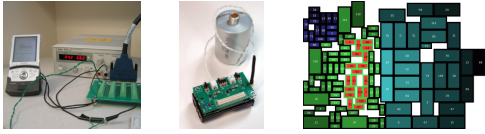


Robert Dick

<http://robertdick.org/esds-two-week>

<http://ziyang.eecs.northwestern.edu/~dickrp/esds-two-week>
 Department of Electrical Engineering and Computer Science
 Northwestern University

Office at Tsinghua University: 9-310 East Main Building



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

Conventional software testing

- Implement and test
- Number of tests bounded but number of inputs huge
- Imperfect coverage

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

Cross-talk

- Shielding
- Bus encoding

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

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

Particle impact

- Temporal redundancy
- Structural redundancy
- Voltage control

Random background offset charge

- Improvements to fabrication
- Temporal redundancy
- Structural redundancy

10 Robert Dick Embedded System Design and Synthesis

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?

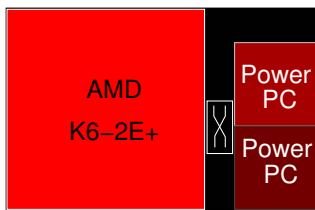
12 Robert Dick Embedded System Design and Synthesis

Manufacturing defects

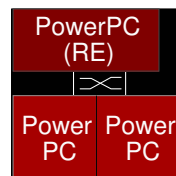
- Spatial redundancy

14 Robert Dick Embedded System Design and Synthesis

Motivating example for reliability optimization



Solution I



Solution II

16 Robert Dick Embedded System Design and Synthesis

Temperature-induced timing faults

- Preemptive throttling
- Global planning

11 Robert Dick Embedded System Design and Synthesis

Electromigration

- Reduce temperature
- Reduce current
- Spatial redundancy

13 Robert Dick Embedded System Design and Synthesis

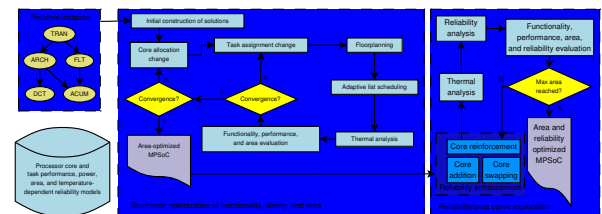
Example lifetime failure aware synthesis flow

Changyun Zhu, Zhenyu Gu, Robert P. Dick, and Li Shang. Reliable multiprocessor system-on-chip synthesis. In *Proc. Int. Conf. Hardware/Software Codesign and System Synthesis*, September 2007. To appear

- Use temperature reduction and spatial redundancy to increase system MTTF
- System MTTF: the expected amount of time an MPSoC will operate, possibly in the presence of component faults, before its performance drops below some designer-specified constraint or it is no longer able to meet its functionality requirements

15 Robert Dick Embedded System Design and Synthesis

Reliability optimization flow

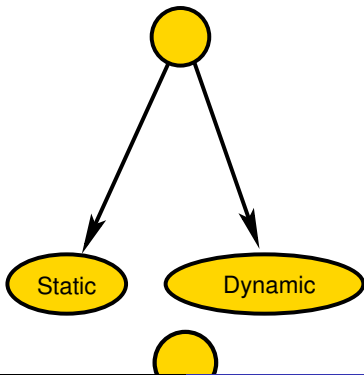


17 Robert Dick Embedded System Design and Synthesis

Lifetime reliability optimization challenges

- Accurate reliability models
- Efficient system-level reliability models
- Efficient fault detection and recovery solutions
- Optimization

Taxonomy of real-time systems



Dynamic

- Task arrival times unpredictable.
- Static (compile-time) analysis possible only for simple cases.
- Even then, the portion of required processor utilization efficiency goes to 0.693.
- In many real systems, this is very difficult to apply in reality (more on this later).
- Use the right tools but don't over-simplify, e.g.,
 We assume, without loss of generality, that all tasks are independent.

If you do this people will make jokes about you.

Hard real-time

- Difficult problem. Some timing constraints inflexible.
- Simplifies problem formulation.

Importance of understanding fault class

- Many reliability techniques attempt to deal with arbitrary fault processes
- However, the properties of the fault process most significant for a particular application may be important
 - Considering them can allow more efficient and reliable designs

Static

- Task arrival times can be predicted.
- Static (compile-time) analysis possible.
- Allows good resource usage (low processor idle time proportions).
- Sometimes designers shoehorn dynamic problems into static formulations allowing a good solution to the wrong problem.

Soft real-time

- More slack in implementation
- Timing may be suboptimal without being incorrect
- Problem formulation can be much more complicated than hard real-time
- Two common (and one uncommon) methods of dealing with non-trivial soft real-time system requirements
 - Set somewhat loose hard timing constraints
 - Informal design and testing
 - Formulate as optimization problem

Periodic

- Each task (or group of tasks) executes repeatedly with a particular period.
- Allows some nice static analysis techniques to be used.
- Matches characteristics of many real problems...
- ... and has little or no relationship with many others that designers try to pretend are periodic.

Periodic → Single-rate

- One period in the system.
- Simple.
- Inflexible.
- This is how a *lot* of wireless sensor networks are implemented.

32

Robert Dick

Embedded System Design and Synthesis

Periodic → Other

- It is possible to have tasks with deadlines less than, equal to, or greater than their periods.
- Results in multi-phase, circular-time schedules with multiple concurrent task instances.
 - If you ever need to deal with one of these, see me (take my code). This class of scheduler is nasty to code.

34

Robert Dick

Embedded System Design and Synthesis

Definitions

- Task
- Processor
- Graph representations
- Deadline violation
- Cost functions

36

Robert Dick

Embedded System Design and Synthesis

Processor

- Processors execute tasks
- Distributed systems
 - Contain multiple processors
 - Inter-processor communication has impact on system performance
 - Communication is challenging to analyze
- One processor type: Homogeneous system
- Multiple processor types: Heterogeneous system

38

Robert Dick

Embedded System Design and Synthesis

Periodic → Multirate

- Multiple periods.
- Can use notion of circular time to simplify static (compile-time) schedule analysis E. L. Lawler and D. E. Wood. Branch-and-bound methods: A survey. *Operations Research*, pages 699–719, July 1966.
- Co-prime periods leads to analysis problems.

33

Robert Dick

Embedded System Design and Synthesis

Aperiodic

- Also called sporadic, asynchronous, or reactive
- Implies dynamic
- Bounded arrival time interval permits resource reservation
- Unbounded arrival time interval impossible to deal with for any resource-constrained system

35

Robert Dick

Embedded System Design and Synthesis

Task

- Some operation that needs to be carried out
- Atomic completion: A task is all done or it isn't
- Non-atomic execution: A task may be interrupted and resumed

37

Robert Dick

Embedded System Design and Synthesis

Task/processor relationship

WC exec time (s)

Tooth	7.7E-6	...	
Road	330E-9	...	
FIR	4.1E-6	...	
Matrix	310E-3	...	

IBM PowerPC 405GP 266 MHz
 IDT79RC32364 100 MHz
 Imsys Cjip 40 MHz

Relationship between tasks, processors, and costs

E.g. power consumption or worst-case execution time

39

Robert Dick

Embedded System Design and Synthesis

Cost functions

- Mapping of real-time system design problem solution instance to cost value
- I.e., allows price, or hard deadline violation, of a particular multi-processor implementation to be determined

40

Robert Dick

Embedded System Design and Synthesis

Central areas of real-time study

- Allocation, assignment and **scheduling**
- Operating systems and **scheduling**
- Distributed systems and **scheduling**
- **Scheduling is at the core of real-time systems study**

42

Robert Dick

Embedded System Design and Synthesis

Allocation, assignment, and scheduling

- In order to efficiently and (when possible) optimally minimize
 - Price, power consumption, soft deadline violations
- Under hard timing constraints
- Providing guarantees whenever possible
- For all the different classes of real-time problem classes

This is what I did for a Ph.D.

44

Robert Dick

Embedded System Design and Synthesis

Distributed systems and scheduling

How does one best dynamically control

- The assignment of tasks to processing nodes...
- ... and their schedules

for systems in which computation nodes may be separated by vast distances such that

- Task deadline violations are bounded (when possible)...
- ... and minimized when no bounds are possible

This is part of what Professor Dinda did for a Ph.D.

46

Robert Dick

Embedded System Design and Synthesis

Back to real-time problem taxonomy: Jagged edges

- Some things dramatically complicate real-time scheduling
- These are horrific, especially when combined
 - Data dependencies
 - Unpredictability
 - Distributed systems
- These are irksome
 - Heterogeneous processors
 - Preemption

41

Robert Dick

Embedded System Design and Synthesis

Allocation, assignment, and scheduling

How does one best

- Analyze problem instance specifications
 - E.g., worst-case task execution time
- Select (and build) hardware components
- Select and produce software
- Decide which processor will be used for each task
- Determine the time(s) at which all tasks will execute

43

Robert Dick

Embedded System Design and Synthesis

Operating systems and scheduling

How does one best design operating systems to

- Support sufficient detail in workload specification to allow good control, e.g., over scheduling, without increasing design error rate
- Design operating system schedulers to support real-time constraints?
- Support predictable costs for task and OS service execution

45

Robert Dick

Embedded System Design and Synthesis

The value of formality: Optimization and costs

- The design of a real-time system is fundamentally a cost optimization problem
- Minimize costs under constraints while meeting functionality requirements
 - Slight abuse of notation here, functionality requirements are actually just constraints
- Why view problem in this manner?
- Without having a concrete definition of the problem
 - How is one to know if an answer is correct?
 - More subtly, how is one to know if an answer is optimal?

47

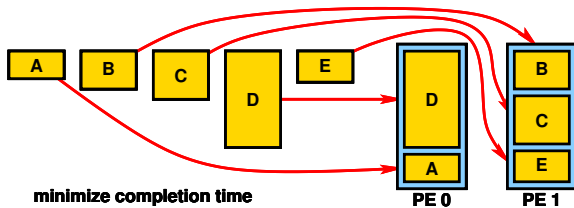
Robert Dick

Embedded System Design and Synthesis

Optimization

Thinking of a design problem in terms of optimization gives design team members objective criterion by which to evaluate the impact of a design change on quality.
 Know whether your design changes are taking you in a good direction

Problem definition



- Given a set of tasks,
- a cost function,
- and a set of resources,
- decide the exact time each task will execute on each resource

Discrete vs. continuous timing

System-level: Continuous

- Operations are not small integer multiples of the clock cycle

High-level: Discrete

- Operations are small integer multiples of the clock cycle

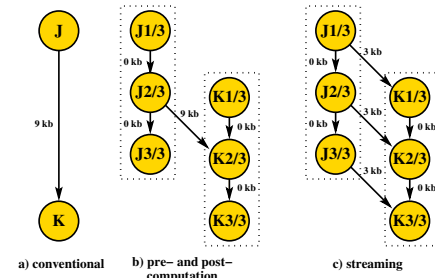
Implications:

- System-level scheduling is more complicated...
- ... however, high-level also very difficult.
- Can we solve this by quantizing time? Why or why not?

Real-time – Best effort

- Why make decisions about system implementation statically?
 - Allows easy timing analysis, hard real-time guarantees
- If a system doesn't have hard real-time deadlines, resources can be more efficiently used by making late, dynamic decisions
- Can combine real-time and best-effort portions within the same specification
 - Reserve time slots
 - Take advantage of slack when tasks complete sooner than their worst-case finish times

Graph extensions



Allows pipelining and pre/post-computation
 In contrast with book, not difficult to use if conversion automated

Types of scheduling problems

- Discrete time – Continuous time
- Hard deadline – Soft deadline
- Unconstrained resources – Constrained resources
- Uni-processor – Multi-processor
- Homogeneous processors – Heterogeneous processors
- Free communication – Expensive communication
- Independent tasks – Precedence constraints
- Homogeneous tasks – Heterogeneous tasks
- One-shot – Periodic
- Single rate – Multirate
- Non-preemptive – Preemptive
- Off-line – On-line

Hard deadline – Soft deadline

Tasks may have hard or soft deadlines

- Hard deadline
 - Task must finish by given time or schedule invalid
- Soft deadline
 - If task finishes after given time, schedule cost increased

Unconstrained – Constrained resources

- Unconstrained resources
 - Additional resources may be used at will
- Constrained resources
 - Limited number of devices may be used to execute tasks

Uni-processor – Multi-processor

- Uni-processor
 - All tasks execute on the same resource
 - This can still be somewhat challenging
 - However, sometimes in \mathcal{P}
- Multi-processor
 - There are multiple resources to which tasks may be scheduled
- Usually \mathcal{NP} -complete

57

Robert Dick

Embedded System Design and Synthesis

Free – Expensive communication

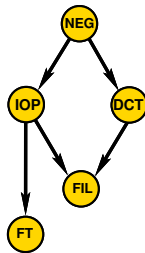
- Free communication
 - Data transmission between resources has no time cost
- Expensive communication
 - Data transmission takes time
 - Increases problem complexity
 - Generation of schedules for communication resources necessary
 - Usually \mathcal{NP} -complete

59

Robert Dick

Embedded System Design and Synthesis

Homogeneous – Heterogeneous tasks



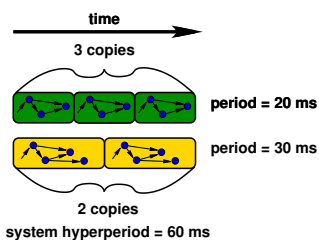
- Homogeneous tasks: All tasks are identical
- Heterogeneous tasks: Tasks differ

61

Robert Dick

Embedded System Design and Synthesis

Single rate – Multirate



- Single rate: All tasks have the same period
- Multirate: Different tasks have different periods
 - Complicates scheduling
 - Can copy out to the least common multiple of the periods (hyperperiod)

63

Robert Dick

Embedded System Design and Synthesis

Homogeneous – Heterogeneous processors

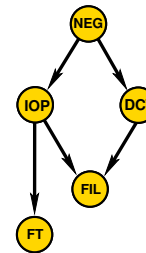
- Homogeneous processors
 - All processors are the same type
- Heterogeneous processors
 - There are different types of processors
 - Usually \mathcal{NP} -complete

58

Robert Dick

Embedded System Design and Synthesis

Independent tasks – Precedence constraints



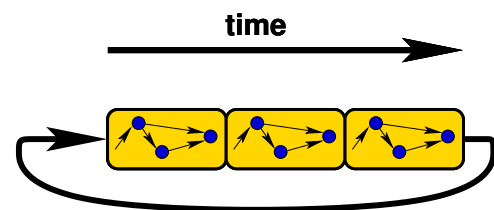
- Independent tasks: No previous execution sequence imposed
- Precedence constraints: Weak order on task execution order

60

Robert Dick

Embedded System Design and Synthesis

One-shot – Periodic



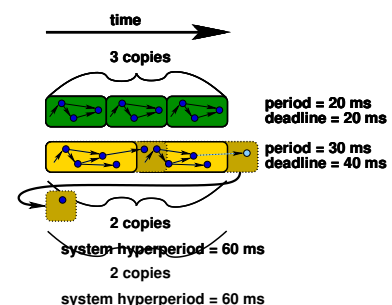
- One-shot: Assume that the task set executes once
- Periodic: Ensure that the task set can repeatedly execute at some period

62

Robert Dick

Embedded System Design and Synthesis

Periodic graphs



64

Robert Dick

Embedded System Design and Synthesis

Aperiodic/sporadic graphs

- No precise periods imposed on task execution
- Useful for representing reactive systems
- Difficult to guarantee hard deadlines in such systems
 - Possible if minimum inter-arrival time known

Aperiodic to periodic

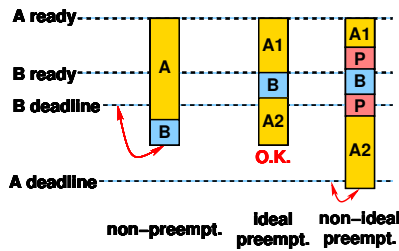
Can design periodic specifications that meet requirements posed by aperiodic/sporadic specifications

- Some resources will be wasted

Example:

- At most one aperiodic task can arrive every 50 ms
- It must complete execution within 100 ms of its arrival time

Non-preemptive – Preemptive



- Non-preemptive: Tasks must run to completion
- Ideal preemptive: Tasks can be interrupted without cost
- Non-ideal preemptive: Tasks can be interrupted with cost

Hardware-software co-synthesis scheduling

Automatic allocation, assignment, and scheduling of system-level specification to hardware and software
Scheduling problem is hard

- Hard and soft deadlines
- Constrained resources, but resources unknown (cost functions)
- Multi-processor
- Strongly heterogeneous processors and tasks
 - No linear relationship between the execution times of a tasks on processors

Periodic vs. aperiodic

Periodic applications

- Power electronics
- Transportation applications
 - Engine controllers
 - Brake controllers
- Many multimedia applications
 - Video frame rate
 - Audio sample rate
- Many digital signal processing (DSP) applications

However, devices which react to unpredictable external stimuli have aperiodic behavior

Many applications contain periodic and aperiodic components

Aperiodic to periodic

- Can easily build a periodic representation with a deadline and period of 50 ms
 - Problem, requires a 50 ms execution time when 100 ms should be sufficient
- Can use overlapping graphs to allow an increase in execution time
 - Parallelism required

The main problem with representing aperiodic problems with periodic representations is that the tradeoff between deadline and period must be made at design/synthesis time

Off-line – On-line

Off-line

- Schedule generated before system execution
- Stored, e.g., in dispatch table. for later use
- Allows strong design/synthesis/compile-time guarantees to be made
- Not well-suited to strongly reactive systems

On-line

- Scheduling decisions made during the execution of the system
- More difficult to analyze than off-line
 - Making hard deadline guarantees requires high idle time
 - No known guarantee for some problem types
- Well-suited to reactive systems

Hardware-software co-synthesis scheduling

- Expensive communication
 - Complicated set of communication resources
- Precedence constraints
- Periodic
- Multirate
- Strong interaction between $\mathcal{N}\mathcal{P}$ -complete allocation-assignment and $\mathcal{N}\mathcal{P}$ -complete scheduling problems
- Will revisit problem later in course if time permits

Behavioral synthesis scheduling

- Difficult real-world scheduling problem
 - Not multirate
 - Discrete notion of time
 - Generally less heterogeneity among resources and tasks
- What scheduling algorithms should be used for these problems?

73

Robert Dick

Embedded System Design and Synthesis

Scheduling methods

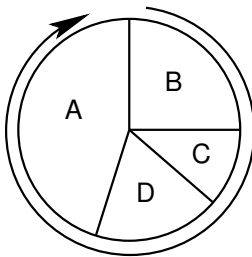
- MILP
- Force-directed
- Frame-based
- PSGA
- RMS

75

Robert Dick

Embedded System Design and Synthesis

Weighted round robin



Time

Weighted round-robin: Time-sliced with variable time slots

77

Robert Dick

Embedded System Design and Synthesis

Priority-driven

- Impose linear order based on priority metric
- Possible metrics
 - Earliest start time (EST)
 - Latest start time
 - Danger! LST also stands for least slack time.
 - Shortest execution time first (SETF)
 - Longest execution time first (LETF)
 - Slack (LFT - EFT)

79

Robert Dick

Embedded System Design and Synthesis

Scheduling methods

- Clock
- Weighted round-robin
- List scheduling
- Priority
 - EDF, LST
 - Slack
 - Multiple costs

74

Robert Dick

Embedded System Design and Synthesis

Clock-driven scheduling

Clock-driven: Pre-schedule, repeat schedule
Music box:

- Periodic
- Multi-rate
- Heterogeneous
- Off-line
- Clock-driven

76

Robert Dick

Embedded System Design and Synthesis

List scheduling

- Pseudo-code:
 - Keep a list of ready jobs
 - Order by priority metric
 - Schedule
 - Repeat
- Simple to implement
- Can be made very fast
- Difficult to beat quality

78

Robert Dick

Embedded System Design and Synthesis

List scheduling

- Assigns priorities to nodes
- Sequentially schedules them in order of priority
- Usually very fast
- Can be high-quality
- Prioritization metric is important

80

Robert Dick

Embedded System Design and Synthesis

Prioritization

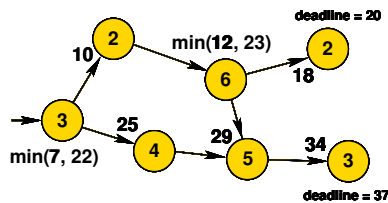
- As soon as possible (ASAP)
- As late as possible (ALAP)
- Slack-based
- Dynamic slack-based
- Multiple considerations

81

Robert Dick

Embedded System Design and Synthesis

As late as possible (ALAP)



- From deadlines, topological sort on the precedence graph
- Propagate execution times, taking the min at reconverging paths
- Consider precedence-constraint satisfied tasks
 - Schedule in order of increasing latest start time (LST)

83

Robert Dick

Embedded System Design and Synthesis

Multiple considerations

- Nothing prevents multiple prioritization methods from being used
- Try one method, if it fails to produce an acceptable schedule, reschedule with another method

85

Robert Dick

Embedded System Design and Synthesis

EDF, LST optimality

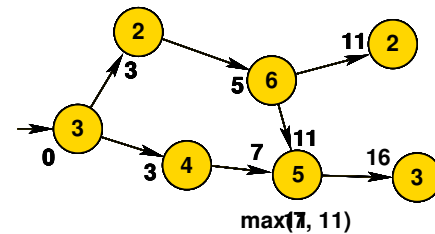
- EDF optimal if zero-cost preemption, uniprocessor assumed
 - Why?
 - What happens when preemption has cost?
- Same is true for slack-based list scheduling in absence of preemption cost

87

Robert Dick

Embedded System Design and Synthesis

As soon as possible (ASAP)



- From root, topological sort on the precedence graph
- Propagate execution times, taking the max at reconverging paths
- Schedule in order of increasing earliest start time (EST)

82

Robert Dick

Embedded System Design and Synthesis

Slack-based

- Compute EFT, LFT
- For all tasks, find the difference, LFT – EFT
- This is the *slack*
- Schedule precedence-constraint satisfied tasks in order of increasing slack
- Can recompute slack each step, expensive but higher-quality result
 - Dynamic critical path scheduling

84

Robert Dick

Embedded System Design and Synthesis

Effective release times

- Ignore the book on this
 - Considers simplified, uniprocessor, case
- Use EFT, LFT computation
- Example?

86

Robert Dick

Embedded System Design and Synthesis

Breaking EDF, LST optimality

- Non-zero preemption cost
- Multiprocessor
- Why?

88

Robert Dick

Embedded System Design and Synthesis

Multi-rate tricks

- Contract deadline
 - Usually safe
- Contract period
 - Sometimes safe
- Consequences?

89

Robert Dick

Embedded System Design and Synthesis

MILP scheduling

P the set of tasks
 t_{max} maximum time
 $start(p, t)$ 1 if task p starts at time t , 0 otherwise
 D the set of execution delays
 E the set of precedence constraints

$$t_{start}(p) = \sum_{t=0}^{t_{max}} t \cdot start(p, t) \text{ the start time of } p$$

91

Robert Dick

Embedded System Design and Synthesis

MILP scheduling

- Too slow for large instances of \mathcal{NP} -complete scheduling problems
- Numerous optimization algorithms may be used for scheduling
- List scheduling is one popular solution
- Integrated solution to allocation/assignment/scheduling problem possible
- Performance problems exist for this technique

93

Robert Dick

Embedded System Design and Synthesis

Self force

F_i all slots in time frame for i
 F_i' all slots in new time frame for i
 D_t probability density (sum) for slot t
 δD_t change in density (sum) for slot t resulting from scheduling
 self force

$$A = \sum_{t \in F_a} D_t \cdot \delta D_t$$

95

Robert Dick

Embedded System Design and Synthesis

Linear programming

- Minimize a linear equation subject to linear constraints
 - In \mathcal{P}
- Mixed integer linear programming: One or more variables discrete
 - \mathcal{NP} -complete
- Many good solvers exist
- Don't rebuild the wheel

90

Robert Dick

Embedded System Design and Synthesis

MILP scheduling

Each task has a unique start time

$$\forall p \in P, \sum_{t=0}^{t_{max}} start(p, t) = 1$$

Each task must satisfy its precedence constraints and timing delays

$$\forall \{p_i, p_j\} \in E, \sum_{t=0}^{t_{max}} t_{start}(p_i) \geq t_{start}(p_j) + d_j$$

Other constraints may exist

- Resource constraints
- Communication delay constraints

92

Robert Dick

Embedded System Design and Synthesis

Force directed scheduling

- P. G. Paulin and J. P. Knight. Force-directed scheduling for the behavioral synthesis of ASICs. *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, 8(6):661-679, June 1989
- Calculate EST and LST of each node
- Determine the force on each vertex at each time-step
- Force: Increase in probabilistic concurrency
 - Self force
 - Predecessor force
 - Successor force

94

Robert Dick

Embedded System Design and Synthesis

Predecessor and successor forces

pred all predecessors of node under consideration
succ all successors of node under consideration
 predecessor force

$$B = \sum_{b \in \text{pred}} \sum_{t \in F_b} D_t \cdot \delta D_t$$

successor force

$$C = \sum_{c \in \text{succ}} \sum_{t \in F_c} D_t \cdot \delta D_t$$

96

Robert Dick

Embedded System Design and Synthesis

Intuition

total force: $A + B + C$

- Schedule operation and time slot with minimal total force
 - Then recompute forces and schedule the next operation
- Attempt to balance concurrency during scheduling

97

Robert Dick

Embedded System Design and Synthesis

Force directed scheduling

- Limitations?
- What classes of problems may this be used on?

99

Robert Dick

Embedded System Design and Synthesis

Problem space genetic algorithm

- Let's finish off-line scheduling algorithm examples on a bizarre example
- Use conventional scheduling algorithm
- Transform problem instance
- Solve
- Validate
- Evolve transformations

101

Robert Dick

Embedded System Design and Synthesis

Rate monotonic scheduling

Main idea

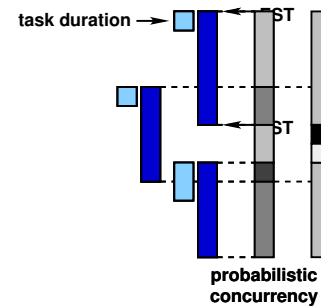
- 1973, Liu and Layland derived optimal scheduling algorithm(s) for this problem
- Schedule the job with the smallest period (period = deadline) first
- Analyzed worst-case behavior on any task set of size n
- Found utilization bound: $U(n) = n \cdot (2^{1/n} - 1)$
- 0.828 at $n = 2$
- As $n \rightarrow \infty$, $U(n) \rightarrow \log 2 = 0.693$
- Result: For any problem instance, if a valid schedule is possible, the processor need never spend more than 31% of its time idle

103

Robert Dick

Embedded System Design and Synthesis

Force directed scheduling



98

Robert Dick

Embedded System Design and Synthesis

Implementation: Frame-based scheduling

- Break schedule into (usually fixed) frames
 - Large enough to hold a long job
 - Avoid preemption
- Evenly divide hyperperiod
- Scheduler makes changes at frame start
- Network flow formulation for frame-based scheduling
- Could this be used for on-line scheduling?

100

Robert Dick

Embedded System Design and Synthesis

Rate monotonic scheduling (RMS)

- Single processor
- Independent tasks
- Differing arrival periods
- Schedule in order of increasing periods
- No fixed-priority schedule will do better than RMS
- Guaranteed valid for loading $\leq \ln 2 = 0.69$
- For loading $> \ln 2$ and < 1 , correctness unknown
- Usually works up to a loading of 0.88

102

Robert Dick

Embedded System Design and Synthesis

Optimality and utilization for limited case

- Simply periodic: All task periods are integer multiples of all lesser task periods
- In this case, RMS/DMS optimal with utilization 1
- However, this case rare in practice
- Remains feasible, with decreased utilization bound, for in-phase tasks with arbitrary periods

104

Robert Dick

Embedded System Design and Synthesis

Rate monotonic scheduling

- Constrained problem definition
- Over-allocation often results
- However, in practice utilization of 85%–90% common
 - Lose guarantee
- If phases known, can prove by generating instance

105

Robert Dick

Embedded System Design and Synthesis

Proof sketch for RMS utilization bound

- Consider case in which no period exceeds twice the shortest period
- Find a pathological case: in phase
 - Utilization of 1 for some duration
 - Any decrease in period/deadline of longest-period task will cause deadline violations
 - Any increase in execution time will cause deadline violations

107

Robert Dick

Embedded System Design and Synthesis

Proof sketch for RMS utilization bound

- Same true if execution time of high-priority task reduced
- $e'_i = p_{i+1} - p_i - \epsilon$
- In this case, must increase other e or leave idle for $2 \cdot \epsilon$
- $e'_k = e_k + 2\epsilon$
- $U'' - U = \frac{2\epsilon}{p_k} - \frac{\epsilon}{p_i}$
- Again, $p_k < 2 \rightarrow U'' > U$
- Sum over execution time/period ratios

109

Robert Dick

Embedded System Design and Synthesis

Notes on RMS

- DMS better than or equal RMS when deadline \neq period
- Why not use slack-based?
- What happens if resources are under-allocated and a deadline is missed?

111

Robert Dick

Embedded System Design and Synthesis

Critical instants

Main idea:

A job's critical instant a time at which all possible concurrent higher-priority jobs are also simultaneously released

Useful because it implies latest finish time

106

Robert Dick

Embedded System Design and Synthesis

Proof sketch for RMS utilization bound

- See if there is a way to increase utilization while meeting all deadlines
- Increase execution time of high-priority task
 - $e'_i = p_{i+1} - p_i + \epsilon = e_i + \epsilon$
- Must compensate by decreasing another execution time
- This always results in decreased utilization
 - $e'_k = e_k - \epsilon$
 - $U' - U = \frac{e'_i}{p_i} + \frac{e'_k}{p_k} - \frac{e_i}{p_i} - \frac{e_k}{p_k} = \frac{\epsilon}{p_i} - \frac{\epsilon}{p_k}$
 - Note that $p_i < p_k \rightarrow U' > U$

108

Robert Dick

Embedded System Design and Synthesis

Proof sketch for RMS utilization bound

- Get utilization as a function of adjacent task ratios
- Substitute execution times into $\sum_{k=1}^n \frac{e_k}{p_k}$
- Find minimum
- Extend to cases in which $p_n > 2 \cdot p_k$

110

Robert Dick

Embedded System Design and Synthesis

Scheduling summary

- Scheduling is a huge area
- This lecture only introduced the problem and potential solutions
- Some scheduling problems are easy
- Most useful scheduling problems are hard
 - Committing to decisions makes problems hard: Lookahead required
 - Interdependence between tasks and processors makes problems hard
 - On-line scheduling next Tuesday

112

Robert Dick

Embedded System Design and Synthesis

Mixing on-line and off-line

- Book mixes off-line and on-line with little warning
- Be careful, actually different problem domains
- However, can be used together
- Superloop (cyclic executive) with non-critical tasks
- Slack stealing
- Processor-based partitioning

113

Robert Dick

Embedded System Design and Synthesis

Mixing on-line and off-line

- Slack stealing
- Processor-based partitioning

115

Robert Dick

Embedded System Design and Synthesis

Example problem: Static scheduling

- What is an FPGA?
- Why should real-time systems designers care about them?
- Multiprocessor static scheduling
- No preemption
- No overhead for subsequent execution of tasks of same type
- High cost to change task type
- Scheduling algorithm?

117

Robert Dick

Embedded System Design and Synthesis

Compression references (for next class)

- Haris Lekatsas, Jörg Henkel, and Wayne Wolf. Code compression for low power embedded system design. In *Proc. Design Automation Conf.*, pages 294–299, June 2000
- Lei Yang, Robert P. Dick, Haris Lekatsas, and Srimat Chakradhar. On-Line Memory Compression for Embedded Systems. *ACM Trans. Embedded Computing Systems*. To appear

120

Robert Dick

Embedded System Design and Synthesis

Vehicle routing

- Low-price, slow, ARM-based system
- Long-term shortest path computation
- Greedy path calculation algorithm available, non-preemptable
- Don't make the user wait
 - Short-term next turn calculation
- 200 ms timer available

114

Robert Dick

Embedded System Design and Synthesis

Bizarre scheduling idea

- Scheduling and validity checking algorithms considered so far operate in time domain
- This is a somewhat strange idea
- Think about it and tell/email me if you have any thoughts on it
- Could one very quickly generate a high-quality real-time off-line multi-rate periodic schedule by operating in the frequency domain?
- If not, why not?
- What if the deadlines were soft?

116

Robert Dick

Embedded System Design and Synthesis

Problem: Uniprocessor independent task scheduling

- Problem
 - Independent tasks
 - Each has a period = hard deadline
 - Zero-cost preemption
- How to solve?

118

Robert Dick

Embedded System Design and Synthesis

Project proposals

- Due 12:00 Sunday
- A one-page project description
- Ideally, you will have some preliminary results or ideas based on reading papers or doing analysis already

121

Robert Dick

Embedded System Design and Synthesis

Next class

- Lecture on data compression in embedded system design
- A real, graded quiz