



Synthesis of Application Specific Schedulers for Heterogeneous Real-Time Systems Guaranteeing Safety & Allowing QoS Extensions

- Introduction
- Goal
- Thread States
- Architecture
- Layers
- System Model
- X Modes
- Synthesis Steps
- SS Reduction
- No Clocks
- Example
- Implementation
- Conclusions
- Future

Christos KLOUKINAS

Christos.Kloukinas@imag.fr

Verimag

(<http://www-verimag.imag.fr/>)

Grenoble, France

Introduction



- ⑥ Heterogeneous Applications
 - ⚡ Periodic, Aperiodic, Sporadic threads
 - ⚡ Difficult to analyse their behaviour
- ⑥ Using PIP/PCP for deadlock avoidance
 - ⚡ Difficult to get priorities right (may need to split threads @ I/O points, at least when modelling)
 - ⚡ PIP/PCP are ***pessimistic***
- ⑥ RMA, EDF, *etc.*
 - ⚡ RMA is only for Periodic (must group non-periodic), max $U = 69\%$
 - ⚡ EDF : overhead, unstable under overload, priority inversion, if not schedulable you don't know why
- ⑥ Extending scheduling with **QoS** constraints ???

- Introduction
- Goal
- Thread States
- Architecture
- Layers
- System Model
- X Modes
- Synthesis Steps
- SS Reduction
- No Clocks
- Example
- Implementation
- Conclusions
- Future

Our Goal



Analyse an application & **synthesise** a scheduler which :

- ⑥ **Guarantees Safety** - no Deadlocks / Deadline Misses
- ⑥ is not as **Pessimistic** as PCP
- ⑥ is easy to **extend** with additional **QoS constraints**

Meant to be used on single-processor systems

- Introduction
- Goal
- Thread States
- Architecture
- Layers
- System Model
- X Modes
- Synthesis Steps
- SS Reduction
- No Clocks
- Example
- Implementation
- Conclusions
- Future

States of Threads




Application threads can be :

- ⑥ **Blocked**
- ⑥ **Ready** to execute = \neg **Blocked**
- ⑥ **Safe** to execute \subseteq **Ready**
- ⑥ **Executing** (*at most one*) \in **Safe**

- Introduction
- Goal
- Thread States
- Architecture
- Layers
- System Model
- X Modes
- Synthesis Steps
- SS Reduction
- No Clocks
- Example
- Implementation
- Conclusions
- Future

Application



Execution Scheduler Stack


TASKS CANDIDATING FOR EXECUTION - \mathcal{C}_{exec}

Ready-Exec




READY TASKS - $\mathcal{R}_{exec} \subseteq \mathcal{C}_{exec}$

Safe-Exec



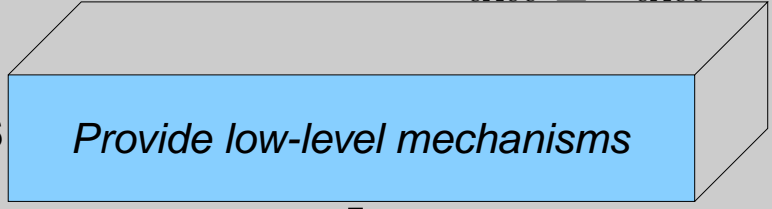
SAFE TASKS - $\mathcal{S}_{exec} \subseteq \mathcal{R}_{exec}$

QoS-Exec



EXECUTING TASKS - $\mathcal{Q}_{exec} \subseteq \mathcal{S}_{exec}$

R-T OS



Choose one among \mathcal{Q}_{exec} for execution

Scheduler Architecture



Scheduler Layers



- ⑥ **Ready-Exec:** which threads are **Ready** ?
(*mutual-exclusion, waiting for notification*)
- ⑥ **Safe-Exec:** which among the **Ready** are **Safe** ?
(*no deadlocks, deadlines honoured*)
- ⑥ **QoS-Exec:** which of the **Safe** meet the **QoS** req. ?
(*speed, memory, energy*)
- ⑥ The Real-Time OS picks one among the latter as the **Executing** thread

We synthesise the **Ready-Exec** & **Safe-Exec** layers, considering a *maximal QoS-Exec* scheduler (*i.e.*, all **Safe** threads meet the **QoS** requirements) — the user provides a different **QoS-Exec** scheduler, if he so wishes.

QoS Policies



- ⑥ What is a QoS policy (our definition) ?
 - ☞ Something that helps increase some “goodness” metric
 - ☞ Something that cannot kill you
So, it’s only applicable on **safe** threads

- ⑥ Examples :
 - ☞ (Locally) Minimise Context-Switches
Helps with optimising Speed, Memory accesses (fewer cache flushes/misses), Energy ...

- ⑥ Exception : (there’s always one ...)
 - ☞ Non-Preemption
Too important to forbid – for network messages, it’s the only way

Scheduler Synthesis - The Casting



- ⑥ The **Good : Scheduler**
- ⑥ The **Bad : Environment** (*i.e.*, Time)
Completely uncontrollable
- ⑥ And the **Ugly : Application**
Only controllable at pre-specified points

Scheduler Synthesis



⑥ Even though in theory we could control each instruction, in practice we can only observe & control instructions of the type:

- ⌚ *lock/unlock* ;
- ⌚ *wait/waitTimed* ;
- ⌚ *waitForNextPeriod* ; and
- ⌚ *timer expirations*

⑥ ***No re-scheduling at notify/notifyAll !***

- ⌚ The notified thread(s) immediately block on the (re-)lock, so no need to re-schedule

Scheduler Synthesis - II




- ⑥ Effectively, a two-player game,
Environment versus **Scheduler** ,
where the **Scheduler** is aided by the **Application**
- ⑥ Each player moves in turn
- ⑥ Scheduler/Controller Synthesis amounts to :
 - ☞ Find a winning strategy for the game
“What action must I (*i.e.*, the **Scheduler**) take to prevent my opponent (*i.e.*, the **Environment**) from winning (*i.e.*, cause the **Application** to miss deadlines) ?”

System Model - The Environment



Alarm $\xRightarrow{\text{alarm}}$ change state & reschedule



$\neg \text{Alarm} \wedge (\text{Compute} \vee \text{wait})$

$\xRightarrow{\text{tick}}$ advance clocks

$\neg \text{Alarm} \wedge \neg (\text{Compute} \vee \text{wait})$


$\xRightarrow{\text{appli}}$ allow application to run

- Introduction
- Goal
- Thread States
- Architecture
- Layers
- System Model
- X Modes
- Synthesis Steps
- SS Reduction
- No Clocks
- Example
- Implementation
- Conclusions
- Future

System Model - The Scheduler

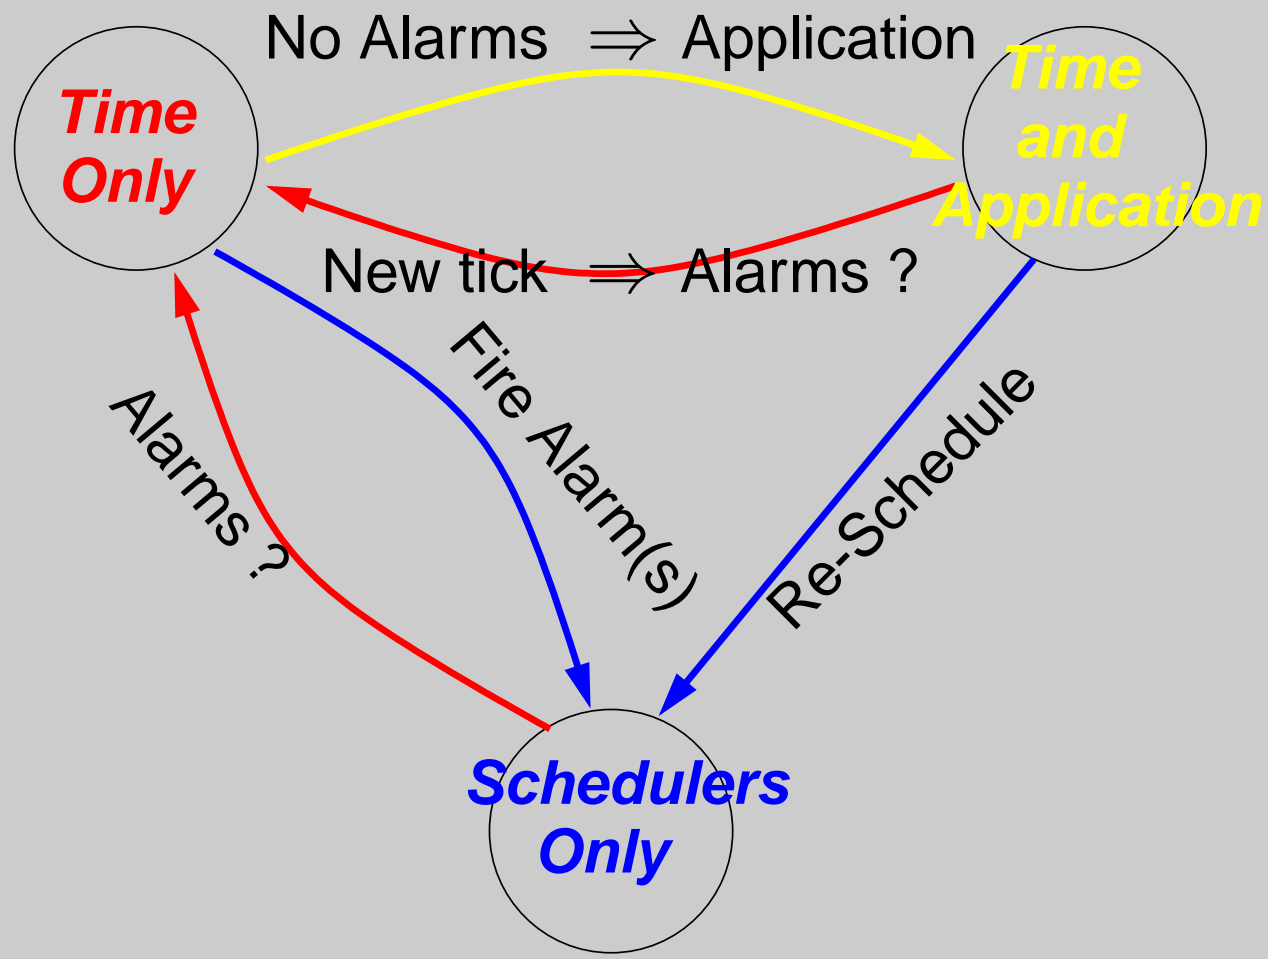
$$t_i \in \mathcal{R}_{\text{exec}} \cap \mathcal{S}_{\text{exec}} \cap \mathcal{Q}_{\text{exec}}$$

$$\text{Choose } t_i \implies T'_{\text{Exec}} = t_i$$



- Introduction
- Goal
- Thread States
- Architecture
- Layers
- System Model
- X Modes
- Synthesis Steps
- SS Reduction
- No Clocks
- Example
- Implementation
- Conclusions
- Future

System Execution Modes



- Introduction
- Goal
- Thread States
- Architecture
- Layers
- System Model
- X Modes
- Synthesis Steps
- SS Reduction
- No Clocks
- Example
- Implementation
- Conclusions
- Future

Scheduler Synthesis Basic Idea



- ⑥ Construct the entire state space
- ⑥ Find **Deadlock** states
 - ⌚ These identify states where the **Application** deadlocked due to a cycle in the shared resources; and
 - ⌚ states where the **Application** missed a deadline.

Scheduler Synthesis Basic Idea - II

- ⑥ From each deadlock state:
 - ⊂ Go backwards in the trace up to the first state where the scheduler could have made this trace impossible (*i.e.*, don't allow some thread to execute)
 - ⊂ All these constraints are effectively the constraints we need for implementing the **Safe** scheduler layer.
 - ⊂ So, **Safe-Exec** is a table, indexed by a system state, containing the **Application** threads which we must not allow to execute at that state
- ⑥ If no thread can be scheduled at some state, add the state to the deadlocked states
- ⑥ If no thread can be scheduled at the initial state, there is no **Safe** scheduler for this application

Synthesis Steps



- ⑥ Use the *untimed* model to remove deadlocks
 - ⌚ Obtain a **Deadlock-Free** scheduler for any combination of underlying hardware and algorithms used for the computations
 - ... avoid **dormant** deadlocks
 - ... be **robust** against wrongly estimated execution durations
 - ... **know** if you got it **really** wrong ...
 - ⌚ Do so using a much smaller model than the **original timed** one
 - ⌚ Use the **Deadlock-Free** scheduler to constrain the **timed** model
 - ⌚ Find additional constraints for avoiding missed deadlines

State Space Reduction



⑥ **Branching Bisimulation Equivalence (bbe)**

reduction

- ④ Series of consecutive ticks are substituted by a single “super” tick
- ④ A **safe reduction**, since we could not perform scheduling anywhere inside such a series in the first place
- ④ We only pay for the ticks we can **observe** ...
- ④ A **74%** reduction of the state space on a small example

State Space Reduction - II

⑥ Non-Preemption

- ③ When an alarm fires, while some thread is computing, we never preempt it
- ③ Conservative reduction - we may fail to find a scheduler
- ③ Sometimes this is the reality (messages in networks)
- ③ If negative, the system is **overloaded** and we identify the set of problematic threads
- ③ A **40%** reduction (on its own), **80%** when combined with *bbe*

⑥ Use the constraints found in the previous steps (deadlocks, non-preemption) to constrain the system when preemption is allowed

Introduction

Goal

Thread States

Architecture

Layers

System Model

X Modes

Synthesis Steps

SS Reduction

No Clocks

Example

Implementation

Conclusions

Future

Synthesised Scheduler



- ⑥ Constraints : $\overrightarrow{PC}_i \wedge \overrightarrow{Clocks}_i \Rightarrow$ Do not execute $\overrightarrow{Thr}_{ij}$
(where i is a state in the state space)
- ⑥ The scheduler needs to examine the PC 's and the Clocks of all the threads to identify the current state i

⑥ Reality Check :

- ⚡ Observing clocks is **costly**
- ⚡ Using clocks (*i.e.*, **watchdogs**) is not for free either
- ⚡ Can we avoid paying this overhead ?
 - ... Try to find a set of constraints which are **independent** of time, *i.e.*,
 $\overrightarrow{PC}_i \Rightarrow$ Do not execute $\overrightarrow{Thr}'_{ij}$

- Introduction
- Goal
- Thread States
- Architecture
- Layers
- System Model
- X Modes
- Synthesis Steps
- SS Reduction
- No Clocks
- Example
- Implementation
- Conclusions
- Future

No Clocks



⑥ Removing clocks from constraints can cause deadlines to be missed

⑥ So, remove clocks & synthesise again (by now the system is reduced ~89%)

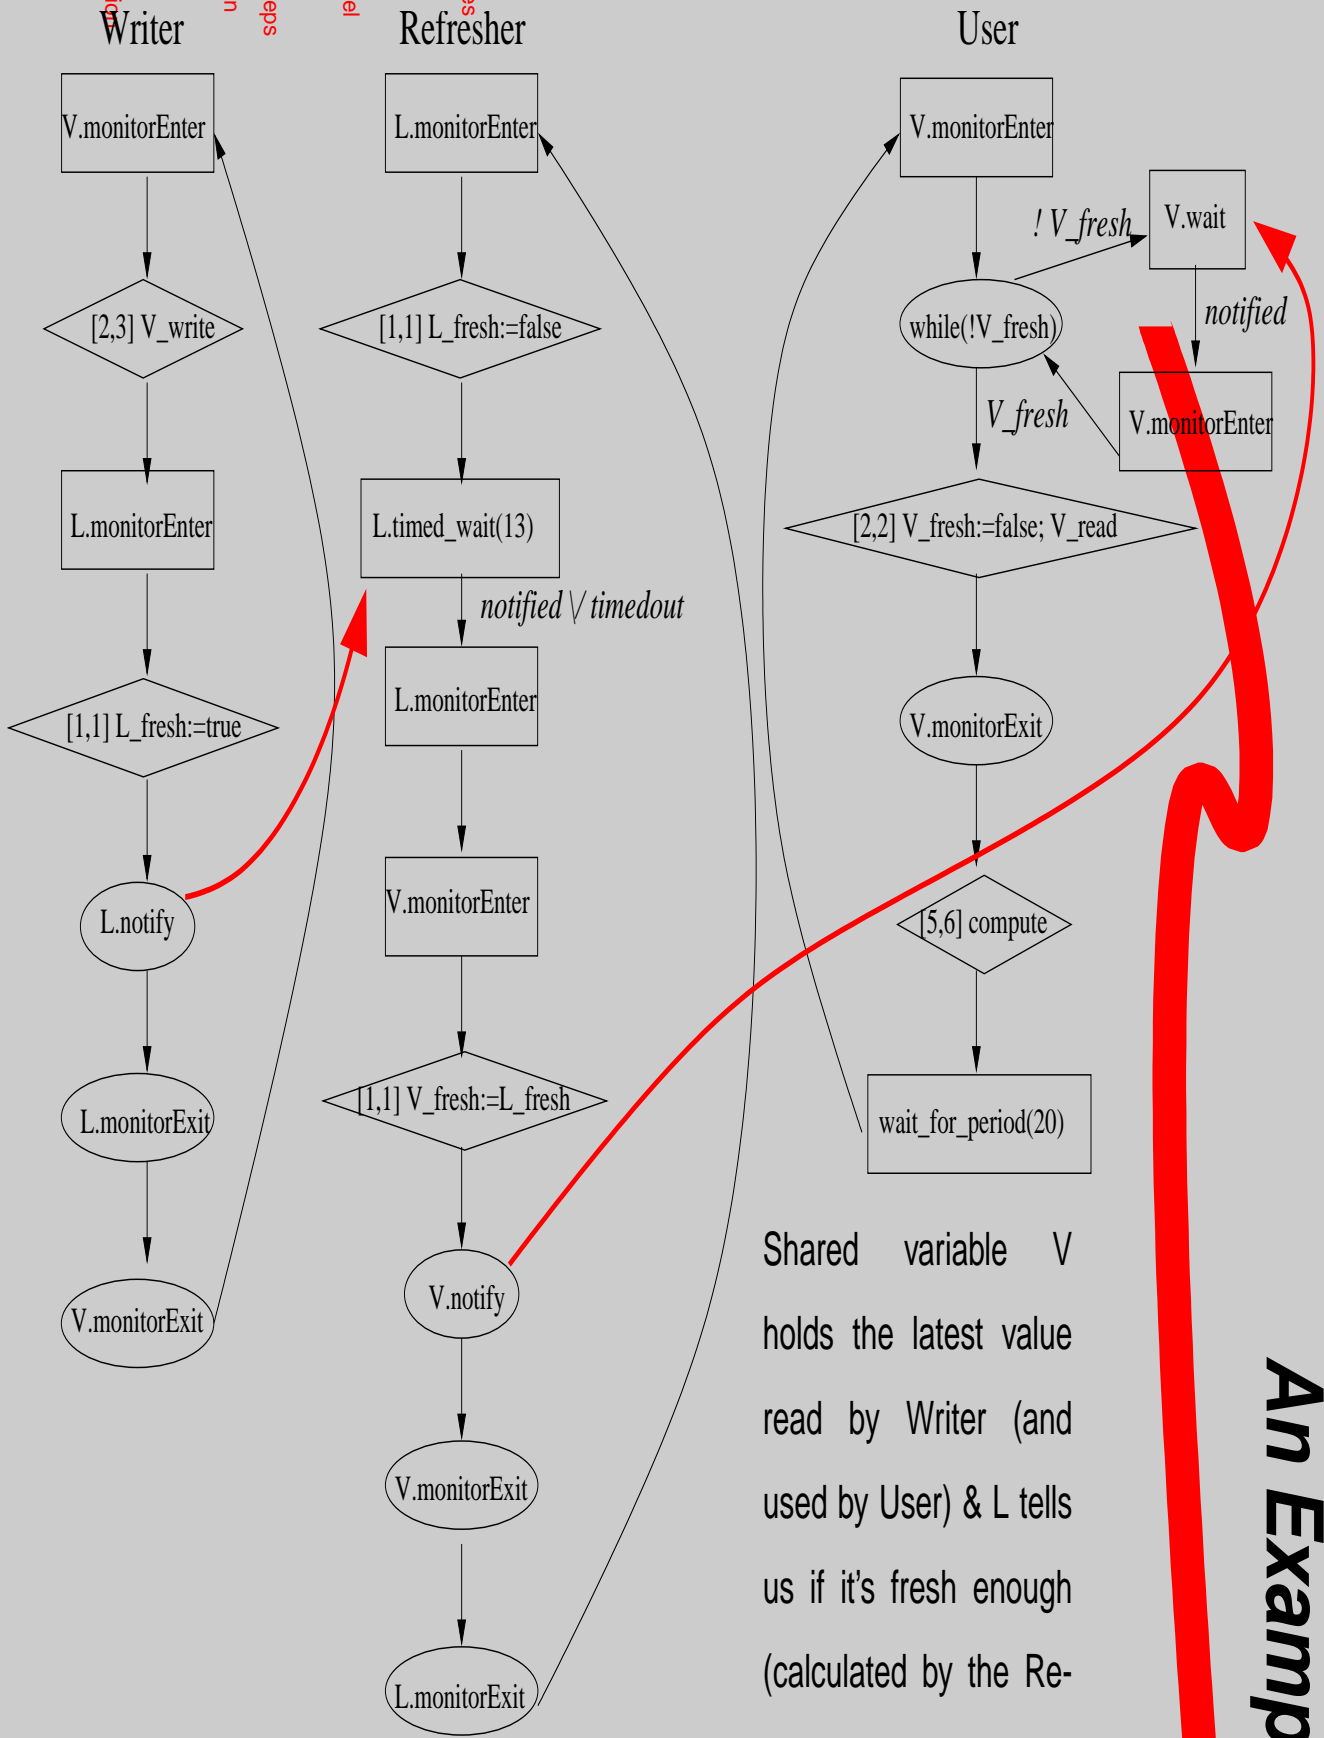
⑥ If it fails what ?

☞ The application is not **robust** with respect to time !

☞ Either :

... Rewrite it (change algos / hardware); or

... Use the safe scheduler which examines the clocks but make sure it ain't too slow !



Shared variable V holds the latest value read by Writer (and used by User) & L tells us if it's fresh enough (calculated by the Refresher)

An Example

Example (continued)

- ⑥ Original Timed Model: 45.5 K States
- ⑥ 10 states in the untimed model are Deadlock states
- ⑥ It Misses a Deadline in 367 states in the timed model
- ⑥ Synthesised :
 - ③ 8 constraints for avoiding Deadlocks
 - ③ 30 constraints for guaranteeing Deadlines when QoS = Non-Preemption
 - ③ 18 additional constraints when allowing Preemption & no longer observing clocks
- ⑥ 56 constraints in total – final **safe** state space has 1.6 K states

Implementation



- ⑥ From Java application to abstract model : By hand
- ⑥ From abstract to detailed model (slicing, etc.) : AWK
- ⑥ CADP tools for the *bbe* reduced state space
- ⑥ Synthesis algorithm : Lisp
- ⑥ Constraints to C code : AWK script
- ⑥ Basic library for eCos: C
- ⑥ 330 MHz Pentium II – Min/Avg/Max/Avg-Dev (in μs)
 - ⌘ **Schedule** : **0.00 / 0.66 / 4.00 / 0.45**
 - ⌘ Context Switch : **0.00 / 0.77 / 1.00 / 0.35**
 - ⌘ Trylock (unlocked) : **0.00 / 0.69 / 2.00 / 0.47**
 - ⌘ Unlock (locked) : **0.00 / 0.75 / 3.00 / 0.47**

- Introduction
- Goal
- Thread States
- Architecture
- Layers
- System Model
- X Modes
- Synthesis Steps
- SS Reduction
- No Clocks
- Example
- Implementation
- Conclusions
- Future

Conclusions

Weak points :

- Initial model hand-constructed – some early results let us believe we can solve this
- We give the CPU to a thread by **stopping** the rest & lock the underlying scheduler at certain points – an implementation with priorities in user space

Strong points :

- Completely automatic analysis – no need for RMA, selecting priorities by hand, *etc.*
- Directly usable with **heterogeneous** applications
- Helps in **analysing** an application – separates constraints for deadlocks, deadlines (under non-preemption, preemption)
- Scheduling is easily **extendible** with QoS policies



- ⑥ Forward, On-the-fly synthesis
- ⑥ Currently away in the mist :
 - ⌚ Scheduling for memory – minimise total memory needs / minimise cache misses
 - ⌚ Scheduling for energy – minimise / stabilise energy consumption
 - ⌚ Multiprocessor systems, with different clock speeds
 - ⌚ Automate WCET estimation of basic blocks
 - ⌚ **Simple** Model-checker – the time automaton can be an observer
- ⑥ **“Optimum”** Schedulers – variables, order, cost, *etc.*