

ReconOS: An RTOS Supporting Hardware and Software Threads

Enno Lübbers and Marco Platzner
Computer Engineering Group
University of Paderborn

`{enno.luebbers, platzner}@upb.de`



UNIVERSITÄT PADERBORN
Die Universität der Informationsgesellschaft

Operating Systems for Reconfigurable Hardware

- traditional approaches integrate hardware accelerators as slave coprocessors

- Linux-based integration of reconfigurable logic
 - u microblaze-ucLinux [Bergmann et. al. 2006]
 - preferred communication through FIFOs
 - u BORPH [So et. al. 2006]
 - file-system based communication between hardware and software

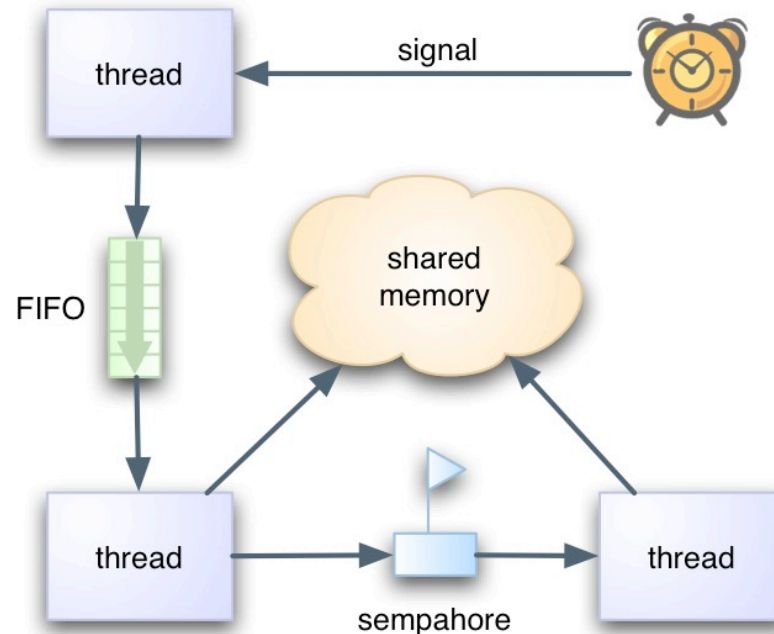
- unified programming model for software and hardware threads
 - u hthreads [Peck et. al. 2006]
 - hardware threads generated from multithreaded C-source
 - OS functionality realized in hardware
 - u **ReconOS**
 - based on software RTOS (eCos)
 - hardware threads are written in VHDL

Outline

- motivation
- programming model
 - u operating system objects
 - u hardware thread design
- execution model
 - u system architecture
 - u OS call delegation
 - u toolchain
- experimental results
 - u operating system overheads
 - u case study
- conclusion & outlook

RTOS-like Programming Model

- applications are modelled with a set of objects
 - u tasks/threads, semaphores, FIFOs, shared memory, timers, etc.



- u these objects, their semantics and possible relationships form the programming model

RTOS-like Programming Model

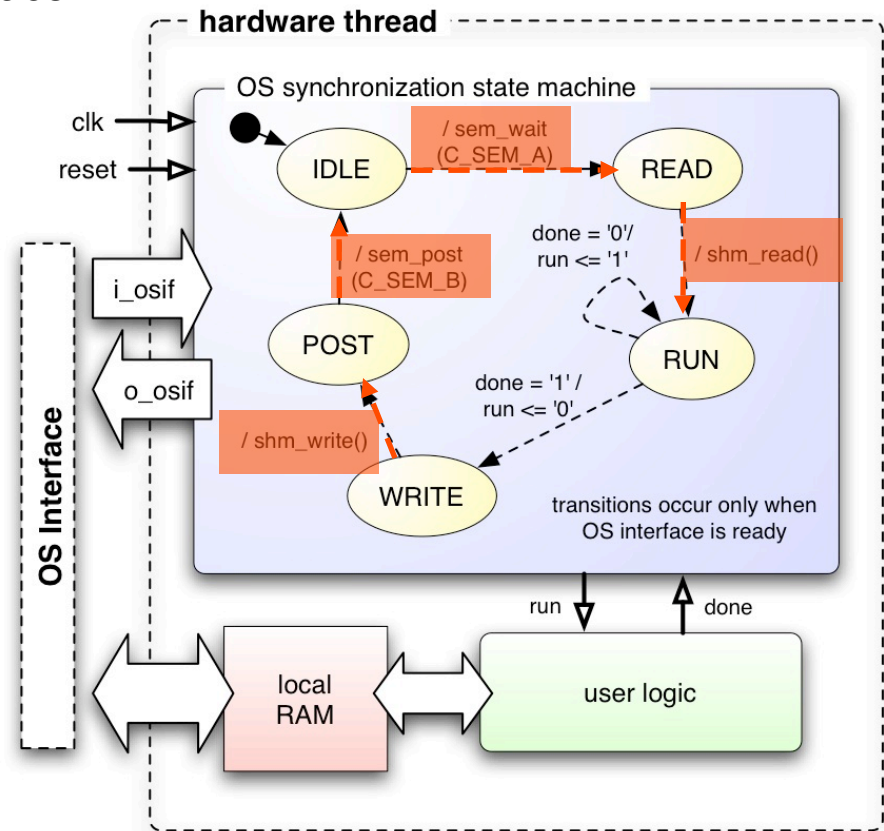
- classic (embedded) software implementation
 - u threads interact with the OS through API functions
 - eg. semaphore_post(), thread_create(), malloc()
 - distinction between blocking and non-blocking calls
 - u sequential execution of threads

- challenges in translating this model to hardware
 - u hardware is inherently parallel
 - "hardware thread" is actually a misleading term
 - hardware has no notion of function calls or even blocking function calls
 - u parallel execution of several hardware threads and one software thread
 - SW-HW and HW-HW synchronization and communication
 - scheduling

- ReconOS approach: extend software RTOS
 - u hardware threads with OS synchronization state machine
 - u delegate threads

ReconOS Hardware Threads

- a hardware thread consists of two parts
 - u an OS synchronization state machine
 - synchronizes thread with operating system calls
 - serializes access to OS objects via the OS interface
 - can be blocked by the OS interface
 - u parallel “user processes”
 - communicate with OS synchronization state machine
 - can directly access local memory blocks
 - are not necessarily blocked



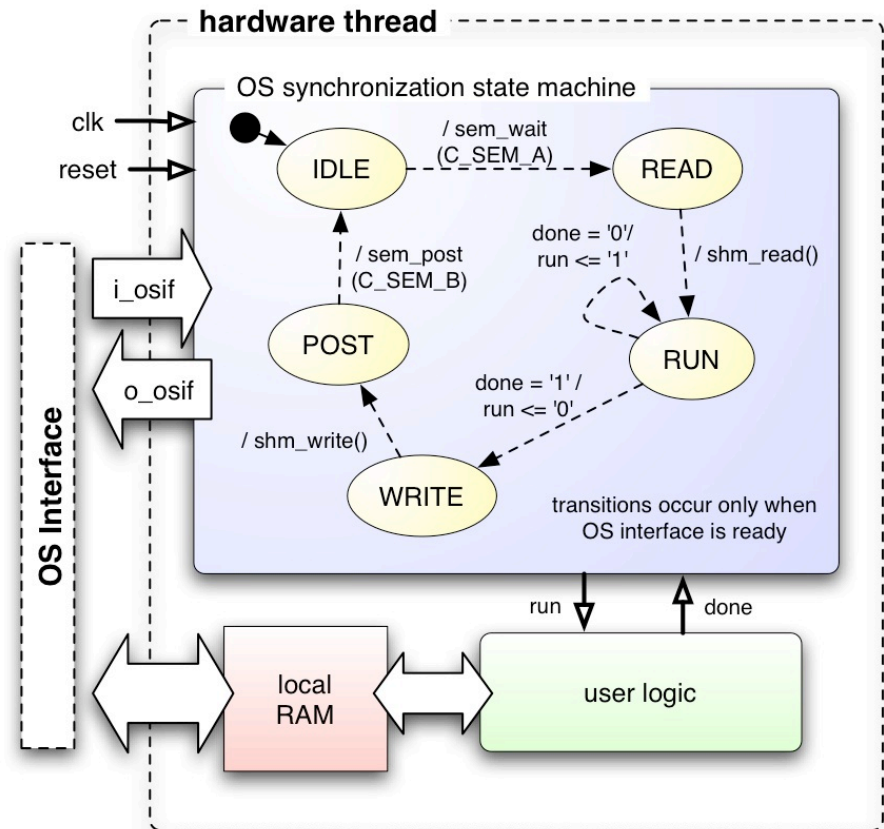
ReconOS API for Hardware Threads

```

1  osif_fsm: process(clk, reset)
2  begin
3    if (reset = '1') then
4      state <= IDLE;
5      run <= '0';
6      reconos_reset(o_osif, i_osif);
7    elsif rising_edge(clk) then
8      reconos_begin(o_osif, i_osif);
9      if reconos_ready(i_osif) then
10     case state is
11     when IDLE =>
12       reconos_sem_wait(o_osif, i_osif, C_SEMA);
13       state <= READ;
14
15     when READ =>
16       reconos_shm_read_burst(o_osif, i_osif,
17         local_address,
18         global_address);
19
20       state <= RUN;
21
22     when RUN =>
23       run <= '1';
24       if done = '1' then
25         run <= '0';
26         state <= WRITE;
27       end if;
28
29     when WRITE =>
30       reconos_shm_write_burst(o_osif, i_osif,
31         local_address,
32         global_address);
33
34       state <= POST;
35
36     when POST =>
37       reconos_sem_post(o_osif, i_osif, C_SEM_B);
38       state <= IDLE;
39
40     when others => null;
41     end case;
42   end if;
43 end if;
44 end process;

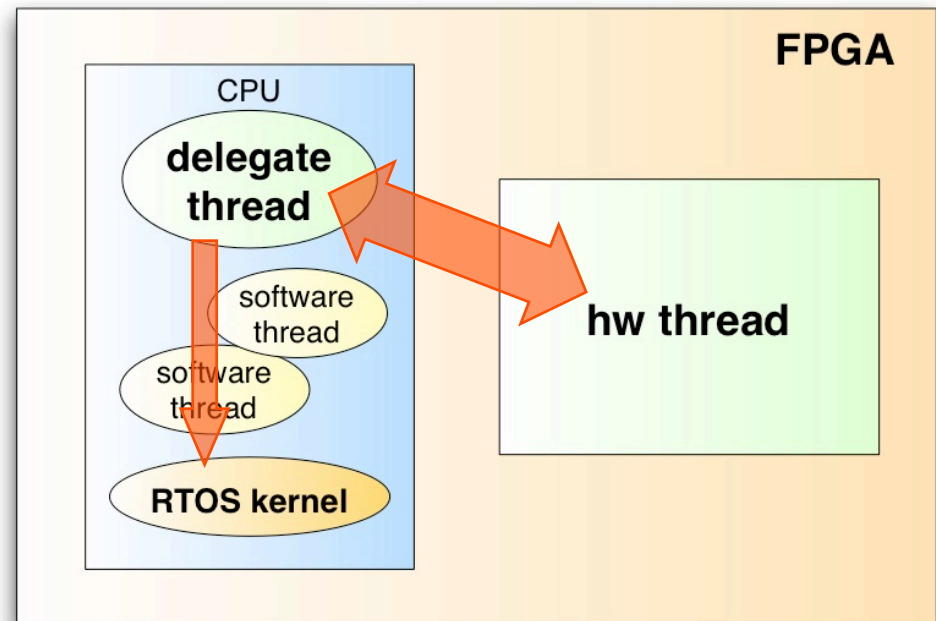
```

- u VHDL function library
- u may only be used in the OS synchronization state machine

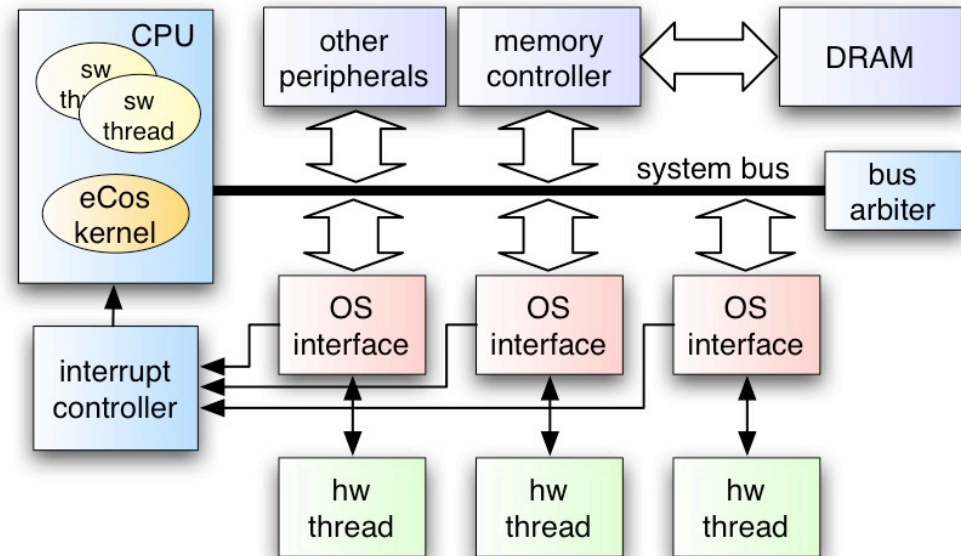



Delegate Threads

- basic mechanism
 - u a delegate thread in software is associated with every hardware thread
 - u the delegate thread calls the OS kernel on behalf of the hardware thread
 - u all kernel responses are relayed back to the hardware thread
- advantages
 - u no modification of the kernel required
 - u extremely flexible
 - u transparent to kernel and other threads
- drawbacks
 - u increased overhead due to interrupt processing and context switch



System Architecture



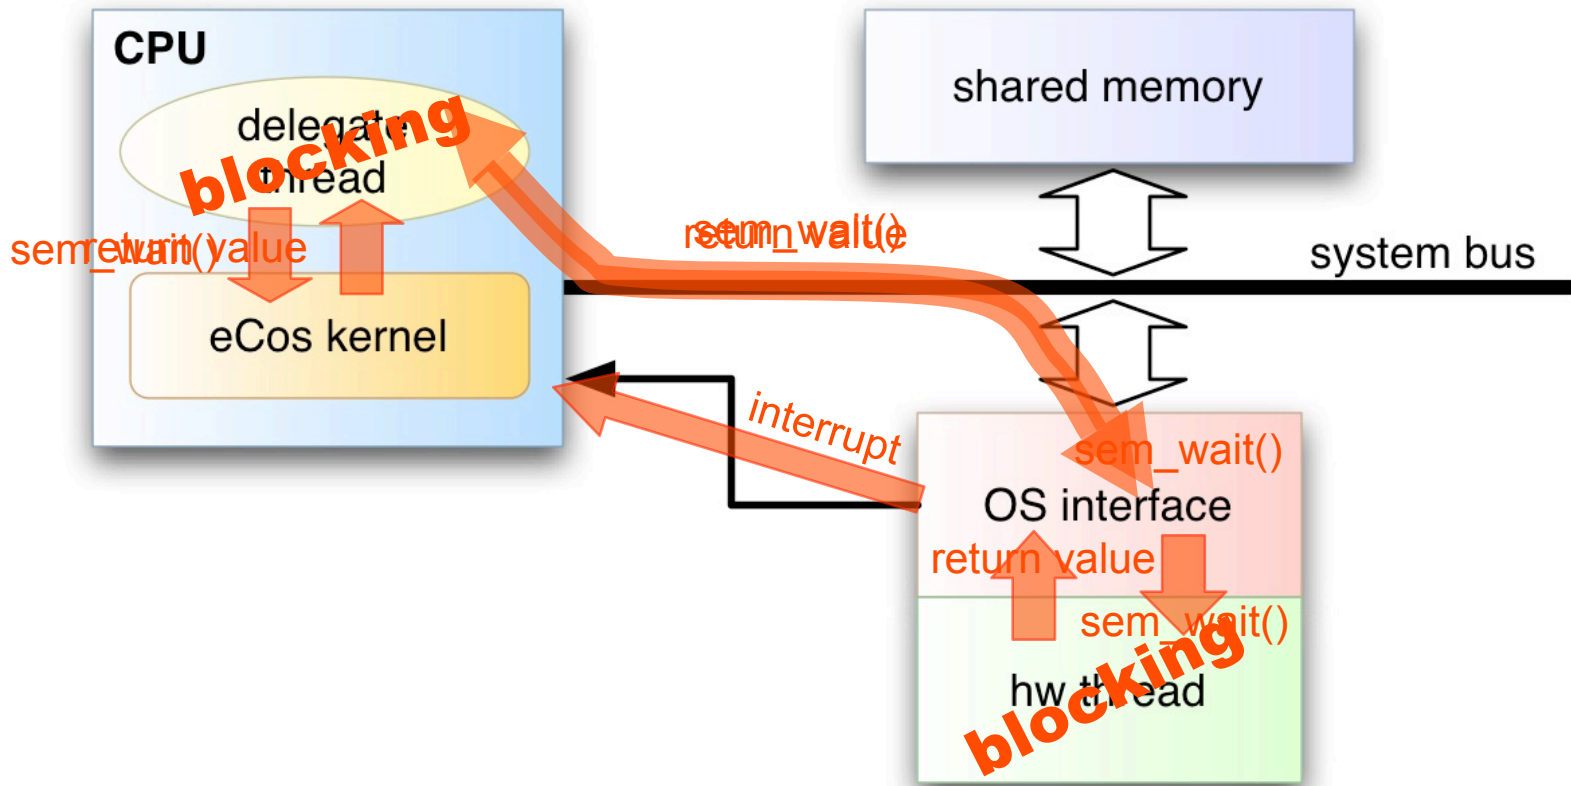
■ development platforms

- u Xilinx ML403 (Virtex-4FX)
- u Xilinx XUPV2P (Virtex-II Pro)
- u embedded PowerPC 405 CPU(s)
- u CoreConnect bus architecture
- u FPGAs support partial reconfiguration

■ real-time operating system

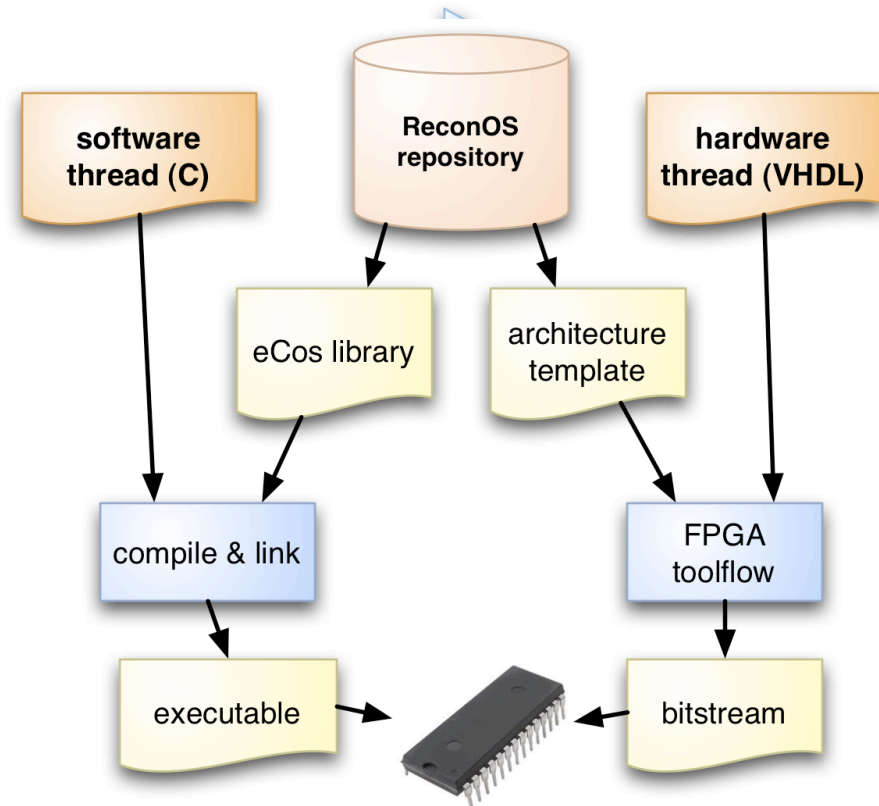
- u eCos for PowerPC ported to development platforms
- u eCos is a widely-used open source RTOS
- u modular, extensible design
- u supplemented with OS interface for hardware threads

OS Call Implementation



Toolchain

- software threads are written in C
 - u using the eCos software API
- hardware threads are written in VHDL
 - u using the ReconOS VHDL API
- architecture generation
 - u automatically inserts OS interfaces and hardware threads into Xilinx EDK platform templates
 - u configures and builds static eCos library
- eCos extensions
 - u hardware thread object encapsulating delegate thread and OS interface “driver”
 - u profiling support to track the state of the hardware threads' OS synchronization state machines

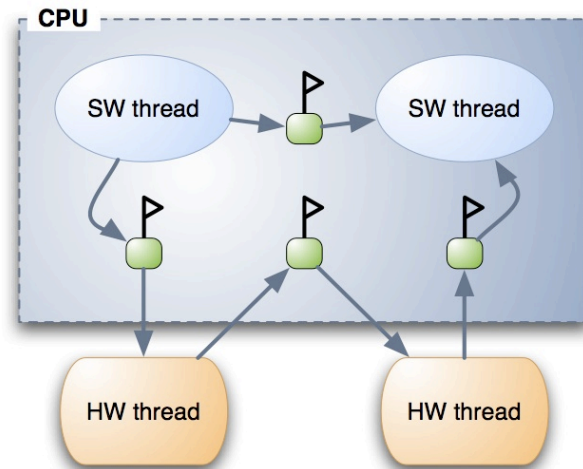


OS Overheads

- synthetic hardware and software threads
 - u semaphore processing time (post → wait)
 - u time for non-blocking OS calls (i.e. `reconos_sem_post()`)
 - u OS interface takes 1051 slices (7% of XC2VP30)

- OS calls involving hardware exhibit higher latencies
 - u additional context switch to delegate
 - u interrupt processing
 - u bus access vs. cache access

- limited impact on system performance
 - u logic resources mainly used for heavy data-parallel processing
 - u less synchronization-intensive control dominated code



Semaphores (post → wait)

SW-to-SW	7.69 μ s
SW-to-HW	13.84 μ s
HW-to-SW	27.13 μ s
HW-to-HW	34.19 μ s
non-blocking OS call	
SW	1.59 μ s
HW	16.51 μ s

Case Study - Image Processing Filter

■ three threads


- u capture image from Ethernet
- u apply LaPlacian filter
- u display image on VGA monitor

■ platform

- u Xilinx XUPV2P (Virtex-II Pro)
- u PPC @ 300MHz, rest @ 100MHz

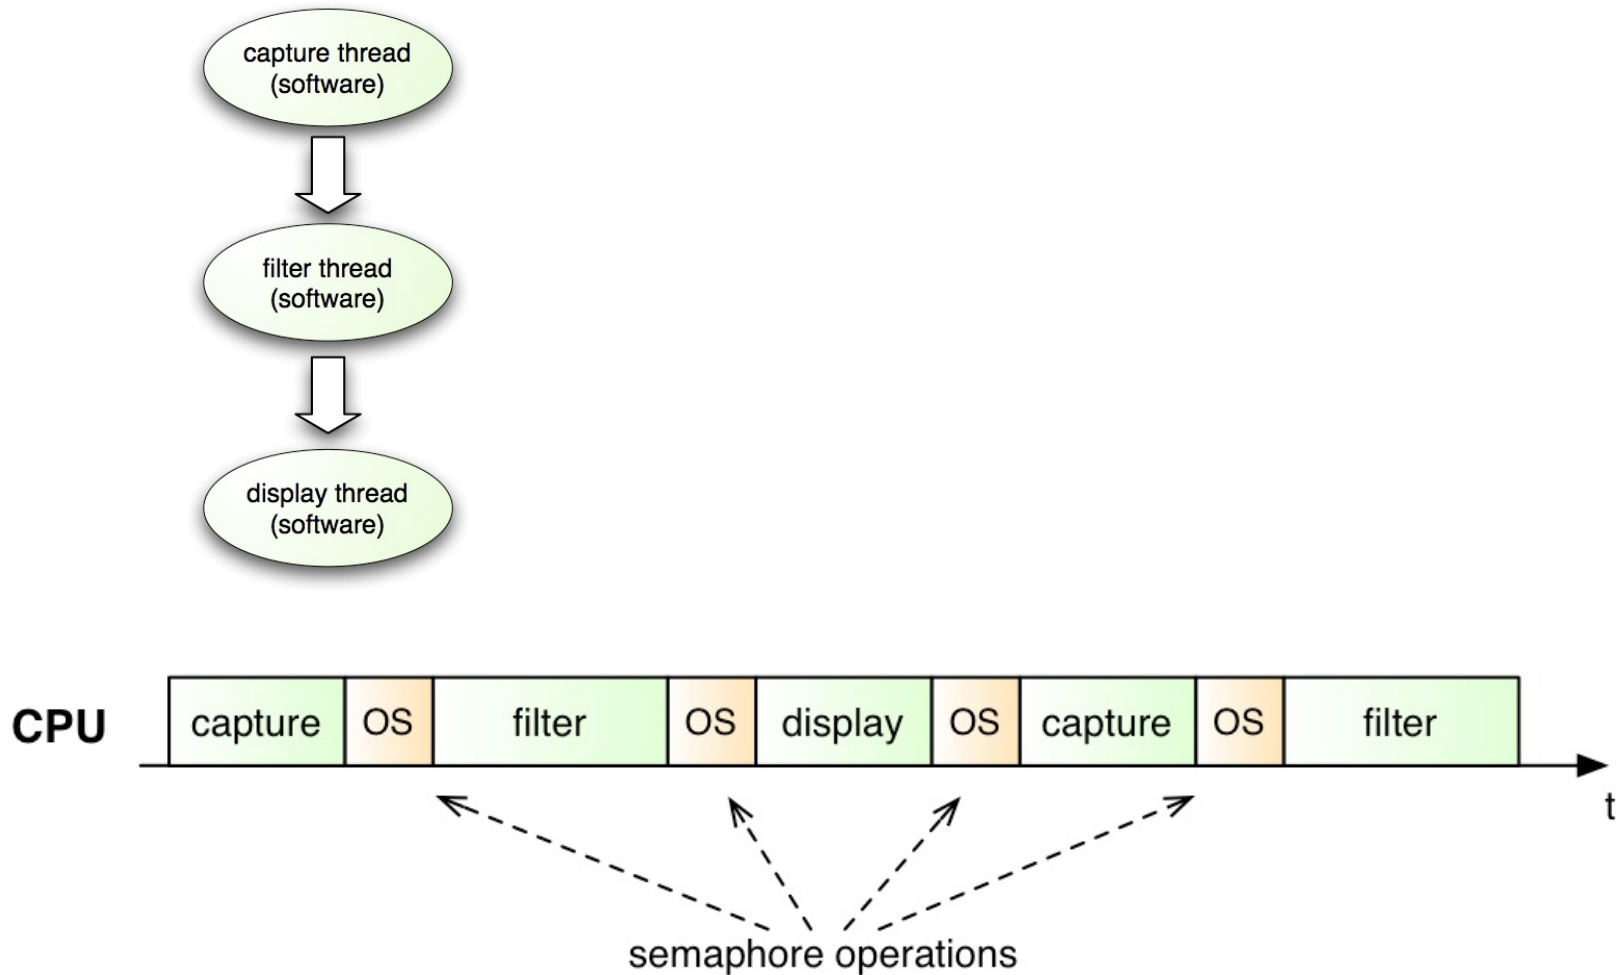
■ threads communicate through shared memory

- u image resolution: 320x240 pixels, 8 bit greyscale
- u image data organized into blocks (e.g. 40 lines = 1 block)
- u a block is protected by two semaphores
 - “ready” semaphore: data can be safely written into this block
 - “new” semaphore: new data is available in this block



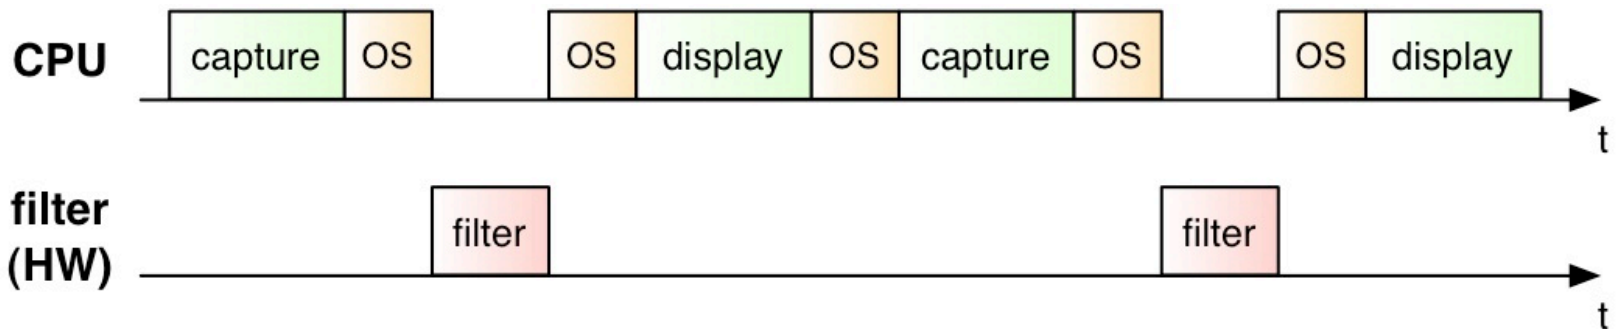
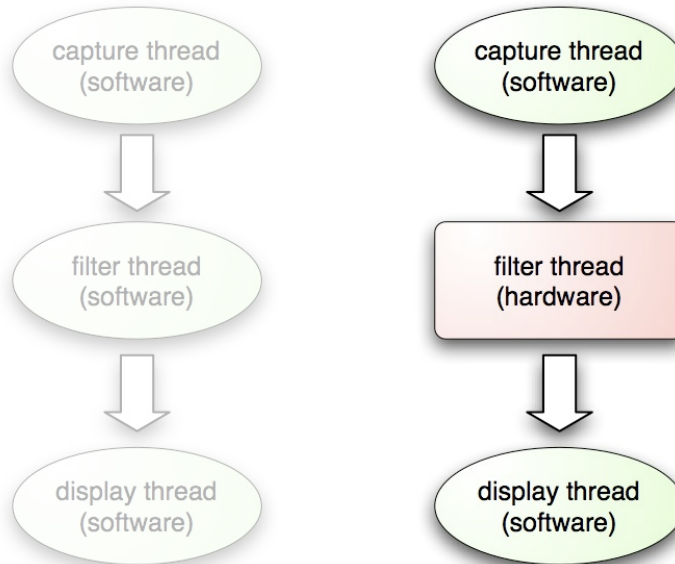
Case Study - Implementation #1

- all threads in software
 - u all computations occur sequentially, with low OS overhead



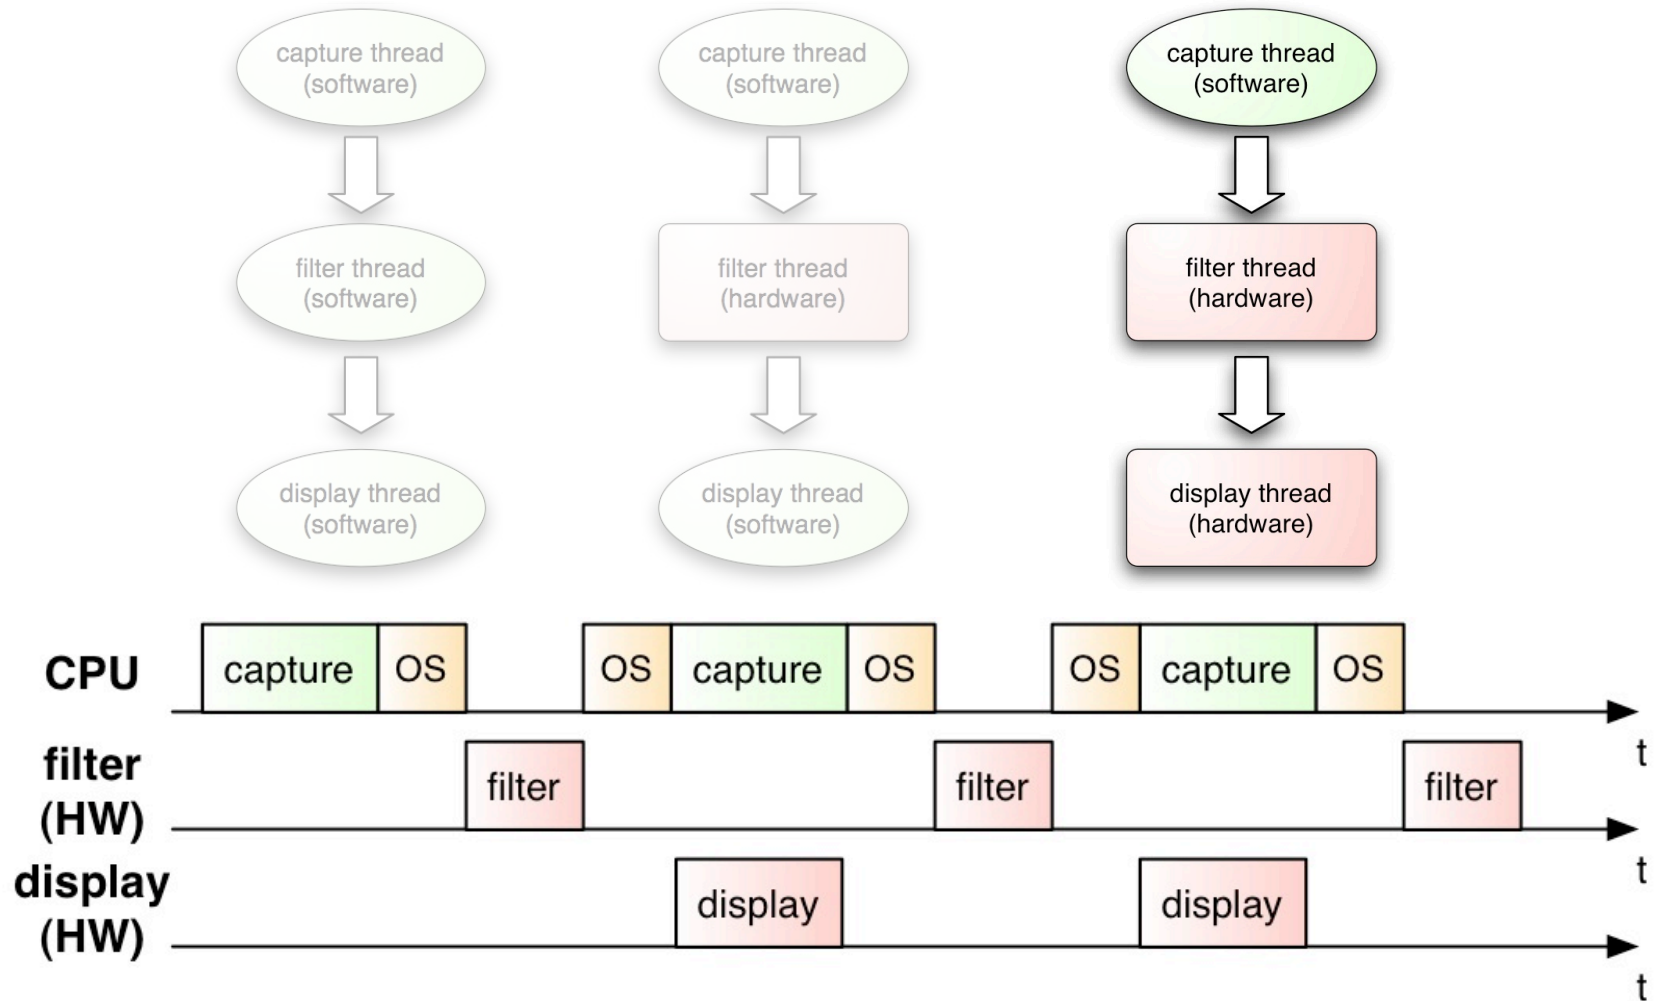
Case Study - Implementation #2

- move filter thread to hardware
 - u convolution filters allow for efficient parallelization



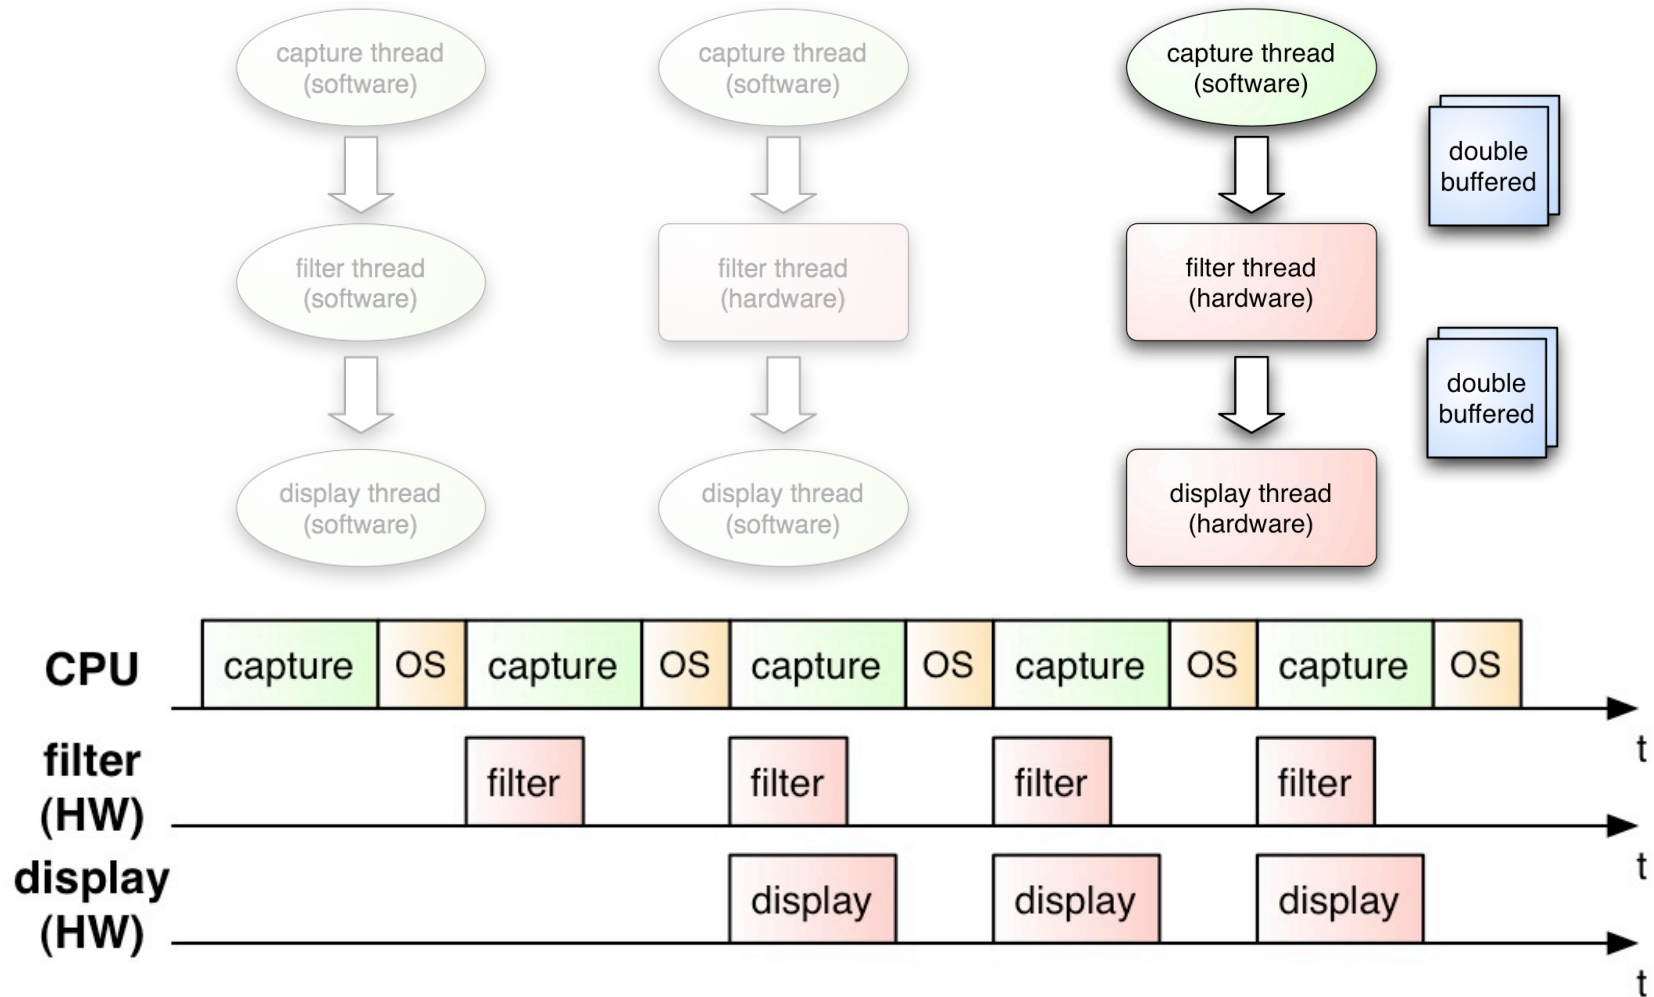
Case Study - Implementation #3

- move also display thread to hardware
 - u display thread can output data concurrently with capture thread




Case Study - Implementation #4

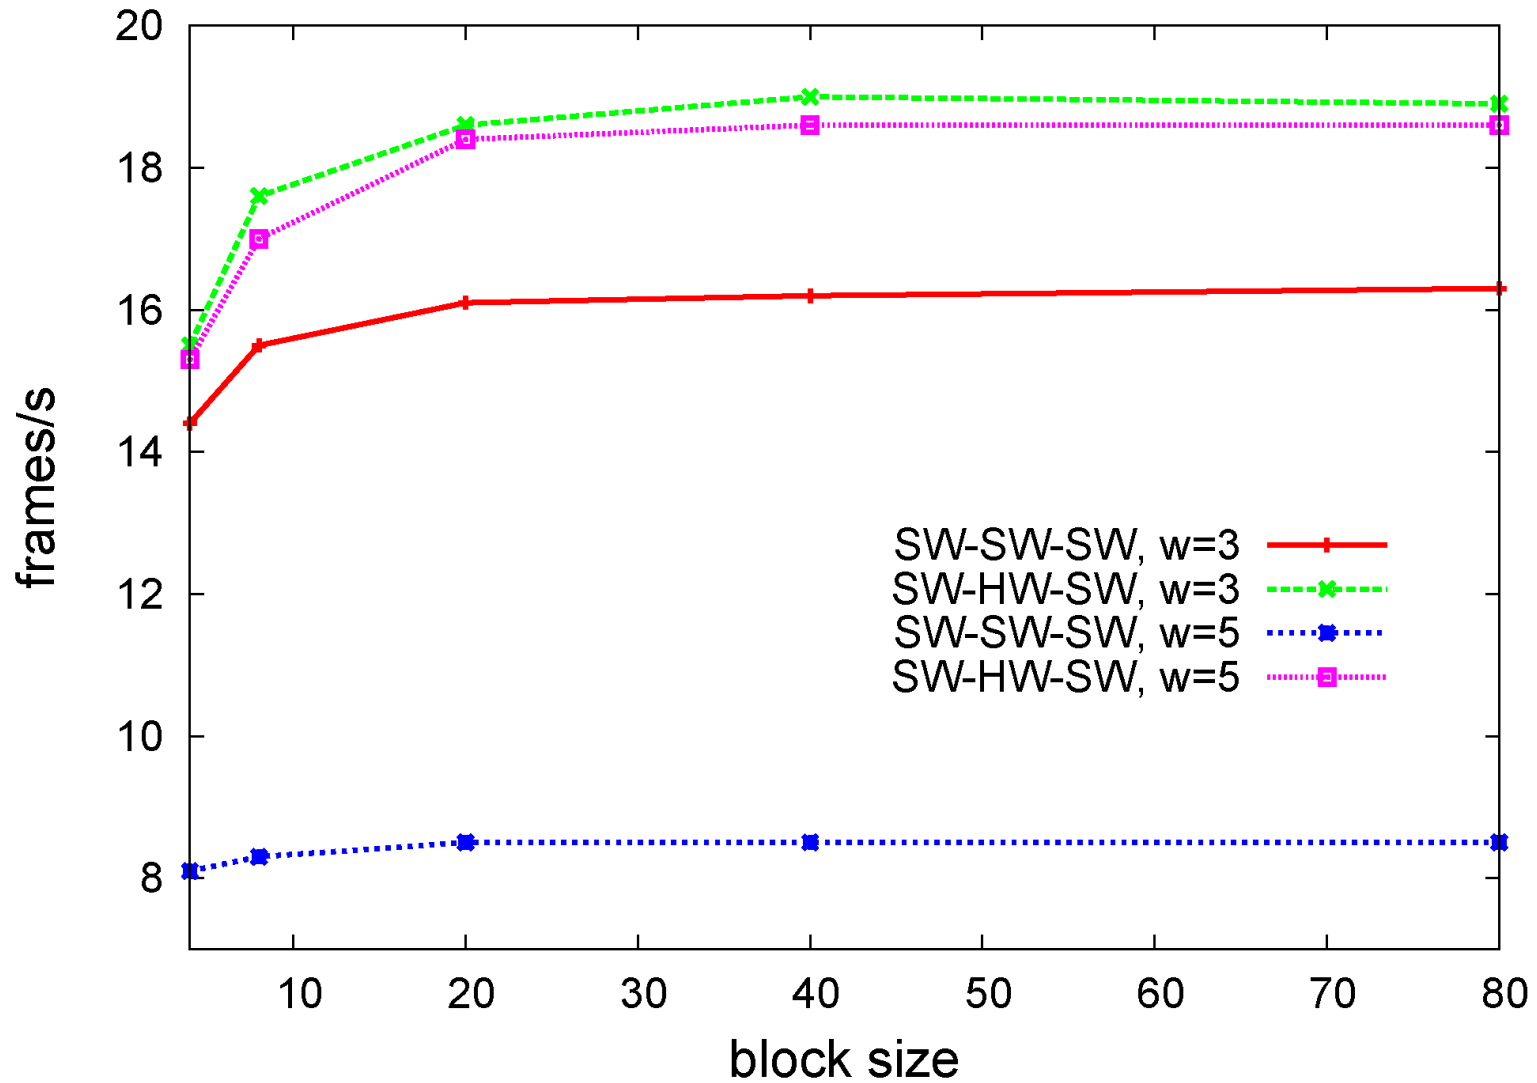
- parallel hardware threads
 - double-buffer image data



Case Study - Results



Case Study - Results



Conclusion & Outlook

- RTOS for hardware and software threads
 - u unified programming model
 - transparent synchronization and communication between hardware and software threads
 - u RTOS-centric execution model
 - extended eCos with support for hardware threads
 - u case study

- ongoing work
 - u include partial reconfiguration
 - extend eCos scheduler
 - preemption, task migration

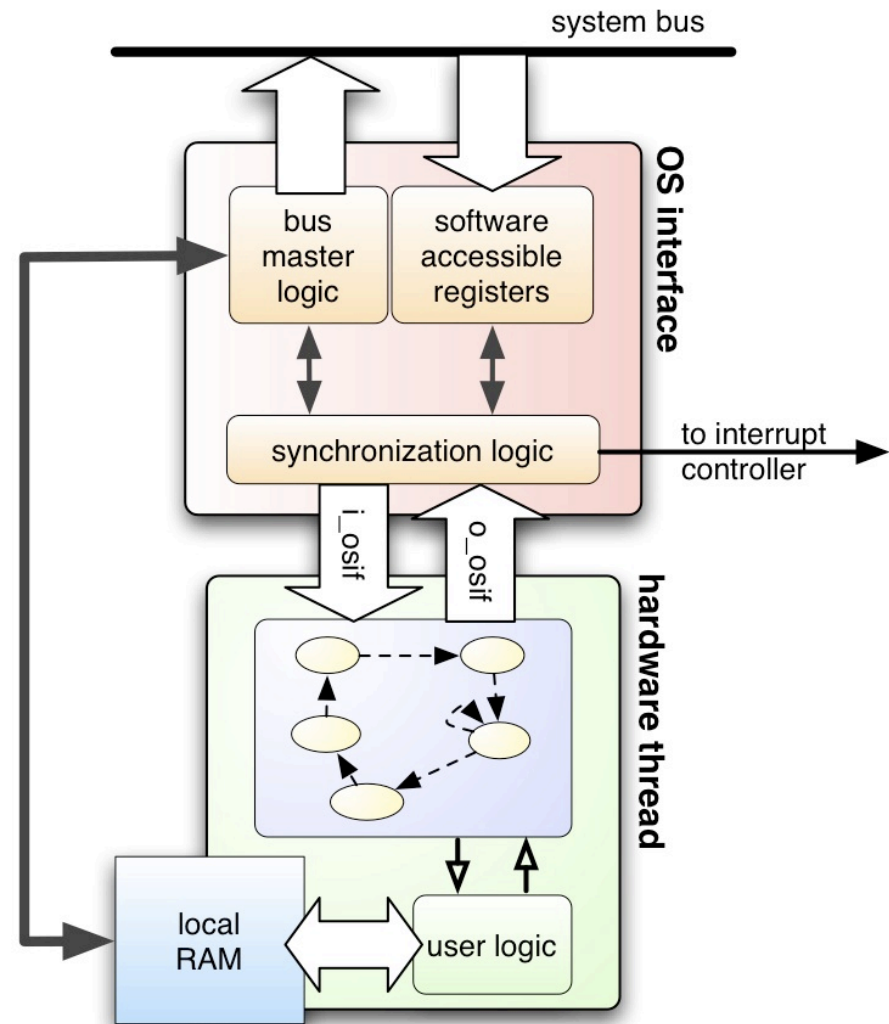
 - u additional platforms
 - Erlangen Slot Machine (ESM)

Thank you

www.reconos.de

OS Interface Implementation

- processes requests from hardware thread
 - u handles blocking and resuming of hardware thread
- master interface
 - u memory accesses are handled directly
 - u single word and burst transfers
 - u direct access to entire system's address space (memory and peripherals)
- slave interface
 - u OS object interactions are relayed to delegate thread
 - u dedicated CPU interrupt
 - u CPU addressable registers
 - u used for OS communication



OS Overheads

- OS call delays exhibit sporadic glitches
 - u due to unpredictable bus arbitration
 - u fix: use separate communication channel for OS calls and memory access

