Scheduling Analysis at Different Levels of Modeling

Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213

Peter H Feiler
June 2009

Multi-Fidelity Resource Budgeting

Resource capacities for processors, memory, bus/networks

- Compute resources: MIPS, MB, bandwidth
- Physical resources: power

Budgets for major subsystems

- Capacity and budget totals
- Early deployment decisions & resource-specific budget totals
- Port group connections & bandwidth budgets

System decomposition & budget refinement

- Budget rollup & re-negotiation

Task & communication refinement

- Rates, WCET and budgets

What If Scheduling Analysis

Explore our options

- Use faster processor
- Add second processor
- Rewrite code to reduce worst-case execution time
- Consider lower signal processing rate for controller
- Leverage operational modes (higher fidelity)

What Are the Scheduling Semantics?

Legacy Ada tasks as “partitions”

- Are scheduled by cyclic executive
- Periodic application tasks scheduled within Ada task as cyclic executive
- Harmonic subrates: finish in frame, manual load distribution

Preemptive partition scheduling on commercial RTOS

- Oxymoron?: ARINC653 specifies static line scheduling

Dispatch by virtual timer

- Virtual timer per legacy Ada task/partition
- All partitions per processor at same rate
- Timer alignment in priority order to reduce context switches

Asynchronous set of processors

- Each processor on its own clock

System-Level Scheduling

Distributed physical & logical resource coordination

Time and data consistent mode transition

Scheduling planning & plan execution

Managing critical end-to-end system response

Software-Based Latency Contributors

Execution time variation: algorithm, use of cache

Processor speed

Resource contention

Preemption

Legacy & shared variable communication

Rate group optimization

Protocol specific communication delay

Partitioned architecture

Migration of functionality

Fault tolerance strategy

Resource Allocation & Management

Resource types

- Physical & logical resources
- Consumable & renewable
- Budgets & capacity

Multi-granularity allocation

- Partition allocations & task allocations

Multiple system-level resource tradeoffs

- Minimized network traffic
- Minimized power
- Minimized response times

Manage Uncertainty & Variability

Stable & variable parameters (rate vs. execution time)

Estimates, benchmarks, actuals vs. dynamic behavior

Sensitivity analysis of key parameters

Utilize operational modes

Predictable timing for stochastic workload (RTQT)

Dynamic workload management (QRAM)

Increased Confidence through Continuous Analysis

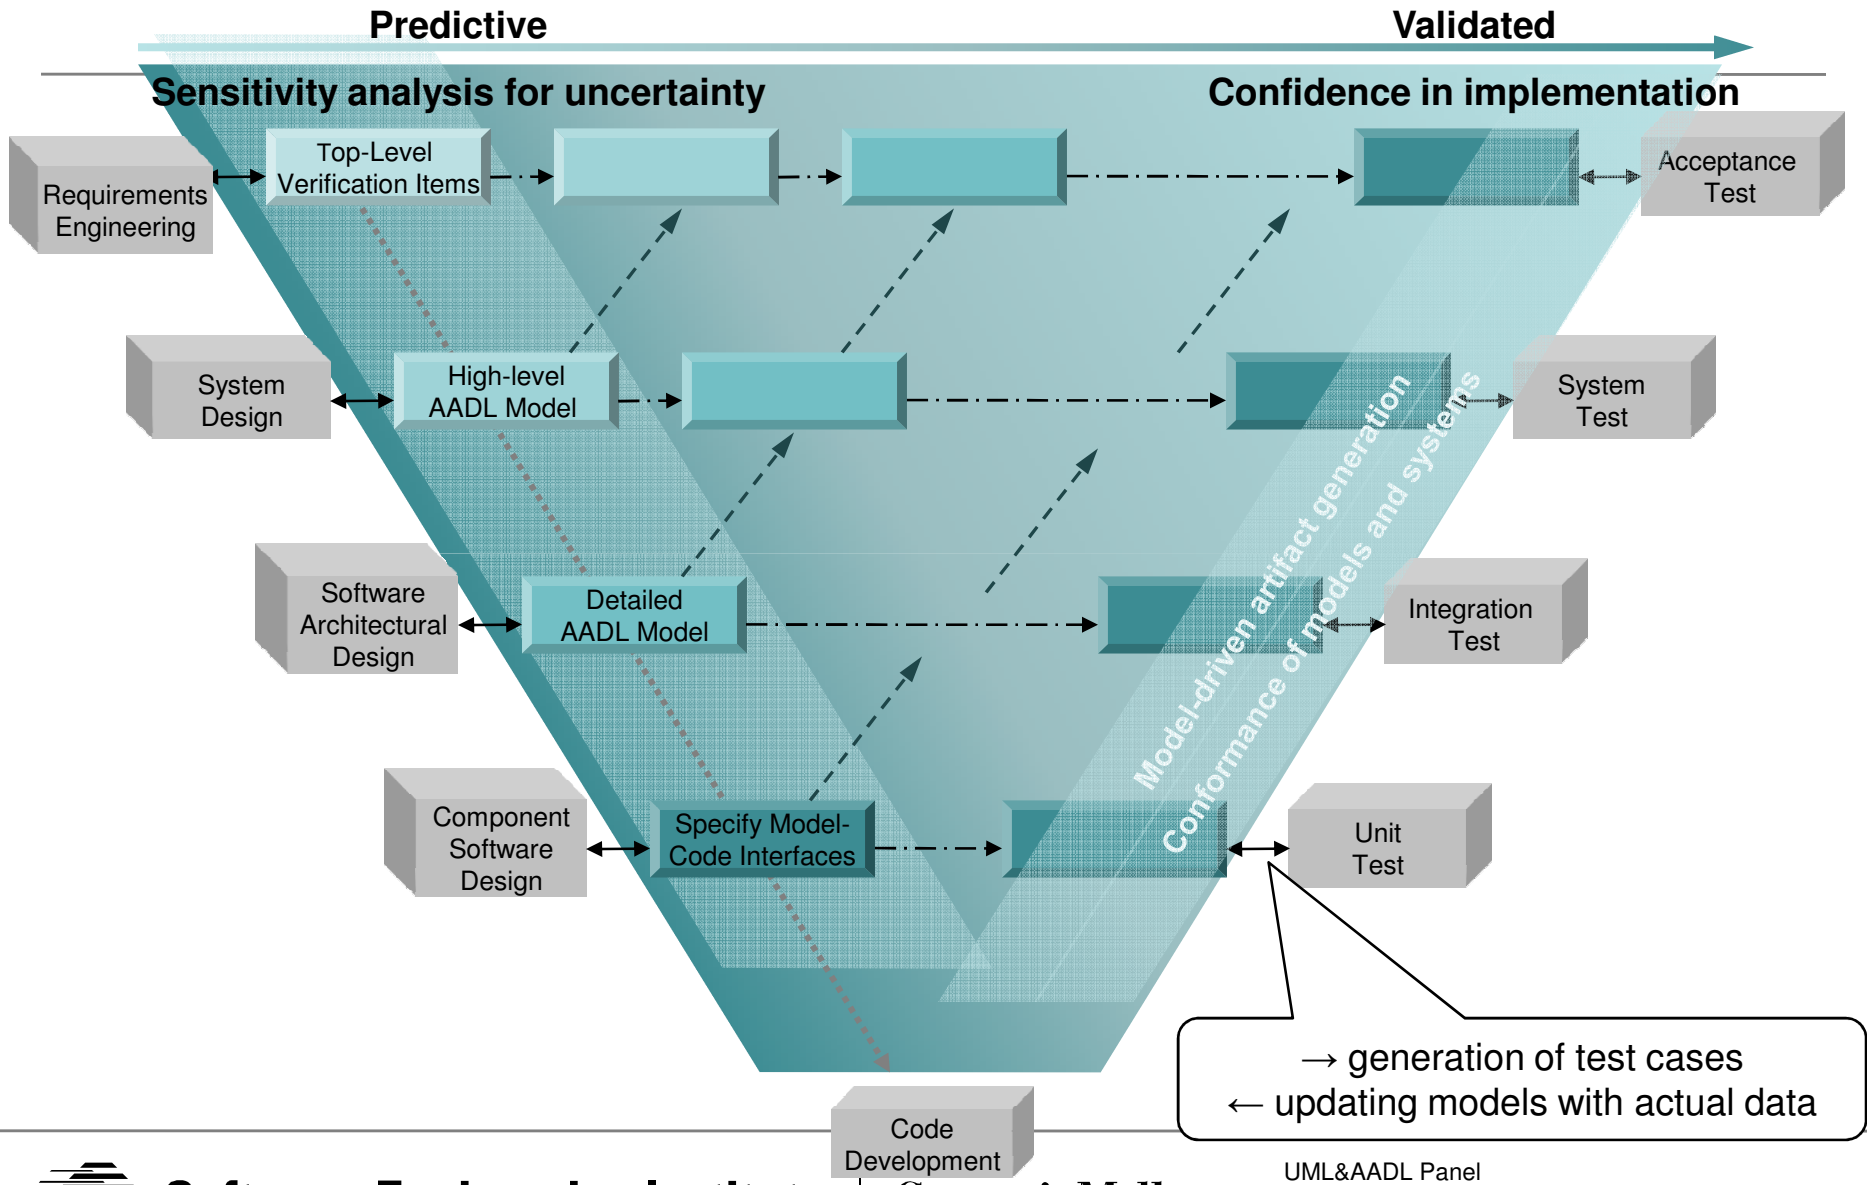

Software Engineering Institute | Carnegie Mellon

Peter Feiler

phf@sei.cmu.edu

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.