screen, and then generates the corresponding entries in the controls table. Such a tool would provide the same functionality as common

WYSIWYG[1] GUI[2] layout tools such those found in Microsoft's "Visual" series of products.

Expansion of the VBS infrastructure

The VBS infrastructure will likely require considerable work as the system grows. Because the system relies on client-side data processing, one aspect of the system stands out as a potential bottleneck: data traffic on the Raytheon grid. Since database query results are transmitted to client computers and displays, the total amount of information transmitted will increase with both the number of clients, and the size of the database. Whereas server capacity can easily be increased by adding additional units, the available bandwidth is limited by the physical cabling in Raytheon facilities, as well as by internet traffic in those cases where enterprise-wide data is being transmitted. If the client-side architecture of VBS is maintained as the system grows, steps will need to be taken to monitor and limit the quantity of data transferred.

6.5 Generalization

Although the project was conducted within the SAW manufacturing area, a generalized architecture was a fundamental goal of the effort, and the system can be used throughout the enterprise. Furthermore, the approach employed is applicable to any situation in which an organization is seeking to bring a set of processes under statistical control. Assuming there is access to a central enterprise database, adapting the system for use in other manufacturing environments would only require adjustment of the labels used in the *Batch* and *Params* tables; the structures could remain the same.

---

[1] WYSIWYG: What You See Is What You Get

[2] GUI: Graphical User Interface

7 | *Conclusions*

Although this work examines the transition of technology from R&D to production in the context of three specific cases, the lessons learned can be generalized into guidelines that are applicable to a broader range of products and scenarios.

This chapter presents the results of this thesis. The main conclusions can be grouped into two categories: (1) the characterization of the transition-to-production process using three distinct phases, and (2) a listing of guidelines for each phase. It is important to note that the guidelines listed in this chapter are by no means exhaustive, nor are they sufficient to guarantee a successful transition. Rather, they are recommendations that, when followed, can help to avoid some of the hidden pitfalls that lie on the path to commercial production.

7.1 The three phases of the transition

The minimization of the risks stemming from process variability is, or at least should be, of primary concern to engineers throughout the entire development lifecycle from initial concept to mature production. However, the focus of the risk mitigation effort changes as the cycle evolves. To better characterize the process whereby a technology is taken from the laboratory to the production floor, three phases were defined. The phases, first introduced in Section 1.2, are repeated below:

Phase I. A new technology is emerging from an R&D center and has been tapped for use in a product. The design is in the final development stages but questions may still exist regarding its achievable performance. The technology may be in the process of being qualified for use in a particular application. The emphasis is still on development.

Phase II. The functionality of the technology has been demonstrated, and the design may be undergoing modifications to make it more practical and/or economical to manufacture. The emphasis is shifting from development to production.

Phase III. The technology is in some level of production. R&D work is considered to be complete, and the emphasis is entirely on production.

To accept the fact that the risk mitigation focus changes in each phase, one must first acknowledge that, in practice, the development cycle rarely follows a theoretically perfect path. In a perfect world, all aspects of process variability would be addressed from the very beginning, and would receive constant attention throughout the cycle. However, technical unknowns, time constraints, political pressures, and other factors often act to prevent this. For example, the technical challenges facing R&D engineers are often so demanding that they adopt an approach of, "first we need to make it work, then we'll worry about how to build it." Such compromises are inevitable, and are the reason for the change in focus throughout the development lifecycle.

Rather than fight to prevent this change in focus, a more pragmatic approach is to understand when the change takes place, how the focus shifts, and what areas require special attention to ensure that no risks are overlooked.

7.2 Phase I guidelines

Shown below are guidelines for minimizing the risks related to process variability during Phase I.

Guideline 1: Technical Due-Diligence

The traditional engineering mode often decouples design from fabrication. Although an awareness of available manufacturing processes might guide the design of a product, the inherent variability of a process is sometimes neglected. Unfortunately, the advancing complexity of high-tech systems is typically matched by increasing sensitivity to variations in physical parameters. Additionally, compiling the list of performance-impacting parameters, and then identifying those among them that are susceptible to process variability, is often far from trivial. For example, a designer might be aware that the diameter of a drilled hole varies by a known amount, and would account for that randomness in his design. But what about variations in material properties in the vicinity of the hole due to increased heating of a dull drill bit? Or perhaps the effects of vibrations on other components in the assembly as the hole is drilled? As the technology becomes more sensitive, the list of factors continues to grow, and the line separating the relevant from the irrelevant becomes harder to discern.

As the Phase I case illustrates, it is critical to identify all the performance-impacting parameters, determine how they will vary during processing, and make the design robust to them. Just as the development risk of the DDSx could have been mitigated by realizing that the design was sensitive to variations in sheet-resistivity, and that the chosen fabrication process could not control the parameter tightly enough, the success rates of other projects can be boosted by accounting for the limitations of the process in the design.

The most direct and effective way to address the problem is to educate system-level design engineers about process variability issues, and train them in techniques for dealing with them. Indeed, Raytheon has already begun pursuing this approach.

Another means for accomplishing this would be to include a *technical due-diligence* step in the design cycle. Such a step would include a formal review of all suspect design parameters, a study of how each is expected to vary during processing, and forecasts of the how multiple variations will interact at the system level (e.g. Monte Carlo simulations).

Guideline 2: Robust organizational structures

There is a growing paradox in modern manufacturing: while the trend toward the outsourcing of manufacturing continues to grow, further divorcing R&D and production groups, the coupling between the two with respect to product performance is stronger than ever. This presents an enormous challenge to the developers of advanced technologies. Although the goal of the technical due-diligence activity mentioned above is to weed out all potential problems related to process variability, it is naïve to expect total success. It is therefore necessary to ensure that mechanisms exist to address unknown issues as they arise.

It is important when considering the case of the DDSx to understand that the threat posed by the design's sensitivity to sheet-resistivity was only part of the problem. Because of the communications limitations that existed between the developers and the fabricators, it was difficult, once the problem was discovered, to determine whether it could be solved through process control or whether a design change was required. This leads to the following conclusion: *not only must the design be robust to variations in parameters, but the organizational structure in which a technology is developed must also be robust.* In the end, companies must do everything they can to prevent unexpected problems, yet simultaneously make provisions for dealing with them when they do.

One of the most effective strategies to strengthen the organizational structure in this regard is to build-in a practice of *continuous improvement.* By constantly working to improve and streamline operations, the organization will stand a much better chance of evolving and keeping pace with environmental changes. Ultimately, the goal of continuous improvement efforts is to circumvent potential problems *before* they occur rather than endlessly responding to them.

Guideline 3: A well-defined interface between developer and customer

In the case of complex systems with long development lifecycles, requirements management can be extremely challenging. The study of the DDSx case revealed that it can be particularly difficult when program management chooses to utilize a technology that was originally designed for a different use.

To ensure that a new technology is indeed suitable for a given application, a formal process is needed to map program needs to device requirements. The process would constitute an interface between the engineering team developing the technology and the personnel responsible for integrating it with the program's systems.

7.3 Phase II guidelines

Shown below is the guideline for minimizing the risks related to process variability during Phase II.

Guideline 4: Process characterization

Before a component can be designed for manufacturability, it is first necessary to fully characterize the fabrication processes that will be involved. As the product becomes more complex and sensitive to variations in design parameters, the extent to which this characterization is conducted also increases. The study of the amplifier CCA soldering process presented in this work illustrates the point: whereas the minute variations in component placement caused by reflow soldering can normally be neglected, the demanding performance requirements of the CCA necessitated an investigation of the process with respect to this parameter. Although the results of the study indicated that a design change was not warranted, such an outcome was by no means guaranteed, and the importance of such characterizations should not be overlooked.

7.4 Phase III guidelines

Shown below are guidelines for minimizing the risks related to process variability during Phase III.

Guideline 5: Data-driven process control

Newly developed technologies are sometimes produced, at least initially, using the same techniques and equipment that were used to prototype the developmental units. Although this practice has several benefits including familiarity and lower startup costs, it is often difficult to control production quality in a lab setting using lab equipment. The

initial production of sensitive high-tech systems can be accompanied by high reject rates, and a quantitative and methodical approach is required to boost process yields.

The capturing and storing of digital process data is essential; data collection provides the information with which the problem can be tackled. Not only can it illuminate sources of variability, but it also enables the implementation of statistical process control solutions to maintain quality conditions.

Guideline 6: Agile data management systems

For large-scale and mature manufacturing operations, a process control system that requires a lot of time and resources to set up can be an acceptable solution. But for new technologies that typically begin their commercial lives in smaller production runs, it can be hard to justify the expense and effort needed to set up such a system, especially if the process is still in flux; but therein lies the catch, for without a data collection and control system, it might not be possible improve and grow the process.

The way out of this trap is to utilize a process control system that is agile, that is, one which can quickly and easily be adapted for use on new and changing processes. Although a company may have a large and impressive data management system, if it is difficult or costly to bring a new process (e.g. one needed for a newly developed technology) under its control, it offers little value.

7.5 Summary

The overarching conclusion reached in this work is the importance of recognizing all process-related risks that pose a potential threat to a developmental technology, not just those that are technical in nature. Organizational structures, communication channels, data management systems, and procedures can have inherent risks in much the same way that manufacturing processes do, and must be likewise addressed.

Appendix A: Test PCB used for experiment[1]

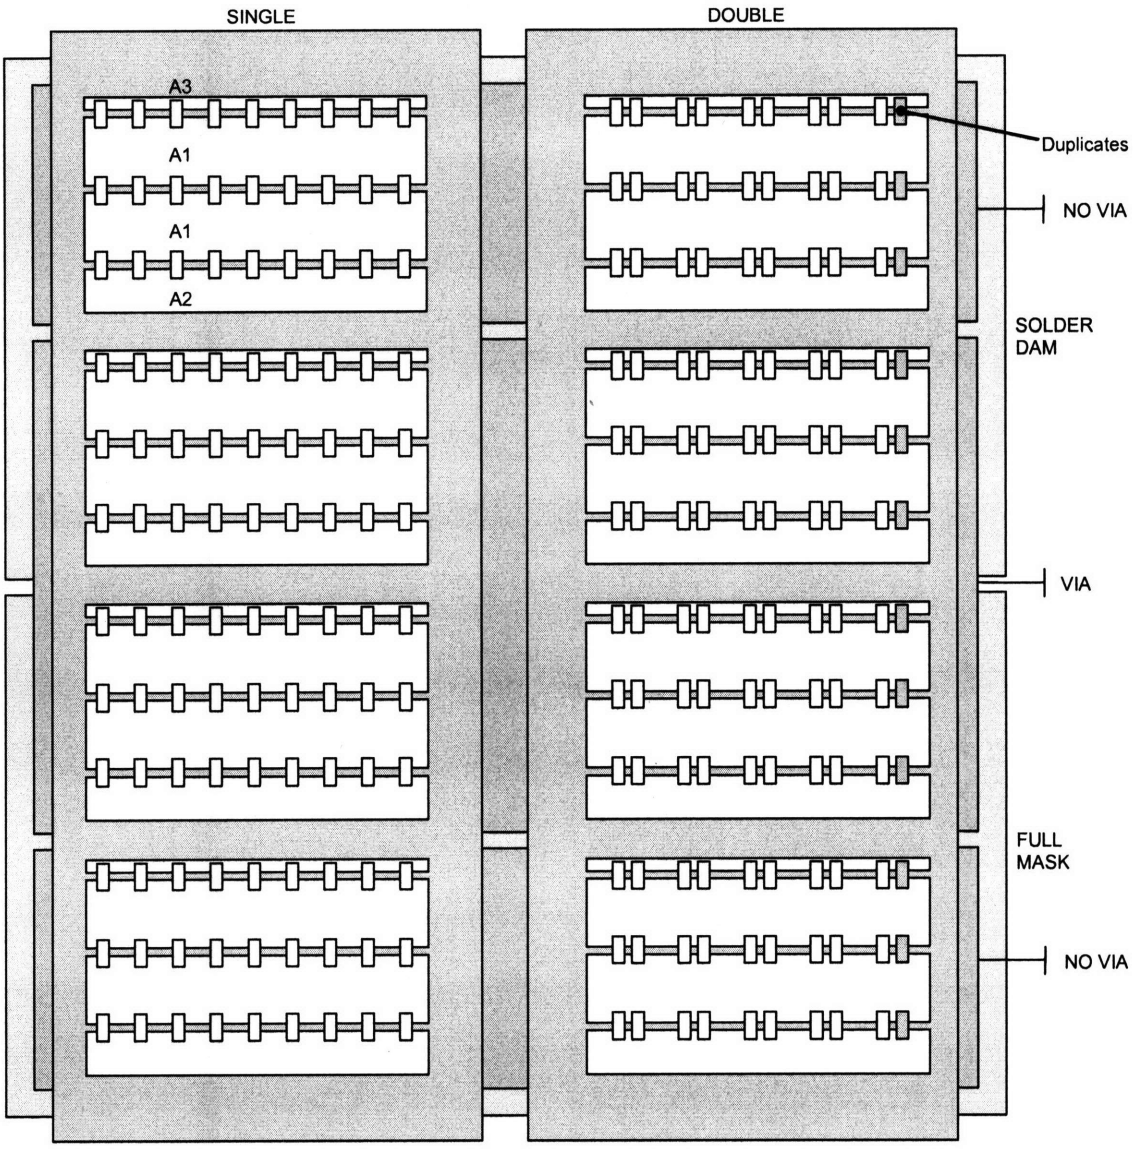

Solder mask dimensions:

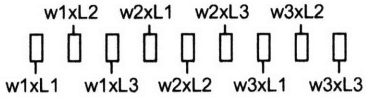

VARIABLES:

On board:

w (3):	**solder pad widths**
L (3):	**solder pad lengths**
AR(3):	**area ratios**
fullMask (2):	**solder mask dam v. full**
double (2):	**isolated component v. adjacent**
via (2):	**via under pad v. no via**

Across boards:

offset (2):	**placement offset, 0 and +33 mils**

[1] Conceptual layout; physical layout differed.

Appendix B: DOE data

Distributions prior to reflow

x and y in mils (thousandths of an inch); θ in degrees.

Figure B-1. Distribution of x prior to reflow

No offset (Offset=0)

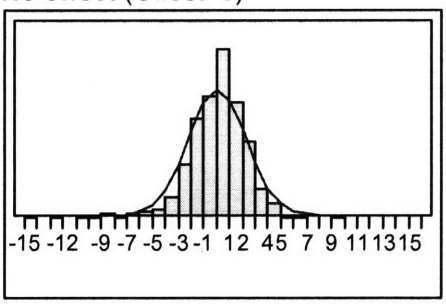

		Prob. outside ±10:
Mean	0.0570689	0.000010
Std Dev	2.2673242	
Std Err Mean	0.0680538	Prob. outside ±10 (centered):
upper 95% Mean	0.1905975	0.000010
lower 95% Mean	-0.07646	
N	1110	

Offset (Offset=1)[1]

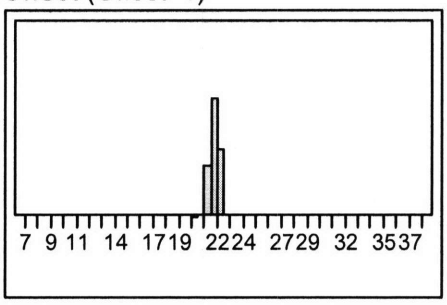

		Prob. outside ±10:
Mean	21.847202	N/A
Std Dev	0.8943934	
Std Err Mean	0.0268452	Prob. outside ±10 (centered):
upper 95% Mean	21.899875	N/A
lower 95% Mean	21.794528	
N	1110	

[1] This data is suspect. It is believed the offset exceeded the measurement range of the AOI equipment.

Figure B-2. Distribution of y prior to reflow

No offset (Offset=0)

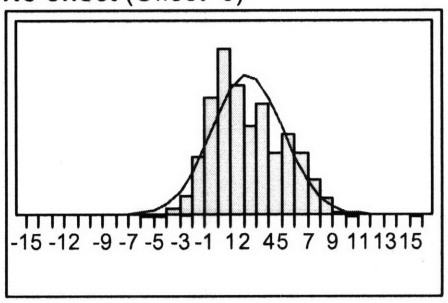

Mean	2.3155281
Std Dev	2.7738378
Std Err Mean	0.0832568
upper 95% Mean	2.4788867
lower 95% Mean	2.1521696
N	1110

Offset (Offset=1)

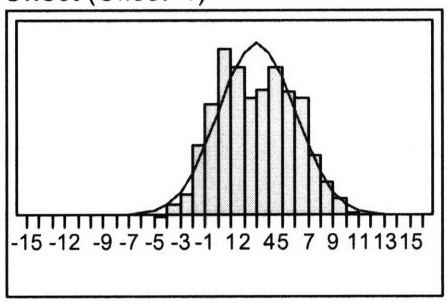

Mean	3.0098248
Std Dev	2.9859731
Std Err Mean	0.089624
upper 95% Mean	3.1856765
lower 95% Mean	2.833973
N	1110

Figure B-3. Distribution of θ prior to reflow

No offset (Offset=0)

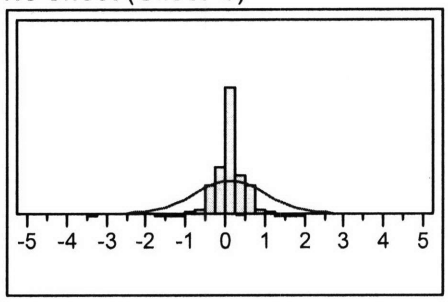

Mean	0.1281335
Std Dev	0.9086042
Std Err Mean	0.0272841
upper 95% Mean	0.1816677
lower 95% Mean	0.0745992
N	1109

Offset (Offset=1)

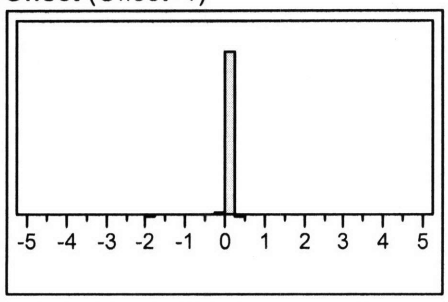

Mean	0.0784296
Std Dev	0.0790325
Std Err Mean	0.0023743
upper 95% Mean	0.0830882
lower 95% Mean	0.073771
N	1108

Distribution of *x* after reflow

x in mils (thousandths of an inch).

Figure B-4. Distribution of *x* after reflow: design values

Design (w=2, l=2, Full Mask=0, Double=0, Via=0, Offset=0)

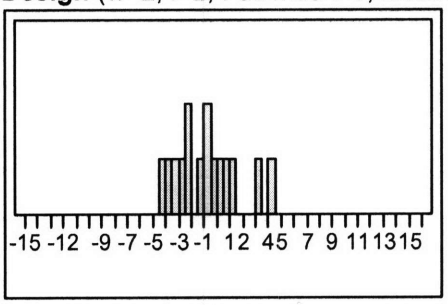

		Prob. outside ±10:
Mean	-0.771654	0.000084
Std Dev	2.4427948	
Std Err Mean	0.6307269	Prob. outside ±10 (centered):
upper 95% Mean	0.5811211	0.000042
lower 95% Mean	-2.124428	
N	15	

Figure B-5. Distribution of *x* after reflow: the effects of an initial *x* offset and vias

Baseline (w=2, Via=0, Offset=0)

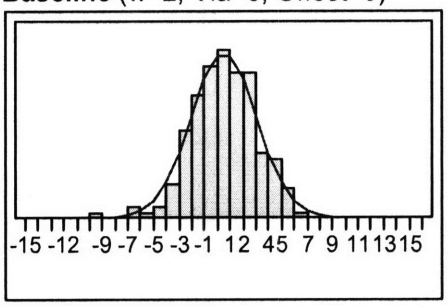

		Prob. outside ±10:
Mean	0.4870954	0.000117
Std Dev	2.5521956	
Std Err Mean	0.1902294	Prob. outside ±10 (centered):
upper 95% Mean	0.8624761	0.000089
lower 95% Mean	0.1117146	
N	180	

With initial offset of 33 mils (w=2, Via=0, Offset=1)

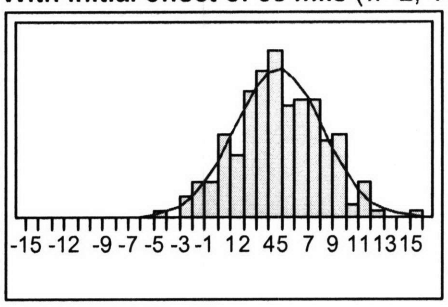

		Prob. outside ±10:
Mean	4.765748	0.061802
Std Dev	3.3991259	
Std Err Mean	0.2533559	Prob. outside ±10 (centered):
upper 95% Mean	5.2656966	N/A
lower 95% Mean	4.2657995	
N	180	

With vias (w=2, Via=1, Offset=0)

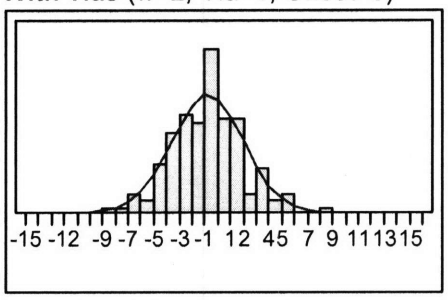

		Prob. outside ±10:
Mean	-0.786308	0.000502
Std Dev	2.775912	
Std Err Mean	0.2069043	Prob. outside ±10 (centered):
upper 95% Mean	-0.378023	0.000315
lower 95% Mean	-1.194593	
N	180	

Figure B-6. Distribution of x after reflow: the effects of different solder pad widths

Baseline (w=2, Via=0, Offset=0)

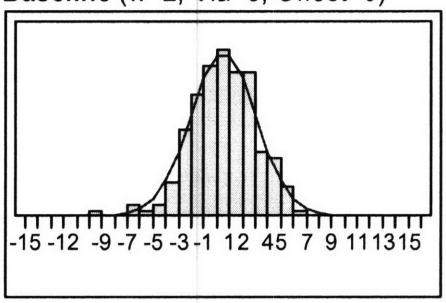

		Prob. outside ±10:
Mean	0.4870954	0.000117
Std Dev	2.5521956	
Std Err Mean	0.1902294	Prob. outside ±10 (centered):
upper 95% Mean	0.8624761	0.000089
lower 95% Mean	0.1117146	
N	180	

With narrower solder pads (w=1, Via=0, Offset=0)

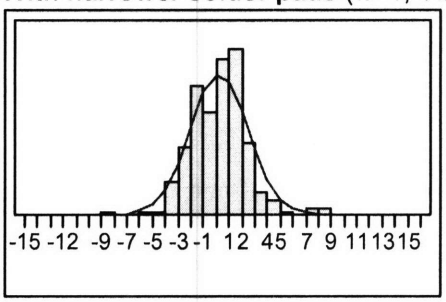

		Prob. outside ±10:
Mean	0.2065992	0.000011
Std Dev	2.2700035	
Std Err Mean	0.1566451	Prob. outside ±10 (centered):
upper 95% Mean	0.5154062	0.000011
lower 95% Mean	-0.102208	
N	210	

With wider solder pads (w=3, Via=0, Offset=0)

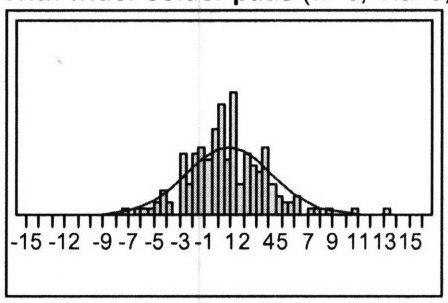

		Prob. outside ±10:
Mean	0.8764217	0.003405
Std Dev	3.3072002	
Std Err Mean	0.2465041	Prob. outside ±10 (centered):
upper 95% Mean	1.3628497	0.002497
lower 95% Mean	0.3899937	
N	180	

Distribution of *y* after reflow

y in mils (thousandths of an inch).

Figure B-7. Distribution of *y* after reflow: design values

Design (w=2, l=2, Full Mask=0, Double=0, Via=0, Offset=0)

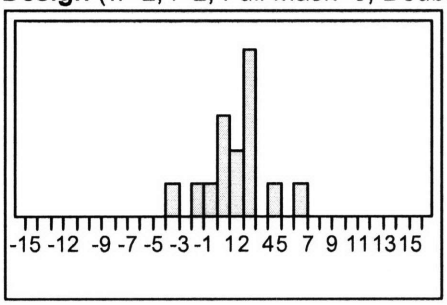

Mean	1.488189
Std Dev	2.3853942
Std Err Mean	0.6159061
upper 95% Mean	2.8091763
lower 95% Mean	0.1672017
N	15

Figure B-8. Distribution of *y* after reflow: the effects of an initial *x* offset and vias

Baseline (w=2, Via=0, Offset=0)

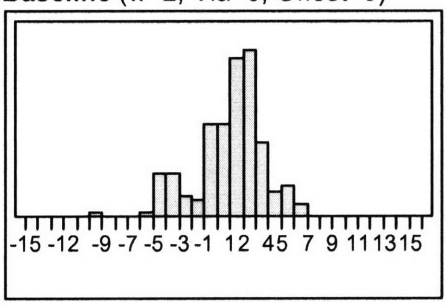

Mean	1.0021872
Std Dev	2.6966222
Std Err Mean	0.2009944
upper 95% Mean	1.3988105
lower 95% Mean	0.605564
N	180

With initial *x* offset of 33 mils (w=2, Via=0, Offset=1)

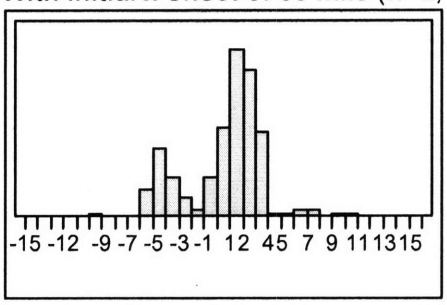

Mean	0.6489501
Std Dev	3.132463
Std Err Mean	0.23348
upper 95% Mean	1.1096775
lower 95% Mean	0.1882228
N	180

With vias (w=2, Via=1, Offset=0)

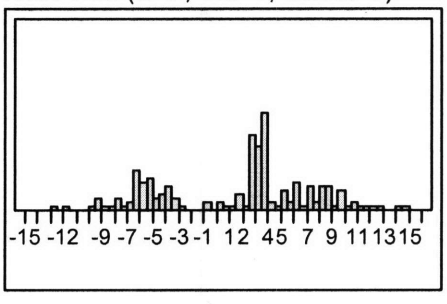

Mean	1.5227472
Std Dev	6.1238974
Std Err Mean	0.4564484
upper 95% Mean	2.4234592
lower 95% Mean	0.6220351
N	180

Figure B-9. Distribution of *y* after reflow: the effects of different solder pad widths

Baseline (w=2, Via=0, Offset=0)

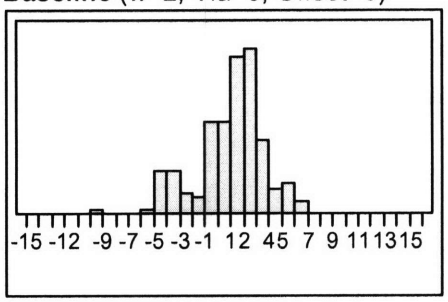

Mean	1.0021872
Std Dev	2.6966222
Std Err Mean	0.2009944
upper 95% Mean	1.3988105
lower 95% Mean	0.605564
N	180

With narrower solder pads (w=1, Via=0, Offset=0)

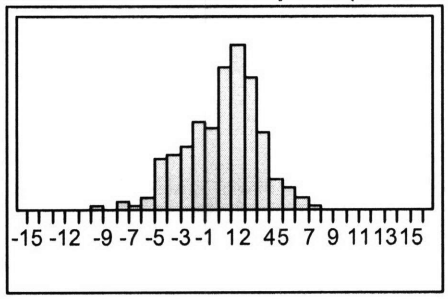

Mean	0.3629546
Std Dev	2.8952377
Std Err Mean	0.1997904
upper 95% Mean	0.7568173
lower 95% Mean	-0.030908
N	210

With wider solder pads (w=3, Via=0, Offset=0)

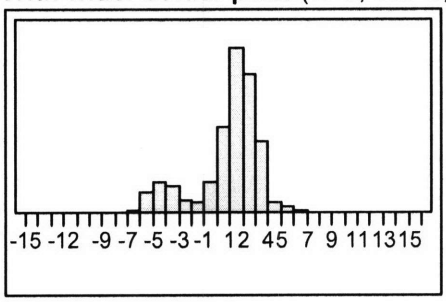

Mean	0.8670166
Std Dev	2.5557621
Std Err Mean	0.1904953
upper 95% Mean	1.242922
lower 95% Mean	0.4911113
N	180

Distribution of θ after reflow

θ in degrees.

Figure B-10. Distribution of θ after reflow: design values

Design (w=2, l=2, Full Mask=0, Double=0, Via=0, Offset=0)

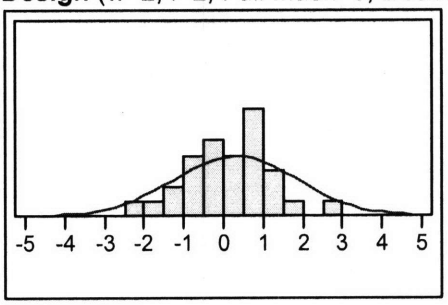

Mean	0.2966667
Std Dev	1.5437266
Std Err Mean	0.2818446
upper 95% Mean	0.8731037
lower 95% Mean	-0.27977
N	30

Figure B-11. Distribution of θ after reflow: the effects of an initial x offset and vias

Baseline (w=2, Via=0, Offset=0)

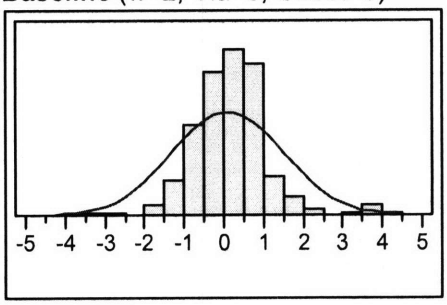

Mean	0.0938889
Std Dev	1.3866951
Std Err Mean	0.1033582
upper 95% Mean	0.2978461
lower 95% Mean	-0.110068
N	180

With initial *x* offset of 33 mils (w=2, Via=0, Offset=1)

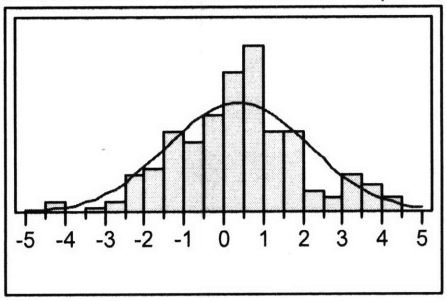

Mean	0.3679775
Std Dev	1.7576427
Std Err Mean	0.1317409
upper 95% Mean	0.6279625
lower 95% Mean	0.1079925
N	178

With vias (w=2, Via=1, Offset=0)

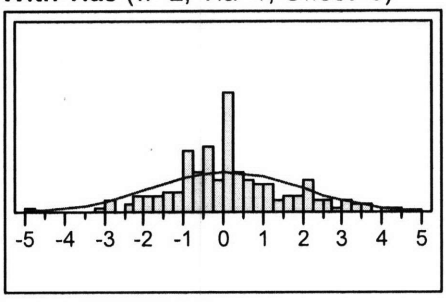

Mean	0.041573
Std Dev	1.8583638
Std Err Mean	0.1392903
upper 95% Mean	0.3164564
lower 95% Mean	-0.23331
N	178

Figure B-12. Distribution of θ after reflow: the effects of different solder pad widths

Baseline (w=2, Via=0, Offset=0)

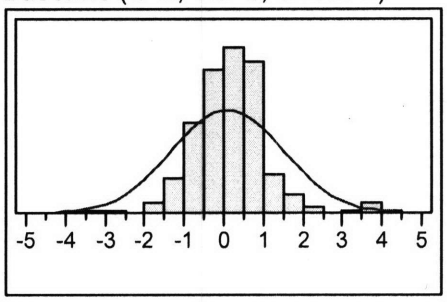

Mean	0.0938889
Std Dev	1.3866951
Std Err Mean	0.1033582
upper 95% Mean	0.2978461
lower 95% Mean	-0.110068
N	180

With narrower solder pads (w=1, Via=0, Offset=0)

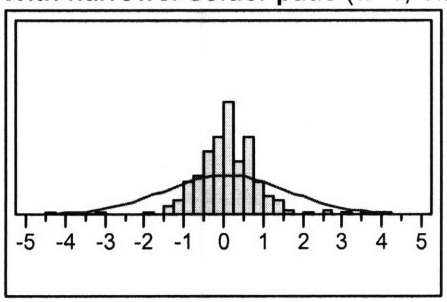

Mean	0.0595238
Std Dev	1.5256122
Std Err Mean	0.1052772
upper 95% Mean	0.2670652
lower 95% Mean	-0.148018
N	210

With wider solder pads (w=3, Via=0, Offset=0)

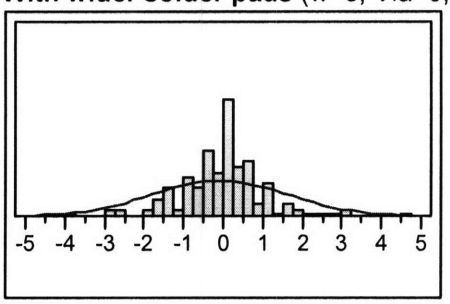

Mean	-0.067039
Std Dev	1.6913453
Std Err Mean	0.1264171
upper 95% Mean	0.18243
lower 95% Mean	-0.316508
N	179

Appendix C: ANOVA of DOE test results

x and *y* in mils (thousandths of an inch); θ in degrees.

Regression analyses were performed on the data from the DOE tests and evaluated using Analysis of Variance (ANOVA) techniques. This appendix includes the results of those analyses.

Model fit for x (after reflow)

Fit to all 7 variables

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	-0.521146	0.362391	-1.44	0.1506
w	**0.7621494**	**0.093238**	**8.17**	**<.0001**
I	-0.131675	0.093238	-1.41	0.1580
AR	0.2988927	0.146083	2.05	0.0409
Full Mask	0.0397248	0.075674	0.52	0.5997
Double	-0.090296	0.096125	-0.94	0.3476
Via	**1.1930668**	**0.07584**	**15.73**	**<.0001**
Offset	**-1.192027**	**0.075674**	**-15.75**	**<.0001**

Summary of Fit

RSquare	0.208375
RSquare Adj	0.20587
Root Mean Square Error	3.565541
Mean of Response	1.162623
Observations (or Sum Wgts)	2220

Analysis of Variance

Source	DF	Sum of Squares	Mean Square	F Ratio
Model	7	7402.215	1057.46	83.1788
Error	2212	28121.343	12.71	Prob > F
C. Total	2219	35523.558		<.0001

Lack Of Fit

Source	DF	Sum of Squares	Mean Square	F Ratio
Lack Of Fit	340	7045.570	20.7223	1.8406
Pure Error	1872	21075.773	11.2584	Prob > F
Total Error	2212	28121.343		<.0001
				Max RSq
				0.4067

Fit to *w*, *via*, and *offset*

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	-0.363808	0.193148	-1.88	0.0598
w	**0.7787039**	**0.092629**	**8.41**	**<.0001**
Via	**1.1899732**	**0.075797**	**15.70**	**<.0001**
Offset	**-1.192027**	**0.075721**	**-15.74**	**<.0001**

Summary of Fit

RSquare	0.205967
RSquare Adj	0.204892
Root Mean Square Error	3.567734
Mean of Response	1.162623
Observations (or Sum Wgts)	2220

Analysis of Variance

Source	DF	Sum of Squares	Mean Square	F Ratio
Model	3	7316.696	2438.90	191.6058
Error	2216	28206.863	12.73	Prob > F
C. Total	2219	35523.558		<.0001

Lack Of Fit

Source	DF	Sum of Squares	Mean Square	F Ratio
Lack Of Fit	8	2072.767	259.096	21.8903
Pure Error	2208	26134.096	11.836	Prob > F
Total Error	2216	28206.863		<.0001
				Max RSq
				0.2643

Model fit for y (after reflow)

Fit to all 7 variables

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	-0.859092	0.46946	-1.83	0.0674
w	0.1280512	0.120786	1.06	0.2892
I	**0.8229758**	**0.120786**	**6.81**	**<.0001**
AR	0.0059566	0.189243	0.03	0.9749
Full Mask	0.0350429	0.098033	0.36	0.7208
Double	**0.3361016**	**0.124526**	**2.70**	**0.0070**
Via	**-0.365337**	**0.098247**	**-3.72**	**0.0002**
Offset	0.174718	0.098033	1.78	0.0748

Summary of Fit

RSquare	0.030222
RSquare Adj	0.027153
Root Mean Square Error	4.618986
Mean of Response	1.026637
Observations (or Sum Wgts)	2220

Analysis of Variance

Source	DF	Sum of Squares	Mean Square	F Ratio
Model	7	1470.708	210.101	9.8477
Error	2212	47193.091	21.335	Prob > F
C. Total	2219	48663.800		<.0001

Lack Of Fit

Source	DF	Sum of Squares	Mean Square	F Ratio
Lack Of Fit	340	10488.092	30.8473	1.5733
Pure Error	1872	36705.000	19.6074	Prob > F
Total Error	2212	47193.091		<.0001
				Max RSq
				0.2457

Model fit for θ (after reflow)

Fit to all 7 variables

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	**0.6325391**	**0.195003**	**3.24**	**0.0012**
w	-0.02981	0.050229	-0.59	0.5529
l	-0.033047	0.050211	-0.66	0.5105
AR	**-0.269455**	**0.078725**	**-3.42**	**0.0006**
Full Mask	0.0863024	0.040738	2.12	0.0342
Double	0.0054137	0.051784	0.10	0.9167
Via	0.0568543	0.040831	1.39	0.1639
Offset	**-0.111002**	**0.040739**	**-2.72**	**0.0065**

Summary of Fit

RSquare	0.014413
RSquare Adj	0.011258
Root Mean Square Error	1.90856
Mean of Response	0.105103
Observations (or Sum Wgts)	2195

Analysis of Variance

Source	DF	Sum of Squares	Mean Square	F Ratio
Model	7	116.4978	16.6425	4.5689
Error	2187	7966.3651	3.6426	Prob > F
C. Total	2194	8082.8629		<.0001

Lack Of Fit

Source	DF	Sum of Squares	Mean Square	F Ratio
Lack Of Fit	340	1408.8883	4.14379	1.1672
Pure Error	1847	6557.4768	3.55034	Prob > F
Total Error	2187	7966.3651		0.0285
				Max RSq
				0.1887

References

Aravamudhan, Srinivasa, Joe Belmonte, and Gerald Pham-Van-Diep. Unpublished. Self-centering of chip components in a Pb-free assembly as a function of component and solder paste print offsets. Speedline Technologies White Paper.

Boning, Duane, Karthik Balakrishnan, Hong Cai, Nigel Drego, Ali Farahanchi, Karen Gettings, Daihyun Lim, Ajay Somani, Hayden Taylor, Daniel Truque, and Xiaolin Xie. 2007. Variation. *IEEE Proceedings: 8th International Symposium on Quality Electronic Design*: 15-20.

Brookner, Eli. 1985. Phased-array radars. *Scientific American* 252(2): 94-102.

Carlson, Gary. 2005. Trusted foundry: the path to advanced SiGe technology. *IEEE Proceedings: Compound Semiconductor Integrated Circuit Symposium (CSIC'05)*: 9-12.

CMMI® for Development, Version 1.2 (CMMI-DEV, V1.2). 2006. Software Engineering Institute, Carnegie Mellon University.

Codd, E. F. 1970. A relational model of data for large shared data banks. *Communications of the ACM* 13(6): 377-387.

Fenn, Alan J., Donald T. Temme, Williams P. Delaney, and William E. Courtney. 2000. The Development of Phased-Array Radar Technology. *Lincoln Laboratory Journal* 12(2): 321-340.

Kaminski, Carl S. 1994. Horizontal R&D Management. *IEEE Proceedings: Engineering Management Conference "Management in Transition: Engineering a Changing World"*: 142-148.

Keller, John. 2006. Defense industry upbeat; military spending to stay healthy over next decade. *Military & Aerospace Electronics*, Nov.

Montgomery, Douglas C. 2001. *Introduction to statistical quality control, 4th Edition.* New York: John Wiley & Sons, Inc.

Neely, Liz. 2007. Raytheon launches ASIC design company. *Electronic News*, http://findarticles.com/p/articles/mi_m0EKF/is_6_48/ai_82571983 (accessed July 2007).

Ottosson, Stig. 1996. Dynamic product development: findings from participating action research in a fast new product development process. *Journal of Engineering Design* 7(2): 151-169.

Padgalskas, Nicholas. 2007. Implementing variation risk management during product development. MS Thesis, Leaders for Manufacturing Program, Massachusetts Institute of Technology.

Raytheon Company. 2006. *Raytheon 2006 Annual Report*, SEC form 10-K.

Schiff, S.C. 2004. Systems engineering of RF system-on-wafer applications in SiGe. *Silicon Monolithic Integrated Circuits in RF Systems, Digest of Papers*: 49-52.

Strauss, Rudolf. 1998. *SMT Soldering Handbook, Second Edition*. Newnes: 248.

Streit, Dwight C., Augusto Guitierrez-Aitken, Michael Wojtowicz, and Richard Lai. 2005. The future of compound semiconductors for aerospace and defense applications. *IEEE Proceedings: Compound Semiconductor Integrated Circuit Symposium (CSIC '05)*: 5-8.

[Unknown author]. 2007. IA in action. *Military Information Technology* 10(1), http://www.military-information-technology.com/article.cfm?DocID=1298 (accessed July 2007).

[Unknown author]. 2007. Defense news top 100. *DefenseNews.com*, http://www.defensenews.com/index.php?S=06top100 (accessed July 2007).