HelenOS

Jakub Jermar

presentation for Solaris RPE

Nov 2, 2006

What is HelenOS

- complete OS w/o some minor features
  - e.g block devices, filesystems, networking :-)
- sources available under BSD and in part under GPL

2001 2002 2003 2004 2005 2006

- original code base (JJ)
- ia32 supported
- support for SMP on ia32
- mere kernel without userspace
- closed source
- school assignment

- mips32 architecture
- OS course assignment
- kernel called SPARTAN

- +5 developers recruited
- HelenOS project started
- open source
- project web site

- amd64 architecture (OP)
- ia64 architecture (JJ, JV)
- ppc32 architecture (MD)

- ia32xen architecture (MD)
- sparc64 architecture (JJ)

- MFF UK OS course now also uses HelenOS

What is HelenOS good for

- still under construction and feature incomplete
- space for development and self-realization
- opportunity to learn about OS principles
- opportunity to learn about different processor archs
- education
- test bed for different ideas
- general purpose OS once the missing parts are in place

Characteristics of HelenOS

- three components: boot, kernel and uspace
- microkernel architecture
- programmed in C and in Assembly, scripts in Python
- ~57000 LOC (~17000 is third party software)
- multiplatform: amd64, ia32, ia32xen, ia64, mips32, ppc32, ppc64, sparc64
- grew up on simulators: bochs, ski, pearpc, simics, qemu, msim, vmware, gxemul
- runs natively on amd64, ia32, ia32xen and sparc64
- SMP supported on amd64, ia32 and sparc64

HelenOS boot component

- GRUB on amd64, ia32 and ia32xen
- SILO + custom loader on sparc64
- custom loader on mips32 and ppc32
- Ski loads ia64 kernel and uspace

HelenOS microkernel

- context switching, scheduling
- exceptions/traps/interruptions and IRQ dispatching
- synchronization
- memory management
- IPC

Scheduling in HelenOS

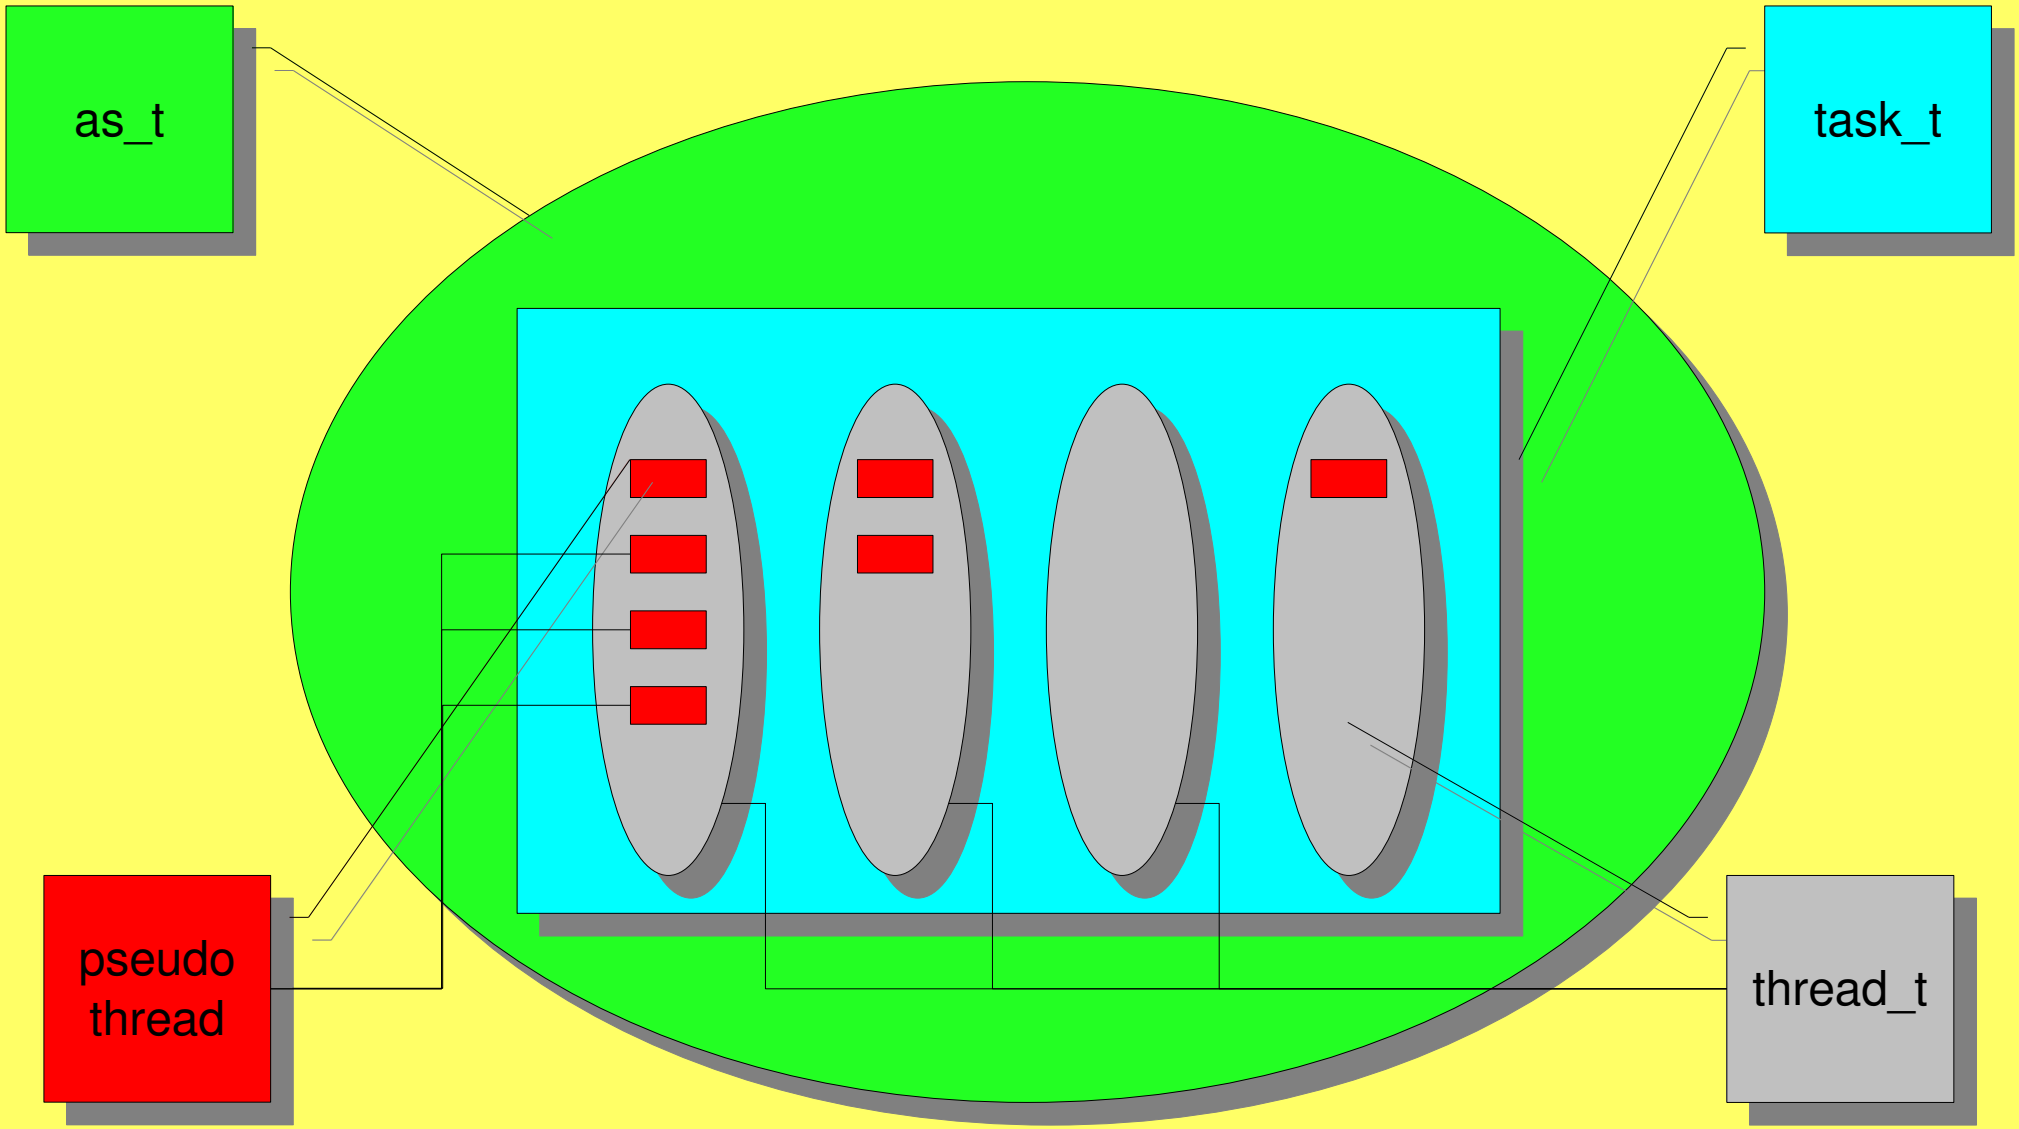

Scheduling in HelenOS

- round robin with multilevel feedback
- each CPU has its own set of runqueues
- load balancing of CPUs
- threads are the schedulable entity
- pseudo threads exist only in uspace
- kernel is preemptible except when holding a spinlock

Synchronization in HelenOS

- spinlocks
- passive synchonization

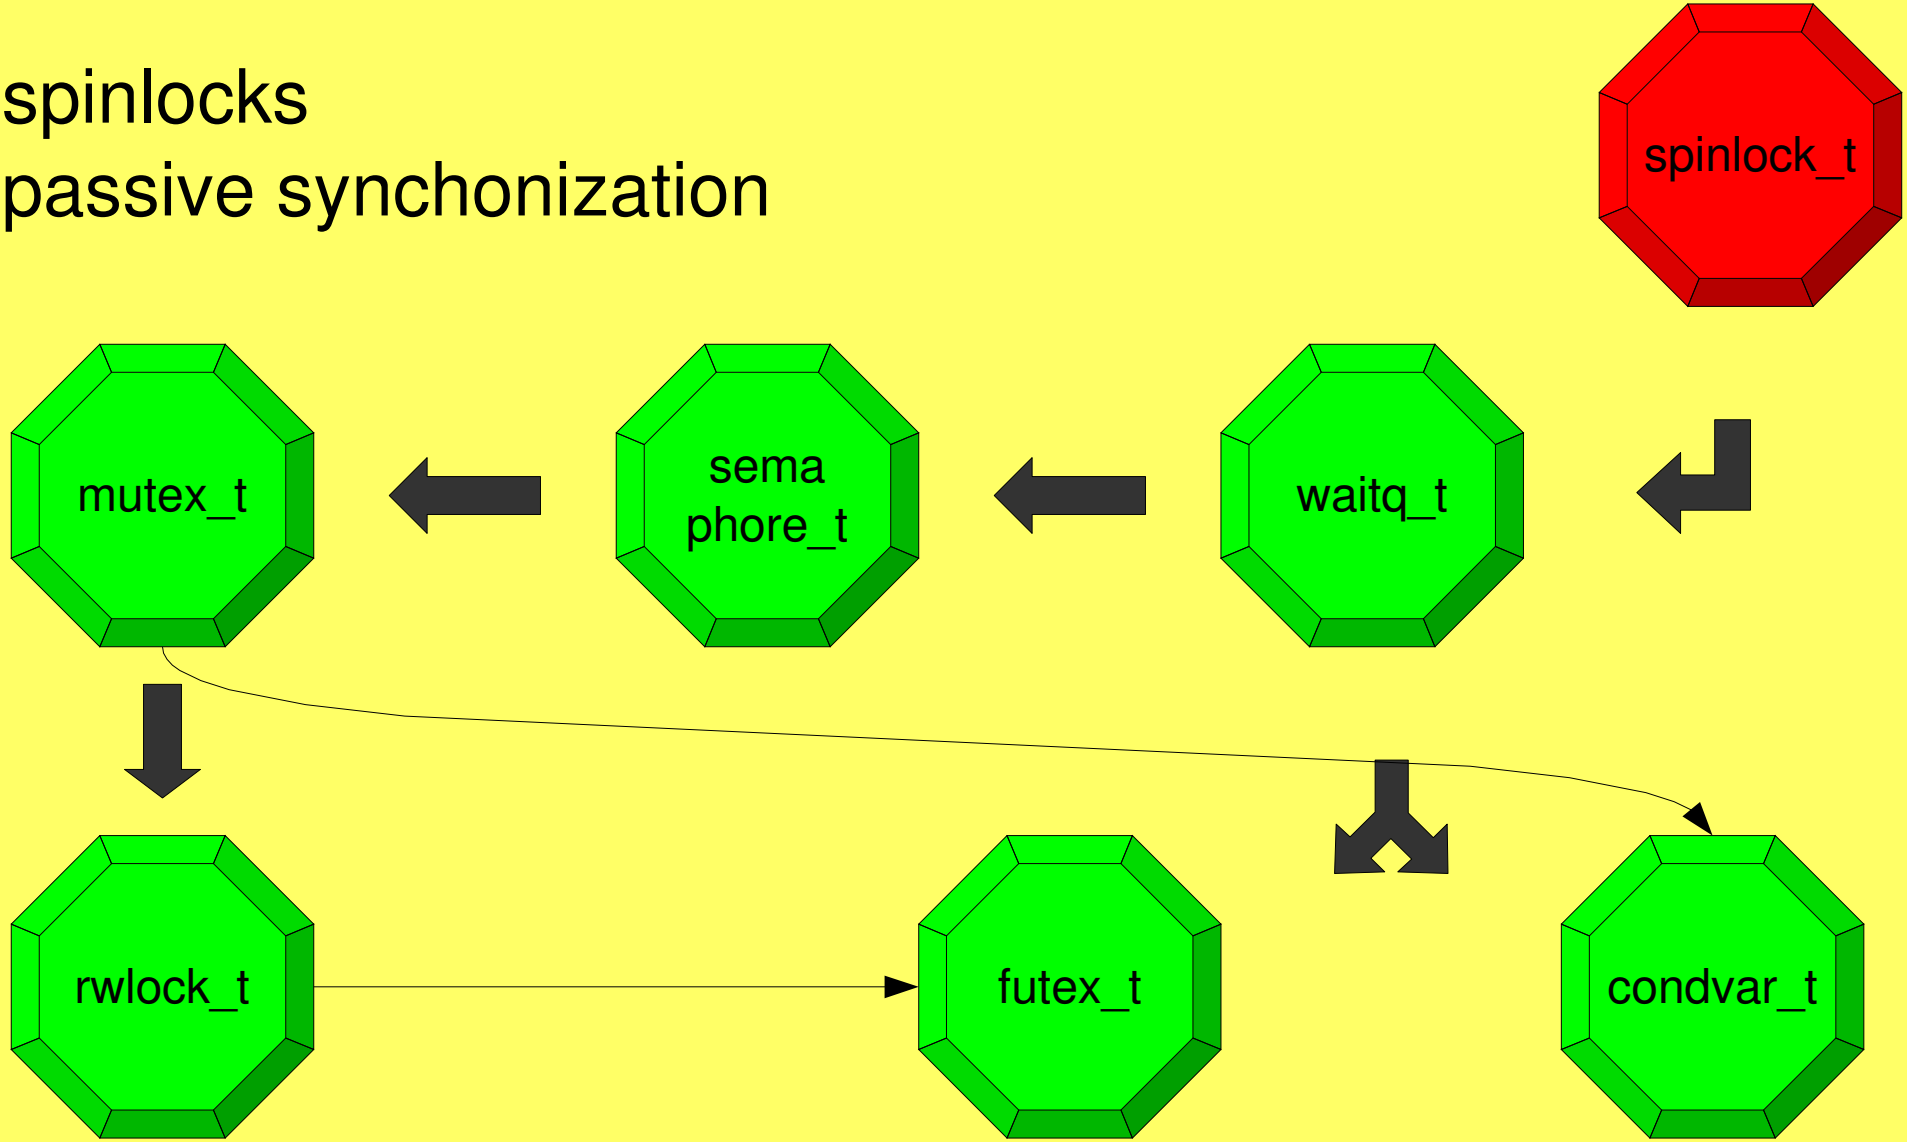

Memory in HelenOS

- „frame allocator“
- physical memory frames organized using buddy system

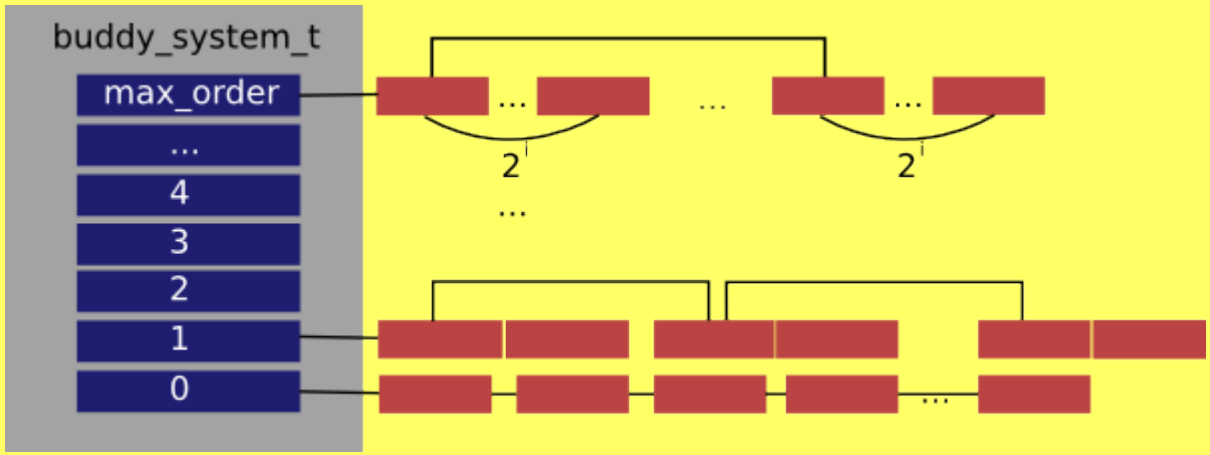

- avoids external fragmentation
- suffers up to 50% internal fragmentation

Memory in HelenOS

- slab allocator
- malloc()/free()

- zero to negligible internal fragmentation
- scalable on SMP

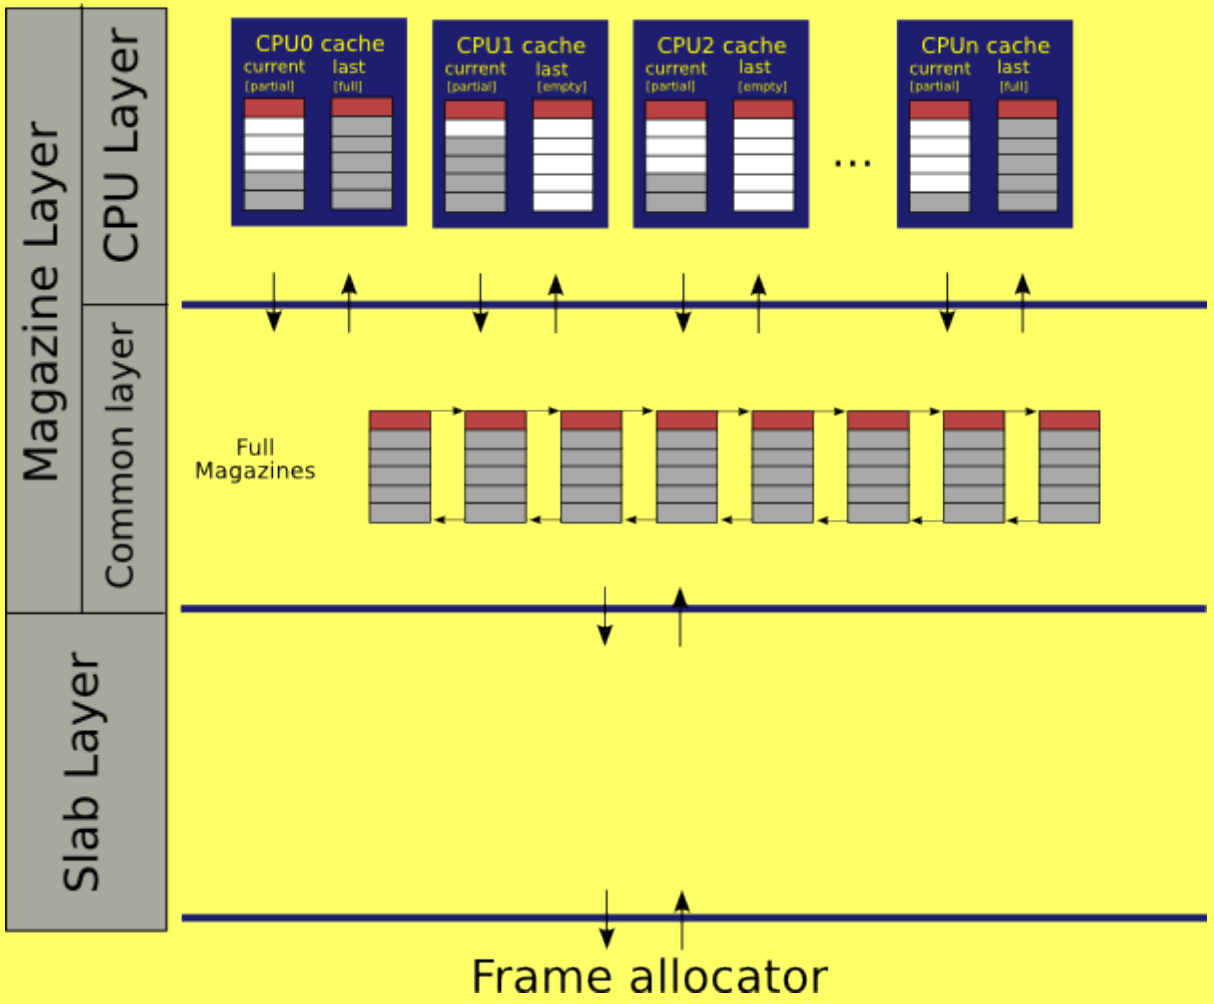

Memory in HelenOS

- virtual memory organized in B+trees

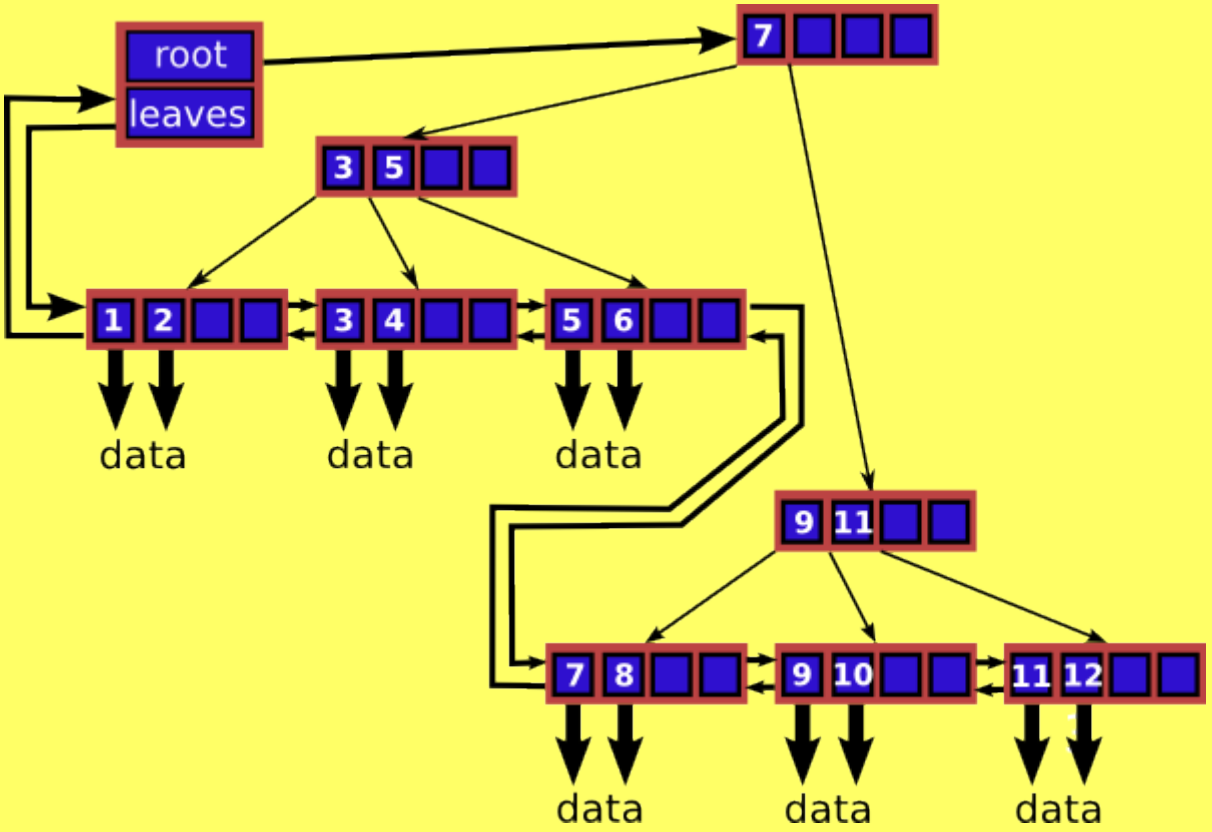

Memory in HelenOS

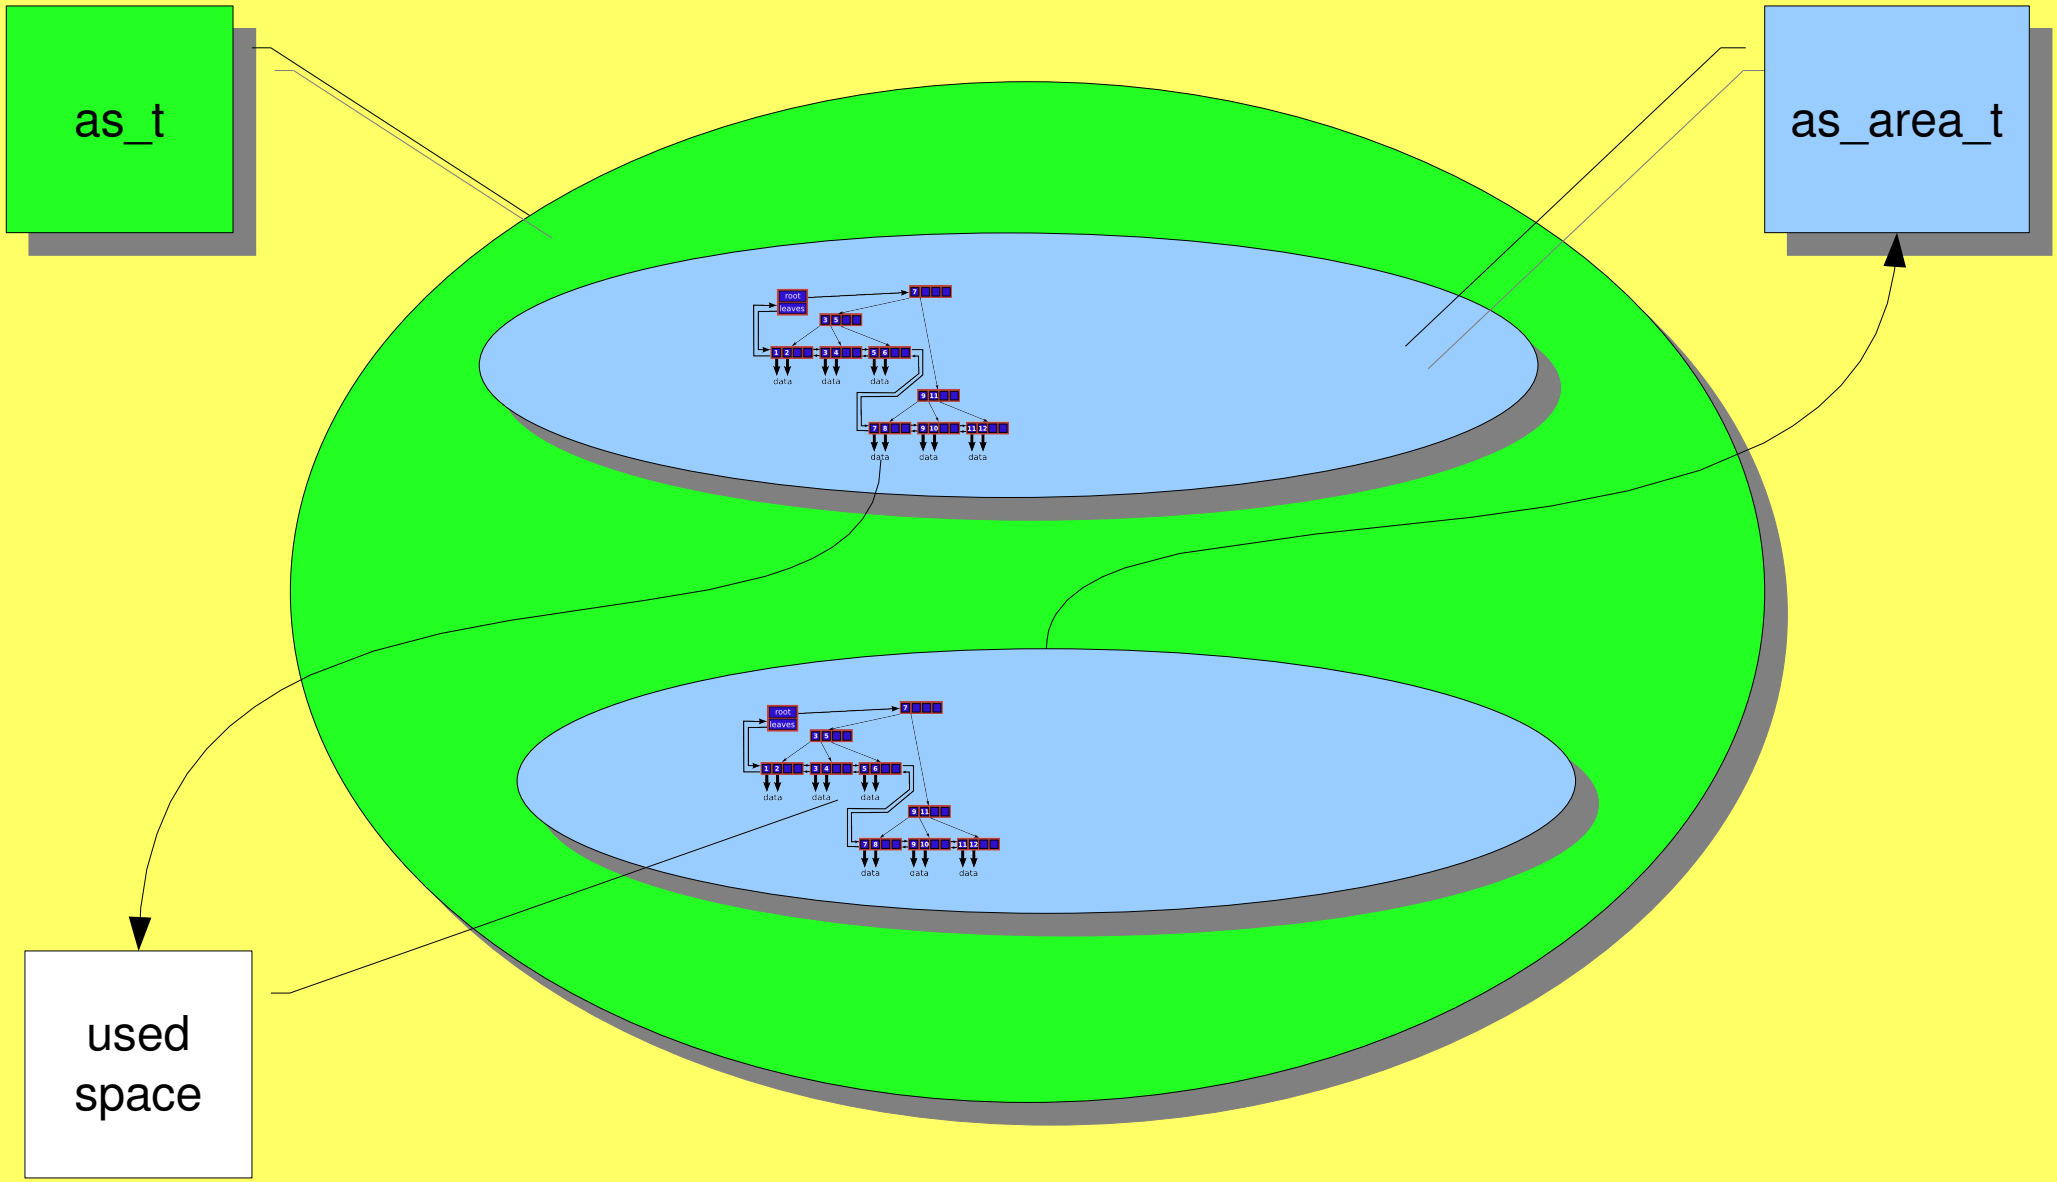

Memory in HelenOS

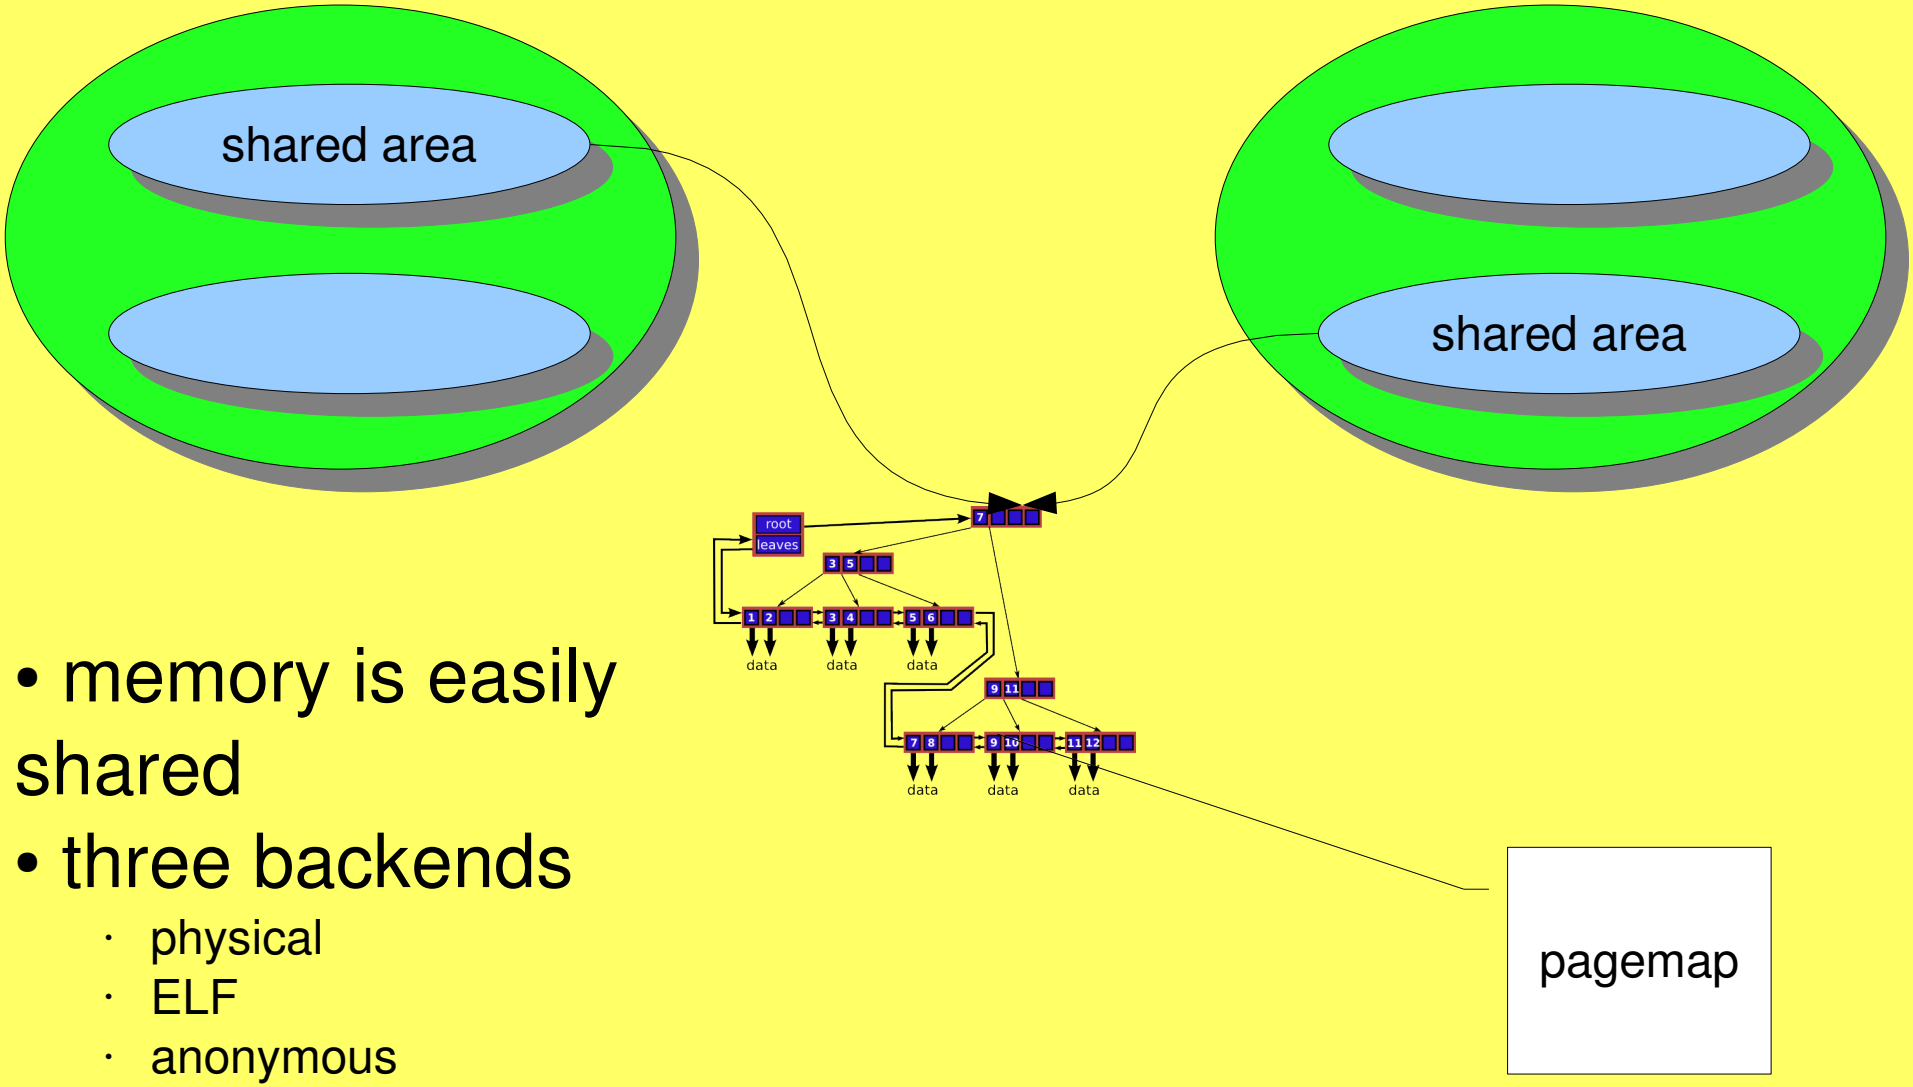

- memory is easily shared
- three backends
  - physical
  - ELF
  - anonymous

Memory in HelenOS

- two mechanisms for mapping pages to frames
- one API

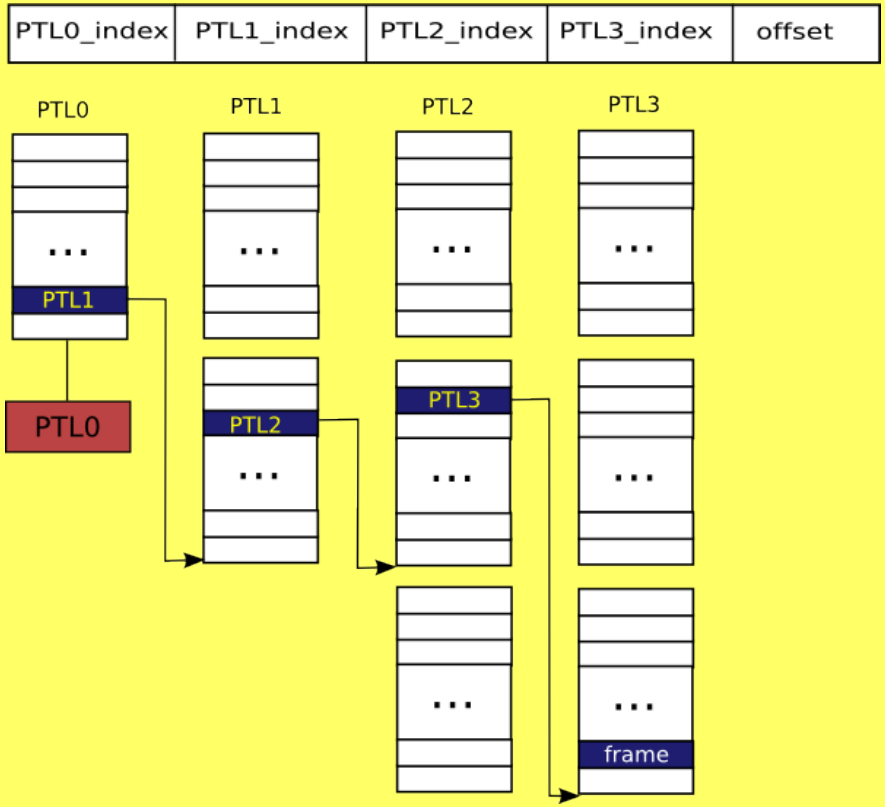

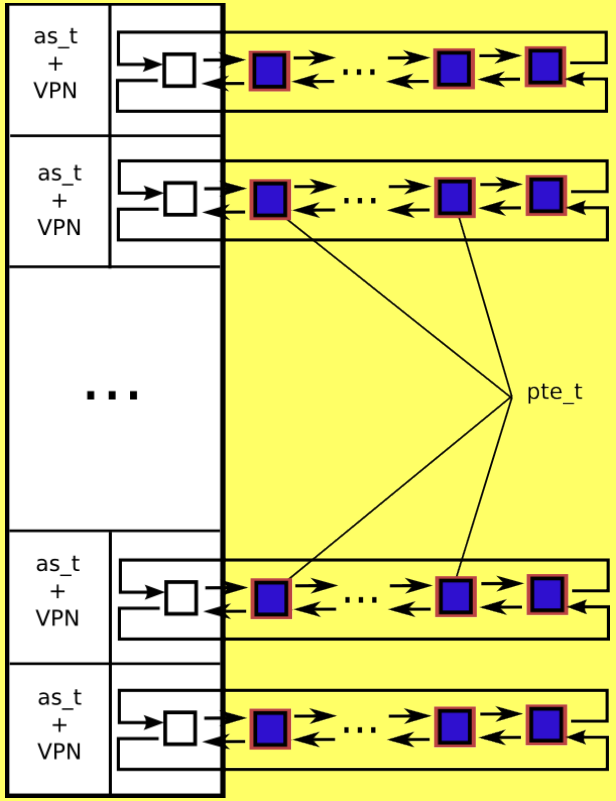

HelenOS IPC

- message passing
- in HelenOS => making calls via phones to answerboxes
- each task has one answerbox
- calls are short: 4 registers (surprisingly not a problem)
- larger data (e.g. images) sent via shared memory
- no copying of messages

HelenOS IPC

- synchronous and asynchronous calls
- problems with connection tracking in multithreaded env.
- worker pseudothreads and manager pseudothreads

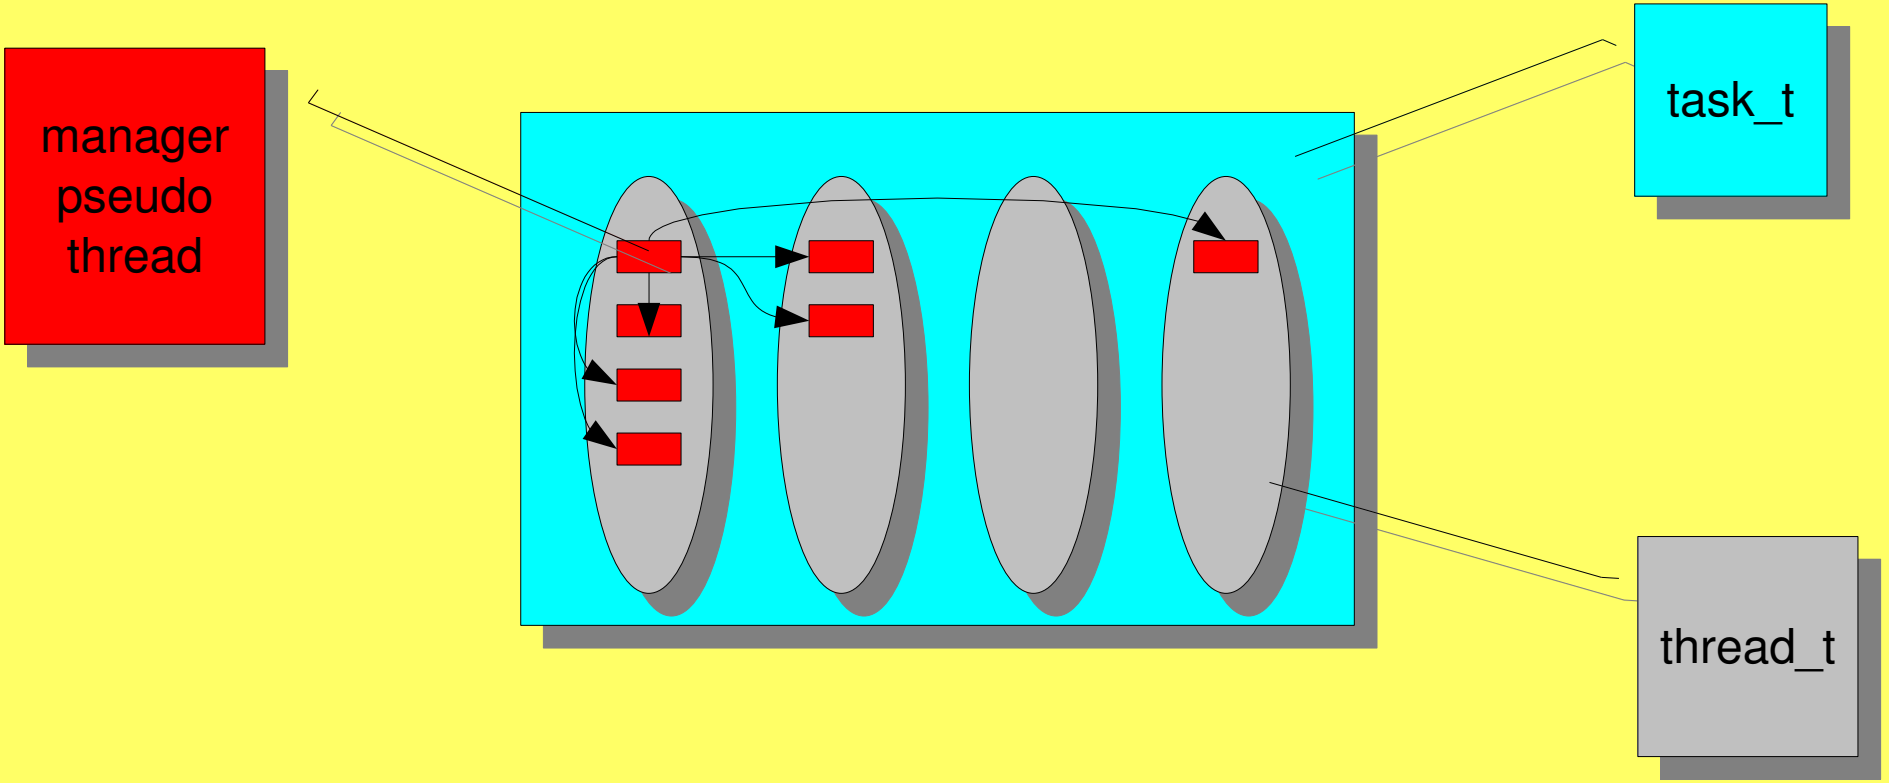

HelenOS userspace

- still pretty rudimentary
- very few POSIX functions
- userspace threads + pseudothreads
- synchronization (futexes)
- IPC + async framework
- ~29 syscalls (11 IPC, 3 vm, 2 synch, 2 sec, 3 ddi, 4 proc)
- ~9 uspace tasks
  - ns
  - console, kbd, fb
  - tetris, klog, ipcc
  - init, pci

Future of HelenOS

- one master thesis before completion (sparc64)
- one master thesis starting (native port to ia64)
- suggested bachelor thesis (libc)
- suggested bachelor thesis (migration and snapshoting)
- suggestion for HelenOS II project
- four students improving timeouts and implementing RCU
- filesystem layer
- (block) device layer
- networking
- old debts: ppc64 + ia32xen
- widening support for all architectures
- the future is in healthy and active community

HelenOS links

- Project homepage: http://www.helenos.eu
- Repository: svn://svn.helenos.eu/HelenOS/trunk
- Browse repository: http://svn.helenos.eu
- Mailing lists:
  - http://lists.modry.cz/cgi-bin/listinfo/helenos
  - http://lists.modry.cz/cgi-bin/listinfo/helenos-devel
  - http://lists.modry.cz/cgi-bin/listinfo/helenos-commits

HelenOS demo

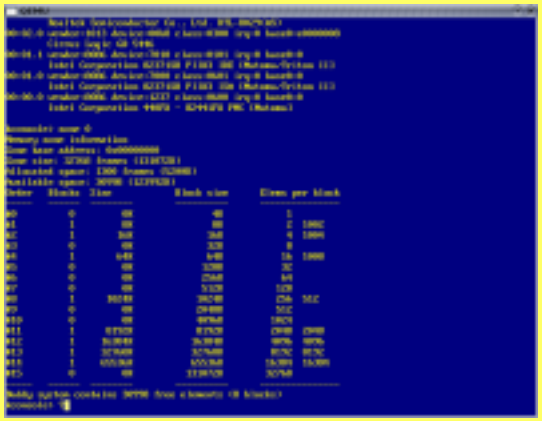

Q&A

Thank you.