



Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Outline

# Let's port together. Debian fun for everyone.

"Most civilised people are out of touch with reality because they confuse the world as it is with the world as they think about it, talk about it and describe it."

Peter De Schrijver  
p2@debian.org

Steve Langasek  
vorlon@debian.org

16th May 2006



# Overview I

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Outline

## Portability issues

why ?

C types

Bitfields

Endianness

Alignment

## Accessing peripheral hardware

Example system architectures

Trends in system design

Out of order transactions

Non-coherent I/O

Userland hardware access

## Fixing porting bugs

Types of build failures

Identifying build failures



# Overview II

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Outline

Failures from unsatisfiable build-deps

Upstream portability problems: pointers and ints

Debian bugs: sudo support

Fix the bug



# Why ?

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

- ▶ Correctness
- ▶ Debian is "The Universal Operating System"
- ▶ Debian is the most used Embedded Distribution
- ▶ Hardware advances will make Debian feasible on new platforms
- ▶ It's enlightening to see and play with other architectures/systems



# C types

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

## ▶ ANSI-C

- ▶  $\text{sizeof(char)} \leq \text{sizeof(short)} \leq \text{sizeof(int)} \leq \text{sizeof(long)}$
- ▶ short and int are at least 16bit
- ▶ long is at least 32bit
- ▶  $\text{sizeof(ptr)} \neq \text{sizeof(int)}$
- ▶ signedness of chars is arch dependent

## ▶ Tips

- ▶ use int as much as possible for computations, loop variables,...
- ▶ use ISO C99 types (`u_int8`, `u_int16`, `u_int32`, ...) for external comms
- ▶ don't abuse chars to 'save memory'
- ▶ use the latest gcc version with `-Wall`



# Bitfields

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

typedef struct bitfields

```
{  
    unsigned char bitfield0:3;  
    unsigned char bitfield1:5;  
}
```

IA32 representation :



PowerPC representation :





# Endianness

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

```
Consider 0x12345678
```

```
Little endian : 0x78 0x56 0x34 0x12
```

```
Big endian : 0x12 0x34 0x56 0x78
```

```
PDP endian : 0x34 0x12 0x78 0x56
```

- ▶ External interfaces
- ▶ Use macros to convert between CPU and specific endianness



# Alignment

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

- ▶ Most RISC cpus require aligned accesses
- ▶ Unaligned accesses are trapped (mostly)
  - ▶ slow
  - ▶ not possible in kernel land
- ▶ Unaligned accesses are seldomly atomic with respect to SMP/other bus masters
- ▶ Better
  - ▶ avoid them
  - ▶ have the compiler generate the code



# Intel style system

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

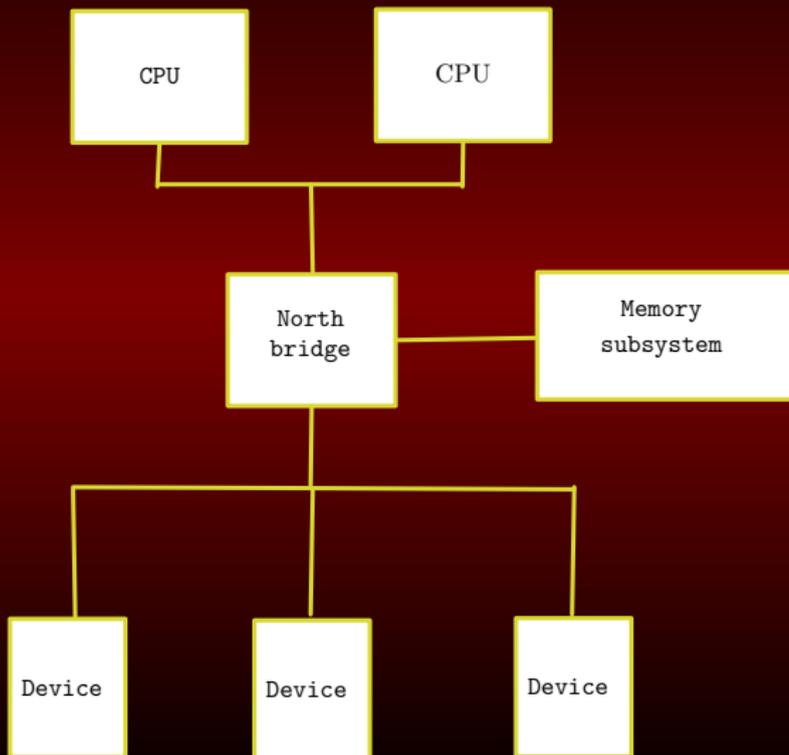
Why ?  
C types  
Bitfields  
Endianness  
Alignment

Peripherals

Architectures  
Trends  
Out of order  
I/O  
Userland

Fixit

Failures  
Identify  
BuildDeps  
PointerPorting  
SudoHappy  
FixitFixit





# Intel style system

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

- ▶ Main components
  - ▶ CPU complex
  - ▶ Northbridge
  - ▶ Southbridge
  - ▶ Memory subsystem
- ▶ Main interfaces
  - ▶ Frontside bus
  - ▶ PCI
  - ▶ AGP

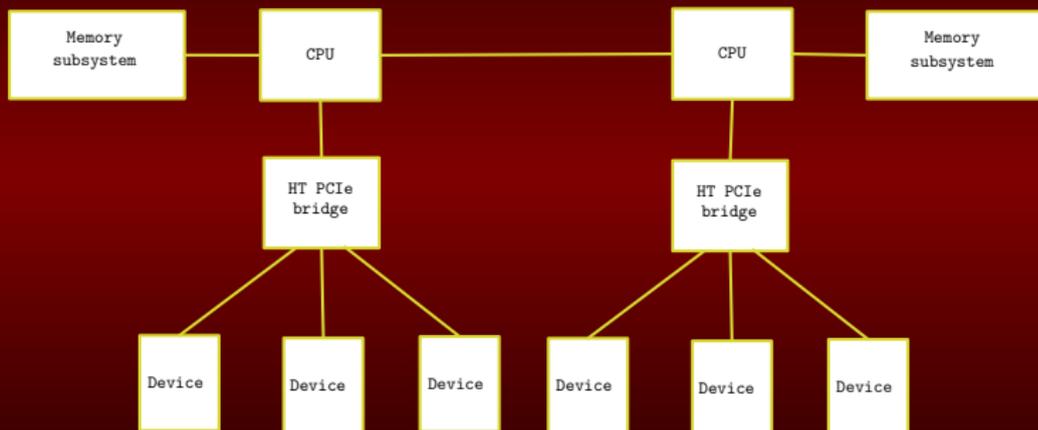


# Opteron style system

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek



Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit



# Opteron style system

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

- ▶ Main components
  - ▶ CPU
  - ▶ Hypertransport - PCIe bridge
  - ▶ PCIe - PCI bridge
- ▶ Main interfaces
  - ▶ Hypertransport
  - ▶ PCIe
  - ▶ PCI



# Trends in system design

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

## ▶ Observations

- ▶ CPUs became much faster then memory
- ▶ bus and memory bandwidth have gone up faster then latencies
- ▶ parallel busses become very hard at high speeds

## ▶ Solutions

- ▶ Caches
- ▶ Burstmode transfers
- ▶ Advanced DMA
- ▶ multiple highspeed serial links



# Out of order transactions

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

- ▶ examples
  - ▶ CPU
  - ▶ bus bridges
- ▶ use read/write barriers
  - ▶ CPU instructions
  - ▶ "magic" reads

possibly out of order :

```
stw r20,0x20(r21)
```

```
stw r22,0x24(r21)
```

always in order :

```
stw r20,0x20(r21)
```

```
eieio
```

```
stw r22,0x24(r21)
```



# Non-coherent I/O

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

- ▶ Some systems do not support "bus snooping"
- ▶ Invalidate cachelines
  - ▶ network traffic
  - ▶ disk buffers
  - ▶ other kinds of streaming I/O
- ▶ non-cacheable memory
  - ▶ microcode
  - ▶ ring buffers



# Addressing

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?  
C types  
Bitfields  
Endianness  
Alignment

Peripherals

Architectures  
Trends  
Out of order  
I/O  
Userland

Fixit

Failures  
Identify  
BuildDeps  
PointerPorting  
SudoHappy  
FixitFixit

- ▶ Virtual addresses
- ▶ Physical addresses
- ▶ Bus addresses
- ▶ Translation Physical to Bus addresses
  - ▶ identity mapped
  - ▶ fixed offset
  - ▶ page based translation
  - ▶ not memory mapped
    - ▶ IA32 I/O ports
    - ▶ PowerPC DCB
- ▶ Always access hardware via special functions



# Transaction atomicity

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

- ▶ Multiple CPUs
- ▶ other busmasters (eg. on PCI)
- ▶ reads and writes are atomic only if aligned
- ▶ atomic read/modify/write is CPU specific
  - ▶ ia32: lock prefix on specific instructions
  - ▶ mips: ll/sc
  - ▶ arm: swap
  - ▶ ppc: lwarx/stwax
- ▶ bridges may break locks



# Userland hardware access

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

- ▶ Hardware access from userland is problematic
- ▶ Seperate command transport from driver logic
  - ▶ Firewire : libraw1394
  - ▶ USB : libusb
  - ▶ SCSI and ATAPI : scsi generic like ioctl
  - ▶ ...
- ▶ Provide abstraction layer for accessing hardware



# Major causes of architecture-specific build failures

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

**Failures**

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

- ▶ software bugs in your package (see previous slides)
- ▶ architecture-specific build-dependency problems
- ▶ temporarily unavailable build-depends
- ▶ missing versioned build-depends + bad timing
- ▶ toolchain-specific bugs: try a new toolchain?
- ▶ build-env-specific bugs (sudo support)
- ▶ buildd-specific bugs



# Step 1: notice the package hasn't built

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

BuildDeps

PointerPorting

SudoHappy

FixitFixit

- ▶ many different resources today to help you track this - use one!
  - ▶ by maintainer:  
<http://people.debian.org/~igloo/status.php?email=<email>>
  - ▶ by package:  
<http://buildd.debian.org/<package>>
  - ▶ by arch:  
<http://buildd.debian.org/~jeroen/status/architecture.php?a=<arch>>
  - ▶ arbitrary list of packages:  
<http://buildd.debian.org/~vorlon/packagesumm.php>
- ▶ drilling down gets you a build log:  
<http://buildd.debian.org/fetch.php?pkg=db4.3&arch=hppa&ver=4.3.29-5&stamp=1147065979&file=log>



## Step 2: read the build log

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?

C types

Bitfields

Endianness

Alignment

Peripherals

Architectures

Trends

Out of order

I/O

Userland

Fixit

Failures

Identify

**BuildDeps**

PointerPorting

SudoHappy

FixitFixit

- ▶ <http://buildd.debian.org/fetch.php?pkg=db4.3&arch=hppa&ver=4.3.29-5&stamp=1147065979&file=log>

```
The following packages have unmet dependencies:  
java-gcj-compat-dev: Depends: ecj-bootstrap-gcj  
(>= 3.1.1) but it is not going to be installed
```

- ▶ obviously not a bug in the upstream source code of db4.3
- ▶ *\*is\** a bug that needs to be addressed for the release!
- ▶ fixing it is out-of-scope for this workshop... porters + maintainers need to agree whether to support java bindings on this arch



## Step 2: read the build log - pointers and ints

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?  
C types  
Bitfields  
Endianness  
Alignment

Peripherals

Architectures  
Trends  
Out of order  
I/O  
Userland

Fixit

Failures  
Identify  
BuildDeps  
PointerPorting  
SudoHappy  
FixitFixit

- ▶ New gccs are nice - they give build failures where before packages would be silently broken!
- ▶ <http://buildd.debian.org/fetch.php?pkg=swt-gtk&arch=alpha&ver=3.1.2-2&stamp=1142711970&file=log>

```
xpcom.cpp: In function 'jint
Java_org_eclipse_swt_internal_mozilla\
_XPCOM_PR_1Malloc(JNIEnv*, _jclass*, jint)':
xpcom.cpp:144: error: cast from 'void*' to
'jint' loses precision
```

- ▶ 64-bit pointers != 32-bit ints
- ▶ straightforward fix in concept, swt-gtk is a special-case due to Java ABIs :(



## Step 2: read the build log

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

Portability

Why ?  
C types  
Bitfields  
Endianness  
Alignment

Peripherals

Architectures  
Trends  
Out of order  
I/O  
Userland

Fixit

Failures  
Identify  
BuildDeps  
PointerPorting  
**SudoHappy**  
FixitFixit

- ▶ some packages fail consistently across alpha, mips, mipsel:  
`dpkg-buildpackage -rsudo`
- ▶ `http://buildd.debian.org/fetch.php?pkg=freetype&arch=alpha&ver=2.2.1-1&stamp=1147731281&file=log`

```
rm -f /debian/tmp/usr/include/freetype2/freetype/  
cache/*
```



## Step 3: start hacking

Let's port  
together.  
Debian fun for  
everyone.

Peter De  
Schrijver

Steve  
Langasek

### Portability

Why ?  
C types  
Bitfields  
Endianness  
Alignment

### Peripherals

Architectures  
Trends  
Out of order  
I/O  
Userland

### Fixit

Failures  
Identify  
BuildDeps  
PointerPorting  
SudoHappy  
FixitFixit

- ▶ No magic formula for fixing all build failures, you have to learn about different kinds of build failures one by one
- ▶ so let's start hacking!