ORNL/LTR-2012/556
NEET/ASI/ORNL/TR-2012/01

OAK RIDGE
NATIONAL LABORATORY
MANAGED BY UT-BATTELLE
FOR THE DEPARTMENT OF ENERGY

COMMON-CAUSE FAILURE MITIGATION PRACTICES AND KNOWLEDGE GAPS

October 2012

Prepared by

Richard Wood
Laura Pullum
Cyrus Smith
David Holcomb
Kofi Korsah
Michael Muhlheim

UT-BATTELLE
ORNL-27 (4-00)

NEET/ASI/ORNL/TR-2012/01

Reactor and Nuclear Systems Division

COMMON-CAUSE FAILURE MITIGATION PRACTICES AND KNOWLEDGE GAPS

Richard Wood
Laura Pullum
Cyrus Smith
David Holcomb
Kofi Korsah
Michael Muhlheim

October 2012

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, Tennessee 37831-6283
managed by
UT-BATTELLE, LLC
for the
U.S. DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

Page intentionally blank

CONTENTS

Page

LIST OF FIGURES

Figure **Page**

Page intentionally blank

LIST OF TABLES

Table **Page**

Page intentionally blank

ACRONYMS

ABB	ASEA Brown Boveri
ABWR	advanced boiling-water reactor
AC	Advant controller
AC	auxiliary cabinet
ACE	actuator control electronics
ADS	automatic depressurization system
AECL	Atomic Energy of Canada, Ltd.
AGR	advanced gas-cooled reactor
AIChE	American Institute of Chemical Engineers
ALU	actuator logic unit
ALWR	advanced light-water reactor
AOO	abnormal operating occurrence
APR	automatic power regulator
APU	acquisition and processing unit
AREMA	American Railway Engineering and Maintenance of Way Association
ARI	alternate rod insertion
ARP	Aerospace Recommended Practice
ASI	Advanced Sensors and Instrumentation
ASIC	application-specific integrated circuit
ATWS	anticipated transients without scram
AVN	Association Vinçotte Nuclear
B-777	Boeing 777
BFS	backup flight system
BNS	Babcock Nuclear Services
BOP	balance-of-plant
BPCS	basic process control system
BTP	Branch Technical Position
C&C	command and control
C&W	caution and warning
CANDU	Canada deuterium-uranium
CCA	common cause analysis
CCF	common-cause failure
CCPS	Center for Chemical Process Safety
CDH	command and data handling
CE	Combustion Engineering
CEDMCS	control element drive mechanism control system
CENELEC	European Committee for Electrotechnical Standardization
CFMS	critical function monitoring system
CFR	Code of Federal Regulations
CIM	communication interface module
CMF	common-mode failure
COTS	commercial-off-the-shelf
CPC	core protection calculator
CPLD	complex programmable logic device
CPU	central processing unit
CSA	Canadian Space Agency
CSF	Compagnie Générale de Télégraphie sans Fil
CSIS	Center for Semicustom Integrated Systems

CSN	Nuclear Safety Council
CUW	reactor water cleanup system
D3	diversity and defense-in-depth
DAL	development assurance level
DAS	diverse actuation system
DBA	design basis accident
DBE	design basis event
DCS	distributed control system
DEC	Digital Equipment Corporation
DIS	digital instrumentation system
DISPO	DIverse Software PrOject
DOE	U.S. Department of Energy
DOE-NE	U.S. Department of Energy Office of Nuclear energy
DOT	U.S. Department of Transportation
DPS	diverse protection system
DRPS	digital reactor protection system
DS&S	Data Systems and Solutions
DSDs	diversity-seeking decisions
DTMs	digital trip modules
ECCS	emergency core cooling system
ECLSS	environmental control and life support system
EdF	Électricité de France
ELAC	elevator and aileron computer
EMS	Energy Management System
EMS	essential multiplexing system
EN	European Norm
EPR	European (or evolutionary) pressurized reactor
EPS	electrical power system
ERA	European Railway Agency
ESA	European Space Agency
ESD	emergency shutdown
ESF	engineered safety feature
ESFAS	engineered safety features actuation system
ESS	emergency shutdown systems
FAA	Federal Aviation Administration
FBW	fly-by-wire
FCPC	Flight Control Primary Computer
FCS	flight control system
FCSC	Flight Control Secondary Computer
FDIR	fault detection, isolation, and recovery
FE	First Energy
FHA	functional hazard assessment
FMCRD	fine motion control rod drive
FMEA	failure mode and effects analysis
FPGA	field-programmable gate array
FRA	Federal Railroad Administration
FTA	fault tree analysis
FWC	feedwater flow control system
FWCS	feedwater control system
GA	General Automation
GDC	general design criteria

GE	General Electric
GEIS	GE Industrial Systems
GNC	guidance, navigation, and control
GPC	general purpose computer
H&B	Hartmann and Braun
HAL/S	High-Order Assembly Language/Shuttle
HBS	hardwired backup system
HFC	Doosan HF Controls
HIACS	Hitachi Integrated Autonomic Control System
HICS	high-integrity control system
HPCF	high-pressure core flooder system
HSE	Health and Safety Executive
HVAC	heating, ventilation, and air conditioning
I&C	instrumentation and control(s)
I/O	input and output
IAEA	International Atomic Energy Agency
IBM	International Business Machine
ICS	integrated control system
IEC	International Electrotechnical Commission
IL	integrity level
ILP	interlocking processor
INL	Idaho National Laboratory
IP	intellectual property
IPLs	independent protection layers
IPS	integrated protection system
IRSN	Institut de Radioprotection et de Sûreté Nucléaire (Institute for Radiological Protection and Nuclear Safety)
ISA	instruction set architecture
ISA	Instrument, System, and Automation Society
ISS	International Space Station
ISTec	Institute for Safety Technology
IV&V	independent verification and validation
JEAG	Japan Electric Association Guideline
JNR	Japanese National Railways
KK	Kashiwazaki-Kariwa (Nuclear Power Station)
KSNP	Korea Standard Nuclear Plant
LBLOCA	large break loss of coolant accident
LCL	local coincidence logic
LOP	lines of protection
LWR	light-water reactor
M-G	motor-generator
MHI	Mitsubishi Heavy Industries
MSIV	main steam isolation valve
NASA	National Aeronautics and Space Administration
NASDA	National Space Development Agency of Japan
NE	Office of Nuclear Energy
NEET	Nuclear Energy Enabling Technologies
NEI	Nuclear Energy Institute
NII	Nuclear Installations Inspectorate
NMS	neutron monitoring system
NPL	nonprogrammable logic

NPP	nuclear power plant
NRC	Nuclear Regulatory Commission
NSSS	nuclear steam supply system
O&M	operation and maintenance
OBC	On-Board Computer
OL-3	Olkiluoto Nuclear Power Station
ORNL	Oak Ridge National Laboratory
OS	operating system
OSHA	Occupational Safety and Health Administration
PAC	priority actuator control
PAS	process automation system
PASS	primary avionics software system
PCS	plant control system
PCS	portable computer system
PDP	Programmed Data Processor
PERFORM.NET	performance-enhanced redundant fiber optic replicated memory network
PESs	programmable electronic systems
PFC	primary flight computer
PFCS	primary flight control system
pfd	probability of failure on demand
PI	process instrumentation
PICS	plant information and control system
PIEs	postulated initiating events
PL/M	program language for microcomputers
PLμS	Programmable Logic Microprocessor System
PLC	programmable logic controller
PLCS	pressurizer level control system
POL	Problem Oriented Language
PPCS	pressurizer pressure control system
PPS	plant protection system
PPS	primary protection system
PRA	probabilistic risk assessment
PRIM	PRIMary flight control computer
PRPS	primary reactor protection system
PS	protection system
PSP	product safety plan
PSSA	preliminary system safety assessment
PWR	pressurized-water reactor
RATP	Régie Autonome des Transports Parisiens (Paris Public Transportation Authority)
RC&IS	rod control and information system
RCIC	reactor core isolation cooling
RCSL	reactor control, surveillance and limitation system
RER	Réseau Express Régional (Paris Rail)
RES	Office of Nuclear Regulatory Research
RFC	recirculation flow control
RHRS	residual heat removal system
RMU	remote multiplexing unit
RPS	reactor protection system
RPT	reactor pump trip
RRS	reactor regulating system

RSA	Russian Space Agency
RSPP	Railroad Safety Program Plan
RTCA	Radio Technical Commission for Aeronautics
RTIF	reactor trip and isolation function
RTSS	reactor trip switchgear system
RTS	reactor trip system
RTU	Remote Terminal Unit
SAAS	severe accidents automation system
SACEM	Système d'Aide à la Conduite, à l'Exploitation et à la Maintenance
SAE	Society of Automotive Engineers
SAIL	Shuttle Avionics Integration Lab
SAR	safety analysis report
SAS	safety automation system
SBP	safety bag processor
SBCS	steam bypass control system
SBPC	steam bypass and pressure control
SC	subcommittee
SC-ABFT	safety critical algorithm-based fault tolerance
SCADA	Supervisory Control and Data Acquisition
SCAP	Système de Contournement à l'AtmosPhère (containment atmospheric control system)
SCAT	Systèmes de Commande des Auxiliaires de Tranche (reactor auxiliary systems control)
SDS1	Shutdown System Number 1
SDS2	Shutdown System Number 2
SEC	spoiler and elevator computer
SICS	safety information and control System
SIL	safety integrity level
SIS	high-integrity safety instrumented system
SKI	Statens Kärnkraftinspektion (Swedish Nuclear Power Inspectorate)
SLCS	standby liquid control system
SME	subject matter expert
SNCF	Société Nationale des Chemins de fer Français (French National Railway Company)
SPIN	Système de protection intégré numérique (Integrated Digital Protection System)
SPPA	Siemens Power Plant Automation
SPS	secondary protection system
SRI	Inertial Reference System
SRM	staff requirements memorandum
SSA	system safety assessment
SSD	safety shutdown systems
SSDE	software development environment
SSLC	system safety logic control
STS	Space Transportation System (Space Shuttle)
STUK	Säteilyturvakeskus (Radiation and Nuclear Safety Authority)
SWAP	size, weight, and power
TC	technical committee
TCS	thermal control system
TLUs	trip logic units
TMR	triple modular redundant
TOSMAP	Toshiba Microprocessor Aided Power System Control

TXS	AREVA Teleperm XS
U.K.	United Kingdom
UA	acquisition units
UATP	acquisition and processing unit for protection
UF	functional units
ULS	logic safeguard unit
UPS	uninterruptible power supply
UTPs	logic processors
UVa	University of Virginia
V&V	verification and validation
V_Frame	Vital Framework
VCP	vital coded processor
VME	VERSAbus-E
VVER	Russian-designed water-cooled water-moderated power reactor
WDPF	Westinghouse Distributed Processing Family
WENRA	Western European Nuclear Regulators' Association
WIPP	Waste Isolation Pilot Plant

ACKNOWLEDGMENTS

The research described in this report was sponsored by the Nuclear Energy Enabling Technologies (NEET) Advanced Sensors and Instrumentation (ASI) Program of the U.S. Department of Energy (DOE) Office of Nuclear Energy. The findings documented in this report leverage information from prior investigations while expanding and enhancing the determination of the state of the practice for common-cause failure (CCF) mitigation. In particular, the authors incorporated into this report technical information generated through their own prior work under U.S. Nuclear Regulatory Commission (NRC) Office of Nuclear Regulatory Research sponsorship.

The authors wish to thank the technical peer reviewers and editorial staff at Oak Ridge National Laboratory for their feedback and assistance in improving this report.

Page intentionally blank

EXECUTIVE SUMMARY

Experience in other industries has shown that digital technology can provide substantial benefits in terms of performance and reliability. However, the U.S. nuclear power industry has been slow to adopt the technology extensively in its instrumentation and control (I&C) applications because of inhibiting factors such as regulatory uncertainty, insufficient technological experience base, implementation complexity, limited availability of nuclear-qualified products and vendors, and inadequate definition of modernization cost recapture. Although there have been examples of digital technology usage in the nuclear power industry, challenges to the qualification of digital technology for high-integrity nuclear power plant (NPP) applications have severely constrained more widespread progress in achieving the benefits that are possible through the transition to digital.

Design criteria for safety-related I&C systems embody principles such as high quality, integrity, reliability, independence, and qualification to ensure that safe conditions are maintained under all operational conditions. Separation and redundancy, physical barriers, and electrical isolation are commonly applied as design measures within a defense-in-depth concept to address potential vulnerabilities related to single failures of equipment and the propagation of failure effects. However, errors, deficiencies, or defects at any stage of a system's life cycle can result in systematic faults that may remain undetected until operational conditions activate the faulted state to result in a failure of a critical function. The potential for common-cause failure (CCF) of multiple systems (or redundancies within a system) constitutes the principal credible threat to defeating the defense-in-depth provisions within I&C system architectures of NPPs. The unique characteristics and inherent complexity of digital I&C systems can exacerbate this vulnerability.

Diversity and defensive design measures are the primary means employed to address CCF vulnerability. However, the value and effectiveness of various strategic approaches for design, implementation, and architecture are not well understood. The lack of technical certainty results in the imposition of complex (and costly) expedient solutions that inhibit the use of digital technology and complicate its regulatory acceptance. Consequently, diversity and defense in depth (D3) has been identified as a high-priority technical issue for the nuclear power industry.

Experience with applying current guidance and practices on CCF mitigation to digital I&C systems has proven problematic, and the regulatory environment has been unpredictable. The impact of CCF vulnerability is to inhibit I&C modernization and, thereby, challenge the long-term sustainability of existing plants. For new plants and advanced reactor concepts, the issue of CCF vulnerability for highly integrated digital I&C systems imposes a design burden resulting in higher costs and increased complexity. The regulatory uncertainty regarding which mitigation strategies are acceptable (e.g., what diversity is needed and how much is sufficient) drives designers to adopt complicated, costly solutions devised for existing plants.

The conditions that constrain the transition to digital I&C technology by the U.S. nuclear industry require crosscutting research to resolve uncertainty, demonstrate necessary characteristics, and establish an objective basis for qualification of digital technology for usage in NPP I&C applications. To fulfill this research need, Oak Ridge National Laboratory is conducting an investigation into mitigation of digital CCF vulnerabilities for nuclear-qualified applications. The outcome of this research is expected to contribute to a fundamentally sound, comprehensive basis for establishing the qualification of digital technology for nuclear power applications.

The starting point for this research involved investigating available documentation on diversity approaches and experience from the international nuclear power industry as well as other industries and organizations, capturing expert knowledge and lessons learned, determining best practices, and evaluating the knowledge gaps that remain. Information on nonnuclear industries and organizations was reviewed to

determine their approaches to avoiding or mitigating the effects of CCF in high-integrity and/or safety-significant systems. This investigation focused on industries that employ similar I&C technologies and have high-consequence applications. For nuclear power, the application of digital technology for I&C systems at international evolutionary NPPs provides a significant resource in determining effective strategies for addressing CCF vulnerability. A review of available standards and guides served as the basis for identifying existing guidance for the treatment of CCF vulnerability. Additionally, recent research efforts into key issues related to CCF mitigation were assessed.

Experience with the impact of CCF in a variety of applications was identified as part of this investigation. In various nonnuclear industries, there have been instances of failures resulting from latent faults triggered by operational conditions. These events serve to illustrate the nature and impact of CCF vulnerabilities. The cited examples are the destruction of Ariane 5 missile Flight 501, intercept failures by Patriot missile batteries, massive radiation overdoses by Therac-25 equipment, the crash of Air France Flight 447, the 2003 electrical blackout of North America, an outage of the AT&T communications network, and a design flaw of the Intel Pentium chip. These events involve instances of flawed software design in which either inadequate requirements were defined or the application was inconsistent with the design basis, an example of the impact of common external conditions, cases illustrating the cascading effect of software failures, and an instance of a platform-specific fault.

Domestic and international nuclear power regulations and regulatory guidance address both defense-in-depth and diversity as means of mitigating single and common-cause failures. U.S. Nuclear Regulatory Commission (NRC) regulations require licensees to incorporate an overall safety strategy for defense-in-depth functions and systems to ensure that abnormal operating occurrences and design basis accidents do not adversely impact public health and safety. The general design criteria, provided in Appendix A of Title 10, Part 50 of the Code of Federal Regulations (10 CFR 50), establish the minimum design requirements for light-water reactors. Aside from the rule requiring mitigation against an anticipated transient without scram (ATWS), most regulatory guidance is found in staff requirements memoranda [SRM on SECY-93-087], branch technical positions [NUREG-0800, Chapter 7, BTP 7-19], and NRC contractor reports [NUREG/CR-6303, NUREG/CR-7007]. Diversity is the preferred mitigation approach for addressing perceived CCF vulnerabilities of I&C system architectures because dissimilarities in technology, function, implementation, and so forth can diminish the potential for common faults. However, the guidance is complex and a subjective judgment is required to determine what diversity usage is adequate to mitigate identified CCF vulnerabilities.

The guidance provided by key international standards for the nuclear industry addresses the basis for a general approach to coping with CCF in I&C systems important to safety. International Electrotechnical Commission (IEC) standard IEC 61513 represents the high-level guidance addressing I&C system architecture considerations. IEC 60880 supplements that guidance by specifically addressing software-based system considerations. IEC 62340 provides a framework for establishing a CCF coping strategy that is consistent with the high-level requirements in IEC 61513 and complementary to the software requirements in IEC 60880. Some nonnuclear industries (aerospace, aviation, chemical process, and rail transportation) were also found to provide standards and guidelines with specific guidance related to CCF mitigation. The fundamental difficulty with the nuclear and nonnuclear standards and guidelines, as with the regulatory guidance, results from the lack of any definitive specification of necessary and sufficient mitigation practices. Consequently, the user (or licensee) is faced with subjective criteria and uncertainties about effectiveness in the determination of CCF mitigation strategies.

Specific examples of CCF mitigation practices can be found at international NPPs. The examples that are described in this report represent a sampling of evolutionary reactors and modernized plants that employ digital technology extensively. In particular, some of the earliest examples of highly integrated digital I&C systems were included in the survey. Specifically, the investigation covered six plants that were commissioned with installed digital I&C systems: Chooz, Darlington, Kashiwazaki-Kariwa, Sizewell, Temelín, and Ulchin. In addition, an example of extensive modernization from analog to digital

technology in an existing plant (Dukovany) was reviewed. Finally, two plants presently undergoing licensing and construction were studied to assess recent trends. These plants are Lungmen and Olkiluoto. The usage of diversity seen in these examples represents clear indication of best practices in the nuclear industry. However, there is also seen a great variety in the extent and type of diversity applied as well as the range of coverage for safety functions.

Several nonnuclear industries were investigated through this research. Many were found to rely primarily on high-quality processes and rigorous hazard identification and resolution. However, a few key safety-critical industries were found to provide clear examples of CCF mitigation approaches. The aerospace industry tends to rely on high-quality processes to minimize the potential for CCF vulnerabilities. The use of failsafe design practices with reduced functionality backups characterizes the prime examples of safety-critical applications for manned space systems (i.e., Space Shuttle, International Space Station). The aviation industry provided several examples of diversity usage, with two prominent approaches. The Airbus approach emphasized diversity of development teams and software, while the Boeing approach emphasized diversity of hardware and implementation tools. The chemical process industry provides guidance that is similar in nature to the nuclear power industry. However, no definitive metrics or specific diversity usage template is provided. The rail transport industry also provided several examples of diversity usage. Early implementations of digital train control systems relied primarily on software diversity. A more hardware-oriented approach based on encoded processors for parallel checking architectures was also seen in key examples.

Recent research into CCF mitigation was assessed as part of this investigation. In particular, NRC research resulted in the development of baseline mitigation strategies that were consistent with acceptable practices based on implementation experience. The key assumption in that research is that qualitative assessment of the impact of diversity attributes and criteria, coupled with insights derived from established practice and key usage examples, provides a valid basis for developing diversity strategies to cope with the potential for CCF. Diversity usage tables and a diversity assessment spreadsheet tool were developed to aid in the evaluation of proposed mitigation strategies. The diversity assessment tool can also be employed for comparative analyses to assess the relative standing of a proposed alternate diversity strategy against the baseline strategies as well as established practices and common usage of the nuclear power and nonnuclear industries. This tool provides a systematic approach to evaluate proposed combinations of diversity criteria. However, the tool is based on subjective weighting of diversity effectiveness derived from engineering judgment and frequency of usage in the limited sample set. Thus, the scoring of strategies should be seen as a qualitative comparison, not an objective measure of CCF mitigation effectiveness.

The findings from diversity research conducted by the nuclear power industry in the United Kingdom confirm that it cannot be conclusively demonstrated with mathematical rigor that intentional or forced diversity will result in independence of failure between systems. Additionally, the effect of diversity usage (individually or collectively applied) cannot be quantitatively determined at present. Basically, it was found that an extensive range of possible diversity-seeking decisions and their combinations are available but there is little definitive guidance for these choices. However, it is clear from qualitative evidence that diversity provides a dependability benefit (i.e., contributes to the mitigation of CCF vulnerabilities through overall system-level fault tolerance) and is a reasonable response to CCF concerns. What are needed are objective measures of digital I&C system characteristics that give indication of the efficacy of various mitigation techniques.

An assessment of the findings from the investigation of the state of the practice for CCF mitigation points to knowledge gaps that should be resolved through further research. The foremost deficiency in knowledge relates to a fundamental understanding of the nature of CCF vulnerability in the context of the nuclear power application domain. In particular, a comprehensive identification is needed of the sources of systematic faults and the triggering conditions that impact safety-related functions in an NPP. These fault-trigger combinations should be mapped to functions and architectural elements (e.g., I&C system

blocks) and related to hazards that could compromise plant safety. In addition, the various diversities and design measures that can mitigate CCF need to be related to the particular kinds of faults, triggers or fault-trigger combinations and to the corresponding failures that can result. The consequence would be better understanding of the impact of each diversity, the value of other defensive design measures, and the synergistic effect of combined mitigation techniques.

The basic knowledge gap can be characterized as a need to establish the effectiveness of various mitigation techniques (e.g., diversity-seeking decisions or DSDs) in addressing specific classes of faults, triggers, or fault-trigger combinations. Essentially, a quantitative characterization of how DSDs diversify failure behavior for parallel systems would enable development of objective decision criteria and provide for a more comprehensive, systematic, and scientifically based determination of what mitigation strategy would be most effective. The questions that need answers include the following: How effective is a particular DSD in resolving a particular CCF vulnerability? Which diversity or design measure is best for certain classes of CCF? How much diversity is adequate? What is the combined effect of multiple DSDs?

To resolve this knowledge gap, more thorough definition of each diversity attribute and defensive design measures should be developed. Various application domains have different characterizations of diversity. In addition, models and metrics are needed to develop systematic methods, quantifiable measures, and objective criteria for evaluating CCF mitigation approaches. Various measures of I&C system characteristics (e.g., quality, reliability, performance, dependability) may be relevant for determining the effectiveness of diversity or design measures in mitigating CCF vulnerabilities. A thorough investigation of potential measures and models to support an aggregate indicator of diverse failure behavior is needed.

It is clear from the investigation of the state of the practice for CCF mitigation that a fundamentally sound basis for acceptable mitigation approaches is needed. Resolving uncertainties and regulatory burden concerning CCF vulnerability can promote elegant, optimal architectures for NPP I&C architectures with a well-defined safety basis, less imposed complexity, and, potentially, reduced cost. Achieving a science-based solution to this key technical challenge can benefit existing plants, new plants, and advanced designs by removing an impediment to more extensive, effective use of digital technology.

1. INTRODUCTION

The U.S. Department of Energy (DOE) Office of Nuclear Energy (NE) established the Advanced Sensors and Instrumentation (ASI) technology area under the Nuclear Energy Enabling Technologies (NEET) Program to coordinate the instrumentation and controls (I&C) research across DOE-NE and to identify and lead efforts to address common needs. As part of the NEET ASI research program, the Digital Technology Qualification project was established based on collaboration between Oak Ridge National Laboratory (ORNL) and Idaho National Laboratory (INL). ORNL is leading the investigation into mitigation of digital common-cause failure (CCF) vulnerabilities for nuclear qualified applications, and INL is conducting an investigation into the suitability of digital alternatives to analog sensors, control loops, and actuators. ORNL is responsible for integrating the technical findings and research products of this collaborative effort.

This technical report documents the findings from first phase of research activities by ORNL. Specifically, the report describes the results of the investigation of CCF mitigation practices and determination of knowledge gaps.

1.1 Background

1.1.1 Technical Issue

Experience in other industries has shown that digital technology can provide substantial benefits in terms of performance and reliability. However, the U.S. nuclear power industry has been slow to adopt the technology extensively in its I&C applications because of inhibiting factors such as regulatory uncertainty, insufficient technological experience base, implementation complexity, limited availability of nuclear-qualified products and vendors, and inadequate definition of modernization cost recapture. Obsolescence of replacement analog components and development of *de facto* standard approaches based on subjective criteria have enabled modest movement toward increasing the use of digital electronics for some command functions (e.g., control or protection algorithms/logic). However, key issues, such as software quality and mitigation of CCF vulnerabilities, have led to the imposition of complex, costly design conventions and implementation practices that challenge the qualification of digital technology for high-integrity nuclear power plant (NPP) applications and constrain the benefits that can be achieved through the transition to digital.

Design criteria for safety-related I&C systems embody principles such as high quality, integrity, reliability, independence, and qualification to ensure that safe conditions are maintained under all operational conditions. Separation and redundancy, physical barriers, and electrical isolation are commonly applied as design measures within a defense-in-depth concept to address potential vulnerabilities related to single failures of equipment and the propagation of failure effects. However, errors, deficiencies, or defects at any stage of a system's life cycle can result in systematic faults that may remain undetected until operational conditions activate the faulted state to result in a failure of a critical function. The potential for CCF of multiple systems (or redundancies within a system) constitutes the principal credible threat to defeating the defense-in-depth provisions within I&C system architectures of NPPs. The unique characteristics and inherent complexity of digital I&C systems can exacerbate this vulnerability.

Diversity and defensive design measures are the primary means employed to address CCF vulnerability. However, the benefits of various strategic approaches for design, implementation, and architecture are not well understood. The lack of technical certainty results in the imposition of complex (and costly) expedient solutions that inhibit the use of digital technology and complicate its regulatory acceptance. Consequently, diversity and defense in depth (D3) has been identified as a high-priority technical issue for the nuclear power industry by both the Digital I&C Steering Committee of the U.S.

Nuclear Regulatory Commission (NRC) and the Industry Digital I&C and Human Factors Working Group of the Nuclear Energy Institute (NEI) [1].

Experience with applying current guidance and practices on CCF mitigation to digital I&C systems has proven problematic, and the regulatory environment has been unpredictable. In a recent license amendment in the United States regarding a digital modernization of plant safety systems, regulatory concerns about CCF vulnerabilities, which were indicated through a D3 analysis, proved difficult to resolve. As a consequence of this issue, the licensing process was considerably delayed, with a substantial time and cost impact as a consequence; it was ultimately determined that implementation of additional diverse systems was also required to mitigate the potential CCF vulnerability, which resulted in greater complexity (and cost) as well. The impact of CCF vulnerability is to inhibit I&C modernization and, thereby, challenge the long-term sustainability of existing plants. For new plants and advanced reactor concepts, the issue of CCF vulnerability for highly integrated digital I&C systems imposes a design burden resulting in higher costs and increased complexity. International regulators in Finland, United Kingdom, and France have expressed concern about the treatment of CCF vulnerability in highly digital I&C architecture for advanced light-water reactor designs. Specifically, the regulatory review of the safety systems for the new third unit under construction at the Olkiluoto NPP in Finland has been complicated because of concerns about the potential susceptibility of the I&C architecture to CCF. The regulatory uncertainty regarding which mitigation strategies are acceptable (e.g., what diversity is needed and how much is sufficient) drives designers to adopt complicated, costly solutions devised for existing plants. The result may be unnecessarily complex I&C architectures that are clearly not optimal solutions and may be inappropriate for advanced reactor designs. Thus, mitigation of CCF vulnerability is an issue of concern for existing plants, new plants, and advanced reactor concepts.

1.1.2 Approaches to Mitigate CCF Vulnerability

There are many techniques for managing digital I&C system faults that have been employed for high-integrity functions within various application domains. They are generally grouped in terms of design evaluation and fault removal, fault tolerance (i.e., detection/masking and recovery), and fault avoidance and mitigation. The techniques indicated involve design approaches, life-cycle actions, technology choices, architectural configurations, and so forth.

Design evaluation and fault removal apply to detailed analyses to identify and eliminate threats to the extent practical, as well as to high-quality processes employed to minimize the potential for faults and remove vulnerabilities as they are discovered. These techniques generally promote fault avoidance at a high level and are primarily oriented toward design approaches and evaluation processes.

Fault tolerance represents specific techniques for accommodating the presence of faults and avoiding consequent failure. Failsafe designs are enabled by these techniques. Detection and masking relate to identifying the presence of a fault or masking its potential effect (i.e., avoiding failure due to the fault). Diagnostics (e.g., fault identification and isolation) and voted redundancies are common techniques. Recovery relates to the response to an activated fault (i.e., failure) and enables continued execution with recapture of the prefailure state.

Fault avoidance and mitigation include design strategies to impede the propagation of the effects of faults (i.e., failures). Separation, independence, and fault containment are techniques for constraining the potential effects of activated faults, while dissimilarity/diversity and checked redundancy are means for mitigating the effect of activated faults by either precluding common faults (in the first case) or detecting and compensating for activated faults (in the second case).

The fault management techniques described above generally relate to the faults themselves and, to some degree, to the triggering conditions that activate the faults to cause failures. These fault management

techniques embody supporting technical and life-cycle methods and approaches on which strategies to cope with CCF vulnerability can be based.

At the outset of I&C system architecture development, design principles are invoked to minimize the use of common elements and to limit failure propagation paths. These design considerations are effective in reducing the potential for CCF susceptibility, but their absolute, across-the-board use can result in extremely complicated, inefficient, and potentially unreliable I&C system architectures. As a result, two principal coping strategies are typically employed in responding to CCF susceptibility: (1) CCF avoidance and (2) CCF mitigation.

The objective of the first strategy is to avoid fault introduction and eliminate potential common triggering conditions to the degree feasible. Comprehensive life-cycle processes with thorough hazard identification and extensive verification and validation activities are employed to yield high-quality systems with the goal of approaching error-free software. Nevertheless, experience confirms that undetected errors can progress through even the most rigorous design process. As an additional aspect of the avoidance approach, design measures can be used to reduce the exposure to anticipated triggering conditions or their concurrent application to multiple systems that may have common faults. Application of such design measures depends upon a well-founded understanding of the types of fault-trigger combinations that may be present and the design conventions that are most effective in preventing concurrent triggering of any common faults that may be present. Examples of these design measures are invariant execution of code and physical separation by barriers into different environmental control zones. However, since there is no assurance that unanticipated common triggering conditions do not exist, use of these measures cannot guarantee sufficient CCF robustness. Thus, the primary goal of this strategy is to minimize the occurrence of common faults and reduce the likelihood of failures.

The objective of the second strategy is to mitigate any vulnerability to CCF through architectural provisions. First, defense-in-depth is employed to compensate for failures in other systems or functions. The International Atomic Energy Agency (IAEA) defines defense-in-depth as “the application of more than one protective measure for a given safety objective, such that the objective is achieved even if one of the protective measures fails” [2]. In practice, several independent systems are implemented to serve as successive barriers to prevent unsafe consequences from occurring. This aspect of the mitigation approach is especially effective against single failures. However, CCF can potentially disable multiple barriers and result in unsafe conditions. Thus, diversity is employed to provide alternate equivalent functionality or systems that are not susceptible to the same CCF as their counterpart(s) within the I&C system architecture. The difficulty occurs in identifying the full range of CCF vulnerabilities that may be present and then selecting the appropriate compensating diversities. Thus, the primary realistic goal of this strategy is to mitigate the impact of a potential CCF by providing alternate or backup functions that are unaffected.

These fault management strategies provide many techniques to reduce the likelihood of faults and failures and to mitigate those vulnerabilities that may exist. The nuclear power industry applies rigorous quality process control to avoid faults, errors, and deficiencies. However, the potential for latent faults persists. Thus, diversity and defensive design measures are employed to mitigate residual CCF vulnerabilities. The trouble, as indicated above, is that great uncertainty remains as to the efficacy of the mitigation strategies employed and the value each provides. Succinctly, the issue that remains is the need to answer the question, “If diversity is required in a safety system to mitigate the consequences of potential CCFs, how much diversity is enough?” Thus, further research is needed to develop comprehensive mitigation strategies to effectively address CCF vulnerabilities without introducing unnecessary complexity and significant cost while also providing a sound scientific basis for establishing an acceptable safety justification.

1.2 Research Approach

Because of the complexity of digital I&C system technology and the necessary reliance on process-driven approaches to software development and quality assurance, there has been an absence of definitive quantitative measures for key digital I&C system characteristics. As a result, it has not been feasible to develop a comprehensive measure of diversity (particularly for software-based systems) that could be used to establish wholly objective acceptance criteria to support diversity reviews.

The research approach employed for this initial effort involved investigating available documentation on diversity approaches and experience from the international nuclear power industry as well as other industries and organizations, capturing expert knowledge and lessons learned, determining best practices, and evaluating the knowledge gaps that remain. Nonnuclear industries and organizations were investigated to determine their approaches to and experience with avoiding or mitigating the effects of CCF in high-integrity and/or safety-significant systems. This investigation focused on industries that employ similar I&C technologies and have high-consequence applications. For nuclear power, the extensive application of digital technology for I&C systems at international evolutionary NPPs provides a significant resource in determining effective strategies for addressing CCF vulnerability. Where available, standards and guides were identified and reviewed. Additionally, recent research efforts into key issues related to CCF mitigation were assessed.

1.3 Report Organization

The report is divided into five major sections: CCF vulnerabilities, guidance on CCF mitigation, examples of CCF mitigation practices, recent research into CCF mitigation strategies, and knowledge gaps. Background information on the nature of CCF and experience with CCF is provided in Chapter 2. Chapter 3 describes regulatory requirements, nuclear power industry standards, and nonnuclear industry guidance related to CCF mitigation. Chapter 4 presents the findings from the survey of approaches to address CCF at international NPPs and within safety-critical nonnuclear industries. Chapter 5 compiles information on other relevant research activities and results. Chapter 6 presents the assessment of knowledge gaps that help to inform research direction.

2. COMMON-CAUSE FAILURE VULNERABILITIES

2.1 Common-Cause Failure of I&C Systems

CCF is defined by the IAEA as a "failure of two or more structures, systems or components due to a single specific event or cause" [3]. The International Electrotechnical Commission (IEC) further adds to the CCF definition by noting that the "coincidental failure of two or more structures, systems or components is caused by any latent deficiency from design or manufacturing, from operation or maintenance errors, and which is triggered by any event induced by natural phenomenon, plant process operation, or action caused by man or by any internal event in the I&C system" [4]. CCF is a class of dependent failures in which the probability of failure is not expressible as the simple product of the unconditional failure probabilities of the individual events. Common-mode failure (CMF) is a subset of CCF and occurs when two or more systems or components fail in the same way.

The basis for a CCF occurrence is described in IEC 62340, "Nuclear power plants—Instrumentation and control systems important to safety—Requirements to cope with common cause failure (CCF)" [4], as corresponding to the systematic incorporation of a latent fault in multiple systems or redundancies followed by the triggering of that common fault to cause a coincidental failure of some or all of the systems or redundancies.

Latent faults can originate at any phase of the digital I&C system life cycle; are typically human induced or technology related; and involve design flaws, performance limitations, or implementation complexity. At a high level, three prominent sources of latent systematic faults are (1) errors in the requirement specification, (2) inadequate provisions to account for design limits (e.g., environmental stress), and (3) technical faults incorporated in the internal system (or architectural) design or implementation.

Quality processes detect and correct many implementation errors. However, as design complexity increases, the feasibility of exhaustive testing or comprehensive formal proof diminishes considerably. Therefore, some residual faults may remain undetected and persist as latent faults within the system. Design errors resulting from flawed, incomplete, ambiguous, or misinterpreted requirements are systematic in nature and are significantly more difficult to detect and correct as the system life-cycle phases progress. These faults and errors are, in and of themselves, not a hazard unless conditions (e.g., operational, environmental, relational, or temporal) activate the faulted state and result in a failure of a critical function.

Triggering conditions that can activate faults and result in failure arise primarily from signal trajectory, human actions, external events, and temporal effects. The signal trajectory for a digital I&C system involves not only current input values but also past input values, the internal state of the system, and the sequence of transitions among internal states. The IEC defines signal trajectory as the "time histories of all equipment conditions, internal states, input signals and operator inputs which determine the outputs of a system" [5]. Failures arising from latent faults activated by signal trajectory triggering conditions clearly correspond to conditions that either were not anticipated or properly addressed during system development and that were not exposed through testing.

Human actions that can induce a CCF include maintenance errors, input mistakes, out-of-sequence commands, and ill-timed or conflicting actions. External events that can pose common cause triggers include transient effects, such as anomalies or failures propagating from other systems or components within the I&C system architecture, and environmental stress, such as seismic, vibratory, electromagnetic and electrical surge, and so forth. Temporal effects that can initiate failures include dependence on calendar-date or time-of-day information, synchronization with a common clock, synchronization of processes or systems, and runtime effects dependent on execution cycle histories (e.g., runtime overflows of buffers or stacks).

The bottom line is that CCF is a credible concern for high-integrity or safety-critical I&C applications that employ complex technologies within complicated system architectures. Both traditional analog-based and more modern digital-based I&C systems are subject to latent systematic faults resulting from design errors or requirements deficiencies. However, because of the complexity of digital I&C systems and the associated inability to execute exhaustive testing, there is increased concern that the potential for latent systematic faults is greater in more fully digital I&C system architectures. In particular, since software (other than the simplest programs) in its coded state or its compiled machine language state cannot be proven to be without error, residual software faults represent a primary CCF concern. As a result, digital I&C systems receive particular emphasis in assessments of CCF susceptibility and the resulting application of techniques for avoiding or mitigating the potential for CCF vulnerabilities.

2.2 Common-Cause Failure Experience

Despite the best efforts of designers, developers, implementers, reviewers, testers, suppliers, and assessors, errors happen. As discussed above, the types of failures that can compromise safety-critical functions typically arise from design mistakes or implementation errors. Failures can also result from undetected internal flaws (i.e., platform faults), system interactions, and external effects. Hazard identification and design measures can minimize the potential for some sources of failure, but unanticipated and untested conditions can still pose a risk.

In various nonnuclear industries, there have been instances of failures resulting from latent faults triggered by operational conditions. These events constitute examples of CCF vulnerabilities that have resulted in significant failures with serious consequences. A selection of these events is reported below to illustrate the nature and impact of CCF vulnerabilities. These events involve instances of flawed software design in which either inadequate requirements were defined or the application was inconsistent with the design basis, an example of the impact of common external conditions, cases illustrating the cascading effect of software failures, and an instance of a platform-specific fault.

2.2.1 Destruction of Ariane 5 Missile Flight 501

The maiden flight of the Ariane 5 launcher, Flight 501, on June 4, 1996, ended in the destruction of the launch vehicle along with four European Space Agency (ESA) satellites, known as Cluster, at a loss of more than $370 million. Only 37 seconds after launch, at an altitude of about 12,000 feet, the launcher veered off its flight path, broke up, and exploded.

Ariane is a series of European civilian expendable launch vehicles developed by the ESA, which is primarily sponsored by France, Germany, and the United Kingdom. The Ariane rockets are launched by a commercial subsidiary of ESA, Ariancespace, from the Centre Spatial Guyanais at Kourou in French Guiana, where the proximity to the equator gives a significant advantage for the launch.

Ariane 1 operated from 1979 to 1986 with 9 of 11 successful launches; Ariane 2 operated from 1986 to 1989 with 5 of 6 successful launches, Ariane 3 operated from 1984 to 1989 with 10 of 11 successful launches, and Ariane 4 operated from 1990 to 2003 with 113 of 116 successful launches. Ariane 1 was a three-stage launcher, derived from military missile weapon technology, and Ariane 2 through 4 were enhancements of the basic vehicle.

Ariane 5 is a nearly complete redesign intended from the beginning to be rated to launch humans as it was designed to launch the manned mini shuttle, Hermes. The two lower stages of the Ariane 1 through 4 are replaced with a single, cryogenic stage using a Vulcain engine, but liftoff requires the additional use of two solid-fuel boosters that are strapped to the sides. The upper stage is restartable and uses a single Aestus engine.

The Ariane 5 Flight Control System is a standard design where the attitude of the launcher and its movements are measured by an Inertial Reference System (SRI). The SRI has a dedicated computer where angle and velocities are calculated from information from a "strap-down" inertial system with laser gyros and accelerometers. Flight information from the SRI is set to the On-Board Computer (OBC), which executes the flight control software and controls the nozzles of the solid boosters and the Vulcain engine. There is redundancy in the SRI system in the form of two SRIs that operate in parallel, but both run identical hardware and software. One SRI is active and the other is in "hot" standby. If the OBC determines that the active SRI has failed, it immediately switches to the other, provided that the standby SRI is operating properly. For further redundancy, there are two OBCs and a number of other units in the Flight Control System (FCS) are also duplicated. It is also important to note that the Ariane 5 SRI is essentially the same as that of the Ariane 4 SRI, especially in regards to the software.

The launcher started to disintegrate about 37 seconds after liftoff because of high aerodynamics loads caused by an angle of attack of more than 20 degrees. These loads led to a separation of the solid fuel boosters from the main stage that triggered the self-destruction system of the launcher to ensure safety. This high angle of attack was caused by full nozzle deflections of both the solid boosters and the main engine. These nozzle deflections were commanded by the OBC based upon information from the active SRI (SRI #2). At that time, the data coming from SRI #2 was not proper flight data, but was instead a diagnostics bit pattern that was being interpreted as flight data. The diagnostic bit pattern was present because SRI #2 had failed due to a software exception. Furthermore, the OBC could not switch to the "hot" standby SRI (SRI #1), because SRI #1 had already experienced the same software exception about 72 milliseconds earlier and declared itself inoperative.

The SRI software exception was caused during execution of a data conversion from 64-bit floating point to 16-bit signed integer value. The floating-point number to be converted was larger than that which could be represented as a 16-bit signed integer. This resulted in an Arithmetic Overflow. This particular data conversion coding was not protected from causing an Arithmetic Overflow, although other conversions of similar variables in the same place in the code were protected. This software error occurred in the software coding that performs alignment of the strap-down inertial system. However, this software only provides a useful calculation prior to liftoff; as soon as the launcher lifts off, this software serves no purpose. Regardless, the alignment function of this software is operative for about 40 seconds after liftoff. This timing is based upon a requirement for the Ariane 4 launcher to be able to "hold" for extensive time periods immediately prior to launch, but is not required at all for the Ariane 5 launcher. The Arithmetic Overflow occurred because of an unexpectedly high value of the horizontal velocity of the launcher. This is actually normal for the Ariane 5 launcher, which has a very different trajectory during the early part of the flight from that of the Ariane 4 launcher.

These conditions arose because the SRI software from the Ariane 4 launcher was reused for the SRI software in the Ariane 5 launcher. However, the flight trajectories for the launchers were different, with the Ariane 5 achieving much greater horizontal acceleration than the Ariane 4. However, the requirement

documents for the Ariane SRI did not address the trajectory data as a functional requirement. The difference in the trajectory data early in the flight path between the Ariane 4 and Ariane 5 launchers caused the crash. The impact on software requirements was actually analyzed, and the coding for four variables of the SRI software was modified to protect an Arithmetic Overflow from occurring with them. However, three variables, including the one that caused the crash, were left unprotected. While the software modifications were examined, reviewed, and approved by project partners at several contractual levels, these three variables were still left unprotected. The rationale for this oversight has been attributed to the culture within the Ariane program of addressing any possible failure using the same methodology as that used for addressing random hardware failures. Thus, failure mitigation was provided through the utilization of backup systems, which were subject to the same software common mode software failure.

2.2.2 Patriot Missile Battery Intercept Failure

Incorrect tracking of an Iraqi Scud missile on the night of February 25, 1991, during Operation Desert Storm of the Gulf War by a Patriot Missile Battery resulted in the death of 28 soldiers and injury of another 98 in Dhahran, Saudi Arabia. The incorrect tracking of the missile was determined to be the result of "software aging." Software aging refers to the progressive performance degradation of a software system due to exhaustion of the operating system resources or accumulation of errors. In this case, an accumulation of numerical calculation errors resulted in the incorrect tracking of the missile. The aging condition arose because the Patriot Missile Battery was being used in a manner that was inconsistent with its design basis.

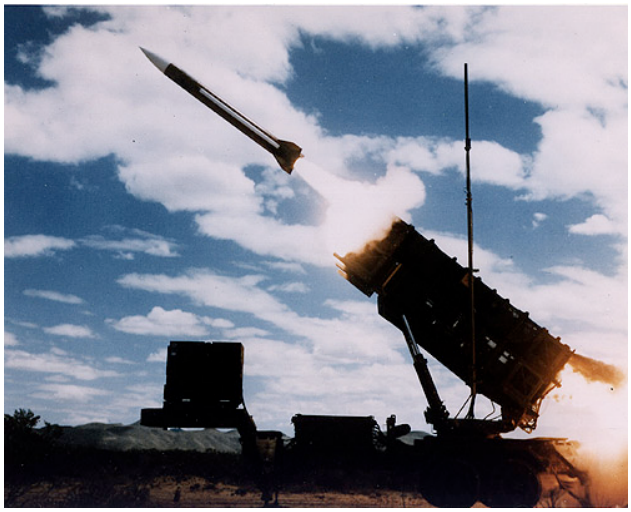

At the time of this failure, the Patriot system and its associated computer software was approximately 20 years old. The Patriot system was originally designed to track and shoot down relatively slow flying (under MACH 2) Soviet medium- to high-altitude aircraft and cruise missiles during the Cold War era. However, in the Gulf War, it was pressed into service to defend against more modern and faster flying (MACH 5) ballistic missiles in use by the Iraqi forces. The Patriot system was also designed as a mobile platform in order to avoid detection and was only specified to operate for a few hours at a time. Contrastingly, in the Gulf War during Operation Desert Shield, Patriot systems were permanently deployed in strategic locations in Saudi Arabia and Israel and allowed to operate for considerable periods of time. At the time of the failure, this particular Patriot system had been operating for approximately 100 hours.

The Patriot system operates by initially performing a "wide area search" for any targets. This Patriot system acquired the incoming Scud ballistic missile during its "wide area search" and calculated a "track" which was an approximation of the path that the Scud missile was expected to follow to impact. Once a potential target is acquired, the Patriot system quickly isolates the target and begins to track it to ensure that it is indeed a threat and not a "false alarm." Part of this tracking and verification of a potential target requires that the Patriot system verify that the radar return shows the object being tracked as being "on course" with the original "track" established during the "wide area search." The Patriot system does this by ensuring that the radar return places the tracked object within a calculated "range gate area" further along the "track." The software aging situation in this Patriot system resulted in the range gate area being about 700 meters away from where the ballistic Scud missile was truly positioned and therefore "off"

from the predicted "track." The Patriot system therefore classified the radar return (i.e., Scud missile) as a "false alarm" and failed to engage the incoming missile.

The software aging bug occurred in the calculation of the next location of the Scud missile where the range gate would be applied. This prediction is calculated based on the missile's velocity and the time of the last radar detection. In the Patriot system, the target velocity is stored as a whole number and a decimal, and time is a continuous integer or whole number measured in tenths of a second (i.e., the longer the system has been running, the larger the number which represents time). The algorithm that predicts the next expected location of the potential target requires both time and velocity be expressed as real numbers. However, the Patriot's computer only has 24 bit fixed-point registers. Because time was measured as the number of tenth-seconds, the value of 1/10, which has a non-terminating binary expansion, was truncated at 24 bits. The error in precision from this truncation grows as time increases and the resulting inaccuracy is directly proportional to the target's velocity. This error in expected location of the Scud missile caused the Patriot system to classify the detection of the missile as a "false alarm."

This example of software aging illustrates the severe consequences of this type of CCF. Typically, software aging problems can be reset, but not eliminated, by rebooting the computer. In multiple redundant computer systems, software aging effects can be mitigated by starting and rebooting the redundant computers at different times (i.e., creating a different internal computer state for each redundant computer).

One point to be emphasized from this example is the impact of functional requirements that do not adequately account for the usage conditions. In effect, the inadequate requirement became the source of common faults and extended operation of the usage profile became the triggering condition. The Patriot system was used in a manner inconsistent with its original requirements; it was designed for relatively slow flying aircraft and missiles, but was utilized against much faster flying missiles. Consequently, the software failure was the result of usage profile deviations as well as inadequate resource management capabilities for the execution environment whereby operating system deficiencies became manifest as temporal phenomena or "aging."

2.2.3 Therac-25 Massive Overdoses

Between June of 1985 and January of 1987, six known accidents involving massive radiation overdoses from a computerized radiation therapy machine, the Therac-25, occurred in the United States and Canada. In three incidents, the injured patients later died from radiation exposure. In each accident, the patient was exposed to approximately 100 times the intended radiation dose. They have been described as the worst series of radiation accidents in the 35-year history of medical radiation therapy machines. These accidents, which occurred early in the development of computerized control systems, highlighted the dangers of software control of safety-critical systems and have become a standard case study in health informatics and software engineering.

The Therac-25 was a radiation therapy machine produced by Atomic Energy of Canada, Ltd. (AECL). It was the third-generation machine built by this company, but the first where computerized control was mainly responsible for the safety features of the equipment. AECL had previously produced the Therac-6 and Therac-20 radiation therapy machines under a partnership with Compagnie General Radiographique (CGR) of France.

In order to understand the CCF of the software that caused the accidents, one must first understand some basic operating principles of the equipment. Medical radiation therapy machines are basically electron accelerators. Medical linear accelerators (linacs) accelerate electrons to create high-energy beams that can treat tumors with minimal impact on the surrounding healthy tissue. Relatively shallow tissue can be treated with only the accelerated electrons; the higher the energy of the electrons, the deeper the "depth

dose". Depth dose is a phenomenon in which the location in the body where the maximum dose buildup occurs deepens as the energy of the electrons increase. With depth dose, the tissue above the target depth is spared. To reach even deeper tissue, the electron beam must be converted into X-rays.

The Therac-25 was a dual-mode linear accelerator that could provide either X-rays at 25 million electron volts (MeV) or electrons at various energy levels. When electrons are utilized for treatment, the computer controls the beam energy from 5 to 25 MeV while scanning magnets spread the beam to a safe, therapeutic concentration. The spread beam also passes through an ion chamber before it reaches the patient in order to measure the strength of the treatment. When X-rays are required for treatment, only one energy level of 25 MeV is available and an electron-beam current of approximately 100 times greater than that used with electron-only treatments is required. To produce the X-rays, the high-energy electron beam must impact a "target". The X-rays produced from the target travel through a "beam flattener" which produces a uniform treatment field. The flattener resembles an inverted ice-cream cone and is an excellent efficient attenuator of energy; therefore, a very energetic electron stream is required in order to produce reasonable X-ray dose rates. When X-rays are produced, the magnitude of the X-ray beam sent to the patient is measured by an X-ray ion chamber, which is also in line with the target and beam flattener.

With a dual-mode radiation therapy machine, like the Therac-25, part of the mechanism for switching between the X-ray treatment mode and the electron treatment mode consists of a physical "turntable" which rotates the equipment necessary for each treatment between the output of the accelerated electron beam and the patient. This physical turntable is a basic hazard of a dual-mode machine; if the turntable is incorrectly positioned, the treatment dose is incorrect. With the Therac-25, the turntable was further complicated in that a third position was added which placed a stainless steel mirror over the patient with a light which illustrated the position and area of the beam to assist with positioning the machine accurately for proper treatment. Thus the turntable of the Therac-25 had three positions: (1) light/mirror for positioning the patient, (2) electron beam treatment with an ion chamber and bending magnets between the beam and patient, and (3) X-ray treatment with target, flattener, and ion chamber between the beam and patient.

With the light/mirror positioned above the patient, no beam should ever be present. For electron beam treatment, the beam power could be varied between 5 MeV and 25 MeV and the bending magnets were utilized to shape the beam for the correct treatment where the radiation dose would be measured by the ion chamber. For X-ray treatment, the electron beam would always produce 25 MeV electrons at a very high current to impact the target, be spread and attenuated by the flattener, and finally have the treatment dose be measured by an ion chamber.

Software is responsible for monitoring the machine status, accepting the treatment information, and setting the machine up for the treatment. If there is a hardware malfunction, the computer is informed and, depending upon the seriousness of the fault, either prevents the treatment from being started or, if the treatment is in progress, pauses the treatment temporarily or permanently suspends the treatment. With the Therac-25, the computer software was also responsible for the safety-critical functions of positioning the turntable to ensure that the correct equipment was in place over the patient and the magnitude of the electron beam was correctly set for the appropriate treatment.

The computer software was executed on a Digital Equipment Corporation PDP-11/23 computer and was written in PDP-11 assembly language. The software for the Therac-25 was developed by a single person over a period of several years and evolved from the Therac-6 software. There was little software documentation produced during development and limited or no software specifications or software test

plan. Apparently there was little unit and software testing at the factory with most effort directed at the integrated system test during installation of the units in the field.

The software had multiple issues, and all issues were never identified. The issue that directly caused the six accidents was related to timing of procedures within the execution of the software where a "race condition" could exist between subroutines that were performing operations to set up the machine for treatment. The timing issue involved the operator's input of treatment parameters when a mistake was made and then quickly corrected. If the operator erroneously selected an X-ray treatment and then within 8 seconds both corrected the treatment to be an electron treatment and told the machine to execute the treatment, the electron beam would mistakenly remain set for an X-ray treatment at 25 MeV and high current, but the turntable would be correctly positioned with either the light/mirror or the ion chamber and inactive bending magnets between the intense electron beam and the patient. This resulted in a highly intense radiation dose being delivered to the patient. The software would recognize a problem, but the error code given to the operator was not informative and the dose shown on the operator's display actually indicated an under-dose because the ion chamber, if present, was saturated or, if not present, was exposed to no beam. Furthermore, the software would only temporarily pause the treatment, not permanently suspend the treatment. Often, since the operator's display indicated an under-dose, the operator would instruct the machine to continue, which resulted in a second excessive dose being delivered to the patient. As the machine's operation was "quirky" in this and other aspects, the operators became accustomed to its unusual behavior and continued treatment without question.

Typical radiation therapy treatment with an electron beam might expose a patient to approximately 200 rads based on a 22-MeV beam over a 10 × 10 cm area for each treatment with the patient receiving 5 to 10 treatments over several weeks. When the Therac-25 malfunctioned, it is approximated that the possible dose was 16,500 to 25,000 rads delivered in one second over an area of only about $1cm^2$. The patients could actually feel this intense radiation and would describe it either as an electrical shock or a burning sensation. One patient who was involved in an accident had already received eight successful Therac-25 treatments. When the accident occurred, he knew something was wrong and began to exit the treatment table when the operator, who could neither hear nor see the patient because of unrelated equipment failures, continued the treatment and exposed the patient a second time to the excessive radiation.

While the development of the Therac-25 software occurred in the early years of computer control of equipment, there are still valuable lessons to be learned from these accidents. During this time period, some companies did not treat software engineering as a priority. The lack of documentation, establishment of software quality assurance practices and standards, design of software audit trails, and extensive testing and formal analysis of software at the module and system level are basic software engineering principles that were violated in the Therac-25. The lack of well-defined and documented software specifications as well as the absence of a comprehensive and rigorous software testing program is clearly a significant oversight. For safety-critical software, special safety analysis and design procedures must be in incorporated. Safety must be ensured at the system level despite software errors.

2.2.4 Air France Flight 447 Crash

Faulty air speed indications led to the crash of Air France Flight 447, an Airbus A330-200 airliner, into the Atlantic Ocean on June 1, 2009, while flying from Rio de Janeiro, Brazil, to Paris, France, and

resulted in the death of all 216 passengers and 12 aircrew. The Airbus 330 series of aircraft utilize a modern, highly redundant, fly-by-wire flight control system.

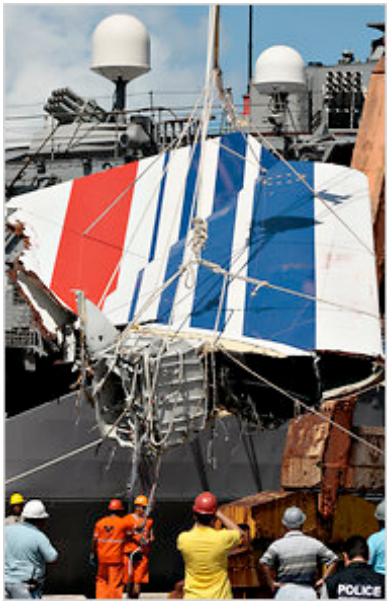

The Airbus 330 flight control system has three primary flight control computers that are responsible for calculations concerned with aircraft control and with sending signals to the actuators associated with the control surfaces and engines and two secondary control computers where control is transferred automatically if the primary computers are unavailable. Hardware diversity in this system is provided by the primary and secondary flight control computers using different processors, being designed and supplied by different companies, and by having processor chips for the different computers being supplied by different manufacturers. Software diversity is provided through development of different channels in each computer (i.e., the command channel and the monitoring channel) by different teams using different programming languages as well as through use of different teams to develop the software for the primary and secondary flight control computers.

Furthermore, the flight control system is dynamically reconfigurable to cope with a loss of system resources. Dynamic reconfiguration involves switching to alternative control software while maintaining system availability. Three operational modes are supported: (1) Normal, (2) Alternative, and (3) Direct. At least two failures must occur before normal operation is lost.

Finally, there is diversity in the implementation of the aircraft control surfaces. The linkages between the flight control computers and the flight surfaces are arranged such that each surface is controlled by multiple independent actuators. Since each actuator is controlled by different computers, the loss of a single actuator or computer does not result in loss of control of that surface. The hydraulic system is three-way replicated, and each redundancy takes a different route though the aircraft.

It seems unimaginable that such a redundant control system could fail with such fatal consequences; however, it did. The Airbus 330 had a history of airspeed indicator problems. The airspeed is determined by pitot probes, and on the Airbus 330, there are three independent systems for calculating and displaying airspeed information: (1) captain, (2) first officer, and (3) standby. Each system uses its own pitot probe, static ports, air data modules (ADMs), air data inertial reference unit (ADIRU), and airspeed indicator. Each pitot probe consists of a tube that projects several centimeters out from the fuselage, with the opening of the tube pointed forward into the airflow. The tube has drain holes to remove moisture, and it is electrically heated to prevent ice accumulation during flight.

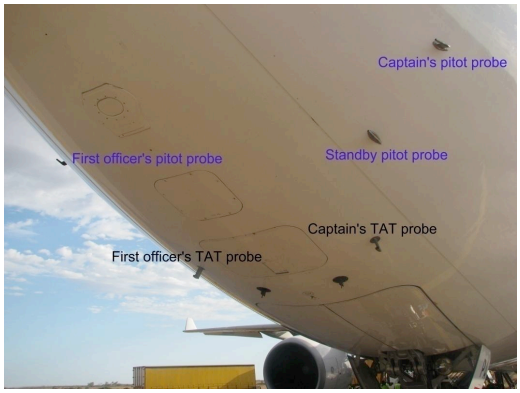

Accurate indications of air speed are crucial to the stability of the airplane. At a cruising height of 35,000 feet, relatively small variations in air speed, either too slow or too fast, can lead to an aerodynamic stall, as they did in the case of Flight 447. Even though there are three independent systems for determining airspeed, the Airbus pitot probes had been experiencing blockage problems when the plane encountered weather in which ice crystals could form. These ice crystals would block the pitot probes, resulting in inaccurate and misleading airspeed indications.

Aircrews from several international airlines had reported evidence suggesting that sudden flash freezing, involving ice crystals, would overcome the heaters in the pitot probes and disrupt airspeed indications. When fed anomalous data from the probes, the Airbus flight control computers, which fly the airplane via autopilot and auto engine throttles, were programmed to shut down. Consequently, the pilots are left to fly out of the emergency manually. The logic governing the computers is called control laws. In nominal operations, the A330 flight control is governed by the normal law. When monitoring triggers a fault, it may be replaced by alternative or direct law. Normal law offers complete protection of the flight envelope (i.e., pitch and bank values are limited, based on expected load factor). In the alternative law, fewer protections exist. In the direct law, the sidesticks control the position of the various control surfaces directly. In alternate or direct laws, angle-of-attack protections are no longer available.

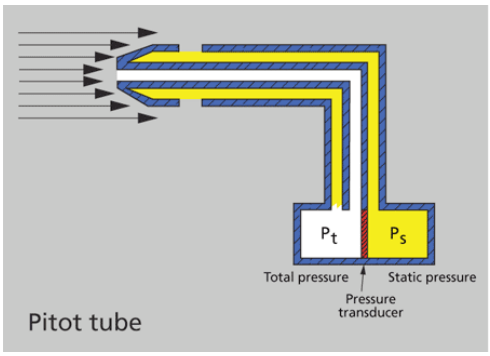

The airspeed presented to the crew is the median value from the three systems. When airspeeds from one of the three airspeed systems deviates too much from the other two, it is automatically rejected and the one presented is the average of the two remaining values. If the difference between these two remaining values becomes too great, they are both rejected and the control law changes to alternate. This was the control law in effect aboard the Air France Flight 447 during the crash. This situation happened under circumstances that challenged the pilots' spatial awareness; it occurred at night and when the airplane was flying through storm turbulence. This common-cause control system failure (failure of multiple airspeed indicators) along with the pilots' inability to accurately perceive and respond to the emergency led to the crash of the aircraft.

This example emphasizes the impact of utilization of a single common element or technology (pitot probes), in even high redundant control systems, in which there is a common external condition applied. Even though there were three pitot tubes connected to three independent airspeed indication systems, there was an overarching condition (ice crystal forming conditions over the entire nose of the aircraft where all three pitot probes were located) that precipitated the failure. This CCF could have been mitigated by the utilization of a different technology for measuring the airspeed of the aircraft in at least one airspeed indication system which would have been unaffected by the formation of ice crystals in the nose area. The common external condition led to formation of ice crystals in all three pitot probes, and the coincidental impact on a critical measurement resulted in the loss of the aircraft.

2.2.5 North American Electrical Blackout of 2003

Beginning on Thursday, August 14, 2003, just before 4:10 PM EDT a widespread power outage occurred in portions of the northeastern and midwestern United States and Ontario, Canada. The power outage affected over 10 million people in Canada and about 45 million people in the United States. Some people were without power for as much as 16 hours, while the majority lost power for only about 7 hours.

The power outage initiated in Ohio within the service area of FirstEnergy (FE), which is centered in the Cleveland-Akron metropolitan area. This area is important to the cause of the blackout because the Cleveland-Akron area is a transmission-constrained load pocket with only limited generation capability. The initiating event, which should have by itself been recoverable, was the shutdown of a power-generating plant in

Eastlake, Ohio, a suburb of Cleveland, in the midst of a high electrical demand at 1:31 PM EDT. The influx of power to replace this generating capacity overheated high-voltage power lines located in distant rural areas. This overheating caused the high-voltage power lines to sag into overgrown trees that resulted in the inability of these power lines to transmit power. As other FE high-voltage power lines attempted to handle the power distribution necessary to compensate for the off-line generator and faulted transmission lines, they too overheated, drooped into overgrown trees, and faulted. This resulted in a cascading effect that eventually forced the shutdown of more than 100 power plants and created simultaneous undervoltage and overcurrent conditions on numerous high-voltage power lines causing their disconnection from the power grid.

The failure of FE to maintain the correct tree height in their transmission right-of-way that eventually caused the faulting of the transmission lines is a CCF resulting from external factors. However, central to the blackout cause was a computer failure at FE that resulted in failure to recognize and understand the deteriorating condition of its power distribution system.

Like all power utilities, FE used supervisory control and data acquisition (SCADA) systems to monitor power system data and control power system equipment. SDADA systems have three types of components: field remote terminal units (RTUs), communication to and between the RTUs, and one or more Master Stations. RTUs installed in generating plants and substations both gather data and control power switching devices. Telephone lines or microwave radio channels provide communication between the RTUs and one or more Master Stations. Often Master Stations are integrated into the control room and serve as the Energy Management System (EMS). The EMS gives transmission system operators visibility regarding their own transmission facilities, and allows them to recognize the impact of their facilities on adjacent power system facilities.

At 2:14 PM EDT, the alarm and logging software on FE's EMS failed while attempting to handle an incoming alarm. The software also failed at this point to notify the EMS operators that it was inoperative. Therefore, neither FE's control room operators nor FE's EMS support personnel were aware that there were no alarms being processed by the EMS. This caused FE EMS operators to run a complex power system without adequate indicators of when key elements of that system were reaching and passing the limits of safe operation. The SCADA system was still receiving and displaying correct information on displays in the EMS control room, but the operators were expecting to be alerted by the malfunctioning EMS alarm system to critical power distribution parameters that required their attention for safe operation of the system and never looked at the relevant displays.

FE's EMS system and their ability to adequately understand and control their power distribution system was further degraded at 2:41 PM EDT when the EMS primary server, the one hosting the failed alarm indicating software, failed either due to the failure of the alarm software, an unrelated "queuing" failure at some remote EMS terminals, or a combination of both. The failed alarm system application and all other EMS software running on the primary server then automatically transferred to the "hot-standby" backup EMS server. However, at 2:54 PM EDT, the backup EMS server also failed because the failed alarm algorithm had transferred to the backup server in the failed condition and remained inoperative on that server as well. With the loss of both the primary and backup EMS servers, the FE power distribution system operators were further inhibited in adequately controlling their system. By 3:08 PM EDT, the servers had been restarted, but the alarm algorithm was still inoperative and both the FE Information Technology (IT) staff and the control room operators continued to be unaware that the alarm system had failed. Because of these computer and software failures, FE controllers were unaware of the collapsing state of their power distribution system and did not take any steps to mitigate the propagation of the effects to adjacent systems.

While physical CCFs were critical to this cascading power outage, software and computer failures were also partly responsible. While the EMS servers had some redundancy and the SCADA system was displaying pertinent information for the operators, the failure of the EMS alarm algorithm was critical to the extent of the power outage.

2.2.6 AT&T Network Outage

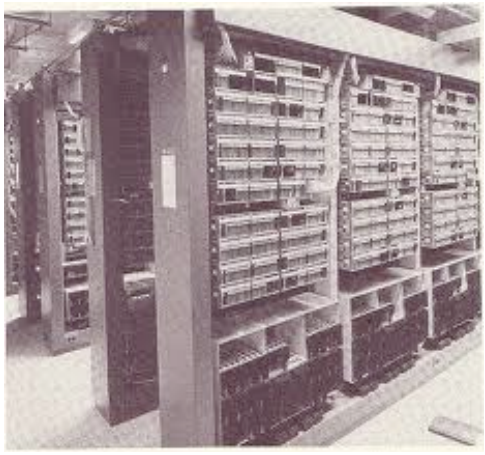

In mid-December 1989, AT&T loaded new software into all 114 of its Class 4 Electronic Switching System (4ESS) switches in the United States. The 4ESS switch was the first telephone digital electronic toll switch for long-distance digital telephone conversation switching. On the afternoon of January 15, 1990 a piece of trunk interface equipment in New York developed an internal problem and notified its associated 4ESS switch that it was having problems and could not correct the problem. The 4ESS switch began running corrective initialization of the trunk interface and halted handling of any new incoming calls while it did so, which takes between 4 to 6 seconds. The New York switch informed all the other switches that it was connected to that it was not taking new calls. Upon successfully initializing the trunk interface equipment, the New York switch began processing new calls.

Prior to the software update, once the first switch began processing calls after the trunk interface re-initialization, it would send another message to all the switches to which it was connected informing them that it was accepting new calls again. After receiving this message, the other switches would confirm that the first switch was indeed working and then accept call routing signals from it. However, the new software changed this protocol. Instead of the first switch messaging the other switches that it was working again with subsequent confirmation, the new software simply had the first switch start sending routing signals to the other switches. Consequently, each switch would interpret the signals as indication that the first switch was again working since they were receiving routing signals from it. This change was made in order for the system to react more quickly and reduce traffic between switches.

The problem occurred when the first switch would quickly send two call routing signals to another switch. When another switch would receive the first new call routing signal, it would begin to update its information about the first switch to note that it was again working. However, if a second call routing signal arrived from the first switch before the second switch had time to completed execute the software to reset itself in respect to the status of the first switch, the second switch would perceive an internal operational error and begin reinitializing itself also notifying all the switches connected to it that it was not accepting new incoming calls. If the second switches did not receive a second call routing from the first switch before they finished resetting themselves, there would have been no problem. If the call routing from the first switch had been further apart in time, it would not have triggered the problem. This software condition caused identical problems in all 4ESS switches around the nation as the chain reaction spread and caused a 9-hour outage of the AT&T long-distance telephone system. AT&T engineers finally loaded the old version of the software into the 4ESS switches to stop the outage, found the coding error which resulted in the problem in the new software, and reduced the messaging load on the internal switch communication network as a final solution.

AT&T finally traced the software problem to an elementary programming error where the programmer misunderstood the effect of the placement of a command within an “if statement” in the coding. This branch of the “if statement” was never exercised during testing or the fault would have been recognized. This example emphasizes the necessity of a comprehensive and complete testing plan for

software routines and the necessity of redundancy in both hardware and software to overcome any overlooked or unforeseen software errors.

2.2.7 Intel Pentium Chip Design Fault

In March 1993, Intel released its fifth-generation processor, the Pentium, to great anticipation and expectation within the computer community. However, by November 1994, Intel was being heavily criticized because they had released a defective product that they knew contained a flaw. The flaw, which became known as the Pentium FDIV bug, produced incorrect results during certain floating point division operations. According to Intel, these incorrect floating point division operations were caused by a few missing entries in the lookup table used by the digital divide operation algorithm. Although actually encountering the flaw in practice was very rare, estimated at one in nine billion floating point divides, Intel's release of the flawed product and its handling of the matter after public disclosure were heavily criticized and ultimately lead to the recall of the defective product.

The underlying cause of the missing entries in the lookup table was a programming error in a script that downloaded these entries from a numerically generated process into the lookup table. There are only five missing entries in the 16 X 48 entry lookup table.

Despite the limited actual impact upon users, in December of 1994, Intel offered to replace all flawed Pentium processors because of the publicity of the problem. The financial impact of this decision was significant, which was estimated by Intel to be approximately $475 million. This case provides an example of a design implementation error that exists in every instance of the computing platforms that are based on this generation of the Pentium microprocessor.

3. GUIDANCE ON CCF MITIGATION

3.1 Nuclear Power Regulatory Guidance on CCF Mitigation

The overall I&C system architecture of an NPP embodies the fundamental safety principle that safe conditions must be maintained under all operational conditions (i.e., normal, abnormal, anticipated operational occurrences, and design basis accidents) as a primary objective of its design and implementation. A key approach to achieving this objective is the provision of multiple means to ensure public health and safety through strategic implementation of defense-in-depth.

Defense-in-depth may be visualized in terms of a concentric arrangement of protective barriers. Before any harmful radiological release could occur to adversely affect the public or the environment, all of the barriers (i.e., fuel rod cladding, reactor coolant system pressure boundary, containment, and emergency response) must be breached. I&C systems have an important role in maintaining the integrity of these barriers. The application of defense-in-depth to the I&C system architecture of an NPP is accomplished by incorporating independent echelons of defense (or lines of defense). Defense-in-depth for I&C systems provides multiple systems to provide independent means to maintain desired operational conditions, prevent accidents, and ensure adequate protection during adverse events (e.g., failures).

Typical I&C echelons of defense are the control system, the reactor trip system (RTS), the engineered safety features actuation system (ESFAS), and the monitoring and indicator system. These echelons can be considered to act as progressively compensating systems with some overlapping capabilities that collectively achieve the safety objectives of an NPP even if one or more of the systems or echelons fail. The means of accomplishing a safety objective for a specific echelon of defense can involve either avoidance of adverse conditions or mitigation of their effects.

Within the protection echelons of defense (i.e., RTS and ESFAS), I&C systems are designed to withstand single failures to ensure accomplishment of safety functions even in the presence of random failures. The single-failure design criterion is generally achieved through the implementation of independent, parallel channels or divisions within a safety system in which redundant safety outputs are voted to determine whether to initiate an appropriate safety action. For these safety systems, functional failure occurs if the output of the voting yields an erroneous result, such as a spurious actuation or failure to act on demand. Thus, functional failures for these systems require multiple redundancies (a voting majority) to fail concurrently in conjunction with a safety demand. CCFs affecting multiple redundancies or systems within or among echelons of defense constitute the principal credible threat to defeating the defense-in-depth provisions within the I&C system architecture of an NPP.

Diversity is the general mitigation approach used for addressing perceived vulnerabilities to CCF of I&C system architectures because dissimilarities in technology, function, implementation, and so forth can mitigate the potential for common faults. Whereas the defense-in-depth approach to ensuring safety employs different functional barriers to compensate for failures in any one or more of the lines of defense, the diversity approach to ensuring safety uses different (i.e., dissimilar) means to accomplish the same or equivalent function, generally within one functional barrier, to compensate for a CCF that disables one or more echelons of defense. Domestic and international nuclear power regulations and regulatory guidance address both defense-in-depth and diversity as means of mitigating single and common-cause failures.

3.1.1 Regulatory Guidance on CCF Mitigation in the United States

U.S. NRC regulations require licensees to incorporate into an NPP an overall safety strategy for defense-in-depth functions and systems to ensure that abnormal operating occurrences (AOOs) and design basis accidents (DBAs) do not adversely impact public health and safety. In particular, the design criteria for NPP safety systems embody principles such as high quality, integrity, reliability, independence, and

qualification. Separation and redundancy, as well as physical barriers and electrical isolation, are generally applied as design measures to address potential vulnerabilities related to a single failure of equipment and the propagation of failure effects [6,7]. These measures tend to minimize shared components or equipment and nonessential interconnections within I&C system architectures. Nevertheless, the potential for CCF vulnerability has long been recognized and diversity is therefore employed as a contributing factor in satisfying safety requirements. For example, the failure of reactor trip functions, which would require the concurrent failure of more than one redundant channel or division in an RTS, is addressed through regulatory requirements for provision of diverse equipment/systems to respond to anticipated transients without scram (ATWS).

The general design criteria (GDC), provided in Appendix A of Title 10, Part 50 of the Code of Federal Regulations (10 CFR 50) [8], establish the minimum design requirements for light-water reactors (LWRs). The introduction to Appendix A explicitly states that "the possibility of systematic, nonrandom, concurrent failures of redundant elements in the design of protection systems and reactivity control systems" needs to be considered. Several of the GDC for protection systems deal with issues that are relevant to mitigation of potential CCF vulnerabilities. Criterion 21, Protection system reliability and testability, requires the capability to withstand any single failure and identifies redundancy and independence as specific design approaches. Criterion 22, Protection system independence, addresses the assurance that the safety function will be provided to accommodate the "effects of natural phenomena, and of normal operating, maintenance, testing, and postulated accident conditions on redundant channels." In particular, GDC 22 requires that "functional diversity or diversity in component design and principles of operation … be used to the extent practical to prevent loss of the protection function." Criterion 23, Protection system failure modes, specifies that a safe state be achieved in response to failures that may result from adverse environments or other anticipated conditions, such as loss of power. Criterion 24, Separation of protection and control systems, invokes separation as a design measure to minimize the prospect of dependencies that could challenge the reliability, redundancy, and independence requirements of the protection system. Criterion 26, Reactivity control system redundancy and capability, requires the provision of two reactivity control systems based on different design principles. Finally, Criterion 29, Protection against anticipated operational occurrences, states that protection system designs must provide an "extremely high probability of accomplishing their safety functions" when challenged by AOOs. Additional relevant design criteria are also provided by the incorporation of Institute for Electrical and Electronics Engineers (IEEE) Std. 603-1991, "Standard Criteria for Safety Systems for Nuclear Power Generating Stations" [6] and IEEE Std. 279-1971, "Criteria for Protection Systems for Nuclear Power Generating Stations" [9], in 10 CFR 50.55a(h).

As seen above, diversity usage for mitigation of potential CCF vulnerabilities is specifically cited in the design criteria as well as being required by regulation (i.e., the ATWS rule in 10 CFR 50.62). The consequence of these regulatory requirements is that diversity approaches, such as the combination of functional and signal diversity, have been extensively employed for conventional (i.e., hardwired) safety systems. These "traditional" diversity strategies remain effective in addressing criteria such as GDC 22. However, the increased potential for CCF vulnerability posed by the unique characteristics of digital technology was found to warrant consideration of additional diversity usage to supplement the traditional diversity strategies. Specifically, the NRC staff expressed its concerns about digital safety systems, including potential CCF vulnerabilities, in SECY 91-292, "Digital Computer Systems for Advanced Light-Water Reactors" [10]. In item II.Q of SECY 93-087, "Policy, Technical, and Licensing Issues Pertaining to Evolutionary and Advanced Light-Water Reactor (ALWR) Designs" [11], the NRC staff documented a four-point position on diversity and defense-in-depth that was subsequently modified in the associated staff requirements memorandum (SRM), dated July 21, 1993 [12].

The NRC four-point position establishes requirements for addressing the potential for CCF vulnerability. The position points are as follows

"1. The applicant shall assess the defense-in-depth and diversity of the proposed instrumentation and control system to demonstrate that vulnerabilities to common-[cause] failures have been adequately addressed.

2. In performing the assessment, the vendor or applicant shall analyze each postulated common-[cause] failure for each event that is evaluated in the accident analysis section of the safety analysis report (SAR) using best-estimate methods. The vendor or applicant shall demonstrate adequate diversity within the design for each of these events.

3. If a postulated common-[cause] failure could disable a safety function, then a diverse means, with a documented basis that the diverse means is unlikely to be subject to the same common-[cause] failure, shall be required to perform either the same function or a different function. The diverse or different function may be performed by a nonsafety system if the system is of sufficient quality to perform the necessary function under the associated event conditions.

4. A set of displays and controls located in the main control room shall be provided for manual, system-level actuation of critical safety functions and monitoring of parameters that support the safety functions. The displays and controls shall be independent and diverse from the safety computer system identified in items 1 and 3 above."

As discussed in SECY 93-087, the four-point position on D3 was generated because hardware design errors, software design errors, and software programming errors are credible sources of CCF for digital safety systems. The safety significance of these potential digital CCFs arises from the prospect that architectural redundancy within a safety system could be defeated and more than one echelon of defense-in-depth could be compromised. The position enhances guidance on addressing the potential for CCF vulnerabilities that arise from conventional (i.e., analog) I&C implementations of safety-related functions (e.g., GDC 22, 10 CFR 50.62) by addressing the unique characteristics and concerns related to digital technology while remaining consistent with that guidance.

It is noted in SECY 93-087 and SECY 91-292 that quality, independence, and diversity are principal factors in defending against CCF vulnerabilities. Criteria for ensuring adequate quality and independence are established in Appendix B of 10 CFR 50 and as part of the design criteria provided in IEEE Std. 603-1991 and IEEE Std. 7-4.3.2, "IEEE Standard Criteria for Digital Computers in Safety Systems of Nuclear Power Generating Stations" [13], which is endorsed in Regulatory Guide 1.152, Revision 2, "Criteria for Use of Computers in Safety Systems of Nuclear Power Plants" [14].

Criteria for assessing adequate diversity are provided within the review guidance given in Branch Technical Position (BTP) 7-19, "Guidance on Evaluation of Defense-in-Depth and Diversity in Digital Computer-Based Instrumentation and Control Systems," [15] in Chapter 7, "Instrumentation and Controls," of NUREG-0800, *Standard Review Plan for Review of Safety Analysis Reports for Nuclear Power Plants* [16]. The objective of BTP 7-19 is to confirm that vulnerabilities to CCFs have been adequately addressed by accomplishing the following:

- verification that "adequate diversity has been provided in a design to meet the criteria established by the NRC's requirements,"
- verification that "adequate defense-in-depth has been provided in a design to meet the criteria established by the NRC's requirements," and
- verification that "the displays and manual controls for critical safety functions initiated by operator action are diverse from computer systems used in the automatic portion of the protection systems."

The review guidance in BTP 7-19 expresses the key concern associated with the potential for CCF vulnerability posed by digital technology. Specifically, "[s]oftware cannot typically be proven to be error-free and is therefore considered susceptible to common-cause failures because identical copies of the software are present in redundant channels of safety-related systems." The D3 assessment method documented in NUREG/CR-6303, *Method for Performing Diversity and Defense-in-Depth Analyses of Reactor Protection Systems* [17], is cited as acceptable for demonstrating that "vulnerabilities to common-cause failures have been adequately addressed" [15].

NUREG/CR-6303 provides guidance on performing a D3 assessment to determine the CCF vulnerability of an NPP I&C system architecture. As a first step in a D3 analysis, a decomposition of the NPP I&C system architecture into a block representation is performed and a determination is made of which blocks are susceptible to a postulated CCF. The assessment of CCF vulnerability involves identification of common elements, interdependencies (e.g., physical, logical), and diversities. Evaluation of defense-in-depth is performed by postulating concurrent failures of identical (or nondiverse) blocks in all redundant divisions or lines of defense while performing "best-estimate" safety analyses. If the estimated plant response exceeds specified limits for any AOO or DBA in the presence of postulated CCF, then a CCF vulnerability exists and corrective action, such as the introduction of additional diversity, should be taken to ensure adequate protection is provided, unless the choice of no corrective action can be otherwise justified. Where the D3 assessment determines that additional diversity is needed to mitigate an identified CCF vulnerability of one or more safety functions, that diversity can be achieved through provision of a separate automatic system to back up the affected safety function(s) or through the introduction of intentional diversity and compensating design measures at the appropriate lower level(s) of the I&C system architecture (e.g., system, divisional redundancies, subsystems, modules, or components).

NUREG/CR-6303 separated diversity attributes into the following six areas to facilitate assessments of adequate diversity in safety systems:

- design diversity,
- equipment diversity,
- functional diversity,
- human diversity,
- signal diversity, and
- software diversity.

The guidance in NUREG/CR-6303 provides a set of recommended criteria for each of the six diversity attributes with several diversity criteria within each attribute. However, because of the number of criteria in each attribute coupled with the number of attributes, the number and complexity of possible combinations of attributes that could be used to achieve adequate diversity in a safety system make the guidance very difficult to use as a safety assessment tool. Consequently, a subjective judgment is required to determine what diversity usage is adequate to mitigate identified CCF vulnerabilities.

3.1.2 Common Regulatory Position on CCF Mitigation in Europe

The Western European Nuclear Regulators' Association (WENRA) invited European safety authorities to contribute to the completion of a common position on the licensing of safety-critical software. The objectives of this effort were determination of best practices concerning key licensing issues posed by computer-based implementations of safety functions at NPPs and establishment of a consensus position. The work group assembled for this effort, which continued a collaborative exchange that began in the mid-1990s, consisted of a group of regulators and safety experts representing seven organizations from six countries. The participating organizations are Association Vinçotte Nuclear (AVN) of Belgium, Säteilyturvakeskus (STUK) of Finland, Bundesamt für Strahlenschutz (Federal Office for

Radiation Protection—BfS) and ISTec of Germany, Consejo de Seguridad Nuclear (Nuclear Safety Council—CSN) of Spain, Statens Kärnkraftinspektion (SKI) of Sweden, and Nuclear Installations Inspectorate (NII) of the United Kingdom. The outcome of this collaborative interaction is a report documenting the common position of the participating safety authorities [18]. The report is directly available from any of the seven organizations.

The common position of the European regulators consists of consensus requirements (based on unanimous agreement) and recommended practices (based on general agreement) addressing key licensing considerations. The clauses that constitute the common position explicitly apply to safety systems and relate to issues arising from the use of digital and programmable technology. These issues address generic and life-cycle-phase aspects of licensing computer-based safety systems. The topics addressing specific stages of the design and development process for digital safety systems are as follows:

- computer-based system requirements,
- hazard analysis,
- safety demonstration,
- reliability targets,
- defense against CCF,
- communication system design,
- fault-tolerant architectures,
- software design and structure,
- coding and programming directives,
- diversification and testing (plans, coverage, and traceability),
- validation and commissioning,
- change control and configuration management, and
- operational requirements.

The topics with general or full life-cycle implications are safety demonstration, safety categories and graded software requirements, reference standards, use and validation of preexisting software, tools, organizational requirements, software quality assurance program and plan, security, use of formal methods, independent assessment, graded requirements for software of safety-related systems, software design diversity, software reliability, and data collection for operational experience. Clearly, the requirements for software design diversity are of particular relevance to this research effort.

For the design of computer-based safety systems, the common position requires that “principles of redundancy, diversity, physical isolation, segregation, and separation between safety functions, safety related functions and functions not important to safety” be applied to computer system architecture design. These principles address considerations such as reliability and independence while providing protection against CCF. Architectural and other design decisions influence the necessity and the nature of the software design diversity employed. The adoption of a simple hardwired system as the diverse alternative to a computer-based safety system can resolve software-related CCF concerns. In fact, this approach is emphasized as a best-practice recommendation for this topic. However, it is recognized that multiple computer-based diverse systems are more likely to be adopted given the increasing prominence of digital technology; therefore, specific requirements for ensuring software design diversity are provided.

The common position on software design diversity addresses design decisions or measures that invoke methods, techniques, and measures to force software design diversity. The goal is to diversify failure behavior among diverse software-based systems. Functional diversity is the foremost design measure identified in the common position, and it is required to be implemented whenever possible for safety system elements that are intended to be diverse. Additionally, the functionally diverse systems are required to be associated with the same safety class and subject to the same graded requirements. Other design decisions or measures specified for the design of computer-based systems are given as follows:

- independence of development teams (with no direct communication between teams);
- different description languages (e.g., specification languages) and notations;
- different programming languages;
- different development methods;
- different development platforms, tools, and compilers;
- different hardware; and
- diverse verification and validation (e.g., back-to-back testing).

It is required that the safety demonstration provide an analysis of potential CCFs with justification of the impact of diversity usage on reliability and CCF potential arising from any commonalities in the product (e.g., systems, redundancies, and components) or process (e.g., life-cycle activities and resources). Simplicity of design and implementation is also emphasized to keep complexity of the system and software to a minimum that is commensurate with satisfying safety requirements. Thus, the common position includes a reliance on sufficiently detailed analysis to determine the need for diversity, confirm the types of diversity providing the appropriate mitigation, and justify omissions of diversity where need is indicated. It is similar to the NRC guidance described above in that it relies on a subjective assessment of when diversity is required and how much diversity is adequate.

3.2 International Nuclear Power Standards

The IEC issues and maintains international normative standards for all electrical, electronic, and related technologies. These standards are developed according to consensus procedures by technical experts supplied by the national committees of participating countries. Subcommittee (SC) 45A, Instrumentation and Control of Nuclear Facilities, of technical committee (TC) 45, Nuclear Instrumentation (TC45/SC45A), has responsibility for standards that apply to I&C systems important to safety in nuclear-energy-generation facilities (e.g., NPPs). These standards cover the entire life cycle of I&C systems at these facilities, ranging from conception through design, manufacture, test, installation, commissioning, operation, maintenance, aging management, modernization, and decommissioning.

The guidance provided in key international standards constitutes the basis for an overall approach to coping with CCF in I&C systems important to safety. IEC 61513 represents the high-level guidance addressing I&C system architecture considerations. IEC 60880 supplements that guidance by specifically addressing software-based system considerations. IEC 62340 provides a framework for establishing a CCF coping strategy that is consistent with the high-level requirements in IEC 61513 and complementary to the software requirements in IEC 60880.

3.2.1 IEC 61513

The IEC standard that covers the system aspects of I&C systems important to safety, including computer-based systems, is IEC 61513, “Nuclear Power Plants—Instrumentation and control for systems important to safety—General requirements for systems” [19]. This top-level standard for I&C systems important to safety at NPPs is the nuclear power industry derivative of the multipart parent document on functional safety of industrial process measurement and control systems (i.e., IEC 61508, “Functional safety of electrical/electronic/programmable electronic safety related systems”). Comparable to the parent standard for general industrial-sector application, IEC 61513 defines a life-cycle process for I&C systems important to safety at NPPs and contains the top-level requirements on system functions, architecture, and I&C system design for application to those I&C systems. These requirements are intended to be independent of technology and apply to hardwired (i.e., analog) and software-based (i.e., digital) systems.

IEC 61513 requires analyses to verify the I&C architecture design at an NPP. A specified analysis that must be conducted is an “evaluation of the effectiveness of measures used to reduce the sensitivity of

the safety groups to CCF" with an emphasis on Category A (i.e., safety) functions. As part of this analysis, common components, identical hardware, and identical software must be determined. Where such commonalities are identified, justification must be provided to demonstrate that the potential for CCF is low.

Correspondingly, IEC 61513 gives requirements for defense against CCF. As noted, the standard emphasizes I&C systems that perform Category A functions in addressing defense against CCF within I&C systems important to safety. Categorization of function is provided in IEC 61226, "Nuclear power plants—Instrumentation and control systems important for safety—Classification" [20], based on the consequence of malfunction. Category A functions are safety functions that play a principal role in the safety of the NPP. These functions are implemented in Class 1 systems (i.e., protection systems, safety actuation systems, emergency power actuation systems).

In IEC 61513, the design goal for defending against CCF is specified as providing "measures against the occurrence of a CCF within I&C systems implementing different lines of defence against the same PIE" [postulated initiating event]. The identified measures include the following:

- design provisions promoting tolerance of hazardous plant events (e.g., external influences and internal hazards),
- design provisions resulting in insensitivity to plant demand design (e.g., decoupling execution from plant status to avoid common triggering conditions),
- design provision to minimize the use of common elements or support systems among lines of defense,
- quality assurance and fault tolerance to minimize the potential impact of systematic faults,
- strategic design decisions to manage complexity, and
- design differences through application of diverse features.

For each design measure, requirements and recommendations are given to guide the usage of these defensive approaches. This guidance is briefly reviewed below, with a subsequent focused treatment of the specific guidance on diversity usage as a CCF defense.

Design provisions enabling hazard tolerance include separation, independence, prevention, and compatibility (e.g., electromagnetic and environmental).

Minimizing the risk of common triggering conditions arising from demand profile involves analysis of I&C components to identify loadings (e.g., electrical, computational) that are demand dependent and reduction of the coupling between I&C system operation and plant conditions.

Avoidance of common elements involves architectural provisions such as independence across different lines of defense for I&C systems protecting against the same PIE, independent monitoring and control capabilities to ensure safety functions in the event of a failure, minimized potential for CCF within independent manual control capabilities that back up automatic safety functions, and arbitration or prioritization of commands for engineered safety feature (ESF) actuation that may conflict during failure conditions.

Measures to reduce the risk due to systematic faults include application of high-quality planning for development and manufacturing life-cycle activities, provision of self-supervision capabilities (e.g., exception-handling routines, watchdog timer, plausibility-checking algorithms), and definition and annunciation of a safe state to be achieved upon detection of failures.

Analysis of the I&C system architecture and individual system designs contributes to managing complexity. Such an analysis involves consideration of the degree to which either computer-based or hardwired technologies are employed and the reliance on human action to ensure that safety functions are maintained.

The design measure of interest for this research is the provision of diversity as an effective means for defending against CCF. Diversities that are identified in IEC 61513 include human diversity, signal diversity, functional diversity, design and test diversity, software diversity, and equipment diversity.

Specific guidance on diversity usage involves recommended practices more than requirements per se. The standard recommends that diversity be used to achieve high reliability when uncertainties exist in the evaluation of a design. Combinations of signal and functional diversities are cited as "particularly effective methods to reduce risk of CCF due to errors in the requirements specifications or in the specification and implementation of application software." For complex I&C systems where there is a limited experience base, equipment diversity is identified as a means to address hardware CCF and contribute to defense against system software faults. Use of diverse methods or procedures for verification and validation is cited as a means to contribute to CCF avoidance without introducing design complexity. Examples of this approach include back-to-back testing with a simulator and use of different testing facilities. Finally, it is required that the effectiveness of any diversity usage that is claimed to minimize the potential for CCF be analyzed and documented with appropriate justification.

3.2.2 IEC 60880

IEC 60880, "Nuclear power plants—Instrumentation and control systems important to safety—Software aspects for computer-based systems performing Category A functions" [21], supplements IEC 61513 by providing "requirements for the software of computer-based I&C systems of NPPs performing functions of safety Category A." The second edition of this standard encompasses both the first edition, issued in 1986, and the supplemental part 2, issued in 2000, along with updated requirements covering the software aspects of the I&C system life-cycle process (as defined in IEC 61513). Additionally, IEC 60880 includes an informative annex on defense against CCFs as well as other annexes on details for the safety software life-cycle process, software requirements and software development, tools for software qualification, and requirements on preexisting software.

In particular, IEC 60880 provides requirements for defense against "software design and coding faults" that can result in the potential for CCF in software-based implementations of Category A functions. The standard states that software "by itself does not have a CCF mode." Instead, CCF is a system failure issue that arises from "faults in the functional requirements, system design, or in the software." Thus, the standard recommends that the potential effects of software CCF be considered in the application of the defense-in-depth principle, with appropriate countermeasures employed throughout the development and evaluation processes. In particular, these countermeasures should be considered in the design, implementation, verification, and validation of each layer of defense and in the assessment of independence and diversity among redundant layers of defense. It is noted that diversity usage may not only reduce the potential for CCF but also enhance reliability of some I&C systems.

The nature of CCF, as described in IEC 60880, is that faults may exist undetected in software until challenged by a specific unanticipated or untested signal trajectory. Thus, the mechanism for CCF is the presence of at least one common latent fault within systems or redundancies that defend against the same PIE and the coincident exposure to specific signal trajectories in a sensitive time frame. IEC 60880 specifically addresses faults arising from the software engineering process.

The standard states that high-quality software engineering practices are the most important defense against software CCF. It is also noted that the use of self-monitoring features can help to limit the potential impact of software CCF. However, since error-free software cannot be ensured in general, IEC 60880 requires an analysis of the potential sources and consequences of software CCF as part of the I&C architecture design assessment. The guidance provided for the analysis is consistent with the guidance on D3 assessments given in NUREG/CR-6303.

Guidance regarding the use of diversity as a countermeasure to address software CCF is given as recommended practices. The primary implementation strategy identified is the use of functional diversity among independent systems. If functional diversity is not feasible, then consideration of system diversity, diverse software features, and diverse design approaches is advised. It is required that justification of the strategy employed be documented. Specific techniques to address the software implementation are identified as diversification of the operational conditions for the software, avoidance of failure propagation paths, mitigation of the impact of CCF, and use of different specifications for different software implementations of the same functional requirements. It is noted that N-version programming is not recommended.

Informative discussion of CCF considerations and diversity options is given in Annex G. This information is not considered part of the normative guidance of the standard. Commonalities that can result in CCF vulnerability are identified as including common software, architecture, algorithms, development methods, tools, implementation methods, staffing, and management. A discussion of the role of signal trajectories in triggering CCFs is provided. Also, the impact of abnormal hardware failures, plant conditions, and events that result from unforeseen signal trajectories, which include unexpected software states, is noted. The annex presents specific diversity features that can be considered for resolving software CCF. These features include the following:

- software diversity features (e.g., functional diversity, different design specifications, and different functional implementations);
- diversity at the system level (e.g., independent diverse actuation systems, different basic technology, different types of computers, hardware modules and major design concepts, and different classes of computers);
- diverse design approaches (e.g., algorithms, system data, hardware for inputs or interfaces, timing and sequencing);
- different design and implementation methods (e.g., languages, compilers, support libraries, software tools, programming techniques, system and application software, software structures, and data);
- diverse testing; and
- diverse management approaches (e.g., separation of design teams, forced diversity between design teams, restricted communication between teams, and different staff).

The potential benefit of functional or software diversity usage is derived from the increased protection against software CCF arising from adequately diverse versions. However, it is noted that potential disadvantages can include greater overall complexity, increased risk of spurious actuation, more complex specifications and design, modification problems (e.g., maintaining diversity during modification), cost, and potential lower quality of diverse versions. Thus, the impact on the reliability of safety functions should be considered in the justification of diversity usage.

3.2.3 IEC 62340

The IEC has recently issued a new standard addressing means to cope with CCFs in I&C systems that perform Category A functions (e.g., safety systems). The standard is IEC 62340, “Nuclear power plants—Instrumentation and control systems important to safety—Requirements for coping with common cause failure (CCF)” [4]. Specifically, IEC 62340 gives requirements regarding the avoidance and mitigation of CCF and provides principles to promote independence among I&C systems.

In providing a strategy to cope with CCF, IEC 62340 discusses the conditions that cause CCF. Basically, the standard adopts the position that a CCF can occur only when two factors are present concurrently:

- a latent systematic fault exists, and
- a corresponding triggering mechanism is activated by a signal trajectory.

The standard defines a "signal trajectory" as the "time histories of all equipment conditions, internal states, input signals and operator inputs which determine the outputs of a system." A "latent fault" presupposes that the fault is not identified by validation testing, self-supervision, or periodic testing in the field. Also, latent systematic faults may originate from any phase of the life cycle (e.g., design phase, manufacturing phase, operational procedures).

Systematic faults within I&C systems may result from human errors in design or implementation (considered to be technology independent) or may arise from physical effects during the manufacturing process (considered to be technology dependent). Common sources of these faults include flaws in the safety function requirements or system specifications, inadequate determination of external (e.g., environmental) stress factors or hardware design limits, and design deficiencies. Systematic faults can also be introduced during maintenance, because of limited analysis and testing during modification. These faults can result from activities such as modification of setpoints, use of revised versions of spare parts, or modernization of I&C system components.

Triggering conditions may be caused by external factors such as common demand profiles (e.g., signal transients), environmental stress, or temporal dependencies (e.g., specific real time or calendar dates). Signal trajectory triggers can involve not only input signal transients but also internal states of digital systems and past execution history. Additionally, the existence of fault propagation mechanisms (e.g., communication interlinks) may propagate failure through mechanisms such as functional dependencies, corrupted data, or failed communication processes to cause consequential failure of other redundancies.

The strategic approach to coping with CCF involves reducing the likelihood of systematic faults being incorporated into independent systems or redundancies, minimizing the presence of failure propagation paths among systems, and reducing the possibility of concurrent exposure to triggering conditions. Accordingly, IEC 62340 provides requirements to establish a coping strategy for CCF. These requirements are grouped in terms of four areas of impact, which are characterized as follows:

- overcoming flaws in the requirements specification,
- preventing coincident failures through design measures,
- tolerating postulated latent software faults, and
- avoiding system failure due to maintenance during operation.

The requirements provided in each area are summarized in the following sections.

3.2.3.1 Overcoming Flaws in the Requirements Specification

It is noted that flawed requirements can lead to systematic faults that create the potential for CCF vulnerability. IEC 62340 states that functional diversity serves as an effective means of coping with the prospect of such faults through the provision of alternate requirements as the basis for diverse systems, subsystems, or redundancies. To enable this coping strategy, an analysis of DBAs and relevant design basis events (DBEs) that are affected by I&C system CCF must be performed. It is noted that most large transients influence nearly all safety parameters in parallel. Thus, the application of functional diversity requires a more detailed analysis of DBEs as a precondition. From this analysis, the subset of DBEs that could cause unacceptable consequences in the presence of I&C system CCF is determined and at least one alternate safety parameter must be identified for each event. On this basis, the specification of diverse safety functions is established and can be implemented through a selected design strategy, subject to demonstration that plant safety targets are achieved. Two prospective design strategies are noted: (1) to group diverse safety functions into independent systems to give full coverage by either system and (2) to

implement the complete set of functions in a primary safety system with a reduced-scope set of functions covered by a lower safety class system based on diverse equipment.

The application of functional diversity in concert with the defense-in-depth principle requires "the identification of those specific safety I&C functions that can ensure independently that the main plant safety targets are met." These diverse safety functions must be allocated to independent I&C systems that are implemented in an architectural arrangement such that plant safety is maintained even in the presence of a postulated failure of one I&C system. Essentially, the failure of one I&C system must not affect the other I&C systems that provide compensating safety functions or lines of defense. The independent performance of the diverse safety functions must be validated and documented.

3.2.3.2 Preventing Coincident Failures through Application of Design Measures

Independence is an essential element of any coping strategy because it enables the impact of CCF to be limited to a single I&C system. The principle of independence is satisfied if a postulated failure of one I&C system does not prevent the other I&C systems from performing their intended safety functions. Effective design principles to defend against CCF begin with requirements that ensure high-quality, high-integrity I&C systems. Adherence to the requirements of existing standards is reinforced in this standard. Specifically, the relevant requirements that must be fulfilled are cited as the following (with the referenced standard identified):

- system design: IEC 61513,
- software design: IEC 60880,
- physical separation: IEC 60709 [22], and
- component qualification: IEC 60780 [23] (environmental) and IEC 60980 [24] (seismic).

In addition to the requirements in the standards above, additional requirements are provided by IEC 62340 to ensure the independent performance of diverse safety functions. Some of these requirements involve analyses of potential CCF mechanisms present in the design. In particular, an analysis of the plant I&C architecture is required to determine whether there exist common mechanisms that could compromise the independence of the diverse I&C systems. It is required that any identified vulnerability be either eliminated or resolved through adequate mitigation. Additionally, an assessment of expected operating conditions for diverse I&C systems must be performed to identify any common triggering conditions to be addressed.

Other design requirements specified in IEC 62340 address particular design measures that are considered effective in promoting independence and coping with CCF. First, "system specific processing paths from sensing the plant status to the actuation of plant safety functions" must be provided without employing any shared components. Second, support systems such as power supplies or heating, ventilation, and air-conditioning (HVAC) must provide sufficiently redundant and separated subsystems. Third, self-supervision must be provided independently for each processing unit. Fourth, functional diversity must be used wherever practical for diverse I&C systems.

In executing the design of independent diverse I&C systems, several design considerations must be addressed. First, the design of these systems must reduce the likelihood that the same input signal transient can initiate a CCF to a level that is not significant at any time during the life of the plant. Essentially, measures must be invoked to ensure that each system is subjected to different signal trajectories. Second, no shared components or services are permitted if their postulated failure can cause a CCF of the independent diverse I&C systems. Third, an analysis of the potential for CCF must be performed to assess the impact of identical hardware or software in independent diverse I&C systems. If the resulting potential for CCF is not negligible, then operation of the systems must be restricted such that they are (1) subjected to different service conditions and operational loads (e.g., input and/or processing demands) or (2) not operationally dependent on the demand profile of the plant process and the

corresponding environmental conditions. Essentially, the diverse I&C systems must either be exposed to different signal trajectories and external influences or be insensitive to those factors. Fourth, if diversification of the demand profile as previously described is not feasible, then qualification for the intended application must be ensured and periodically tested. Alternately, equipment diversity may be analyzed for consideration.

For software-based I&C systems, it is required that each software module of the application, as well as the associated signal trajectories, be assessed for potential CCF vulnerability. In particular, functional diversity is required to diversify the input signal component of the signal trajectories and introduction of other diversities to the system designs must be considered to diversify the internal state component of the signal trajectories. Additionally, independent diverse I&C systems must not perform identical application functions since the possibility exists that "coincidental, quasi-synchronized failure of these systems maybe triggered from the same input signal transient."

Regarding the treatment of system communications, requirements are given to ensure that failure propagation through communication paths is avoided. Specifically, communication is not permitted between independent I&C systems that are provided to protect against the impact of CCF. Additionally, requirements addressing internal propagation paths within safety systems are stated. These design measures include detection of data correctness on receipt, exclusion of faulty data from processing, physical separation of redundant subsystems, and protection of safety functions from the effects of communication failure (e.g., failure of the transaction or failure of the subsystem handling communications). In particular, system operation must not be jeopardized by failure of any central subsystems that require communication to more than one redundancy of a safety system to accomplish their information exchange function. For example, these subsystems "may provide information to the main control room for display or may support modification of parameters derived from the plant process." Furthermore, it is required that all software functions provided for the transfer of messages be implemented in a manner that ensures that the correct execution of these transfer mechanisms cannot be compromised by the information content (e.g., data values) being communicated.

The potential for system failure to be induced by maintenance activities must also be addressed in the design of independent I&C systems. Specifically, the safety system design must be analyzed to ensure that maintenance and test activities are properly accommodated by (1) means to prohibit spurious actuation due to maintenance and (2) provisions to limit the simultaneous impact of maintenance or testing on multiple safety functions.

Additional design measures addressed in IEC 62340 include system integrity, independence from external dates or messages, and assurance of physical separation and environmental robustness. Provisions to ensure system integrity through self-supervision (as required in this standard and IEC 60880) must include determination of a predefined state to invoke on failure detection. This "failed" state must be based on failsafe principles. Requirements regarding avoidance of dependencies address precautions against dependence on external time and provisions for access security (which are referenced from IEC 60880). Finally, other standards are cited for requirements on separation and isolation (IEC 60709), equipment qualification (IEC 60780), and electromagnetic compatibility (IEC 61000-4) [25].

3.2.3.3 Tolerating Postulated Software Faults

It is noted that in accordance with IEC 61513, digital safety systems should be designed to "operate internally without dependence on the demand profile." The software-specific requirements given in IEC 60880 are supplemented by additional requirements in IEC 62340. These requirements, which are consistent with IEC 60880, are intended to "reduce the possibility that assumed latent software faults may be triggered from data which depend on transients of the plant process." In particular, it is required that application and system software be separated such that "the algorithmic processing of plant process data

is entirely performed by the application software." Additionally, execution of system software functions "should not be influenced by any data which directly or indirectly depends on the plant status." To satisfy this requirement, IEC 60880 is cited along with the following design measures: "invariant cyclic processing of the application functions," "invariance of processing load and communication load," and "avoidance of interrupts triggered by process data."

Other software-related coping requirements address tolerance of invalid input signals and spurious signal transients, online identification of invalid or faulty input signals, protection of other safety functions in the presence of single function failure due to invalid input signals, and provision of a safe action in response to multiple CCF or input signal failures. It is cautioned that the signal validation by comparison of redundant information can introduce dependencies between redundancies that must be analyzed for CCF possibilities.

3.2.3.4 Avoiding System Failure Due to Maintenance during Operation

IEC 62340 addresses the prospect that CCF can be induced by maintenance activities during operation. Specifically, it is required that simultaneous activities are limited to "a single redundancy to avoid a resulting failure of more than one of the redundant trains, channels, or subsystems." Additionally, an analysis must confirm that the prospective impact of maintenance activity during power operation cannot induce failure of other nonrelated systems performing safety functions. Finally, it is required that the useful lifetime of components be determined to limit the potential effect of aging degradation and that replacement components be adequately qualified and their compatibility be sufficiently verified to avoid introduction of new failure modes or reduction of system reliability.

3.3 Nonnuclear Industry Guidance on CCF Mitigation

Within high-value, high-integrity, and safety-significant industries, failure avoidance and mitigation approaches are ubiquitously employed to decrease the likelihood of I&C system failure. These nonnuclear, high-failure-consequence industries, which employ similar I&C applications, have almost completely transitioned to digital technology. The guidance from these other industries can serve to identify useful practices for CCF mitigation within digital I&C system architectures.

None of the other high-consequence industries is directly analogous to the nuclear power industry. Both inherent technical and regulatory oversight differences exist between these other industries and the nuclear power domain. For example, flight control systems within the aviation industry typically do not have a readily accessible safe shutdown state, have short-term potential catastrophic control trajectories, and make frequent significant adjustments to control elements. These inherent characteristics make the requirements for probability of failure on demand more stringent for aviation than nuclear power generation. Another technical difference relates to the nature of the safety-critical functions that are characteristic of some nonnuclear industries. In some cases, the systems of concern embody continuous control functions rather than as-needed protection functions. Consequently, the demand profile for these applications is extensive and the actions of the systems are expected, continuously occurring, and actively monitored (both automatically and manually) for correct behavior. In contrast, nuclear safety functions are characterized by a sparse demand profile and actions are rare. While conditions that would initiate safety systems action are monitored, the safety system action is unusual and, for fast-acting events, often unexpected. Additional technical differences arise from dissimilar constraints posed by the unique conditions associated with some nonnuclear industries. For instance, size, weight, and power (SWAP) represent significant constraints in some of the nonnuclear application domains investigated. Not only must the prospective impact on feasibility, cost-effectiveness, and risk burden be considered, but a real potential exists for safety to be compromised if systems are too large, heavy, or consume too much power. Finally, the nuclear power industry has significantly greater regulatory oversight (particularly in terms of prior approval of system changes) than comparable high-failure-consequence industries. Thus,

economic factors such as cost, efficiency, and investment protection may have more impact on design and implementation strategies for the less comprehensively regulated industries. The differences in domain context and the nature of the safety-critical applications must be considered in evaluating the CCF mitigation approaches contained in the guidance from these other industries.

Several industries were investigated as part of this research. Many were found to rely primarily on high-quality processes and rigorous hazard identification and resolution. However, four industries in particular were found to have specific guidance related to CCF mitigation. The application domains that provided the most significant information are the aerospace, aviation, chemical process, and rail transportation industries.

3.3.1 Aerospace Industry

The U.S. government aerospace organization is the National Aeronautics and Space Administration (NASA). NASA performs scientific investigation and exploration of space through manned and robotic missions. With such diverse applications, I&C architectures and requirements for NASA spacecraft vary considerably. However, safety-critical guidance to support manned space operations addresses treatment of CCF vulnerabilities.

NASA requires CCFs to be "considered" and "assessed," but specific CCF mitigation approaches, such as diversity, are not explicitly required. The NASA Safety Manual [26] more specifically addresses redundancy as a means to achieve fault tolerance. The level of protection required is a function of the hazard severity and probability, and may be achieved by a combination of availability, reliability, maintainability (restorability), and redundancy. Use of redundancy to achieve failure tolerance requires specification of acceptable reliability and provision of sufficient redundancy to tolerate two failures or operator errors where loss of life or mission failure could occur and tolerate one failure or operator error (failsafe) where system loss/damage or personal injury could occur. Where there is sufficient time between the occurrence of a failure and the manifestation of its effect, failure tolerance can be achieved through design enabling restoration to safe operation based on ("hot" or "cold") spares, operational procedures, or maintenance. Where there is not sufficient time for recovery, functional redundancy must be provided. Functional redundancy is defined as "situation where a dissimilar device provides safety backup rather than relying on multiple identical devices" [26]. Nevertheless, the use of redundancy to achieve failure tolerance requires verification that any assumption of failure independence is not invalidated by CCFs.

The NASA Software Safety Standard [27] states that nonsafety critical and safety-critical software may reside on the same processor, although design provision must ensure that the safety-critical function cannot be disabled or impaired. Software within a safety-critical system is generally presumed to be safety critical and is treated accordingly. If nonsafety critical software resides in the same system (i.e., on the same processor) with safety-critical software, the partition or isolation method is treated as safety critical, but the isolated nonsafety code is not. This requirement on the treatment of software is particularly important for the incorporation of commercial-off-the-shelf (COTS) software. Software design and code implementation may not compromise any safety controls or processes, cannot create any additional undocumented or unresolved hazards, and must maintain the system in a safe state during all modes of operation. Catastrophic hazards must be able to tolerate two hazard control failures (two-fault tolerant), while critical hazards must be able to tolerate a single hazard control failure (single-fault tolerant) [26,28].

Human-rated systems require an assessment of CCF vulnerabilities and manual override capability. Human-Rating Requirements for Space Systems [29] requires that flight software shall, at a minimum, be tested using a flight-equivalent avionics test-bed operating in real-time. Space systems are required to be designed so that no two failures result in crew or passenger fatality or permanent disability. The space system relies upon operators as a diverse control system by requiring the crew (and ground control) to

have the capability to manually override higher-level software. The space system is also required to provide the capability for autonomous operation of critical functions. The crew (and ground control) can initiate, override, or abort automatic initiation sequences. As a defense against CCFs, use of dissimilar redundancy or backups is required to be assessed. Dissimilar redundancy can be characterized in terms of "additional functional capability (hardware and associated software) to provide at least two [different] means of performing the same task" [30].

3.3.2 Aviation Industry

The civil aviation industry within the United States is regulated by the Federal Aviation Administration (FAA) under the U.S. Department of Transportation (DOT). As part of its regulatory oversight responsibilities, the FAA certifies the airworthiness of aircraft avionics. The Society of Automotive Engineers (SAE) publishes the Aerospace Recommended Practice (ARP) standard ARP 4754, "Certification Considerations for Highly-Integrated or Complex Aircraft Systems" [31]. SAE ARP 4754 addresses certification aspects of highly integrated or complex systems intended for installation on aircraft while accounting for the overall aircraft operating environment and functions. SAE ARP 4754 defines the full engineering life cycle, which includes planning, development, testing, and certification. SAE ARP 4754 also establishes guidelines for assigning Development Assurance Levels (DALs) to a system, its components, and any software based on the most severe failure conditions associated with the corresponding part. These DALs are assigned according to failure conditions classifications (i.e., catastrophic, hazardous/severe major, major, minor, and no safety effect).

The standard SAE ARP 4754 relates to aircraft system development. Additional guidelines for software development and hardware development are provided by Document (DO) 178B, "Software Considerations in Airborne Systems and Equipment Certification" [32], and DO-254, "Design Assurance Guidance for Airborne Electronic Hardware" [33]. These guidelines are published by the Radio Technical Commission for Aeronautics (RTCA).

A safety assessment process is described in SAE ARP 4752 to generate evidence of compliance with airworthiness requirements. The primary processes involve a Functional Hazard Assessment (FHA), Preliminary System Safety Assessment (PSSA), System Safety Assessment (SSA), and Common Cause Analysis (CCA). A CCA is required for systems assigned DALs A (Catastrophic) or B (Hazardous/Severe Major). The CCA begins after applicable separation and isolation requirements are identified to minimize commonalities and interdependencies. The CCA proceeds with Zonal Safety Analysis to identify location-specific challenges to independence, Particular Risks Assessment to common external events or influences of concern, and Common Mode Analysis to confirm assumptions of independence. This latter analysis addresses the potential effects of "design, manufacturing, and maintenance errors and the effects of common component failures" [31]. Categories for common-cause faults are identified in terms of software design error, software coding error, requirements error, repair process error, environmental factors, hardware failure, hardware design error, compiler error, production process error, installation error, operational error, and cascading failures.

As a means to resolve the findings of the safety analyses, SAE ARP 4754 identifies the use of system architectural features such as redundancy, partitioning, or dissimilarity to eliminate or contain the degree to which an item contributes to a specific failure condition. However, SAE ARP 4754 does not use the same definitions of key terms as the nuclear industry. The aviation industry terms of redundancy, partitioning, and dissimilarity are comparable to the nuclear industry concepts of redundancy, isolation, and diversity.

Redundancy is the provision of more than one means for accomplishing a function. For example, redundancy can involve additional separate equipment to perform the same function as a primary piece of equipment. The redundant elements may be parallel or backup, active or passive, and/or of similar or dissimilar designs. SAE ARP 4754 indicates that redundancy is necessary to provide failsafe design

protection from catastrophic failure conditions. SAE ARP 4754 further indicates that redundancy also may be necessary to meet the requirements associated with other severe failure conditions.

SAE ARP 4754 describes partitioning as a "design technique for providing isolation to contain and/or isolate faults and to potentially reduce the effort necessary for the system verification processes." Partitioning is a similar concept to isolation as used in IEEE 603.

The concept of dissimilarity as used in aircraft design is similar to the concept of diversity as used in the nuclear environment. The following excerpt from Sect. 5.4.1 of SAE ARP 4754 indicates that the use of dissimilarity or diversity is encouraged:

"For all but the simplest systems, it is practically impossible to guarantee the correctness and completeness of requirements or the correctness of all necessary assumptions. An architectural strategy incorporating dissimilarity can be a powerful means of reducing the potential for errors in requirements or in design implementation to cause serious effects… ." Additionally, the standard states that "[w]hen dissimilarity is used as a means of design error containment, the degree of credit should be related to the type and scope of design errors shown to be covered by the dissimilarity… . Assuming adequate independence can be shown, dissimilar design implementations of dissimilar functions can provide containment coverage for both implementation and function requirements errors."

SAE ARP 4754, Sect. 5.4.1.2, "Dissimilar, Independent Designs Implementing an Aircraft-Level Function," also includes the following:

"To be considered within this category, there must be substantial differences between the designs in terms of the means of preventing the top level failure condition(s), the methodology by which the designs are created, the technology through which the designs are implemented, and the operations through which the functions are used. Validation of any assumptions of independence is of particular importance in demonstrating compliance… ."

Alternate architectures are also identified in the standard if dissimilar independent designs cannot be achieved. These include backup parallel designs, active-monitor parallel designs, and primary/secondary designs. The final case corresponds to dissimilar designs implementing a function with a primary portion satisfying the highest DAL associated with the most severe conditions and the secondary portion at a DAL that is one level lower than the primary portion.

3.3.3 Chemical Process Industry

The chemical process industry regularly produces, stores, transforms, and consumes highly toxic, explosive, highly flammable, and carcinogenic materials in large quantities. Moreover, it employs physically (e.g., temperature, pressure) and chemically aggressive environments to perform its basic functions. Hazardous material processing takes place in many chemical process plants. These plants share many features with NPPs. Modern chemical plants feature a main control room that presents information about the plant and process status to the operator. Local control loops may also be employed to control particular aspects of the process operation. The primary control functions are performed by the basic process control system (BPCS), while protective functions are provided by separate, high-integrity safety instrumented systems (SISs). An SIS is "composed of sensors, logic solvers, and final control elements whose purpose is to take the process to a safe state when predetermined conditions are violated" [34]. Typical SISs include emergency shutdown systems (ESD or ESS), safety interlock systems, protective logic systems, and safety shutdown systems (SSD). Although SISs traditionally involve physical (e.g., pneumatic and hydraulic) and electrical (e.g., direct wired, electromechanical, and solid-state relay) systems, programmable electronic systems (PESs) are becoming prevalent. Common PES platforms include programmable logic controllers (PLCs), distributed control systems (DCSs), or application-specific stand-alone microcontrollers.

The Occupational Safety and Health Administration (OSHA) under the U.S. Department of Labor enforces health and safety regulations for industrial processes and has regulatory oversight responsibility for workplace safety and worker health. Additionally, the American Institute of Chemical Engineers (AIChE) established the Center for Chemical Process Safety (CCPS) to develop and disseminate voluntary guidance for use in the prevention of chemical accidents.

The 1992 OSHA rule on Process Safety Management of Highly Hazardous Chemicals [35] is the federal regulation for the chemical processing industry most directly comparable to Chapter 7 of the NRC's Standard Review Plan (NUREG-0800). The OSHA rule was developed to prevent and mitigate hazardous releases of the regulated chemicals. The rule was adopted following several catastrophic chemical and petrochemical incidents causing multiple deaths and extensive property damage. In particular, a toxic gas leak at a chemical process plant in Bhopal, India, directly caused the death of over 2000 people. Inadequate safety design and nonfunctioning safety systems due to poor maintenance (a CCF contributor) were identified as contributing factors [36].

The OSHA rule addresses process hazard assessment, risk control measures, and consequence evaluation for system failures, as well as documentation and maintenance requirements. The central OSHA requirements are founded on a process hazard analysis that identifies, evaluates, and specifies the controls for the hazards of a particular process. Of note, the rule requires "that equipment complies with recognized and generally accepted good engineering practices" as opposed to prescriptively specifying particular equipment design and performance requirements. There are no specific requirements regarding consideration of the potential for CCF vulnerability or the use of diversity.

Following the 1985 Bhopal disaster, the AIChE/CCPS developed a series of guidelines providing technical information and recommendations for chemical process safety. In particular, the CCPS Guidelines for Safe Automation of Chemical Processes [37] provides the most extensive guidance on design practices for SISs. Most of the information presented in this section is drawn from the guidance in this document. Additional guidance and standards considered include the CCPS guide, "Guidelines for Safe and Reliable Instrumented Protective Systems" [38]; the Instrument, System, and Automation Society (ISA) standard S84.01-1996, "Application of Safety Instrumented Systems" [39]; and IEC 61511, "Functional Safety: Safety Instrumented Systems for the Process Industry Sector" [40].

Starting with the safe design, defense-in-depth is generally employed for chemical processes through provision of successive independent protection layers (IPLs). As a result, thorough separation between BPCS and SIS layers and among individual systems is encouraged to promote independence. Depending on the risk, each SIS is assigned an integrity level (IL) from among three distinct safety performance levels. Redundancy of components and signal paths, along with the extensive use of active diagnostics, provides degrees of fault tolerance associated with each IL level. Specific techniques to minimize faults include software quality assurance practices, use of watchdog timers, pulsed outputs to detect failures, and fault-tolerant configurations (e.g., triple modular redundant with two-out-of-three voting). Additionally, SIS interlocks are designed to be failsafe.

Nevertheless, recognizing CCF to be a significant concern for control and safety systems (especially those employing PESs), the CCPS recommends diversity in protective systems for hazardous processes. Diversity is identified as referring to "factors that make two components (e.g., devices, subsystems, systems, software systems, communications systems, sensors, or final control elements) different in a way that minimizes common mode fault" [37]. The CCPS further states that diversity "may include the use of different physical methods, technology, manufacturers, installation, maintenance personnel and/or environment" [38].

Identifying the degree of risk in the chemical process, and thereby determining the SIS diversity needs, begins with a detailed process analysis. After process hazards have been identified, process modifications to reduce the overall risks are then considered. Next a basic process control strategy is identified. Process risks are then assessed, through probabilistic risk assessment, by considering accident

likelihood and consequences coupled with predicted safety equipment performance probabilities. A minimum safety performance integrity level is then associated with a particular process based upon the identified risk and the available IPL. Higher risks are associated with higher ILs and increase the required amount of engineering rigor in the process control and safety system design. For the highest integrity level, the CCPS recommends that "Diversity should be considered and used where appropriate" [37].

The CCPS safety evaluation model, relatively speaking, maintains a considerable degree of correspondence with that of the nuclear power industry. The CCPS safety evaluation model employs independent protection layers—roughly in accord with the "echelons of defense" of the nuclear industry. The CCPS endorses separation (lack of direct communication), independence (no common components or collocation), and diversity of each layer of the control and protection system(s). The CCPS also provides guidelines for necessary exchange of information among separate safety channels (e.g., for voting) and for buffered intercommunication to other components. Employing multiple, independent protection layers is also provided as an example of increasing safety system diversity.

Acknowledging the significant functional difference between the process control and safety systems, most of the CCPS diversity recommendations adopt that difference as a basic diversity attribute. Essentially, the chemical industry notes that the functions of the control system and the safety system are different. Consequently significant diversity is thus inherently obtained by having independent safety and control systems.

The CCPS does not provide detailed guidance on how much diversity is required for a particular process risk. In fact, the CCPS specifically places the responsibility for determining the appropriate amount of safety engineering on the plant owners. The minimum number of IPLs required to address a process risk can be derived from the user company's safety policy.

However, the CCPS guidance does provide high-level recommendations on the use of diversity, depending on the IL associated with each SIS. Diversity usage recommendations include the use of different technologies, different manufacturers (or products from different vendors), and different application programming teams. For hardware diversity, different sensors and logic equipment are identified as options. For system software diversity, different controller/logic platforms and smart sensor devices are recommended. For application software diversity, development of different programs is recommended. It is explicitly noted in the guidance that diversity "can cause serious problems when reliability is sacrificed to achieve diversity" [37]. Therefore, diversity is recommended only where reliable components are available.

Finally, the CCPS does provide cautionary guidance about the difficulties in eliminating CCF throughout the system life cycle. The CCPS indicates that CCFs are frequently of human origin with system maintenance, testing, and design being prime common failure sources. The elimination of these vulnerabilities is thus difficult to achieve in SIS. Further, while the CCPS does endorse both passive and active diagnostics of the plant control and safety systems (e.g., internal and external watchdog timers) at the same time the guidance notes that the additional complexity engendered by the diagnostics increases the possibility for system failure from a separate source. Thus, a balance between adequate diagnostic coverage and minimizing system complexity is required.

3.3.4 Rail Transportation Industry

The Federal Railroad Administration (FRA) under the U.S. Department of Transportation promulgates and enforces rail safety regulations as a central element of its mission to oversee domestic rail transportation. In Europe, the European Railway Agency (ERA) was established in 2004 to facilitate an integrated railway system by reinforcing safety and interoperability.

The FRA safety regulations are found in Title 49 of the Code of Federal Regulations. The Signal and Train Compliance Manual is formed by Parts 233–236 of Title 49. Of particular relevance is Subpart H,

"Standards for Processor-Based Signal and Train Control Systems," of Part 236, "Rules, Standards, and Instructions Governing the Installation, Inspection, Maintenance, and Repair of Signal and Train Control Systems, Devices, and Appliances" [41], which establishes regulations addressing the use of microprocessors in signal and train control systems. Other subparts address requirements for interlocking systems, traffic control systems, and automatic train stop, train control, and cab signal systems.

The regulations in 49 CFR 236 Subpart H require the establishment of a Railroad Safety Program Plan (RSPP) based on product safety plans (PSPs). The RSPP must address system requirements and concepts, design for verification and validation (V&V), design for human factors, and configuration management controls. In particular, a safety analysis must be included which describes the critical behavioral characteristics, risk assessment procedures, any safety precedence applied, and the safety assessment process. In addition to containing the aforementioned risk assessment, the PSP must also provide a hazard mitigation analysis and V&V plan as part of a complete description of the safety assessment. Within the regulations, practices developed by the American Railway Engineering and Maintenance of Way Association (AREMA) for the application of vital electronic/software-based equipment are adopted. Coded processors represent a high-integrity implementation approach that contributes to addressing the potential for CCF vulnerabilities. This approach will be described in the subsequent examples of diversity usage.

As indicated previously, the ERA began in 2004 through publication of the European Rail Safety Directive (2004/49/EC),[*] which forms the basis of the European rail safety scheme. However, this directive is at a high level, emphasizing overall system quality and not focused on implementation methodologies. Subsequently, the ERA issued the initial Common Safety Methods and guidance on the development of Common Safety Targets. The Common Safety Methods, developed as recommendations in late 2007, primarily address the use of risk assessment, based on hazard identification and consensus assessment principles, as a means of establishing safety requirements [42]. In April of 2008, the first recommendations on a framework of methods to be used for calculation, assessment, and enforcement of Common Safety Targets were issued. Generation of the first set of Common Safety Targets is anticipated in 2009.

The principal European railway standard that addresses digital safety-critical systems is the CENELEC (European Committee for Electrotechnical Standardization) European Norm (EN) 50128 "Railway applications—Communications, signaling and processing systems—Software for railway control and protection systems" [43], which draws heavily from IEC 61508 [44]. The norm provides guidance on software safety integrity levels, personnel and responsibilities within the software life cycle, life-cycle documentation, requirements specifications, architectures, design and implementation, verification and testing, software/hardware integration, validation, assessment, quality assurance, and maintenance.

Within EN 50128, the guidance on software assessment highly recommends a Common Cause Failure Analysis. The informative Annex B describes methods of CCF Analysis as "general quality control, design reviews, verification and testing by an independent team, and analysis of real incidents with feedback of experience from similar systems" [43]. The norm also contains specific guidance regarding software architectures that addresses means to mitigate CCF. These include defensive programming, safety bag techniques, and diverse programming. Defensive programming techniques include approaches to check for control or data anomalies, such as plausibility checks for data or control flow sequence checking for code execution. The safety bag approach is based on the concept of a safety envelope (or "bag") surrounding the application to ensure only safe actions are authorized (see Annex B of Ref. 43). Safety bag techniques involve an external monitoring application on an independent computer with the application based on a different specification from the safety-critical application. The purpose of the safety bag processor is to confirm that the actions/commands of the safety-critical application are

[*] http://www.era.europa.eu/Document-Register/Pages/Agency-Regulation.aspx

"safe, not necessarily correct, actions" [43]. Given detection of a potentially hazardous state for the safety-critical application, the safety bag processor enforces a safe state. Diverse programming involves N-version programming with arbitration based on either complete agreement or majority voting. For software of the highest safety integrity level (SIL 4), defensive and diverse programming techniques are highly recommended while safety bag techniques are recommended.

4. EXAMPLES OF CCF MIGITATION PRACTICES

4.1 International Nuclear Power Industry Examples of CCF Mitigation

Specific examples of CCF mitigation practices can be found at international NPPs. The examples that are described represent a sampling of evolutionary reactors and modernized plants that employ digital technology extensively. In particular, five of the earliest examples of highly integrated digital I&C systems that have been implemented at new installations were included in the survey. These plants are Chooz, Darlington, Kashiwazaki-Kariwa, Sizewell, Temelín, and Ulchin. An example of extensive modernization for an existing plant (Dukovany) based on digital I&C technology was investigated as well. Finally, two plants currently undergoing licensing and construction were studied to assess recent trends. These plants are Lungmen and Olkiluoto.

4.1.1 Chooz B (France)

The Chooz B Nuclear Plant Unit 1, commissioned in 1996, is the prototype of the standardized N4-class pressurized-water reactors (PWRs) supplied by Framatome (now AREVA NP) [45,46–48]. The microprocessor-based safety system for N4 reactors was jointly developed by Framatome, Electricité de France (EdF), and Schneider Electric/Merlin Gerin (now Data Systems and Solutions—DS&S, a subsidiary of Rolls Royce) and is designated as version two of the Système de protection intégré numérique (Integrated Digital Protection System—SPIN). Diverse compensating functions to back up the safety system for a limited set of PIEs are provided by the Class 2E (i.e., safety-related) ATWS system. Thus, Chooz employs a primary and secondary diverse system architectural approach for CCF mitigation. Additionally, functionally diverse subsystems are employed within the primary safety system.

At the system level for automatic control and protection, the reactor protection system (SPIN) is grouped within the Class 1E CO3 system (COntrôle-COmmande COuer or I&C system for the reactor core), which also contains the nuclear instrumentation system and the control rod drive system. The safety support systems are provided by the Class 1E CS3 system (Contrôle des Systèmes Support de Sauvegarde or safeguards control system) and SCAP system (Système de Contournement à l'AtmosPhère or containment atmospheric control system). General automation is provided by SCAT (Systèmes de Commande des Auxiliaires de Tranche or reactor auxiliary systems control), which is implemented on the Contronic-E platform supplied by Hartmann and Braun (H&B). The Class 2E ATWS functions are incorporated into SCAT.

As noted, the SPIN system is the primary safety system that provides the reactor trip and emergency cooling functions. It consists of four divisions of measurement and calculation equipment and two trains of redundant logic equipment. Figure 4.1 illustrates the configuration of the system. Each division contains multiple processors. In particular, the ensemble of two acquisition units (UA) and five functional units (UF) constitutes the Acquisition and Processing Unit for Protection (UATP). In general, the measurements are quadruple redundant with each sensor set being connected to one of the four divisions. Within each of the divisions, two acquisition unit processors (UA1 and UA2) acquire the signals. Signals are distributed from the acquisition units to five functional unit processors (UF1 through UF5) using two separate, redundant protection data networks (called NERVIA). The functional units perform the required "partial trip" determinations.

Trip data from the UATP of each division are transmitted across redundant, isolated branches of the protection data network for distribution to the two trains. These data from each division are collected and retransmitted on two separate protection networks supporting the Logic Safeguard Unit (ULS) associated with each train. Thus, there are ten protection data networks consisting of eight UATP networks (two per division) feeding into two ULS networks (each collecting data from one set of four divisional UATP networks).

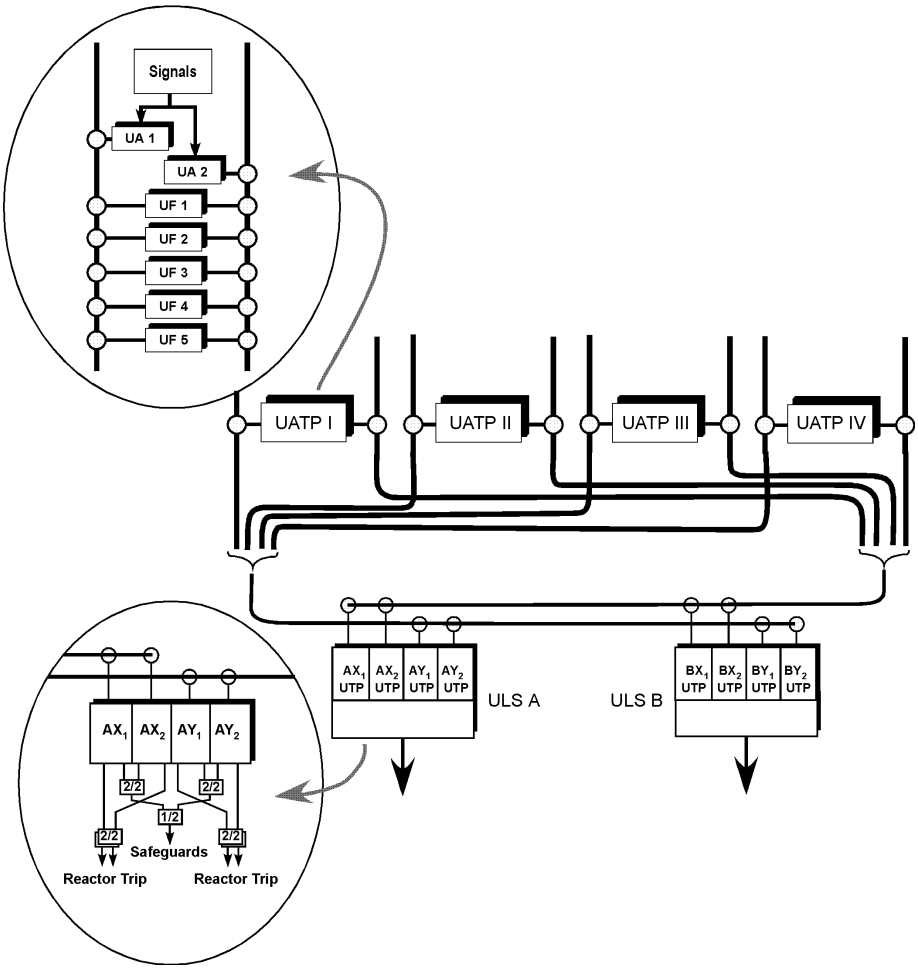

Fig. 4.1. SPIN architecture. (Adapted from Ref. 47.)

Each ULS train (A or B) contains four logic processors (UTPs) that are divided into two pairs (X or Y). Each pair is connected to a different ULS protection network. Based on the trip data from the UATPs, the ULS performs two-out-of-four specific coincidence logic and safety features system level logic. Basically, the SPIN system provides two-out-of-four voting for reactor trip. For emergency core cooling actuation, a logical operation is included that provides an "OR" operation between dual two-out-of-two voters. In both cases, the SPIN design provides protection against a single failure.

As noted above, the N4 design provides ATWS functions in the SCAT system to provide protection against high-frequency events should the SPIN system fail. The ATWS functions are treated as Class 2E, so the system adheres to enhanced quality requirements. The probabilistic safety analysis for the N4 plant showed that loss of secondary feedwater is particularly important in the event of SPIN failure, so an ATWS protection signal based on low-steam-generator level was implemented. Thus, the ATWS scope offers very limited coverage against the full range of PIEs addressed by the safety system. Consequently, the ATWS system constitutes a reduced functionality backup system where the ATWS and safety systems have a different purpose and utilize different functions and logic. Since SCAT, specifically ATWS, and SPIN command some common actuation equipment, priority logic is implemented to arbitrate among these signals, including manual actuation initiation signals. The priority logic is implemented using relays.

Regarding the implementation of the two diverse systems, the DS&S SPIN platform, which is based on the Motorola 68000 microprocessor, serves as the Chooz safety system. The H&B Contronic-E platform used for the ATWS system employs the Intel 80286 microprocessor with an Intel 80287

co-processor. The software for SPIN was written in C, while a proprietary graphical programming language was used for ATWS. It is not known whether this language involved function blocks that may have been written in C or generated C code, so this form of diversity cannot be confirmed.

Finally, additional diversity is provided within the Chooz safety system through functional and signal diversity to provide diverse actuation initiation criteria corresponding to each DBE. The diverse functions are distributed within each division by function among the five UF microprocessors within each divisional UATP, with each unit responsible for one or more protection functions. Consequently, the algorithms and program architecture among these units incorporate some differences.

4.1.2 Darlington (Canada)

The Darlington Nuclear Generating Station is the site of four Canada deuterium-uranium (CANDU) reactors supplied by AECL [49,50]. Units 1 and 2 were commissioned in 1990 as the first CANDU plants to employ "fully" digital I&C systems. There are two diverse digital shutdown systems within each unit at Darlington, with each capable of independently shutting down the reactor in response to detection of any PIE.

The two shutdown systems, designated as Shutdown System Number 1 (SDS1) and Shutdown System Number 2 (SDS2), are functionally independent and physically separate from each other, and from the plant control systems that support normal operation. Specifically, functional independence between the shutdown systems is provided through the use of different means for safety actuation based on diverse physical principles: mechanical (solid) shutoff rods for SDS1 and direct liquid poison injection into the moderator for SDS2. Where feasible, each shutdown system has two diverse trip parameters corresponding to each PIE. Thus, additional functional diversity is provided internally within each system through diverse actuation initiation criteria as well as between shutdown systems through the diverse actuation mechanisms. Diverse trip parameters are available between shutdown systems in a few cases (e.g., low flow and low Δp for loss of flow events). Separate sensors are used for each shutdown system, and where feasible, diverse measurements are employed for the same parameter (e.g., in-core neutron flux).

Figure 4.2 shows the architectural arrangement of computers for SDS1 and SDS2. Each shutdown system contains three physically separate but identical divisions composed of trip computers. The inputs to each trip computer consist of measured parameters and test signals/commands, while the outputs are trip signals and display data. Communication links shown as dotted lines are normally disabled by hardware interlocks. The human-system interfaces and monitoring computers are also shown on the figure, including the Display/Test computers for each division with their associated video display units.

Redundancy in the form of duplication, triplication, and division voting are used to address single failures. Initiation of shutdown action is based on two-out-of-three coincidence among division trip decisions within a shutdown system. SDS1 depends on general coincidence among divisions for trip voting (i.e., two divisions indicating trip without regard to correspondence between the particular actuation initiation criterion), while SDS2 employs local coincidence among divisions for software-based division trip voting (i.e., two divisions indicating trip for the same actuation trip criterion). Final system trip voting is performed with relay logic.

The diversity established between SDS1 and SDS2 begins with the use of computers from different manufacturers as the base platform for each system. The two platforms are based on different computer chip families and have different board layouts. Additionally, development of each system employed separate compilers, computer languages, and development software and was accomplished by different development teams. Specifically, SDS1 uses General Automation (GA) Model 220 machines (based on the GA-16/220 microprocessor) with the application software programmed in FORTRAN and GA assembler. The trip computers for SDS2 are Digital Equipment Corporation (DEC) Programmed Data Processor (PDP) computers based on the LSI-11/23 microprocessor, and the application software is

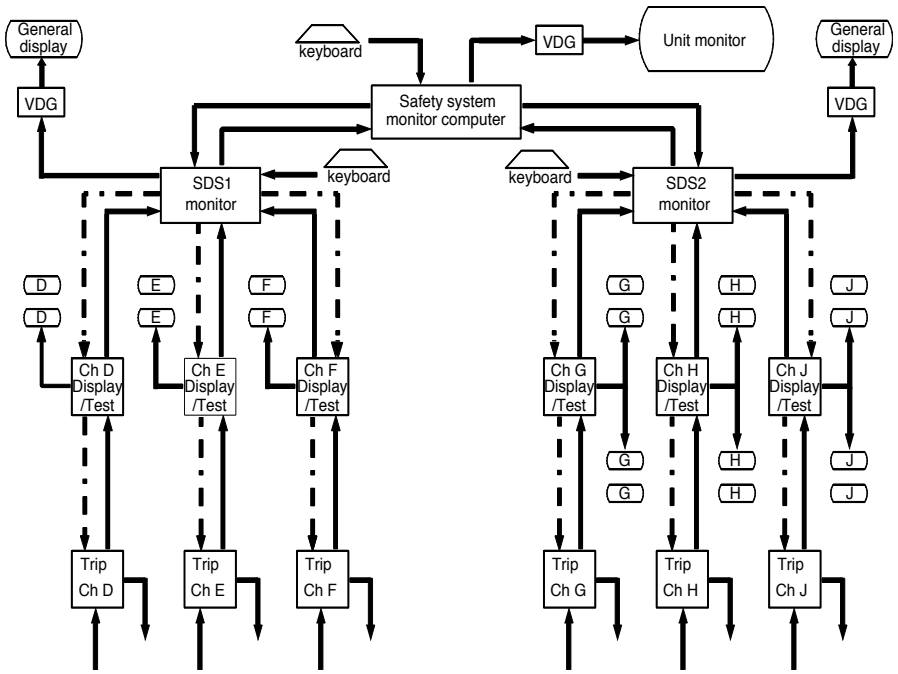

Fig. 4.2. Fully computerized shutdown system.

programmed in Pascal and MACRO assembler. All three divisions within a shutdown system contain identical software (except for division identification), but as noted, the software for each shutdown system is different.

4.1.3 Dukovany (Czech Republic)

The Dukovany Nuclear Power Plant is a four-unit power station based on the design of the Russian water-cooled water-moderated power reactor (VVER) [51]. A modernization program for each unit was initiated in 2002 with phased implementation spanning several outages. In 2005, Unit 3, which began operation in 1986, was the first to have its main upgrade projects completed. The modernization was accomplished using SPINLINE 3, which was developed jointly by Schneider Electric and Framatome (now Rolls Royce DS&S and AREVA NP, respectively). SPINLINE 3 was used to upgrade the RTS, ESFAS, emergency load sequencer, reactor limitation system, and reactor control system. For Dukovany, the digital reactor protection system (DRPS) fulfills the roles of the RTS, ESFAS, and reactor limitation system. Within the DRPS, separate Lines of Protection (LOP) are established based on functionally diverse subsystems employing diverse signals and separate trains of actuation equipment.

The Dukovany plant, like other VVERs, is only able to support instrumentation for three divisions of protection logic. The voting is consequently two out of three. Within each of the three divisions, the SPINLINE 3 design implements the functionally diverse subsystem approach in a manner similar to that accomplished at Chooz using the SPIN system. As previously described, at least two parameters are identified as event indicators associated with each PIE. These diverse actuation initiation criteria are grouped and processed by separate subsystems, LOP A and LOP B (as shown in Fig. 4.3). The digital instrumentation system (DIS) performs the data acquisition and safety comparison processing for each division. The diverse parameters are distributed to separate pairs of processors corresponding to the two LOP. Partial trip results are transmitted across separate NERVIA networks to each division of the DRPS where the two-out-of-three voting is accomplished in two trains corresponding to the two LOP. The

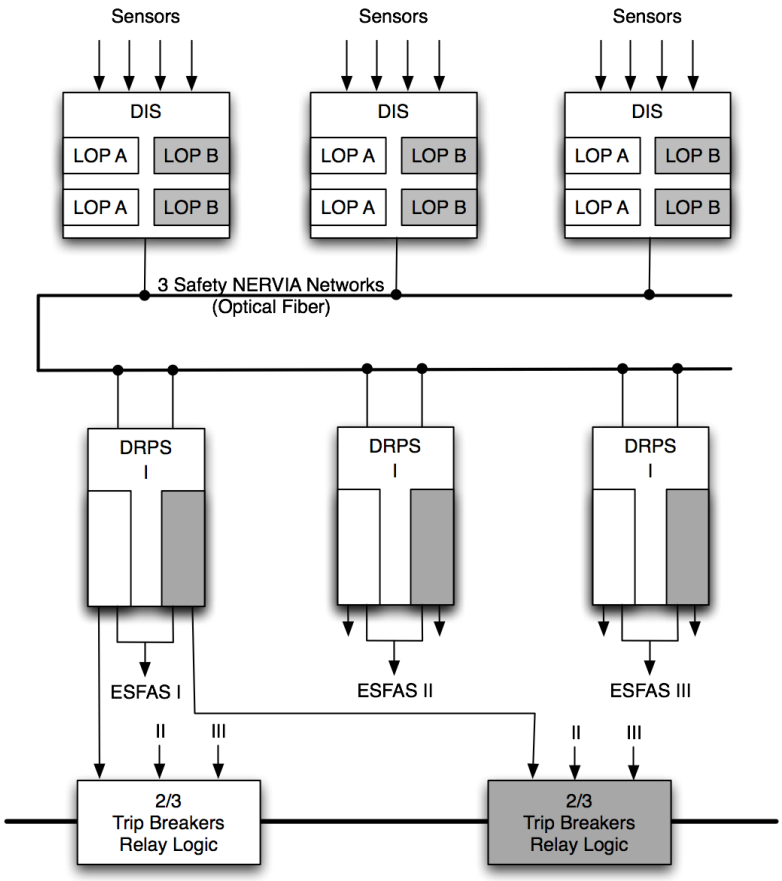

Fig. 4.3. Digital safety system at Dukovany Nuclear Power. (Adapted from Ref. 51.)

voting logic implementation is similar to that described above for the ULS in SPIN at Chooz (see Fig. 4.1). Both LOP trains control the ESFAS actuators associated with the division, while each LOP train drives separate diverse trip breakers based on two-out-of-three general coincidence logic. The diverse trip breakers are supplied by different manufacturers.

While the subsystems utilize separate processing units, the platform and communication network (SPINLINE 3 and NERVIA, respectively) associated with each subsystem are identical. The SPINLINE 3 platform is based on the Motorola 68040 microprocessor. NERVIA is a high-bandwidth token ring network that utilizes broadcast messaging for data transfer. The application software is designed based on formal programming language techniques using a graphical data-flow-oriented development environment called CLARISSE. The CLARISSE system and software development environment (SSDE) provides automatic C code generation for analysis or compilation into binary code for direct implementation.

Since no information on the implementation of ATWS functionality was available, the provision of a diverse backup system could not be confirmed for Dukovany. It was found that the safety (RTS and ESFAS), limitation, and control lines (i.e., echelons) of defense are all implemented on the SPINLINE 3 platform. The principal diversity argument for functionally diverse subsystems arises from the diversification of input profiles and execution of different software applications (i.e., different signal trajectories) such that the diverse subsystems of the RTS and ESFAS should not share any common stimuli other than the initiating event. Cyclic, invariant execution of functions is used to avoid common demand dependencies. The impact of time dependency is addressed as a potential common stimulus by requiring asynchronous operation, static memory and program configuration, no external interrupts, and no operations requiring accumulation or functions of time.

4.1.4 Kashiwazaki-Kariwa 6 and 7 (Japan)

Units 6 and 7 of the Kashiwazaki-Kariwa Nuclear Power Station (KK-6/KK-7) are the first operating advanced boiling-water reactors (ABWRs) [46,52,53]. The units were constructed by Hitachi, Toshiba, and General Electric (GE). GE supplied the turbine/generators for both units, while Hitachi and Toshiba alternated by unit as the lead contractors for either the nuclear steam supply system (NSSS) or the balance-of-plant (BOP) systems. Toshiba supplied the control and safety systems for the KK-6 NSSS, while Hitachi supplied those I&C systems for KK-7. Commercial operation of KK-6 began in 1996, and KK-7 connected to the electric grid in 1997.

The I&C systems for NSSS control and protection throughout either KK-6 or KK-7 are implemented on a common microprocessor-based platform using a similar software development environment (e.g., design methods, implementation tools, symbolic language). The protection and control systems of KK-6 were implemented on Toshiba Microprocessor Aided Power System Control (TOSMAP) platforms, which are based on Intel microprocessor-family central processing units (CPUs), while the KK-7 systems were implemented on Hitachi Integrated Autonomic Control System (HIACS) platforms, which are based on Motorola microprocessor-family CPUs. The application of diversity at KK-6/KK-7 focused on backup capabilities provided by limited ATWS functionality and manual controls.

Figure 4.4 shows an overview schematic of the I&C systems at KK-6/KK-7. Safety functions are implemented in the reactor protection system (RPS) and emergency core cooling system (ECCS). Each safety system consists of four redundant divisions and employs two-out-of-four voting. Anticipated transient without scram mitigation logic drives the automatic Reactor Pump Trip (RPT) and Alternate Rod Injection (ARI) system as an alternate shutdown means using analog circuits. Automatic control for NSSS systems is provided by I&C systems such as the rod control and information system (RC&IS), recirculation flow control system (RFC), feedwater flow control system (FWC), and automatic power regulator (APR). In the figure, communication links correspond to multiplexed connections (thick lines) or hardwired cables (thin lines) where optical multiplexing of field data is performed by remote multiplexing units (RMUs).

In Japan, the application of digital technology in NPPs progressed systematically from auxiliary systems, dedicated control loops, and monitoring systems in the 1980s to nonsafety control systems and then safety systems in the 1990s. The long-term experience gained by the Japanese nuclear power industry from this phased introduction of digital technology is credited through confidence in the efficacy of consensus practices (e.g., design measures and software qualification) [54] to reduce the potential for software CCF vulnerability. In particular, a symbolic language (Problem Oriented Language—POL) is used to provide an intuitive structured representation of the software specifications (interlock block diagrams) that is implemented through graphically driven coding tools. Additionally, simplicity of software structure is promoted through simple logic, cyclical execution, static resource usage, and avoidance of external interrupts. Thus, the Japanese nuclear power industry emphasizes consensus software development practices that are intended to facilitate software verification and validation as a primary means for minimizing the potential for systematic software faults.

In keeping with the analog heritage of NPP I&C architectures, conventional diversity approaches are incorporated in Japanese NPPs. In KK-6/KK-7, diversity across lines of defense (RPS, ECCS, automatic control) results from the different purpose and functional relationships that are the bases of each system. Functional diversity is also provided through diverse means for safety actuation. Specifically, the RPS has three reactor shutdown initiation mechanisms (i.e., two ways to depressurize scram accumulators and a fast actuation mode for electric control rod drive mechanism) and the ECCS has two high-pressure injection systems (i.e., the high-pressure core flood system and the reactor core isolation cooling system) as well as one low-pressure flooding system. An automatic depressurization system is also provided to transition to low pressure should a small break event occur.

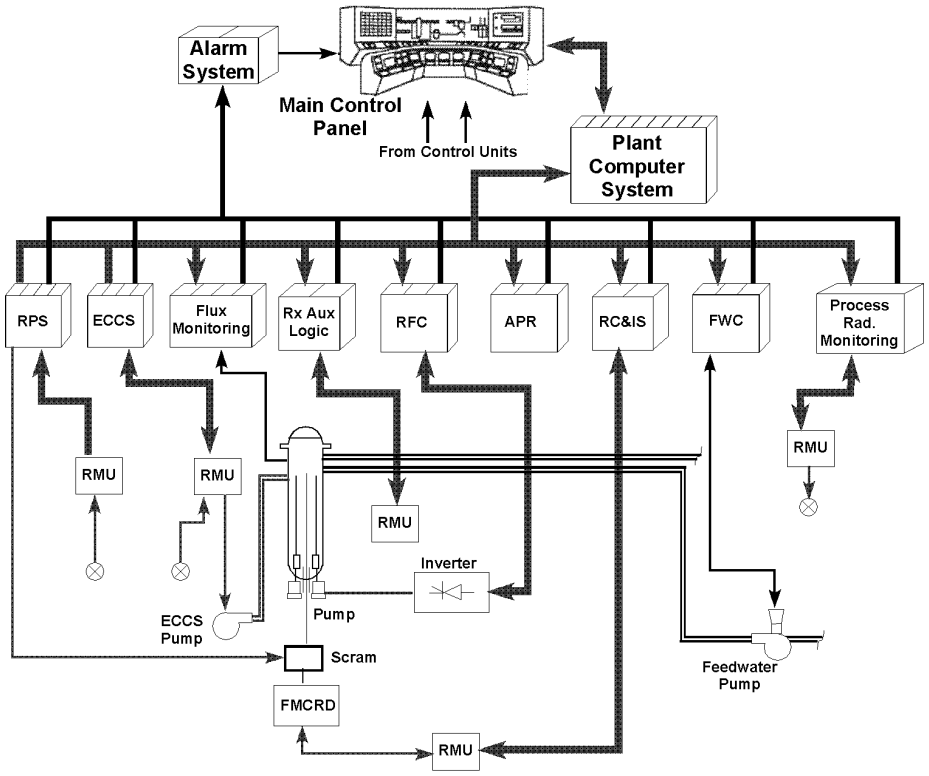

Fig. 4.4. Overview of Kashiwazaki-Kariwa I&C systems. (Adapted from Refs. 46 and 52.)

To cope with any remaining potential for digital CCF vulnerability, manual safety function initiation capabilities are provided in the main control room to serve as a diverse backup. Manual safety action is initiated through hardware switches and hardwired logic circuits, which bypass the digital automatic safety systems. These manual actions include scram, main steam isolation valve actuation, and high-pressure core flood system initiation. Diverse displays of essential parameters are also provided. These essential measurements consist of reactor-pressure-vessel water level, reactor pressure, main steam isolation valve (MSIV) status, reactor water cleanup system (CUW) isolation valve status, reactor core isolation cooling (RCIC) valve status, and high-pressure core flood system status. The manual trip signal de-energizes the power to every divisional trip relay so reactor scram is initiated by a diverse mechanism from that used for automatic trip actuation.

4.1.5 Lungmen (Taiwan)

The Lungmen Nuclear Power Station is a two-unit ABWR plant currently under construction by the by GE for the Taiwan Power Company (Taipower) [55,56]. The control, information, and safety systems are all implemented digitally for Lungmen. Figure 4.5 illustrates the principal control and safety systems. The plant employs six main vendors with several subcontractors to provide the integrated systems. The primary system suppliers are GE, DRS Technologies (formerly Eaton Corporation), GE Industrial Systems (GEIS), Invensys Process Systems, Hitachi, and Mitsubishi Heavy Industries (MHI). The use of multiple vendors and digital platforms results in significant system diversity among the echelons of defense. The systems that constitute these echelons utilize different platforms and perform different functions that provide some level of backup or complementary mitigation for the primary safety functions. Thus, the backup and compensating functions introduced across lines of defense provide significant diversity across the board. In particular, ATWS mitigation logic is provided to serve as the principal backup in the event of a CCF in the safety system. This backup functionality utilizes several diverse systems within the Lungmen I&C architecture.

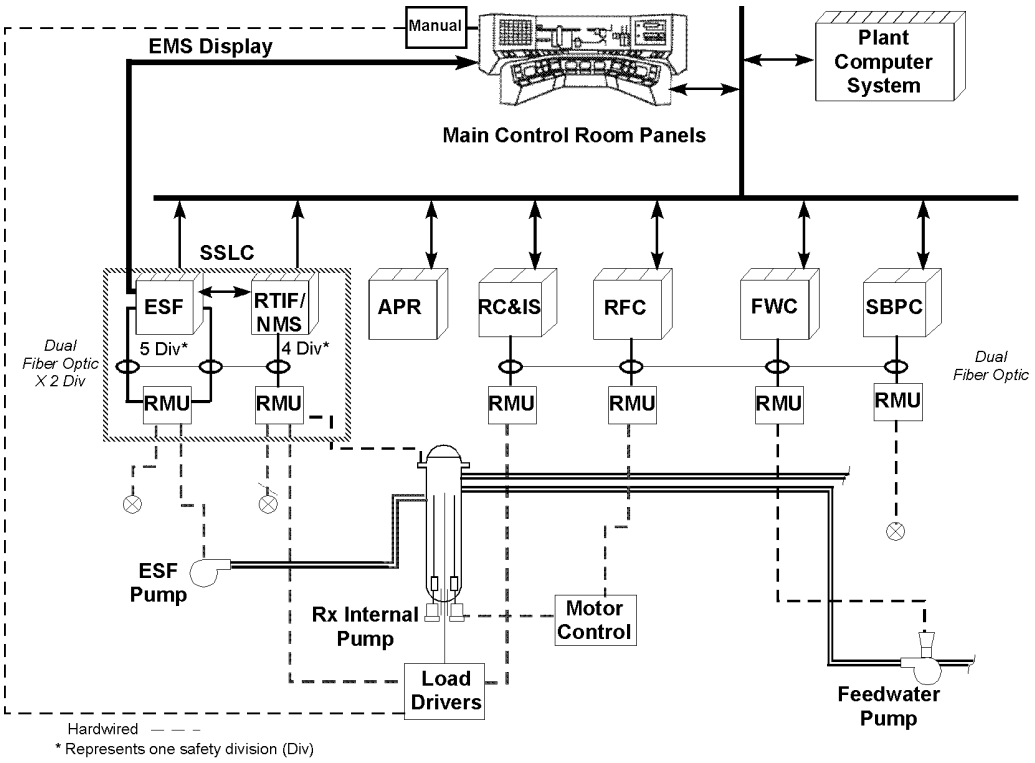

Fig. 4.5. Overall architecture of Lungmen I&C systems. (Adapted from Ref. 55.)

The main Class 1E safety systems for the plant constitute the System Safety Logic Control (SSLC). These systems include the RPS, ESF system, and neutron monitoring system (NMS). These safety systems within the SSLC are supplied primarily by two vendors, GE and DRS Technologies. DRS supplies the ESF system, and GE supplies the RPS, NMS, and associated isolation function systems.

The RPS combines functions for the reactor shutdown via rod scram and the isolation of the reactor system by closing the main steam isolation valves. It is sometimes identified as the Reactor Trip and Isolation Function (RTIF) system.

The ESF system operates the emergency core cooling system (ECCS) and other cooling and post-accident protective functions. The ECCS systems include the High-Pressure Core Flooder System (HPCF), the Automatic Depressurization System (ADS), the Reactor Core Isolation Cooling (RCIC) System, and the low-pressure flooder mode of the Residual Heat Removal System (RHRS). The ECCS provides a series of diverse and redundant systems to provide cooling to the fuel following a design basis accident.

The RPS is implemented using the GE NUMAC platform. It is configured as a quadruple redundant system that consists of distributed processing elements. The main modules that comprise a safety division are RMUs, digital trip modules (DTMs), and trip logic units (TLUs). These modules are configured in a logical pathway from measurement to actuation with downstream interfaces provided via optical communication links. The RMUs communicate multiplexed field data to the DTMs, which perform safety calculations. The partial safety actuation results from the DTMs are communicated to the TLUs in all four divisions. The TLUs perform two-out-of-four-voting to establish divisional trip results.

The ESF system is composed of five divisions with a distributed modular structure similar to that of the RPS. The ESF modules are implemented using the DRS Technologies Programmable Logic Microprocessor System (PLμS) based on the 32 bit PLμS 32 microprocessor. Four divisions constitute the dedicated ESF system for a reactor unit, while the fifth division serves as the unit interface to manage a spare swing set of emergency diesel generators that service both units of the plant. The four primary

divisions communicate within the ESF system, with the RPS, and to ESF actuation devices across the essential multiplexing system (EMS). As is common, the digital safety actuation logic implements two-out-of-four voting.

The EMS provides five separate serial ring networks (i.e., four for the ESF system and one supporting the fifth division). Each division is connected to two EMS rings. The EMS is treated as two divisions consisting of two rings connected to two ESF divisions. The ring network is implemented based on the DRS Technologies performance-enhanced redundant fiber optic replicated memory network (PERFORM.NET). The two EMS divisions are linked to each other and the RPS through redundant communication interface modules (CIMs).

The main process control systems at Lungmen are implemented on fault-tolerant control platforms. In particular, the Feedwater Control (FWC) System, Steam Bypass and Pressure Control (SBPC) System, Recirculation Flow Control (RFC) System, and Automatic Power Regulator (APR) are implemented as triple modular redundant (TMR) controllers using the GEIS Mark VIe platform. This TMR platform is based on the Freescale 8349 (i.e., PowerPC) microprocessor. These systems act to maintain operating conditions in an acceptable range and also provide actuation mechanisms that serve to backup the safety systems.

To enable that backup capability, the Lungmen I&C architecture provides a separate system for ATWS mitigation logic as an alternate means for safe shutdown and cooling of the plant. The ATWS system is primarily a non-Class 1E backup that utilizes several control systems and alternate, diverse shutdown means, such as the Standby Liquid Control System (SLCS), ARI, and Fine Motion Control Rod Drive (FMCRD). However, some of the ATWS logic is implemented in diverse modules within the SSLC cabinets. The system provides diversity in its sensors, hardware, and software. The ATWS system is conceived as a simple, safe recovery system to protect the plant in the event that the safety systems should fail to function due to CCF.

The Lungmen ATWS system consists of redundant logic to initiate diverse automatic actuation of safety or compensating functions. The system contains a simple, reduced set of automatic actuations compared with either the RPS or ESF system. The system also provides analog displays and manual inputs that connect through a minimum set of equipment to the actuated equipment to give the operator a diverse means of manual control. The ATWS controls are available in the main control room and the Remote Shutdown System in the standby control room.

The protective actions provided by the ATWS system include backup scram of the safety rods, liquid poison injection, speed trip or runback of the recirculation pumps, and feedwater runback. The logic for these actions is implemented within the RFC and other systems and on ATWS logic modules in the SSLC. The logic for backup scram is implemented in the TMR RFC system. These actions include two-out-of-three logic for actuation of the safety rods, the FMCRD, and the ARI. Additional logic implemented in the RFC addresses internal pump runback and reactor pump trip. The logic utilizes measurement and status inputs from the FWC, SBPC, SSLC, and manual initiation to provide signal diversity. Mitigation logic to initiate SLCS injection and feedwater runback as well as inhibit ADS actuation is implemented on ATWS logic processor modules in the SSLC cabinets. Specific details on the ATWS logic processor was not addressed in available information resources.

Another aspect of the diversity usage at Lungmen involves the dissimilarity of the safety functions applied in each division of the ESF system. Basically, the software for the safety applications of the ESF is not identical in all divisions. Specifically, the ESF interlock logic is different in each division. The inputs and outputs vary in number and type. Redundant sensors have data messages with unique identifications and time-tags in each division. The intent is to promote differences in the software that may reduce the potential for CCF vulnerabilities that depend on coincident timing or execution. The system is designed so that modules operate asynchronously and thus a common clock or timing signal cannot be a source of CCF. Nevertheless, certain errors depend on the same operation occurring in all

modules at the same or close to the same time. The differences in the division software are believed to reduce the likelihood of such errors from occurring or from occurring simultaneously in all divisions.

4.1.6 Olkiluoto-3 (Finland)

The EPR is an advanced evolutionary PWR supplied by AREVA NP (formerly Framatome) [57,58]. It is currently under construction in Finland as Unit 3 of the Olkiluoto Nuclear Power Station (OL-3), with an expected commissioning in 2016.

The EPR provides an extensive, highly integrated digital I&C architecture based on the AREVA Teleperm XS (TXS) and Siemens Power Plant Automation (SPPA) T2000 (formerly Teleperm XP) platforms. The I&C architecture for OL-3 provides a reduced functionality digital backup for the primary safety system and a "hardwired" backup system (HBS), based on FPGAs, to mitigate the potential for CCF vulnerabilities within the microprocessor-based systems. Thus, OL-3 conforms to a primary and secondary diverse system architectural approach.

Major I&C systems are shown in Fig. 4.6. The I&C architecture includes the Safety Information and Control System (SICS), the Plant Information and Control System (PICS), the Protection System (PS), the Reactor Control, Surveillance and Limitation System (RCSL), the Severe Accidents Automation System (SAAS), the Safety Automation System (SAS), and the Process Automation System (PAS). Priority Actuator Control (PAC) modules are provided as interfaces to shared actuation devices. The reactor trip and ESF functions are contained within the quadruple-redundant PS system.

The PS is implemented on the TXS platform, which is based on the AMD K6-E2 microprocessor. The nonsafety I&C systems are implemented using the SPPA-T2000 platform, which is based on dual SIMATIC S7-400H microprocessors. Of those nonsafety systems, the dual-redundant SAS provides diverse digital backup of the PS safety function for high-frequency PIEs, such as AOOs. The HBS is a quadruple redundant system that provides automatic backup of all reactor trip functions. Although a specific design had not been reported at the time of this investigation, the expectation is that AREVA would develop the diverse FPGA-based HBS. Hardwired manual initiation capabilities are also provided as an additional backup.

In addition to the multiple layers of diverse backups to mitigate the potential impact of CCF vulnerability for the PS, the I&C architecture of OL-3 (and the EPR in general) also employs functionally diverse subsystems within each division of the PS. This strategic diversity usage, as for other international NPPs (e.g., Sizewell, Chooz, Temelín, Dukovany), assigns diverse safety parameters to different subsystem diversity groups, A and B, within each division. A high degree of functional diversity is achieved because diverse signals and some actuated devices are assigned to different subsystems.

The configuration of the PS involves four divisions, each consisting of five acquisition and processing units (APUs) and four actuator logic units (ALUs). Within a division, each APU is assigned to one of the two functionally diverse subsystem groupings. Each APU communicates its safety actuation results to the corresponding subsystem grouping of ALUs in each of the other three divisions. Each subsystem within a division also provides dual ALUs for redundant voting per subsystem using the shared safety actuation signals from across divisions. For the ESF logic, these redundant voters are connected via an "OR" operator. In comparison to a design with single voter, this architecture increases the division reliability by the capability to generate an ESF signal when a single voter fails. The reactor trip logic also contains redundant voters, but these voters are connected with an "AND" operator. This logic provides protection against a spurious reactor trip. The reactor trip signals from the voted subsystem groupings drive different trip breakers. The voted ESF actuation signals from the grouped subsystems are assigned to primary and alternate ESF mechanisms (e.g., emergency feedwater system and safety injection system, which can both provide core cooling) where feasible.

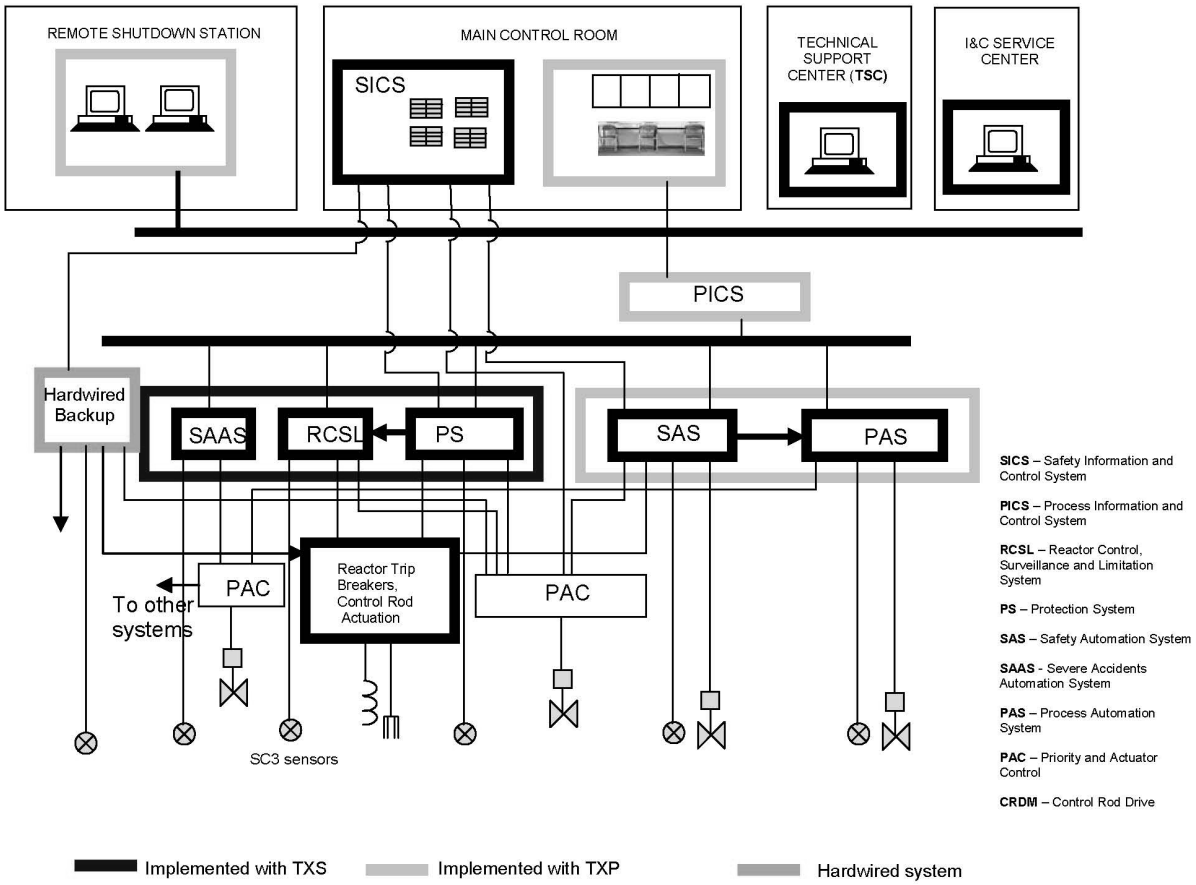

Fig. 4.6. Olkiluoto-3 I&C architecture. (Adapted from Ref. 57.)

The potential for CCF vulnerabilities between the functionally diverse subsystem groupings (i.e., subsystems A and B) is expected to be minimized because the subsystems employ different parameters associated with each PIE based on diverse functional relationships. Essentially, the application software is diversified because the protection functions and parameter/sensor inputs are different. The subsystems do not share any common safety functions. The subsystems are electrically independent and are not connected by any communication links. Nevertheless, common equipment is used for the subsystems and software diversity is limited because the subsystems share the same system software and function block modules. Equipment and logic diversity are achieved in OL-3 by a reduced functionality mitigation capability in the form of the digital backup functions that are implemented as part of the SAS and through the nonsoftware-based backup functions provided by the HBS. The SAS is implemented via a diverse platform using the SPPA-T2000 equipment, while the HBS provides additional, more technologically distinct, diversity through the FPGA-based backup trip functions. The SAS employs a limited set of measurements corresponding to the reduced set of PIEs in its scope. The HBS uses separate measurements of the same parameters for backup trip functions as those used by the PS. The SAS is a nonsafety system with enhanced quality, while the HBS is a safety-related system of a lower safety class than the PS.

In addition to the functional diversity provided by the subsystems A and B within the PS as well as the mitigation arising from the diverse backup systems, there is additional defense in depth provided in the I&C architecture. Specifically, the RCSL system provides control, surveillance, and limitation functions to reduce reactor trips and safety system challenges. Basically, the RCSL supplies soft protection by avoiding safety system challenges by limiting plant conditions. For example, actions such as a power runback are means by which it restores normal operating conditions in response to transients.

Finally, a potential source of CCF vulnerability for protection systems is commonality or sharing of the final actuation device. In the EPR design, a PAC module serves as the interface to ESF actuators and pumps drivers. Its purpose is to manage the use of the actuation resource by arbitrating commands from different sources (e.g., safety, control, and manual commands) while also providing resource protection (i.e., limiting demands to saturated or failed equipment). Dual-use equipment, such as the ESF cooling systems, provides both safety and normal operating functions. Selecting between the input signals requires a final signal arbiter to enforce priority based on safety goals. In OL-3, the PAC is not a simple set of relays (e.g., Chooz or Sizewell) but it is a more complex device providing FPGA-based priority logic and communication interfaces to nonsafety systems. The PAC prioritizes the various sense and command inputs and distributes an output that reflects the plant licensing requirements and operational preferences. In addition, it monitors checkback (or surveillance) signals from the actuators and other devices to protect those resources. The checkback feature limits actuation at the saturation limits. For example, the PAC inhibits demands to a valve to prevent driving it past the full in or full out position. Multi-use actuators are interfaced through a PAC module.

4.1.7 Sizewell (United Kingdom)

The Sizewell B Nuclear Power Station is the only PWR in the United Kingdom (U.K.) [45,59–62]. The Sizewell PWR, supplied by Westinghouse, began commercial service in 1995. The characteristic that distinguished Sizewell from most other PWRs at the time was its extensive use of digital I&C technology. In fact, Sizewell is the first plant at which the Westinghouse Integrated Protection System (IPS) was installed. The IPS architectural approach provides an integrated structure of microprocessor-based subsystems using the Westinghouse Eagle series platform. Features such as the safety functions that are supported, the configuration of safety divisions into quadruple redundancies (designated as guardlines), and the provision of two-out-of-four voting logic are generally the same as those found in conventional analog safety systems at other Westinghouse PWRs. The primary distinction for Sizewell is that it was commissioned with control and safety system implementations based on microprocessor technology and digital data links (e.g., networks or optical fiber links). Thus, the Sizewell B plant serves as a pioneering example of the continuing trend toward more highly integrated digital I&C systems.

Sizewell uses primary and secondary diverse systems to address CCF concerns. However, as described below, functionally diverse subsystems are also employed within the digital primary protection system (PPS). The PPS implements the reactor trip and ESF functionality needed to respond to the full range of DBEs. A diverse secondary protection system (SPS) based on hardwired modules is also provided. As is the case for the PPS, the SPS is arranged in quadruple-redundant guardlines to enable two-out-of-four voting. Both systems are assigned to the highest safety class, and no communication interconnection is permitted between them. British Energy and GEC [General Electric Company plc, now Babcock Nuclear Services (BNS)] developed the SPS while Westinghouse supplied the remainder of the I&C systems for the Sizewell NSSS, including the PPS.

At the time Sizewell was designed and the licensing process was initiated, concern over CCF vulnerability attributed to software was emerging in the international nuclear power industry. As a result, several design measures, including diversity, and various regulatory review approaches were actively discussed to address the potential threat posed by digital CCF. To resolve these concerns, the NII within the U.K. Health and Safety Executive (HSE) employed a special case procedure using a risk-based safety analysis for software-based systems.

A specific determination of the risk-based regulatory assessment was that the SPS must employ thoroughly diverse protection technology to sufficiently reduce the risk contribution associated with a common fault in the system requirements or software design and thereby achieve the required safety goals. To satisfy this requirement, Laddic technology, which had been developed for use in protection systems at British gas reactors, was selected as the basis for the Sizewell SPS. Basically, a SPS guardline is composed of analog trip units for signal processing and Laddic modules for safety actuation voting.

Laddic devices perform logic calculations using pulsed currents through magnetic cores. The underlying physical mechanism for Laddic logic processing is clearly fundamentally diverse from logic processing based on integrated circuit electronics. Additionally, given the long history of operation for these devices in Magnox and advanced gas-cooled reactors (AGRs), Laddic hardware had a well-documented reliability record in nuclear applications in contrast to the very limited experience with digital technology at the time.

Additional, conventional diversity usage was incorporated within the Sizewell protection systems. Specifically, Westinghouse implemented functionally diverse subsystems as part of the digital PPS at Sizewell. Specifically, more than one parameter measured by different types of sensors was identified to cover each PIE. Two alternate groupings of these actuation initiation criteria were assigned to separate subsystems, each of which consists of dedicated computing resources and input/output electronics.

In each of the four guardlines (i.e., divisions), two sets of functionally diverse subsystems were established, with one set corresponding to the two diverse groupings of termination functions (i.e., reactor trip) and the other set providing the two diverse groupings of mitigation functions (i.e., ESF). Keeping termination and mitigation functions separate is intended to ensure that the echelons of defense remain distinct. Figure 4.7 illustrates the separation of functionally diverse subsystems for the reactor trip and ESF within one guardline.

Correspondingly, diverse sensors are provided in the plant design to enable the functional diversity. Additional signal diversity is provided between the PPS and SPS with the selection of sensors from different vendors for each system. Similarly, different vendors were used to supply the reactor trip breakers associated with each protection system. In addition to the eight breakers that are configured in pairs to give two-out-of-four general coincidence logic for reactor trip, the SPS can also remove power from the rod control system bus as a backup means of tripping the reactor.

Other features of interest for the Sizewell I&C architecture include command prioritization, application of failsafe principles, and digital platform differences for protection and control. Sizewell contains some safety actuation equipment (e.g., valves) that can receive control commands from the PPS, the SPS, and the High-Integrity Control System (HICS). As a result, relay-based “priority” interfaces to safety components are employed to arbitrate among commands that originate in the different systems. The logic is based on achieving a safe state in the presence of conflicts.

To better cope with component failures at the system level, the Laddic logic modules can be configured to fail to a preferred state on loss of power. Thus, a failsafe design was implemented for the SPS in which a known safe “failed” state is established by design. As is common for digital safety systems, the PPS employs watchdog timers and self-diagnostics to detect faulted states and enforce a known state as the fault recovery action. Determining the efficacy of this digital failsafe solution depends on the confidence that can be achieved through a systematic assessment of whether the self-diagnostics are comprehensive and without faults of their own.

To promote a failsafe reactor trip interface, a dynamic trip bus was developed to provide dynamic logic units corresponding to each trip parameter. The bus is designed to fail to a safe state if a continuous stimulus is removed due to failure (i.e., the breakers trip unless they remain actively energized).

The PPS for Sizewell is implemented on the Westinghouse Eagle 2000 platform, while the HICS is implemented on IPS and Integrated Control System (ICS) hardware. HICS provides automatic control for the NSSS, manual control of safety components, and data management for safety displays. The PPS is based on the Intel 80286 microprocessor, while the HICS CPUs are Intel 80386 microprocessors. Balance-of-plant control is implemented using the Westinghouse Distributed Processing Family (WDPF) platform, which also is based on the Intel 80286 microprocessor. The hardware architecture for each

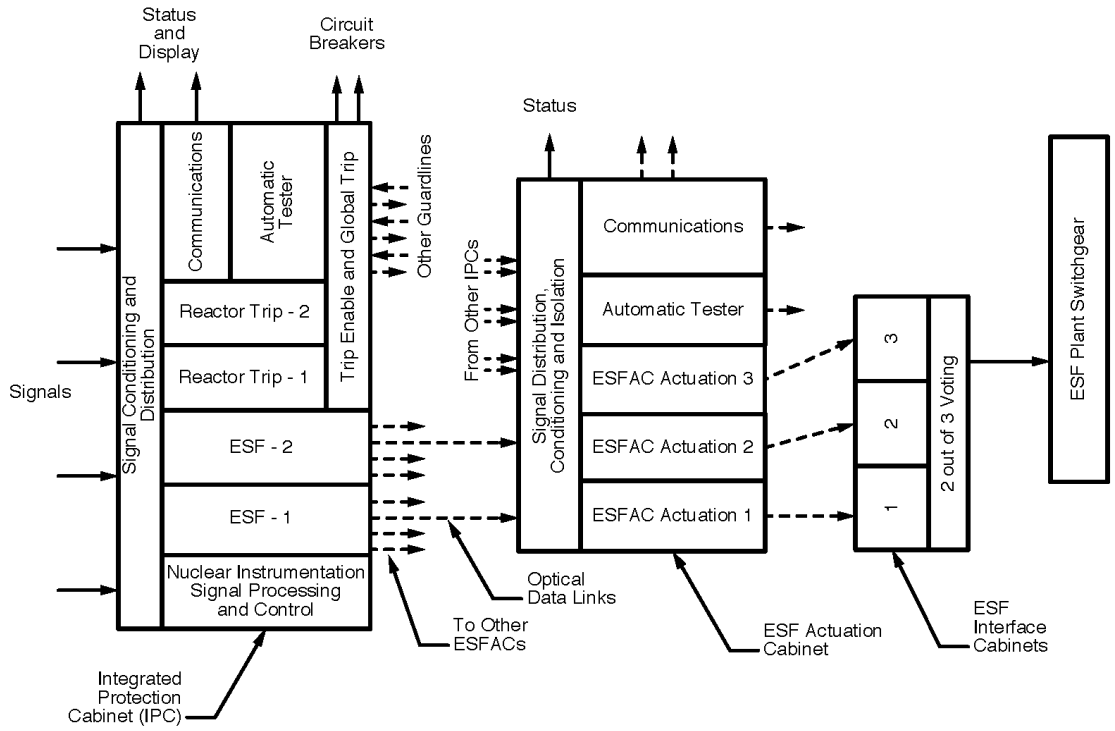

Fig. 4.7. Functionally diverse subsystems for Sizewell PPS. (Adapted from Ref. 60.)

computer subsystem uses the Multibus I internal data bus. The PPS application software was primarily implemented using a high-level structured program language. Use of assembly language was avoided except where required by timing or hardware interface constraints. For the Sizewell PPS, PL/M-86 was the software language employed.

4.1.8 Temelín (Czech Republic)

The Temelín Nuclear Power Plant is a two-unit Russian-designed VVER generating station [45,63–65]. Construction based on the VVER-1000/320 design began in 1982 but was suspended at the end of the decade. Following resumption of construction, a modernization program was initiated to replace the original I&C systems with digital technology. The modernized Unit 1 was commissioned in 2002.

The Sizewell protection system design was adopted as a reference, and the Westinghouse IPS was chosen as the basis for the modernization of the primary reactor protection system (PRPS) at Temelín. The Temelín architectural design adhered to the Sizewell example of providing a diverse protection system. However, instead of a secondary system based on the fundamentally diverse Laddic technology, it was decided to use a microprocessor-based system based on a different platform for the Temelín diverse protection system (DPS). The Westinghouse Ovation digital control modules were selected as the platform for the Temelín DPS. The principal requirement driving the incorporation of a diverse system is that the overall plant safety system must be capable of mitigating an event concurrent with a postulated CCF in either PRPS or DPS, but not both simultaneously.

The primary and secondary diverse protection systems at Temelín are essentially equivalent in safety classification. The PRPS is fully Class 1E and the DPS consists of Class 1E and dedicated equipment. The Temelín DPS assumes the same role as the SPS at Sizewell by serving as a backup safety system for AOOs that are estimated to occur with frequency greater than 10^{-3} events per year. Other than the use of digital technology for the DPS, the primary difference between the I&C system architecture at Temelín

and that at Sizewell is the constraint of three rather than four divisional sensor sets to conform to the original Russian-designed configuration of the VVER I&C systems. Thus, both the PRPS and DPS are implemented as triple redundant systems, and each employs two-out-of-three voting logic for actuation. An additional feature of the I&C system at Temelín is the availability of an additional line of defense through the presence of a separate reactor limitation system, which was also modernized.

As noted, the PRPS is divided into three identical, redundant divisions. Each division communicates its partial trip status to the other divisions for two-out-of-three specific coincidence voting by the microprocessor systems. Subsequent general coincidence voting logic is implemented at the circuit breakers, which are configured into three trains of actuation logic. The PRPS is implemented using the Westinghouse Eagle 2000 platform. As with Sizewell, separate functionally diverse subsystems based on alternate actuation initiation criteria (e.g., parametric diversity arising from signal diversity) are provided within each division. Each subsystem incorporates a "host" (or main) processor and a number of supporting processors for communication, input/output, and auxiliary processes. The Eagle processors are implemented using Intel 80486 microprocessors and supporting integrated circuits. The PRPS application software is written in a combination of PL/M 86 and ASM86 assembler.

The DPS provides a secondary automatic means to shut down and cool the plant should the PRPS fail to take appropriate action in response to a reduced set of events (i.e., high-frequency PIEs). The system also uses two levels of two-out-of-three voting (by the microprocessors and relays). In addition, a second set of breakers is provided for the DPS. These breakers are separate from the breakers used by the PRPS and are supplied by a different vendor. As stated above, the three divisions for the DPS are implemented on Ovation equipment, which is based on Motorola 68000 microprocessors. The DPS application software is written in Ada.

The Ovation platform provides a compact design in which the processor module, as well as the I/O modules, resides on the same VME (VERSAbus-E) bus. Thus, the functionally diverse subsystems within the DPS are not as distinctly separate as for the PRPS using the IPS/Eagle platform.

Other differences between the Eagle and Ovation platforms include different bus architectures (Multibus vs VME, respectively), different network communication technology (proprietary token bus vs reflective memory bus), and different I/O handling (proprietary vs VME-based). Finally, different development teams, development processes, development platforms, and tools were used for each system while different V&V teams were established as well.

The integration of the primary and diverse safety systems at the actuated device level for Temelín required a more complicated priority logic module than the relay-based logic at Sizewell. While the presence of multiple systems (PRPS, DPS, limitation, control, and manual initiation) issuing commands that must be arbitrated has an impact, the previously identified requirement, in which either safety system must compensate for loss of the other due to CCF, drives the need for a robust prioritization capability. Thus, Westinghouse developed nonprogrammable logic (NPL) equipment to implement command priority logic for safety valves and pumps that are affected by multiple systems. Additionally, a portion of the diesel generator sequencing logic is also implemented in NPL equipment. The equipment performing prioritization of safety commands is qualified as Class 1E. Nevertheless, the priority module is a common point at which both the primary and secondary diverse protection systems connect to the final actuated device so the potential for CCF vulnerability must be considered. Consequently, the NPL design is intended to provide a very simple, highly reliable component that is more fully testable than a software-based module.

4.1.9 Ulchin (Korea)

The Ulchin Nuclear Power Plant is a six-unit power station [66,67]. Units 5 and 6 are based on the Korea Standard Nuclear Power Plant (KSNP) design and were commissioned in 2004 and 2005, respectively. For these units, the main I&C systems are implemented on digital computer-based platforms. Figure 4.8 shows the schematic configuration of I&C systems at Ulchin 5&6.

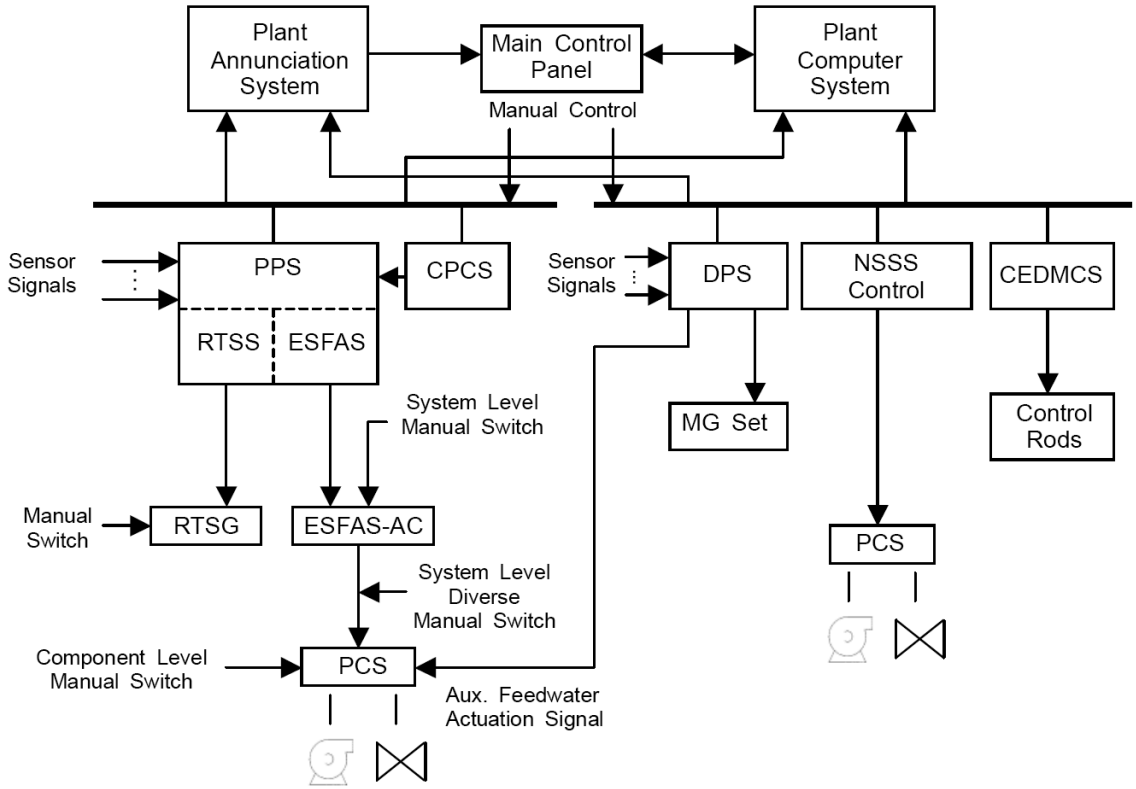

Fig. 4.8. Overview of I&C systems at Ulchin 5&6.

The safety system at Ulchin is composed of the Plant Protection System (PPS), Core Protection Calculation System (CPCS), Engineered Safety Feature Actuation System-Auxiliary Cabinet (ESFAS-AC), Plant Control System (PCS), and Process Instrumentation Cabinet (PI). The nonsafety control system consists of the NSSS control system, which includes the Reactor Regulating System (RRS), Feedwater Control System (FWCS). Steam Bypass Control System (SBCS), Control Element Drive Mechanism Control System (CEDMCS), and Pressurizer Pressure/Level Control System (PPCS/PLCS). The information and annunciation systems include the Plant Computer System, Plant Annunciation System and Critical Function Monitoring System (CFMS). A Diverse Protection System (DPS) is installed to mitigate the consequence of ATWS events in the presence of a potential CCF of the PPS.

The PPS is comprised of four redundant channels that perform the necessary bistable, coincidence, initiation logic and associated maintenance/test functions. Four redundant channels are provided to satisfy single-failure criteria and improve plant availability. The Bistable Processor in each PPS channel receives process sensor analog inputs, discrete and analog signals from the ex-core detector systems, and discrete signals from the CPCS to perform the bistable trip functions. A Reactor Trip or ESFAS initiation signal is generated whenever two-out-of-four redundant bistable trip conditions are sensed in the Local Coincidence Logic (LCL) processor for a particular function. The PPS produces discrete output signals from each channel including trip signals used for the Reactor Trip Switchgear System (RTSS) and actuation signals for each ESF, which are used for initiation of ESFAS.

The ESFAS-AC consists of two independent and redundant trains of equipment housed in separate auxiliary cabinets. The system-level ESFAS initiation signals are received from PPS, and the ESFAS-AC performs the selective two-out-of-four actuation logic. Based on the result of this logic, ESF component level initiation signals are distributed to the PCS.

The DPS augments the PPS to address the requirements for reduction of risk from an ATWS event, as required by regulation. The DPS utilizes independent and diverse logic to initiate reactor trip and auxiliary feedwater actuation. The DPS is a two-channel control-grade system that uses a two-out-of-two logic to initiate a reactor trip when pressurizer pressure exceeds a predetermined value, or to initiate auxiliary feedwater actuation when a steam generator level drops to a predetermined level.

In Ulchin 5&6, the PPS and ESFAS-AC configurations are based on the Advant Controller 160 (AC160) programmable logic controller (PLC), which was supplied by ASEA Brown Boveri–Combustion Engineering (ABB-CE) [now ABB Group]. The CPCS and the PCS vendors were Concurrent Computer and Doosan HF Controls (HFC), respectively. The nonsafety control systems are implemented on digital processors, such as an OMRON PLC or a Foxboro SPEC 200 Micro controller. The DPS configuration is based on a Modicon PLC, which is now supplied by Schneider Electric. The use of multiple vendors and digital platforms promotes system diversity among the echelons of defense.

For the KSNP, there are four echelons of defense. The echelons are the control systems, the reactor trip system, the ESFAS, and the monitoring and indication system. For Ulchin 5&6, the reactor trip system and ESFAS share the same digital processors at the system level. Therefore, any disabling of the digital PPS and ESFAS-AC is assumed to fail all of their output signals in a credible manner. However, the individual component actuation logic for ESF functions is implemented at Ulchin using a different digital processor from the system-level ESFAS processor. This design, based on different processors between system and component levels, enables the component level control for ESF to continue even if the digital PPS and ESFAS-AC functions are disabled due to CCF.

From the diversity point of view, all critical safety functions at Ulchin (e.g., reactivity control, inventory control, and heat removal) can be controlled by both the control system and the protection system. These systems are functionally diverse, as are the fluid/mechanical systems they control. In addition, Ulchin employs both hardware and software diversity between the control and protection I&C systems to minimize the potential for CCF vulnerability. Specifically, the protection system is based on the AC160 microprocessor (i.e., Motorola CPU), the DPS uses the Modicon PLC (i.e., Intel CPU), and other control systems employ the OMLON PLC (i.e., a vendor-specific CPU). Hardwired manual actuation measures for reactor trip and ESF system/component level actuation are also provided. These hardwired manual controls are connected directly to the reactor trip switchgear, digital ESFAS-AC cabinet output, or individual ESF component input. Therefore, the DPS and the hardwired manual control features are available as a means to cope with a postulated CCF that could disable the digital PPS and ESFAS-AC.

4.2 Nonnuclear Industry Examples of CCF Mitigation

Several nonnuclear industries were investigated through this research. Many were found to rely primarily on high-quality processes and rigorous hazard identification and resolution. However, a few key safety-critical industries were found to provide clear examples of CCF mitigation approaches. The application domains that provided the most significant information are the aerospace, aviation, and rail transportation industries. The findings from these industries are presented in this section.

4.2.1 Aerospace Industry

The manned space operations of NASA provide the most prominent examples of safety-critical applications for the aerospace industry. In particular, the I&C architectures of the Space Shuttle and the International Space Station (ISS) provide relevant examples of mitigation approaches for CCF vulnerabilities.

Beyond the adherence to rigorous quality assurance practices, redundancy, fault tolerance, and backup use of human operators are NASA's primary means for achieving highly reliable systems. Mission control and the ISS use a "law of large numbers" type approach; if one system/computer fails,

there are still many computers available for control. The ISS and Space Shuttle use reduced functionality backup systems as a means for improving the probability of mission success in the event of primary software failure. The command and control architecture for manned missions uses commercially available software and hardware. “Fault protection” software routines provide the ability to recover from failures.

The FCS of the Space Shuttle or Space Transportation System (STS) and the station command and data-handling (CDH) system of the ISS provide prominent examples of safety-critical I&C applications for human-rated space missions. These systems provide the most relevant cases of diversity usage by NASA.

4.2.1.1 Space Shuttle

Prior to the Space Shuttle, manned spacecraft computers were programmed at the machine level using assembly language. The delays and expense of the Apollo software development, along with the realization that the Shuttle software would be many times as complex, led NASA to encourage the development of a language that would be optimal for real-time computing. The result was HAL/S (or High-Order Assembly Language/Shuttle). Using the HAL/S language, IBM developed the Primary Avionics Software System (PASS), which is the principal software used to operate the Space Shuttle during a mission. The PASS software is priority-interrupt-driven, or asynchronous—it performs computations on demand and in strict observance to a predefined order of importance [68].

PASS is a quadruple redundant avionics system that is implemented on IBM AP-101S general purpose computers (GPCs). For the first generation avionics system, a fifth GPC was provided on board the Shuttle as a spare. The spare GPC is no longer flown. The functional design for PASS is based on fail operational/failsafe principles. The four GPCs are synchronized at every process initiation and each subsequent input and output (I/O) action. All vital sensors are quadruple redundant as well but the input data for each GPC is equalized using median selection with threshold monitoring. The operational approach is to require agreement among the output of all four active PASS computers. A detected disagreement would result in the dissenting GPC being voted out of the set, with the action being annunciated. When significant degradation occurs, the crew takes manual action (e.g., engages the backup flight system). As previously noted, the application code was implemented using HAL/S. The priority-driven operating system (OS) was written in assembly language. For the Space Shuttle program, NASA used an independent verification and validation (IV&V) team to enhance its software assurance [69,70].

To protect against the prospect of a latent programming error in the PASS software that could render the Space Shuttle uncontrollable during a critical flight phase, NASA contracted with Rockwell and Intermetrics to develop a backup flight system (BFS). This system has its own set of requirements based on reduced functionality flight control laws. In addition, programmers could not reuse any of the code developed for PASS. Nevertheless, like IBM, Rockwell elected to use HAL/S as the programming language. A cyclical time-slice OS was developed for the BFS [68]. The BFS is implemented on a fifth IBM AP-101S computer. The BFS also contributes to the output comparison among the PASS computers. It also serves as the reduced functionality backup during critical flight stages should failure of the PASS be detected [69,70].

The philosophy taken for the BFS was to develop a very simple and straightforward software program and then exhaustively test it. The result was a program that contained only 12,000 words of executable instructions, including the ground checkout and built-in test for the computer. The actual flight control portion of the software consisted of approximately 6,000 words. The remainder of the code was for the systems management functions [71,72].

Two CCFs of the digital I&C have been identified in operation for the Space Shuttle. During a mission, solder shorted out some CPU boards, causing two control computers to crash. Another CCF mode was discovered during simulator testing when crewmembers discovered that all four control computers locked up when executing an abort sequence. The cause was a counter that did not reset during

interrupts. This fault in turn caused the code to encounter values outside the expected range for some variables that resulted in an erroneous branch to an untested execution path [73].

4.2.1.2 International Space Station

The ISS is a cooperative endeavor among NASA, the Russian Space Agency (RSA), the Canadian Space Agency (CSA), the National Space Development Agency of Japan (NASDA), and the European Space Agency (ESA). Over 100 separate computers provide data collection, processing, communication, and control functions for the ISS. The primary station management system is the Command and Data Handling (CDH) system. Boeing provided the CDH system as the prime contractor/supplier to NASA.

The function of the CDH system is to provide command and control of the ISS. The CDH system is implemented in a three-tiered architecture of 25 computers, which are based on Intel 80386SX CPUs. These computers are interconnected by data buses and are accessed by the ISS crew via IBM Thinkpad 760XD laptops, as known as the Portable Computer System (PCS) [74].

Figure 4.9 illustrated the hierarchy of the CDH system. Tier 1 (or the control tier) consists of the Command and Control (C&C) computers, which serve as the ISS station-wide control system and interface access point for the ISS crew through the PCS. Tier 2 (or the local tier) consists of subsystem level functions for the Electrical Power System (EPS), Guidance, Navigation, and Control (GNC) system, Environmental Control and Life Support System (ECLSS), and Thermal Control System (TCS). The purpose of the Tier 2 computers is to execute system-specific application software. Tier 3 (or the user tier) computers provide the direct interface for sense and control components (i.e., the sensors and effectors/actuators) [74].

As indicated in the figure, information flow involves command proceeding down the hierarchy and telemetry (or data) proceeding up the hierarchy. Thus, the station level (Tier 1) C&C computers initiate a command, which proceeds through the subsystem level (Tier 2) to the equipment level (Tier 3) for actuation. As noted, the ISS crew interfaces the system at the control tier. Direct interaction with the lower tiers is not provided. Data queries and commands must proceed through the hierarchy [75].

The three Tier 1 computers are configured to be triple redundant to provide two-fault tolerance. The redundancy is implemented such that one computer is operational, another is a “warm” backup, and the third is powered off in “cold” standby. There are five pairs of dual redundant Tier 2 computers, with the second computer in each pair powered off for cold standby (except for the GNC pair in which a “warm” backup is provided) to provide single-fault tolerance. The twelve Tier 3 computers are not generally implemented in a redundant configuration although some software redundancy is provided through duplicate functions or component interfaces [74].

Software for the CDH computers is written in Ada and includes Caution and Warning (C&W) capabilities at each level. The primary functionality implemented in each computer can be characterized as telemetry, commands, time synchronization, and fault detection, isolation, and recovery (FDIR). The FDIR functions address the computer and the data bus [75].

On April 25, 2001, an independent computer failure coupled with a common-cause failure of the other two first-level control (Tier 1) computers resulted in the failure of all three C&C computers on board the ISS [76]. More specifically, the three Tier 1 computers failed a few days after a new software package was installed. The Tier 2 and 3 computers continued operating to keep many basic functions, such as the primary life support systems, in operation. The Tier 2 computers activated a reduced functionality failsafe mode that triggered backup functions and issued a reboot demand to the Tier 1 computers. Subsequently, one of the disabled computers came back on line. Analysts uncovered an error in the software load that was believed to be the source of the problem [77].

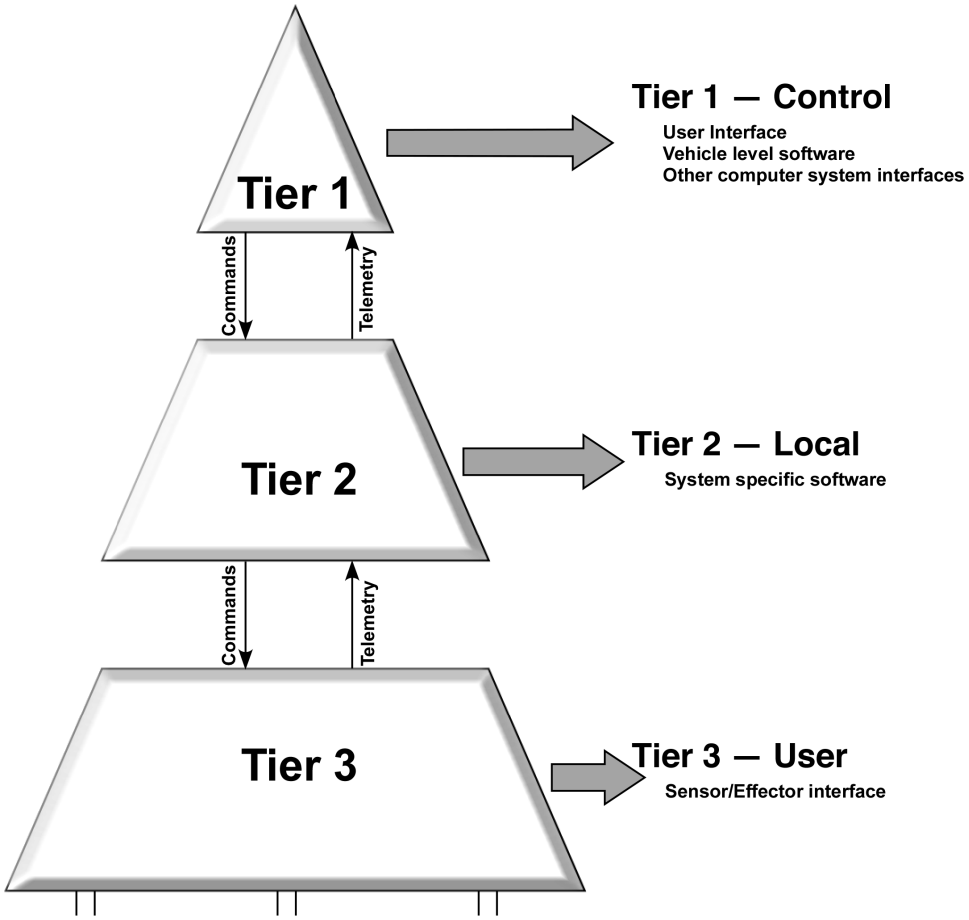

Fig. 4.9. Three-tiered architecture for the CDH system on the ISS. (Adapted from Ref. 74.)

4.2.2 Aviation Industry

Aircraft avionics includes all systems that enable the aircraft to fly safely or have direct control over the aircraft (i.e., high-integrity or safety-critical systems), as well as equipment that supports those systems. The FCS represents the one of most significant high-integrity avionics systems since it provides command and control for the primary flight control surfaces of the aircraft and its proper functioning is essential for commercial airliner flight. It is essential that the FCS be designed so that it avoids systematic faults and tolerates single failures. Consequently, the FCS provides the most relevant examples of CCF mitigation approaches within the aviation industry.

Aircraft manufacturers Airbus Industrie and Boeing provide the most extensive examples for digital fly-by-wire (FBW) FCSs that have been developed for the commercial aviation industry. Airbus A320 serves as one of the earliest implementations and is included in this survey. Successor Airbus flight controllers and the Boeing 777 (B-777) FCS are also presented to capture the evolution of diversity usage in modern FBW systems.

4.2.2.1 Airbus A320

The A320, which was certified in 1988, represents a pioneering use of digital FBW FCSs in commercial aircraft [78]. The overall FCS is composed of diverse redundant primary and secondary control systems. The primary FCS is the elevator and aileron computer (ELAC) while the secondary FCS is the spoiler and elevator computer (SEC). The ELAC and SEC are physically and electrically separated with their own redundant sensors, communication (e.g., data/command links), and power supplies.

Each FCS consists of a self-checking pair based on two channels composed of a control computer and a separate monitor computer. These paired computers form redundant modules within each system, with the ELAC being duplicated and the SEC being triplicated. The redundant modules control redundant actuators. While one module is active, the other module is in standby mode and the redundant actuators are not active. Figure 4.10 illustrates the general architecture employed for the A320.

The pairing of control and monitor computers for the ELAC and SEC systems results in four functionally diverse implementations [79]. The control computers in each system supply flight commands based on normal laws for controlling the assigned actuators. Functional diversity arises because the control elements for the primary and secondary FCS are different. The SEC also provides a reduced functionality backup to the ELAC based on alternate flight control laws. Additionally, manual control based on direct control laws is provided through direct electrical linkage and is backed up mechanically as well.

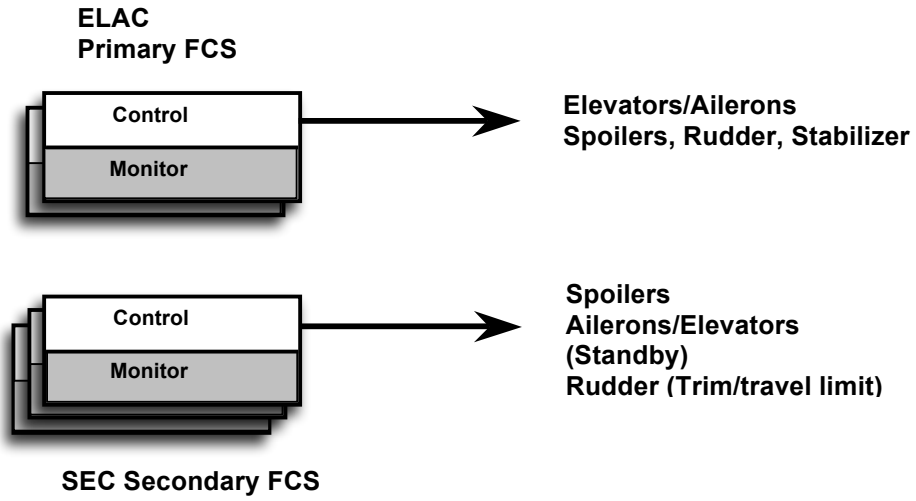

Fig. 4.10. Airbus A320 architecture. (Adapted from Ref. 80.)

The monitor computers implement similar functionality to the control computers, with each derived from the same functional requirements, to support comparison against the control computer outputs. This approach allows controller performance to be monitored for failures. When a failure of a control computer is detected by a monitor computer, primary control is transferred to a redundant module or, if unavailable, to a secondary FCS module. Diversity between the monitor and control computer applications within a module is promoted through use of different development teams, with some measure of forced diversity provided due to different design and implementation tools. Thus, the functions for the control and monitor computers are not necessarily identical, and the programming for each provides some diversity due to personnel and tool set differences.

Different companies supplied the ELAC and SEC modules for the A320. Thomson-CSF (Compagnie Générale de Télégraphie sans Fil—CSF) supplied the ELAC modules, which are based on Motorola 68010 CPUs, while SFENA and Aerospatiale supplied the SEC modules, which are based on Intel 80186 CPUs [81]. For each FCS, different teams programmed each channel in different computer languages and then implemented using different compilers. For the ELAC, the control computer was programmed in Pascal while the monitor computer was programmed in C. For the SEC, the control computer was programmed in MACRO assembler while the monitor computer was programmed in Pascal [82].

Therefore, the primary and secondary FCS for the A320 used two different design and manufacturing teams with different microprocessors (and associated circuits), different software

architectures, and different functional specifications [83]. Within each FCS, separate design teams using different languages, compilers, and other design tool sets were used for the control and monitor channels [84].

4.2.2.2 Airbus A340

The Airbus A340, certified in 1992, represents an evolution of the digital FCS developed for the A320. As with the A320, the overall FCS for the A340 also employs diverse redundant primary and secondary control systems. These systems are the Flight Control Primary Computer (FCPC) and the Flight Control Secondary Computer (FCSC). The primary FCS is also identified as PRIM (meaning PRIMary flight control computer), and the secondary FCS is identified as SEC (meaning SECondary flight control computer). In a manner that is consistent with the architecture established for the A320, both PRIM and SEC are composed of control/monitor computer pairs. These pairs constitute separate parallel channels within each system. These paired channels are replicated to provide three PRIM modules and two SEC modules. Within each module, the control channel generates the flight commands while the monitor channel generates comparative "commands." Both computers within a module compare differences in their outputs based on common input signals. If differences between the outputs exceed a threshold and persist for a sufficient interval, the module is automatically disconnected from the control path to provide a "fail fast" scheme. Control is automatically transferred to a redundant "standby" module that serves as a hot spare [85].

Functional diversity between the primary and secondary FCS is achieved through different control laws, control elements, and reduced functionality. The PRIM system implements elaborate flight control laws for fully functional flight control, while the SEC system implements simpler, more robust flight control laws (i.e., less functions aimed at ensuring smoother flight and greater passenger comfort) [83]. Other diversity usage between the PRIM and SEC systems includes different microprocessors (Intel 80386 for PRIM and Intel 80186 for SEC), different hardware suppliers (Aerospatiale for PRIM and Sextant Avionique [formerly Thomson-CSF and SFENA]), and different application development teams within the common system supplier (Aerospatiale) employing different design approaches and implementation tools (e.g., different high-level specification languages, coding techniques, programming languages, and compilers/translators were employed for the different channels within the different systems) [84]. Specifically, the PRIM control channel was coded automatically in assembly language while the SEC control channel was coded manually in assembly language. Additionally, the PRIM monitor channel was programmed using PL/M (program language for microcomputers) while the SEC monitor channel was programmed using Pascal [83,84].

4.2.2.3 Airbus A380

In 2007, the progressive development of the Airbus digital FCS continued with the certification of the A380. The overall FCS for the A380 also employs diverse redundant primary and secondary control systems. The FCPC and FCSC of the A380 provide similar functionality to those of the A340. The primary differences between the overall FCS approaches for the two aircraft involve architectural changes. The A380 does not provide any mechanical backup for the electronic control linkages, and the FCPC and FCSC are both triple redundant (i.e., three modules of control/monitor channels) for the A380.

The diversity usage for the A380 is similar to that employed for the A320 and A340. The FCPC is supplied by Thales Avionics (formerly Thompson-CSF and later Sextant Avionique), while the FCSC is supplied by Diehl Aerospace, which is a joint venture of Thales Avionics and Diehl Group. The FCPC is based on the Motorola Power PC CPU, while the FCSC is based on a different CPU (identified as a "SHARE" processor in Ref. 83). As before, the functional requirements for the FCPC and FCSC are different (i.e., based on different control laws for different primary control elements with a standby reduced-functionality backup provided by the FCSC). Similarly to the previous generations of the Airbus

FCS, the data flow within the system is loosely synchronized between pairs and modules. Thus, slight differences arise in data values.

Finally, the different suppliers for the FCPC and FCSC result in the use of different development teams for the two systems. Within each organization, different development teams are provided for the control and monitor channels and the use of different automatic code generation tools based on different languages and different compilers/translators is enforced [83].

4.2.2.4 Boeing 777

The Boeing 777 was certified in 1995. It represents the initial and foremost example of the Boeing approach to digital FCS. The B-777 primary flight control system (PFCS) was supplied by GEC-Marconi Avionics Ltd. [86]. The PFCS consists of three parallel channels that are physically and electrically separate. Each channel contains an identical primary flight computer (PFC). The PFCs are the central controllers of the PFCS. The PFCs are connected to data buses to enable transmission of commands to four Actuator Control Electronics units (ACEs) and also to permit information exchange among the controllers.

The channels share their data for equalization to permit direct comparison of consistent computational outputs. In addition, the channels conduct a median selection among their shared outputs to validate each final actuation command [87]. The need for agreement among the channels creates the potential for Byzantine faults [88]. However, the chosen implementation approach provides Byzantine-fault tolerance through bus and data synchronization to address asymmetric faults in communication and command validation to address asymmetric values in functional outputs [85].

In addition to satisfying numerical reliability targets for the PFCS, Boeing also addressed deterministic goals in its design. These goals were as follows: (1) "[n]o single fault, including common [cause] hardware fault, regardless of probability of occurrence, should result in an erroneous … transmission of output signals without a failure indication" and (2) "[n]o single fault, including common [cause] hardware fault, regardless of probability of occurrence, should result in loss of function in more than one PFC" [89].

Consequently, the concept of triple modular redundancy is employed for all hardware resources of the PFCS (e.g., computing systems, airplane electrical power, hydraulic power, and communication pathways). In particular, triple modular redundancy is used in the design of each PFC through the provision of three internal computational lanes [90]. Essentially, the PFCS consists of three identical channels composed on three dissimilar (or diverse) lanes. Thus, the design constitutes a "Triple-Triple Redundancy" architecture.

As shown in Fig. 4.11, three-version dissimilarity is integrated into the design through the use of different hardware (i.e., three different microprocessors). Each PFC consists of three dissimilar computational lanes, with each lane containing its own microprocessor, power supply, and communication interface [91]. The three identical PFCs are designated as Left, Center, and Right. Each PFC receives data from all three of the flight control data buses, but transmits only on its associated bus as shown in the figure. Cross-channel comparisons for median selection are performed based on communications across different buses (e.g., the Left PFC compares its current command against Center and Right commands received across the "C" and "R" buses).

The functionality of each lane is the same, but each is assigned a separate operational role. The three modes of operation are command, monitor, and standby. The command lane is the active controller for the channel. The monitor lane performs the same calculation as the command lane and shares its output for comparison. The standby lane is effectively in hot standby mode as it also performs the same calculations. However, the standby lane does not transmit its output unless it is activated due to a failure of the command lane. The cross-lane comparisons involve transmissions across the same bus (e.g., the left [L] bus for the Left PFC) [92].

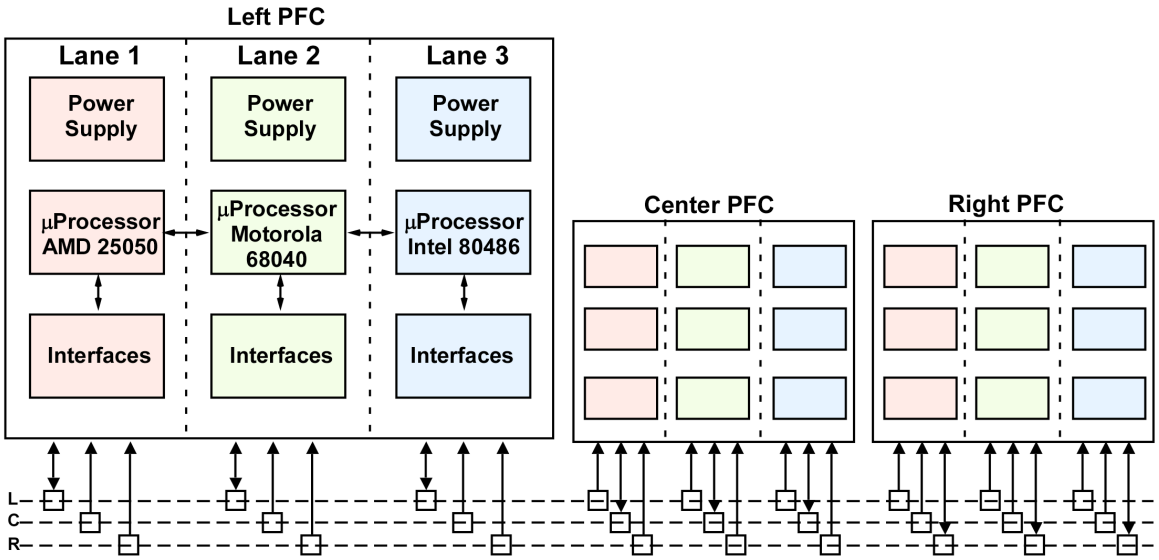

Fig. 4.11. Triple-triple redundancy architecture of the primary flight computer. (Adapted from Ref. 91.)

The initial design approach proposed for the B-777 PFCS was to implement significant diversity among the lanes of each PFC through the use of different design teams and different software implementations [93]. However, Boeing decided against such extensive diversity due to concerns that (1) the management of multiple development teams would be onerous and prone to error [86], (2) effort to maintain independence among the development teams would restrict communication among software and system engineers and prevent correction of requirements errors [94], and (3) adoption of N-version programming would not be effective in avoiding common programming errors [91]. Thus, a single development team for the software application was used to generate the control software as common Ada source code. To enhance software quality, formal methods (e.g., static and dynamic analyses) were applied to PFCS algorithms [86]. Nevertheless, a different Ada compiler for the software implementation in each lane was used to enhance the triple dissimilarity [91].

4.2.3 Rail Transportation Industry

Collision avoidance is a key operational safety concern for railway and train control. Signaling, interlocks, and train speed control are critical functions for ensuring railway safety. Ensuring unobstructed routes and track circuits while controlling train traffic requires a distributed system of sensing and control elements, both embedded on the trains and stationed along the tracks. By stopping or slowing trains to inhibit access to occupied tracks, railways have a readily accessible safe state. Thus, systems can be designed to fail to a local de-energized "stop" configuration. Essentially, a failsafe condition can be achieved in which all trains stop [95]. This failsafe approach results in a practical emphasis for rail safety system on identifying faulted conditions and stopping the affected trains until the hazard can be cleared or the system can be fixed.

Deployed automatic train control, as described in the literature, employs varying amounts of system diversity, which is generally achieved through software, to achieve CCF mitigation. Three examples of CCF mitigation approaches in safety-critical rail transportation applications involve either the safety bag approach or vital coded processors. These techniques are illustrated in the following examples.

4.2.3.1 Austrian Federal Railways

The Elektra railway interlocking control system was designed by Alcatel Austria for the Austrian Federal Railways. The Elektra system was first implemented in 1989. The architecture employs diversity

through its dual channel checking system [96,97]. One channel, designated as the interlocking processor (ILP), executes the interlocking control software (i.e., functional computation), while the other channel, designated as the safety bag processor (SBP), executes the monitoring software (i.e., checking computation). Intercommunication between the processors is implemented via the Votrics communication layer to provide a fault-tolerant message passing architecture for functional monitoring and communications monitoring [98]. Both channels employ redundancy and active replication with identical triple-redundant hardware and software internal to a channel. The replicated, redundant modules perform synchronized computations and require complete agreement to support failure detection. Any faulty component that is detected is reset and resynchronized. The monitoring channel's purpose is to check whether each interlocking control function places the system in a safe state. Actions are only performed if the second channel agrees that the results proposed by the first channel do not violate any safety conditions. If this is not the case, a transition to a safe state is invoked.

For the Elecktra interlocking system, the safety bag implementation employed diverse development teams within Alcatel. Additionally, diverse functional requirements were the basis for the channels, with different functions programmed using different tools in each channel [99]. Specifically, the software specifications for the monitoring channel were derived from the railway authority's operating regulations, while the software specifications for the interlocking control channel were based on functional requirements. Additionally, the monitoring channel was programmed according to a rule-based paradigm, while the interlocking control channel was programmed according to a procedural paradigm. Thus, different languages (i.e., Pamela for SBP and CHILL for ILP) and compilers were used.

4.2.3.2 Paris Rail

In 1988, the Paris Public Transportation Authority (Régie Autonome des Transports Parisiens—RATP) and the Société Nationale des Chemins de fer Français (French National Railway Company—SNCF) engaged a consortium of railway equipment manufacturing companies (GEC Alstom Transport, MATRA Transport, and CSEE Transport, now part of Ansaldo Trasporti) to develop a microprocessor-based automatic train control system. The resulting Système d'Aide à la Conduite, à l'Exploitation et à la Maintenance (SACEM) fault-tolerant train speed control system was first implemented on the Paris Rail line A (RER A or Réseau Express Régional A). SACEM was characterized by development of the vital coded processor (VCP) approach [100]. In 1998, an application of the VCP for the unmanned automatic subway, Météor, enhanced the use of a formal development process to reduce the potential for design errors. The VCP approach is currently supported by manufacturers such as Siemens and Alstom for applications that include the Canarsie Line in New York and the North East Line in Singapore.

The basic premise behind the VCP is to provide a hardwired comparator to confirm the proper execution of the safety or control function in the computer system by comparing expected (i.e., pre-determined) properties of the code against observed or generated properties of the code. The principle of encoding is based on expressing information about the application program and its execution using an arithmetical code, an operational signature, and a dynamic or temporal code (i.e., "technique of dynamisation") [101]. The process for implementing the VCP proceeds as follows [102]. Using a formal process (based on the B formal language for the examples cited), an implementation (i.e., abstract model) is first developed from the software specification in the formal language and is subsequently translated into code. As part of this process, the implementation undergoes formal proof during the development process. The translation of the implementation into code is based on two diversely developed translators (i.e., different teams, designs, and programming languages) [102] to yield two distinct versions of the code. One version is compiled to become the safety application object code. The other version is processed to create reference signatures of the code execution. The VCP is implemented with a hardwired checker that compares precomputed signatures against the actual signatures corresponding to the runtime values. If a discrepancy is detected, an error has occurred and a failsafe condition is enforced.

For the SACEM example, formal methods were used at a later stage of the development effort than for the Météor example. Essentially, the code (written in Modula 2) was developed, inspected, tested, subjected to formal proof, and then processed through formal re-expression (i.e., a formal specification was generated after the fact). Separate teams were used at each stage of the process: design, safety assessment, validation, and formal re-expression. Additionally, separate sub-teams were used for validation [103]. In the case of the Météor application, a formal specification was developed up front [102]. In this case, the concept of separate teams at different life-cycle stages persisted with separate formal support, development, testing, and validation teams. It should be noted that this use of separate teams at different stages of the life-cycle does not necessarily provide life-cycle diversity at the latter stages because the final system is an individual integrated system rather than two separate systems. Thus, there are not two separate, parallel developments by teams that can remain separated, as would be customary usage for life-cycle diversity.

4.2.3.3 Los Angeles Metro Green Line

In the mid-1990s, the Center for Semicustom Integrated Systems (CSIS) at the University of Virginia (UVa) teamed with the Advanced Technology Group of Union Switch and Signal (now a part of Ansaldo Trasporti) to develop the Vital Framework (V_Frame) [104]. The V-Frame is a fault-tolerant safety-critical platform to support the use of COTS hardware and software. The V-Frame can be seen related to the VCP approach except that it does not depend on formal development of the initial code or application-specific implementation of dedicated hardware.

The V_Frame embodies the safety-critical algorithm-based fault tolerance (SC-ABFT) approach developed at UVa [105]. SC-ABFT provides a method for verifying whether applications are executed correctly within a certain probability. In this approach, an application or algorithm is decomposed to its fundamental operations or primitive blocks so that the sequence of execution for those operations can be verified though the generation and confirmation of a check-stream. To avoid the paradox of a self-referencing system, a separate checking device accomplishes the verification of the check-stream [104].

The decomposed algorithm can be represented in terms of a data flow graph that captures key attributes of the set of equations. In particular, the data flow graph uniquely identifies each operator, each input and output object, and the temporal relationship among operators in the execution sequence. Based on this deconstruction, code words can be generated to construct the check-stream representing the correct execution of the algorithm. Subsequently, the correct operation of each primitive block can be precomputed and stored in a look-up table. The blocks themselves are simple enough to allow proof of correctness. Having precomputed, proven blocks enables checking the correct execution of each block in real time by comparing the results of the look-up table versus those of the code calculation. Additionally, corresponding check-streams can be established to enable verification of correct execution in the field. This checking capability is implemented either in “a redundant processor executing software or a low-complexity custom hardware device” [105]. This type of system relies on having primarily discrete, as opposed to continuous, variables to allow the control system to be decomposed into a finite set of states.

The V_Frame implementation of the SC-ABFT was demonstrated in prototype form simulating the Los Angles Metro Green Line. The first demonstration at UVa involved a COTS-based test system using a Motorola 68040 processor card with supporting I/O implemented in a VME-based chassis [105]. Also, the check process was performed via a check algorithm that executed on a Motorola 68040 processor card. Later prototypes involved the use of FPGAs, with a commercial platform being subsequently developed by Ansaldo [106].

5. RECENT NUCLEAR POWER INDUSTRY RESEARCH INTO CCF MITIGATION STRATEGIES

5.1 NRC Research on Diversity Strategies

The NRC has established regulatory guidance addressing a method for assessing the D3 provided by the I&C system architecture at an NPP. This method enables determination of whether vulnerabilities to CCF have been adequately addressed. The guidance is provided within BTP 7-19. This guidance provides a method for determining the need for diversity. However, no definitive guidance was available specifying how much diversity is sufficient to mitigate CCF vulnerabilities that may arise from digital safety system designs. Thus, ORNL was engaged to develop a technical basis for establishing acceptable mitigating strategies that address the potential for digital CCF vulnerabilities [107]. The specific objective of this research effort was to identify and develop diversity strategies, which consist of combinations of diversity attributes and their associated criteria, by leveraging the experience and practices of other industries and the international nuclear power community. Effectively, these baseline sets of diversity criteria constitute appropriate mitigating diversity strategies that can adequately address potential CCF vulnerabilities in digital safety systems. The strategies are suitable for use by regulatory staff as comparative templates or guides to support confirmation of acceptable diversity usage in addressing CCF vulnerabilities that are identified via a D3 analysis. The purpose of this report is to document the diversity strategies developed through this research and describe the supporting technical basis.

5.1.1 Research Approach and Methods

The diversity strategies developed through the NRC research are composed of combinations of diversity criteria that are adapted from the attributes and criteria defined in NUREG/CR-6303. The guidance separates diversity attributes into the following six areas to facilitate assessments of adequate diversity in safety systems:

- design diversity,
- equipment diversity,
- functional diversity,
- human diversity,
- signal diversity, and
- software diversity.

To better reflect the nature of specific diversities, the attributes were expanded and clarified. The "human" diversity attribute is designated the "life-cycle" diversity attribute to account for its true nature and to avoid the erroneous inference that this attribute involves plant operator diversity or human-versus-machine diversity. In fact, the human (i.e., life-cycle) diversity attribute relates to addressing human-induced faults throughout the system development life-cycle process (e.g., mistakes, misinterpretations, errors, configuration failures) and is characterized by dissimilarity in the execution of life-cycle processes. Additionally, the "equipment" diversity attribute is subdivided into two new attributes to reflect the differences related to the manufactured equipment source (i.e., the manufacturer or supplier) and the differences related to logic processing components (e.g., computational or processing elements such as CPU, printed circuit board, bus architecture for microprocessor-based equipment). Thus, the single "equipment" diversity attribute is treated as two diversity attributes: "equipment manufacturer" and "logic processing equipment." Finally, the "software" diversity attribute is designated the "logic" diversity attribute to account for the different means of representing and executing functions that diverse technologies provide (e.g., software for microprocessors, hardwired logic in programmable devices, electronic circuitry for analog modules).

The guidance in NUREG/CR-6303 provides a set of recommended criteria for each of the diversity attributes. However, the number of criteria in each attribute, coupled with the number of attributes, creates a large number and complexity of possible combinations of attributes and criteria that could be used to achieve adequate diversity in a safety system, making the guidance difficult to use as a safety assessment tool. Nevertheless, it is possible to define effective diversity strategies based on consensus practices and experience within other application domains.

The research approach for establishing diversity strategies involved capturing expert knowledge and lessons learned, determining best practices, and assessing the nature of CCFs and compensating diversity attributes. The basis for these strategies centers on practices derived from examples of diversity usage by the international nuclear power industry and several nonnuclear industries with high-integrity and/or safety-significant I&C applications. The approaches to diversity identified from international NPPs serve as representative examples of the strategies. While the examples identified from nonnuclear industries are relevant because of the safety significance of the functions and the use of comparable technology, context differences in the usage domains limit their direct applicability. Thus, key insights are derived from these examples to inform the development of diversity strategies in this research. The resulting diversity strategies address considerations such as the effect of technology choices, the nature of CCF vulnerabilities, and the prospective impact of each diversity type. In particular, the impact of each attribute and criterion on the purpose, process, product, and performance aspects of diverse systems are considered.

Based on the findings of the investigation into diversity usage practices and the establishment of acceptable diversity strategies, a prototypic comparative tool was developed to provide a resource to consistently confirm that the amount of diversity in a safety system design is sufficient relative to a predetermined acceptance threshold region. As part of the research effort, usage information on diversity attributes was collated consistent with modified NUREG/CR-6303 diversity attributes and criteria. The diversity attributes and criteria were then weighted using the available set of usage examples. The weights and supporting algorithms were translated into a worksheet format to allow users to evaluate the relative amount of diversity in a system design, independent of the technology employed. Subsequently, common trends in diversity attributes and related criteria usage were identified as the basis for a process to comparatively evaluate diversity in safety system designs. Development of this basis involved establishment of an acceptance threshold region derived from previously accepted diversity usage and the baseline strategies defined through this research. The spreadsheet tool provides the capability to quantitatively compare candidate diversity usage strategies against a range of predefined acceptable diversity strategies.

5.1.2 Definition of Diversity Strategies

The study of diversity in nonnuclear industries identified different approaches that range from no diversity (e.g., the almost total reliance on redundancy of high-quality modules and defense-in-depth layers with no “intentional” diversity) to minimal diversity (e.g., reduced functionality backups with limited diversity) to more extensive diversity (e.g., combinations of techniques for fault management addressing high-consequence failures with “encouraged” but not fully specified diversity). The primary diversities cited for establishing sufficient application independence are functional, signal, software, and life-cycle (associated with the application software). While some examples of diversity usage have been noted in other industries, there have been little explicit guidance and infrequent dependence on this approach. The less-common utilization of diversity as a mitigating strategy for several nonnuclear industries appears to be driven by considerations such as fundamental reliance on high-quality practices and procedures within an application domain, the nature of the applications and behavior of the processes, implementation constraints (e.g., size, weight, power, and cost), and acceptability of some risk.

For evolutionary NPPs with significant use of digital systems, a common diversity usage approach involves a systematic subdivision of the protection functions into versions A and B and an assessment of the degree of diversity between the two versions based on a pair-wise comparison of the individual

mitigation characteristics. The result is identification of the categories of the diversity attributes that can be used to show that the diverse systems do not have some common vulnerability that could cause a protective function to fail. Most digital I&C system architectures identified in the investigation make the claim of diversity, but they differ in overall approach. The approaches to diversity usage in the reported case histories can be grouped into three broad categories: coequal diverse systems, primary/secondary diverse systems, and functionally diverse subsystems. Of these examples, functional diversity is the most common.

By employing the findings from the diversity usage investigation, baseline combinations of diversity attributes and criteria were formulated to establish acceptable diversity strategies. To facilitate the development of the strategies, a framework for classifying strategic approaches to diversity usage was devised. Technology, which corresponds to the design diversity attribute of NUREG/CR-6303, is chosen as the principal system characteristic by which the strategies are grouped. The rationale for this classification framework involves consideration of the profound impact that technology-focused design diversity provides. Basically, instances of design diversity are readily observable and most of the other diversity attributes are strongly affected by the design/technology choice. Specifically, NUREG/CR-6303 states that "the clearest distinction between two candidate subsystems would be design diversity."

The classification of diversity strategies developed in this research consists of three families of strategies: (1) different technologies—Strategy A, (2) different approaches within the same technology—Strategy B, and (3) different architectures within the same technology—Strategy C. Using this convention, the essential characteristics of the three strategy families are summarized as follows:

- **Strategy A** focuses on the use of fundamentally diverse technologies as the basis for diverse systems, redundancies, or subsystems. The Strategy A baseline, at the system or platform level, is illustrated by the example of analog and digital implementations providing design diversity. This choice of technology inherently contributes notable equipment manufacturer, processing equipment, functional, life-cycle, and logic diversities. Intentional application of life-cycle and equipment manufacturer diversities is included in the baseline, while the traditional use of functional and signal diversities is also adopted. The use of a microprocessor-based primary protection system and an analog secondary protection system represents the principal example of Strategy A drawn from the survey findings.
- **Strategy B** involves the use of distinctly different technology approaches as the basis for diverse systems, redundancies, or subsystems. The Strategy B baseline can be described in terms of different digital technologies, such as the distinct approaches represented by programmable logic devices and general-purpose microprocessors. This choice of technology inherently contributes some measure of equipment manufacturer, processing equipment, functional, life-cycle, and logic diversities. Intentional application of logic processing equipment, life-cycle, and equipment manufacturer diversities is included in the baseline, while the traditional use of functional and signal diversities is also adopted. The application of different digital technologies (i.e., CPUs vs FPGAs) as the basis for a primary safety system and a diverse backup (or checker) system represents the principal example of Strategy B drawn from the survey findings.
- **Strategy C** represents the use of architectural variations within a technology as the basis for diverse systems, redundancies, or subsystems. An example of the Strategy C baseline involves different digital architectures, such as the diverse microarchitectures provided by different CPUs. This choice of technology inherently contributes some limited degree of equipment manufacturer, life-cycle, and logic diversities. Intentional application of equipment manufacturer, logic processing equipment, life-cycle, and logic diversities is included in the baseline, while the traditional use of functional and signal diversities is also adopted. The use of diverse microprocessors as the basis for primary safety systems and diverse backup systems, such as (ATWS) or (DAS), constitutes the principal examples of Strategy C drawn from the survey findings.

As noted, each of the strategy families is characterized by combinations of diversity criteria that provide adequate mitigation of potential CCF vulnerabilities when combined with the traditional

diversities generally employed for conventional hardwired systems. In addition to the baseline strategy within each family, acceptable variants of each baseline were also developed. Implementation of a diversity strategy (e.g., baseline or identified variant) from any of the three families serves to minimize the opportunities for common systematic faults, concurrent execution profiles, and similar responses to external influences that can contribute to the potential for CCF vulnerabilities in digital I&C systems.

Table 5.1 provides an overview of the three baseline strategies in terms of criteria adapted from NUREG/CR-6303. The basis for the strategy classifications was the technology employed, given that this fundamental difference between systems provides an identifiable, easily recognizable diversity characteristic of system design. Acceptable variants of these three strategies were also developed.

Table 5.1. Overview of baseline diversity strategies

Diversity attribute	Strategy[a]		
	A	B	C
Design			
Different technologies	x	–	–
Different approach—same technology	–	x	–
Different architectures	i	i	x
Equipment Manufacturer			
Different manufacturer—different design	x	x	–
Same manufacturer—different design	–	–	–
Different manufacturer—same design	–	–	x
Same manufacturer—different version	–	–	–
Logic Processing Equipment			
Different logic-processing architecture	i	i	x
Different logic-processing versions in same architecture	–	–	–
Different component integration architecture	i	x	x
Different data-flow architecture	i	–	–
Functional			
Different underlying mechanisms	i	i	–
Different purpose, function, control logic, or actuation means	x	x	x
Different response-time scale	–	–	–
Life-cycle			
Different design organizations/companies	x	x	x
Different management teams within same company	–	–	–
Different design/development teams (designers, engineers, programmers)	i	i	i
Different implementation/validation teams (testers, installers, or certification personnel)	i	i	i
Logic			
Different algorithms, logic, and program architecture	i	x	x
Different timing or order of execution	i	i	–
Different runtime environment	i	i	x
Different functional representation	i	i	x
Signal			
Different parameters sensed by different physical effects	x	x	x
Different parameters sensed by same physical effects	x	x	x
Same parameter sensed by a different redundant set of similar sensors	x	x	x

[a]Intentional diversity (x), inherent diversity (i), not applicable (–).

5.1.3 Implementation Approach to Facilitate Assessment of CCF Mitigation

The grouping of diversity combinations according to Strategies A, B, and C facilitates a systematic organization of strategies into families that are readily amenable to evaluate. The classification of strategies enables a consistent representation of the comparative use of diversity between systems, redundancies, subsystems, modules, or components. As a consequence, this research leads to a systematic evaluation process for reviewing the application of diversity strategies to address CCF vulnerabilities identified through a D3 assessment.

The principal elements of the diversity evaluation process, which is applicable to confirm the response to any CCF vulnerabilities identified via a D3 assessment, include the following steps:

1. Classify the diversity strategy—identify what technology is employed.
2. Confirm inherent diversity credit—ensure that intrinsic benefits of technology differences are not compromised.
3. Identify intentional diversity usage—verify which intentional diversities are explicitly employed to address CCF.
4. Categorize diversity usage as a function of one of the following:
   - Strategy A, B, or C;
   - one of the variants of A, B, or C; or
   - alternate strategy.
5. Assess the diversity strategy—The diversity usage tables and diversity assessment tool developed through this research provide support for comparative evaluations against the baseline diversity strategies.
6. Determine if the diversity strategy is adequate—A conclusion that a proposed diversity strategy adequately addresses CCF mitigation needs, as identified via a D3 assessment, can be based upon either conformance to one of the three baseline strategies (or an accepted variant) or determination that the strategy reasonably ensures CCF mitigation comparable to that provided by a baseline strategy (i.e., an acceptable rationale is provided to support mitigation claims).

The evaluation process for diversity strategies is intended to appropriately credit the inherent diversities arising from the chosen technologies while emphasizing identification of the intentional diversities explicitly employed to address the potential CCF vulnerabilities. In assessing the rationale for an alternate diversity strategy, the impact of each diversity criteria on purpose, process, product, and performance aspects of the diverse systems should be considered. The objective is to confirm that the diversity strategy provides sufficient CCF mitigation capability by adequately minimizing the opportunity for common systematic faults, reducing the occurrence of concurrent execution profiles, and lessening the likelihood of similar responses to external influences.

5.1.4 NRC Research Conclusions

The results of this research effort have identified and developed diversity strategies, which consist of combinations of diversity attributes and their associated criteria, by leveraging the experience and practices of nonnuclear industries and the international nuclear power community. Effectively, these baseline sets of diversity criteria constitute appropriate mitigating strategies that adequately address potential CCF vulnerabilities in digital safety systems. The strategies represent guidance on acceptable diversity usage and can be applied directly to ensure that CCF vulnerabilities identified via a D3 assessment have been adequately resolved. Alternately, the strategies can serve as comparative norms, in combination with the diversity usage tables and/or diversity assessment tool, to support confirmation that equivalent CCF mitigation capability is provided.

5.2 British Research on Diverse Software

Under British law, the nuclear power industry in the United Kingdom must fund safety research annually as part of the U.K. Nuclear Research Programme. This research has included investigations into methods for characterizing software diversity, employing statistical testing approaches for validating software quality, and qualifying smart sensors. The research into software diversity [108] is of particular relevance to the investigation into effective CCF mitigation practices and identification of knowledge gaps. Principally, the software diversity research was conducted by the Centre for Software Reliability at City University London and the Critical Systems Research Centre at Bristol University under research contracts established for the DIverse Software PrOject (DISPO). The DISPO projects began in 1996 and were conducted initially over 3-year periods. Later projects were conducted on an annual basis with the most recent project, DISPO5, covering 2006 and 2007.

The primary characteristics of the DISPO research are

- detailed problem parsing,
- careful progression of research topics,
- cautious logic about overextending conclusions, and
- reliance upon probabilistic models to understand the effects of commonality or separation influences on producing diverse versions of a system for a diverse redundant configuration.

The DISPO research concentrates on the use of diversity in digital systems. The basic application of diversity within an NPP I&C architecture composed of software-based systems involves parallel redundant systems or subsystems (e.g., versions, channels, redundancies) that perform the same or equivalent functions and are arranged in a one-out-of-*N* or voted configuration. The simplest example is a diverse redundant pair of systems (or redundancy versions) that are implemented in a one-out-of-two (logical "OR") configuration.

The DISPO research focus involves the use of diversity as a means to achieve system dependability, with particular emphasis on accounting for the likely presence of faults in software. Dependability is defined as the "[t]rustworthiness of a delivered service (e.g., a safety function) such that reliance can justifiably be placed on this service." Attributes of dependability include reliability, availability, and safety. The findings of the research contribute to understanding the relationship between diversity and failure independence, identifying life-cycle decisions that encourage diversity, and assessing the qualitative impact of diversity.

An additional consideration introduced toward the later stages of the DISPO research is the application of diversity in the assessment of dependability. Essentially, this aspect of the current investigation addresses the use of "diverse arguments." This study involves the consideration of diversity to address weaknesses in the arguments that are used to support dependability claims (i.e., diverse bases for multiple or "multi-legged" arguments). The findings suggest that the potential increase in confidence for a claim depends crucially on the degree of dependence between arguments (e.g., between their assumptions).

5.2.1 General Findings of the DISPO Program

The relevant measure of interest for system dependability is the probability of failure on demand (pfd) for the system. A common erroneous assumption is that different versions resulting from "independent" (or, more accurately, separate) development processes result in independence of failure behavior [109]. Experimental studies by Knight and Leveson [110] showed that "independently" developed software versions did not necessarily fail independently. In essence, the failures identified in the investigation were not statistically independent but rather showed a strong positive correlation between the failure behaviors of the different versions solving the same problem. It is noted that even

with separate development teams and processes, common influences, assumptions, understandings, and mistakes may be present and there may be only conditional independence of version failures. In effect, dependent failure sets (i.e., the set of demands that result in failure due to the presence of faults) may exist. There is often an erroneous assumption that the common pfd of two versions is zero and that the pfd of the diverse redundant system (composed of the separately developed versions) is exactly equal to the product of the probabilities associated to each of the two failures sets. As this assumption does not generally prove true, a conclusive mathematical basis often cannot be established to demonstrate that increasing diversity between versions will increase diversity against failure.

As a consequence of knowledge captured from system development experience, advances in reliability modeling of diverse systems, and experiments conducted during the multiyear research program, a principal DISPO finding is that claims for statistical independence between failures of diverse versions have not been reasonably demonstrated. Of particular note, claims of independence for diverse system failures cannot be sustained even in the case of applied functional diversity. These findings clearly indicate that the provability of dependability for an overall system based on design diversity is limited. The research shows that independent development by itself is not sufficient to ensure the version failures are independent for a randomly chosen demand. Nevertheless, it is observed that increasing diversity may increase reliability for separate developments. In those instances, overall system reliability may be enhanced by strong diversity enforcement mechanisms.

The DISPO research has shown that confidence in a dependability claim can be increased through the use of design diversity. The conclusion is that "forced diversity is a good thing," although individual or collective effects are difficult to quantify. The research has also shown that some forms of dependence or interaction (e.g., shared knowledge about requirement deficiencies) may bring substantial benefit not only to the development process but also to the resulting system reliability.

Two key technical issues investigated by the DISPO research team involve the achievement of dependability and the assessment of dependability. Regarding the former, the application of diversity in digital I&C systems can be encouraged by invoking decisions in the management of the system design process. These choices are described as diversity-seeking decisions (DSDs). The effect of such decisions is to promote a high degree of fault diversity. The remaining challenge arises because the effect of these decisions on failure diversity (i.e., achieving reduced correlation between failure behaviors of different versions) is indirect. Figure 5.1 illustrates the relationship. There is insufficient knowledge to definitively guide the choice among DSDs to effectively produce the desired failure diversity and thus, in turn, quantify improvement in system dependability. However, there is clear qualitative evidence of the benefit of applying these DSDs individually.

Regarding the latter technical issue investigated by the DISPO research team, assessment of dependability involves establishing assurance that critical (or safety) functions are protected against CCF through diversity. Assessment involves both oversight of the development processes to ascertain that diversity is present and understanding of the associated impact on the pfd corresponding to each diverse system. The DISPO research has contributed to improved reliability assessment for diverse systems in terms of "independent fault" models. However, since there is not a known definitive relationship between the DSD-induced product/fault independence and the desired failure independence, dependencies are to be expected. Nevertheless, research findings indicate that if the separate development processes are managed to enforce diversity, then independence (or possibly negative correlation) between failures of design-diverse versions can be enhanced. However, the application of such measures is still insufficient to conclusively justify claims of independence. The problem remains that even when the presence of diversity is established, there are no quantifiable measures to determine its efficacy and there is no means of assessing the system reliability (or the impact on safety) from such knowledge. The bottom line is that the use of diversity (particularly forced diversity) as a means of improving dependability of software-based systems through fault tolerance is beneficial, but there remain real difficulties in assessing what the quantitative effect on reliability for specific systems.

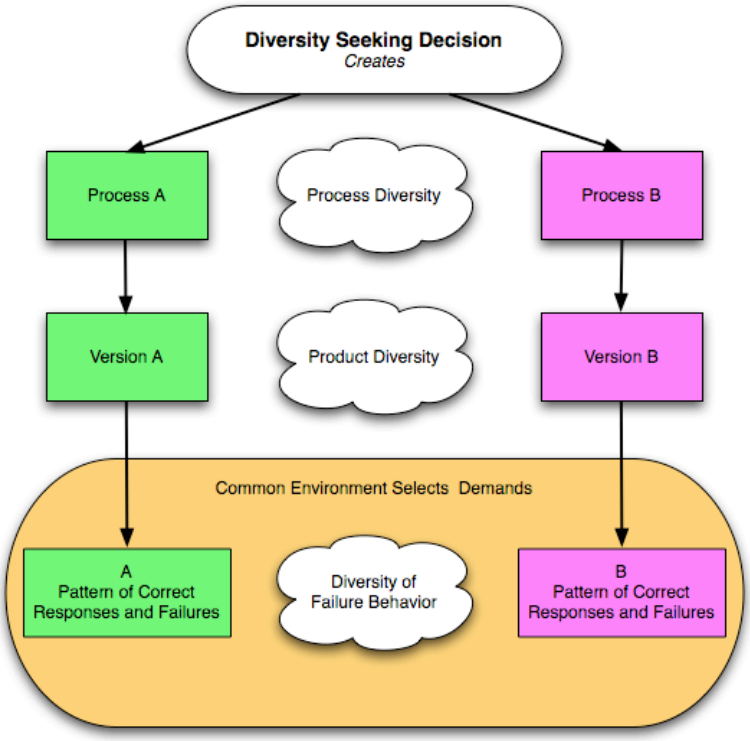

Fig. 5.1. The different facets of diversity and their interdependence. (Adapted from Ref. 111.)

5.2.2 Practices for Achieving Diversity

As a fundamental element of this research program, the DISPO research team investigated the effect of variation of difficulty in which the concept of "difficulty functions" was developed [112]. A difficulty function is described as the probability that a randomly chosen version would fail on a given demand, which indicates the difficulty in achieving the desired response (or conversely, the ease of making a mistake in implementing the desired response to that demand). The presumption is that mistakes correlate with the difficulty posed by demands or requirements. Essentially, the idea is that it is possible to develop dissimilar software versions by employing different processes (e.g., different software engineering practices and procedures), leading to different difficulty functions over the space of demands. The desired result is negative covariance between the difficulty functions for different versions, which means that demands that are "hard" to satisfy (or cause difficulty) for one version are not the same as those demands that are difficult for another version. Difficulty variation may be achieved by decoupling development activities that are essentially the same. Substituting different influence factors (e.g., management directives, shared communication, resource availability, etc.) for each version or forcing diversity within development activities through the intentional use of different processes, methods, tools, etc., are examples of means by which to accomplish this goal.

Key findings have been published regarding practices that have the potential for increasing the extent of diversity between redundant implementations of software or software-based systems [111]. Limitations of existing knowledge preclude definitive diversity recommendations based on quantitative estimation of the combined effect of specific practices. However, useful indications of the qualitative effect have been observed and provide some measure of justification, beyond intuitive argument, for decisions that contribute to diversity. The presentation of DISPO research findings in Ref. 111 summarizes prospective means for achieving failure diversity with respect to design faults that induce failures, discusses expected advantages for each method, and identifies available anecdotal or experimental evidence.

A means of "forcing" diversity is to impose constraints on the software-based system development process to introduce development differences between two versions of a diverse redundant architecture. The desired benefit is that the difficulties presented to each development team will differ, leading to the credible prospect that common faults would be unlikely to occur in the two versions. Based on the research, examples of DSDs include the following: "using different development environments, different tools and languages at every level of specification, design and coding, implementing each function with different algorithms, applying different V&V methods, etc." Table 5.2 summarizes the identified DSDs that can contribute to achieving the goal of failure independence of software-based systems through data diversity, design diversity, and functional diversity. In this context, data diversity refers primarily to input differences achieved by measurement dissimilarity, stochastic signal behavior, analytical variation (reexpression), and loose coupling between functional instantiations (e.g., asynchronous execution of function in separate systems). Design diversity embodies all of the design options that can engender diversity in the development of parallel systems that provide the same or similar input-to-output function. Thus, design diversity in the DISPO research context is more expansive in scope than for its usage in NUREG/CR-6303, which focuses primarily on technology choice and usage.

Table 5.2. Overview of diversity-seeking decisions from U.K. DISPO research

Diversity-seeking decision	
DSD type	**Variant**
Data Diversity	
Diverse inputs	— Stochastic input variations — Reexpression of inputs — Different signal sources (with functional diversity)
Design Diversity	
Separate developments	
Diverse development teams	
Diverse descriptions, programming languages, and notations	
Diverse requirements or specifications	— Different expressions of identical requirements — Different required properties/constraints providing the same behavior — Different required behavior
Diverse development methods	
Diverse verification, validation, and/or testing	
Diverse code (automatic code transformation)	
Diverse development platforms	— Different tools — Different compilers
Diverse support platform (runtime platform)	— Separation and loose coupling — Different timing — Different (dissimilar) hardware — Different operating systems or runtime executives
Functional Diversity	
Diverse functionality	

The design diversity discussed in the DISPO research includes differences in the system life-cycle process (e.g., resources, methods, tools) as well as different implementations of the functionality. Establishing "cognitive" diversity for the designers, implementers, testers, and so forth is central to

minimizing the potential for common mistakes, errors, and misunderstandings that can lead to systematic faults. Functional diversity involves the establishment of different functional relationships (e.g., diverse parameter and initiation criteria to protect against the same PIE) as the basis for diversity. Signal diversity as discussed in the DISPO research is necessary to enable functional diversity.

Two forms of forced diversity are discussed extensively in the DISPO research: (1) "normal" forced diversity and (2) functional diversity. The first approach to forced diversity involves imposed differences in development activities leading to different design versions that are based on the same underlying physical relationships that correspond to use of the same or similar inputs to indicate each specific event. The second approach to forced diversity involves employing alternate underlying physical relationships and results in different design versions utilizing different inputs to provide indication of each event. It is noted that claims of independence among functionally diverse systems are not absolutely justified, and there is a distinct possibility of correlated failures [113]. Thus, functional diversity, while very effective at addressing key CCF vulnerabilities, can benefit from application of other DSDs.

5.2.3 Qualitative Impact of Diversity

For many of the plausible DSDs, the mechanism by which dependability improvement can be achieved is not fully understood (i.e., "how" specifically they work and, as a result, "why" or "whether" they will work). Most of these DSDs address the likelihood of faults rather than failures, so the recommendations do not directly resolve CCF potential. Nevertheless, it is concluded that taking action to reduce the potential common faults is a reasonable approach. Additionally, the research team notes that diverse runtime platforms are considered to be "the only form of diversity that is generally and absolutely necessary, as the system designers usually have no other effective defence against platform faults."

The DiSPO project categorized DSDs of various types (data, design, and functional diversity) and provided a list of problems to which the DSDs could reasonably be applied and the potential costs of doing so. Table 5.3 is a concise presentation of this information. The researchers note however, the following caveats [111]:

- Probable mechanism of action and the problems tolerated: The degree to which these types of faults are tolerated is not usually known. The reader would use this column to match DSDs to perceived threats, and check that all threats against which diversity is the preferred defense are actually "covered" by that DSD. In general, combining DSDs will "cover" the union of the sets of threats against which the DSDs are individually effective.
- Efficacy of DSDs against a specific threat: It is generally unclear how much combining two or more DSDs believed to help against a given threat will improve the defenses vs adopting just one of the DSDs. It is possible that this improvement is quite limited.
- Considerations on cost, efficacy, and practical experience: The considerations of cost included in the table are limited to factors of additional cost caused by DSDs and omits the obvious cost of duplication of all stages of development subsequent to that for which the DSD is applied. In general, the costs include the cost of replicated activities, the cost of coordinating the replicated activities, and savings in some activities such as testing. Generic cost models have also been published to assist in these determining these estimates [114–116].

Table 5.3. Overview of DSD, problems tolerated, and cost-efficacy considerations (from Ref. 111)

DSD Type	Probable Mechanism of Action; Problems Tolerated	Cost, Efficacy, and Practical Considerations
Data Diversity		
Using random perturbations of inputs, or Using algorithm specific re-expression of inputs	Ensuring that if the input to one channel is within a failure region, the input to another identical channel may not be. It should be more effective with failure regions that are small or irregularly shaped.	Generally cheap as no diverse versions required. Efficacy proven in experiments, very variable between faults. "Random" data diversity is often obtained *gratis* as a side effect of other decisions in fault-tolerant design. Special re-expression algorithm may imply additional costs.
Design Diversity		
Separate ("independent") developments	Protection of developments against all *unnecessary* common influences that may lead to common failures	Most basic DSD, necessary precondition (and necessary cost) for applying most others. In some situations may be the most effective DSD.
Diverse development teams	"Forced" diversity via team selection is an appealing idea, but not proven in practice	Appears very desirable, if difficult to implement, *within* version developments; *between* version developments, serious issues of "diversity vs version reliability"
Diversity in description/programming languages and notations – Overall	Probable defense against some slips, and cognitive diversity against mistakes in higher-level problem-solving Advantages affect both writers/verifiers of documents/programs and their users: implementors of next-stage, more detailed document Diverse programming languages also usable to promote diversity in demands on platform	Efficacy must depend heavily on "how different" the languages are (e.g., functional vs imperative) With very diverse languages, issues of "diversity vs version reliability" With diverse specs, it may be possible to use appropriate language for each version and have very different languages

Table 5.3. (continued)

DSD Type	Probable Mechanism of Action; Problems Tolerated	Cost, Efficacy, and Practical Considerations
Design Diversity (continued)		
Diversity in description/programming languages and notations – Diverse requirements or specifications	At all stages in development: cognitive diversity Advantages affect both writers/verifiers of documents/programs and their users: implementors of next-stage, more detailed document	Wide range of options, from purely aesthetic differences to specifying completely different behaviors, at system or subsystem level. The latter is presumably most effective, but increases system design effort
Diversity in description/programming languages and notations – Different expressions of substantially identical requirements	cf "Diversity in description/programming languages and notations - overall"	Often cheap
Diversity in description/programming languages and notations – Different required properties implying the same behavior	Cognitive diversity benefiting both writers and verifiers of the specification and implementors of the specification Special cases: reduced-functionality secondary version, with scope for higher reliability checker-only channel, with greater scope for cognitive diversity	Applicability varies Some problems may be easily specified via alternative, equivalent sets of required properties, some cannot Issues of in-process vs between processes diversity
Diversity in description/programming languages and notations – Requiring different behaviors from the diverse versions	Cognitive diversity benefiting both specifiers and implementors See also "functional diversity"	Applicability varies with availability of alternative algorithms for achieving same goal System design and version specification complications to ensure that diverse version exhibit consistent enough behavior

Table 5.3. (continued)

DSD Type	Probable Mechanism of Action; Problems Tolerated	Cost, Efficacy, and Practical Considerations
Design Diversity (continued)		
Diverse development methods	Cognitive diversity benefiting those applying the methods and those applying its results	With complete, "packaged" methods there is little chance of redesigning DSDs in detail: limitation but also probable savings When designing differences of detail (e.g. applying different methods in requirement elicitation), issues of in-process vs between-processes diversity
Diverse verification, validation, testing	Both inherent differences in defects covered, and cognitive diversity	Issues of in-process vs between processes diversity, "diversity vs version reliability"
Automatic code transformation	Producing different inputs to compilers to tolerate their faults	Cheap and obvious, but may be defeated by compiler optimization
Diverse development platforms: Diverse tools	Diversity in: limitations of tools in preventing mistakes; defects in tools that may cause software faults; presentation of problems to users (cognitive diversity)	A mixed bag of heterogeneous possibilities: wide range of costs, many practical constraints as tools are limited to applying specific methods and notations. In this sense, diverse tools may improve separation between developments
Diverse compilers (also applicable to replicas of a single version)	Diverse compiler bugs, so that any failure points they introduce are hoped to be different in different versions; diverse executable which may tolerate faults in run-time platforms	Usually cheap, popular due to compilers being in universal use, complex and known to have bugs

Table 5.3. (continued)

DSD Type	Probable Mechanism of Action; Problems Tolerated	Cost, Efficacy, and Practical Considerations
Design Diversity (continued)		
Diverse support platforms: Run-time platform	Diverse platforms should exhibit different bugs possibly failing on different demands; diverse robustness with respect to application failures; possibly diverse requirements on application behavior	Important as run-time platforms are known to have design faults and are often outside the control of application system designers. Also, platform faults may cause common failures on specific [classes of] demands irrespective of details of application
Diverse support platforms: Separation and loose coupling Diverse timing	Data diversity for both applications and platforms, in addition to better tolerance to upsets from EMI, etc.	Usually decided on grounds of system design philosophy, hardware fault tolerance: advantages for software fault tolerance are *gratis* With more complex adjudication than wired-OR, loose-coupled redundancy requires special care in design
Diverse support platforms: Diverse hardware	Different bugs, different compiler bugs	Comparatively inexpensive as mostly about buying diverse off-the-shelf components. Virtually mandatory as microprocessors (and probably complex ASICs) should be expected to have design faults
Diverse support platforms: Diverse operating systems or run-time executives	Different OS bugs; different requirements on applications, hence some cognitive diversity for application level developers; different demands on hardware and thus some tolerance of hardware faults	There is evidence that even different COTS implementations of the 'same' operating system specifications (e.g., POSIX standard) exhibit some failure diversity There is the attractive though unproven possibility of running application versions on radically different OSs, e.g., event-triggered vs time-triggered

Table 5.3. (continued)

DSD Type	Probable Mechanism of Action; Problems Tolerated	Cost, Efficacy, and Practical Considerations
Design Diversity (continued)		
Diverse support platforms: "Partial" diversity, limited to subsystems	The subsystems that are diversified benefit from the effects of the DSDs applied to them; they may produce beneficial data diversity for the other (non-diverse) subsystems	May be a way of containing the cost of diversity, focusing resources on the more critical subsystems
Functional Diversity		
Functional Diversity	Cognitive diversity at all stages in development, and differences in limitations of sensors and physical models in regions of controlled system's state space May tolerate all kinds of errors, including specification errors and even gaps in understanding of controlled system's behavior It should not be assumed to ensure failure independence between versions; common causes of mistakes causing common-mode failures may re-appear in later stages of development despite diversity at requirements level	Widely used throughout safety-critical applications to tolerate both physical and design faults Intuitively appealing: maximum possible degree of cognitive diversity between developments, at the cost of developing separate specifications for the diverse channels; enforces separation of developments in later stages Increases system design effort to ensure consistency of top-level requirement on all versions (lesser problem for simple protection systems)

5.2.4 Using Dependence to Decrease Correlation between Faults in Multiple Versions

The DISPO project identified some forms of dependence that may decrease the correlation between faults in two version of software. This is accomplished by reducing the covariance (over the space of possible values of the various influencing factors) between mistakes of the two teams affecting the same demand. These approaches tend to ensure that if a demand appears likely to be a failure point for one version, it will be made less likely to be a failure point for the other version.

Examples include:

- monitoring the progress of the two teams' designs and intervening to keep them diverse.
- monitoring the developed code and assessing the defect proneness of the various modules. If a module in version A appears to have become especially complex or otherwise at risk of faults, team B could be required to apply special precautions to those module(s) serving similar purposes in version B;

- monitoring testing and if a test case reveals a fault in one version, instead of fixing it, concentrating the testing effort for the other version on that demand and similar ones. This is not a likely approach in safety-critical software.

These examples depend on explicit monitoring of the versions' development, verification and validation.

Additional DISPO findings on achieving diversity between (among) two or more versions include the following:

- Combining diversity in design with diversity in fault removal can cost-effectively improve robustness against CCFs.
- The optimum combination of diversity attributes and diversity criteria is unknowable in general, given the inadequate understanding of the direct relationship between minimizing common faults with the goal of minimizing common failure.
- The prospective increase in reliability through the use of diverse separate development processes can only be characterized realistically as an expectation that the uncertainty associated with the system reliability will be greater if the introduction of diversity through DSDs is not carefully administered to encourage independence.

5.2.5 DISPO Research Conclusions

The DISPO project made significant advances in software diversity theory for digital I&C software, namely in the identification of DSDs, their individual threat tolerance potential, and their cost-efficacy based on a qualitative assessment of using individual DSDs.

A summation of the major accomplishments of the DISPO research includes the following:

- identification of coarse-grained rules-of-thumb about considerations driving decisions about the application of diversity;
- conduct of a more abstract analysis of the factors that determine the effectiveness of decisions in enhancing effective diversity and/or system reliability to assist practitioners in selecting among decision rules proposed by others and developing their own decisions in their specific circumstances;
- identification of relevant diversity-related questions that have "no purely mathematical answer and for which it is difficult to define clearly which answers are scientifically justified" while taking into account known empirical results and the diverse expert opinions in the field; and
- clarification of which decision criteria are based on more solid bases, and on which bases the more uncertain conclusions are founded.

A key conclusion of the research is that the demonstration that functional diversity does not, of itself, guarantee the validity of claims for failure independence. In addition, it is acknowledged that the research findings are limited in the degree to which objective decision criteria can be established. Specifically, it was found that an extensive range of possible diversity-seeking decisions and their combinations are available but there is little definitive guidance for these choices. As stated below, the DISPO research team concluded that further research was necessary to achieve the desired practical guidance [117].

> "[W]e do not hide that many of these results are only useful for better understanding these complex, counter-intuitive problems: they do not lead to simple, general recipes for design and assessment. Such understanding is necessary before it is possible to begin engineering diverse fault-tolerant systems with dependability assurance founded on formal models."

A synopsis of the researchers' more recent work includes a) research on assessment of system dependability, most directly relevant to safety cases; and b) research on achievement of diversity clarifying the roles of various diversity-seeking decisions in the design process. In a recent paper [118] updating the state of the practice for validating ultra-high dependability for software-based systems, Littlewood and Stringini noted the lack of progress in the field and suggested that more emphasis should be given to the role of confidence and epistemic uncertainty.

Page intentionally blank

6. KNOWLEDGE GAPS

Throughout most safety-critical, high-integrity industries such as nuclear power, quality processes and design measures that promote fault avoidance are employed to address the impact of single, random failures and reduce potential sources of common faults. These measures are effective for known hazards. However, they cannot guarantee freedom from CCF vulnerabilities. As noted, errors occur in spite of the best efforts of designers, developers, implementers, reviewers, testers, suppliers, and assessors. Hazard identification and design measures can minimize the potential for some sources of failure, but unanticipated and untested conditions can still pose a risk. Quality processes can detect and correct many implementation errors. However, as design complexity increases, some residual faults may remain undetected and persist as latent faults within the system. Consequently, the increasing use of highly complex digital technologies in modern I&C system designs poses additional concern that common systematic faults may persist undetected in spite of rigorous, high-quality life-cycle processes.

The experience with CCF in various applications, as described in Chapter 2, gives evidence to the significance of the safety threat that can arise if adequate mitigation of CCF vulnerabilities is not provided. The discussion of CCF in I&C systems identifies key elements of these vulnerabilities: latent faults and triggering conditions. The latent undetected faults that are at the heart of CCF vulnerabilities constitute a form of unknown hazards. The impact of unanticipated conditions or unexpected dependencies that can trigger a CCF poses another form of unknown hazards. There is a clear need for better definition of the nature of CCF in terms of the sources by which systematic faults (e.g., flaws, deficiencies, misunderstandings, mistakes, errors, defects) are introduced in I&C systems and the triggering mechanisms (e.g., common states, conditions, external influences) for common failure.

A rigorous identification of fault types and triggering conditions could support a thorough, systematic evaluation of CCF susceptibilities and allow for comprehensive determination of effective design measures to substantially reduce the CCF potential. Unfortunately, the complexity of the technology and the limited understanding of direct causal effects (especially for human-induced life-cycle-process-initiated faults) challenge the ability of designers and assessors to rely upon such an approach. As a result, more-subjective assessments and best-practice remediation are employed to provide reasonable assurance that adequate CCF mitigation is provided.

Chapter 3 documents the existing regulatory guidance and relevant standards for the nuclear power industry, as well as identifying the comparable guidance in several key safety critical industries. For the nuclear power industry, the use of diversity as a mitigating strategy to resolve CCF concerns supplements the quality assurance practices employed to satisfy safety requirements. In particular, diversity usage is specified in the design criteria for NPP safety systems as well as being required by regulation for NPPs. Regulatory guidance and international nuclear power standards cite diversity usage as the primary mitigation approach for addressing CCF threats. Guidance for key nonnuclear industries with high-integrity and/or safety-significant I&C applications also invokes diversity usage as a recommended CCF mitigation approach.

The application of diversity as a mitigation strategy is intended to minimize the prospect of common faults and the likelihood of common triggering conditions. Traditional diversity strategies (e.g., functional and signal diversity) have been commonly employed in the nuclear power industry for hardwired safety systems, with an emphasis on addressing commonalities and design-basis uncertainties. Unfortunately, because of the complexity of digital I&C systems and the associated inability to execute exhaustive testing, there is increased concern that the potential for latent systematic faults is greater in more fully digital I&C system architectures. In particular, since software (other than the simplest programs) in its coded state or its compiled machine language state cannot be proven to be without error, residual software faults represent a primary CCF concern. As a result, digital I&C systems receive particular emphasis in the assessment of CCF susceptibility and more extensive mitigation approaches are employed. However,

comprehensive guidance and objective acceptance criteria have not been established to resolve the effectiveness of defensive design measures, or specific types or combinations of diversity. There are many combinations of diversities that can be implemented and much uncertainty exists concerning which usage strategy is most effective and how much intentional diversity is enough. At best, only subjective criteria and best practices are available to provide a qualitative basis for determining whether adequate protection is provided. Basically, as design complexity increases, the challenge of providing sufficient evidence to establish reasonable assurance that the potential for CCF vulnerability has been adequately addressed becomes more difficult.

Several examples of CCF mitigation strategies from international NPPs and selected nonnuclear industries are described in Chapter 4. The approaches to diversity identified from international NPPs serve as representative examples of strategies that have been deemed to be acceptable. However, recent experience with I&C modernization in the United States and for new plant licensing internationally have shown that regulatory uncertainty remains about what constitutes acceptable CCF mitigation. In particular, there is still no definitive guidance available to objectively determine what kinds and how much diversity is sufficient to mitigate CCF vulnerabilities that may arise from digital safety system designs. Without quantitative measures of the efficacy of various mitigation techniques, ad hoc mitigation strategies persist. The resulting regulatory uncertainty leads to inhibition of technology modernization within the industry or drives costly, complex architectural solutions that may decrease dependability.

While the examples identified from nonnuclear industries are relevant because of the safety significance of the functions and the use of comparable technology, context differences in the usage domains limit their direct applicability. Thus, key insights can be derived from these examples of CCF mitigation practices to inform the development of a strategic approach suitable for the nuclear power industry. In particular, nonstandard architectural approaches (e.g., primary-checker configurations) with much more extreme functional diversity may warrant consideration by the nuclear power industry.

The NRC research described in Chapter 5 resulted in the development of baseline mitigation strategies that were consistent with acceptable practices based on implementation experience. The key assumption in that research is that qualitative assessment of the impact of diversity attributes and criteria, coupled with insights derived from established practice and key usage examples, provides a valid basis for developing diversity strategies to cope with the potential for CCF. These baseline strategies address considerations such as the effect of technology choices, the nature of CCF vulnerabilities, and the prospective impact of each diversity type. In particular, the impact of each attribute and criterion on the purpose, process, product, and performance aspects of diverse systems was evaluated based on an engineering assessment. Diversity usage tables and a diversity assessment spreadsheet tool were developed to aid in the evaluation of proposed mitigation strategies. The diversity assessment tool can also be employed for comparative analyses to assess the relative standing of a proposed alternate diversity strategy against the baseline strategies as well as established practices and common usage of the nuclear power and nonnuclear industries. This tool provides a systematic approach to evaluate proposed combinations of diversity criteria. However, the tool is based on subjective weighting of diversity effectiveness derived from engineering judgment and frequency of usage in the limited sample set. Thus, the scoring of strategies should be seen as a qualitative comparison, not an objective measure of CCF mitigation effectiveness.

The findings from the British diversity research program confirm that it cannot be conclusively demonstrated with mathematical rigor that intentional or forced diversity will result in independence of failure between systems. Additionally, the effect of diversity usage (individually or collectively applied) cannot be quantitatively determined at present. Basically, it was found that an extensive range of possible diversity-seeking decisions and their combinations are available but there is little definitive guidance for these choices. However, it is clear from qualitative evidence that diversity provides a dependability benefit (i.e., contributes to the mitigation of CCF vulnerabilities through overall system-level fault

tolerance) and is a reasonable response to CCF concerns. What are needed are objective measures of digital I&C system characteristics that give indication of the efficacy of various mitigation techniques.

An assessment of the findings from the investigation of the state of the practice for CCF mitigation points to knowledge gaps that should be resolved through further research. The foremost deficiency in knowledge relates to a fundamental understanding of the nature of CCF vulnerability in the context of the nuclear power application domain. In particular, a comprehensive identification is needed of the sources of systematic faults and the triggering conditions that impact safety-related functions in an NPP. These fault-trigger combinations should be mapped to functions and architectural elements (e.g., I&C system blocks) and related to hazards that could compromise plant safety. In addition, the various diversities and design measures that can mitigate CCF need to be related to the particular kinds of faults, triggers or fault-trigger combinations and to the corresponding failures that can result. The consequence would be better understanding of the impact of each diversity, the value of other defensive design measures, and the synergistic effect of combined mitigation techniques.

For example, the impact and benefits of diversity attributes and their associated criteria can be identified in terms of common fault sources (purpose and process), location of vulnerabilities (product), and common triggering conditions (performance). Essentially, the effect of each diversity attribute should be characterized according to the resultant capability to minimize the introduction of common faults, mitigate the presence of corresponding vulnerabilities, manage commonality in usage (i.e., execution), and reduce similarity in susceptibility to external factors. In turn, these diversity effects can be expressed in terms of minimized prospects for common systematic faults, reduced occurrence of concurrent execution profiles, and/or lessened likelihood of similar responses to external influences; thereby they provide a fundamental, scientific basis for decisions leading to diversified failure behavior among I&C systems.

The basic knowledge gap can be characterized as a need to establish the effectiveness of various mitigation techniques (e.g., diversity-seeking decisions or DSDs) in addressing specific classes of faults, triggers, or fault-trigger combinations. Essentially, a quantitative characterization of how DSDs diversify failure behavior for parallel systems would enable development of objective decision criteria and provide for a more comprehensive, systematic, and scientifically-based determination of what mitigation strategy would be most effective. The questions that need answers include the following: How effective is a particular DSD in resolving a particular CCF vulnerability? Which diversity or design measure is best for certain classes of CCF? How much diversity is adequate? What is the combined effect of multiple DSDs?

To resolve this knowledge gap, more thorough definition of each diversity attribute and defensive design measures should be developed. Various application domains have different characterizations of diversity. For example, functional diversity in the nuclear power industry corresponds to a parametric diversity of similar functions while the rail transportation industry employs much more fundamentally different functions as the basis for this form of diversity.

Models and metrics are needed to develop systematic methods, quantifiable measures, and objective criteria for evaluating CCF mitigation approaches. Various measures of I&C system characteristics (e.g., quality, reliability, performance, dependability) may be relevant for determining the effectiveness of diversity or design measures in mitigating CCF vulnerabilities. A thorough investigation of potential measures and models to support an aggregate indicator of diverse failure behavior is needed.

It is clear from the investigation of the state of the practice for CCF mitigation that a fundamentally sound basis for acceptable mitigation approaches is needed. Resolving uncertainties and regulatory burden concerning CCF vulnerability can promote elegant, optimal architectures for NPP I&C architectures with a well-defined safety basis, less imposed complexity, and, potentially, reduced cost. Achieving a science-based solution to this key technical challenge can benefit existing plants, new plants and advanced designs by removing an impediment to more extensive, effective use of digital technology.

The crosscutting research conducted under this project began with the investigation of the state of the practice, with a focus on recent nuclear plant experience and current approaches in nonnuclear industries. Based on the knowledge gap assessment, the next step is to develop a taxonomy characterizing the nature of CCF vulnerabilities. Models and metrics will be developed in subsequent research activities to establish methods for quantifying the impact of mitigation strategies involving diversity, defensive design measures and/or other DSDs that can be invoked during the life cycle of a system. Subsequently, case studies will be performed to test and evaluate the tools and techniques that are developed. This research will proceed toward a full-scope benchmark demonstration of a systematic approach to addressing CCF vulnerability based on verifiable, quantifiable methods.

7. REFERENCES

1. U.S. Nuclear Regulatory Commission, "Summary Of March 28-29, 2006, EPRI and NEI Workshop on Digital Instrumentation and Controls (I&C) and Control Room Licensing Issues," NRC **A**gencywide **D**ocuments **A**ccess and **M**anagement **S**ystem (ADAMS) Accession Number ML070590059, March 7, 2007.
2. International Atomic Energy Agency, *IAEA Safety Glossary*, Ed. 2.0, Vienna, Austria, 2006.
3. International Atomic Energy Agency, "Radiation Aspects of Design for Nuclear Power Plants," IAEA S-G-1.3, Vienna, Austria, 2005.
4. International Electrotechnical Commission, "Instrumentation and Control Systems Important to Safety—Requirements to Cope with Common Cause Failure (CCF)," IEC 62340, Geneva, Switzerland, 2008.
5. International Electrotechnical Commission, "Nuclear Power Plants—Instrumentation and Control Systems Important to Safety—Software Aspects for Computer-Based Systems Performing Category A Functions," IEC 60880, Ed. 2.0, Geneva, Switzerland, 2006.
6. Institute of Electrical and Electronics Engineers, "Standard Criteria for Safety Systems for Nuclear Power Generating Stations," IEEE Std 603-1991, Piscataway, New Jersey, 1991.
7. Institute of Electrical and Electronics Engineers, "Application of the Single-Failure Criterion to Nuclear Power Generating Station Safety Systems," IEEE Std. 379-2000, Piscataway, New Jersey, 2000.
8. *U.S. Code of Federal Regulations*, Title 10, Part 50, "Domestic Licensing of Production and Utilization Facilities," U.S. Nuclear Regulatory Commission, Washington, D.C..
9. Institute of Electrical and Electronics Engineers, "Criteria for Protection Systems for Nuclear Power Generating Stations," IEEE Std. 279-1971, Piscataway, New Jersey, 1971.
10. U.S. Nuclear Regulatory Commission, "Digital Computer Systems for Advanced Light-Water Reactors," SECY 91-292, Washington, D.C., September 26, 1991.
11. U.S. Nuclear Regulatory Commission, "Policy, Technical, and Licensing Issues Pertaining to Evolutionary and Advanced Light-Water Reactor (ALWR) Designs," SECY-93-087, Washington, D.C., April 2, 1993 (ADAMS Accession No. ML003708021).
12. U.S. Nuclear Regulatory Commission, "Policy, Technical, and Licensing Issues Pertaining to Evolutionary and Advanced Light-Water Reactor (ALWR) Designs," Staff Requirements Memorandum on SECY-93-087, Washington, D.C., July 21, 1993 (ADAMS Accession No. ML003708056).
13. Institute of Electrical and Electronics Engineers, "Standard Criteria for Digital Computers in Safety Systems of Nuclear Power Generating Stations," IEEE Std. 7-4.3.2-2003, Piscataway, New Jersey, 2003.
14. U.S. Nuclear Regulatory Commission, "Criteria for Use of Computers in Safety Systems of Nuclear Power Plants," Regulatory Guide 1.152, rev.2, Washington, D.C., 2006 (ADAMS Accession No. ML053070150).
15. U.S. Nuclear Regulatory Commission, "Guidance for Evaluation of Diversity and Defense-in-Depth in Digital Computer-Based Instrumentation and Control Systems," Branch Technical Position 7-19, Washington, D.C., 2007 (ADAMS Accession No. ML070550072).

16. U.S. Nuclear Regulatory Commission, *Standard Review Plan for the Review of Safety Analysis Reports for Nuclear Power Plants, Instrumentation and Controls*, NUREG-0800, Chapter 7, rev.5, Washington, D.C., 2007.

17. U.S. Nuclear Regulatory Commission, *Method for Performing Diversity and Defense-in-Depth Analyses of Reactor Protection Systems*, NUREG/CR-6303, December 1994.

18. Licensing of safety critical software for nuclear reactors. Common position of seven European nuclear regulators and authorized technical support organizations, HSE, 2007, http://www.hse.gov.uk/nuclear

19. International Electrotechnical Commission, "Nuclear power plants—Instrumentation and control for systems important to safety—General requirements for systems," IEC 61513, March 2001.

20. International Electrotechnical Commission, "Nuclear power plants—Instrumentation and control systems important for safety—Classification," IEC 61226, ed. 2.0, February 2005.

21. International Electrotechnical Commission, "Nuclear power plants—Instrumentation and control systems important to safety—Software aspects for computer-based systems performing category A functions," IEC 60880, ed. 2.0, May 2006.

22. International Electrotechnical Commission, "Nuclear power plants—Instrumentation and control systems important to safety—Separation," IEC 60709, November 2004.

23. International Electrotechnical Commission, "Nuclear power plants—Electrical equipment of the safety system—Qualification," IEC 60780, October 1998.

24. International Electrotechnical Commission, "Recommended practices for seismic qualification of electrical equipment of the safety system for nuclear generating stations," IEC 60980, June 1989.

25. International Electrotechnical Commission, "Electromagnetic compatibility (EMC)—Part 4-1: Testing and measurement techniques—Overview of IEC 61000-4 series," IEC 61000-4-1, October 2006.

26. National Aeronautics and Space Administration*, NASA General Safety Program Requirements*, NPR 8715.3, April 2007.

27. National Aeronautics and Space Administration, *Software Safety Standard*, NASA-STD-8719.13B, Change No. 1, July 2004.

28. National Aeronautics and Space Administration, *Safety Policy and Requirements For Payloads Using the Space Transportation System*, NSTS 1700.7B, Change No. 17, June 2004.

29. National Aeronautics and Space Administration, *Human-Rating Requirements for Space Systems*, NPR 8705.2A, February 2005.

30. National Aeronautics and Space Administration, *NASA Software Safety Guidebook*, NASA-GB-8719.13, March 2004.

31. Society of Automotive Engineers, "Certification Considerations for Highly-Integrated or Complex Aircraft Systems," SAE ARP 4754, SAE International, Warrendale, Pennsylvania, 1996.

32. Radio Technical Commission for Aeronautics, "Software Considerations in Airborne Systems and Equipment Certification," DO-178B, RTCA, Inc., 1992.

33. Radio Technical Commission for Aeronautics, "Design Assurance Guidance for Airborne Electronic Hardware," DO-254, RTCA, Inc., 2000.

34. *The Automation Systems and Instrumentation Dictionary*, 4th edition, ISA—The Instrumentation Systems and Automation Society, Research Triangle Park, North Carolina, 2003, p. 433.

35. OSHA, "Process safety management of highly hazardous chemicals," 29 CFR Part 1910.119, Washington, D.C. (1992).

36. I. Eckerman, The Bhopal Saga. Causes and consequences of the world's largest industrial disaster, Universities Press (India) Private Ltd., Hyderabad, 2004.

37. Center for Chemical Process Safety, *Guidelines for Safe Automation of Chemical Processes*, American Institute of Chemical Engineers, New York, New York, 1993.

38. Center for Chemical Process Safety, *Guidelines for Safe and Reliable Instrumented Protective Systems*, American Institute of Chemical Engineers, New York, New York, 2007.

39. Instrumentation, Systems and Automation Society, "Application of Safety Instrumented Systems (SIS) for the Process Industry," ANSI/ISA S84.01-1996, Research Triangle Park, North Carolina, 1996.

40. International Electrotechnical Commission, "Functional Safety: Safety Instrumented Systems for the Process Sector," IEC 61511, Geneva, Switzerland, 2003.

41. FRA, "Standards for Processor-Based Signal and Train Control Systems," 49 CFR Part 236, Subpart H, Washington, D.C. (2005).

42. European Railway Agency, "European Railway Agency Recommendation on the 1st set of Common Safety Methods," ERA-REC-02-2007-SAF, http://www.era.europa.eu/public/core/Safety/Documents/our%20products/cst-csm/ERA-REC-02-2007-SAF.pdf

43. European Committee for Electrotechnical Standardization, "Railway applications—Communications, signaling and processing systems—Software for railway control and protection systems," EN 50128, Brussels, Netherlands, 2001.

44. International Electrotechnical Commission, "Functional safety of electrical/electronic/programmable electronic safety-related systems," IEC 61508, Geneva, Switzerland, 1999.

45. U.S. Nuclear Regulatory Commission, *Advanced Reactor Licensing: Experience with Digital I&C Technology in Evolutionary Plants*, NUREG/CR-6842, April 2004.

46. International Atomic Energy Agency, *Modern Instrumentation and Control for Nuclear Power Plants: A Guidebook*, IAEA Technical Report Series No. 387, 1999.

47. J. D. White et al., *WTEC Panel Report on European Nuclear Instrumentation and Controls*, Loyola College, Baltimore, Maryland, 1991.

48. B. Fride, J. Y. Henry, and S. Manners, "Safety Assessment of Computerized Instrumentation and Control for Nuclear Power Plants," International Conference on Probabilistic Safety Assessment Methodology and Applications (PSA-95), November 26–30, 1995, Seoul, Korea.

49. J. R. Popovic and G. J. Hinton, "CANDU Computerized Safety System," presented at the *Advanced Computer Technology for the Power Industry*, Scottsdale, Arizona, EPRI, December 4–6, 1989.

50. *ACR-1000 Technical Summary: An Evolution of CANDU*, Atomic Energy of Canada Limited, Mississauga, Ontario, Canada.

51. J.-P. Burel, F. Dalik, K. Wagner, Miroslav RIS, and J.-P. Mauduit, "Modernization of I&C systems for the ANP Dukovany by the use of computer-based equipment," *CNRA/CSNI Workshop*

on Licensing and Operating Experience of Computer-Based I&C Systems Workshop Proceedings, NEA/CSNI/R(2002)1/Vol. 2, September 2001, Hluboka nad Vltavou, Czech Republic.

52. S. Kunito, "Construction and Operation Experience of Digitalized Safety Systems of Japanese ABWR," IAEA Technical Meeting on Common-Cause Failures in Digital Instrumentation and Control Systems of Nuclear Power Plants, June 20, 2006, Bethesda, Maryland.

53. S. Makino, "Operating Experience of Digital Safety Related System of Kashiwazaki-Kariwa Unit Nos. 6 and 7," CNRA/CSNI Workshop on Licensing and Operating Experience of Computer-Based I&C Systems Workshop Proceedings, NEA/CSNI/R(2002)1/Vol. 2, Hluboka nad Vltavou, Czech Republic.

54. Japan Electric Association Guideline, "Application Criteria for Programmable Digital Computer System in Safety–Related System of Nuclear Power Plants," JEAG 4609, 1999.

55. C-F. Chuang and Y-B. Chen, "Regulatory Overview of Digital I&C in Taiwan Lungmen Project," NRC 19th Annual Regulatory Information Conference, March 13–15, 2007, Rockville, Maryland.

56. C-K. Lee, "The Network Architecture and Site Test of DCIS in Lungmen Nuclear Power Station," 5th International Topical Meeting on Nuclear Plant Instrumentation Control and Human Machine Interface Technology (NPIC & HMIT 2006), November 12–16, 2006, Albuquerque, New Mexico.

57. J. Hyvärinen, "Presentation Slides: OL3 I&C Review Status," ASN/IRSN-NRC-STUK Meeting, March 22, 2007, Paris, France.

58. U.S. EPR Pre-Application Review Meeting: U.S. EPR Digital Protection System Topical Report, presentation by AREVA NP, Inc., to the NRC, March 1, 2007, Rockville, Maryland.

59. G. W. Remley, B. M. Cook, and P. A. Loftus, "Sizewell B Integrated Control and Instrumentation System: A Vision Becomes Reality," Conference Record of the 1992 IEEE Nuclear Science Symposium and Medical Imaging Conference, Vol. 2, Orlando, Florida, Oct. 25–31, 1992, pp. 736–738.

60. A. Johnson, "The implementation of Sizewell B automatic control systems," International Conference on Electrical and Control Aspects of the Sizewell B PWR, London, UK, Sept. 14–15, 1992, pp. 143–148.

61. C. Percival and D. Bradbury, "The engineering specification, design and implementation of the Sizewell B reactor secondary protection system," International Conference on Electrical and Control Aspects of the Sizewell B PWR, London, UK, Sept. 14–15, 1992, pp. 232–244.

62. G. B. Moutrey and G. Remley, "Sizewell B power station primary protection system design application overview," International Conference on Electrical and Control Aspects of the Sizewell B PWR, London, UK, Sept. 14–15, 1992, pp. 221–231.

63. W. C. Gangloff and C. L. Werner, "I&C Modernization for VVER Reactors," *IEEE Transactions on Nuclear Science*, **40**(4), 819–825 (August 1993).

64. P. Závodsky, "Independent Assessment of the Temelín Safety System Software," *CNRA/CSNI Workshop on Licensing and Operating Experience of Computer-Based I&C Systems Workshop Proceedings*, NEA/CSNI/R(2002)1/Vol. 1, September 2001, Hluboka nad Vltavou, Czech Republic.

65. R. G. Orendi, "Human Factors Experience in Designing a Modern Control Room for a VVER-1000 Nuclear Plant," IEEE Sixth Annual Human Factors Meeting, 1997, Orlando, Florida.

66. H. S. Park, "Regulatory Review of the Test Features of the Digital Plant Protection System for Ulchin Nuclear Power Plant Units 5 & 6," 4th International Topical Meeting on Nuclear Plant

Instrumentation Control and Human Machine Interface Technology (NPIC & HMIT 2004), September 19–22, 2004, Columbus, Ohio.

67. C. H. Jeong, "Suitability Review of Reliability Analysis of the Digital Plant Protection System for Ulchin Nuclear Power Plant Units 5 & 6," 4th International Topical Meeting on Nuclear Plant Instrumentation Control and Human Machine Interface Technology (NPIC & HMIT 2004), September 19–22, 2004, Columbus, Ohio.

68. J. E. Tomayko, *Computers in Spaceflight: The NASA Experience*, NASA CR 182505, March 1988.

69. J. E. Tomayko, Computers Take Flight: A History of NASA's Pioneering Digital Fly-By-Wire Project, NASA SP-2000-4224, April 2000.

70. J. F. Hanaway and R. W. Moorehead, *Space Shuttle Avionics System*, NASA SP-504, January 1989.

71. J. E. Tomayko, Computers Take Flight: A History of NASA's Pioneering Digital Fly-By-Wire Project, NASA SP-2000-4224, April 2000.

72. J. F. Hanaway and R. W. Moorehead, *Space Shuttle Avionics System*, NASA SP-504, January 1989.

73. P. Ladkin, "Excerpt from the Case Study of the Space Shuttle Primary Control System," http://www.rvs.uni-bielefeld.de/publications/Incidents/DOCS/ComAndRep/Ariane/shuttle.html

74. National Aeronautics and Space Administration, *International Space Station Familiarization*, NASA TD9702A, Houston Texas, July 1998.

75. P. Robinson et al., "Applying Model-Based Reasoning to the FDIR of the Command & Data Handling Subsystem of the International Space Station," *Proceedings of the 7th International Symposium on Artificial Intelligence, Robotics and Automation in Space* (ISAIRAS03), Detroit, Michigan, May 19–23, 2003.

76. "ISS Status Report: ISS 01-11," May 2001, http://www.astronautix.com/details/iss52449.htm

77. "Ailing Computers Critical to ISS," May 2001, http://www.space.com/news/spacestation/sts100_iss_computers_010501.html

78. I. Moir and A. Seabridge, Aircraft Systems: Mechanical, electrical, and avionics subsystems integration, 3rd Edition, John Wiley & Sons, Ltd., 2008.

79. J. Voas, A. Ghosh, F. Charron, and L. Kassab, "Reducing uncertainty about common-mode failures," *Proc. of Eighth International Symposium on Software Reliability Engineering*, Albuquerque, New Mexico, November 2–5, 1997, pp. 308–319.

80. P. Traverse, "Dependability of Digital Computers on Board Airplanes," *Dependable Computing for Critical Applications*, Vol. 4, A. Avizienis and J. C. Laprie, eds., 1991, pp. 134–152.

81. J. E. Tomayko, "Computers Take Flight: A History of NASA's Pioneering Digital Fly-By-Wire Project," NASA SP-2000-4224, Washington, D.C., 2000.

82. G. Mauri, "Integrating Safety Analysis Techniques, Supporting Identification of Common Cause Failure," Ph.D. Dissertation, The University of York, September 2000.

83. P. Traverse, I. Lacaze and J. Souyris, "Airbus fly-by-wire: A total approach to dependability," *IFIP World Computer Congress*, Toulouse, France, August 2004.

84. D. Briere and P. Traverse, “AIRBUS A320/A330/A340 Electrical Flight Controls: A Family of Fault-Tolerant Systems,” *Digest of Papers FTCS-23: The Twenty-Third International Symposium on Fault-Tolerant Computing*, June 1993, pp. 616–623.

85. D. P. Siewiorek and P. Narasimhan, “Fault-Tolerant Architectures for Space and Avionics Applications,” http://ic-www.arc.nasa.gov/projects/ishem/Papers/Siewiorek_Fault_Tol.pdf

86. Health and Safety Commission, “The Use of Computers in Safety-Critical Applications - Final Report of the Study Group on the Safety of Operational Computer Systems,” London, UK, November 1998.

87. W. Torres-Pomales, “Software fault tolerance: A tutorial,” *NASA Technical Report NASA/ TM-2000-210616*, October 2000.

88. L. Lamport, R. Shostak, and M. Peace, “The Byzantine Generals Problem,” *ACM Trans. on Programming Languages and Systems*, Vol. 4, No. 3, July 1982.

89. Y. C. Yeh, “Safety critical avionics for the 777 primary flight controls system,” *IEEE Conference on Digital Avionics Systems*, Vol. 1, Daytona Beach, Florida, October 2001, pp. 1–11.

90. Y. C. Yeh, “Dependability of the 777 Primary Flight Control System,” *Proc. of the 5th IFIP Int’l Working Conf. on Dependable Computing for Critical Applications (DCCA-5)*, Urbana-Champaign, Illinois, September 1995.

91. Y. C. Yeh, “Triple-Triple Redundant 777 Primary Flight Computer,” *Proc. of the 1996 IEEE Aerospace Applications Conference*, Vol. 1, Aspen, Colorado, February, 1996, pp. 293–307.

92. R. W. Pratt, Flight Control Systems - Practical Issues in Design and Implementation. Institution of Engineering and Technology, 2000.

93. J. H. Lala and R. E. Harper, “Architectural Principles for Safety-Critical Real-Time Applications,” *Proceedings of the IEEE*, **82**(1), January 1994, pp. 25–40.

94. Y. C. Yeh, “Design Considerations in Boeing 777 Fly-By-Wire Computers,” *3rd IEEE High-Assurance Systems Engineering Symposium (HASE),* Washington, D.C., IEEE Computer Society Press, 1998, pp. 64–73.

95. W. L. Heimerdinger and C. B. Weinstock, *A Conceptual Framework for System Fault Tolerance*, Technical Report CMU/SEI-92-TR-033 (ESC-TR-92-033), October 1992.

96 H. Kantz and C. Koza, “The ELEKTRA Railway Signalling-System: Field Experience with an Actively Replicated System with Diversity,” *Twenty-Fifth International Symposium Fault-Tolerant Computing (FTCS-25)*, Pasadena, California, June 27–30, 1995, pp. 453–458.

97. A. Denault, “Fault Tolerance in Railway Signalling System: A study of the Elektra Interlocking Systems,” http://www.adinfo.qc.ca/alex/wp-content/Elektra/elektra.pdf

98. G. Wirthumer, “Votrics—Fault Tolerance Realised in Software,” *IFAC Proceedings SAFECOMP 89*, Vienna, Austria, December 1989, pp. 135–140.

99. D. Powell et al., “Architectural Approaches for using COTS Components in Critical Applications,” www.laas.fr/~dpowell/slides/0005%20EWDC11.pdf

100. P. Forin, “Vital Coded Microprocessor: Principles and Application for Various Transit Systems,” *Proc. IFAC-GCCT*, Paris, France, September 1989, pp. 79–84.

101. C. Hennebert and G. Guiho, “SACEM: a fault tolerant system for train speed control,” *Twenty-Third International Conf. on Fault-Tolerant Computing (FTCS-23)*, Toulouse, France, 1993.

102. D. Dollé, “Vital software: Formal method and coded processor,” Third Embedded Real Time Conference (ERTS 2006), Toulouse, France, January 25–27, 2006.

103. G. Guiho and C. Hennebert, “SACEM Software Validation,” 12th ICSE, IEEE Computer Society Press, Mars 1990, pp. 186–191.

104. J. A. Profeta III et al., “Safety-Critical Systems Built with COTS,” *Computer*, **29(**11), 54–60 (November 1996).

105. D. T. Smith et al., “An Algorithm Based Fault Tolerance Technique for Safety-Critical Applications,” *1997 Proceedings of the Annual Reliability and Maintainability Symposium*, Philadelphia, Pennsylvania, January 1997.

106. B. Johnson, University of Virginia, private communication with R. T. Wood, October 2007.

107. U.S. Nuclear Regulatory Commission, *Diversity Strategies for Nuclear Power Plant Instrumentation and Control Systems*, NUREG/CR-7007, February 2010.

108. B. Littlewood et al., “DISPO Project at City University,” Centre for Software Reliability, City University, London, 2006.

109. B. Littlewood, P. Popov, and L. Strigini, “DISPO project: A summary of CSR work on modelling of diversity,” Centre for Software Reliability, City University, London, UK, 2006.

110. J. C. Knight and N. G. Leveson, “Experimental evaluation of the assumption of independence in multiversion software,” *IEEE Trans Software Engineering*, **12**(1), 96–109 (1986).

111. B. Littlewood and L. Strigini, “A discussion of practices for enhancing diversity in software designs,” DISPO LS-DI-TR-04, Centre for Software Reliability, City University, London, 2000.

112. B. Littlewood and D. R. Miller, “Conceptual Modelling of Coincident Failures in Multi-Version Software,” *IEEE Transactions on Software Engineering*, **15**(12), 1596–1614 (1989).

113. B. Littlewood, P. Popov, and L. Strigini, “A note on modelling functional diversity,” *Reliability Engineering and System Safety*, vol. 66, no. 1, pp. 93–95, 1999.

114. G. E. Migneault, “The Cost of Software Fault Tolerance,” NASA Technical Memoranum 84586, NASA Langley Research Center, Hampton, Virginia, September 1982.

115. J.C. Laprie, J. Arlat, C. Beounes, and K. Kanoun, “Definition and analysis of hardware and software fault-tolerant architectures,” *Computer*, **23**(7), 39–51 (1990).

116. U. Voges, “Software diversity,” *Reliability Engineering System Safety* **43**(2), 103–110 (1994).

117. B. Littlewood, P. Popov, L. Strigini, “Design Diversity: an Update from Research on Reliability Modelling,” *Proc of the 9th Safety-critical Systems Symposium*, Bristol 2001.

118. B. Littlewood and L. Strigini, “ ‘Validation of ultra-high dependability…’ – 20 years on,” *Safety Systems, The Safety-Critical Systems Club Newletter*, 2011.