

Bus-Independent Device Accesses

Matthew Wilcox

`matthew@wil.cx`

Alan Cox

`alan@redhat.com`

Bus-Independent Device Accesses

by Matthew Wilcox

by Alan Cox

Copyright © 2001 by Matthew Wilcox

This documentation is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

For more details see the file COPYING in the source distribution of Linux.

Table of Contents

1. Introduction.....	1
2. Known Bugs And Assumptions	1
3. Memory Mapped IO.....	2
3.1. Getting Access to the Device	2
3.2. Accessing the device.....	2
3.3. ISA legacy functions.....	3
4. Port Space Accesses	4
4.1. Port Space Explained	4
4.2. Accessing Port Space.....	4
5. Public Functions Provided	5

Chapter 1. Introduction

Linux provides an API which abstracts performing IO across all busses and devices, allowing device drivers to be written independently of bus type.

Chapter 2. Known Bugs And Assumptions

None.

Chapter 3. Memory Mapped IO

3.1. Getting Access to the Device

The most widely supported form of IO is memory mapped IO. That is, a part of the CPU's address space is interpreted not as accesses to memory, but as accesses to a device. Some architectures define devices to be at a fixed address, but most have some method of discovering devices. The PCI bus walk is a good example of such a scheme. This document does not cover how to receive such an address, but assumes you are starting with one. Physical addresses are of type unsigned long.

This address should not be used directly. Instead, to get an address suitable for passing to the accessor functions described below, you should call `ioremap`. An address suitable for accessing the device will be returned to you.

After you've finished using the device (say, in your module's exit routine), call `iounmap` in order to return the address space to the kernel. Most architectures allocate new address space each time you call `ioremap`, and they can run out unless you call `iounmap`.

3.2. Accessing the device

The part of the interface most used by drivers is reading and writing memory-mapped registers on the device. Linux provides interfaces to read and write 8-bit, 16-bit, 32-bit and 64-bit quantities. Due to a historical accident, these are named byte, word, long and quad accesses. Both read and write accesses are supported; there is no prefetch support at this time.

The functions are named `readb`, `readw`, `readl`, `readq`, `writew`, `writel` and `writeq`.

Some devices (such as framebuffer) would like to use larger transfers than 8 bytes at a

time. For these devices, the `memcpy_toio`, `memcpy_fromio` and `memset_io` functions are provided. Do not use `memset` or `memcpy` on IO addresses; they are not guaranteed to copy data in order.

The read and write functions are defined to be ordered. That is the compiler is not permitted to reorder the I/O sequence. When the ordering can be compiler optimised, you can use `__readb` and friends to indicate the relaxed ordering. Use this with care. The `rmb` provides a read memory barrier. The `wmb` provides a write memory barrier.

While the basic functions are defined to be synchronous with respect to each other and ordered with respect to each other the busses the devices sit on may themselves have asynchronicity. In particular many authors are burned by the fact that PCI bus writes are posted asynchronously. A driver author must issue a read from the same device to ensure that writes have occurred in the specific cases the author cares. This kind of property cannot be hidden from driver writers in the API.

3.3. ISA legacy functions

On older kernels (2.2 and earlier) the ISA bus could be read or written with these functions and without `ioremap` being used. This is no longer true in Linux 2.4. A set of equivalent functions exist for easy legacy driver porting. The functions available are prefixed with 'isa_' and are `isa_readb`, `isa_writeb`, `isa_readw`, `isa_writew`, `isa_readl`, `isa_writel`, `isa_memcpy_fromio` and `isa_memcpy_toio`

These functions should not be used in new drivers, and will eventually be going away.

Chapter 4. Port Space Accesses

4.1. Port Space Explained

Another form of IO commonly supported is Port Space. This is a range of addresses separate to the normal memory address space. Access to these addresses is generally not as fast as accesses to the memory mapped addresses, and it also has a potentially smaller address space.

Unlike memory mapped IO, no preparation is required to access port space.

4.2. Accessing Port Space

Accesses to this space are provided through a set of functions which allow 8-bit, 16-bit and 32-bit accesses; also known as byte, word and long. These functions are `inb`, `inw`, `inl`, `outb`, `outw` and `outl`.

Some variants are provided for these functions. Some devices require that accesses to their ports are slowed down. This functionality is provided by appending a `_p` to the end of the function. There are also equivalents to `memcpy`. The `ins` and `outs` functions copy bytes, words or longs to the given port.

Chapter 5. Public Functions Provided