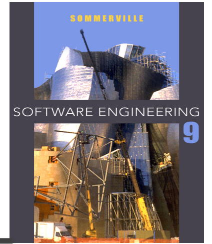

Chapter 13 – Dependability engineering

Topics covered

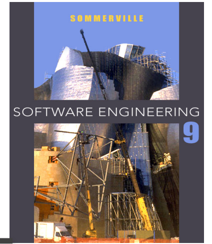

- Redundancy and diversity
  - Fundamental approaches to achieve fault tolerance.
- Dependable processes
  - How the use of dependable processes leads to dependable systems
- Dependable systems architectures
  - Architectural patterns for software fault tolerance
- Dependable programming
  - Guidelines for programming to avoid errors.

Software dependability

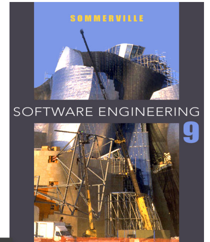

- In general, software customers expect all software to be dependable. However, for non-critical applications, they may be willing to accept some system failures.
- Some applications (critical systems) have very high dependability requirements and special software engineering techniques may be used to achieve this.
  - Medical systems
  - Telecommunications and power systems
  - Aerospace systems

Dependability achievement

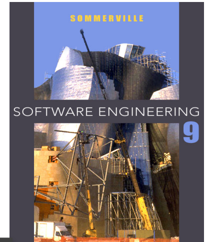

- Fault avoidance
  - The system is developed in such a way that human error is avoided and thus system faults are minimised.
  - The development process is organised so that faults in the system are detected and repaired **before** delivery to the customer.
- Fault detection
  - **Verification** and **validation** techniques are used to discover and remove faults in a system before it is deployed.
- Fault tolerance
  - The system is designed so that faults in the delivered software do not result in system failure.

The increasing costs of residual fault removal

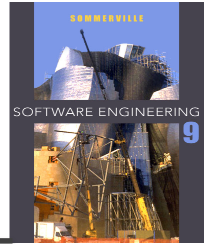

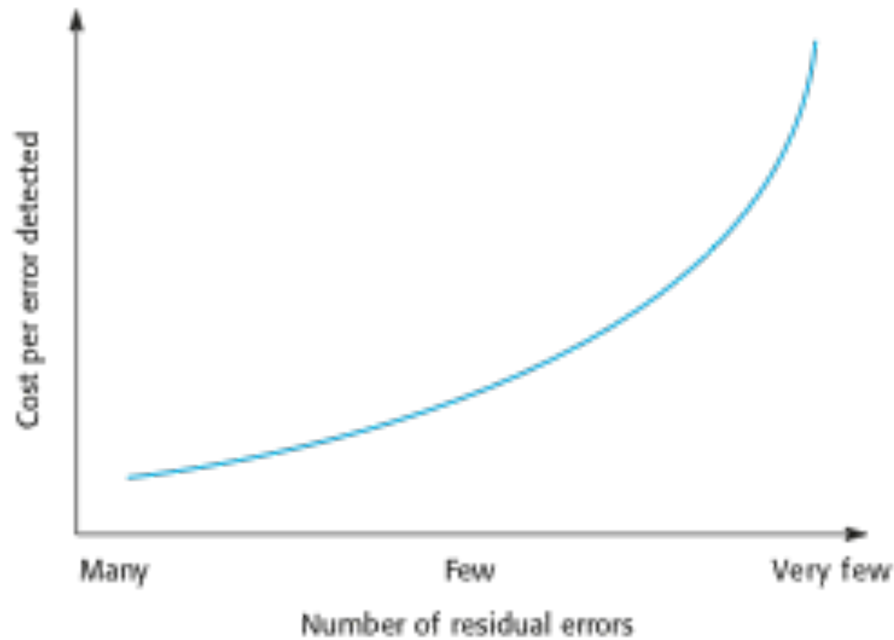

Regulated systems

- Many critical systems are regulated systems, which means that their use must be approved by an external regulator before the systems go into service.
  - Nuclear systems
  - Air traffic control systems
  - Medical devices
- A safety and dependability case has to be approved by the regulator. Therefore, critical systems development has to create **the evidence** to convince a regulator that the system is dependable, safe and secure.

Diversity and redundancy

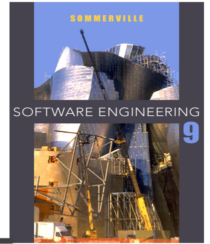

- Redundancy
  - Keep more than 1 version of a critical component available so that if one fails then a backup is available.
- Diversity
  - Provide the same functionality in different ways so that they will not fail in the same way.
- However, adding diversity and redundancy adds complexity and this can increase the chances of error.
- Some engineers advocate simplicity and extensive V & V is a more effective route to software dependability.

Diversity and redundancy examples

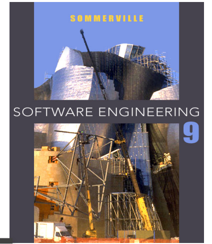

- Redundancy. Where availability is critical (e.g. in e-commerce systems), companies normally keep backup servers and switch to these automatically if failure occurs.
- Diversity. To provide resilience against external attacks, different servers may be implemented using different operating systems (e.g. Windows and Linux)

Process diversity and redundancy

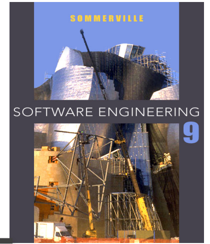

- Process activities, such as validation, should not depend on a single approach, such as testing, to validate the system
- Rather, multiple different process activities the complement each other and allow for cross-checking help to avoid process errors, which may lead to errors in the software

Dependable processes

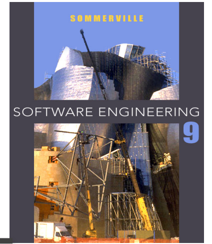

- To ensure a minimal number of software faults, it is important to have a well-defined, repeatable software process.
- A well-defined repeatable process is one that does not depend entirely on individual skills; rather can be enacted by **different** people.
- For fault detection, it is clear that the process activities should include significant effort devoted to verification and validation.

Attributes of dependable processes

Process characteristic	Description
Documentable	The process should have a defined process model that sets out the activities in the process and the documentation that is to be produced during these activities.
Standardized	A comprehensive set of software development standards covering software production and documentation should be available.
Auditable	The process should be understandable by people apart from process participants, who can check that process standards are being followed and make suggestions for process improvement.
Diverse	The process should include redundant and diverse verification and validation activities.
Robust	The process should be able to recover from failures of individual process activities.

Validation activities

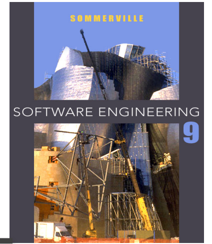

- ✧ Requirements reviews.
- ✧ Requirements management.
- ✧ Formal specification.
- ✧ System modelling
- ✧ Design and code inspection.
- ✧ Static analysis.
- ✧ Test planning and management.
- ✧ Change management, discussed in Chapter 25, is also essential.

Fault tolerance

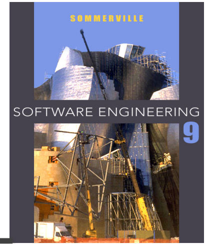

- In critical situations, software systems must be fault tolerant.
- Fault tolerance is required where there are high availability requirements or where system failure costs are very high.
- Fault tolerance means that the system can continue in operation in spite of software failure.
- Even if the system has been proved to conform to its specification, it must also be fault tolerant as there may be specification errors or the validation may be incorrect.

Dependable system architectures

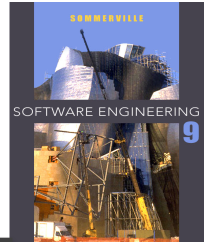

- Dependable systems architectures are used in situations where fault tolerance is essential. These architectures are generally all based on redundancy and diversity.
- Examples of situations where dependable architectures are used:
  - Flight control systems, where system failure could threaten the safety of passengers
  - Reactor systems where failure of a control system could lead to a chemical or nuclear emergency
  - Telecommunication systems, where there is a need for 24/7 availability.

Protection systems

✧ A specialized system that is associated with some other control system, which can take emergency action if a failure occurs.

- System to stop a train if it passes a red light
- System to shut down a reactor if temperature/pressure are too high

✧ Protection systems independently monitor the controlled system and the environment.

✧ If a problem is detected, it issues commands to take emergency action to shut down the system and avoid a catastrophe.

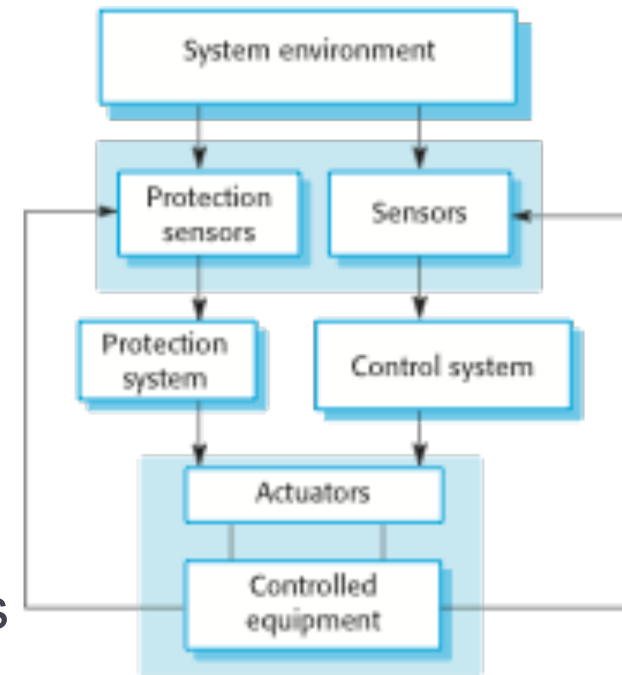

Protection system functionality

SOMMERVILLE
SOFTWARE ENGINEERING 9

- Protection systems are redundant because they include monitoring and control capabilities that replicate those in the control software.
- Protection systems should be diverse and use different technology from the control software.
- They are simpler than the control system so more effort can be expended in validation and dependability assurance.
- Aim is to ensure that there is a low probability of failure on demand for the protection system.

Self-monitoring architectures

- Multi-channel architectures where the system monitors its own operations and takes action if inconsistencies are detected.
- The same computation is carried out on each channel and the results are compared. If the results are identical and are produced at the same time, then it is assumed that the system is operating correctly.
- If the results are different, then a failure is assumed and a failure exception is raised.

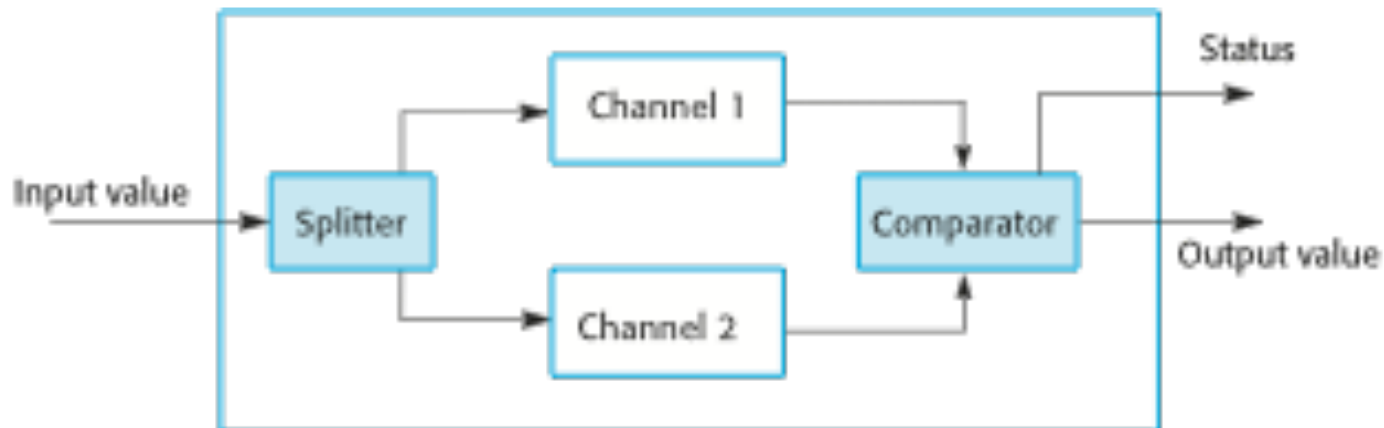

Self-monitoring systems

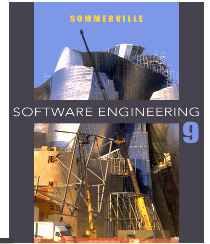

- ✧ Hardware in each channel has to be diverse so that common mode hardware failure will not lead to each channel producing the same results.
- ✧ Software in each channel must also be diverse, otherwise the same software error would affect each channel.
- ✧ If high-availability is required, you may use several self-checking systems in parallel.
  - This is the approach used in the Airbus family of aircraft for their flight control systems.

Airbus flight control system architecture

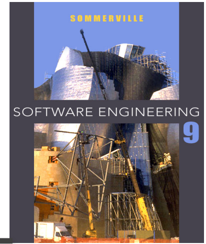

Input value

Primary flight control system 1

Input

Splitter

Channel 1

Channel 2

Comparator

Status

Output

Filter

Output

Primary flight control system 2

Status

Output

Filter

Primary flight control system 3

Status

Output

Filter

Secondary flight control system 1

Splitter

Channel 1

Channel 2

Comparator

Status

Output

Filter

Secondary flight control system 2

Status

Output

Filter

- The Airbus FCS has 5 separate computers, any one of which can run the control software.
- Extensive use has been made of diversity
  - Primary systems use a different processor from the secondary systems.
  - Primary and secondary systems use chipsets from different manufacturers.
  - Software in secondary systems is less complex than in primary system – provides only critical functionality.
  - Software in each channel is developed in different programming languages by different teams.
  - Different programming languages used in primary and secondary systems.

N-version programming

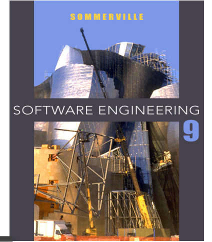

- Multiple versions of a software system carry out computations at the same time. There should be an odd number of computers involved, typically 3.
- The results are compared using a voting system and the majority result is taken to be the correct result.
- Approach derived from the notion of triple-modular redundancy, as used in hardware systems.

Hardware fault tolerance

- Depends on triple-modular redundancy (TMR).
- There are three replicated identical components that receive the same input and whose outputs are compared.
- If one output is different, it is ignored and component failure is assumed.
- Based on most faults resulting from component failures rather than design faults and a low probability of simultaneous component failure.

Triple modular redundancy

N-version programming

✧ The different system versions are designed and implemented by different teams. It is assumed that there is a low probability that they will make the same mistakes. The algorithms used should but may not be different.

✧ There is some empirical evidence that teams commonly misinterpret specifications in the same way and chose the same algorithms in their systems.

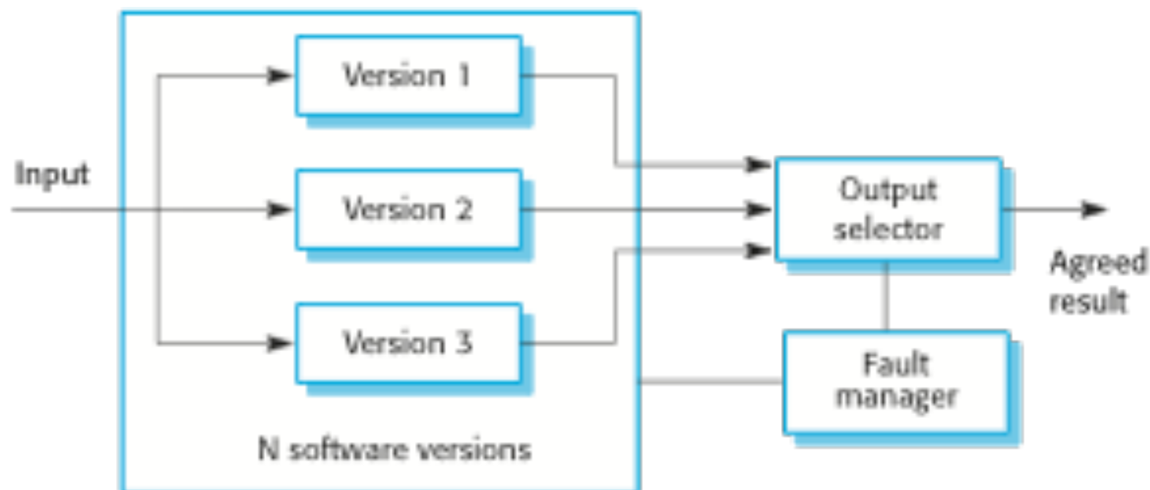

Software diversity

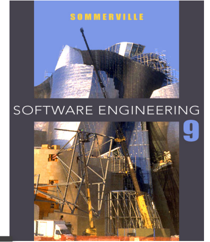

- Approaches to software fault tolerance depend on software diversity where it is assumed that different implementations of the same software specification will fail in different ways.
- It is assumed that implementations are (a) independent and (b) do not include common errors.
- Strategies to achieve diversity
  - Different programming languages
  - Different design methods and tools
  - Explicit specification of different algorithms

Problems with design diversity

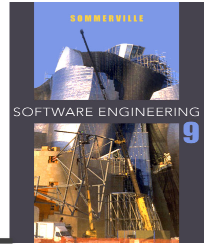

- Teams are not culturally diverse so they tend to tackle problems in the same way.
- Characteristic errors
  - Different teams make the same mistakes. Some parts of an implementation are more difficult than others so all teams tend to make mistakes in the same place;
  - Specification errors; if there is an error in the specification then this is reflected in all implementations;
  - This can be addressed to some extent by using multiple specification representations.

Specification dependency

- Both approaches to software redundancy are susceptible to specification errors. If the specification is incorrect, the system could fail
- This is also a problem with hardware but software specifications are usually more complex than hardware specifications and harder to validate.
- This has been addressed in some cases by developing **separate software specifications** from the same user specification.

Improvements in practice

- In principle, if diversity and independence can be achieved, multi-version programming leads to very significant improvements in reliability and availability.
- In practice, observed improvements are much less significant but the approach seems leads to reliability improvements of between 5 and 9 times.
- The key question is whether or not such improvements are worth the considerable extra development costs for multi-version programming.

Dependable programming

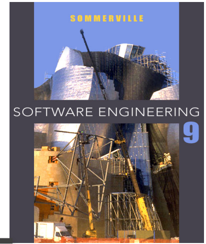

- Good programming practices can be adopted that help reduce the incidence of program faults.

Dependable programming guidelines
1. Limit the visibility of information in a program
2. Check all inputs for validity
3. Provide a handler for all exceptions
4. Minimize the use of error-prone constructs
5. Provide restart capabilities
6. Check array bounds
7. Include timeouts when calling external components
8. Name all constants that represent real-world values

Good practice guidelines for dependable programming

Control the visibility of information in a program

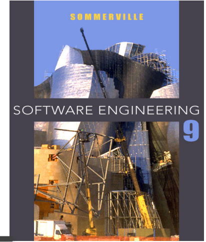

- Program components should only be allowed access to data that they need for their implementation.
- This means that <u>accidental corruption</u> of parts of the program state by these components is <u>impossible</u>.
- You can control visibility by using abstract data types where the data representation is private and you only allow access to the data through predefined operations such as get () and put ().

Check all inputs for validity

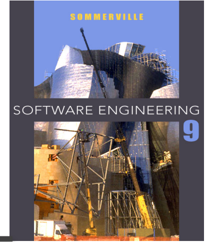

- All program take inputs from their environment and make assumptions about these inputs.
- However, program specifications rarely define what to do if an input is not consistent with these assumptions.
- Consequently, many programs behave unpredictably when presented with unusual inputs and, sometimes, these are threats to the security of the system.
- Consequently, you should always check inputs before processing against the assumptions made about these inputs.

Validity checks

- ✧ Range checks
  - Check that the input falls within a known range.
- ✧ Size checks
  - Check that the input does not exceed some maximum size e.g. 40 characters for a name.
- ✧ Representation checks
  - Check that the input does not include characters that should not be part of its representation e.g. names do not include numerals.
- ✧ Reasonableness checks
  - Use information about the input to check if it is reasonable rather than an extreme value.

Provide a handler for all exceptions

- A program exception is an error or some unexpected event such as a power failure.
- Exception handling constructs allow for such events to be handled without the need for continual status checking to detect exceptions.
- Using normal control constructs to detect exceptions needs many additional statements to be added to the program. This adds a significant overhead and is potentially error-prone.

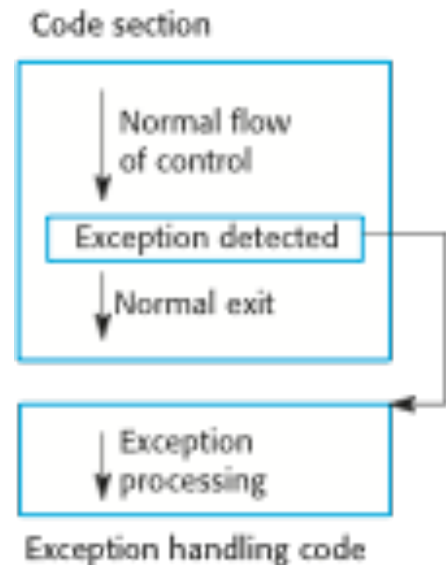

Exception handling

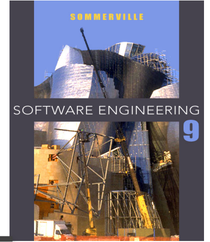

- ✧ Three possible exception handling strategies
  - Signal to a calling component that an exception has occurred and provide information about the type of exception.
  - Carry out some alternative processing to the processing where the exception occurred. This is only possible where the exception handler has enough information to recover from the problem that has arisen.
  - Pass control to a run-time support system to handle the exception.
- ✧ Exception handling is a mechanism to provide some fault tolerance

Minimize the use of error-prone constructs

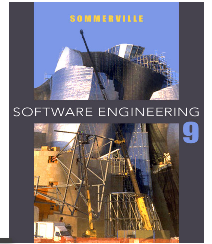

- ✧ Program faults are usually a consequence of human error because programmers lose track of the relationships between the different parts of the system
- ✧ This is exacerbated by error-prone constructs in programming languages that are inherently complex or that don't check for mistakes when they could do so.
- ✧ Therefore, when programming, you should try to avoid or at least minimize the use of these error-prone constructs.

Error-prone constructs

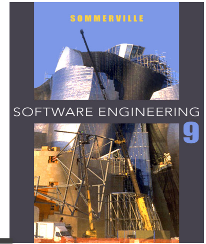

- Unconditional branch (goto) statements
- Floating-point numbers
  - Inherently imprecise. The imprecision may lead to invalid comparisons.
- Pointers
  - Pointers referring to the wrong memory areas can corrupt data. Aliasing can make programs difficult to understand and change.
- Dynamic memory allocation
  - Run-time allocation can cause memory overflow.

Error-prone constructs

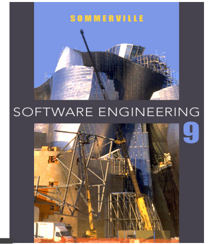

- ✧ Parallelism
  - Can result in subtle timing errors because of unforeseen interaction between parallel processes.
- ✧ Recursion
  - Errors in recursion can cause memory overflow as the program stack fills up.
- ✧ Interrupts
  - Interrupts can cause a critical operation to be terminated and make a program difficult to understand.
- ✧ Inheritance
  - Code is not localised. This can result in unexpected behaviour when changes are made and problems of understanding the code.

Error-prone constructs

- Aliasing
  - Using more than 1 name to refer to the same state variable.
- Unbounded arrays
  - Buffer overflow failures can occur if no bound checking on arrays.
- Default input processing
  - An input action that occurs irrespective of the input.
  - This can cause problems if the default action is to transfer control elsewhere in the program. In incorrect or deliberately malicious input can then trigger a program failure.

Provide restart capabilities

- For systems that involve long transactions or user interactions, you should always provide a restart capability that allows the system to restart after failure without users having to redo everything that they have done.
- Restart depends on the type of system
  - Keep copies of forms so that users don't have to fill them in again if there is a problem
  - Save state periodically and restart from the saved state

Check array bounds

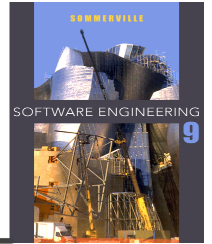

- In some programming languages, such as C, it is possible to address a memory location outside of the range allowed for in an array declaration.
- This leads to the well-known 'bounded buffer' vulnerability where attackers write executable code into memory by deliberately writing beyond the top element in an array.
- If your language does not include bound checking, you should therefore always check that an array access is within the bounds of the array.

Include timeouts when calling external components

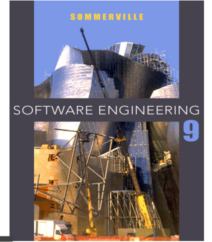

- ✧ In a distributed system, failure of a remote computer can be 'silent' so that programs expecting a service from that computer may never receive that service or any indication that there has been a failure.
- ✧ To avoid this, you should always include timeouts on all calls to external components.
- ✧ After a defined time period has elapsed without a response, your system should then assume failure and take whatever actions are required to recover from this.

Name all constants that represent real-world values

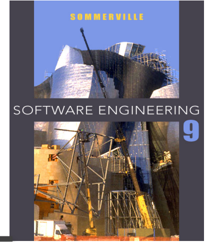

- Always give constants that reflect real-world values (such as tax rates) names rather than using their numeric values and always refer to them by name
- You are less likely to make mistakes and type the wrong value when you are using a name rather than a value.
- This means that when these 'constants' change (for sure, they are not really constant), then you only have to make the change in one place in your program.

Key points

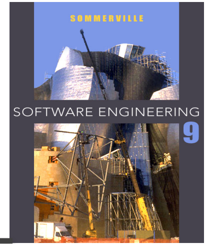

- ✧ Dependability in a program can be achieved by avoiding the introduction of faults, by detecting and removing faults before system deployment, and by including fault tolerance facilities.
- ✧ The use of redundancy and diversity in hardware, software processes and software systems is essential for the development of dependable systems.
- ✧ The use of a well-defined, repeatable process is essential if faults in a system are to be minimized.
- ✧ Dependable system architectures are system architectures that are designed for fault tolerance. Architectural styles that support fault tolerance include protection systems, self-monitoring architectures and N-version programming.

Key points

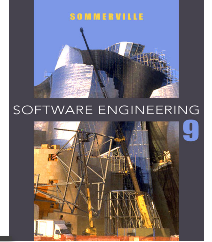

- Software diversity is difficult to achieve because it is practically impossible to ensure that each version of the software is truly independent.
- Dependable programming relies on the inclusion of redundancy in a program to check the validity of inputs and the values of program variables.
- Some programming constructs and techniques, such as goto statements, pointers, recursion, inheritance and floating-point numbers, are inherently error-prone. You should try to avoid these constructs when developing dependable systems.