HelenOS

http://d3s.mff.cuni.cz

Department of Distributed and Dependable Systems

D3S

Martin Děcký

decky@d3s.mff.cuni.cz

CHARLES UNIVERSITY IN PRAGUE

faculty of mathematics and physics

What is HelenOS?

HelenOS

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

Department of Distributed and Dependable Systems D3S

What is HelenOS?

HelenOS

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

Department of Distributed and Dependable Systems D3S

What is HelenOS?

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

Not a Linux/*BSD/etc. distribution
SPARTAN microkernel
Custom user space
http://www.helenos.org

Department of Distributed and Dependable Systems

What is HelenOS?

HelenOS

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

Department of Distributed and Dependable Systems D3S

What is HelenOS?

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

3-clause BSD license
bzr://bzr.helenos.org/mainline
http://trac.helenos.org
http://www.ohloh.net/p/helenos

Department of Distributed and Dependable Systems

What is HelenOS?

HelenOS

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

Department of Distributed and Dependable Systems D3S

What is HelenOS?

HelenOS

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

Rather breadth-first than depth-first
Desktop
Server
Embedded

Department of Distributed and Dependable Systems D3S

What is HelenOS?

HelenOS

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

Department of Distributed and Dependable Systems

What is HelenOS?

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

IA-32 and AMD64 (PC)
IA-64 (Itanium)
ARM (Neo FreeRunner, BeagleBoard, BeagleBone, *Raspberry Pi*)
MIPS (Malta)
PowerPC (iMac G4)
UltraSPARC (Ultra 60, Enterprise T1000)

What is HelenOS?

HelenOS

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

Department of Distributed and Dependable Systems D3S

What is HelenOS?

HelenOS

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

Mostly familiar "common sense" API
Unicode (no legacy character sets)
No skeletons in the closet (`strcpy`, signals, `fork` & `exec`)
`libposix` as an optional emulation layer

Department of Distributed and Dependable Systems

What is HelenOS?

HelenOS

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

Department of Distributed and Dependable Systems D3S

What is HelenOS?

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

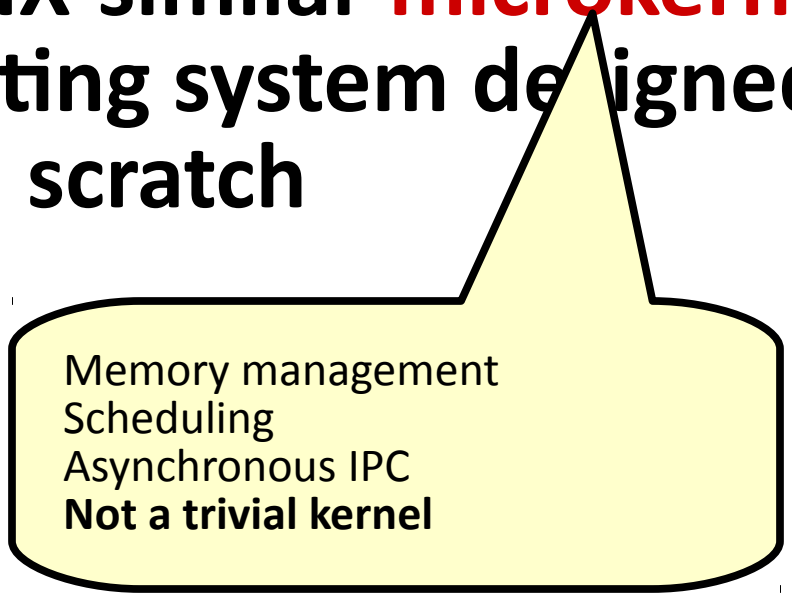

What is HelenOS?

HelenOS

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

Department of Distributed and Dependable Systems D3S

What is HelenOS?

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

Component-based design
Separate and isolated user space tasks (servers)
Naming service, VFS, file system drivers, Location service, device drivers, network layers, graphics stack layers, etc.

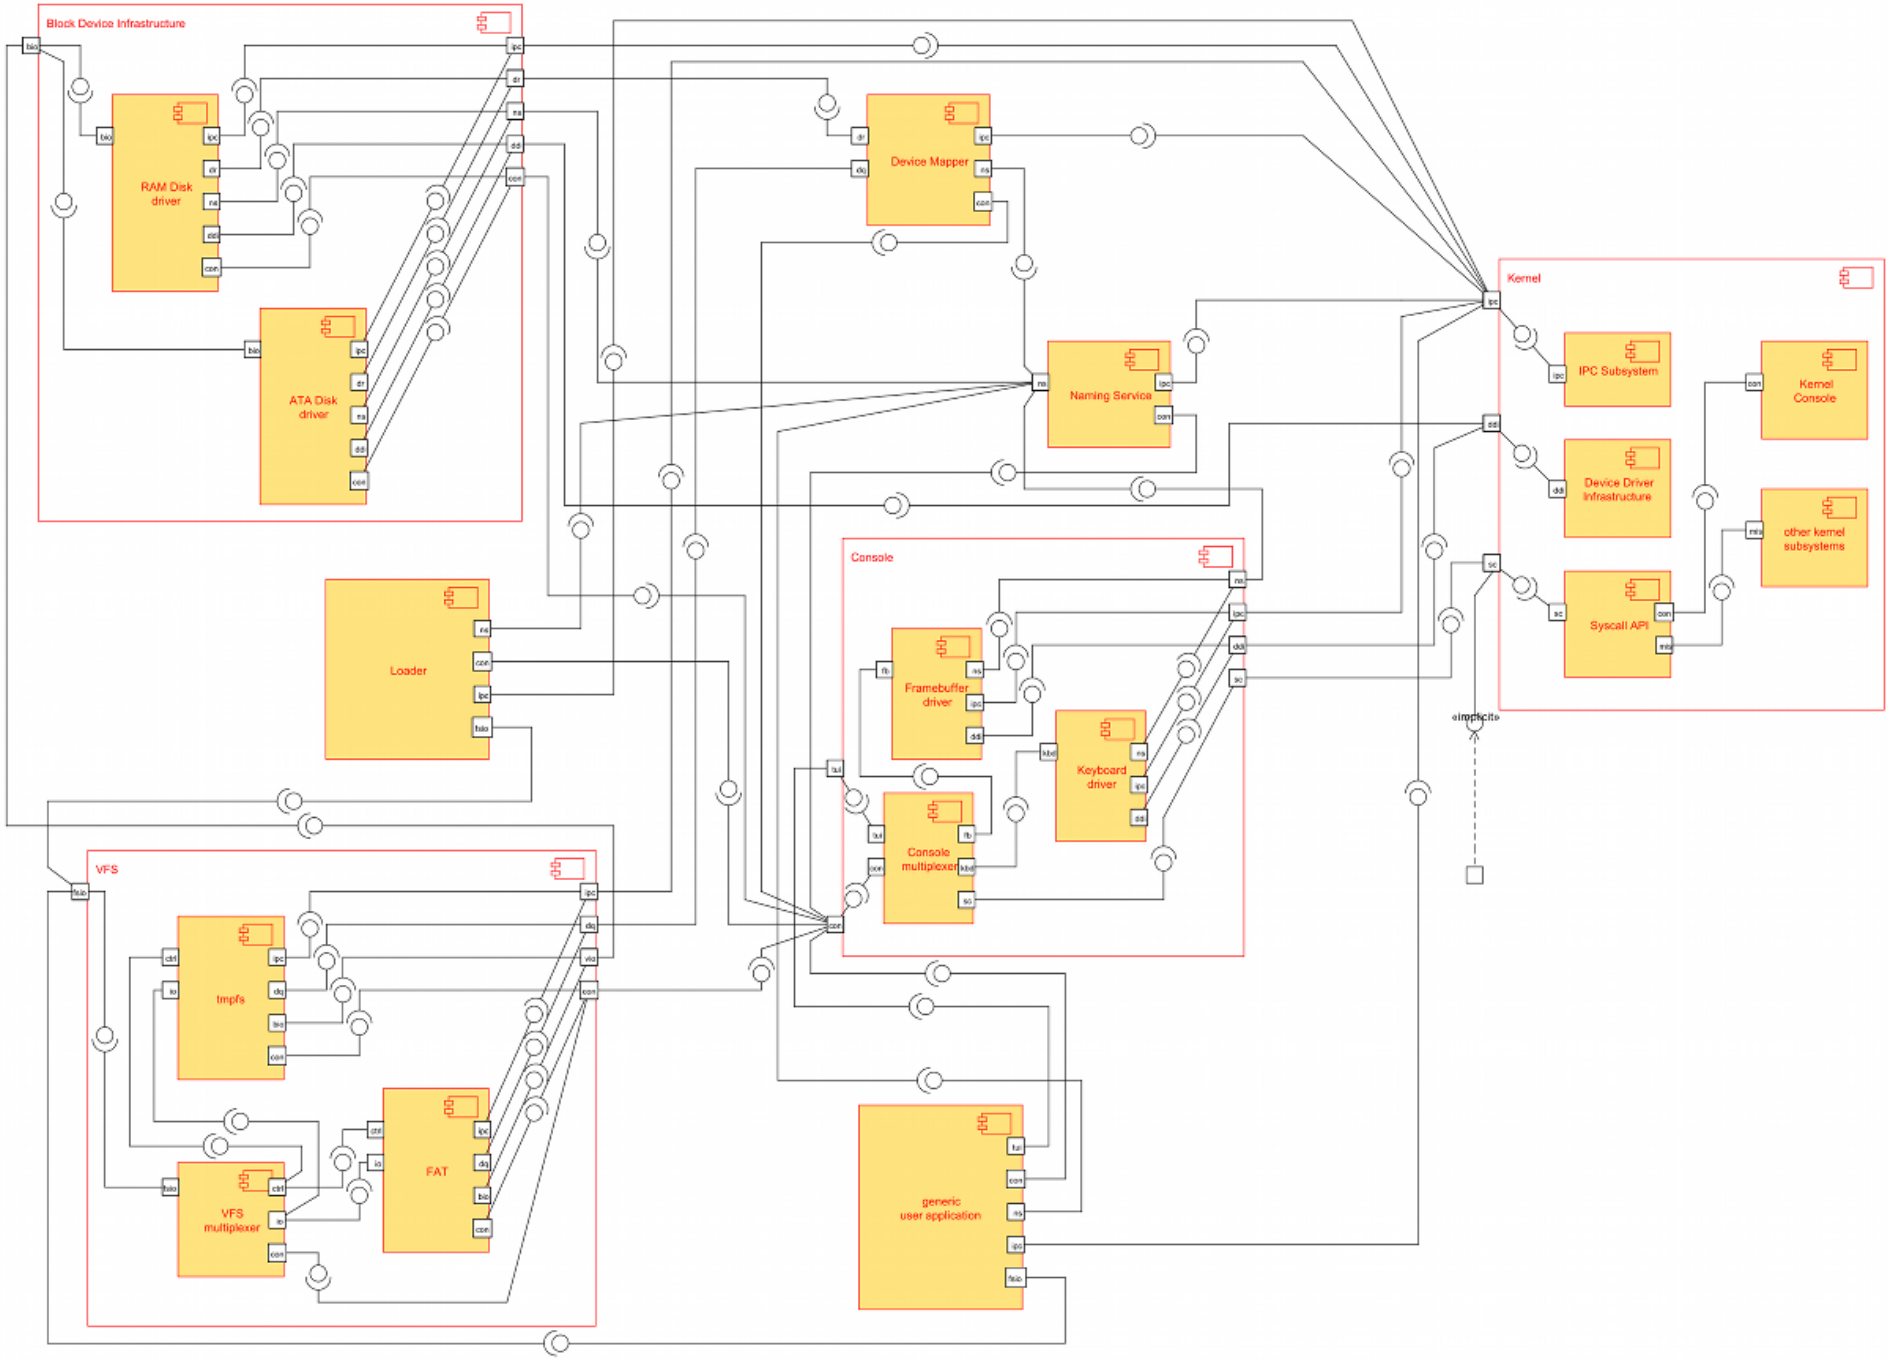
Block Device Infrastructure
RAM Disk driver
ATA Disk driver
Device Mapper
Naming Service
Kernel
IPC Subsystem
Kernel Console
Device Driver Infrastructure
other kernel subsystems
Syscall API
«implicit»
Loader
Console
Framebuffer driver
Keyboard driver
Console multiplexer
VFS
tmpfs
FAT
VFS multiplexer
generic user application

What is HelenOS?

HelenOS

open source general-purpose multiplatform POSIX-similar microkernel multiserver operating system designed from scratch

Department of Distributed and Dependable Systems D3S

Designed from scratch

HelenOS

- **SPARTAN kernel since 2001**
- **HelenOS since 2005**
- **Latest release 0.5.0 (Fajtl) in August 9th 2012**
- **Some 330K lines of code in the mainline branch**
  - About 93 % in C, 3 % in assembler
  - About 30 % of comments
    - "Very well-commented source code" [ohloh.net]
- **Some 45 contributors**
  - About 20 active during the last year (only 10 of them at least 10 commits)

Department of Distributed and Dependable Systems

```
 * @param session Session.
 *
 * @return New exchange or NULL on error.
 *
 */
async_exch_t *async_exchange_begin(async_sess_t *sess)
{
	if (sess == NULL)
		return NULL;

	async_exch_t *exch;

	fibril_mutex_lock(&async_sess_mutex);

	if (!list_empty(&sess->exch_list)) {
		/*
		 * There are inactive exchanges in the session.
		 */
		exch = (async_exch_t *)
		    list_get_instance(list_first(&sess->exch_list),
		    async_exch_t, sess_link);

		list_remove(&exch->sess_link);
		list_remove(&exch->global_link);
	} else {
		/*
		 * There are no available exchanges in the session.
		 */

		if ((sess->mgmt == EXCHANGE_ATOMIC) ||
		    (sess->mgmt == EXCHANGE_SERIALIZE)) {
			exch = (async_exch_t *) malloc(sizeof(async_exch_t));
			if (exch != NULL) {
				link_initialize(&exch->sess_link);
				link_initialize(&exch->global_link);
				exch->sess = sess;
				exch->phone = sess->phone;
			}
```

Commits

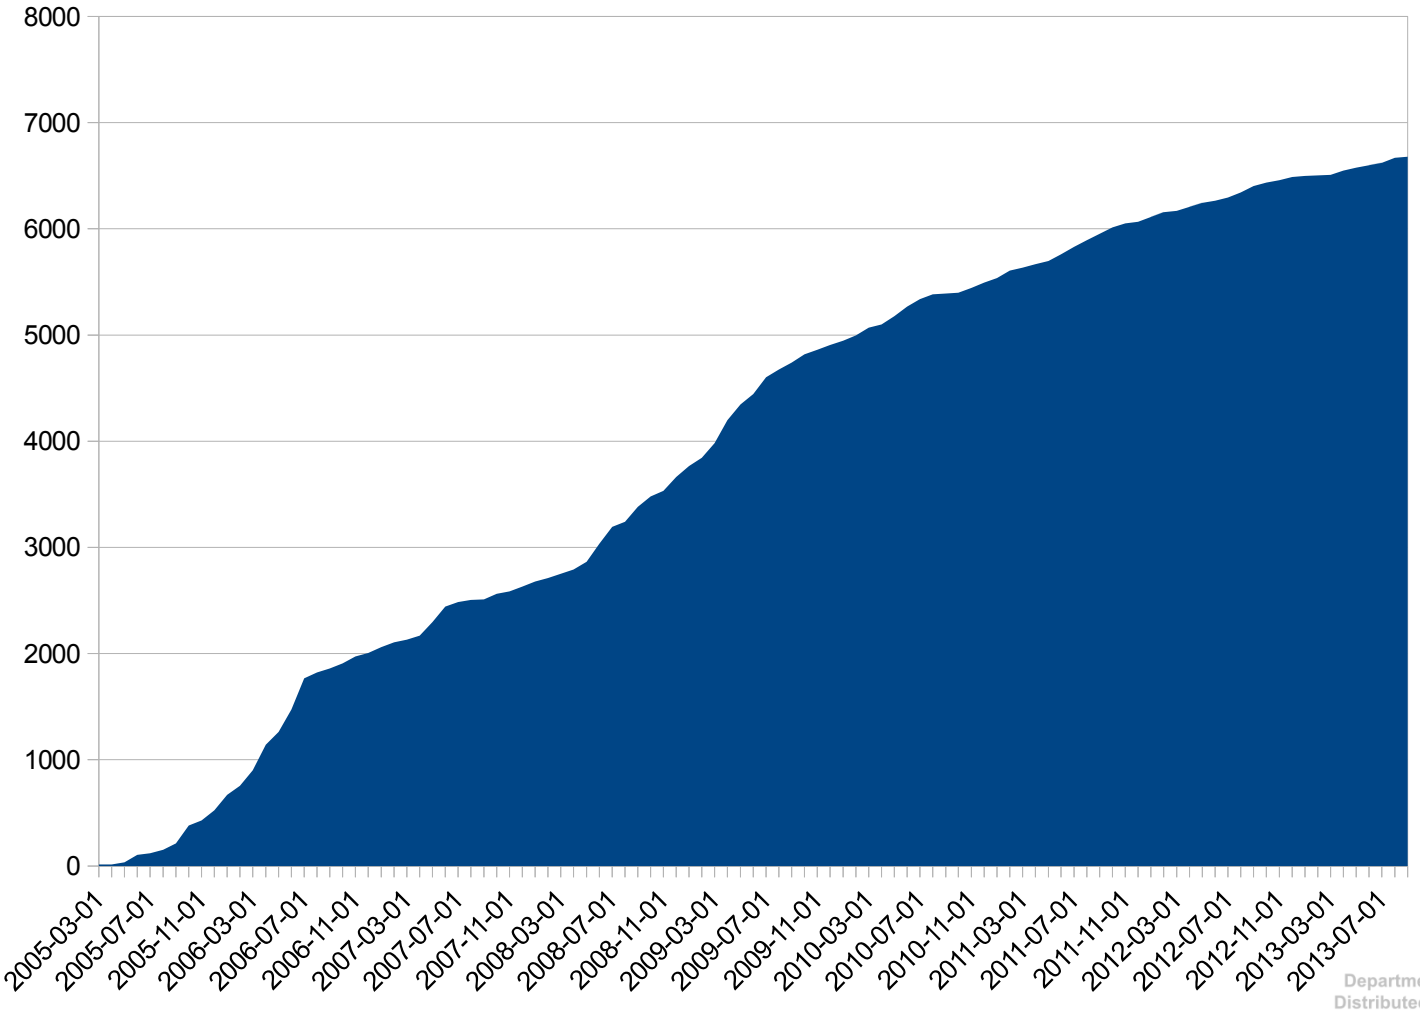
8000
7000
6000
5000
4000
3000
2000
1000
0
2005-03-01
2005-07-01
2005-11-01
2006-03-01
2006-07-01
2006-11-01
2007-03-01
2007-07-01
2007-11-01
2008-03-01
2008-07-01
2008-11-01
2009-03-01
2009-07-01
2009-11-01
2010-03-01
2010-07-01
2010-11-01
2011-03-01
2011-07-01
2011-11-01
2012-03-01
2012-07-01
2012-11-01
2013-03-01
2013-07-01

Lines of code

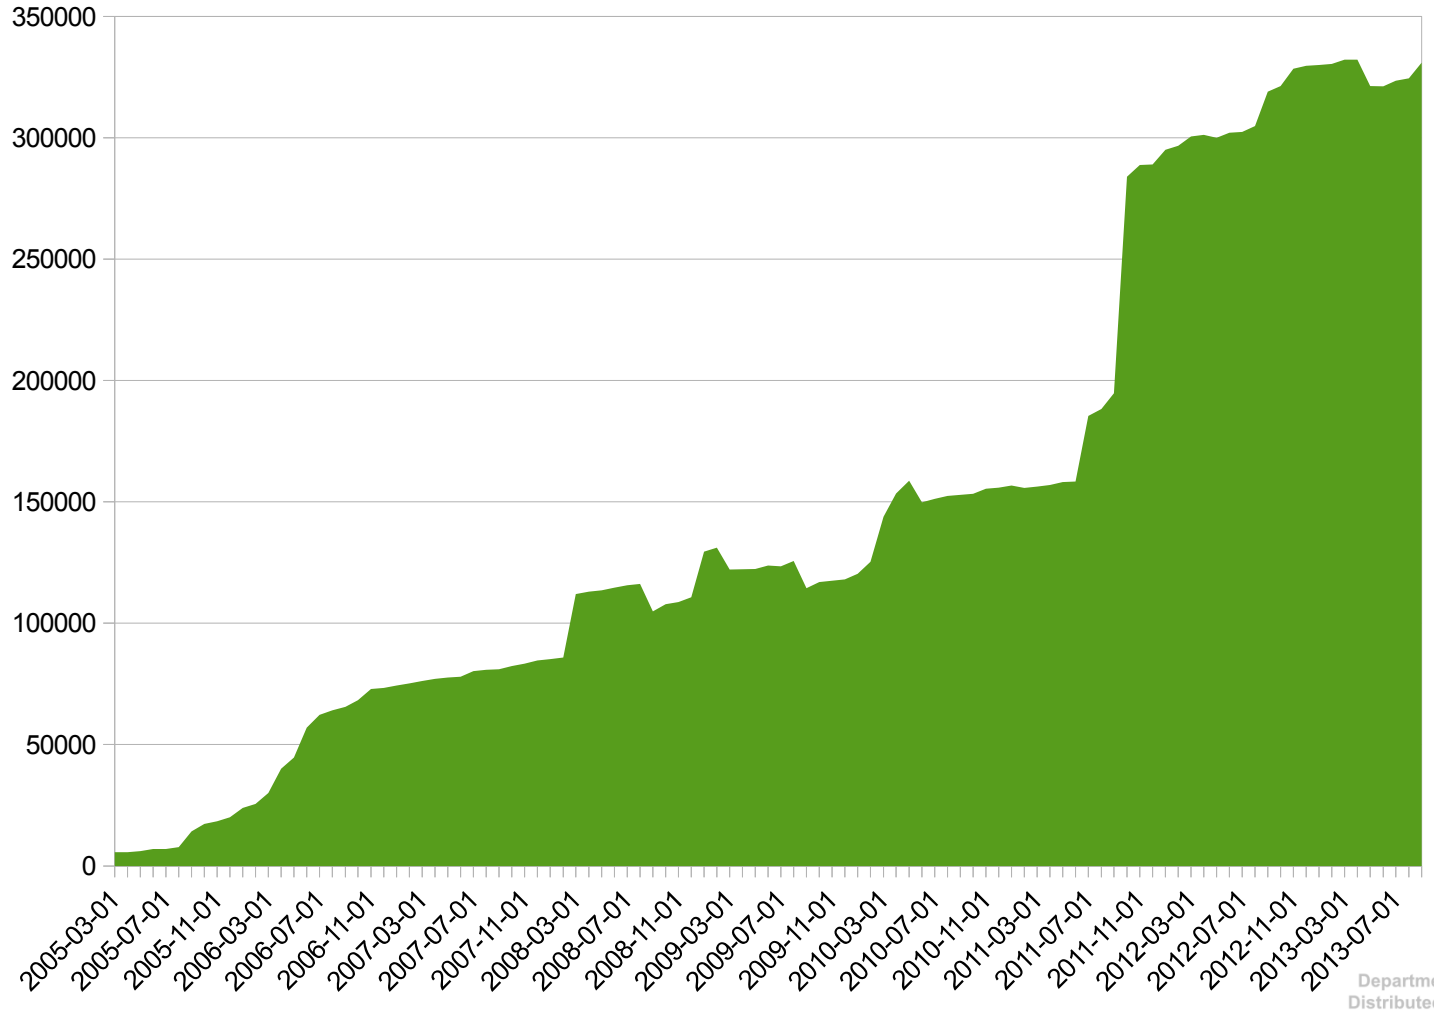

Design principles

HelenOS

- Microkernel design principle
- General-purpose design principle
- Non-fundamentalistic design principle
- Full-fledged design principle
- Multiserver design principle
- Split of mechanism and policy design principle
- Encapsulation design principle
- Portability design principle

Department of Distributed and Dependable Systems D3S

Why?

HelenOS vs. Linux

~~HelenOS vs. Linux~~

HelenOS and Linux


Windows
A fatal exception 0E has occurred at 0028:C562F1B7 in VXD ctpci9x(05)
+ 00001853. The current application will be terminated.
* Press any key to terminate the current application.
* Press CTRL+ALT+DEL again to restart your computer. You will
lose any unsaved information in all applications.
Press any key to continue

Why?

HelenOS

Reliability
Robustness
Dependability
Natural design

Department of Distributed and Dependable Systems D3S

Dependability

- **IEEE definition**
  - *"Dependability is a measurable and provable degree of system's availability, reliability and its maintenance support"*
- **In other words**
  - Formal verification of correctness and quality of service with respect to predefined specification/criteria

Dependability (2)

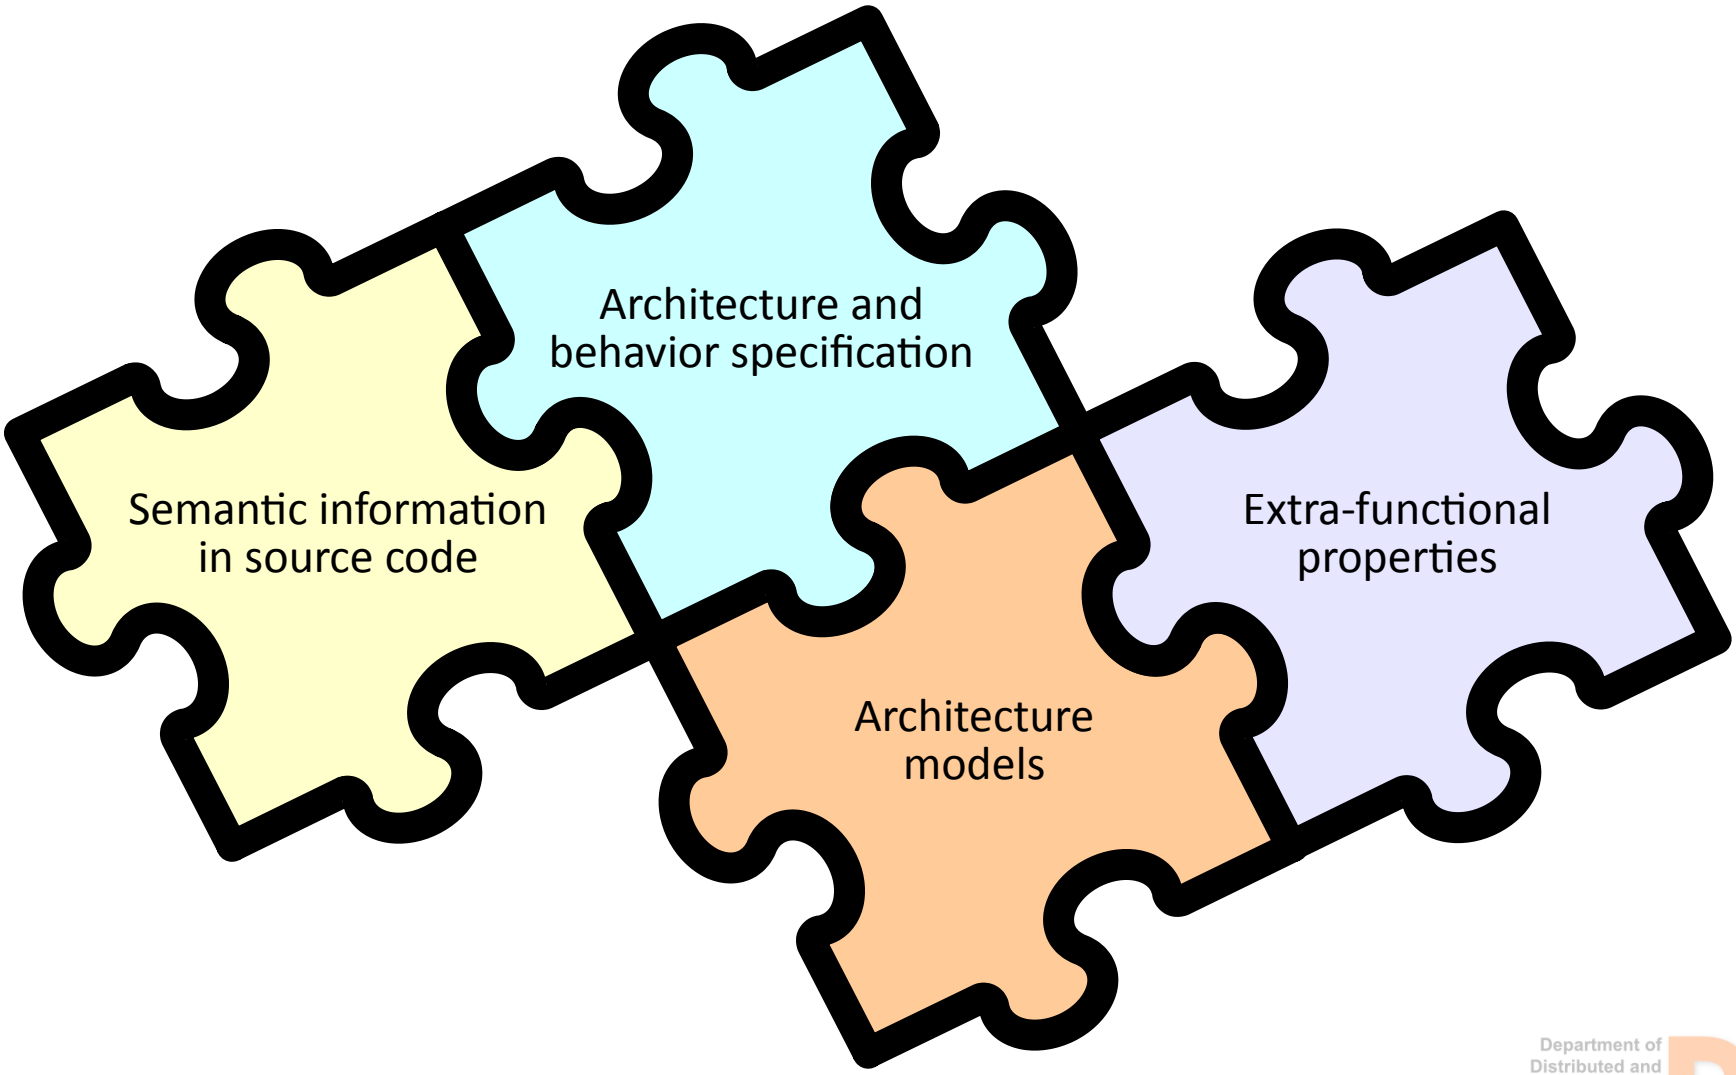

Dependability (2)

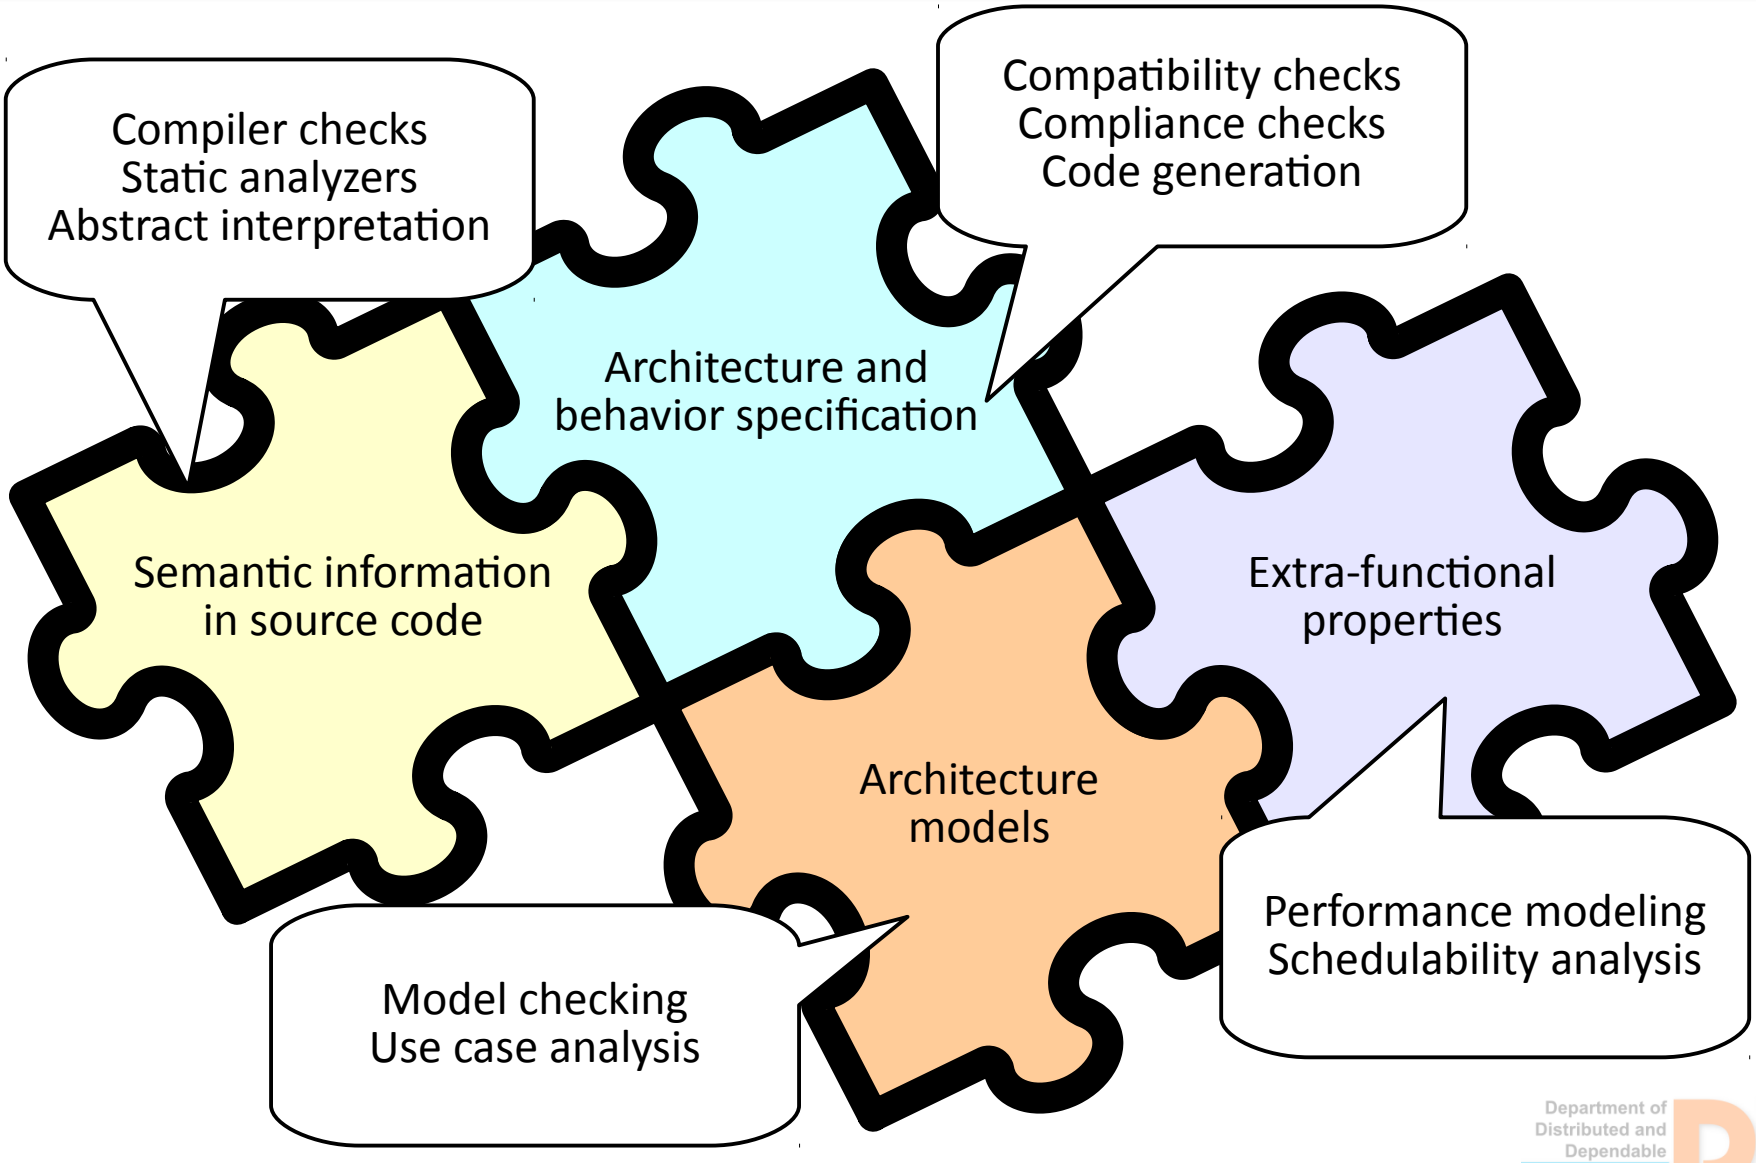

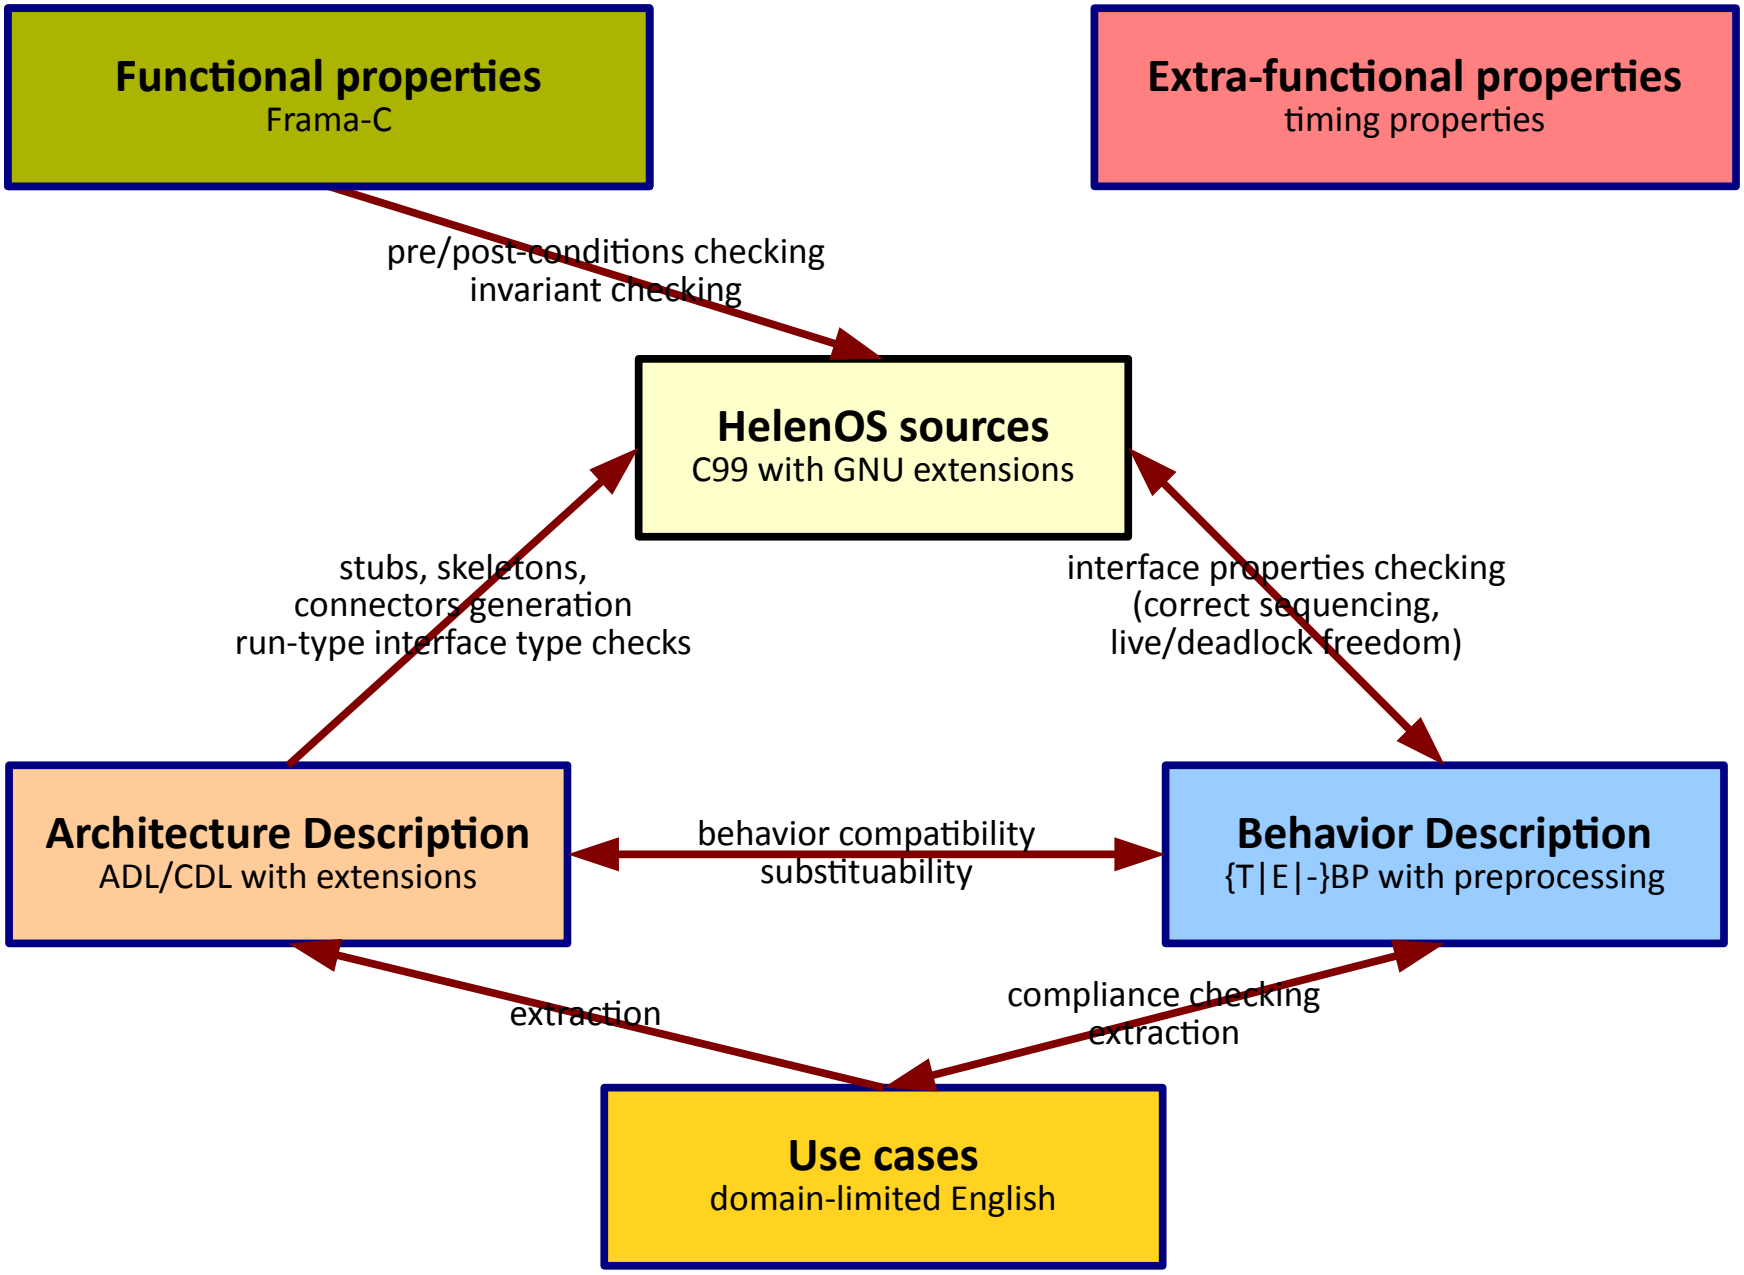
Functional properties
Frama-C
Extra-functional properties
timing properties
pre/post-conditions checking
invariant checking
HelenOS sources
C99 with GNU extensions
stubs, skeletons,
connectors generation
run-type interface type checks
interface properties checking
(correct sequencing,
live/deadlock freedom)
Architecture Description
ADL/CDL with extensions
behavior compatibility
substituability
Behavior Description
{T|E|-}BP with preprocessing
extraction
compliance checking
extraction
Use cases
domain-limited English

Microkernel OS community

- **FOSDEM 2012**
  - February 4th – 5th 2012, Brussels, Belgium
  - Université libre de Bruxelles
  - Jakub Jermář chaired the *Microkernel OS Devroom*
    - Participation of HelenOS, Genode Labs, NOVA, MINIX, Hurd
- **FOSDEM 2013**
  - February 2nd – 3rd 2013, Brussels, Belgium
  - Genode Labs chaired the *Microkernel and Component-based OS Devroom*

FOSDEM 2012

HelenOS

Home › Schedule › Tracks › Microkernel OS Devroom

Microkernel OS Devroom

Room: K.3.201

ical | xcal

Sunday 2012-02-05

Event	Speaker	Room	When
Welcome! or Why do we meet here today	Jakub Jermář	K.3.201	09:00-09:25
Introduction to the NOVA kernel API	Julian Stecklina	K.3.201	09:30-10:00
Introduction of the Genode OS Framework	Norman Feske	K.3.201	10:10-10:55
Introduction to HelenOS	Jakub Jermář	K.3.201	11:10-11:55
The microkernel overhead	Martin Děcký	K.3.201	13:00-14:00
The agony of choice - the diversity of microkernels in Genode	Stefan Kalkowski	K.3.201	14:10-14:55
Dive into HelenOS Device Drivers	Jiří Svoboda	K.3.201	15:05-15:50
Panel discussion	Julian Stecklina , Jakub Jermář , Ben Gras , Christian Helmuth	K.3.201	16:00-17:00

Department of Distributed and Dependable Systems D3S

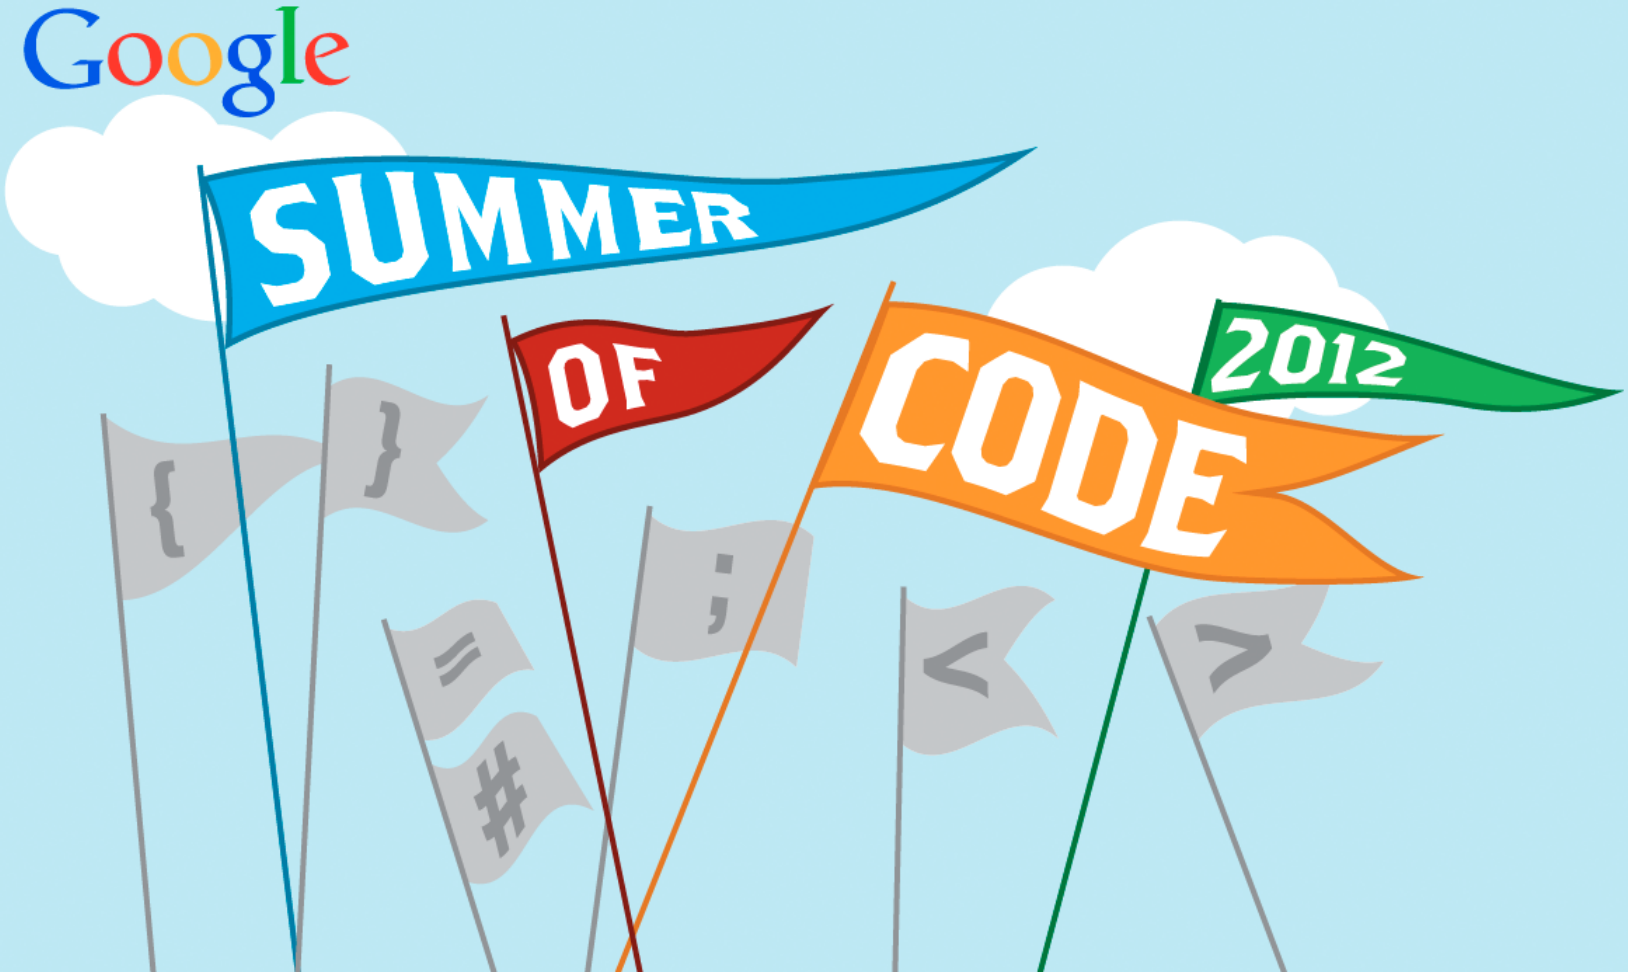
Google
SUMMER
OF
CODE
2012
{
}
;
=
#
<
>

Learning by doing

- Keep the Moore's law in mind early during the design phase
- Do not put the kernel in charge of purely user space namespaces
- Too much synchronization spoils the kernel

Features

- **File systems**
  - ext4, FAT, exFAT, ISO 9660, UDF, MFS
- **Networking**
  - IPv4 & IPv6, NE2000, E1000, RTL8139
- **GUI**
  - Composing desktop
- **USB 1.1 (UHCI, OHCI)**
  - HID, mass storage
- **PATA, SATA (AHCI)**
- **GNU binutils, Portable C Compiler (PCC), MIPS simulator**

What next?

What next?

- **Towards self-hosting**
  - GCC, Clang
- **Go**
- **VFS2**
- **FUSE**
- **BIRD, KnotDNS**

Department of Distributed and Dependable Systems D3S

What next? (2)

- **User space driven system-wide scheduler**
- **User space driven SMP management**
- **New RCU algorithms**
  - AP-RCU (highly portable, decently scalable PaR)
  - AH-RCU (highly scalable, microkernel-friendly)
- **Implicitly shared resources management**
  - De-duplicated caching, future usage prediction (read-ahead), resource pressure evaluation (out-of-memory conditions)

Join us!

HelenOS

www.helenos.org

Department of Distributed and Dependable Systems D3S