

NAVAL POSTGRADUATE SCHOOL MONTEREY, CALIFORNIA



THESIS

**MINIATURIZATION OF A MICROCONTROLLER
FOR THE
TACTILE SITUATIONAL AWARENESS SYSTEM**

by

Terrence L. Wood

June 1999

Thesis Advisor:

Douglas J. Fouts

Approved for public release; distribution is unlimited.

DTIC QUALITY INSPECTED 4

19990716 028

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188.	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE June 1999	3. REPORT TYPE AND DATES COVERED Master's Thesis		
4. TITLE AND SUBTITLE Miniaturization of a Microcontroller for the Tactile Situational Awareness System		5. FUNDING NUMBERS		
6. AUTHOR(S) Terrence L. Wood				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSORING / MONITORING AGENCY REPORT NUMBER		
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.		12b. DISTRIBUTION CODE		
13. ABSTRACT (maximum 200 words) Loss of Situational Awareness (SA) is a leading cause of pilot-related mishaps, resulting in numerous fatalities and costing the Department of Defense an estimated \$300 million annually in destroyed aircraft. Loss of SA can occur when a pilot incorrectly perceives the attitude, altitude, or motion of their aircraft. As one solution to the SA problem, the Naval Aerospace Medical Research Laboratory has developed the Tactile Situational Awareness System (TSAS). The primary objective of TSAS is to enhance pilot performance and reduce SA-related aircrew/aircraft losses by providing continuous non-visual information using the normally-underutilized sensory channel of touch. Using vibrotactile stimulators, TSAS applies information taken from the aircraft's instruments to the pilot's torso. Prototypes have been built and flight-tested with positive results. However, the current implementation of TSAS is a research system that is not compatible with the crowded cockpits of modern aircraft. This thesis presents a design of a miniature microcontroller for the TSAS that is compatible with tactical environments. This new microcontroller system incorporates the functionality of the research TSAS into a palm-sized device.				
14. SUBJECT TERMS Electronics, Embedded Computing, Microcontroller, Tactile Situational Awareness System, TSAS		15. NUMBER OF PAGES 207		
		16. PRICE CODE		
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

Approved for public release; distribution is unlimited.

**MINIATURIZATION OF A MICROCONTROLLER FOR THE
TACTILE SITUATIONAL AWARENESS SYSTEM**

Terrence L. Wood
Major, Canadian Air Force
B.Eng. (EE), Royal Military College, Kingston, Ontario, 1988


Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN ELECTRICAL ENGINEERING

from the

**NAVAL POSTGRADUATE SCHOOL
June 1999**

Author:




Terrence L. Wood


Approved by:



Douglas J. Fouts, Thesis Advisor



Randy L. Wight, Second Reader



Jeffrey B. Knorr, Chairman

Department of Electrical and Computer Engineering

ABSTRACT

Loss of Situational Awareness (SA) is a leading cause of pilot-related mishaps, resulting in numerous fatalities and costing the Department of Defense an estimated \$300 million annually in destroyed aircraft. Loss of SA can occur when a pilot incorrectly perceives the attitude, altitude, or motion of their aircraft. As one solution to the SA problem, the Naval Aerospace Medical Research Laboratory has developed the Tactile Situational Awareness System (TSAS). The primary objective of TSAS is to enhance pilot performance and reduce SA-related aircrew/aircraft losses by providing continuous non-visual information using the normally-underutilized sensory channel of touch. Using vibrotactile stimulators, TSAS applies information taken from the aircraft's instruments to the pilot's torso. Prototypes have been built and flight-tested with positive results. However, the current implementation of TSAS is a research system that is not compatible with the crowded cockpits of modern aircraft. This thesis presents a design of a miniature microcontroller for the TSAS that is compatible with tactical environments. This new microcontroller system incorporates the functionality of the research TSAS into a palm-sized device.

TABLE OF CONTENTS

I. INTRODUCTION.....	1
A. SITUATIONAL AWARENESS AND SPATIAL DISORIENTATION.....	2
B. THE TACTILE SITUATIONAL AWARENESS SYSTEM	4
C. OTHER POSSIBLE APPLICATIONS FOR TSAS	8
D. GOAL OF THE THESIS	13
E. THESIS OUTLINE.....	13
II. TSAS MICROCONTROLLER DESIGN METHODOLOGY	15
A. TSAS MICROCONTROLLER REQUIREMENTS	15
B. DESIGN METHODOLOGIES	17
1. Commercial-Off-The-Shelf Design.....	17
2. Custom Design.....	18
3. Hybrid Design	19
C. DESIGN METHODOLOGY FOR THE TSAS MICROCONTROLLER.....	19
1. Commercial Off The Shelf Design.....	19
2. Custom Design.....	27
3. Hybrid Design	32
D. DESIGN METHODOLOGY SUMMARY	35
E. DESIGN OF THE TSAS MICROCONTROLLER	36
III. TSAS MICROCONTROLLER CIRCUIT DESIGN.....	39
A. CARDPC INTERFACE.....	39
B. MIL-STD-1553 DATA BUS INTERFACE.....	44
C. TIC INTERFACE.....	55

D. RS-232 INTERFACE	60
E. SOLID STATE DISK CIRCUITRY	63
F. DC TO DC POWER CONVERSION	65
G. INTERNAL SIGNAL GENERATION	69
H. MISCELLANEOUS DESIGN COMMENTS	74
I. SUMMARY.....	75
IV. TSAS MICROCONTROLLER PCB LAYOUT DESIGN.....	77
A. PCB LAYOUT OVERVIEW	77
B. OVERVIEW OF PCB LAYOUT SOFTWARE REQUIREMENTS.....	81
C. PRINTED CIRCUIT BOARD COMPOSITION AND DESIGN CONSTRAINTS	84
1. PCB Layer Stackup Assignment	84
2. Pads	85
3. Traces	87
4. Vias	90
5. Other Considerations	91
D. CARDPC/CPU INTERFACE COMPONENTS.....	92
1. EASI Connector.....	92
2. Capacitors	93
3. Resistors	94
4. PCB Layout.....	94
E. MIL-STD-1553 DATA BUS INTERFACE	94
1. BU-61588/65178 Mini-ACE.....	94
2. 16MHz Clock.....	95
3. TST-9007 Twin Stacked Transformers.....	95
4. MIL-STD-1553 Twinax Connectors.....	96

5. Mini-ACE 8x2 Header	97
6. Capacitors	97
7. PCB Layout.....	97
F. TIC INTERFACE	99
1. TL16C550C UART	99
2. 1MHz Clock.....	99
3. 74HCT08	100
4. Modular Jack.....	100
5. Capacitors	100
6. PCB Layout.....	100
G. RS-232 INTERFACE	101
1. RS-232 Transmitter/Receiver.....	101
2. Capacitors	101
3. 9-pin Subminiature D-Connector	102
4. PCB Layout.....	102
H. SOLID STATE DISK CIRCUITRY	103
1. DiskOnChip 2000.....	103
2. 32-pin "J" Lead Socket.....	104
3. Capacitor	104
4. PCB Layout.....	104
I. DC TO DC POWER CONVERSION	104
1. DC Power Jack	105
2. On/Off Switch.....	105
3. MAX1626 DC-DC Controllers.....	105
4. Current-Sense Resistors.....	105
5. Inductors.....	106
6. External Switching Transistors	106
7. Diodes	106
8. Capacitors	106

9. PCB Layout.....	107
J. INTERNAL SIGNAL GENERATION	108
1. PLD	108
2. 74HCT08	108
3. Capacitors	108
4. PCB Layout.....	108
K. POWER PLANES.....	108
1. Ground Planes.....	109
2. VCC Plane	109
L. SUMMARY.....	109
V. CONCLUSIONS AND RECOMMENDATIONS	111
1. TSAS Microcontroller Design Advantages	111
2. Further Development Efforts.....	112
3. Recommendations	113
APPENDIX A. MOTOROLA MPC860 POWERQUICC OVERVIEW	115
A. INTRODUCTION.....	115
B. METHOD OF APPROACH.....	115
C. RESULTS.....	116
1. Overview.....	116
2. Embedded PowerPC Core (EPPC)	118
3. System Interface Unit (SIU)	130
4. Communications Processor Module (CPM)	132
5. Implementations.....	136
D. SUMMARY	137
E. CONCLUSIONS	138

APPENDIX B. TSAS MICROCONTROLLER SCHEMATICS.....	141
APPENDIX C. TSAS MICROCONTROLLER PLD LOGIC.....	151
APPENDIX D. TSAS MICROCONTROLLER PCB LAYOUT	153
APPENDIX E. TSAS MICROCONTROLLER BILL OF MATERIALS.....	171
LIST OF REFERENCES	173
INITIAL DISTRIBUTION LIST	179

LIST OF FIGURES

Figure 1. Tactile Situational Awareness System.....	5
Figure 2. TSAS in Roll Situation	6
Figure 3. TSAS-equipped UH-60.....	7
Figure 4. Astronauts Performing Space Shuttle EVA	9
Figure 5. RPV and Control Station [Ref. 3].....	11
Figure 6. Basic Mechanical Dimensions, 16-bit PC/104 Module [Ref. 16].	20
Figure 7. How PC/104 Modules Stack [Ref. 16].....	21
Figure 8. Using PC/104 Modules on an Application Baseboard [Ref. 16].....	22
Figure 9. Using PC/104 Modules as "Standalone" Stacks [Ref. 16].....	22
Figure 10. Cell Computing Pentium CardPC [Ref. 17].....	24
Figure 11. Parvus Corp. Complete Scalable CPU [Ref. 18].....	25
Figure 12. SBS Avionics' ABI-PC104 1553 Interface [Ref. 19]	26
Figure 13. Cell Computing Pentium CardPC Block Diagram [Ref. 17].....	33
Figure 14. Typical Data Bus Architecture [Ref. 29].....	46
Figure 15. ACE Block Diagram [Ref. 31]	49
Figure 16. BU-61588/65178 Interface to 1553 Bus [Ref. 31]	54
Figure 17. CardPC ISA Bus 16-bit Memory Read/Write Cycle [Ref. 25].....	57
Figure 18. Harris HIN211 Internal Circuit Overview [Ref. 43]	62
Figure 19. M-Systems DiskOnChip 2000 [Ref. 45]	64
Figure 20. Maxim MAX1626 Typical Operating Circuit [Ref. 48]	68
Figure 21. CardPC Memory Map [Ref. 25]	70
Figure 22. CardPC I/O Map [Ref. 25].....	73
Figure 23. Mounting Technologies: (a) Thru-hole, (b) Surface Mount [Ref. 49].....	78
Figure 24. CardPC Heat Sink [Ref. 25]	81
Figure 25. TSSOP24 Symbol (Not to Scale)	83
Figure 26. Example "Ratsnest"	84
Figure 27. MPC860 Block Diagram [Ref. 68].....	117
Figure 28. Block Diagram of the Embedded PowerPC Core (EPPC) [Ref. 68].....	119

Figure 29. LSU Functional Block Diagram [Ref. 20].....	121
Figure 30. Instruction Cache Organization [Ref. 20].....	124
Figure 31. MMU Flow Diagram [Ref. 20].....	127
Figure 32. Effective-to-Physical Address Translation for 4 Kbyte Pages [Ref. 20]	127
Figure 33. Two-Level Translation Table [Ref. 20].....	128
Figure 34. System Interface Unit Block Diagram [Ref. 68]	131
Figure 35. Communications Processor Module Block Diagram [Ref. 20]	133
Figure 36. Example MPC860 Application [Ref. 68]	138
Figure 37. TSAS Microcontroller Circuit Schematics Index	142
Figure 38. CardPC Interface Circuit Schematic.....	143
Figure 39. MIL-STD-1553 Bus Interface Circuit Schematic	144
Figure 40. TIC Interface Circuit Schematic	145
Figure 41. RS-232 Interface Circuit Schematic	146
Figure 42. Solid State Disk Circuit Schematic.....	147
Figure 43. DC-to-DC Power Conversion Circuit Schematic.....	148
Figure 44. Signal Generation Circuit Schematic.....	149
Figure 45. PLD Logic Schematic	151
Figure 46. TSAS Microcontroller PCB, Showing Both Sides and All Layers of Etch.....	153
Figure 47. TSAS Microcontroller PCB, Top and Bottom Side, Without Etch.....	154
Figure 48. TSAS Microcontroller PCB, Top Side.....	155
Figure 49. TSAS Microcontroller PCB, Bottom Side	156
Figure 50. TSAS Microcontroller, Showing TOP Signal Layer Etch	157
Figure 51. TSAS Microcontroller, Showing S2 Signal Layer Etch.....	158
Figure 52. TSAS Microcontroller, Showing S3 Signal Layer Etch.....	159
Figure 53. TSAS Microcontroller, Showing BOTTOM Signal Layer Etch.....	160
Figure 54. CardPC Interface Component Layout.....	161
Figure 55. MIL-STD-1553 Data Bus Interface Component Layout.....	162
Figure 56. TIC Interface Component Layout	163
Figure 57. RS-232 Interface Component Layout	164
Figure 58. Solid State Disk Component Layout	165

Figure 59. DC-DC Power Conversion Component Layout.....	166
Figure 60. Internal Signal Generation Component Layout.....	167
Figure 61. TSAS Microcontroller Ground Plane Layout.....	168
Figure 62. TSAS Microcontroller VCC3 and VCC5 Planes Layout.....	169
Figure 63. TSAS Microcontroller VCC Split-Plane Layout.....	170

LIST OF TABLES

Table 1. Typical PC/104 Module Specifications [Ref. 16].....	23
Table 2. Cell Computing Pentium CardPC Specifications [Ref. 16].	34
Table 3. I/O Device Memory Mapping	71
Table 4. Recommended Trace Widths for External Conductors [Ref. 54]	88
Table 5. Recommended Trace Widths for Internal Conductors [Ref. 54]	88
Table 6. Before and After Exceptions [Ref. 20].	123
Table 7. Level-1 Descriptor Format [Ref. 20].....	129
Table 8. Level-2 Descriptor Format [Ref. 20].....	130
Table 9. MPC860 Derivatives [Ref. 71].	136

LIST OF SYMBOLS, ACRONYMS AND ABBREVIATIONS

Ω	Ohms
0WS	Zero Wait State
1553	Military Standard Specification 1553 Aircraft Internal Data Bus
ACE	Advanced Communication Engine
ADS	Application Development System
ASIC	Application Specific Integrated Circuit
ATM	Asynchronous Transfer Mode
BALE	Bus Address Latch Enable
BC	Bus Controller
BGA	Ball Grid Array
BIOS	Basic Input/Output System
BM	Bus Monitor
bps	bits per second
CAD	Computer-Aided Design
CD	Carrier Detect
CMOS	Complementary Metal-Oxide Semiconductor
COMx	Communication Port 'x' (e.g. COM2)
COTS	Commercial Off The Shelf
CP	Communications Processor
CPM	Communications Processor Module
CPU	Central Processing Unit
CS	Chip Select
CTS	Clear To Send
CQ	Completion Queue
DC	Direct Current
DCD	Data Carrier Detect
DCE	Data Circuit-terminating Equipment
DDC	Data Device Corporation

DIMM	Dual In-line Memory Module
DIP	Dual In-line Package
DMA	Direct Memory Access
DoD	Department of Defense
DOS	Disk Operating System
DRAM	Dynamic Random Access Memory
DTR	Data Terminal Ready
DSR	Data Set Ready
DTE	Data Terminal Equipment
EA	Effective Address
EASI	Embedded All-in-one System Interface
EEPROM	Electrically Erasable Programmable Read Only Memory
EIA	Electronic Industries Association
EMI	Electromagnetic Interference
EPPC	Embedded PowerPC Core
EPROM	Erasable Programmable Read Only Memory
EVA	Extra Vehicular Activity
F	Farads
FCC	Flight Control Computer
FDD	Floppy Disk Drive
FPGA	Field Programmable Gate Array
FPU	Floating Point Unit
G	Gravity
GCA	Ground Collision Avoidance
GPR	General Purpose Register
GPS	Global Positioning System
H	Henries
HCT	High-speed CMOS, TTL compatible
Hz	Hertz
I/O	Input/Output

PC	Inter-integrated Controller
IBM	International Business Machines
IC	Integrated Circuit
IFF	Identification Friend or Foe
INS	Inertial Navigation System
IOR	I/O Read
IOW	I/O Write
IRQ	Interrupt Request
ISA	Industry Standard Architecture, and Instruction Set Architecture
KB	Keyboard
LED	Light Emitting Diode
LPT	Line Printer
LSB	Least Significant Byte
LSU	Load/Store Unit
MByte	MegaByte
MC	Mission Computer
MCM	Multi-Chip Module
MEMR	Memory Read
MEMW	Memory Write
MHz	MegaHertz
MIL-STD-1553	Military Standard Specification 1553 Aircraft Internal Data Bus
Mini-ACE	Miniature Advanced Communication Engine
MMU	Memory Management Unit
MPC860	Motorola MPC860 PowerQUICC
MS	Mouse
MSB	Most Significant Byte
MSTCLR	Master Clear
MT	Monitor Terminal
NAMRL	Naval Aerospace Medical Research Laboratory
NASA	National Aeronautics and Space Administration

NPS	Naval Postgraduate School
NRE	Non-Recurring Engineering
NRZ	Not Return to Zero
NRZI	Not Return to Zero Inverted
ns	nanosecond
OS	Operating System
PBGA	Plastic BGA
PC	Personal Computer
PC/104	A standard based on the IBM PC
PCB	Printed Circuit Board
PCMCIA	Personal Computer Memory Card International Association
PGA	Pin Grid Array
PLCC	Plastic Leaded Chip Carrier
PLD	Programmable Logic Device
PowerQUICC	PowerPC Quad Integrated Communications Controller
QFP	Quad Flat Pack
QUICC	Quad Integrated Communications Controller
RAM	Random Access Memory
RD/WR	Read/Write
READYD	Ready D
RESETDRV	Reset Drive
RISC	Reduced Instruction Set Computer
ROM	Read Only Memory
RPV	Remotely Piloted Vehicle
RT	Remote Terminal
RTC	Real Time Clock
RTOS	Real Time Operating System
RTS	Request To Send
RWR	Radar Warning Receiver
RXD	Receive Data

SA	Situational Awareness, and System Address
SCC	Serial Communications Controller
SD	Spatial Disorientation
SEAL	Sea, Air, Land
SIMM	Single In-line Memory Module
SINS	Seal Inshore Navigation System
SIU	System Interface Unit
SMEMR	System Memory Read
SMEMW	System Memory Write
SO	Small Outline
SOIC	Small Outline IC
SRAM	Static RAM
SSOP	Shrink Small Outline Package
STRBD	Strobe Data
TDM	Time-Division Multiplexing
TIC	Tactor Interface Chip
TLB	Translation Lookaside Buffer
TSAS	Tactile Situational Awareness System
TSSOP	Thin Shrink Small Outline Package
TST	Twin Stacked Transformers
TTL	Transistor-Transistor Logic
TXD	Transmit Data
UART	Universal Asynchronous Receiver Transmitter
USART	Universal Synchronous-Asynchronous Receiver Transmitter
UV	Ultraviolet
V	Volts
VGA	Video Graphics Array
VLSI	Very Large Scale Integration
VRAM	Video RAM
VSWMCM	Very Shallow Water Mine Countermeasures

ACKNOWLEDGEMENT

I would like to extend my sincerest gratitude to Professor Douglas Fouts for his willingness to sponsor this (currently) unfunded thesis topic, and for his experience and guidance throughout its development. His enthusiasm for research and belief in the concept of tactile communications have been inspirational and instrumental in advancing the technology to the point where it is ready for implementation. I am strongly convinced that others will soon recognize the importance of this work, and provide funding such that this technology is brought out of the lab and into cockpits where it can start saving lives.

I would like to thank Ron Phelps of the Naval Postgraduate School for his prompt resolution of numerous, frustrating Cadence software package "bugs," and for his many tutorials on printed circuit board design.

I thank Randy Wight of the Naval Postgraduate School for his review of this document, and the worthwhile suggestions that he made.

I thank the two previous research students involved with the Tactile Situational Awareness System, LT Brian Luke and LCDR Jeffery Link, for their diligent efforts. Their ground-breaking work has made my task much easier than it would have otherwise been.

Lastly, this thesis would not have been possible without the help and support of my wife, Jody Thomson. I dedicate this work to her.

I. INTRODUCTION

A Canadian F-18 takes off in low cloud. On climb-out, the pilot notices that his G-suit has come unplugged. As he reaches over to plug the suit back in, the aircraft inadvertently and imperceptibly goes inverted. The pilot then receives an aural warning that he has passed through 1000 feet; believing that he is still upright, he instinctively pulls back on the stick. The aircraft and pilot plunge into a lake and are lost.

A US Navy F-14 pilot takes off in low cloud. After a short 55 degree pitch-up at full military power, he cuts the afterburners and pushes the nose of the aircraft down to level off. He proceeds to hold the stick forward until he perceives that the aircraft is level. In fact, the aircraft is in a nose-down attitude – acceleration and gravity forces have combined to provide a false sensation of what “level” is. Both the aircraft and pilot are lost.

The above are but two examples of Spatial Disorientation (SD). SD is defined as the incorrect perception of attitude, altitude, or motion of one's aircraft relative to the earth or other significant objects [Ref. 1]. The Naval Aerospace Medical Research Laboratory (NAMRL) estimates that numerous lives and over \$300 million in destroyed aircraft are lost annually due to SD. Only an estimate of the scope of the problem is possible given that, regrettably, the only person that could explain what really happened is almost invariably lost in the accident. Investigators are left to try to piece together the scenario using flight data recorders, an accident scene, eyewitness accounts, and intuition. This was the case in the two crashes described above.

A. SITUATIONAL AWARENESS AND SPATIAL DISORIENTATION

A loss of Situational Awareness (SA) is defined as the failure of a pilot to adequately know their position in space or what is occurring around them [Ref. 2]. There are three components to SA [Ref. 3]:

1. Spatial orientation of the pilot and their platform relative to the earth.
2. Orientation relative to friendly and enemy forces.
3. Orientation relative to geography and the battlefield environment.

In aviation, the greatest threat to SA is the loss of spatial orientation, termed SD. Interestingly, the SD problem seems to have worsened with the advent of fly-by-wire aircraft. In fact, modern aircraft are becoming so aerodynamically efficient at all airspeeds that the pilot is left with little or no sensory cues (e.g., buffet, cockpit noise) relative to stall speed [Ref. 4].

This lack of physical feedback in fly-by-wire aircraft has been compensated for with increased and improved instrumentation such as the gyro-stabilized artificial horizon, but this has led to its own problems in that the pilot has to continually look inside the cockpit to establish attitude and bearings. In response, the Heads-Up Display (HUD) was invented to allow pilots to keep their heads out of the cockpits a large portion of the time. Additionally, flashing, colored lights and audible warnings have been employed to indicate emergencies.

Why, then, do SD mishaps continue to occur?

The answer to this question requires a slightly deeper understanding of the physiological underpinnings of SD. (A scientific explanation will not be attempted here, but the reader's indulgence is asked in the exegesis that follows.) The human brain's

perception of motion and attitude is the result of the continual interpretation of three independent, redundant, and concordant sensory systems: vision, the vestibular system or inner ear, and the somatosensory system (skin, joint, and muscle sensors) [Ref. 5]. These systems are not perfect: provided with the right inputs the brain can be fooled, and convincingly so, into believing that the body is experiencing certain motions in defiance of physical reality. Virtual reality rides capitalize on this basic fact of human physiology. One may rationalize that the inputs of a virtual reality "ride" are illusionary and therefore attempt to ignore them, but most cannot fight the brain's physiological interpretation.

Man's "normal" environment is on land, and the vestibular and somatosensory systems have adapted to utilize the force of gravity as a means of maintaining spatial orientation. In the aviation environment however, constantly changing acceleration and gravity forces dictate that a pilot often has to rely on vision as the primary means of maintaining spatial orientation. If the visual system is compromised (e.g., clouds, darkness, increased workload, distraction, etc.), the natural reaction is to subconsciously resort to the redundant vestibular and somatosensory systems, which may be in error. The results can be catastrophic.

Loss of SA through SD is widely recognized as a primary cause of aircrew/aircraft losses. For example, of the 15 US Navy aircraft lost to noncombatant action in the Desert Shield/Storm conflict, 7 were SD mishaps [Ref. 6]. The US Navy has placed a very high priority on abating SA, as evidenced by a new Mission Need Statement for systems to enhance aircrew SA [Ref. 2]. Current systems, based on visual displays and aural warnings, are insufficient. So what else can be done to fight the powerful effects of SD?

B. THE TACTILE SITUATIONAL AWARENESS SYSTEM

NAMRL's response to this complex problem is one of "fighting fire with fire" - by providing the pilot with strong somatosensory inputs based on true flight information, it is hoped that misleading inputs will be overcome. Displaying attitude information tactually will also allow the pilot to keep their head out of the cockpit for a greater portion of time, thereby allowing visual horizon references to be maintained.

The sense of touch is an avenue of communication that to this point has remained largely untapped in the area of aviation. Its effects, however, can be very powerful. The gag is old, but illustrates the point: you are immersed in conversation when someone sneaks up behind you, taps you on a shoulder, and moves to the other side. You instinctively turn your head to the direction of the touched shoulder. The point is, while it may be relatively easy to ignore a visual or aural message, when you are touched the offending source usually gets your full attention.

The idea of using the sense of touch as a way of providing feedback to the pilot is not new; in fact, a vibratory display system for improved aircraft flight control (pitch and roll) had been suggested as early as 1954 [Ref. 7]. A paper written by M. Zlotnik of Northrop Corporation in 1988 [Ref. 4] suggested using an array of vibrotactile stimulators, or "tactors" (as commonly found in a vibrating pager) placed around the throttle arm of fighter pilots to convey airspeed information. With each tactor programmed to correspond to a specific airspeed, tactile interpretation would be possible by noting the position of the stimulus relative to the arm.

NAMRL's idea is more sophisticated in that it involves an array of tactors mounted in a vest to be worn around the pilot's torso. It is called the Tactile Situational Awareness System (TSAS). Figure 1 illustrates the concept. Aircraft sensors feed flight information to the aircraft's main computer, which in turn feeds the information to the tactor microcontroller. The microcontroller analyses the flight information, and determines which tactor to vibrate. For example, if the aircraft is in straight and level flight, no sensations are




Figure 1. Tactile Situational Awareness System

felt; in a dive sensations are felt in the stomach, in a climb sensations are felt along the back, and in a roll sensations are felt along the side. Figure 2 depicts a roll situation. In essence, the pilot is presented with a three-dimensional tactile display of attitude information.




Figure 2. TSAS in Roll Situation

It should be noted that the TSAS is not just a concept - prototype systems have been built and flight-tested, with positive results. For example, a test pilot in the back seat of a T-34 was blindfolded to block any visual cues. The test pilot, using only tactile inputs from the TSAS, successfully flew straight and level for five minutes and performed climbing and descending turns, loops and aileron rolls, and ground-controlled approaches. He was also able to recover when placed in unusual attitudes by the front seat pilot. Similarly, a pilot was successfully able to hover a UH-60, Figure 3, using tactile inputs. The test pilots have described TSAS as intuitive. Ref. 6 describes these and other test results.



Figure 3. TSAS-equipped UH-60

The main hardware shortcomings of the current implementation of NAMRL's TSAS are the size of the tactor controller and the size of the cable assembly required to drive the tactor array. Both are too large for the crowded cockpits of modern military aircraft. The Naval Postgraduate School is addressing these shortcomings on two fronts. LCDR Jeffrey Link has just completed a Very Large Scale Integration (VLSI) technology chip design, termed the Tactor Interface Chip (TIC) [Ref. 8]. The TIC essentially creates a "smart" tactor, which can receive and interpret commands from a serial communications bus. Whereas a microcontroller driving tactors without a TIC would require one power line for each tactor in the array, the microcontroller/TIC combination will require only 4 wires: power, ground, clock and data. As the next generation of the TSAS is predicted to have at

least 40 tactors, the TIC will result in a substantial reduction in cable assembly size. As for the size of the tactor controller, this thesis presents the design of a palm-sized microcontroller that can fit into a breast pocket.

C. OTHER POSSIBLE APPLICATIONS FOR TSAS

To this point, discussion has centered around the use of TSAS to display attitude information to a pilot. There are many other possible applications for TSAS. In the aviation environment, any aircraft sensor capable of providing an input that can be processed into a tactile cue could be used, such as [Ref. 1]:

1. Gyroscopes
2. Global Positioning Systems (GPS)
3. Inertial Navigation Systems (INS)
4. Radar Warning Receivers (RWR)
5. Identification Friend or Foe (IFF)
6. Altimeters
7. Range Finders
8. Sonars

For example, TSAS could be used to display enemy aircraft locations for a fighter pilot. If there is a bogey at 5 o'clock high, then a tactor in the vicinity of the upper back and right shoulder would vibrate. Relative proximity to the target could be indicated by the strength and frequency of the pulse (i.e. the closer the target, the stronger and quicker the pulse).

In the space environment, astronauts performing extravehicular activities (EVA) can benefit from tactile feedback. Astronauts performing EVAs have limited natural feedback as to grip strength, resulting in overexertion. In 1986, NASA and the University of Wisconsin investigated providing artificial tactile cues to astronauts to relate grip strength by placing pressure sensors in the gloves and tactile stimulators on the arm [Ref 4]. Tactile feedback can also help astronauts maintain spatial orientation during EVAs. The low gravity of space renders the body's vestibular and somatosensory sensors ineffective, and thus the visual system must be relied upon to maintain spatial orientation. As shown in Figure 4, Space Shuttle EVAs do not present a particular problem for astronauts as they are usually confined to the cargo bay. Here, the astronauts are able to maintain a reference point as they are either tethered to the cargo bay railings or attached to the robotic arm.




Figure 4. Astronauts Performing Space Shuttle EVA

They are also in full view of crew members inside of the shuttle. As pointed out by J. Rochlis in Ref. 9 however, it is expected that International Space Station EVAs will present an SD problem. Due to the size of the space station structure (approximately 110m by 75m), astronauts working on the space station may not be in view of the airlock, intervehicular crew members, or each other; the common reference frame that the shuttle previously afforded has been removed.

Another application for TSAS is in the area of Very Shallow Water Mine Countermeasures (VSWMCM). Navy SEAL (Sea, Air, and Land) divers currently perform VSWMCM with the use of a visual navigation device that directs them through a set sweep pattern. Research was conducted using the SEAL Inshore Navigation System (SINS) as a sensor source for the TSAS. In six dive tests, divers felt that the tactor technology was easier to use and provided better navigation than the visual-only displays [Ref. 10].

Tactile communication also has possible applications in the piloting of Remotely Piloted Vehicles (RPV). Even more drastic than the space environment, the pilot of an RPV (as in Figure 5) has no vestibular or somatosensory inputs as to the attitude of their craft. In this case, tactile inputs can provide an artificial feel of the RPV for the pilot.

The possibilities for TSAS seem almost unlimited. For example, other possible extensions are for silent communications amongst army platoon members and as an alert for air traffic controllers. The reader will have noticed that although a lot of research has been expended in the area of tactile communication as related above, none of the work has resulted in fielded systems. One may rightfully ask why, if tactile communication is such a great idea and has been studied for so long, there are currently no systems in use. In fact,




Figure 5. RPV and Control Station [Ref. 3]

there are a number of systems fielded in the commercial world, as will be discussed in the next paragraph. As for military applications, it is suggested that electronics technology is finally at the point where TSAS can become operationally feasible. Indeed, the TIC developed by Jeffery Link and the microcontroller described in this thesis are a potent combination: they are small, light-weight, low-power, and computationally powerful.

Commercial applications of tactile communication are numerous. Tactile communication devices for the blind, which directly translate spatial-temporal visual

information to the skin, include the Optacon (OPTical-to-TACTile CONversion), the TVSS (Tactile Vision Substitution System), and the Kinotact (KINesthetic, Optical and TACTile display). Users of the Optacon can typically read at a rate of about 50 words per minute. Tactile communication devices for the deaf, which use positional encoding of frequency information, include the Felix and Tactaid VII systems. Virtual reality devices are also using the sense of touch to convey meaning. For example, some devices such as the PHANTOMTM (SensAble Technologies Inc., Cambridge, MA) and ImpulseEngineTM (Immersion Corp., San Jose, CA) monitor the position of a contact point (such as the ball of a joystick), and deliver a force based on the penetration of the point into a virtual object [Ref. 11].

Navigation devices for the blind that could benefit from tactor technology include the MoBIC (Mobility of Blind and elderly people Interfacing with Computers) [Ref. 12] and the Nomad [Ref. 13]. These devices allow the user to do trip planning at home, and then utilize satellite-based navigation systems to guide the user along his intended route of travel. Information is provided to the user through synthesized speech, for example, "The nearest telephone is 50 yards ahead of you on your left." Rather than rely solely on audio messages, which may be interfered with by ambient street noise, a user could be guided directly to a destination through the use of tactor pulses.

In summary, there are many exciting possibilities for the use of tactile communications. A large volume of independent research has concluded that utilizing the sense of touch to convey information is not only feasible, it is effective. The TIC and TSAS microcontroller finally miniaturize the hardware required to support tactile communications

to the extent that it may be comfortably and unobtrusively worn by humans and yet still have the processing power necessary to perform its intended function.

D. GOAL OF THE THESIS

This thesis details a design for a palm-sized microcontroller for the TSAS. LT Brian Luke had already completed a design for a TSAS microcontroller [Ref. 3], and originally the intent of this thesis was to review his design for Electromagnetic Interference (EMI) issues, thermal considerations, cross-talk, coupled capacitance, design flaws etc., such that the design was refined to the point that it was ready for fabrication. For a number of reasons, as will be discussed in the next chapter, it was decided that it would be best if a new design were implemented.

The goal of this thesis will therefore be to present the design and rationale for a new microcontroller for the TSAS. Architectural design and Printed Circuit Board (PCB) layout decisions will be explained in detail.

E. THESIS OUTLINE

The remainder of this thesis is organized as follows: Chapter II discusses various options for developing hardware to meet the requirements of NAMRL's TSAS; Chapter III discusses the architectural design of a microcontroller for the TSAS; Chapter IV discusses the Printed Circuit Board (PCB) layout of this microcontroller; and Chapter V contains conclusions and suggestions for future work.

Five appendices are provided for reference: Appendix A contains an overview of the Motorola MPC860 communications processor; Appendix B contains the circuit schematics of the TSAS microcontroller design; Appendix C details the address decoding and device

selection logic for the Programmable Logic Device (PLD); Appendix D contains PCB layout pictures; and Appendix E contains a parts list with manufacturer information for the TSAS microcontroller as designed.

II. TSAS MICROCONTROLLER DESIGN METHODOLOGY

This chapter presents an overview of the how the TSAS microcontroller detailed in this thesis was designed. The microcontroller requirements will first be presented, followed by various options for meeting the requirements. The pros and cons of each method will be discussed. The rationale for the selected design methodology will be presented, and an explanation of why the TSAS microcontroller developed by LT Brian Luke [Ref. 3] was superceded will be given. The chapter will conclude with a discussion of the steps that were followed in designing the microcontroller.

A. TSAS MICROCONTROLLER REQUIREMENTS

As of the writing of this thesis, the TSAS project is unfunded. All previous monies have been expended building the TSAS prototype. Given that there is no specific customer per se driving requirements, "best guesses" at what potential customers might desire of the TSAS have had to be made. Based on discussions with NAMRL [Ref. 14 and Ref. 15], the following list of requirements was drafted:

1. The microcontroller must be able to interface with the Tactor Interface Chip.
2. The system must be able to control a minimum of 40 tactors independently.
3. The system must activate the tactors using a user-defined voltage and frequency.
4. The size of the microcontroller should be as small as possible, preferably palm-sized.
5. The system should be optimized for low-power operation and should be capable of running off of a battery or a plug-in DC power source.
6. The microcontroller must be able to communicate with the MIL-STD-1553 aircraft data bus.

7. The system must have a standard RS-232 and/or RS-422 interface to allow communications with other serial navigation devices and computers.
8. The microcontroller must be capable of quickly processing Euler angle transformations for 3-D vectoring purposes.
9. The system must be capable of being programmed in a high-level language such as C++ or ADA.
10. The microcontroller should have as much memory as is feasible.
11. The system must not generate so much heat as to be uncomfortable to wear.

It is quickly noticed that many of the above requirements are in direct conflict with each other. For example, the requirements that the microcontroller be as small as possible and yet have as much memory as possible, and the requirements that the system be low power and that the CPU be powerful. Immediate questions come to mind; for example, just how much memory is enough? The requirements are broadly stated however, so as not to be too restrictive on the use of the new technology and the microcontroller design. As discussed in Chapter I, there are many possible applications for TSAS, and some have greater processing demands than others. The palm-sized microcontroller presented in this thesis has been designed to be capable of easily handling the most demanding TSAS application (that of EVAs, as discussed at Ref. 15). It will therefore be over-designed for some TSAS applications; however, as will be discussed later, the microcontroller is scalable. That is to say, if more (or less) memory or CPU processing power is required then different modules can be inserted, all within the same form factor. In this way, the microcontroller can accommodate different requirements based on different applications, while allowing for the lowest power consumption and cost possible.

B. DESIGN METHODOLOGIES

Given the requirements for the TSAS, the task is now to develop hardware that will meet these requirements. There are three fundamental methodologies that can be followed:

1. Use a Commercial-Off-The-Shelf (COTS) product that meets the specifications;
2. Build a custom system from the ground up using VLSI technology and/or basic components; or
3. Use a hybrid of Methods 1 and 2 (e.g., utilize commercial Multi-Chip Modules (MCM) or other pre-packaged functions along with basic components/VLSI).

There are, of course, some gray areas in these three broad categories. For example, for the purposes of this discussion a slight modification of a COTS product will be considered to still fall under design method 1.

There are pros and cons to each of the above methods. They will be discussed in turn.

1. Commercial-Off-The-Shelf Design

There are many compelling reasons to go with a COTS product if there is already something available that meets the requirements, almost meets the requirements, or meets them with slight modification. In fact, this is now the preferred DoD method of systems procurement. By purchasing an already-existing product, Non-Recurring Engineering (NRE) costs are avoided, which can be a large percentage of overall cost if the volume of units desired is small. If the COTS product is mature, then a lot of the bugs have been worked out already. One can therefore generally field the product much quicker than is possible with a custom design. Technical support is usually available. To top it all off, the

customer can benefit from the competitive nature of the commercial electronics industry: over time, the product will likely become more capable, while decreasing in cost.

The disadvantage in going with a COTS product is that it is usually not exactly what is desired. There are likely some additional features in the system that have been paid for, but can not be utilized. Besides cost, these extra features can add unwanted weight, size, and power consumption.

2. Custom Design

By contrast, there is usually very little excess in the custom design – one can design a device to contain only those features that are required. Depending on the technology used, this can result in a substantial reduction in size and weight. Utilizing VLSI will result in the greatest reduction. Design with discrete components and Integrated Circuits (IC) can still result in some excess though (e.g., if only two out of four inverters in a 7404 package are utilized). Also, there is not the reduction in the number and size of ICs that is possible with VLSI. However, this is usually far better than having complete functional blocks go unutilized, as can be the case when using a COTS product.

The disadvantages of going with a custom design are the up-front NRE costs, the retaining fees for engineer(s) capable of maintaining the system, the time it can take to get the product fielded, and the risk of outright failure due to design complexity. These costs are usually higher for VLSI design than IC design. In either case, the costs can be significant. However, if size and weight are of critical importance to the success of a design, there may be no other choice than to go with a custom design.

3. Hybrid Design

A hybrid design involves using a COTS device to perform some primary function(s) and then adding application-specific functionality through custom design. Naturally, the hybrid design gains from the benefits of both the COTS and custom design methods but also suffers from their disadvantages. Often this methodology is chosen if it is desired to take advantage of the benefits of COTS but a product that exactly suits the requirements cannot be found; one therefore selects the product that fulfills the largest subset of the requirements and then custom-designs the rest of the functionality.

C. DESIGN METHODOLOGY FOR THE TSAS MICROCONTROLLER

This section will consider each of the design methodologies discussed above as applied to the TSAS microcontroller, including a discussion of pros and cons. The rationale for the selected design method will then be presented.

1. Commercial Off The Shelf Design

Not surprisingly, there are no existing COTS products that will exactly fulfill the requirements of a microcontroller for the TSAS. This is due in large part to the unique interface needs of the custom-built TIC. The TIC requires a 1 MHz clock signal in addition to clocked serial data in Universal Asynchronous Receiver Transmitter (UART, discussed in Chapter III) form. However, after extensive research, many products were found that came close to meeting the requirements of the TSAS microcontroller. Virtually all of these products are based on the PC/104 (PC – Personal Computer, 104 – the number of interface pins) standard, so this standard will be briefly discussed now.

Over the past decade, the PC architecture has become an increasingly popular choice for embedded controllers. Utilizing standardized hardware and software around the broadly-supported PC architecture reduces development costs, risk, and time to market for manufacturers, analogous to the COTS design methodology discussed above. PCs are being used as controllers for vending machines, laboratory instruments, and communications devices, to name a few examples. The standard PC bus form factor (12.4" x 4.8") is too large for most embedded control applications, thus a smaller standard, the PC/104, was adopted. This standard is supported by the PC/104 Consortium, a group of over 160 manufacturers. PC/104 offers full architecture, hardware, and software compatibility with the PC bus in compact (3.6" x 3.8", Figure 6) stackable modules [Ref. 16]. A depiction of how the modules stack together is shown in Figure 7.




Figure 6. Basic Mechanical Dimensions, 16-bit PC/104 Module [Ref. 16].




Figure 7. How PC/104 Modules Stack [Ref. 16].

The PC/104 concept is quite simple: the user simply selects a PC/104 single board computer based on their processing requirements, and then selects additional PC/104 modules as required to support the functionality of their particular device. PC/104 single board computers are available with processors ranging from i386's to Pentiums with MMX. Clock speeds of over 200 MHz are supported. A wide variety of PC/104 modules are available, offering functions such as hard drives, floppy drives, video displays, modems, audio, ethernet, power conversion, etc. The user has two options for layout, depending on the form factor desired: the CPU (Central Processing Unit) module and expansion modules can be laid out in a flat configuration, as per Figure 8, or they can be stacked, as per Figure 9.




Figure 8. Using PC/104 Modules on an Application Baseboard [Ref. 16].




Figure 9. Using PC/104 Modules as "Standalone" Stacks [Ref. 16].

The following list details some of the environmental specifications that are typical of many of the PC/104 modules available on the market today:

Factors Critical to Mobile & Portable Applications	Typical Characteristics of PC/104 Modules
Size	3.550 x 3.775 x 0.6 in.
Weight	2 to 3.5 oz.*
Power consumption	1 to 5W*
Shock	50G 3-axis peak (per MIL-STD-202F, Method 213B, Table 213-1, Condition A)
Vibration	11.95G 3-axis RMS at 100-1000Hz (per MIL-STD-202F, Method 214A, Table 214-I, Condition D)
Operating temperature	0 to +70°C, standard -40 to +85°C, extended
Storage temperature	-55 to +85°C
Humidity	5 to 95%, non-condensing
EMI compliance	EN 55022 Class B (radiated & conducted emissions)
EMC and ESD compliance	IEC 801-2 (electrostatic susceptibility) IEC 801-3 (E/M field susceptibility) IEC 801-4 (fast transient susceptibility)
MTBF	ground mobile, at 55°C: 30,000 to 70,000 hrs* ground fixed, at 55°C: 150,000 to 650,000 hrs* (per MIL-HDBK-217)
* These values vary according to the specific module.	

Table 1. Typical PC/104 Module Specifications [Ref. 16].

These are impressive specifications and are likely well-suited to any TSAS application. It should be noted that many manufacturers have fielded products for the aerospace industry. As utilizing COTS PC/104 modules appeared to be a desirable method of designing a TSAS microcontroller, and is in keeping with the DoD-preferred method of systems acquisition, this option was researched thoroughly. In short, there are some very good options that would require only minor hardware additions. An example of a possible design is presented next.

There are numerous manufacturers that market PC/104 CPU modules built around a specific processor. Parvus Corporation has a unique idea in that they employ COTS themselves in the design of some their PC/104 CPU modules. Parvus utilizes a so-called CardPC, as pictured in Figure 10, as the heart of their CPU module. CardPC's, currently made by three different manufacturers, provide the essentials of a PC architecture: the CPU, ROM (Read Only Memory), RAM (Random Access Memory), an I/O (Input/Output) controller, a graphics controller, a keyboard controller, etc. Pentium processors are




Figure 10. Cell Computing Pentium CardPC [Ref. 17]

available at clock speeds up to 233 MHz. Parvus then adds in the necessary headers, connectors, discrete components, etc. that are required to provide a functioning system. The advantage for the user is that the CardPC's are swappable - if processing needs change, then a faster or slower CardPC can be inserted instead of redesigning the whole device. A

picture of Parvus's Complete Scalable CPU is shown in Figure 11 (the CardPC is mounted on the back of the board). It should be noted that RS-232 interface capability is provided by the CardPC.




Figure 11. Parvus Corp. Complete Scalable CPU [Ref. 18]

The TSAS microcontroller is also required to interface with the MIL-STD-1553 aircraft data bus. There are a few manufacturers that produce PC/104 expansions modules for just this purpose. Perhaps the best known is SBS Avionics Technologies. Shown in Figure 12 is SBS's ABI-PC104 1553 Interface. This device features one dual redundant 1553 channel along with Bus Controller (BC), Remote Terminal (RT), and Bus Monitor (BM) capability. Software drivers and libraries are provided at no additional cost [Ref. 19].

With the RS-232 and MIL-STD-1553 interfaces taken care of, all that really remains to be designed is a DC-DC power conversion circuit (for battery operation) and the TIC




Figure 12. SBS Avionics' ABI-PC104 1553 Interface [Ref. 19]

interface. Unfortunately, there is no room on either of the two PC/104 modules to place additional circuitry. Thus, a third (custom-designed) PC/104 module would be required. The total stack height of the modules (before encasing) would be approximately 2 inches. There are a number of manufacturers that provide protective cases for PC/104 module stacks (e.g., refer to Ref. 18), so this is quickly taken care of as well.

In summary, COTS appears to be a viable solution for the design of the TSAS microcontroller. As there is only relatively minor custom design required, a complete hardware design could be readied in a short period of time.

The disadvantages of employing this design method: as PC/104 CPU modules are intended to exactly replicate the functionality of larger PC motherboards, there is a lot of functionality built-in that the TSAS microcontroller does not require. For example, a large percentage of the board space is used by headers for video displays, floppy drives, parallel

and serial ports, etc. As there is no room left over on either the CPU modules or MIL-STD-1553 modules for designing DC-DC power supplies or a TIC interface, a third PC/104 module is necessitated. A possibility might be to ask the CPU module manufacturer to depopulate the board of functions that are not required, but presumably this would add extra expense and time to the design.

2. Custom Design

As discussed in section B.2 above, there are a few different ways the TSAS microcontroller can be custom-designed: VLSI technology, discrete components plus ICs, or perhaps a mixture of the two. In previous work, LT Brian Luke designed a microcontroller for the TSAS utilizing the custom design methodology [Ref. 3]. His work will be used to illustrate several points in this section.

It should be mentioned at the outset that the following discussion is in no way meant to discredit the work performed by Luke. He undertook a large, complex task with the best of intentions, and produced a good design given the time he was allocated. The points mentioned below are intended to illustrate some of the disadvantages of the custom design method and are not meant to be criticisms of his work. On the contrary, the work that Luke performed has proved invaluable to the current work - without his contributions, the design of the TSAS microcontroller would not be as far along as it is.

The approach that Luke took, that of custom-designing the TSAS microcontroller, appears to be logical given the mutually exclusive requirements of building a controller that is as powerful as possible and yet as small as possible. His work has been classified as a custom design because he designed everything to the IC and discrete component level,

including the main processing function. To reduce overall controller size and simplify the design, Luke chose to employ a couple of key ICs that employ MCM philosophy: the Motorola MPC860 PowerQUICC (Power PC Quad Integrated Communications Controller), and the ILC Data Device Corporation BU-61588 Mini-ACE (Advanced Communication Engine). The utilization of these devices is not sufficient to warrant the classification of Luke's work as a hybrid design however, given the complexity of the remaining design. These devices and other aspects of his design will now be discussed.

For the main processor of his microcontroller Luke chose the MPC860 PowerQUICC. Specifically designed for communications and networking applications, the MPC860 actually has two processors: a 32-bit PowerPC and a 32-bit RISC (Reduced Instruction Set Computing) processor. The purpose of the 32-bit RISC processor is to offload peripheral tasks from the PowerPC core. Thus, the chip reportedly consumes less power overall. For example, the core is interrupted only upon frame reception or transmission, rather than on a per-character basis [Ref. 20].

The MPC860 is a powerful device and is well-suited to communications. For example, its Serial Communications Controllers (SCC) are capable of transmitting and receiving UART protocol and its System Interface Unit allows generation of clock signals for external devices. (An overview of the MPC860's capabilities is presented in Appendix A.) In terms of the TSAS microcontroller's interface with the TIC, these features negate the requirement for two ICs: a UART, and an oscillator. The MPC860 also has a number of other desirable features, to include glue-less interfacing to many types of memory and

peripherals, easy-to-understand PowerPC architecture, and solid customer support such as example system designs (e.g., Ref. 21).

However, the MPC860 does have some limitations, particularly as applied to the TSAS microcontroller. Unlike the majority of PowerPC chips being produced today, which are superscalar, the MPC860 is single-issue. It also has no Floating Point Unit (FPU) and thus does not support any PowerPC floating point instructions. Lastly, the MPC860 currently has a top clock speed of only 66 MHz. While these performance specifications may be fine for most communications applications, the TSAS microcontroller is required to perform Euler angle transformations using sine and cosine functions [Ref. 15]. Although floating point instructions could be emulated in software, doing so would cost additional clock cycles. As the current prototype TSAS controller, an i486SX25 (which has an FPU), can take up to 5 minutes to perform some EVA calculations [Ref. 15], doubts about the MPC860's suitability for this task have arisen. (As an aside, it should be mentioned that none of the software for the TSAS has been developed by the Naval Postgraduate School. Therefore, the requirement for floating point calculations was not known to NPS until discussions at Ref. 15).

To interface with the MIL-STD-1553 aircraft data bus, Luke chose the ILC Data Device Corporation's BU-61588 Mini-ACE. This is also a very powerful chip that employs MCM philosophy by integrating transceivers, protocol, memory management, processor interface logic, and RAM into one package. In one square inch, the chip performs much of the basic functionality of the PC/104 module pictured in Figure 12. Also required for a complete interface are an external transformer, and twinaxial cables to tie into the 1553 data

bus. The ACE User's Manual [Ref. 22] issues a strong warning on page 248 that the transformer and ACE be located as physically close to each other as possible. Doing so limits the voltage drops in the analog signal traces when transmitting and helps minimize cross-talk from other signals on the board. In Luke's design, the transformer and ACE are separated by over 2 inches (with the MPC860 in between), which in all likelihood would cause noise problems. Also, the 10-pin connector utilized is inappropriate; a pair of twinaxial cables, or four coaxial cables, are required.

A major factor in the printed circuit board layout of any design is, of course, the physical dimensions of the included components. In Luke's design, a large number of ICs are utilized, the majority of which are not included in NPS's Cadence (a proprietary Computer-Aided Design (CAD) tool) libraries. The physical dimension information on these parts therefore has to be manually entered into the computer, a tedious and error-prone task. As one might expect, there were some errors in some of the parts created. For example, the Atmel AT27C800 8 Mega-bit UV (Ultra Violet) EPROMs (Erasable Programmable Read Only Memory) are in actuality about 50% larger [Ref. 23] than depicted in Figure 60 of Ref. 3. Given the density of Luke's design, this unfortunate oversight meant that the properly-dimensioned devices cannot be "squeezed" into the existing layout; significant rearrangement of components is necessary.

The TIC interface portion of Luke's design is based on an RS-232 standard: data from an MPC860 SCC is fed into an RS-232 transceiver, which in turn leads to a 9-pin header. The TIC (whose design was not finalized until 6 months after the completion of Luke's design) however, requires 4 inputs: a 5V power line, a ground, a 5V data line, and a

1MHz clock signal. The TIC is not capable of converting the 12V received from the RS-232 transceiver to 5V, and is not provided a clock signal with which to synchronize the input data and drive the TIC circuitry.

One of the requirements for the TSAS microcontroller is that it should be capable of running off of a battery or a plug-in DC power source (see section A of this chapter). Luke's design satisfies the latter portion of this requirement in that it will work with a regulated 5V DC source. It cannot run from a battery source however, as batteries tend to dip significantly in voltage as they are discharged (some of the 5V IC devices have a strict $5V \pm 0.25V$ tolerance). With the addition of a second DC-DC converter, the TSAS microcontroller would be capable of running from a battery, in addition to a plug-in DC source. Additionally, Luke chose to use a linear regulator as this was what was suggested by Ref. 21. As astutely pointed out by Prof. D. Fouts, a linear regulator is much less efficient and therefore produces much more heat than a switching regulator. Switching regulator circuitry is a little more complex and requires more board space but in light of the benefits the additional costs appear to be worth the price.

There are a number of other issues with the custom design of the TSAS microcontroller that will not be mentioned here. The above design points are intended to illustrate the point that, with a custom design, the increased complexity often results in oversights. These oversights can be difficult to identify and can wind up dramatically delaying the fielding of the design or jeopardize its implementation altogether. When compounded with the fact that this was LT Luke's first board design, and was produced utilizing an advanced CAD tool (Cadence's Allegro), the mechanics of which are not taught

by NPS, the design should probably be considered "high risk" (even with the rectification of known errors). It should also be pointed out that one of the primary reasons for building a custom design is the possibility of a smaller form factor than is possible with a COTS or hybrid solution. However, as the CardPC vividly demonstrates, it is very difficult to keep pace with the competitive electronics industry. Experienced engineers in numerous companies are constantly employing the latest in laser etching techniques, surface mount devices, etc. to reduce the size of existing products. If the custom design will not afford a size advantage in our situation, then a COTS or hybrid solution should probably be pursued.

The decision to go with a completely new design was not the result of any particular fault of Luke's design, but was rather an accumulation of smaller issues that required re-designing. This was not a decision that was taken lightly, as clearly a lot of effort had been expended in the original design, work that would have to be repeated in any new design. It was felt, however, that it would require less effort to produce a new design (albeit utilizing a different design methodology) than it would be to redesign Luke's work. Additionally, a new design allows the introduction of improved capabilities for the microcontroller.

3. Hybrid Design

While conducting research into the feasibility of utilizing a COTS product for the TSAS microcontroller, it occurred to the author that use could be made out of a CardPC in much the same way that Parvus uses it for their PC/104 CPU module (Figure 11). A block diagram of Cell Computing's Pentium Card PC is shown in Figure 13. As can be seen, the CardPC provides all of the basic essentials of an IBM (International Business Machines) PC. Depending on user requirements, 32, 64, or 128 Mbytes of RAM can be inserted in the

DIMM (Dual In-line Memory Module). A list of the CardPC's specifications for some of the different models available is shown in Table 2.




Figure 13. Cell Computing Pentium CardPC Block Diagram [Ref. 17]

ITEM	P166 CardPC	P200 CardPC	P233 CardPC
CPU	Pentium 166MHz w/ MMX	Pentium 200MHz w/ MMX	Pentium 233MHz w/MMX
CHIPSET	Intel 430TX	Intel 430TX	Intel 430TX
POWER			
Power Supply	3.3 & 5V	3.3 & 5V	3.3 & 5V
Operating – Observed Max ¹	6.6 W	7.2W	TBD
DOS Prompt	5.1W	5.6W	TBD
Windows 95 ²	4.2W	4.5W	TBD
Suspend To RAM Mode	20mW	20mW	TBD
OPERATING TEMP			
Ambient (T _A)	0-50°C	0-50°C	0-50°C
Case (T _C)	0-85°C	0-85°C	0-85°C
MEMORY			
DIMM	16, 32, 64, 128MB ³	16, 32, 64, 128MB	16, 32, 64, 128MB
Flash ROM	512KB	512KB	512KB
DISPLAY			
Graphics Controller	Trident Cyber9320	Trident Cyber9320	Trident Cyber9320
Bus Interface	PCI	PCI	PCI
VRAM	1MB	1MB	1MB
Max CRT Resolution	1024x768, 256 col.	1024x768, 256 col.	1024x768, 256 col.
Max LCD Resolution	800x600, 16-bit col.	800x600, 16-bit col.	800x600, 16-bit col.
PERIPHERAL I/O			
Keyboard	Mitsubishi M38802	Mitsubishi M38802	Mitsubishi M38802
FDD	National PC 87334	National PC 87334	National PC 87334
IDE	National PC 87334	National PC 87334	National PC 87334
Serial Port	National PC 87334	National PC 87334	National PC 87334
Parallel Port	National PC 87334	National PC 87334	National PC 87334
PHYSICAL			
Height (without fan/ with fan)	12.7mm / 21.5mm	12.7mm / 21.5mm	12.7mm / 21.5mm
Footprint	54mm x 85.6mm	54mm x 85.6mm	54mm x 85.6mm
Weight	55g	55g	55g
ENVIRONMENTAL			
MTBF	100,000 hrs	100,000 hrs	100,000 hrs
Vibration	15G	15G	15G
Shock	50G	50G	50G
Electrostatic Resistance	15kV, 100pF, 1.5kΩ	15kV, 100pF, 1.5kΩ	15kV, 100pF, 1.5kΩ
Connector Insertion/Removal	1000 cycles	1000 cycles	1000 cycles
DIMM Insertion/Removal	10 cycles	10 cycles	10 cycles
Bend	2kg, one side, 1min	2kg, one side, 1min	2kg, one side, 1min
Drop	75cm drop	75cm drop	75cm drop
Twist	12.6kg/cm, 5minx5cyc	12.6kg/cm, 5minx5cyc	12.6kg/cm, 5minx5cyc

Table 2. Cell Computing Pentium CardPC Specifications [Ref. 16].

It is evident that, as is typical with COTS products, much of the functionality provided by the CardPC will not be utilized by the TSAS microcontroller. The difference here, however, as compared to the Parvus Scalable CPU PC/104 module, is that the unused functions will not be finished (i.e., they will not be furnished with headers/interconnects and supporting circuitry). With the CardPC as the heart of the microcontroller, all that remains is to design the required I/O peripheral capabilities and power generation. As will be shown in Chapter IV, the resulting design is approximately half the size of the COTS design discussed in section C.1 above (approx. the same width and length but about half the height).

D. DESIGN METHODOLOGY SUMMARY

In summary, the above discussion has illustrated a few different possibilities for designing the TSAS microcontroller. Theoretically, the custom design should lead to a product that is smaller than that available in the marketplace. It is difficult, however, to keep pace with a competitive commercial industry that is constantly employing the latest in circuit-reduction technology. In addition, the custom design is time-consuming and prone to errors. With minor additions, there is a viable COTS solution for the TSAS microcontroller, albeit at the cost of a larger-than-required form factor. However, it is believed that in this case, the optimum solution is a hybrid design based around a CardPC. Therefore, this is the design methodology that was followed in developing the TSAS microcontroller presented in this thesis. The resulting design will be shown to be significantly smaller than both the COTS solution and the custom design of Ref. 3, and yet still be every bit as capable.

E. DESIGN OF THE TSAS MICROCONTROLLER

This section is intended to briefly chronicle the steps that were followed in the development of the TSAS microcontroller presented in this thesis. It is hoped that doing so will provide insight into why certain paths were followed, in (the hopefully likely) case the TSAS project becomes funded and this design is implemented.

As mentioned in Chapter I, the original intent of this thesis was to review the TSAS microcontroller design of LT Brian Luke; the goal was to ensure that the design was perfected to the point that it was ready for production. Thus, the research began with a thorough study of the Motorola MPC860, the processor of choice for the Ref. 3 TSAS microcontroller design. This study culminated in the production of Appendix A.

Next ensued a study of printed circuit board design. To be honest, the NPS library is lacking in material on the subject. Fortunately, there is a wide range of information now available on the Internet. Somewhat concurrently, a study of the Ref. 3 thesis was conducted. It was at this time that many questions about the requirements of the TSAS microcontroller surfaced: What sort of processing capability was required? What sort of calculations would be performed? Was a Floating Point Unit required? Just how much memory was needed? NAMRL provided answers where possible (Ref. 15) but without a specific customer they could not give definitive specifications. The lack of hard specifications was quite disturbing at first, as it makes it much more difficult to verify whether the design will satisfy its intended requirements. In the end, the TSAS microcontroller came to be seen as being similar to a PC; it is capable of performing many different tasks and can satisfy a potentially wide array of users.

While reviewing the design, it was also necessary to use the CAD software to examine details not published in the thesis. It was at this point that it became clear just how much work had been put into Luke's thesis and how much work would be required to make all the necessary changes. The Cadence Allegro tool is a sophisticated package with a multitude of options to facilitate the process of laying out a printed circuit board for the experienced designer. It is not, however, a package that simplifies the task of laying out a board for the beginner. There are literally thousands of pages of documentation with no example designs or tutorials to demonstrate how to effectively use the tool. NPS does not teach the use of Allegro in any of its courses and there is only one individual on campus, Mr. Ron Phelps, that is knowledgeable in its use. Thankfully, Mr. Phelps was able to dedicate the time necessary to ensure the TSAS microcontroller design's success.

Once it became clear that wholesale changes to the Ref. 3 design would be required, an intensive search for possible alternatives was initiated. As discussed above, a number of viable solutions were found. After deciding to use a hybrid design with a CardPC, selection of devices with which to complete the interface and power functions was required. The Internet was invaluable in this regard. Almost all major electronics manufacturers now have comprehensive web sites, offering everything from catalogs of their products to data sheets and technical support. The task of finding out who makes a particular type of product is now greatly simplified – in most cases, a Web search will quickly point to several manufacturers. Once a manufacturer is selected, one can download their data sheet on the product and begin design immediately - no longer does one have to wait for post mail, only to discover that the device will not meet the requirements.

After the selection of all devices, the circuit schematics were designed and the components laid out on a printed circuit board as will be discussed in the next two chapters.

III. TSAS MICROCONTROLLER CIRCUIT DESIGN

This chapter will delineate the details and rationale behind the circuitry of the new TSAS microcontroller. The treatment is comprehensive; while this does not make for interesting reading, it should greatly facilitate the review of the design for those who follow on with the work. There are many physical packaging issues to be discussed, but these will be treated in the next chapter. The discussion here will center on why certain devices were selected over others and why they are wired up as they are. The interface to the CardPC will be discussed first, followed by the peripheral interfaces, the solid state disk, the DC-DC power conversion circuitry, and the generation of internal signals (to include the PLD address decoding and device selection logic).

In the following discussion, “*” will denote an active-low signal; for example, “MEMR*” means that MEMR (Memory Read) is asserted low.

A. CARDPC INTERFACE

There are currently three different manufacturers of CardPCs. The TSAS microcontroller design is based on a Cell Computing CardPC. Cell Computing was chosen because they had the most comprehensive information available on the Internet and they are the only CardPC reseller based in the United States (the others are in Europe and Japan). Therefore, “CardPC” in the discussion below will refer specifically to Cell Computing’s CardPC. In theory, however, the different CardPCs should be interchangeable. This adds robustness to the design in that if the chosen manufacturer ceases support of the CardPC, there are still likely to be other sources available.

The CardPC squeezes all of the functions of an IBM PC/AT motherboard into a credit card-sized package measuring 2.125 inches by 3.375 inches. The CardPC is billed as the world's smallest Pentium-based motherboard [Ref. 17]. It contains Intel's Pentium processor with MMX, Intel's 430TX core logic chipset, a 144-pin DIMM socket, a graphics controller, 1MByte of Video RAM, 512KBytes of Flash ROM, a keyboard and mouse controller, and National Semiconductor's PC87334VJG Super I/O chip. The graphics controller and Video RAM, keyboard and mouse controller, and many functions of the Super I/O chip are not likely to be used in any TSAS application; however, they will undoubtedly be useful for system development and maintenance.

Of special note is the fact that the Cell Computing CardPC's flash memory can be updated using a DOS (Disk Operating System) utility program, thereby simplifying the job of reprogramming systems already out in the field.

It should be stressed at this point that the CardPC is not a single board computer – it requires external power supplies, headers, connectors, and discrete components to be supplied by the user. Cell Computing does market a single board computer, the Mighty Mite, that utilizes their CardPC. The 4 inch by 5 ¾ inch Mighty Mite, however, is no more suitable for use as a TSAS microcontroller than the PC/104 CPU modules discussed in Chapter II: unusable headers consume a large percentage of the overall board space. Notwithstanding, information provided by Cell Computing engineers on component selection and use in the Mighty Mite design has proven to be valuable in the design of the TSAS microcontroller [Ref. 24].

Access to the many functions pictured in Figure 13 is via a 236-pin Embedded All-in-one System Interface (EASI) connector, which plugs into an EASI connector socket mounted on the target PCB. This connector is represented in the circuit schematic of Figure 38. The function of each of the pins is clearly described in Cell Computing literature [Ref. 25], and will not be repeated here. Although many of the pins are not used, the majority of them are not required to be terminated high or low, as the CardPC takes care of this internally. There are some exceptions however, as detailed in Appendix C of Ref. 25. These exceptions are discussed next.

Referring to the Figure 38 schematic, the inputs of the Floppy Disk Drive (FDD), mouse (MS) and keyboard (KB) functions are pulled high, as they are not required for the TSAS microcontroller. Similarly, pins of the parallel port (LPT, which stands for Line Printer) and the second serial communications port (COM2) are tied either high or low as appropriate. As specified on page 107 of Ref. 26, it is possible to use a single pull-up or pull-down resistor to serve multiple unused inputs. This fact has been taken advantage of to save board space. The $200\text{K}\Omega$ resistor used for COM2 DTR and COM2 RTS may seem excessive, but as per the note at the bottom of page 93 of Ref. 25, these signals also serve as reset configuration signals for the internal controller. Therefore, a resistor in this range must be used to prevent interference with the reset function of these signals. As a final note on the pull-up/pull-down resistors, the pins as depicted in the Figure 38 body diagram are not necessarily adjacent (the small numbers near the pins are the actual physical pin numbers). Pairs and trios of pins that were physically adjacent to each other were selected

to share a resistor in order to cut down on unnecessary wiring. This explains the crossed circuit lines in the schematic.

The power connections for the CardPC are shown at the bottom of the Figure 38 body diagram. The Pentium CardPC has three 3.3V power planes:

1. VCC3 - Power supply for CPU and Chipset.
2. STRVCC3 - Power supply for VRAM, DRAM, VGA & DRAM controllers.
3. RSMVCC3 - Power supply for Suspend Control Logic.

The necessity for the different 3.3V power planes is for support of the CardPC's various power management functions (STR in item 2 above stands for "Suspend To RAM"). As the TSAS controller is required to process information in real time, it was decided that these modes would be of no utility. Based on information provided by Cell Computing engineers, the three planes were tied together.

In addition to the 3.3V power planes, there is a 5V plane and RTCVCC (RTC – Real Time Clock). RTCVCC is the backup power supply system for the RTC and CMOS (Complementary Metal-Oxide Semiconductor) RAM. As backup power is not provided in the design (to be discussed later in the power section), RTCVCC is connected to the 5V power supply.

There are a number of capacitors used in the Figure 38 schematic. These capacitors are used for bypassing, a technique for reducing unwanted noise, timing errors, and Electromagnetic Interference (EMI). As this technique is used throughout the design, it will be discussed here. The following explanation borrows heavily from Ref. 27.

In high-frequency circuits, noise from current transients caused by rapid switching (i.e., fast rise times) can be a significant source of problems. These current transients flow

on traces and planes. Due to the finite inductance of these traces and planes, a current pulse with a fast rise time flowing through the inductance creates a voltage transient. These voltage transients can radiate from any antenna-like source to create EMI problems, and can be beyond the noise margins of logic devices.

As an example of the problems that can occur, consider a Schottky Transistor-Transistor Logic (TTL) device, which typically has a rise time between 2 to 3 nanoseconds (ns). If the device switches logic levels and requires current as a result, it is supplied from the power or ground planes. Due to the inductance associated with these planes, they simply cannot respond quickly enough. Some of the possible results of this situation are:

1. Timing errors. The chip may simply slow down, switching only as fast as the charge can be supplied. The resulting logic-state uncertainty may cause logic errors and timing problems.
2. Noise Errors. If the current level is not sufficient, output levels may fall out of specification and noise margins may significantly decrease, resulting in logic level uncertainty and logic errors.
3. EMI radiation. Brief voltage transients caused by current flow through inductance can cause noise spikes in adjacent traces, or radiate through antennae.

To summarize the problem, there may be chips that require additional electron flow in a very short period of time and sources (power or ground planes) that cannot supply electron flow as quickly as is required. A solution is to provide a supply of charge close to where it is needed, in the form of a bypass capacitor. There are some problems with this approach, however. Capacitors large enough to supply sufficient charge may have too much lead inductance to be able to respond quickly enough and capacitors with low enough

lead inductance to respond quickly may not be able to supply enough charge. Therefore, more than one size and type of bypass capacitor may be required. This is the reason for the two values of capacitance shown in Figure 38. Although grouped together on the schematic, the capacitors are physically spread along the EASI connector.

As a final comment on the Figure 38 schematic, it will be noticed that a number of the pins on the upper-right side of the body diagram have been given signal names. The uppermost group is composed of Industry Standard Architecture (ISA) bus pins, which are used to connect peripherals to the Pentium. The lower group, COM1, is a serial communications port that will be used to develop an RS-232 interface. The function of these pins will be explained as they are used in later sections.

B. MIL-STD-1553 DATA BUS INTERFACE

MIL-STD-1553 describes a standard for a high-speed (1Mbps, bits per second) serial data bus. Since its inception in 1973, MIL-STD-1553 has evolved into the predominant, internationally accepted networking standard for the integration of military platforms. Not used solely for US Airforce and Navy aircraft, this standard now encompasses applications for tanks, ships, satellites, missiles, and the International Space Station [Ref. 28]. In the aircraft application, data from aircraft sensors is communicated to various on-board computers via the MIL-STD-1553. Ref. 29, SBS Avionics' *An Interpretation of MIL-STD-1553B*, provides an excellent overview of the 1553 data bus and is available on the Internet. A very brief explanation of the 1553, derived from this reference, is now provided.

MIL-STD-1553 defines requirements for digital, command/response, time division multiplexing techniques for a 1MHz serial data bus. It also specifies the data bus, its interface electronics, and the concept of operation and information flow on the bus. A typical multiplex data bus architecture is shown in Figure 14. Some definitions of terms are necessary:

1. Data bus – all the hardware necessary to provide a single data path between the bus controller and all the associated remote terminals, to include screened twisted pair cables, isolation resistors, transformers, etc.
2. Command/Response – operation of a data bus system such that remote terminals receive and transmit data only when commanded to do so by the bus controller.
3. Bus Controller (BC) – the terminal assigned the task of initiating information transfers on the data bus.
4. Bus Monitor (MT) – the terminal assigned the task of listening to bus traffic and extracting selected information to be used at a later time.
5. Remote Terminal (RT) – all terminals not operating as the bus controller or as a bus monitor.

The data bus functions asynchronously in a command/response mode. The data is encoded in Manchester II bi-phase level format, which consists of a self-clocking waveform appropriate for use on short transformer-coupled local buses. Sole control of information transmission on the bus resides with the Bus Controller, which initiates all transmissions. Information flow is comprised of messages, which are in turn composed of three types of words: command, data, and status. Each Remote Terminal is assigned to a unique address.




Figure 14. Typical Data Bus Architecture [Ref. 29]

In an aircraft, typically the Flight Control Computer (FCC) or Mission Computer (MC) would be the Bus Controller. All other systems, such as the Radar Warning Receiver, would be Remote Terminals. Clearly, if the TSAS microcontroller is to display flight information to a user, it must have access to the information provided by the aircraft sensors. This information can be requested from the 1553 Bus Controller through a MIL-STD-1553 interface. Such an interface has been designed for the TSAS microcontroller, and is described now.

In his Ref. 3 design, LT Luke elected to utilize the ILC Data Device Corporation (DDC) BU-61588 Miniature Advanced Communication Engine (Mini-ACE) to realize the 1553 interface. After research, this appears to be a solid choice. DDC bills it as the smallest integrated MIL-STD-1553 BC/RT/MT in the industry, and true to the advertising, nothing else seems to come close. The compactness comes at a price, however - a DDC

sales representative quoted each BU-61588 to be approximately \$550 when purchased in small quantities. The good news is, the same chip was about \$770 a year and a half ago when priced by LT Luke. Hopefully this price trend will continue. Nevertheless, the chip appears to be worth the cost; a non-integrated interface solution would require a much larger PC board and more wiring. The use of multiple components is also likely to increase the number of routing layers the TSAS microcontroller requires [Ref. 28].

It should be stressed that although the BU-61588 Mini-ACE is capable of functioning as a Bus Controller, there is NO intent to have it perform this function (for example, as a redundant system should the main BC fail). Doing so would greatly increase the integration engineering and flight testing that the TSAS microcontroller would have to be subjected to and could compromise the whole project. There is another version of the Mini-ACE, the BU-65178, that does not have BC capability. However, it also does not have MT capability. MT capability may be required depending on the software algorithm chosen to cull information from the 1553 data bus (e.g., monitor the flow of information and react when a relevant piece of data is passed, or continually poll for updated information). At any rate, the BU-65178 comes in exactly the same form factor and pin-out as the BU-61588 and the two devices may be used interchangeably in the TSAS microcontroller design. The BU-65178 is approximately 10% less expensive than the BU-61588.

The specifics of how the Mini-ACE is integrated into the TSAS microcontroller design will now be covered. It should be noted that the user's manual for the ACE, Ref. 30, is 500 pages long. An attempt to summarize this manual will not be made here. It is

expected that those interested in a comprehensive understanding of the following description of the Figure 39 schematic will have a copy of the user's manual to refer to (it is available on the Internet).

The ACE comprises a complete integrated interface between a host processor (i.e., the CardPC Pentium) and a MIL-STD-1553 bus. A block diagram of the ACE is shown in Figure 15. The processor and memory interface logic block contains the internal address latches and bi-directional data buffers to provide a direct interface to a host processor bus. This is a feature very useful to the TSAS microcontroller design, as the CardPC data and address lines are 3.3V, while those of the Mini-ACE are 5V. The RT mode of the Mini-ACE provides three data structures for buffering data. These structures, combined with the Mini-ACE's interrupt capabilities, serve to ensure data consistency while offloading the host processor. The Mini-ACE appears as a memory device to the host processor – data pulled from the 1553 data bus by the Mini-ACE is placed in the shared RAM where it can be read by the host processor. Up to 128KBytes of RAM is offered in the Mini-ACE but the TSAS microcontroller design assumes the base 8KBytes (4K x 16bits) of RAM. This is assumed to be more than sufficient for the needs of the TSAS microcontroller, and also allows the Mini-ACE to be memory mapped in the lower 1MByte of memory (i.e., system memory).

The Mini-ACE can be driven at two clock speeds - 12MHz and 16MHz. For the fastest possible throughput, a 16MHz clock has been selected, as pictured in Figure 39. As will be discussed later, this clock speed marries extremely well with the requirements of the TIC interface circuit.




Figure 15. ACE Block Diagram [Ref. 31]

The Mini-ACE requires some external logic in order to interface with the CardPC. This logic has been implemented in the PLD, a schematic of which is shown in Appendix C. The logic was derived from Refs. 30 and 33. The various inputs and outputs of the Mini-ACE will now be discussed individually, starting at the upper left of Figure 39:

1. **IRQ10*.** Two Interrupt Request (IRQ) lines are required for the TSAS microcontroller design: one for the MIL-STD-1553 interface and one for the TIC interface. It was decided to give the 1553 interface the higher-priority IRQ, as time-consuming processing of 1553 data bus information is required before it can be displayed. Ref. 25, p.12, lists IRQ 7, 10, 11, and 15 as being unused. According to Ref. 32, the relative priority of these lines is 10, 11, 15, and then 7. Thus, IRQ 10 was chosen as the IRQ line for the Mini-ACE.

2. RESETDR*. The (active-high) signal RESETDR (Reset Drive) is an ISA bus reset signal used to reset or initialize ISA bus devices at power-up time. As the MSTCLR (Master Clear) input of the Mini-ACE is active-low, the RESETDR signal is inverted in the PLD logic, Figure 45, and then connected directly to MSTCLR.
3. SD<15..0>. The 3.3V 16-bit System Data (SD) bus can be connected directly to the Mini-ACE's 5V data lines, due to built-in transceivers.
4. READYD*. READYD* (Ready Data) is a Mini-ACE handshaking output to the host processor. During a read access, it signals that data is available to be read on D15 through D0 when asserted. During a write cycle, it signals that data has been transferred to a register or RAM location. This signal is used to derive the CardPC's OWS* (Zero Wait State) signal input, as the Figure 45 schematic shows. OWS* indicates to the ISA bus controller within the CardPC that the present ISA bus cycle can be completed without inserting any additional wait states.
5. MEMW*. This signal stands for Memory Write, and is asserted by the ISA bus controller when a memory write operation is taking place on the ISA bus. This in turn instructs the Mini-ACE to accept data off the ISA bus.
6. SA13. This input to the MEM/REG* pin of the Mini-ACE selects whether this is a memory access (active high) or a register access (active low). Exactly why SA13 has been connected will be discussed in the memory-mapping portion of the PLD section.

7. STRBD*. Strobe Data – this input into the Mini-ACE is used in conjunction with SELECT* to initiate and control the data transfer between the host processor and the ACE. As indicated in Figure 45, it is asserted whenever a memory read or memory write operation is taking place on the ISA bus.
8. ACE_CS*. This active low Chip Select signal is the result of address decoding and other circuitry as depicted in Figure 45. When asserted, the Mini-ACE is enabled.
9. SA<12..1>. These 12 lines of the System Address bus are required to access the Mini-ACE's registers and 8Kbytes of RAM. The ACE is a 16-bit machine, thus SA0 is not required ($2^{12} = 4K$ of 16-bit words).
10. TRANS/BUFF*. The Transparent/Buffered* input to the ACE is used to select between the different modes available for host processor interfacing. The buffered mode is most commonly used and is selected here by tying the signal to ground. (The transparent, or Direct Memory Access (DMA) mode, requires a DMA controller.)
11. 16/8*_BIT. This input allows the ACE to be used with 8 bit host processors. As a 16-bit ISA bus is being employed, this line is tied high.
12. MSB_LSB. Most Significant Byte/Least Significant Byte. This input to the ACE is only used in 8-bit buffered mode and can be left unconnected [Ref. 30, p.323].

13. ZERO_WAIT*. This input is used to select between the zero wait state mode (ZERO_WAIT* = 0) and the nonzero wait state mode (ZERO_WAIT* = 1).
As per Ref. 33, this signal has been tied high.
14. TRIGGER_SEL. Trigger Select. This signal has no operation in the 16-bit buffered mode (which is the TSAS microcontroller case) and can be left unconnected [Ref. 30, p.240].
15. POLARITY_SEL. In 16-bit buffered mode, the Polarity Select input signal is used to control the logic sense of the RD/WR* signal. It is pulled high here so that a logic 0 on RD/WR* signifies a write operation, in congruence with the ISA bus MEMW* signal.
16. RTAD4 – RTAD0. Remote Terminal Address inputs. These five lines allow the selection of a unique RT address for the ACE when interfacing with a 1553 data bus (all RTs connected to a 1553 bus must have unique addresses).
As the addresses of the other RTs in any particular application are unknown and will vary, RTAD4 – RTAD0 have been connected to a header so that the RT address of the ACE may be easily changed. These pins, as well as RTADP (described below), are internally pulled up to 5V [Ref. 30, p.164].
One half of the header is therefore connected to ground.
17. RTADP. Remote Terminal Address Parity. This signal must provide odd parity with RTAD4 – RTAD0 in order for the RT to respond to non-broadcast commands. This signal is also attached to the header for easy manual selection.

18. RT_AD_LAT. RT Address Latch. This input is internally hard-wired in the BU-65178/61588 and should be left unconnected.
19. TX_INH_A and B. Transmitter Inhibit, Channels A and B. The ACE has dual redundant 1553 data bus interface channels. This allows the switching of channels if the error rate on the one being used exceeds a specified value. Depending on the application, a redundant channel may not be required or desired. Flexibility to inhibit either the A or B channels has therefore been provided by connecting these lines to a header. Note that the signals are active high. Therefore, they must be connected to ground through the use of the header if it is desired to utilize them.

The TX_RX_A/B and TX_RX_A*/B* (Transmit and Receive, Channels A and B) pins are used to interface with the MIL-STD-1553 data bus. Data is encoded in digital Manchester II bi-phase (negative and positive voltages) form. Stubbing is the method by which a separate cable is connected between the primary data bus and a terminal. The preferred method of connecting stubs to the data bus is to use transformer coupling (vice direct coupling), as this method provides increased DC isolation, increased common mode rejection, and less mix-matching to the data bus [Ref. 29, p. 41]. As well, in a direct-coupled terminal the main bus is not protected against a short circuit in the stub cabling, whereas transformer coupling protects the bus against such a fault. Figure 16 shows the two possible methods of completing the interface. The preferred transformer-coupled method is implemented in the TSAS microcontroller design.




Figure 16. BU-61588/65178 Interface to 1553 Bus [Ref. 31]

Beta Transformer Technology Corporation, a subsidiary of ILC DDC, manufactures transformers specifically for use with the ACE family of products. The required 1:1.79 turns ratio can be supplied by the TST-9007 (TST – Twin Stacked Transformers). This device also saves space in that it has two transformers stacked on top of each other. Primary and secondary turns-ratio pins have been selected in accordance with the manufacturer's data sheet [Ref. 34]. As per the note on p.245 of Ref. 30, the transformers' center taps on the ACE side have been grounded so that only one transformer leg is driven at any point in time.

For the tap into the 1553 bus itself, a senior MIL-STD-1553 specialist at ILC DDC suggested utilizing twinaxial (twinax) cables and connectors [Ref. 33]. A single twinax cable allows for the shielded transmission of both a negative and a positive voltage. Coaxial (coax) cables could also be used but twice as many are required. To provide the

TSAS microcontroller with the greatest robustness, it was decided to provide both A and B channel capability. Twinax cables are employed to cut the connectors required to two, vice the four coax connectors that would have been required.

C. TIC INTERFACE

The Tactor Interface Chip (TIC) has specific microcontroller interface requirements. First, it requires a 5V power line and a ground line. These can be provided by the TSAS microcontroller without problem. The TIC design also specifies a 1MHz clock signal for clocking its internal circuitry, as well as for synchronizing input serial data. The serial data itself must be in UART format and must meet the clock speed (i.e., 1Mbps). These last two requirements require special treatment, and will be discussed here.

The National Semiconductor PC87334 Super I/O chip on the CardPC supports two PC16550 UART buffered serial ports. The 16550 UART is very commonly used for serial communications in PCs. Regrettably, as is the case with most 16550s, the output of the Super I/O chip's 16550 UARTs is limited to 115.2 Kbps [Ref. 25, p.44]. In addition, there is no easy way to synchronize the serial output data of the Super I/O. It was, therefore, decided that an external UART should be provided.

There are a number of different types of UARTs. The original UART chip shipped with the IBM PC was the 8250, which had a one byte buffer and was limited to 9600 baud (bits per second). It was replaced by the 16450, which had the same architecture but a higher maximum baud specification. The 16550, 16650, and 16750 UARTs have 16, 32, and 64 byte buffers respectively. The larger buffers allow for higher baud rates – the devices can store more data before their interrupts have to be serviced [Ref. 35]. For the

TSAS microcontroller, the smallest-sized UART that would do the job was desired. As TIC command lengths are typically 4 bytes [Ref. 8, p.10 and p.85], the buffer size of the 16550 should be more than adequate.

There are a few high-performance 16550 UARTs on the market. One that was capable of 1Mbps utilizing a 16MHz input clock was preferred, so that the same clock that drives the Mini-ACE could also drive the UART. Most UARTs do not fit these requirements. For example, the National PC16550D is capable of 750Kbps or 1.5Mbps, but not 1Mbps, and requires a 24MHz external clock to achieve these transfer rates. After some searching, a suitable UART was found, the Texas Instruments TL16C550C Asynchronous Communications Element (ACE). There is potential for confusion between this acronym and that of ILC DDC's ACE. Therefore, this device will simply be referred to as the TL16C550C. This UART is capable of 1Mbps with a 16MHz input clock. Figure 40 shows how the TL16C550C was utilized to design a TIC interface. The details of the interconnections shown in this schematic will now be covered.

The TL16C550C is an 8-bit device, thus its data pins are connected to the lower 8 system data lines (SD<7..0>). The first three system address lines are connected to A2-A0 to allow the selection of which UART register to read from or write to. These address lines are latched through the use of ADS* (Address Strobe). When ADS* is active (i.e., low), A2-A0, CS2*, CS1 and CS0 drive the TL16C550C's internal select logic directly. When ADS* is high, the register select and chip select signals are held at the logic levels they were in when the low-to-high transition of ADS* occurred [Ref. 36]. To strobe the UART, the ISA bus signal BALE (Buffered Address Latch Enable) is used. BALE is an ISA bus signal

used to latch valid addresses and memory decodes during CPU transfers. Devices should latch the System Address bus on the falling edge of BALE [Ref. 37]. As depicted in the ISA-bus 16-bit memory read/write cycle timing diagram shown in Figure 17, BALE goes low after the System Address lines are valid and remains low for the rest of the cycle. Thus, BALE is suitable for strobing the TL16C550C.




Figure 17. CardPC ISA Bus 16-bit Memory Read/Write Cycle [Ref. 25]

Of the three chip select lines on the TL16C550C, CS1 and CS0 are tied to SA5 and SA3 respectively. The reasons for these connections will be explained in the internal logic

generation section below. The third chip select signal, CS2*, is generated by the PLD. The memory-mapping encoded in the PLD will also be covered in the same section.

One of the free CardPC interrupt request lines, IRQ7, has been tied to the INTRPT (Interrupt) pin of the TL16C550C. As discussed earlier, IRQ7 has a lower priority than IRQ10, which is dedicated to the BU-61588 Mini-ACE.

The TL16C550C has two sets of Read and Write lines: RD1* and WR1*, and RD2, and WR2. The purpose of the two sets, one active low and the other active high, is to facilitate interfacing with the host CPU's Read and Write signals. Thus, IOR* (I/O Read) and IOW* (I/O Write) may be connected directly, without requiring inversion. The TIC does not transmit any data to the TSAS microcontroller, it only receives data [Ref. 8]. Thus, only the SOUT (Serial Out) pin of the UART is connected to the TIC.

During the initial phase of the TIC interface design, it was thought best to use the 3.3V (vice 5V) version of the TL16C550C due to its lower power consumption. To accomplish this, the 3.3V SOUT signal was boosted to 5V through the use of a 5V 74HCT08 AND gate (which would act as a buffer) before connection to the TIC data line. A HCT (High-speed CMOS, TTL compatible) device was selected due to the high switching speed required and the CMOS technology involved on both sides of the gate. An AND gate was selected because there are 4 AND gates per 74HCT08 package, and three other AND gates were required elsewhere in the design. One input of the AND was simply pulled high.

Upon further research, it was decided to use the 5V version of the TL16C550C for two reasons. First, the HCT gate can only output 4.5V vice the required 5V. This may not

be critical for the eventual factor selected but it was decided to error on the side of caution. Second, the typical current consumption of the TL16C550C is so small (10mA) that the difference in the power consumed by the two versions of the device is almost negligible. There may be an issue with the 3.3V ISA bus not having enough voltage to drive the 5V TL16C550C, although it is not believed that there will be a problem since the UART is a CMOS device. If this turns out to be a problem, the original plan may be reverted to, as there is still one AND gate available for use in the 74HCT08 package.

The TIC does not require any handshaking from the UART because it utilizes the externally-provided 1MHz clock signal to synchronize the data. Note that during testing of the TIC interface, a slower oscillator can be installed. The transmission speed of the TL16C550C can be programmed to match many different clock speeds [Ref. 36].

Given that the UART does not receive any data, and given that the TIC does not require any handshaking, none of the TL16C550C's modem lines are required. They are looped back onto each other in what is termed a "null modem" configuration. The aim is to make the UART believe that it is talking to another UART. The way this works is as follows [Ref. 38]: Data Terminal Ready (DTR) is looped back to Data Set Ready (DSR) and Data Carrier Detect (DCD). Thus, when DTR is asserted, DSR and DCD immediately become asserted. This serves to make the UART believe that the virtual UART to which it is connected is ready and has detected the carrier signal of the real UART. The Request To Send (RTS) and Clear To Send (CTS) lines are also connected. When the UART wishes to send data it asserts the RTS and immediately gets a reply through the asserted CTS that it is okay to do so. The Ring Indicator (RI) line can be left unconnected.

The XIN (External Clock) input is tied to the 16MHz clock output of the Figure 39 MIL-STD-1553 interface circuit. The BAUDOUT output is connected to the RCLK (Receiver Clock) input. Although the receiver section is not used, this keeps RCLK from floating.

D. RS-232 INTERFACE

The Electronics Industries Association (EIA) RS-232 standard is the most common asynchronous serial line standard [Ref. 39]. The TSAS microcontroller has been equipped with an RS-232 interface to allow communications with other serial devices (e.g., a modem, keyboard, or Global Positioning System (GPS)). The Figure 41 schematic shows how this interface is implemented. Before discussing this circuit, some background information on the RS-232 standard will be provided.

RS (Recommended Standard) -232 has been around for many years. New versions of the standard have been released over time, with RS-232E being the latest. (RS-232E is equivalent to the International Telecommunications Union's V.28 standard.) RS-232 specifies the electrical interface between Data Terminal Equipment (DTE), such as computer terminals, and Data Circuit-terminating Equipment (DCE), such as modems. It defines a number of parameters concerning voltage levels, loading characteristics, and timing relationships. The essential feature of RS-232 is that all data signals are carried as single voltages referred to a common ground [Ref. 40]. RS-232 supports simplex, half-duplex, and full-duplex type channels. In a simplex channel, data only travels in one direction (this is the case in the TSAS microcontroller/TIC interface). In a half-duplex channel, data may travel in either direction but at any given time the data can only travel in

one direction. In a full-duplex channel, data may travel in both directions simultaneously [Ref. 41]. Half-duplex and full-duplex channels require the use of handshaking lines, which RS-232 provides for. RS-232 specifies the gender and pin use of connectors but not their physical type. 25 and 9-pin D-type connectors are common. DTE should have a male connector, as the TSAS microcontroller does, and DCE a female connector [Ref. 39].

Lastly and most importantly for the TSAS microcontroller RS-232 interface, RS-232 signals have higher voltage levels than TTL. This is to allow for the reliable transmission of data over a long distance (up to 100 feet). Therefore, the TTL-level data from UARTs must be converted to RS-232 levels by transceivers. Note that RS-232 transceivers have an inverting action. Thus, a TTL logic high is translated to between -3 to -12V and a TTL logic low is translated to between 3V to 12V [Ref. 42].

The CardPC's National PC87334 Super T/O chip provides two 16550 UART serial communications ports. One of these ports, COM1, has been used for the RS-232 interface. The data and handshaking lines necessary for an RS-232 interface are provided by the UART; what remains is to convert the voltage levels to the RS-232 standard and to provide a connector. Figure 41 shows the selection of the Harris HIN211 RS-232 Transmitter/Receiver (i.e., transceiver) for voltage conversion. The Maxim MAX211 is an identical alternative. Requiring only a 5V external supply, this device has onboard charge pump voltage converters which generate +10V and -10V. The HIN211 also features a low power shutdown mode to conserve energy in battery-powered applications [Ref. 43]. The internal circuit overview of the HIN211, Figure 18, will facilitate the understanding of the Figure 41 schematic discussion that follows.




Figure 18. Harris HIN211 Internal Circuit Overview [Ref. 43]

Of the signals listed on the left side of the Figure 41 schematic, three are Transmitter inputs (DTR*, TXD (Transmit Data), and RTS*). These inputs are buffered and inverted, and the corresponding output pin of the HIN211 is connected to the appropriate RS-232 connector pin. The fourth input pin, T4IN, is not required and can be left unconnected due to the internal pull up resistor. The five Receiver inputs (RI*, CTS*, RXD (Receive Data), DSR*, and DCD*) from the RS-232 side of the interface have their voltage levels stepped down to TTL levels and are output to the UART.

Due to the very low power that this device consumes in shutdown mode (approximately 1µA [Ref. 43]), it is possible to leave this device permanently enabled.

Doing so saves on the amount of PLD decode logic and wiring required. Thus, EN* (Enable) and SD (Shutdown) are both tied to ground.

As a final note, 220pF capacitors are fitted to each signal line between the transceiver and the RS-232 connector based on a recommendation from Cell Computing [Ref. 44, p.7]. Their purpose is to reduce noise from long RS-232 cable lines.

E. SOLID STATE DISK CIRCUITRY

Of the 512KBytes of Flash ROM provided by the CardPC, System BIOS (Basic Input/Output System) and Video BIOS image files occupy 256KBytes. While the remaining space is more than sufficient for storage of most real-time operating systems (RTOS), it is not sufficient for Windows CE. It is also clearly not sufficient for non-RTOS operating systems such as Windows 95 or 98. To provide flexibility, it was decided to integrate a solid state disk into the TSAS microcontroller design. Doing so also provides additional capability in that the solid state disk will not only serve as a program storage device, it can also function as a non-volatile data storage device (much like a hard drive). For example, during TSAS microcontroller employment the solid state disk could be used to record all aircraft sensor inputs that were received during a mission and which tactors were correspondingly activated. After the mission, the socketed solid state disk can simply be extracted from the TSAS microcontroller and placed into a PC ground station where the stored data could be analyzed.

During PC/104 research, it was noticed that a number of manufacturers were employing M-Systems' DiskOnChip2000 (Figure 19). Upon further investigation, it was determined that the TSAS microcontroller design would greatly benefit from this device.

The DiskOnChip2000 is available in capacities ranging from 2 – 72 MBytes (with 144MBytes expected soon), all in a single, standard 32-pin Dual Inline Package (DIP) form factor. This allows the amount of memory provided to any particular TSAS microcontroller end-application to be scaled according to requirements. With no moving parts, the power consumption is low and the reliability is high. Also, there are no access, seek, or spin-up delays that are typically encountered when using mechanical disk drives.



Figure 19. M-Systems DiskOnChip 2000 [Ref. 45]

The DiskOnChip is mapped into the expansion BIOS memory space of the PC and requires only 8KBytes. During the boot process, the DiskOnChip loads its software into the PC's memory and installs itself as an additional disk drive in the system. The DiskOnChip can also be used as the only disk in the system, allowing the operating system (OS) to boot from it. Many operating systems are supported, including VxWorks, Windows CE, Windows NT 4.0, and Windows 95/98 [Ref. 45]. It is envisioned that the DiskOnChip will be utilized in the following way: the OS files will be stored on and booted from the

DiskOnChip. Once loaded, the OS will start the TSAS application software, also stored on the DiskOnChip.

Figure 42 shows the TSAS microcontroller DiskOnChip 2000 interface. Only thirteen address lines are required to access the memory in any DiskOnChip, regardless of capacity. It is byte-addressable - thus the connection of the lower eight System Data lines. The CE* (Chip Enable) pin is tied to DISK_CS*, an output of the PLD memory-map encoding (to be discussed later). OE* (Output Enable) and WE* (Write Enable) are connected to MEMR* and MEMW*, as per the design example provided in Ref. 46.

As a final note, the DiskOnChip 2000 is available in a 3.3V device. In addition to the benefit of lower power consumption, this fact allows the direct connection of the 3.3V System Data lines. Were the 5V DiskOnChip to be used, an external transceiver might have to be supplied to buffer the data lines.

F. DC TO DC POWER CONVERSION

As mentioned in the requirements listed at the beginning of Chapter II, it is desirable that the TSAS microcontroller be capable of running off a battery and/or a DC plug-in source. The TSAS microcontroller has two power requirements: 5.0V $\pm$ 0.25V, and 3.3V $\pm$ 0.15V. If a regulated DC plug-in source provides one of these voltage levels, then a buck or boost regulator can be employed to generate the other. However, in most TSAS microcontroller applications there are no convenient DC plug-in sources. A battery cannot be used to directly supply either the 5V or the 3.3V due to the voltage tolerances required. This is because of the tendency of Nickel Cadmium (NiCd), Nickel Metal Hydroxide (NiMH), alkaline, and sealed lead-acid battery voltages to dip as they are discharged. For

most types, the ratio of charged to discharged voltage is at least 1.5-to-1 [Ref. 47]. To allow the greatest flexibility, it was decided to design both a 5V and 3.3V power supply so that both batteries and DC plug-in sources could be used. This design is discussed next.

Many embedded designs use linear regulators to achieve desired output voltage levels. They are simple and compact, and can be quickly integrated. LT Luke's design used a linear regulator to generate 3.3V from 5V based on the reference design contained in Ref. 21. Originally the present design used a linear regulator for the same purpose based on suggestions in Ref. 44. As pointed out by Prof. D. Fouts, however, linear regulators can be very inefficient. The difference between the input voltage and desired voltage is dissipated as heat. This not only results in quicker dissipation of batteries, it creates thermal management problems as well. On Prof. Fouts' suggestion, two switching regulators have been employed. There are a number of available devices to choose from but the device decided on was the Maxim MAX1626. This regulator achieves greater than 90% efficiency over a wide load-current range and is capable of supplying more than 2A. It also requires less supporting circuitry than many other switching regulators.

The amount of current to supply at each voltage level was determined conservatively. The exact power consumption of the TSAS microcontroller is not known at this time due to the fact that the type and number of factors to be employed has not yet been finalized and the fact that Cell Computing's CardPC current consumption ratings (as specified in Chapter 7 of Ref. 25) were measured while the system was driving a video display. Nevertheless, estimates of the absolute maximum current required were made as follows:

1. 5V. The maximum observed current drawn by the CardPC at 5V is 630mA [Ref. 25, p.58]. There are only a few devices in the rest of the TSAS microcontroller circuitry that require 5V. The MiniACE is the most power-hungry and requires 0.5A if a 50% duty cycle is assumed. This is conservative, the actual duty cycle is likely to be much less. The rest of the 5V devices consume only 70mA for a total of 1.2A. Although the exact type of tactor to eventually be used is not known, the maximum current draw of the types evaluated is 300mA. It is assumed that not more than three tactors will be pulsed simultaneously (due to sensory overloading), thus the tactors will draw less than 900mA. Total maximum current drawn at 5V: $<2.1A$.
2. 3.3V. The absolute maximum current drawn by the CardPC at 3.3V is 1.2A [Ref. 25, p.58]. The rest of TSAS microcontroller 3.3V circuitry draws a combined absolute maximum of less than 140mA. Total absolute maximum current drawn at 3.3V: $<1.4A$.

Based on these figures, the MAX1626, which is capable of providing more than 2A, should be more than adequate for each voltage level.

The typical operating circuit for the MAX1626 is depicted in Figure 20. The MAX1626 is capable of converting input voltages up to 16.5V. To retain maximum flexibility, this same maximum input was used when making calculations for component values as detailed in Ref. 48. Not surprisingly, component values matched those depicted in Figure 20.




Figure 20. Maxim MAX1626 Typical Operating Circuit [Ref. 48]

The schematic in Figure 43 shows the DC to DC power conversion circuit for the TSAS microcontroller. Input voltage is fed into the DC IN jack and can be controlled by the single-pole double-throw (SPDT) switch. The fact that there are two MAX1626s being employed is taken advantage of in that only one set of bypass capacitors for the input voltage is supplied. This results in a large reduction in the board space required as each of these capacitors is much larger than the actual IC itself. As per the data sheet, SHDN (the Shutdown input) is connected to ground for normal operation. The 3*/5 input is tied low for the left MAX1626 (i.e., the output is set to 3.3V) and tied high for the right MAX1626 (i.e., the output is set to 5V). A Light Emitting Diode (LED) has been added, shown in the lower left corner, to indicate to the user that the TSAS microcontroller is turned on. Otherwise, the circuit depicted in Figure 43 is a direct application of Figure 20.

G. INTERNAL SIGNAL GENERATION

A number of signals have been generated internally for logic and device selection purposes. The Figure 44 schematic shows the use of a PLD in this regard. The PLD selected is the Atmel ATF22LV10C [Ref. 50]. This chip is fabricated using a CMOS process, providing high-density and fast speed. The ATF22LV10C is a low voltage (3.3V) device but its inputs and outputs are TTL compatible. There are 12 inputs and 10 I/O macrocells. Unused product terms are automatically disabled to decrease power consumption. Typical power consumption is 60mA.

The schematic in Appendix C, Figure 45, shows the logic that is to be burned into the PLD. The logic will now be explained, commencing with the memory mapping of devices.

There are three devices in the design that require memory mapping and subsequent address decoding: the BU-61588 Mini-ACE, the TL16C550C UART, and the DiskOnChip 2000 solid state disk. These devices require 8KBytes, 8Bytes, and 8KBytes of memory respectively. Figure 21 shows the memory map of the CardPC. ISA memory between 1MByte and 16MBytes is available for customer use. However, anything mapped in this area will require a large number of address lines to be decoded and there are precious few I/O pins available for use in the PLD. Luckily, there is an 80KByte hole in the System Memory (which is the lower 1MByte) between C8000h and DC000h that is also available for customer use. The key to reduced decoding logic is the fact that the two ISA bus signals, SMEMR* (System Memory Read) and SMEMW* (System Memory Write), serve to indicate that the current read or write cycle is occurring in the lower 1MByte of memory.




Figure 21. CardPC Memory Map [Ref. 25]

Therefore, there is no requirement to decode address lines higher than SA19 if SMEMR* and SMEMW* are used in the decoding scheme.

The BU-61588 Mini-ACE actually requires two select signals, one to indicate that the Mini-ACE has been selected (ACE_CS*) and one to select between its 8KBytes of RAM or 32 bytes of registers (MEM/REG*). $8\text{ K} = 2^{13}$, thus SA12 thru SA1 are required to decode down to the 16-bit word level, as shown in Figure 39. It was decided to map the RAM of the Mini-ACE device to the 8KByte block from CA000h – CBFFFh and the Mini-ACE registers to C8000h – C801Fh. As shown in Table 3, this allows address line SA13 to be used to distinguish between a RAM or register access. Thus, SA13 is fed directly into the MEM/REG* pin of the Mini-ACE (Figure 39). The logic for the generation of ACE_CS* is shown in Figure 45.

Bit / Addr	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0
C8000	1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
C801F	1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1
CA000	1	1	0	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0
CBFFF	1	1	0	0	1	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1
CC000	1	1	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
CDFFF	1	1	0	0	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1	1
DBFFF	1	1	0	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1
3E8	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	0	1	0	0	0
3EF	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	0	1	1	1	1

Table 3. I/O Device Memory Mapping

The DiskOnChip has been mapped to memory commencing at CC000h. As shown in Table 3, this selection of addressing allows address line SA14 to be used to distinguish between a Mini-ACE or DiskOnChip access.

Note that with the address decoding scheme described above, an 8KByte block of memory is reserved for the use of the Mini-ACE's 32 bytes of registers and a 16KByte block of memory is reserved for the use of the DiskOnChip (only 8KBytes is required). While this is not ideal, it saves on decode logic – whereas 13 address lines would have otherwise required decoding, only 7 need be decoded. As will be described shortly, these savings allow the use of a smaller PLD and decrease the amount of etch required. In addition, the three devices being memory-mapped still only use approximately 32KBytes out of the 80KBytes available, the remainder of which will not be used anyway. It may appear that aliasing of the DiskOnChip is a concern. However, as per Ref. 45 p.4, due to

the internal anti-alias algorithm of the DiskOnChip software, its window size can be larger than 8KBytes.

The memory mapping of the TL16C550C UART is a little simpler - PC architecture reserves the address space from 0h to 3FFh for the mapping of common I/O devices. In much the same way that SMEMR* and SMEMW* serve to indicate a system memory access, the ISA bus signals IOR* (Input/Output Read) and IOW* (Input/Output Write) indicate that a mapped I/O device access is occurring. Therefore, only the address space 0h – 3FFh need be further decoded (requiring SA9 – SA0). Figure 22 shows the CardPC's I/O map. Of the four serial ports provided for, COM1 and COM2 are utilized by the built-in Super I/O chip. COM3 was therefore chosen as the serial port for the TL16C550C. The three address lines connected directly to the UART, SA2 – SA0 (as shown in Figure 41), allow the selection of which of the eight bytes to read from or write to. Referring to Table 3, the active-high UART chip select pins are connected to SA5 and SA3 as these two address lines are always high over the range 3E8h – 3EFh. SA7 and SA6 are also always active high and are combined to form one signal, SA7_6, as shown in Figure 44. The remaining address lines, SA9, SA8 and SA4, are connected to the PLD as shown in the same figure.




Figure 22. CardPC I/O Map [Ref. 25]

In addition to memory mapping, the PLD also performs some logic switching necessary for handshaking with the BU-61588 Mini-ACE. RESETDRV and IRQ10 are simply inverted. The logic depicted immediately below RESETDRV in Figure 45 was derived from Ref. 30 and Ref. 33. It serves to ensure that no more wait states than are necessary are inserted by the CardPC. The ACE_CS* output is also fed back to the ISA bus

signal IOCS16* (Input/Output 16-bit Chip Select), which signifies to the CardPC that the current cycle is a 16-bit I/O operation.

Also depicted in Figure 44 are three 74HCT08 AND gates. As mentioned earlier, an earlier version of the present TSAS microcontroller design utilized a 3.3V UART device, thereby requiring that its serial output be buffered to 5V for the TIC interface. A 74HCT08 was chosen to perform this task so that the remaining three gates in the package could be used to generate internal logic, thereby off-loading the PLD. Although an AND gate is no longer used in the UART circuit, the remaining three gates are still used here. When combined with the strategic memory-mapping discussed above, this technique allowed the remaining logic to just fit into a 28-pin PLD. This has resulted in size, power, and routing density savings.

In either a read or write to its address space, the Mini-ACE should be selected. Therefore, SMEMR* and SMEMW* are OR'ed to form the one signal SMEM*. MEMR* and MEMW* are similarly ORed to form STRBD* (Strobe Data).

H. MISCELLANEOUS DESIGN COMMENTS

This section is used for design comments that do not fit neatly into any of the other sections.

It was decided not to implement a power supervisor in this design to save on board real estate. A power supervisor monitors input voltage levels. When the input dips below a preset voltage level, the supervisor switches to battery backup so that RAM can be saved to memory. Power supervisors are very common in computer system designs. However, as

the TSAS microcontroller processes data in real time, it is of little utility to store any RAM data as it will have been superseded by the time the controller is back to operational status.

One of the stated requirements for the TSAS microcontroller is to have an RS-232 and/or RS-422 interface. The RS-422 standard is very similar to the RS-232 standard but it employs two data lines, a negative voltage and its positive image, for noise rejection purposes. The possibility of having both standards implemented in the TSAS microcontroller was investigated but it would have meant an increase in the overall board size. Therefore, it was therefore decided to go with the more common of the two standards, the RS-232.

I. SUMMARY

In summary, the hybrid design methodology described in Chapter II has been used to design the TSAS microcontroller. A COTS product, Cell Computing's CardPC, has been utilized to provide the main processing function. The remaining I/O peripheral functionality has been custom-designed utilizing commonly-available ICs and discrete components. The resulting TSAS microcontroller design meets the requirements specified on pages 15 and 16. It is a robust design in that the majority of the components are available from a number of different manufacturers. The amount of memory and CPU processing power available may also be easily scaled to user requirements, without changing the design. As will be seen in the next chapter, all of this functionality is provided in a palm-sized device that may be placed in a flight suit breast pocket.

IV. TSAS MICROCONTROLLER PCB LAYOUT DESIGN

This chapter will describe the selection of package-types and the PCB layout of the various components of the TSAS microcontroller. Most components are available in more than one package type. For example, resistors can be thru-hole, surface mount, mounted in IC-type DIPs, or mounted in single in-line packages (SIPs). The rationale for the selection of particular component packages will be explained. Design constraints and the application of PCB layout principles will be discussed where appropriate.

Two basic methods for organizing this chapter were considered: discussion of PCB layout by component type, or by function. Organizing by type would allow the quick location of particular components and eliminate redundancies because some components are used in more than one function. However, many of the discrete components have been selected due to specific requirements of the functions they support. Thus, it was felt that the latter method would make for more coherent reading. An overview of the layout process will be presented first, including a description of how software was applied to accomplish the task.

A. PCB LAYOUT OVERVIEW

One of the most critical requirements of the TSAS microcontroller is that it be as small as possible, preferably palm-sized. From the outset, it was clear that the design would have to be double-sided (i.e., components would have to be placed on both sides of the printed circuit board). As the CardPC would occupy most of one side of the board, it was also equally clear that maximum use of surface-mount technology (vice thru-hole) would

have to be made. For those not familiar with these technologies, Figure 23 shows examples of surface mount and thru-hole devices for a resistor and IC. As surface mount devices are soldered onto the same side of the board on which they are placed (vice thru-hole devices), components can be mounted directly opposite each other on both sides of the board (as long as there are no thermal issues to contend with). Not all components are available in surface mount technology, however, and this fact has forced certain layout decisions to be made as will be discussed below. Additionally, some surface mount devices can actually take up almost as much room in extra lead length as they save by allowing mounting on both sides. In these cases, it may be better to stick with the solder joint security of thru-hole devices; in a couple of instances this was done.




Figure 23. Mounting Technologies: (a) Thru-hole, (b) Surface Mount [Ref. 49]

To help orientate the reader with the design, the finished product is presented first. Figure 46 in Appendix D shows all facets of the design and is to scale. The figure is transparent (i.e., both sides of the board are simultaneously presented). The Reference Designators (e.g. Y02 or U10) of components mounted on what will be referred to as the bottom (CardPC) side of the board are mirrored. As this diagram is difficult to interpret, Figure 47 shows an enlarged picture without the interconnect. The two following diagrams, Figure 48 and Figure 49, show only the top and bottom side of the board respectively. These diagrams, and the rest of the Appendix D diagrams, are larger than the actual end-product so that more detail is legible.

Before immersing into the details of component location, a brief overview of the placement philosophy will now be given. The TSAS microcontroller is designed to reside in an open breast pocket mounted on the tactor vest. Connectors to external components have therefore been located at the top of the device, as shown in Figure 46. The MIL-STD-1553 twinax connectors have been located to the far left, closest to the cockpit wall, and the on/off switch to the far right, closest to the right hand of the pilot. The TIC connector has been located at the bottom of the device, to relieve congestion at the top. As the tactor wires are internal to the vest, this presents no problem. The remaining components were placed based on routing density, routing length, noise, and thermal considerations.

Routing density and length was minimized through effective use of the "ratsnest," as will be discussed in the next section. Noise was minimized through the use of power and ground planes, bypass capacitors, and by physically separating analog and digital signals.

Heat generation is an important consideration when designing a system to be as dense as possible. The three most “lossy” components in the TSAS microcontroller are the CardPC and the two power regulators. As discussed in Chapter III, the regulators were chosen to be of the switching type for greater efficiency and therefore less heat generation. They were also not placed directly opposite the CardPC, but rather were placed near the top of the device, next to the connectors. As for the Pentium CardPC, it has a feature called “thermal throttling”. This feature slows the CPU clock down to $1/4^{\text{th}}$ of its normal speed when the CPU temperature exceeds specified limits. Also, the 166 MHz Pentium CardPC, the P166, is envisioned as the CardPC of choice for the TSAS microcontroller for two reasons. First, because it should more than meet all TSAS microcontroller processing requirements and second, because Cell Computing states that the P166 does not need active cooling in most situations. This is due to the fact that the aluminum framing acts as a heat sink and thermal conductor. The P166 actually consumes less power than the P133 – 6.6W vice 9.7W, under worst case operating conditions [Ref. 25, p.5]. Presumably, this is as a result of a die shrink for the Pentium 166 chip. If it turns out that cooling is required for the P166, or if a faster CardPC is desired, then two cooling options are available, a heat sink or a fan. The passive heat sink, Figure 24, would probably be preferred, if it provided sufficient cooling. The heat sink would add an additional 0.6 inches in height to the end product.




Figure 24. CardPC Heat Sink [Ref. 25]

B. OVERVIEW OF PCB LAYOUT SOFTWARE REQUIREMENTS

As mentioned in Chapter III, the circuit schematics for the TSAS microcontroller design were built using Cadence's *Concept Schematic*. The PCB layout was completed with Cadence's *Allegro*. To facilitate understanding of the rest of this Chapter, the interaction of these two programs, and their specific requirements, will now be discussed.

Concept Schematic allows an electronic design to be pictorially displayed in a form that is understandable to humans. A wide variety of analog and digital devices can be copied from Cadence-provided standard libraries and inserted into a design. In the event that the standard library does not contain the parts that are required for the design (as was the case with almost all components in the TSAS microcontroller design), the user must

supply information about the device (e.g., the number of pins, their names, type of output, etc.). The Cadence program *RapidPart* can facilitate this task. The pins on the resulting representations of the devices can then be interconnected, as shown in the schematic diagrams in Appendix B. If simulation models for the devices are available, then the functionality of the design may be verified through simulation. If the design is to be physically implemented on a printed circuit board, then a separate software program (Allegro, in the case of the TSAS microcontroller) is required to perform the layout.

The primary purpose of Concept Schematic, as far as Allegro is concerned, is to create what is called a “netlist” – a description of how the pins on the devices involved are interconnected. The netlist has no bearing to the physical layout of the design as portrayed in the schematic diagram – it simply indicates which pins are connected to which pins. In fact, Allegro can use the netlists of many other CAD tools besides Concept Schematic. In addition to the netlist, Allegro requires information about the physical dimensions of devices. As with Concept Schematic, for many common devices, this information is available in Cadence’s standard libraries. If not available, it must be supplied by the user (the program *Allegro Symbol* has been created to simplify this task). Therefore, in the case of the TSAS microcontroller, physical dimension information from manufacturers’ data sheets was used to create symbols for each device used. These symbols also specify “pad” dimensions – pads are the board contact points used for soldering the pins of devices. Figure 25 shows an example of a symbol for a 24-pin Thin Shrink Small Outline Package (TSSOP) device, complete with pads. (The pads are shown in black here. In some of the Appendix D diagrams, pads are shown as transparent so as not to block other detail.)




Figure 25. TSSOP24 Symbol (Not to Scale)

Once a physical symbol is available for each device used, the design may be “packaged.” Packaging refers to the association of components in a schematic design with a particular type of physical package. For example, the AND gates (and their symbolic equivalents) depicted in Figure 44 are packaged into a single 14-pin 74HCT08 TSSOP. Each AND gate is assigned three specific pin numbers based on the information that has been provided about the 74HCT08 TSSOP. In the case of Allegro, packaging may be accomplished using an intermediary front-end tool program called *fet2a*. This program takes the netlist from Concept Schematic, packages it using Cadence’s *PackagerXL*, and then associates it with an Allegro PCB layout designated by the user.

Once this association has been completed, the user may commence laying out the board. As each device is placed, a “ratsnest” is displayed – this serves to indicate how the pins of the device being placed are interconnected to pins of the devices already placed. An example of a ratsnest is shown in Figure 26. This tool greatly aids the user in minimizing the density and amount of etch required to connect the devices. Once all devices are placed, the user can replace the ratsnest with etch. After silkscreening, the design is converted to a format readable by PCB manufacturers (a Gerber plot), and is sent for fabrication.




Figure 26. Example "Ratsnest"

C. PRINTED CIRCUIT BOARD COMPOSITION AND DESIGN CONSTRAINTS

1. PCB Layer Stackup Assignment

Before etching a PCB layout, the engineer must decide on how many etch layers to utilize. The number of layers required is predicated based on two issues: routing density, and the acceptable level of EMI. Given the added weight and expense of multi-layer designs, one cannot arbitrarily add extra layers simply to ease routing work. Careful device placement has allowed the component-dense TSAS microcontroller design to be etched

using only four routing layers. With the addition of power and ground planes, the total number of layers for this design could have been as little as six; however, as stated in Ref. 51, p. 17, an eight layer board with 4 routing layers and 4 power planes is optimal for reducing EMI. Based on this information, the TSAS microcontroller PCB layer stackup assignment has been designed as per the stackup recommended in Ref. 51. More specifically, the stackup is as follows, from top to bottom:

1. TOP (first signal layer)
2. GND1
3. S2 (second signal layer)
4. GND2
5. VCC (split into two planes, one each for 3.3V and 5V)
6. S3 (third signal layer)
7. GND3
8. Bottom (fourth signal layer)

The additional two ground layers will be an added expense (less than \$100/unit); however, the requirement to contain EMI in cockpits was deemed to override this cost.

2. Pads

To electrically connect the pins of different components, the pins must first be electrically connected to the board itself. This is accomplished through the use of pads. Pads are an area of exposed copper on the board to which the pins can be soldered. There are different types of pads for the different types of components being mounted. In the TSAS microcontroller design, both thru-hole and surface-mount devices are utilized. A pad for thru-hole devices consists of a hole drilled through the board into which the pin can be inserted and an area of exposed copper around the hole. Pads for surface mount devices

generally consist of a rectangular area of exposed copper large enough for the lead to sit on and be soldered to the board. Both types of pads may additionally specify an outer area of protective soldermask.

There are no hard and fast rules about what dimensions pads should be. Generally speaking, a pad should be big enough to account for variations in pin size and to make a good solder joint, but not so big as to unnecessarily consume board space or make it hard to solder (due to heat dispersion). While all manufacturers publish detailed physical dimensions of their devices, only a few provide recommended pad dimensions (or "land areas/patterns"). There are many recommendations in the literature as to how to dimension thru-hole pads; Ref. 52 was found to be the best:

- Finished hole size = nominal pin diameter " (0.007" to 0.016")
- Pad diameter = finished hole size + (0.020" or more)
- Soldermask diameter = pad diameter + (0 to 0.010")

These recommendations have been followed in the cases where the manufacturer did not specify pad dimensions (i.e., the zener diodes, TST-9007, and header).

Information on recommended surface-mount pad dimensions was harder to come by. For the ICs, pad sizes were determined by mimicking (and scaling as necessary) manufacturer-specified pads for leads that were identical, or very close to, the leads in question. Pad sizes for the miniature capacitor and resistor surface-mount components are a little more critical, given that they are likely to be placed and soldered by a machine. Kemet Electronics Corporation publishes a document, Ref. 53, that specifies recommended pad dimensions for the sizes of components used. The Kemet document is in turn based upon the Institute for Interconnecting and Packaging Electronic Circuits (IPC) document, IPC-

SM-782 Revision A. The recommendations in the Kemet document were followed closely when designing pads for the EIA 0603 and EIA 3528-size capacitors and resistors. The pads for the remaining capacitors and resistors were designed in accordance with specific OEM recommendations.

3. Traces

Traces are the lines of exposed copper ("etch") that electrically connect pins. A number of factors need to be considered with respect to traces. These factors will be discussed here.

The purpose of etch is to carry electrical signals from an origin to a destination. As such, they must be capable of carrying the current demanded of them. Table 4 and Table 5, taken from Ref. 54 (and in turn derived from IPC document IPC-D-275), provide guidelines for external and internal trace widths. It is expected that the TSAS microcontroller will run moderately hot (over 60°C), due to its compact design. As well, the majority of traces carry currents well under 0.2A. Therefore, the default Cadence trace width of 0.005 inches was accepted. Selected trace widths, such as those of the DC-DC conversion circuit, were increased to 0.010 inches. Additionally, most of high-current etch has been routed on the external (TOP and BOTTOM) signal layers.

1 Oz. Plating, External Conductors					
Width	10°C	20°C	30°C	45°C	60°C
.005"	400 mA	500 mA	650 mA	800 mA	1.0 A
.010"	800 mA	1.0 A	1.3 A	1.6 A	1.9 A
.015"	1.2 A	1.5 A	1.8 A	2.1 A	2.8 A
.020"	1.5 A	1.7 A	2.0 A	2.5 A	3.1 A
.025"	1.7 A	2.2 A	3.0 A	3.5 A	4.0 A
.050"	3.2 A	3.9 A	4.8 A	5.7 A	6.5 A
.100"	4.8 A	6.2 A	8.0 A	9.5 A	10.4 A
.150"	6.0 A	8.5 A	11.0 A	12.6 A	13.5 A

Table 4. Recommended Trace Widths for External Conductors [Ref. 54]

1 Oz. Plating, Internal Conductors				
Width	10°C	20°C	30°C	45°C
.005"	200 mA	225 mA	250 mA	275 mA
.010"	400 mA	450 mA	600 mA	750 mA
.015"	550 mA	600 mA	750 mA	1.0 A
.020"	650 mA	700 mA	800 mA	1.2 A
.025"	750 mA	1.0 A	1.2 A	1.7 A
.050"	1.5 A	1.7 A	2.2 A	2.8 A
.100"	2.2 A	3.1 A	3.7 A	4.5 A
.150"	3.0 A	4.0 A	5.2 A	6.1 A

Table 5. Recommended Trace Widths for Internal Conductors [Ref. 54]

In high-speed designs, the length of traces can become a concern. This is because the traces begin to take on the characteristics of transmission lines; if the trace is not terminated some way in its characteristic impedance, then reflections will occur which can cause noise voltages and false signals. The point at which this happens is usually defined as when the two-way delay is more than the rise time of the pulse. This translates to approximately 3 inches for a 1 ns rise time in FR4 [Ref. 55]. The fastest device in the TSAS microcontroller is the 74HCT08, which may have a rise time as fast as 1 ns [Ref. 66]. The trace lengths of all of its leads, however, are much smaller than 3 inches. In fact, due to the compact design of the TSAS microcontroller, there appear to be no traces that would act as transmission lines. There are some high clock speeds in the TSAS microcontroller but they are relegated to the CardPC. The system clock of the ISA bus runs at only 8MHz. The highest clock speed in the rest of the design is that of the 16MHz clock and its trace lengths are very short. Therefore, it appears that transmission line effects should not be a problem in the TSAS microcontroller design.

Also of potential concern is the distance between traces. If traces are placed in close proximity and run parallel to each other over a distance, then crosstalk may become a problem. The distances involved are similar to those discussed above for transmission line effects [Ref. 56]. Crosstalk may be reduced and controlled through the effective use of ground and power planes [Ref. 55]. Therefore, given the short trace lengths and the use of power and ground planes between every signal layer in the TSAS microcontroller design, it is not expected that crosstalk will constitute a problem. The Cadence default distance of 0.005 inches between traces has been used.

The last issue to be discussed with respect to traces is that of right-angle routing. This is a controversial area in the electronics industry, with some saying that there is no problem with using right-angle corners and others saying that they must be avoided and that 45 degree miter joints must be used instead. The reasons for not using right-angle corners seem to vary. For example, Ref. 57 states that right-angles cause traces to become similar to antennae thereby causing EMI problems while Ref. 51 states that their use causes the capacitance per unit length to increase and the inductance per unit length to decrease thereby causing signal distortion. To avoid any possible problems, no right-angle corners were utilized in the routing of the TSAS microcontroller PCB.

Figures 50-54 of Appendix D show the etch for the TOP, S2, S3, and BOTTOM signal layers respectively.

4. Vias

The density of signal traces in this design has dictated the use of multiple routing layers. A method of connecting signals on these different routing layers and for connecting pins to power and ground planes is required. This is accomplished through the use of vias. A via is a hole drilled into the printed circuit board and then plated so that one signal layer may be electrically connected to another. There are three types of vias: thru-vias, blind vias, and buried vias. The thru-via, as the name implies, goes all the way through the board and can be used to connect signals on the top and bottom layers, or any combination in between. A blind via goes from the top or bottom layer to a layer somewhere in the middle. A buried via connects two middle layers. Theoretically, the thru-via is all that is required to connect all signals on any board. However, indiscriminate use of thru-vias will waste routing space

by placing unnecessary "obstacles" around which etch must be diverted. Also, there is debate in the industry as to the effects of vias on signal integrity, with some advocating the use of as few vias as possible. Therefore, blind and buried vias have been used whenever appropriate.

In addition to the number of vias, the diameter of the holes is another factor to consider. The plated hole must be capable of carrying the electrical current through it. As stated in Ref. 52, the greater the number of different hole sizes in a design, the greater the expense to manufacture the board. In the TSAS microcontroller design, the majority of the signal traces carry a very small amount of current. However, some (e.g., the switching regulator traces) carry larger amounts of current. Therefore, two different vias were used in the design, a small (0.012 diameter) via and a large (0.018 diameter) via. Dimensions for the small via were based on recommendations contained in Ref. 20. The large via is a scaled version of the small.

Vias have also been used to electrically connect the three ground layers in the design. Normally, the ground planes would be tied together through the connection of thru-hole ground pins. Due to the large number of surface-mount devices in the TSAS microcontroller design, this method is insufficient. As suggested by Ref. 51, buried vias have been sparingly and uniformly placed around the board where space allows to electrically connect the ground planes.

5. Other Considerations

In order to manufacture the PCB, the fab plant will require two or three tooling holes [Ref. 52]. Two holes have been added to the TSAS microcontroller PCB specifically

for this purpose and are located at the top and bottom corners of the right side of the board as shown in Figure 46. In addition, there are two holes drilled through the board to mount the CardPC. These are located on the left side of the PCB and can also be used as tooling holes before the CardPC is mounted.

Once all components are mounted on the PCB, the board will require encasement for protection. An aluminum case with louvered sides for air circulation is envisioned, although the case was not designed as part of this thesis. In order to precisely position I/O components, it was assumed that the case would be 1/16" thick and separated by 0.050 inches from the edges of the PCB. It is also envisioned that the four tooling holes mentioned above will be used to mount the assembled PCB onto posts built into the case.

D. CARDPC/CPU INTERFACE COMPONENTS

This section will describe the physical packages selected for the components depicted in the Figure 38 schematic diagram.

1. EASI Connector

The CardPC interfaces with other components on a host PCB via an EASI connector. There are two types of EASI connectors to choose from: the C2I-EZ-TH thru-hole device, available directly from Cell Computing, or the ICMA-236LMY surface-mount device, available from Honda Connectors. The first instinct was to use the surface-mount device. However, after careful measurements, it was evident that the surface-mount device would save very little board area overall, approx. 0.25 square inches due to its larger dimensions [Refs. 58 and 44]. Given the importance of this connector in particular, it was

decided that the solder-joint security of the thru-hole device outweighed the minimal board space that the surface-mount device would save.

Cell Computing also manufactures a clip that retains the CardPC in its connector socket, the C2I-EZ-CLIP-1. Cell recommends the use of this device to prevent accidental dislodging of the CardPC. Given the possible vibratory environments the TSAS microcontroller may be exposed to (e.g. helo operations), it was felt prudent to include this clip in the design.

The overall dimensions of the CardPC, EASI connector, and clip combination are approximately 2.7 inches by 3.6 inches [Ref. 44]. These dimensions have played the most critical role in determining the overall board size. The board had to be increased in width and length to allow for thru-hole devices and I/O connectors. If not for this fact, the board could have been even smaller as there is still some unused space available.

As a final note, a special plastic sheet is required between the CardPC and the PCB for signal isolation, as specified in Ref. 44, p.13.

2. Capacitors

There are a large number of power pins on the EASI connector. To reduce noise, timing errors, and EMI, 10 μ F and 0.01 μ F bypass capacitors have been uniformly distributed in the vicinity of these pins. The 0.01 μ F capacitors are ceramic, surface-mount EIA 0603 types, and the 10 μ F capacitors are tantalum surface mount EIA 3528 types, in accordance with information provided at Ref. 24. Pads have been designed according to information contained in Ref. 53.

3. Resistors

The CardPC contains many pull-up resistors internally, thus the majority of unused inputs do not have to be externally terminated. There are some exceptions, however, as detailed in Appendix C of Ref. 25. These exceptions have been terminated with surface-mount EIA 0603 1/16 Watt resistors, in accordance with information provided in Ref. 24.

4. PCB Layout

Figure 54 shows the layout of the CardPC, capacitors and resistors on the PCB. As can be seen, the resistors and capacitors are tightly squeezed in behind the 236 pins of the EASI connector. This has been done so as not to block etch routes on the other side of the connector. The resistors have been placed in close proximity to the pins they are pulling up or down and the capacitors have been uniformly spread across the connector, in close proximity to their respective voltage pins (either 3.3V or 5V).

At 0.51 inches in height, the CardPC is the determining factor for overall height on the bottom side of the board.

E. MIL-STD-1553 DATA BUS INTERFACE

This section will describe the physical packages selected for the components depicted in the Figure 39 schematic diagram.

1. BU-61588/65178 Mini-ACE

The Mini-ACE family is available in a choice of pin grid array (PGA), quad flat pack, or gull lead packages. Due to the very large amount of surface area consumed by the quad flat pack, it was quickly ruled out as an option. The gull lead Mini-ACE is a surface mount device but the long length of its leads almost doubles the amount of PCB area that it

requires compared to the main ceramic package (1.91 square inches vice 1.0) [Ref. 31]. Unfortunately, the Mini-ACE is not currently available in a Ball Grid Array (BGA) package. (A BGA device has an array of short pins (balls) that are soldered to the same side of the board as the device itself is mounted on. The PGA has a similar arrangement of pins, but is a thru-hole device). The disadvantage of the PGA Mini-ACE package is that because the CardPC covers most of the surface of one side of the PCB, an increase in overall board dimensions would be necessary to mount the PGA to the side of the CardPC. It was therefore decided to utilize the gull lead Mini-ACE package.

2. 16MHz Clock

Clock oscillators come in two standard sizes, full size and half size. Surface-mount, half-size oscillators in a standard pin orientation were selected for the TSAS microcontroller. These four-pin devices measure 0.5" by 0.5" and are made by a number of different manufacturers; Vishay Dale was selected. This 16MHz oscillator consumes 25mA max., and has a standard frequency stability of 0.01% [Ref. 59].

3. TST-9007 Twin Stacked Transformers

The receiver/transmitter channels of the ILC DDC Mini-ACE require external transformers to connect to the MIL-STD-1553 bus. DDC has a subsidiary, Beta Transformer Technology Corporation, that manufactures transformers for just this purpose. The TST-9000 series of transformers have been specifically designed to conserve board space when two transformers are required. The two transformers are stacked on top of each other (TST stands for Twin Stacked Transformers). Thus, two transformers may be mounted in the surface area normally required of one, albeit at the expense of height. The

maximum height of the device is only 0.280 inches which does not make it a determining factor.

Similar to the Mini-ACE, the TST-9000 series is available in a thru-hole, quad flat pack, or gull-lead package. The same considerations apply: the quad flat pack consumes too much PCB surface area and the gull lead package requires triple the surface area compared to the main body of the package in this case (1.27 square inches vice 0.39) [Ref. 34]. Thus, it was decided to utilize the thru-hole version of the TST-9007.

4. MIL-STD-1553 Twinax Connectors

As discussed in Chapter III, a senior MIL-STD-1553 specialist at ILC DDC suggested utilizing twinaxial (twinax) cables and connectors to tap into the 1553 data bus [Ref. 33]. A single twinax cable allows for the shielded transmission of both a negative and a positive voltage. Coaxial (coax) cables could also be used but twice as many are required. As the TSAS microcontroller provides redundant 1553 bus channels, two twinax connectors are required. Based on a recommendation from the DDC specialist, Trompeter Electronics Inc. (a manufacturer of cables and connectors) was contacted. A specialist at Trompeter suggested that their 450 series of concentric twinax connectors be investigated. The 450 series have been specifically designed for airborne digital data bus applications utilizing the MIL-STD-1553 data bus [Ref. 60]. After further investigation, these connectors did appear to be slightly smaller than similar offerings from other manufacturers. The BJ450 was selected, as this rear-mounting bulkhead jack minimizes the distance of the solder pot over the circuit board.

5. Mini-ACE 8x2 Header

Every Remote Terminal (RT) in a MIL-STD-1553 data bus system requires a unique address. As the addresses of the other RTs in any TSAS microcontroller end-application will vary, a means of adjusting the TSAS microcontroller's RT address is required. A header has been provided for this purpose. A metallic clip placed over two opposite pins in a header electrically connects them. Sets of pins can be adjoined to form a header as long as is required. Eight were needed in this case.

The type of header selected has a 2mm distance from pin to pin, the smallest that could be found. Surface mount versions of 2mm headers are available but consume almost twice the surface area of the main body due to the size of their leads. The thru-hole version was selected for greater solder joint security.

A number of manufacturers make 2mm headers; Samtec was selected.

6. Capacitors

Ref. 31 calls for 10 μ F tantalum and 0.1 μ F capacitors to be used for bypassing. Small form-factor EIA 3528 and EIA 0603 surface mount capacitors were chosen, respectively.

7. PCB Layout

A diagram of the MIL-STD-1553 Data Bus circuit PCB layout is shown in Figure 55. There are a number of PCB layout constraints in this circuit, which will now be discussed.

The TX_RX, CH A and CH B signals depicted in Figure 39 are analog signals. As stated in Ref. 30, p. 250, ground planes should not be placed under these traces to prevent

analog noise transmission. As shown in Figure 61, the ground planes have been cut from these areas. Ref. 30, p.248 also states that the external transformers should be placed as close as possible to the TX_RX pins. This was accomplished, while also keeping these signals to the side of the board, away from the CardPC.

Ref. 30, p.250, states that the optimal circuit layout is to have a single ground for both the analog and digital sections. As a ground plane could not be placed under the transformers (as mentioned above), etch was used to connect the ground pin of the transformers to the ground planes.

Ref. 30, p.254, states that the external clock source should be mounted as close to pin 30 of the Mini-ACE as possible. This was accomplished – the 16MHz Clk pin is located directly across from pin 30.

Note 12 on page 3 of Ref. 31 states that 10 μ F tantalum and 0.1 μ F capacitors should be placed as close as possible to pins 20 and 72 of the Mini-ACE PGA package, as well as a 0.1 μ F capacitor at pin 37. This was accomplished.

To provide additional room at the top end of the TSAS microcontroller for other I/O connectors, it was decided to mount the two twinax connectors directly opposite each other (i.e., one was mounted to each side of the board). The 0.438" diameter of the jacks are not a consideration on the bottom side of the board due to the height of the CardPC but will be the determining factor (by about 0.1") on the top side of the board. As clearance above the PCB components is required for air circulation anyway, this was considered acceptable. Taking into account the longitudinal dimensions of the jacks, the thickness of the

encasement and its separation from the PCB, it was calculated that the jacks would extend 0.212" into the PCB.

F. TIC INTERFACE

This section will describe the physical packages selected for the components depicted in the Figure 40 schematic diagram.

1. TL16C550C UART

As described in Chapter III, the Texas Instrument TL16C550C UART was selected as the heart of the TIC interface. The smallest form-factor that this device is available in is a surface-mount, 44-pin, Plastic J-Leaded Chip Carrier (PLCC-J) device. The symbol for this commonly-used form-factor was one of the few that was available for use in the Cadence standard libraries. The symbol was used directly with no modifications to it or the associated pads.

2. 1MHz Clock

Similar to the Mini ACE oscillator, the clock selected for the TIC interface is a half-size CMOS surface-mount device. It was originally thought that placing a thru-hole version of this device in a socket would facilitate the insertion of slower clocks for testing purposes. However, as stated in Ref. 51, sockets should never be used with clocks as they add lead length inductance and provide an additional path for RF currents and harmonics to radiate or couple. The clock has, therefore, been left unsocketed. Thus, any desired replacement will have to be soldered.

3. 74HCT08

Three of the four 2-input AND gates available in a 7408 package have been used in the TSAS microcontroller design, as pictured in Figure 40. A 5V, 14-pin HCT (High-speed CMOS, TTL compatible) TSSOP device has been selected, the Fairchild Semiconductor 74VHCT08. This high-speed device will minimize propagation delays while ensuring compatibility with CMOS signals.

4. Modular Jack

A robust, 4-pin connector is required for the TIC interface. A surface mount modular jack (similar to a phone jack) from Molex was selected. This jack will allow easy connection/disconnection of the tactor vest and TSAS microcontroller.

5. Capacitors

There are two bypass capacitors depicted in Figure 40; both are ceramic, surface-mount EIA 0603 capacitors.

6. PCB Layout

Figure 56 shows the PCB layout of the TIC interface circuit. Given the high number of interconnections required between the UART, EASI pins and PLD, the UART was placed approximately mid-way between these devices and the TIC modular jack. The length of the data trace from the UART to the TIC jack is less than two inches, which should not present any propagation delay problem. The 5V and ground connections of the TIC jack are made directly to the respective power planes. The Clk pin of the oscillator is located in close proximity to the TIC Clk pin.

G. RS-232 INTERFACE

This section will describe the physical packages selected for the components depicted in the Figure 41 schematic diagram.

1. RS-232 Transmitter/Receiver

The data generated by the Super I/O UART must be brought to and from TTL voltage levels from RS-232 levels. This is accomplished via an RS-232 transceiver. The Harris HIN211 was selected; the Maxim MAX211 is identical. These devices require only 0.1 μF external capacitors whereas other transceivers require as much as 10 μF capacitors. They also consume very low power (typically 5mA), and have a power-saving shutdown function. Both come in DIP, Small Outline IC (SOIC), or Shrink Small Outline Packages (SSOP) [Ref. 43]. Of the two surface mount devices, the SSOP occupies only 40% of the area that the SOIC does. Thus, the SSOP was chosen as the RS-232 transceiver package for the TSAS microcontroller.

2. Capacitors

There are a number of capacitors required for this circuit. Four 0.1 μF bypass capacitors are required for the RS-232 transceiver. These have been selected as surface-mount EIA 0603 ceramic capacitors.

Eight 220pF capacitors are also utilized on the lines to the 9-pin D-connector. These capacitors help to suppress noise. Surface-mount EIA 0603 ceramic capacitors have also been utilized here.

3. 9-pin Subminiature D-Connector

9-pin subminiature D-connectors are commonly used in RS-232 serial communications circuits (e.g., for PC mice, keyboards, and modems). A number of other devices, such as some GPS models, also use this connector. It was therefore selected as the connector for the TSAS microcontroller RS-232 circuit. There are a number of manufacturers who make these connectors. A connector from Positronic, an established, ISO 9001-certified company, was selected.

Although surface-mount subminiature D-connectors are available from Positronic, a company representative recommended that they not be used because they are very hard to establish secure solder joints with. Thru-hole D-connectors are available with the leads bent to 90 degrees, which saves on the amount of surface area consumed by the connector. Of the 90° thru-hole devices, there are two basic types: the long mount and the short mount. A length is chosen depending on how far the bulkhead (i.e., encasement) is from the PCB; a short mount connector, the HDC905B3000, was selected for the TSAS microcontroller. Ref. 61, p. 46, gives all connector dimensions and contact hole patterns.

4. PCB Layout

Figure 57 shows the PCB layout of the RS-232 circuit. All capacitors are mounted in close proximity to the pins they are designed to support. The D-connector is 0.339" in length from the inside edge of its front plate to the middle of its PCB mounting hole. Thus, the connector was placed such that the middle of the mounting hole was 0.226" into the PCB ($0.339'' - 0.050'' - 1/16'' = 0.226''$).

H. SOLID STATE DISK CIRCUITRY

This section will describe the physical packages selected for the components depicted in the Figure 42 schematic diagram.

1. DiskOnChip 2000

The DiskOnChip2000 solid state disk is only available in one package – the 32-pin thru-hole DIP depicted in Figure 19. This is a large device, consuming about 1.2 square inches of PCB area. It is preferable to convert the DiskOnChip to a surface mount device by using a socket for three reasons. First, an increase in overall board dimensions would be necessary if the DiskOnChip was to be mounted as a thru-hole device. Second, being able to remove the DiskOnChip from a socket allows stored data to be extracted at a PC ground station as described in Chapter III. Third, a socket gives the user the flexibility to easily increase the amount of memory available in the system at a later date.

Using a socket will increase the height of this component by about 0.1 inches. The Disk On Chip actually comes in two heights, depending on the memory capacity: the 0.216 inch-high MD-2200 (capacities of 24MBytes and below) and the 0.464 inch-high MD-2201 (capacities of 40MBytes and above). In addition to having a lower profile, the MD-2200 also consumes less power [Ref. 45]. It is assumed that 24MBytes of RAM is more than sufficient for any current TSAS microcontroller end-application. Therefore, the total height of the socket and DiskOnChip would be approximately 0.32 inches. This is about 0.1" lower than the height of the twinax connector and is not a determining factor. Of course, if more than 24Mbytes is required, the higher-profile MD-2201 may be used, at the expense of a larger protective enclosure. The socket chosen is described next.

2. 32-pin "J" Lead Socket

A number of manufacturers make sockets for converting 32-pin thru-hole DIPs to surface mount devices. Two pin terminations are available: gull wing, where the leads extend outward from the device and "J" lead, where the leads curve underneath the device. Gull wing leads are easier to hand-solder but consume more PCB surface area. Given the extremely small form factor of some of the surface mount components in the TSAS microcontroller design, it is expected that most if not all components will be machine placed and soldered. The J lead socket was, therefore, selected. An Andon Series 220 socket was utilized.

3. Capacitor

A 0.1 μF bypass capacitor has been added to this circuit as suggested by the example design provided at Ref. 46, p.3. The capacitor selected is a standard EIA 0603 ceramic surface mount component.

4. PCB Layout

Figure 58 shows the solid state disk circuit layout. The DiskOnChip 2000, U06, was placed near the bottom of the board due to the high number of connections between this device and adjacent EASI connector pins.

I. DC TO DC POWER CONVERSION

This section will describe the physical packages selected for the components depicted in the Figure 43 schematic diagram.

1. DC Power Jack

There is a wide assortment of DC power jacks available on the market. One that was as small as possible, given the supply current requirements, and that could be securely mounted to the TSAS microcontroller encasement, was desired. A 2.1mm diameter jack with two bulkhead-mount screw holes from CUI Stack was selected.

2. On/Off Switch

An On/Off switch that was small and yet easy to operate with a gloved-hand was desired. A single-pole, double-throw switch with two bulkhead-mount screw holes from C&K Components was selected. C&K provide the option of various heights on their switches, which will allow flexibility to meet end-user desires.

3. MAX1626 DC-DC Controllers

The two Maxim MAX1626 DC-DC switching regulators come in a Small Outline (SO)-8 package. This device consumes only 0.048 square inches of PCB area. Its supporting components, however, consume much more.

The components required to support the MAX1626 are described next.

4. Current-Sense Resistors

Each MAX1626 utilizes a current-sense resistor to minimize the peak switching current, thereby increasing efficiency and reducing the size of external components. As recommended by Ref. 48, a Dale Power Metal Strip Resistor, model WSL-1206, was selected. The size of this device conforms to EIA 1206. As the surface mount resistor supplied in the Cadence standard library, *smdres*, also conforms to EIA 1206, this symbol was used in the design.

5. Inductors

In accordance with the recommendation given by Ref. 48, Coilcraft surface mount power inductors, Series D03316, were selected for the regulator circuits. Pad sizes were drawn as per the suggestions made in Ref. 62.

6. External Switching Transistors

Each MAX1626 circuit requires a P-channel enhancement-mode MOSFET. A device that meets the requirements and comes with two transistors in one package was found, the Siliconix dual P-Channel MOSFET, Si4963DY [Ref. 63]. The device comes in a standard SO-8 package.

7. Diodes

The high switching frequency of the MAX1626 requires a high-speed rectifier. Schottky diodes are recommended by Ref. 48. Motorola surface mount Schottky power rectifiers, model MBRS340T3, were selected. Package dimensions were derived from Ref. 64.

A Light Emitting Diode (LED) has also been supplied to indicate power state to the user. A standard thru-hole LED was selected.

8. Capacitors

As shown in Figure 43, the MAX1626 requires four different types of capacitors: a 68 μF low-Equivalent Series Resistance (ESR) tantalum, a 220 μF low-ESR tantalum, a 0.47 μF , and a 0.1 μF capacitor.

In accordance with the suggestion made at Ref. 48, the low ESR capacitors are Sprague 595D Series solid tantalum chip capacitors. A case size of C (6.3V) was selected

for the 220 μF capacitors and a case size of D (20V) for the 68 μF capacitors. Solder pads were dimensioned according to the table provided at Ref. 65, p.52.

The 0.47 μF capacitor was chosen to be a surface mount EIA 1206 ceramic capacitor. The 0.1 μF capacitors were selected as surface mount EIA 0603 ceramic capacitors.

9. PCB Layout

Figure 59 shows the PCB layout of the DC-DC power conversion circuit. The circuit was placed near the top of the board to allow for greater air circulation. Effective use of both sides of the PCB has been made. Most of the etch is located on the TOP and BOTTOM signal layers to allow for greater cooling and greater current-carrying capacity. Page 12 of Ref. 48 recommends a placement and routing scheme for the current-sense resistor and the 0.01 μF and 0.47 μF bypass capacitors; this recommendation was followed. The transistor SO-8 is placed directly between the two MAX1626 SO-8's.

The On/Off switch and the DC power jack have been located on the bottom (CardPC side) of the PCB due to their height. Note that due to the style of connectors on these two devices (post-and-hole), they do not come in direct contact with the PCB. Surface mount pads have been placed on the PCB for these connectors so that wire can be soldered to the connectors and the pads. This arrangement allows for flexibility in the placement of the switch and power jack. For example, they may be placed horizontally instead of vertically as they are now.

J. INTERNAL SIGNAL GENERATION

This section will describe the physical packages selected for the components depicted in the Figure 44 schematic diagram.

1. PLD

Due to the addressing scheme used, a PLD with only 22 I/O pins is required to perform the internal logic generation depicted in Appendix C. The Atmel ATF22LV10C 24-pin TSSOP device has been selected.

2. 74HCT08

As described in Chapter III, a quad 2-input AND 74HCT08 has been utilized for some of the internal signal generation. Fairchild Semiconductor's TSSOP14 74VHCT08 has been selected to perform this function.

3. Capacitors

To provide bypassing for the PLD and the 74HCT08, two 0.1 μ F surface mount EIA 0603 ceramic capacitors have been selected.

4. PCB Layout

Figure 60 shows the PCB layout of the signal generation circuit. The 74HCT08 and ATF22LV10C, U03 and U04 respectively, were situated on the board based on routing density.

K. POWER PLANES

This section will describe the four power planes in the TSAS microcontroller design: GND1, GND2, GND3, and VCC.

1. Ground Planes

There are three ground planes in the PCB: GND1, GND2, and GND3. The ground planes have been electrically connected through the use of plated thru-holes, and vias. Figure 61 shows how the ground planes are laid out (all three are identical). As discussed in the MIL-STD-1553 section above, the analog section of the Mini-ACE circuit (in the upper right portion of the diagram) does not have any embedded ground planes so as to prevent the transmission of analog noise into the digital circuits.

To reduce the vertical size of vias, component ground pins were generally connected to the closest ground plane. A few exceptions were made so as to ensure that the three ground planes were electrically connected across the surface of the board.

2. VCC Plane

The TSAS microcontroller has two voltage level requirements, 3.3V and 5V. Figure 62 shows how the one VCC plane has been split into two planes. The area inside the irregular shape on the left half of the board is the 3.3V plane and the area outside is the 5V plane. Its shape has been determined by the pins that it must support. For example, most of the EASI 3.3V pins are concentrated toward the middle of the connector. Figure 63 is a clearer representation of the split between the two planes. The line separating the two planes is 0.015" of anti-etch, termed a "void." The void ensures electrical separation of the two planes.

L. SUMMARY

The TSAS microcontroller design utilizes some of the smallest, most power efficient components available on the market today. These components have been mounted

on both sides of an eight-layer printed circuit board. The end-result is that the TSAS microcontroller printed circuit board is only 4.2" x 3.5" x 1.0". The encased product will be slightly larger, approximately 4.3" x 3.6" x 1.2". The TSAS microcontroller described herein is truly palm-sized and will easily fit into a flight suit breast pocket.

V. CONCLUSIONS AND RECOMMENDATIONS

Spatial Disorientation is a significant problem that costs the DoD dearly each year in lost lives and equipment. Much research has been conducted in the area of tactile communications as a means of combatting SD, with very positive results. Applying this research, the Naval Aerospace Medical Research Laboratory has specified a requirement for a palm-sized, light-weight, battery-powered microcontroller capable of interfacing with sensors and displaying the received information via an array of tactors mounted in a vest. The Tactile Situational Awareness System microcontroller presented in this thesis delivers exactly this product. Small and light-weight, and yet computationally powerful, the TSAS microcontroller can interface with the MIL-STD-1553 data bus and a wide variety of serial communication protocol devices. Extremely easy to operate (just plug in the cables and turn it on), the device may be unobtrusively worn by the operator. The TSAS microcontroller / Tactor Interface Chip combination are indeed a powerful duo that are poised to take the TSAS concept out of the prototype stage and place this important communications technology into the hands of the war fighter.

This chapter will discuss the conclusions and recommendations that resulted from the present work.

1. TSAS Microcontroller Design Advantages

There are many advantages to the TSAS microcontroller design presented in this thesis. First, the TSAS microcontroller utilizes a COTS processor as the heart of its system,

thereby speeding development time and reducing risk. The design also takes advantage of some of the smallest, most power-efficient supporting components available on the market.

The TSAS microcontroller design is also based on the well-known and understood IBM PC architecture, which has many advantages in and of itself. First, the PC architecture is supported by a very large number of software and hardware manufacturers, and therefore enjoys heightened price competition and continual product improvements. Second, finding expertise to maintain the system is greatly simplified because the PC architecture is so widely utilized. Third, as the current TSAS software has been developed on x86-based machines, porting the code to the TSAS microcontroller should not present any problems.

The TSAS microcontroller design is very robust: almost all of the components (with the exception of the Mini-ACE) are made by multiple manufacturers, thereby increasing product availability and competitiveness. The design is also scalable; the end-user can add or subtract memory and processing power to suit their particular needs, with no changes to the design.

The end result is a design that meets the specifications, and yet is affordable: it is estimated that the final cost per unit will be less than \$2000. Compared to the previous design [Ref. 3], the TSAS microcontroller is 62 percent smaller in board surface area, and is at least three times faster.

2. Further Development Efforts

The work presented in this thesis represents the vast majority of theoretical hardware development required to make the TSAS microcontroller a reality. Additional time could perhaps be invested to improve the physical layout and routing of the various

components of the design so as to reduce EMI even further. Little more can be accomplished beyond this - it is time to purchase components and begin the actual hardware development. There is also significant software development required to drive the various components (e.g., the UART and Mini-ACE).

3. Recommendations

Hopefully, this thesis has conveyed a convincing argument as to the feasibility of tactile communications, particularly as embodied in the TSAS microcontroller/TIC combination. Any further development will require funding to support salary and purchase of hardware. It is not recommended that any further work be performed on the TSAS microcontroller design without funding in place - there is little more that can be done without purchasing component hardware and supporting development equipment.

Once funding has been received, it is recommended that the components listed in Appendix E be purchased in sufficient quantities to build two TSAS microcontrollers. To facilitate and expedite development, it is recommended that Cell Computing's CardPC Evaluation Board also be purchased. Although much time and effort has gone into the PCB layout of the TSAS microcontroller, it is not recommended that the PCB be manufactured and utilized for initial development work. The TSAS microcontroller design is complex, and trying to uncover the source of any errors via an eight-layer circuit board would be very difficult indeed. It is recommended that the design be breadboarded first, one function at a time. It is also recommended that software be developed concurrently, preferably by the same individual/team that is assembling the microcontroller.

APPENDIX A. MOTOROLA MPC860 POWERQUICC OVERVIEW

A. INTRODUCTION

The Motorola MPC860 PowerPC QUad Integrated Communications Controller (PowerQUICC) and its derivatives are one-chip integrated microprocessor and peripheral combinations that are designed primarily for data communications and networking applications. Previous work on a microcontroller for the Tactile Situational Awareness System [Ref. XX] had utilized the MPC860 as the main CPU of the design. To verify that the MPC860 met the requirements of the TSAS microcontroller, and to ensure that the miniaturized controller system design utilized the MPC860 correctly, this processor was studied extensively. Although it turns out that the MPC860 is not suitable as a CPU for the TSAS microcontroller, it is nonetheless a very powerful and capable chip. As its study comprised a portion of this thesis, the results of the research are recorded here.

The MPC860 is a complex microprocessor, with many and varied capabilities (the User's Manual alone is over 1100 pages). Thus, a detailed presentation of its capabilities will not be possible in this brief overview; the focus here will be on the primary elements of the Instruction Set Architecture (ISA) and the current implementations of the processor.

B. METHOD OF APPROACH

Numerous documents (Refs 20, 67, 68, 69, 70, and 71) were studied in preparation for the writing of this report. Ref. 68, a CD version of Motorola's 4-day training course on the MPC860, was particularly useful. The characteristics of the ISA and implementations

of this chip that are discussed in this overview reflect their applicability to the TSAS microcontroller. Given the factual nature of the architectural specification, this report paraphrases or reproduces the material given in some of the referenced Motorola documentation.

C. RESULTS

1. Overview

Motorola bills the MPC860 as the next-generation MC68360. It is therefore useful to include a quick description of what the MC68360 was intended for and is capable of.

The MC68360 Quad Integrated Communications Controller (QUICC) is an integrated microprocessor and peripheral chip that was designed for communications applications. There are four serial communications controllers (thus the "quad"), two serial management controllers, and one serial peripheral interface on this device. The chip has three main modules: the CPU32+ core (a 32-bit version of the CPU32), the System Integration module (SIM60), and the Communications Processor Module (CPM). Some key features of the 68360:

- Up to 32 address lines (at least 28 available)
- Up to 32-bit data bus (dynamic bus sizing for 8 and 16 bits)
- Slave mode to disable CPU32+ (enables use with external processors)
- Memory controller (eight banks)
- Four independent timers
- Two independent DMAs
- Four baud rate generators
- Plus the controllers and interface mentioned above.

The MPC860, by contrast, retains all of the capability of the MC68360 but yields higher performance, integration, and functionality. True to industry form, one of the main advertising features of the MPC860 is its backwards compatibility with the MC68360. The MPC860 still has three main modules, as depicted in Figure 27: the Embedded PowerPC Core (EPPC), the Communications Processor Module (CPM), and the System Interface Unit (SIU). All three modules utilize the 32-bit internal bus. As can be seen, the MPC860 is comprised of two processors (i.e., it is a dual processor architecture): the 32-bit PowerPC, and the 32-bit RISC in the CPM. The purpose of the 32-bit RISC processor is to offload




Figure 27. MPC860 Block Diagram [Ref. 68]

peripheral tasks from the PowerPC core; thus, the chip reportedly consumes less power overall.

Each of the three main modules will be discussed in turn.

2. Embedded PowerPC Core (EPPC)

The adoption of the PowerPC processor over the CPU32+ processor is undoubtedly the largest single revision from the MC68360 architecture. The PowerPC core is compliant with the *Book 1 Specification for the PowerPC Architecture*, so it possesses much of the same functionality as the PowerPC-based chips commonly utilized in personal computers. There are, however, some significant differences that will be mentioned.

Figure 28 depicts a block diagram of the EPPC. It is a 32-bit implementation of the PowerPC architecture. Unlike most PowerPC chips, which are superscalar, this is a single issue machine - only one instruction is fetched into the 4-deep instruction queue each clock cycle. Branches can be processed and executed in parallel with instructions that must be dispatched to an integer or load/store execution unit, but the branches must also enter the instruction queue (to preserve program order). Due to the four-stage pipeline (fetch, decode, read + execute, and write-back), an instruction may be dispatched to an execution unit every clock cycle. In order to be dispatched, however, there must be room for the instruction (including branches) in the six-deep Completion Queue (CQ). The CQ maintains a history of executed instructions, complete with processor state information that can be affected by the instruction's execution. This ensures a precise exception model (to be described later), and allows out of order completion of instructions. An instruction can only be retired after all instructions ahead of it have been retired, and after the instruction




Figure 28. Block Diagram of the Embedded PowerPC Core (EPPC) [Ref. 68]

has updated its destination register without causing an exception. Although the CQ can hold six instructions, only four can be integer instructions (the other two can be condition code or branch instructions). Long latency instructions (such as integer divide instructions or load/store instructions) may cause the CQ to stall, thereby stalling dispatch until the long

latency instruction vacates the completion queue.

Static branch prediction is implemented to help reduce latency caused by branch misprediction. Instructions fetched under a predicted branch are flushed if the branch was mispredicted, and the correct target stream fetched.

The MPC860 has thirty-two 32-bit General Purpose Registers (GPRs), but notably lacks any Floating Point Registers (floating point arithmetic is not supported). The lack of a Floating Point Unit in the MPC860 is cause for concern, as TSAS microcontroller application software must perform floating point calculations [Ref. 15]. There are some workarounds for this limitation (e.g., FPU software emulation), but they lead to slower processing and increased design complexity.

Register renaming is employed. The registers required to support the PowerPC user instruction set (minus floating point instructions) are implemented, to include the Condition Register (CR), Count Register (CTR), and Link Register (LR) (contained in the Branch Processing Unit), and the integer exception register (XER) (contained in the Integer Unit). The MPC860 also employs a number of other PowerPC-compliant supervisor level and implementation-specific registers.

The Integer Unit implements the usual types of integer instructions:

- Arithmetic instructions
- Compare instructions
- Trap instructions
- Logical instructions
- Rotate and shift instructions

The Load/Store Unit (LSU) transfers data between the GPR's and the internal bus. By implementing the LSU as a separate execution unit, stalls in the memory pipeline do not affect the master instruction pipeline (unless there is a data dependency). There is a two-entry load/store instruction address queue and a two-entry 32-bit integer data queue, as depicted in Figure 29. Data is written directly back to the GPR's via a dedicated writeback bus.




Figure 29. LSU Functional Block Diagram [Ref. 20]

The MPC860 implements most PowerPC instructions in its Instruction Set (with the exception of Floating Point instructions, as previously mentioned). All PowerPC operand conventions are adhered to.

Exception handling is transparent to user software, and is handled in supervisor mode. The MPC680 implements a precise exception model, as specified in *PowerPC Family: The Programming Environments for 32-Bit Implementations*. The precise exception model ensures that the out-of-order execution of instructions does not affect data integrity during the processing of an exception. This is accomplished by allowing instructions in the CQ to complete when an exception is encountered before instruction processing transfers to the appropriate exception vector. If one of these completing instructions generates another exception, then that exception is handled first. Before control passes to the exception handler, the machine state is preserved by storing the instruction address at which fetching should resume in Save/Restore Register 0 (SRR0) and Machine State Register (MSR) bits in SRR1. When the exception is taken, the instruction causing the exception might not have started executing, could be partially executed, or may have completed, depending on the exception and instruction types as per Table 6. Partially completed instructions can be re-executed after the exception is handled, thereby simplifying exception processing because software does not have to save the machine's internal states.

Exception Type	Instruction Type	Before/After	Contents of SRR0
Hard reset (caused by HRESET or SRESET)	Any	NA	Undefined
System reset	Any	Before	Next instruction to execute
Machine check	Any	Before	Faulting instruction
TLB miss/error ¹	Any	Before	Faulting fetch or load/store
Other noninstruction-related exceptions	Any	Before	Next instruction to execute
Alignment	Load/store	Before	Faulting instruction
Privileged instruction	Any privileged instruction	Before	Faulting instruction
Trap	tw, twi	Before	Faulting instruction
System call	sc	After	Next instruction to execute
Trace	Any	After	Next instruction to execute
Debug I- breakpoint ¹	Any	Before	Faulting instruction
Debug L- breakpoint ¹	Load/store	After	Faulting instruction + 4
Software emulation ¹	NA	Before	Faulting instruction
Floating-point unavailable	Floating-point	Before	Faulting instruction

¹ Implementation-specific exceptions not defined by the PowerPC architecture

Table 6. Before and After Exceptions [Ref. 20].

The MPC860 contains separate 4-Kbyte, two-way set associative instruction and data caches. Figure 30 depicts the layout of the instruction cache (the data cache is similar).

Both caches use a Least-Recently Used (LRU) replacement algorithm. The caches are physically addressed, with the physical address tag stored in the cache directory. Cache lines are 16 bytes long. During a cache miss, the burst-fill is performed critical-word first. Both caches can be disabled, invalidated, or locked by writing to their respective control registers. Locking a cache lets the programmer ensure that frequently used instructions or data are resident in the respective cache.




Figure 30. Instruction Cache Organization [Ref. 20]

The MPC860 does not support instruction cache snooping, and therefore software must maintain instruction cache coherency. The instruction cache has a 32-bit interface to the instruction sequencer, thereby allowing high throughput to the four-entry instruction queue. Each instruction cache line is loaded from a four-word boundary (i.e., address bits A[28-31] of the effective address are 0). A[21-27] index a set, and A[28-29] select a word within the cache line. The tags consist of the top 21 bits of the physical address (PA[0-20]).

There are three special purpose registers to control the instruction cache: the instruction control and status register, the instruction cache address register, and the instruction cache data port register.

There are three data cache coherency states: modified-valid, unmodified-valid (or 'exclusive'), and invalid; however, the MPC860 does not provide support for snooping external bus activity. Therefore, coherency between the data cache and external memory or I/O devices must be maintained by software. The data cache is directly coupled to the LSU to allow efficient movement of data to and from the GPR's. Each data cache line is loaded from a four-word boundary (i.e., address bits A[28-31] of the effective address are 0). A[21-27] index a set, and A[28-31] select a byte within the cache line. The tags consist of the top 21 bits of the physical address (PA[0-20]). There are three special purpose registers to control the data cache: the data control and status register, the data cache address register, and the data cache data port register.

The MPC860 utilizes Memory Management Units (MMU) to implement virtual memory management (cache control, storage access protections, and effective to real address translation). There are two MMU's, one each for data and instructions. Unlike the PowerPC architecture, which specifies 4 KByte pages only, the MMU's of the MPC860 support 4, 16, and 512 KByte, and 8 MByte pages. Further protection flexibility is offered through the support of 1 KByte subpage protection when using 4 KByte pages. The data and instruction translation lookaside buffers (DTLB and ITLB, respectively) are 32-entry and fully associative. Four entries in each TLB can be locked, thereby ensuring fast translation for selected pages. In the event that no more than 32 pages of instructions or data are

required for an application, the TLB's can be initialized from reset. Instruction and data address translation can be disabled separately, in which case the effective address is also the physical address. In terms of performance, a data or instruction cache hit will take one clock cycle if the access is from the same page as the last access. If an instruction access is from a different page than the last access, but still hits in the ITLB, then the access will take two cycles. Segment registers are not defined.

Figure 31 depicts the flow diagram for the MMU's (assuming that translation is enabled). The EPPC generates a 32-bit effective address (EA) for memory accesses. As discussed earlier, if the current access is on the same page as the last access, then the work is already done and the algorithm is exited almost immediately (not depicted above). If the current access refers to a different page than the last access, the MMU checks to see if there is an address match in the TLB. If there is a hit, and there is a protection match, then the MMU asserts the real address. If there is not a hit, then a TLB miss interrupt is generated, and a tablewalk results. If the protection does not match, a TLB error results.

Figure 32 shows how effective addresses are translated to physical addresses in the case of 4 KByte pages. If comparison of the PR (privilege level) bit of the Machine State Register (MSR), the MMU Current Address Space ID Register (M_CASID), and the upper 20 bits of the EA hit in the TLB, then the contents of the physical page number are concatenated with the 12 least significant bits (in the case of 4 KByte pages) of the EA to form the physical address that is sent to the cache and memory system. Note that the M_CASID is only compared with the TLB entry's ASID when the matching entry is programmed as unshared.




Figure 31. MMU Flow Diagram [Ref. 20]




Figure 32. Effective-to-Physical Address Translation for 4 Kbyte Pages [Ref. 20]

The MMU hardware supports a two-level software tablewalk as depicted in Figure 33 below. The tablewalk procedure is relatively straightforward, and is similar to many other popular processors. An MPC860-specific register, the MMU Tablewalk Base Register (M_TWB) provides the base address of the Level-1 Table Base. This 1024-entry table is indexed by the most significant 10 bits of the EA. 20 bits of the 4-byte Level-1 Table entry point to the base address of the Level-2 Table Base, which is indexed by the next 10 bits of the EA for a 4 KByte page.




Figure 33. Two-Level Translation Table [Ref. 20]

Tables 7 and 8 below detail the full contents of the Levels 1 and 2 Descriptors.

Bits	Name	Description
0–19	L2BA	Level-2 table base address. Bits 18–19 should be 0b00 unless MD_CTR[TWAM] = 1.
20–22	—	Reserved
23–26	APG	Access protection group
27	G	Guarded memory attribute for entry 0 Nonguarded memory 1 Guarded memory
28–29	PS	Page size level one 00 Small (4 Kbyte or 16 Kbyte) 01 512 Kbyte 11 8 Kbyte 10 Reserved
30	WT	Writethrough attribute for entry 0 Copyback cache policy region (default) 1 Writethrough cache policy region
31	V	Level-one segment valid bit 0 Segment is not valid 1 Segment is valid

Table 7. Level-1 Descriptor Format [Ref. 20].

Bits	Name	4-Kbyte Pages AND Mx_CTR[PPM] = 1		Larger than 4-Kbyte Pages OR Mx_CTR[PPM] = 0	
0–19	RPN	Physical (real) page number			
20–21	PP	Protection for 1st 1-Kbyte subpage in a 4-Kbyte page	For Instruction Pages <u>Supervisor</u> <u>User</u> 00 No access No access 01 Executable No access 1x Executable Executable For Data Pages <u>Supervisor</u> <u>User</u> 00 No access No access 01 R/W No access 10 R/W R/O 11 R/W R/W	For Instruction Pages <u>Supervisor</u> <u>User</u> Extended encoding: 00 No access No access 01 Executable No access 1x Reserved PowerPC encoding: 00 Executable No access 01 Executable Executable 1x Executable Executable	For Data Pages <u>Supervisor</u> <u>User</u> Extended encoding: 00 No access No access 01 R/O No access 1x Reserved PowerPC encoding: 00 R/W No access 01 R/W R/O 10 R/W R/W 11 R/O R/O
22	PP	2nd 1-Kbyte subpage		0 Bits 20–21 contain PowerPC encoding 1 Bits 20–21 contain extended encoding	
23				C—Change bit for entry. Set to 1 by default if change tracking functionality is not desired. 0 Unchanged region (write-protected) 1 Changed region, write allowed	
24–25	3rd 1-Kbyte subpage	MD_CTR[PPCS] = 0. For 1 Kbyte pages in mode 3, set to the appropriate subpage validity (see Section 9.5, "Protection Resolution Modes"). Otherwise, set to 0b1111.		MD_CTR[PPCS] = 1 1000 Hit (supervisor accesses only) 0100 Hit (user accesses only) 1100 Hit for both	
26–27	4th 1-Kbyte subpage				
28	SPS	Clear		Small page size. Valid only when L1 descriptor[PS] = 00 0 4 Kbyte 1 16 Kbyte	
29	SH	Shared page 0 Entry matches only if a TLB entry's ASID field matches the value in M_CASID. 1 ASID comparison is disabled for the entry.			
30	CI	Cache-inhibit attribute for the entry.			
31	V	Page valid bit			

Table 8. Level-2 Descriptor Format [Ref. 20].

3. System Interface Unit (SIU)

The MPC860 SIU performs a system integration and control function. A block diagram of the SIU is given in Figure 34. The SIU controls system startup, initialization and operation, and protection. It controls memory devices, and provides an interface between the internal Unified Bus and the external system bus. The SIU also provides power management functions, clock and timer functions, debug support, and a PCMCIA Interface.




Figure 34. System Interface Unit Block Diagram [Ref. 68]

The memory controller supports up to eight memory banks with glueless interfaces to many different memory devices (DRAM, SRAM, SSRAM, EPROM, Flash EPROM, SRDRAM, and EDO) and peripherals. It can provide 0 to 15 wait states for each bank of memory, and has 4 byte enable signals for varying-width devices. It also has an output enable signal, and a boot chip select signal that is available on reset. The DRAM interface can support port sizes of 8, 16, and 32-bits. A limitation to the “glueless” interface is that only one bank of DRAM can be supported without requiring external buffers.

The system configuration and protection function controls the overall system. System configuration and protection features include a bus monitor, a software watchdog timer, interrupt control, a periodic interrupt timer, a PowerPC timebase counter (in compliance with PowerPC architecture) and decrementer, and a real time clock with alarm

register. A clock synthesizer, which allows the user to multiply a low input clock frequency provided by a crystal or external oscillator, generates clock signals for the EPPC, CPM, and external devices.

The SIU supports a variety of power management features to include Full On, Doze, Sleep, Deep Sleep, and Low Power Stop. Each level has less functionality and therefore consumes less power than the level above it, but requires a longer wake-up time.

Lastly, there is a PCMCIA host adapter module that provides control logic for a PCMCIA socket interface. The designer must provide external analog power and buffering. Two sockets (ports) are supported. If only one PCMCIA port is being used, the other port may be used as a general-purpose input with interrupt capability.

4. Communications Processor Module (CPM)

The last of the three main modules of the MPC860, the CPM, is designed specifically for facilitating the use of the MPC860 in communications and networking applications. It supports multiple communications channels and protocols. The CPM offloads many tasks from the PowerPC core. For example, the core is interrupted only upon frame reception or transmission, rather than on a per-character basis. Figure 35 is a block diagram of the CPM. The CPM is comprised of three main submodules: the Communications Processor (CP), two independent DMA controllers, and 4 general purpose timers. Each submodule will be discussed in turn.




Figure 35. Communications Processor Module Block Diagram [Ref. 20]

The CP can communicate with the EPPC in a number of different ways: it can utilize the dual-port RAM to exchange parameters, it can issue commands to the core through the CP Command Register (CPCR), it can generate interrupts via the CPM interrupt controller (CPIC), and it lets the core read the CPM status/event registers at any time.

The CP communicates with peripherals via the peripheral bus. It is optimized for serial communications, and is comprised of a 32-bit RISC processor, four Serial Communication Controllers (SCC), two Serial Management Controllers (SMC), one Serial Peripheral Interface (SPI), one Inter-Integrated Circuit (I²C) Interface, 5 KBytes of dual-port

(EPPC – CPM) RAM, an interrupt controller, a time slot assigner, three parallel ports, a parallel interface port, and four independent baud rate generators. There are also sixteen virtual serial DMA (SDMA) channels to support the transmit and receive channels of the SCCs, SMCs, SPI, and I²C.

The four SCCs, which are the most powerful of the eight communications devices, can be independently configured to implement a variety of protocols for bridging functions, routers, or gateways, or to interface with a wide variety of standard networks. The SCCs have separate receive and transmit FIFOs; those of SCC1 are 32 bytes each, and those of SCC2-SCC4 are 16 bytes each. The following protocol modes are supported:

- HDLC/SDLC
- HDLC Bus
- UART/Synchronous UART
- Asynchronous HDLC
- IrDA
- BiSync
- Signaling System #7
- AppleTalk/LocalTalk
- Transparent Mode (or on some of the MPC860 family members, Ethernet)

The two SMCs are not as flexible as the SCCs - they can be independently configured to support one of only three protocols: UART, Transparent, or General Circuit Interface (GCI). The UART protocol can be used to provide a debug/monitor port, thus freeing up an SCC. The SMCs support loopback and echo modes for testing purposes. The

SMC receiver and transmitter are doubled-buffered, creating an effective FIFO size of two characters.

The SPI is a full-duplex, synchronous, character-orientated channel. It supports a four-wire interface: receive, transmit, clock, and slave select. The SPI allows the MPC860 to communicate with other MPC860 chips, or a variety of other processors and peripheral devices. During an SPI transfer, data is sent and received simultaneously. The I²C controller allows the MPC860 to communicate with other I²C devices, to include SIMMs and DIMMs. It uses a synchronous, multimaster bus that can connect several integrated circuits on a board.

The CP has four general-purpose parallel I/O ports, labeled A, B, C, and D. All of the ports are bi-directional. Port B can be used as a parallel interface port, thereby allowing easy connection to Centronics interfaces.

The second submodule of the CPM is comprised of two physical serial DMA (SDMA) channels. These two physical channels are used by the CP to implement sixteen virtual SDMA channels – one each for the serial controllers' (SCCs, SMCs, SPI and I²C) transmitters and receivers. The two physical SDMA channels are also used by the CPM to emulate two general purpose independent DMA (IDMA) channels. The IDMA channels can be used for memory to memory and peripheral to memory transfers.

The last submodule of the CP is comprised of four general purpose timers. The timers are each 16 bits, and two can be cascaded to form a 32-bit timer. At 25MHz, the timers have a maximum period of 10.7 seconds, and a 40ns resolution.

5. Implementations

An attempt to obtain physical characteristics (e.g. number of transistors, process size etc.) on the MPC860 was made, but information was noticeably lacking in this regard. What has been determined is that the basic architecture of the MPC860 has remained unchanged since it was introduced in 1995. The only improvement to speak of has been the clock speed of the device; originally available at 25MHz, there are now 40, 50, and 66MHz versions (the 66MHz version having been released in June 1998). Rather than improving the basic architecture, Motorola seems to have focussed its efforts on derivatives of the MPC860 that offer varied capability to support peripherals. In this way, one can select the derivative that most closely meets one's intended application. The following table lists the derivatives of the MPC860 that are currently available:

Device	Ethernet Support	Number of SCCs	64-Channel HDLC Support
MPC860DC	SCC1	2	N/A
MPC860DE	Yes	2	N/A
MPC860DH	Yes	2	Yes
MPC860	N/A	4	N/A
MPC860EN	Yes	4	N/A
MPC860MH	Yes	4	Yes
MPC860SAR	Yes	4	Yes
MPC860T	10/100	4	Yes
MPC860DT	10/100	2	Yes

Table 9. MPC860 Derivatives [Ref. 71].

Of note, all of the above versions of the MPC860 are available in Ball Grid Array (BGA) packaging, thereby allowing a higher pin-count in a more compact format than a standard quad flat pack. The BGA for the MPC860 has 357 pins, and is approximately one square inch in size.

D. SUMMARY

The MPC860 and its family are versatile chips that are useful in numerous communications and networking applications. Some of the current embedded applications include:

- T1/E1 interface lines
- simultaneous voice and data
- LAN to WAN connections
- Ethernet bridges and routers
- PBX switches
- fault-tolerant LANs
- smart Ethernet hubs
- remote access routers
- ATM line card controllers
- branch office and multiprotocol routers
- cellular base stations
- central office switches
- primary rate ISDN protocol termination
- set top boxes

The MPC860 has functions similar to many other mainstream processors, but its dual-processor architecture and communications functions make it unique. The following diagram is an example application of the MPC860, taken from Ref 68. The details of this

What is an Example Application?




Figure 36. Example MPC860 Application [Ref. 68]

particular application will not be discussed here; it is included only to give a flavor of how the MPC860 can be used.

E. CONCLUSIONS

The MPC860 is a powerful and versatile chip that will meet a wide variety of communications needs. For example: its Serial Communications Controllers (SCC) are capable of transmitting and receiving UART protocol, and its System Interface Unit allows generation of clock signals for external devices. In terms of the TSAS microcontroller's interface with the TIC, these features negate the requirement for two ICs (a UART, and an oscillator). The MPC860 also has a number of other desirable features, to include glueless

interfacing to many types of memory and peripherals, easy-to-understand PowerPC architecture, and solid customer support such as example system designs (e.g., Ref. 21).

The MPC860 does have some limitations however, particularly as applied to the TSAS microcontroller. Unlike the majority of PowerPC chips being produced today, which are superscalar, the MPC860 is single issue. It also has no Floating Point Unit (FPU), and thus does not support any PowerPC floating point instructions. Lastly, the MPC860 currently has a top clock speed of only 66 MHz. While these performance specifications may be fine for most communications applications, the TSAS microcontroller is required to perform Euler angle transformations using sine and cosine functions [Ref. 15]. Although floating point instructions could be emulated in software, doing so would cost additional clock cycles. As the current prototype TSAS controller, an i486SX25 (which has an FPU), can take up to 5 minutes to perform some EVA calculations [Ref. 15], doubts about the MPC860's suitability for this task have arisen. It has therefore been decided not to utilize the MPC860 in the TSAS microcontroller.


APPENDIX B. TSAS MICROCONTROLLER SCHEMATICS

This appendix includes all of the schematic diagrams for the TSAS microcontroller design. The first diagram is an index page for the rest of the diagrams.

	4	3	2	1	
D	TACTILE SITUATION AWARENESS SYSTEM (TSAS) MICROCONTROLLER DESIGN PAGE 2 - CARDPC INTERFACE PAGE 3 - MIL-STD-1553 BUS INTERFACE PAGE 4 - TIC INTERFACE PAGE 5 - RS-232 INTERFACE PAGE 6 - SOLID STATE DISK PAGE 7 - DC TO DC POWER CONVERSION PAGE 8 - SIGNAL GENERATION				D
C					C
B					B
A					A
	4	3	2	1	


NAVAL POSTGRADUATE SCHOOL, MONTEREY CA	
TITLE: TSAS MICROCONTROLLER	DATE: 19 APR 99
ENGINEER: TERRENCE WOOD	PAGE: 1 OF 8

Figure 37. TSAS Microcontroller Circuit Schematics Index



TITLE: CARDPC INTERFACE		DATE: 19 APR 99
ENGINEER: TERENCE HOOD		PAGE: 2 OF 8

Figure 38. CardPC Interface Circuit Schematic






Figure 40. TIC Interface Circuit Schematic




Figure 42. Solid State Disk Circuit Schematic




Figure 43. DC-to-DC Power Conversion Circuit Schematic




Figure 44. Signal Generation Circuit Schematic

APPENDIX C. TSAS MICROCONTROLLER PLD LOGIC

The following schematic diagram shows the logic that is to be burned into the Atmel ATF22LV10C programmable logic device.




Figure 45. PLD Logic Schematic

APPENDIX D. TSAS MICROCONTROLLER PCB LAYOUT

This appendix presents a number of different views of the TSAS microcontroller PCB layout. The first figure, Figure 46, is to scale and details all facets of the design. The figure is transparent; that is to say, both sides of the board are simultaneously presented. The Reference Designators (e.g. Y02 or U10) of components mounted on what will be referred to as the bottom (CardPC) side of the board are therefore mirrored.




Figure 46. TSAS Microcontroller PCB, Showing Both Sides and All Layers of Etch




Figure 47. TSAS Microcontroller PCB, Top and Bottom Side, Without Etch




Figure 48. TSAS Microcontroller PCB, Top Side




Figure 49. TSAS Microcontroller PCB, Bottom Side




Figure 51. TSAS Microcontroller, Showing S2 Signal Layer Etch




Figure 52. TSAS Microcontroller, Showing S3 Signal Layer Etch




Figure 53. TSAS Microcontroller, Showing BOTTOM Signal Layer Etch




Figure 54. CardPC Interface Component Layout




Figure 55. MIL-STD-1553 Data Bus Interface Component Layout

TSAS MICROCONTROLLER
NAVAL POSTGRADUATE SCHOOL
COPYRIGHT 1999




Figure 56. TIC Interface Component Layout




Figure 57. RS-232 Interface Component Layout

TSAS MICROCONTROLLER
NAVAL POSTGRADUATE SCHOOL
COPYRIGHT 1999




Figure 58. Solid State Disk Component Layout




Figure 59. DC-DC Power Conversion Component Layout




Figure 60. Internal Signal Generation Component Layout




Figure 61. TSAS Microcontroller Ground Plane Layout




Figure 62. TSAS Microcontroller VCC3 and VCC5 Planes Layout

APPENDIX E. TSAS MICROCONTROLLER BILL OF MATERIALS

This appendix lists specific information on all of the components used in the TSAS microcontroller design.

Component	Qty	Value	Manufacturer	Part Number
C01-C03, C06-C10	8	220pF	Kemet	C0603C221J5GAC
C19, C23, C27, C29-C31, C33, C35, C37, C38, C40, C42, C47	13	0.01 μ F	Kemet	C0603C103K5RAC
C12-C16, C24, C26, C34, C44-C46	11	0.1 μ F	Kemet	C0603C104K4RAC
C05	1	0.47 μ F	Kemet	C1206C474K4RAC
C20-C22, C25, C28, C32, C36, C39, C41, C43	10	10 μ F	Kemet	T491B 106K016 AS
C04, C11	2	68 μ F	Vishay Sprague	595D686X0016C2T
C17, C18	2	220 μ F	Vishay Sprague	595D227X06R3C2T
D01, D02	2		Motorola	MBRS340T3
J01, J04	2		Trompeter	BJ450
J02	1		Samtec	TMM-108—01-G-D-TH
J03	1		Positronic	HDC9M5B3000
J05	1		Molex	85502-0009
L01, L02	2	22 μ H	Coilcraft	D03316
LED01	1		Any	Thru-Hole
P01	1		CUI Stack	PJ-010A
R16, R17	2	0.04 Ω	Vishay Dale	WSL-1206
R15	1	470 Ω	Vishay Dale	CRCW0603471JRT1
R01-R07, R09-R14	13	4.7K Ω	Vishay Dale	CRCW0603472JRT1
R08	1	200K Ω	Vishay Dale	CRCW0603204JRT1
S01	1		C&K	1101M1S3ZQE2
T01	1		ILC DDC Beta	TST-9007
U01	1		Harris	HIN-211IX
U02	1		ILC DDC	BU-61588G3-212
U03	1		Fairchild	74VHCT08MTC
U04	1		Atmel	ATF22LV10C-15XI
U05	1		Texas Instrmnt	TL16C550C-FN
U06	1		M-Systems	MD-2200-D24-V3-X-PI
U07, U09	2		Maxim	MAX1626ESA
U10	1		Cell Computing	C2I-EZ-TH
Y02	1	1MHz	Vishay Dale	XOSM-52BR1
Y01	1	16MHz	Vishay Dale	XOSM-52BR16

LIST OF REFERENCES

1. Naval Aerospace Medical Research Laboratory and Coastal Systems Station, *Tactile Situational Awareness System (TSAS)*, a presentation posted at www.tsas.namrl.navy.mil.
2. US Navy, *Mission Need Statement for Systems to Enhance Aircrew Situational Awareness*, Serial Number MO93-88-98, August 1998.
3. Luke, B., "Design of a Microelectronic Controller with a MIL-STD-1553 Bus Interface for the Tactile Situational awareness System," Master's Thesis, Naval Postgraduate School, Department of Electrical and Computer Engineering, September 1998.
4. Zlotnik, M., "Applying Electro-Tactile Display Technology to Fighter Aircraft -- Flying With Feeling Again," *Proceedings of the National Aerospace and Electronics Conference*, IEEE, pp. 191-197, 1988.
5. Gillingham, K. and Wolfe, J., "Spatial Orientation in Flight," USAF document USAFSAM-TR-85-31, USAF School of Aerospace Medicine, Brooks AFB, TX, December 1986.
6. McGrath, B., et al., "The Role of Intelligent Software in Spatial Awareness Displays," *Proceedings for the Third Annual Symposium and Exhibition on Situational Awareness in the Tactical Air Environment*, Piney Point, MD, June 2, 1998.
7. Ballard, J.W., and Hessinger, R.W., "Human-Engineered Electromechanical Tactual Sensory Control System," *Electrical Manufacturing*, Vol. 54, pp. 118-121, 1954.
8. Link, J., "Design of a Serial Communication Protocol and Bus Interface Chip for Tactile Communications," Master's Thesis, Naval Postgraduate School, Department of Electrical and Computer Engineering, March 1999.
9. Rochlis, J., "A Vibrotactile Display for Aiding Extravehicular Activity (EVA) Navigation in Space," Master's Thesis, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Boston, MA, 1998.
10. McTrusty, T., "Swimmer Inshore Navigation System (SINS) Tactile Situational Awareness System (TSAS) Test Report," Coastal Systems Station, Panama City, FL, August 2, 1997.

11. Tan, H.Z., and Pentland, A, "Tactual Displays For Wearable Computing," *Proceedings of the First International Symposium on Wearable Computers*, IEEE, pp. 84-89, 1997.
12. Fritz, S. et al., "User Interface Design for a Travel Aid for Blind People," University of Magdeburg, Germany, 1995.
13. Nottingham University Department of Psychology, Blind Mobility Research Unit, *TouchBlasterNomad*, information posted at <http://www.psyc.nott.ac.uk/bmru/tbnom.html>.
14. Discussions between Mr. Braden McGrath/CDR Angus Rupert of the Naval Aerospace Medical Research Laboratory, and LT Brian Luke of the Naval Postgraduate School, June 20-21, 1997.
15. Discussions between Braden McGrath of the Naval Aerospace Medical Research Laboratory, and MAJ Terrence Wood of the Naval Postgraduate School, September 1998.
16. Lehrbaum, R., "Designing with PC/104 – A Tutorial," Ampro Computers Inc., document posted at <http://www.ampro.com>.
17. Cell Computing, Inc., document posted at <http://www.cellcomputing.com>.
18. Parvus Corporation, The, document posted at <http://www.parvus.com>.
19. SBS Avionics Technologies, document posted at <http://www.sbs-avionics.com>.
20. Motorola, Inc., *MPC860 PowerQUICC User's Manual*, Version 1.0, Motorola, Inc., 1998.
21. Motorola, Inc., "MPC860 Reference Design Board (SAMBA) Release Guide," Revision A, Motorola Applications Engineering, May 12, 1998.
22. ILC Data Device Corp., "ACE Series BC/RT/MT Advanced Communication Engine Integrated 1553 Terminal User's Guide," Bohemia, NY, 1992.
23. Atmel Incorporated, *8-Megabit UV Erasable EPROM Data Sheet*, Atmel Corporation, 1997.
24. E-mail messages between Mr. David Cowan of Cell Computing, Inc., and MAJ Terrence Wood of the Naval Postgraduate School, January and February 1999.
25. Cell Computing, Inc., *Pentium CardPC Hardware Reference Manual*, Document Number C2i-P5-RM-021, Version 2.1, September 1998.

26. Wakerly, J., *Digital Design Principles and Practices*, 2nd edition, Prentice-Hall Inc., Englewood Cliffs, NJ, 1994.
27. Brooks, D., "Maintaining clean power, part one," *Printed Circuit Design Magazine*, April 1997.
28. ILC Data Device Corporation, *MIL-STD-1553 Designer's Guide*, 4th ed., ILC Data Device Corporation, Bohemia, NY, 1993.
29. SBS Avionics Technologies, *An Interpretation of MIL-STD-1553B*, document posted at <http://www.sbs-avionics.com>.
30. ILC Data Device Corp., "ACE Series BC/RT/MT Advanced Communication Engine Integrated 1553 Terminal User's Guide," Bohemia, NY, 1992.
31. ILC Data Device Corp., *BU-61588/65179*/65178/61688*/61689* Miniature Advanced Communications Engine (Mini-ACE) Mini-ACE Plus Data Sheet*, Bohemia, NY, 1992.
32. Kozierok, C., *The PC Guide*, Version 1.8.1, a document posted at <http://www.pcguide.com/index.htm>, March 1999.
33. E-mail messages between Mr. Mike Glass, Senior Staff Applications Engineer at ILC Data Device Corporation, and MAJ Terrence Wood of the Naval Postgraduate School, January 1999.
34. Beta Transformer Technology Corporation, *TST-9000 Series Data Sheet*, ILC Data Device Corporation, 1994.
35. Byte Runner Technologies, document posted at <http://www.byterunner.com/why.html>.
36. Texas Instruments Incorporated, *TL16C550C, TL16C550CI Asynchronous Communications Element with Autoflow Control Data Sheet*, Texas Instruments, Dallas, TX, 1998.
37. Sokos, M., *The ISA Bus*, a document posted at <http://www.controlled.com/isa.txt>, 1997.
38. Peacock, C., *Interfacing the Serial/RS-232 Port*, Version 5.0, a document posted at http://desun1.epfl.ch/~conti/gnr/pc_doc1/serial.htm, July 1997.
39. Howe, D., *Free On-Line Dictionary of Computing*, document posted at <http://wombat.doc.ic.ac.uk/foldoc/foldoc.cgi?query=RS-232>, 1999.

40. Sangoma Technologies, Inc., *RS232, RS422 and V.35 Interface Tutorial*, a document posted at <http://165.247.199.3/signal.htm>, Markham, Ontario, August 1998.
41. Desmond, A., *Do You need Information About RS-232C?*, a document posted at <http://www.epl.co.uk/info16.htm>, EPL Ltd., May 1998.
42. E-Mags, "Electronics on the Web", Issue #1, Feature 7, a document posted at <http://www.u-net.com/epr/electron/issue1/feat0710.htm>, June 1996.
43. Harris Semiconductor, *HIN211 thru HIN213 Data Sheet*, Harris Corporation, Melbourne, FL, 1999.
44. Cell Computing, Inc., *CardPC Design Guide*, Document Number C2i-CARDPC-DG-010, Version 1.0, December 1998.
45. M-Systems, Inc., *DiskOnChip2000 MD2200, MD2201 Data Sheet*, Document Number 91-SR-002-00-8L Rev.2.1, M-Systems, June 1998.
46. M-Systems, Inc., *Designing with the DiskOnChip2000*, Document Number 91-SR-002-01-7L Rev.2.0, M-Systems, July 1997.
47. Strassberg, D., "A surfeit of power-supply voltages plagues design of compact products," *EDN Magazine*, May 1994.
48. Maxim Integrated Products, *MAX1626/MAX1627 Data Sheet*, Maxim Integrated Products, Sunnyvale, CA, June 1996.
49. Landers, T.M. et al., *Electronics Manufacturing Processes*, Prentice Hall, Upper Saddle River, NJ, 1994.
50. Atmel Corporation, *ATF22LV10C Data Sheet*, Ver. 0780G, Atmel Corporation, Colorado Springs, CO, February 1999.
51. Montrose, M.I., *Printed Circuit Board Design Techniques for EMC Compliance*, IEEE Press, New York, NY, 1996.
52. Patrick, G., *The PCB Design Process*, a document posted at <http://www.arcanet.com/~gpatrick>, October 1996.
53. Kemet Electronics Corporation, *Surface Mount – Mounting Pad Dimensions and Considerations*, Document Number F-2100D, Kemet Electronics Corporation, Greenville, SC, March 1998.
54. Patrick, G., *PCB Design Information*, a document posted at <http://www.aracnet.com/~gpatrick/pcb.html>, November 1996.

55. Brooks, D., "Ground Plane 101," Printed Circuit Design Magazine, October 1997.
56. Brooks, D., "Crosstalk, Part 2: How Loud is Your Crosstalk?," Printed Circuit Design Magazine, December 1997.
57. Brooks, D., "Right-angle Corners," *Printed Circuit Design Magazine*, February 1997.
58. Honda Connectors, FAX of ICMA-236LMY Dimensions, Honda Connectors, Mountain View, CA, 3 Mar 99.
59. Vishay Dale, *Models XOSM-52 Data Sheet*, Vishay Dale, Yankton, SD.
60. Trompeter Electronics, Inc., *Trompeter Products Catalog*, Number T20, Westlake Village, CA, December 98.
61. Positronic Industries, Inc., *Catalog of Industrial and Military Application Subminiature-D Connectors*, Positronic Industries, Inc., Springfield, MO, May 1997.
62. Coilcraft, Inc., *SMT Power Inductors – D03316 Series Data Sheet*, a document posted at <http://www.webcom.com/liebman/us/do3316.html>, Coilcraft, Cary, IL, January 1999.
63. Siliconix, Inc., *Si4963DY Dual P-Channel 2.5-V (G-S) Rated MOSFET Data Sheet*, Document S-51649 (Rev. A), Temic Semiconductors, April 1997.
64. Motorola, Inc., MBRS340T3 Surface Mount Schottky Power Rectifier Data Sheet, Document Number MBRS340T3/D (Rev. 2), Motorola, Inc., 1996.
65. Vishay Sprague, *Type 595D Solid Tantalum Chip Capacitors Data Sheet*, document posted at <http://www.vishay.com>.
66. Fairchild Semiconductor Corporation, *74VHC08 - 74VHCT08 Quad 2-Input AND Gate Data Sheet*, Fairchild Semiconductor Corporation, January 1998.
67. Motorola, Inc., *Application Note – Getting Started With the MPC860: A Design Guide*, Motorola Inc., 1998.
68. Motorola, Inc., *MPC860 PowerQUICC Training Compact Disc*, Motorola Inc., 1998.
69. Motorola, Inc., *MPC860 PowerQUICC Product Brief*, Motorola Inc., 1998.
70. Motorola, Inc., *MC68360 User's Manual*, Motorola Inc, 1993.
71. Motorola, Inc., "PowerQUICC PowerPC Integrated Communications Controllers – MPC860 Family," a document posted at <http://www.mot.com/netcomm>, July 1998.

INITIAL DISTRIBUTION LIST

	No. Copies
1. Defense Technical Information Center 8725 John J. Kingman Rd., STE 0944 Ft. Belvoir, VA 22060-6218	2
2. Dudley Knox Library Naval Postgraduate School 411 Dyer Rd. Monterey, CA 93943-5101	2
3. Chairman, Code EC Department of Electrical and Computer Engineering Naval Postgraduate School Monterey, CA 93943-5121	1
4. Curricular Officer, Code 34 Engineering & Technology Curricular Office Naval Postgraduate School Monterey, CA 93943-5109	1
5. Professor Douglas J. Fouts, Code EC/Fs Department of Electrical and Computer Engineering Naval Postgraduate School Monterey, CA 93943-5121	1
6. Randy L. Wight, Code EC/Wr Department of Electrical and Computer Engineering Naval Postgraduate School Monterey, CA 93943-5121	1
7. Director Professional Education and Development NDHQ MGen Pearkes Building Ottawa, ON K1A 0K2 CANADA	2
8. DAEPM (FT) NDHQ MGen Pearkes Building Ottawa, ON K1A 0K2 CANADA	2

9. Directorate of Flight Safety 2
NDHQ
MGen Pearkes Building
Ottawa, ON K1A 0K2
CANADA
10. MAJ Terrence L. Wood 2
c/o Capt. J. Daly
30 Dubhe Dr.
Ottawa, ON K1K 0L5
CANADA