

EECS 507: Introduction to Embedded Systems Research Reliability, Testing, and Formal Methods

Robert Dick

University of Michigan

Outline

1. Deadlines and announcements
2. Motivation
3. Sources of reliability problems
4. Fault detection and correction
5. Testing and formal methods
6. Thermal analysis

Deadlines and announcements

11 March: Midterm exam.

Will have project presentations near end of semester.

Will read a few other general papers associated with them, then.

For example, security.

Switch to embedded machine learning after midterm exam.

16 March: R. P. Dick, L. Shang, M. Wolf, and S.-W. Yang, “**Embedded Intelligence in the Internet-of-Things**,” *IEEE Design & Test of Computers*, Dec. 2019.

Outline

1. Deadlines and announcements
2. **Motivation**
3. Sources of reliability problems
4. Fault detection and correction
5. Testing and formal methods
6. Thermal analysis

Example medical device

Real-time.

Extreme consequences for failures.

Coupled physical, computation systems, which may be complex.

Subject to use by experts in other areas who are naïve about the technology employed.

Require high rate signal processing and decision making.

Tight power consumption constraints.

Regulatory process stressing informal testing techniques.

Therac-25

AECL Medical (Medical device company owned by Canadian government).

Designed from start for software control using PDP 11.

Hardware interlocks eliminated, relied on software.

Custom real-time executive.

- No explicit synchronization.
- Non-atomic test and set.
- Shared memory.

User interface making failure difficult to detect, and encouraging ignoring warnings.

Malfunction 54

Common cryptic error messages, most of which did not indicate danger to patients.

Operators learned to ignore warnings.

Operators taught that there are “so many safety mechanisms” that it is virtually impossible to harm a patient. Milgram + Monderman.

Consequences and responses

Holes burned in people.

Paralyzed arm, legs, vocal cords, and lung, followed by death.

Destroyed brain.

Victims were told by technicians that this is “impossible”, at least until their skin fell off.

Blamed on mechanical errors, additional checks applied, deemed “corrected”.

Causes?

System checked for edits only if bending magnet flag is set.

Edits made within 8 seconds, while bending magnets were being adjusted, not detected.

Real causes?

“We know there are many safety codes, guides, and regulations to guide them.” Rawlinson

Outline

1. Deadlines and announcements
2. Motivation
3. Sources of reliability problems
4. Fault detection and correction
5. Testing and formal methods
6. Thermal analysis

Types of reliability

Algorithm correctness

Does the specification have the desired properties?

Robustness in the presence of transient faults

Can the system continue to operate correctly despite temporary errors?

Robustness in the presence of permanent faults and process variation

Can the system continue to operate correctly in the presence of permanent errors?

Definitions

There are various definitions of these terms, mostly depending on domain such as software development or hardware design.

Fault: out-of-specs operation of component.

Failure: out-of-specs operation of system.

Error: the specific cause of a failure.

Not all faults result in failures, even without correction.

Specification errors

The specification doesn't match the needs of the customer / intention of the designer.

Common.

Hard to identify especially for informal specifications with implicit assumptions.

Implementation errors

The implementation doesn't match the specification.

Can sometimes be automatically checked if specification is formal.

Common.

Hard to manually identify. Complex to automatically identify.

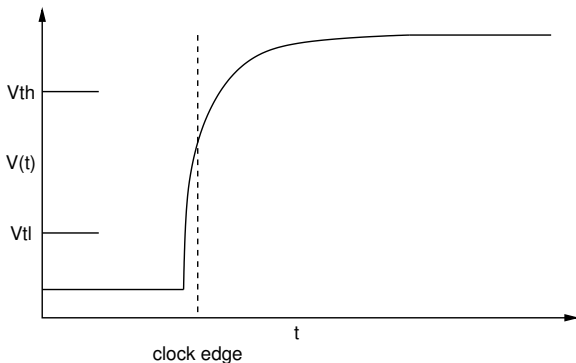
An error per ~ 100 lines of code is considered a very good rate.

Transient and intermittent component faults

Transient faults: occur briefly and generally without bias toward future faults in the same component.

- Timing faults.
- R drop.
- dl/dt .
- C or L crosstalk.
- Particle strikes.
- Random background offset charge effects.


Timing faults



Something slows combinational logic propagation enough to miss clock edge.

Danger increases with process scaling.

Single-event upsets




High-energy neutrons generated by interaction of cosmic rays and atoms in upper atmosphere.

Altitude-dependent.

Credit to Silvaco for image.

Technology dependence of fault models




Single-electron tunneling transistor.

Conducts, electrons in single-file.

Turn on by setting gate minimal energy remainder to 0.5 electrons.

SET transfer function



Random background offset charge effects.

Defect at boundary of crystalline structure can trap electron(s).

Infrequent tunneling in and out moves charge relative to gate.

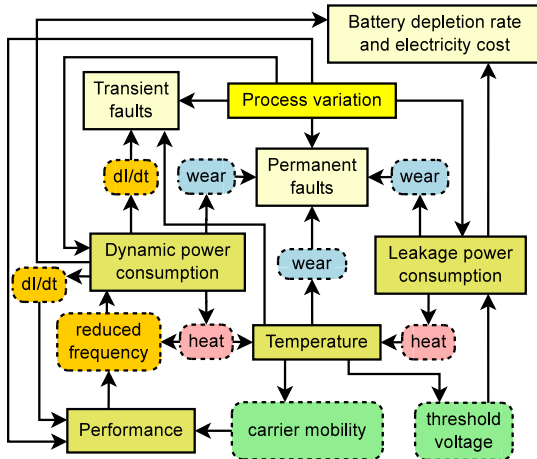
Intermittently shifts entire I-V curve.

Permanent faults

Once they occur, they are generally consistent.

- Electromigration.
- TDDDB.
- (N/P)BTI.
- Thermal cycling.


Relationships among power, temperature, and reliability



Wear mechanisms I

Electromigration.

Dislocation of metal atoms caused by momentum imparted by electrical current in wires and vias.



Credit to Taychatanapat, Bolotin, and Kuemmeth for figure.

Wear mechanisms II

Time-dependent dielectric breakdown.

Deterioration of the gate dielectric layer: formation of conductive path.

Stress migration

Directionally biased motion of atoms in wires due to mechanical stress.

Negative bias temperature instability

- Electric field dependent disassociation of Si-H bonds at Si-SiO₂ interface.
- Increases threshold voltage.
- Significant for PMOSFETs under negative bias.
- Partially recovers when negative bias is removed.

Wear mechanisms III

Thermal cycling

- Mechanical stress resulting from mismatched coefficients of thermal expansion for adjacent material layers.
- Special class of memory: depends on recent temperature history, not just wear state and environment.

Lifetime estimation of the failure mechanisms

Most mechanisms

Arrhenius equation:

$$MTTF = j_1 e^{\frac{j_2}{T}}$$

- j_1 and j_2 : wear process dependent constants.
- T : temperature.

Thermal cycling

Generalized Coffin-Manson eq.:

$$N = \frac{k_1}{(\delta T - T_{th})^{k_2}} e^{\frac{k_3}{T_{max}}},$$

- k_1 , k_2 , and k_3 : constants.
- N : cycles to failure.
- δT : thermal cycle amplitude.
- T_{th} : temperature change threshold.
- T_{max} : maximum temperature during the cycle.

Outline

1. Deadlines and announcements
2. Motivation
3. Sources of reliability problems
4. Fault detection and correction
5. Testing and formal methods
6. Thermal analysis

Section outline

4. Fault detection and correction

- Fault detection

 - Responses to transient and intermittent faults

 - Responses to permanent faults

Fault detection

Timing-based

Memory EDC/ECC

Redundancy

State introspection

Section outline

4. Fault detection and correction

Fault detection

Responses to transient and intermittent faults

Responses to permanent faults

Responses to transient and intermittent faults

Improved fabrication process.

ECC / EDC.

(Active) shielding.

Voltage and temperature control.

Preemptive thermal throttling.

Temporal redundancy.

Temporal redundancy: checkpointing

A tool for robustness in the presence of transient faults.

Periodically store system state.

On fault detection, roll back to known-good state.

Should system-wide or incremental, as-needed restores be used?

When should checkpoints be taken?

Structural redundancy

Add redundant components.

Dual modular redundancy allows detection.

Triple modular redundancy allows correction.

Other responses

Re-assignment: migrate tasks to still-working components.

Built-in self repair.

Fallback operating modes.

Section outline

4. Fault detection and correction

Fault detection

Responses to transient and intermittent faults

Responses to permanent faults

Temperature-dependant wear

Design-time and run-time planning to control spatial and temporal thermal patterns.

Improved cooling solutions.

Improved materials.

Electromigration

Reduce temperature.

Reduce current.

Spatial redundancy.

Manufacturing defects

Improved fabrication processes.

Improved layout.

Spatial redundancy.

Perhaps with BISR.

Outline

1. Deadlines and announcements
2. Motivation
3. Sources of reliability problems
4. Fault detection and correction
5. Testing and formal methods
6. Thermal analysis

Hardware testing

Combinational

- Subject logic to input test vectors, perhaps designed for coverage.
- Verify that outputs are correct.
- Full coverage impractical: 64-bit adder: 2^{128} test vectors.)

Sequential

- First, get system into desired state.
- Then, test many inputs.
- Achieving high coverage much harder than combinational.
- Can use scan-flops to simplify.

Conventional software testing

Implement and test.

Number of tests bounded but number of inputs huge.

Imperfect coverage.

Model checking

Use finite state system representation.

Use exhaustive state space exploration to guarantee desired properties hold for all possible paths.

Guarantees properties.

Difficulty with variables that can take on many values.

Symbolic techniques can improve this

Difficulty with large number of processes.

Critical barriers to use

For simple systems, manual proofs possible.

For very complex systems, state space exploration intractable.

May require new, more formal, specification language.

Overcoming barriers to use

Automatic abstraction techniques permitting use on more complex systems.

Difficult problem.

Target moderate-complexity systems where reliability is important.



- Medical devices.
- Transportation devices.
- Electronic commerce applications.

Give users a high-level language that is actually easier to use than their current language, and provide a path to a language used in existing model checkers.

Outline

1. Deadlines and announcements
2. Motivation
3. Sources of reliability problems
4. Fault detection and correction
5. Testing and formal methods
6. Thermal analysis

R(C) model



Partition into 3-D elements (diagram 2-D for simplicity)

Thermal resistance \leftrightarrow Resistance

Heat flow \leftrightarrow Current

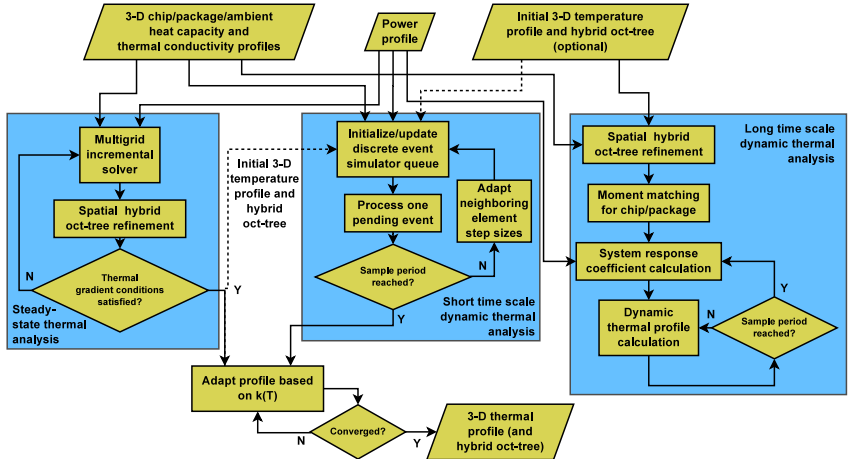
For dynamic: Heat capacity \leftrightarrow Capacitance

Problem definition

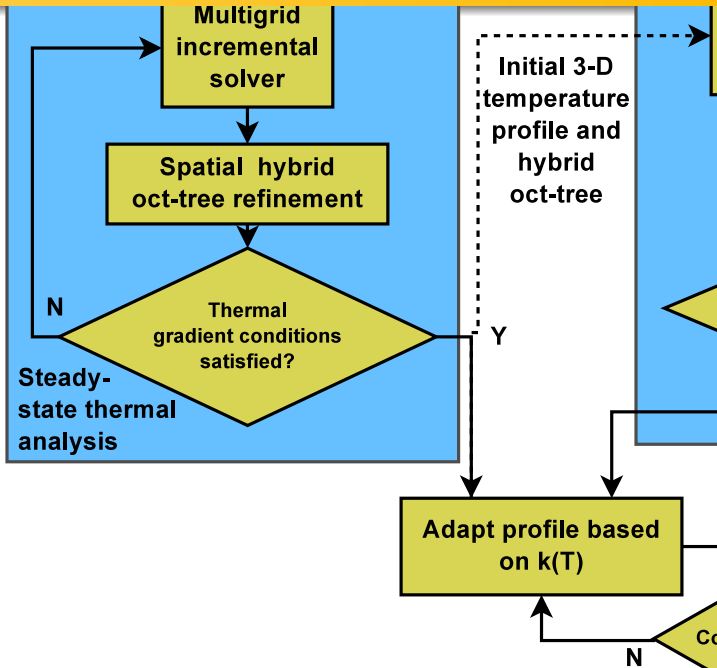
$$\mathbf{C} \frac{d\mathbf{T}(t)}{dt} = \mathbf{A}\mathbf{T}(t) + \mathbf{P}U(t)$$

- **A** is the thermal conductivity matrix
- Steady-state: Initial temperature and **C** unnecessary
- Dynamic: Transient temperature analysis, must also consider heat capacity

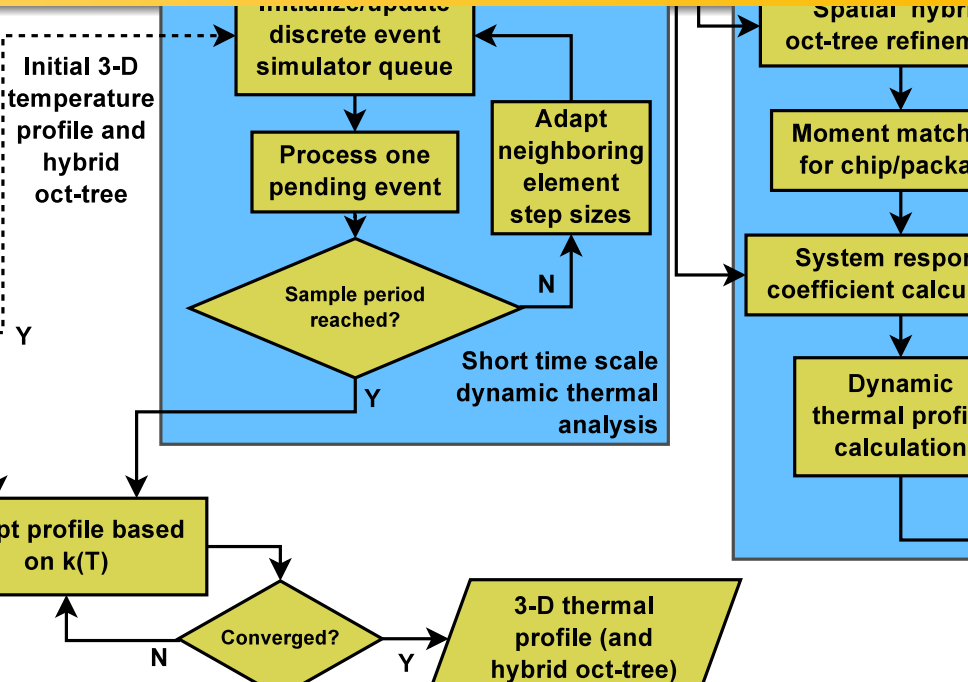
Thermal analysis infrastructure overview



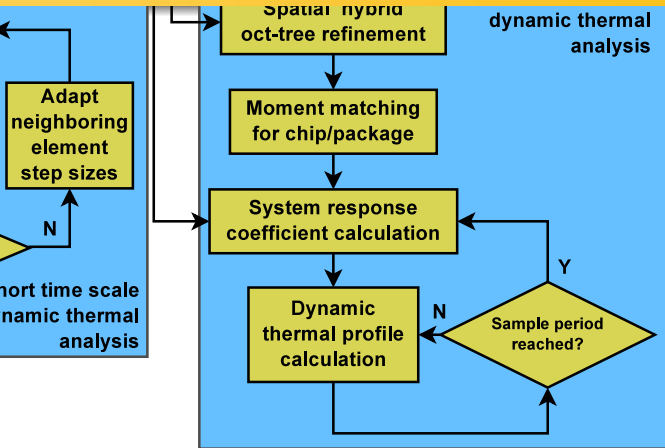
Thermal analysis infrastructure overview



Thermal analysis infrastructure overview




Thermal analysis infrastructure overview



3-D thermal profile (and hybrid oct-tree)

Thermal analysis infrastructure overview




Steady-state thermal analysis

Basis: Multigrid analysis

Fast, multi-resolution relaxation method for matrix solving.

- 1 Iterative solver (relaxation) on fine grid.
 - 2 Coarsen and propagate residual upward.
 - 3 Iterative solver for error at coarser level.
 - 4 Correct fine-grained solution based on coarse-grained error.
 - 5 Iterative solver for error at fine level.
- Main challenge: Too slow for repeated use on large structures, especially 3-D chip-package modeling.
 - **Observation: Steepness of thermal gradients vary across IC.**


Neighbor temperature difference histogram



Spatial adaptation can improve performance w.o. loss of accuracy.

Hybrid oct-tree

- Reduce element count by merging when $\Delta T < \epsilon$
- Conventional oct-tree inefficient for chip-package model
- Anisotropic thermal gradients
- We generalize to **hybrid oct-tree**
- Arbitrary partitioning on each axis



Hybrid oct-tree

