

Open Firmware

Recommended Practice:

Interrupt Mapping

Version 0.9

Unapproved DRAFT

Published by the Open Firmware Working Group

This document is a voluntary-use recommended practice of the Open Firmware Working Group. The Open Firmware Working Group is an ad hoc committee composed of individuals interested in Open Firmware as defined by IEEE 1275-1994, related standards, and their application to various computer systems.

The Open Firmware Working Group is involved both in IEEE sanctioned standards activities, whose final results are published by IEEE, and in informal recommendations such as this, which are published on the Internet at:

`http://playground.sun.com/1275`

Membership in the Open Firmware Working Group is open to all interested parties. The working group meets at regular intervals at various locations. For more information send email to:

`p1275-wg@risc.sps.mot.com`

Revision History

1. Introduction

This recommended practice defines a mapping mechanism between bus-specific interrupt values, as reported via the "interrupts" property of the bus's child nodes, and a system platform's "native" interrupt facility.

1.1. Purpose

The base Open Firmware specification (IEEE Std 1275-1994) defines a generic bus-specific property ("interrupts") that is used to report *interrupt specifiers* used by that bus's devices. An *interrupt specifier* usually has a component that represents a number that corresponds to an input of an interrupt controller; it may also supply other information about the interrupt (e.g., an indication of whether the interrupt is edge or level sensitive). Each bus binding to Open Firmware must define the format and interpretation of its "interrupts" properties.

Platforms (e.g., CHRP) contain *interrupt controllers* to which interrupts from devices are wired. In some cases, the interpretation of how an *interrupt specifier* relates to an *interrupt controller* is simple. However, many platform architectures allow somewhat arbitrary "wiring" of interrupts. The *interrupt specifiers* (i.e., the "interrupt" property values) do not contain enough information in and of themselves to convey the information that an operating system needs to manage the handling of interrupts.

This recommended practice defines an architecture that provides the additional platform-specific information about interrupts and how they are wired.

1.2. Scope

2. References and Definitions

2.1. References

[1] *IEEE Standard for Boot (Initialization Configuration) Firmware: Core Requirements and Practices*, published by IEEE as IEEE Std 1275-1994.

[2] *PCI Local Bus Specification, Revision 2.1*, published by the PCI Special Interest Group.

[3] *PowerPC™ Microprocessor Common Hardware Reference Platform: A System Architecture*, published by Morgan Kaufmann.

[4] *PowerPC™ Microprocessor Common Hardware Reference Platform (CHRP™) System binding to: IEEE Std 1275-1994*, published by the Open Firmware Working Group.

2.2. Definitions

3. The interrupt tree model

The model presented by this recommended practice is the representation of an *interrupt tree* that describes how interrupts are wired, cascaded, etc. on a platform, and the manner in which a client (e.g., an operating system) can use this information.

1 The interrupt tree is represented by means of new properties, defined by this recommended prac-
2 tice, contained within the Open Firmware device tree for a platform. Note that the interrupt tree is
3 represented (and searched) by properties from leaf (device) nodes "upwards" towards the root of
4 the interrupt tree; this is unlike the device tree that has explicit links from the root "downwards"
5 towards leaf nodes.
6

7 In order to represent arbitrary platform wiring, where interrupts may be "distributed" among mul-
8 tiple interrupt controllers, the term *interrupt tree* is technically incorrect. In such cases, the struc-
9 ture is more correctly called an *interrupt graph*. However, since many platforms have a simpler
10 model, we will generically refer to the structure as a tree.
11

12 Each sub-tree of the interrupt tree represents interrupts within an *interrupt domain* that defines the
13 format and interpretation of "interrupts" properties for devices that are members of the do-
14 main. Since multiple busses may exist in a given platform, each of which has its own interrupt
15 domain, the interrupt tree consists of multiple interrupt domains. The root of the interrupt tree de-
16 fines the platform's interrupt domain.
17

18 The base document defines the value of each "interrupts" property entry to be an *interrupt*
19 *specifier*. In practice, however, this value may have to be interpreted with respect to the device's
20 *unit address*. For example, most PCI devices will have an "interrupts" value of 1, as required
21 by the [2]. However, the wiring of interrupts, at least for plug-in devices, is determined by the par-
22 ticular slot in which the device is plugged. On PCI, the slot is implied by the device's *unit address*
23 (which contains the relevant **device#** information).
24

25 Because of this general coupling of *unit address* and *interrupt specifier*, the term *unit interrupt*
26 *specifier* is used when discussing a value that consists of the pair (*unit address*, *interrupt specifier*).
27 For nodes that represent devices, the number of cells to represent a *unit interrupt specifier* is the
28 sum of the "#address-cells" and "#interrupt-cells" properties; for nodes that do
29 not represent devices, there is no relevant "#address-cells" value, so that the number of cells
30 is solely determined by the "#interrupt-cells" value. The latter case exists due to the na-
31 ture of representing interrupt mapping outside the context of the normal device tree.
32
33

34 Each nexus in an interrupt tree represents where some interpretation and/or transformation of an
35 "interrupts" property value might be done. This interpretation is either direct because the
36 node is the interrupt tree's root, it represents an *interrupt controller*, or it requires a "mapping"
37 of a *unit interrupt specifier* in one interrupt domain into a *unit interrupt specifier* in another domain.
38 In some cases (e.g., PCI-PCI bridge), the domains are essentially the same, but some mapping
39 might be necessary because of wiring.
40

41 The result of mapping an interrupt to the top of the interrupt tree results in a platform-specific val-
42 ue; the platform's binding must define the interpretation of this value (e.g., source number and
43 sense for an Open PIC interrupt controller).
44

45 3.1. "#interrupt-cells" property

46 The assumption of the base document was that the interpretation of "interrupts" values, in-
47 cluding their format, was totally specified by a bus binding. However, this recommended practice
48 is designed to allow the traversal of the interrupt tree without such implicit knowledge. This is af-
49 farded by means of a new property ("#interrupt-cells") that explicitly defines the number
50 of cells required for the representation of a single *interrupt specifier*.
51
52
53
54
55
56
57

3.2. "interrupt-parent" property

Since the interrupt tree may not match the physical bus tree (which is what the Open Firmware device tree represents), a new property ("interrupt-parent") is introduced that can denote the interrupt tree hierarchy from device-tree nodes upwards within the interrupt tree.

If a device node does not specify an explicit interrupt parent (i.e., does not have the "interrupt-parent" property), and it is not an interrupt controller (i.e., it does not have an "interrupt-controller" property), that node's interrupt parent is assumed to be its device tree parent.

3.3. "interrupt-map" property

At any level in the interrupt tree, a mapping may need to take place between the child interrupt domain and the parent's. This is represented by a new property called "interrupt-map". This property defines the mapping of *unit interrupt specifiers*, as reported by the *unit address* and *interrupt specifiers* of devices on a bus (as defined by that bus's binding), or as transformed by the "interrupt-map" property of the interrupt child of the node containing this property, to *unit interrupt specifiers* in some other interrupt domain. The "interrupt-map" property can represent wiring conventions (e.g., PCI cards with on-board PCI-to-PCI bridges, or platform routing of PCI interrupt pins) or might simply indicate a change of interrupt domain representation.

The "interrupt-map" property is a table, each entry of which consists of a child *unit interrupt specifier*, an interrupt parent *phandle* and a parent *unit interrupt specifier*. The "interrupt-map" table is looked up by matching a *unit interrupt specifier* (as masked by the "interrupt-map-mask" defined below) against child components. When a match is found, the lookup proceeds up the interrupt tree by traversing upwards to the interrupt parent (as specified by the interrupt parent *phandle* component of the matching entry) with the parent *unit interrupt specifier* component as the working interrupt value.

Note that since the interrupt parent of each entry may be a different interrupt tree node, with different values for "#address-cells" and "#interrupt-cells", the interrupt parent *phandle* of each entry must be used to determine the number of cells in the parent *unit interrupt specifier* component.

3.4. "interrupt-map-mask" property

When performing interrupt mapping via the "interrupt-map" property, not all of the bits of a *unit interrupt specifier* may be relevant to the lookup. For example, a PCI *unit interrupt specifier* consists of 4 cells (3 for the *unit address* and 1 for the *interrupt specifier*). However, only the **device#** component of the *unit address* is relevant. The "interrupt-map-mask" property defines a mask that is applied to the *unit address specifier* before using it to look up values in the "interrupt-map" table.

When the "interrupt-map-mask" property is present, the device's *unit interrupt specifier* (i.e., the concatenation of a device's *unit address* and *interrupt specifier*) is the starting point for generating a lookup value. The "interrupt-map-mask" property is a bit-mask that specifies which bits are to be considered in the looking up of "interrupt-map" entries (i.e., the child component). The *unit interrupt specifier* value is masked by ANDing each of its cells with the corresponding cell of the "interrupt-map-mask" value. The resulting value is matched against child values in the "interrupt-map" table.

3.5. "interrupt-controller" property

A sub-tree root is denoted by the presence of an "interrupt-controller" property. In general, such a node represents where some sort of processing is required to respond to an interrupt. The programming model of this interrupt controller is determined by the "name", "device_type" and/or "compatible" properties of the node.

The root of the interrupt tree is determined when the traversal of the interrupt tree reaches an interrupt controller node that does not have an explicit interrupt parent (i.e., does not have an "interrupt-parent" property).

On some platforms, the interrupts from one interrupt controller are “cascaded” into another interrupt controller higher up the interrupt tree. In these cases, the sub-tree’s interrupt controller node will, itself, have an "interrupts" property that is interpreted in the interrupt domain of its interrupt parent.

Some platforms may not have a single interrupt controller into which lower-level interrupt controllers are wired. E.g., independent interrupt controllers may report interrupts by sending interrupt “messages” directly to processors. In this case, a single device tree node can still be defined as the root of the platform’s interrupt tree in order to define the interrupt domain of the platform. I.e., the interrupt tree root does not necessarily represent a physical interrupt controller.

4. Interrupt tree parent properties

Note: interrupt controllers that are cascaded are both the parents of its child nodes and the child of an interrupt controller higher up the interrupt tree. Therefore, they will have both parent and child properties.

The following properties are defined for interrupt nexus nodes by this recommended practice.

"#interrupt-cells" S

Standard *property-name* to define the number of cells in an *interrupt specifier* within an interrupt domain.

prop-encoded-array:

An integer, encoded as with **encode-int**, that denotes the number of cells required to represent an *interrupt specifier* in its child nodes.

"interrupt-map-mask" S

Standard *property-name* to define the transformation of *unit interrupt specifiers* of child nodes into values that correspond to *child-interrupt* entries of the "interrupt-map" table entries.

prop-encoded-array:

An array of integers, each encoded as with **encode-int**.

The value of this property is a bit-mask that is applied to the concatenation of *unit address* and *interrupt specifier* for a device. The number of cells of this property is the sum of the values of the "#address-cells" and "#interrupt-cells" for this interrupt domain.

1 | "interrupt-map" S

2 |

3 | Standard *property-name* to define *interrupt specifier* mappings.

4 |

5 | *prop-encoded-array*:

6 | Arbitrary number of interrupt mapping entries.

7 |

8 | Each mapping entry consists of a 3-tuple of (*child-interrupt*, *interrupt-parent*, *parent-interrupt*). The number of cells for the *child-interrupt* specifier is determined by the "#address-cells" and "#interrupt-cells" property of this node. The number of cells for the *parent-interrupt* value is determined by the "#address-cells" and "#interrupt-cells" property values of this node's *interrupt-parent*.

9 |

10 |

11 |

12 |

13 |

14 | "interrupt-controller"

15 |

16 | Standard *property-name* to indicate an interrupt (sub-)tree root.

17 |

18 | *prop-encoded-array*.

19 | None. The presence of this property indicates that this node represents an interrupt controller.

20 |

21 |

22 | The interpretation of an *interrupt specifier* within the interrupt domain defined by the interrupt controller is defined by other properties of this node (e.g., its "device_type").

23 | 5. Interrupt tree child properties

24 |

25 |

26 |

27 |

28 |

29 | The following property is defined for children of an interrupt nexus node.

30 |

31 | Note: interrupt controllers that are cascaded are both the parents of its child nodes and the child of an interrupt controller higher up the interrupt tree. Therefore, they will have both parent and child properties.

32 |

33 |

34 |

35 | "interrupt-parent" S

36 |

37 | Standard *property-name* to denote the interrupt tree parent of this node.

38 |

39 | *prop-encoded-array*:

40 | An integer, encoded as with **encode-int**, that is the *phandle* of the interrupt nexus node that is the interrupt parent of the node.

41 |

42 |

43 | The absence of an "interrupt-parent" property in an interrupt controller node (i.e., a node that has the "interrupt-controller" property) indicates that this node represents the platform's interrupt tree root. The absence of an "interrupt-parent" property in a device node indicates that the interrupt tree parent is the device tree parent of this node.

6. Examples

6.1. PCI bus

The PCI binding defines an "interrupts" property to consist of one cell, which encodes whether the PCI device's interrupt is connected to the PCI connector's INTA#...INTD# pins, with values 1...4, respectively (assuming that the device is on a plug-in PCI card).

However, platforms typically wire the interrupts between connectors in a manner that attempts to distribute the interrupts from multiple cards across different interrupt inputs to its interrupt controller. An operating system does not obtain much information by just looking at the "interrupts" property of a device.

The "interrupts" value is insufficient to be used as the child lookup value in an "interrupt-map" table, since the device's device number (which determines the card connector into which the device is plugged) must be included in the mapping. So, for "pci" bus nodes, an "interrupt-map-mask" property must be used. The addition of this property allows the device number component of a device's *unit address* to participate in the child *interrupt specifier* lookup value.

The "interrupt-map-mask" value for "pci" would be (in hex):

0000F800	00000000	00000000	00000007
<u> </u>	<u> </u>	<u> </u>	<u> </u>
<i>unit address</i>	<i>mask</i>	<i> </i>	<i>interrupt specifier</i>
	<i>mask</i>		<i>mask</i>

The bits in the *phys.hi* component of the *unit address* mask masks off the **device#** field; the *interrupt specifier* mask masks off the low-order 3 bits, which is sufficient to cover the values 1-4.

6.2. CHRP platform

This section gives an example of how a typical CHRP ([3],[4]) platform's device tree would represent its interrupt tree. In order to understand the example, a basic introduction to the CHRP interrupt controllers is presented below.

6.2.1. Open PIC interrupts

The CHRP platform defines the platform's primary interrupt controller to be the Open PIC. This controller has a number of interrupts, each of which is represented by a *source number* and a *sense*.

Each Open PIC interrupt source is described by its source number that corresponds to a register pair for that interrupt within the Open PIC. Various fields within the register pair allow the setting of the interrupts priority, masking of the interrupt, etc. One important piece of information that must be programmed for an interrupt source is its sense; i.e., whether the interrupt is considered triggered by a positive edge or a low level.

Therefore, to represent interrupts within the domain of Open PIC, two cells are used. The first cell represents the interrupt source number, while the second specifies whether the interrupt is positive edge triggered (0) or active low level triggered (1).

6.2.2. ISA interrupt controller

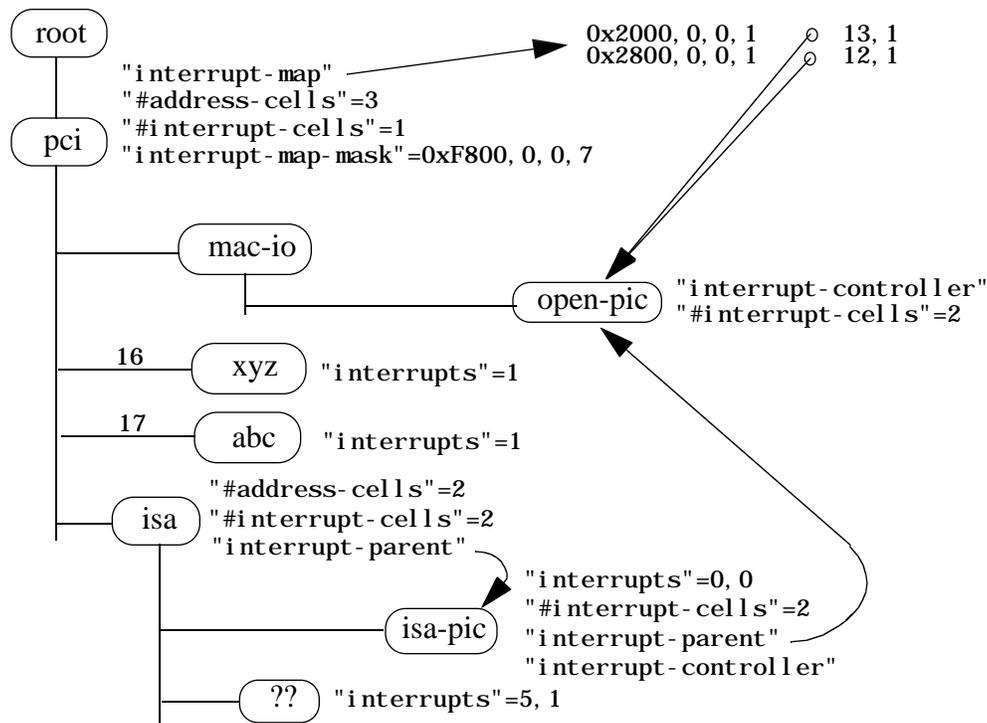
In addition to the Open PIC, a CHRP platform also has a "legacy" ISA interrupt controller. This

controller (basically, an 8259) is cascaded into the Open PIC. I.e., when an ISA interrupt occurs, an Open PIC interrupt will be generated. The processing of this interrupt will discover that the source was from the ISA interrupt controller, which must then be accessed to determine the original source of the interrupt.

The ISA binding defines the format of ISA interrupts, which consists of two cells, where the first is the interrupt level (0...15) and the second cell indicates the type (e.g., positive edge).

6.2.3. CHRP platform example

The following diagram shows how the interrupt tree for a CHRP platform would look. This example assumes that the Open PIC is contained within the "mac-io" chip. It is not important for this example to understand the details of the "mac-io"; rather, it is the structure of the interrupt tree that is being conveyed. The only important pieces of information about the "open-pic" device within "mac-io" is that it represents interrupts with 2 cells per interrupt (indicated by its "#interrupt-cells" property value) and that it is marked as an interrupt controller (by means of the "interrupt-controller" property).



The PCI devices (mac-io,xyz,abc) have an implicit interrupt parent that is the "pci" host bridge. The interrupt domain of PCI has *interrupt specifiers* that are 1 cell each, as indicated by the "#interrupt-cells" value in the "pci" node. The mac-io node does not generate PCI interrupts; it has no "interrupts" property. However, xyz and abc both generate interrupts on INTA#, as indicated by their "interrupts" values of 1.

The "pci" node contains a "interrupt-map-mask" property that indicates that the lookup of child *interrupt specifiers* is done by a combination of bits from their *unit address* and *interrupt specifiers*. The number of cells of the "interrupt-map-mask" property is 4, which is the sum of its "#address-cells"(3) and "#interrupt-cells"(1) properties.

xyz's **device#** is 16, which corresponds to a *phys.hi* of 0x2000. The corresponding entry in

1 the "interrupt-map" table indicates that its interrupt gets mapped to 13, 1 (i.e., interrupt
2 source 13, active low level) in the "open-pic". Likewise, abc's interrupt is mapped by masking
3 its *unit address* with 0xF800, 0, 0 giving 0x2800, 0, 0 and its *interrupt specifier* with 7 giving
4 a combined child lookup value of 0x2800, 0, 0, 1. The corresponding entry in the "inter-
5 rupt-map" table denotes the "open-pic" node as its parent with a *unit interrupt specifier* of
6 12, 1. Note that the "open-pic" node does not have a "#address-cells" property, so that
7 the number of cells for the parent *unit interrupt specifiers* is 2 (which is the value of its "#inter-
8 rupt-cells" property).

10
11 If the platform wired all of the INTA#s together, the interrupts would be shared. To represent this,
12 the "interrupt-map" table would show that the different PCI interrupts mapped to the same
13 *interrupt specifier* value. For example, the "interrupt-map" table above could have entries
14 for (0x2000, 0, 0, 1 open-pic 12, 1) and (0x2800, 0, 0, 1 open-pic 12, 1) indicating
15 that the INTA#s for both xyz and abc are shared.

16
17 The nodes under the "isa" bridge have interrupts encoded as specified by the ISA binding, which
18 states that "interrupts" properties consist of 2 cells, where the first cell is the interrupt number
19 (0...15) and the second cell indicates type (i.e., low level, rising edge, etc.); the "#interrupt-
20 cells" property of the bridge indicates the number of cells required to represent interrupts (i.e.,
21 2).

22
23 All of the ISA devices (with the exception of the "isa-pic", which is an interrupt controller) are
24 children of the "isa" node which becomes the default interrupt parent, since none of its children
25 have explicit "interrupt-parent" properties. Since an "interrupt-map" property is not
26 present in this node, no transformation of *interrupt specifiers* is made when traversing the interrupt
27 tree to its parent.

28
29 The interrupt parent of the "isa" node is the "isa-pic" interrupt controller. The presence of
30 the "interrupt-controller" property in the "isa-pic" node indicates that this is an in-
31 terupt controller, thus defining an interrupt sub-tree root. Since it has an "interrupt-par-
32 ent" property, this interrupt controller is cascaded into its parent, which is the platform's Open
33 PIC. The "isa-pic" interrupt is presented as interrupt source 0 (positive edge triggered) of the
34 "open-pic", as indicated by its "interrupts" value.

37 38 7. Interrupt mapping algorithm

39
40 The following psuedo-code shows the steps necessary to translate an interrupt specification from
41 a device into a platform specific value.

42
43 The psuedo-code is written at a very high level, with some liberty taken with details. The argu-
44 ments to the procedure are the device's device-node pointer and its interrupt-specifier. Note that,
45 in general, separate calls to map each interrupt specifier would be required.

46
47 The code is basically a large loop that terminates when an *interrupt-controller* node is found that
48 has no interrupt-parent.

```

1  procedure map-interrupt( device-node, interrupt-specifier )
2
3  unit-address = valueof( "reg"[0], device-node )
4  unit-interrupt-specifier = cat( unit-address, interrupt-specifier )
5  this-node = device-tree-parent( device-node )
6
7  begin      \ loop up tree until we reach the root
8
9      if present( "interrupt-controller", this-node )
10         if present( "interrupt-parent", this-node )
11             parent-node = valueof( "interrupt-parent", this-node )
12             if present( "#address-cells", parent-node )
13                 unit-address = valueof( "reg"[0], this-node )
14             else
15                 unit-address = NULL
16             then
17                 interrupt-specifier = valueof( "interrupts", this-node )
18                 unit-interrupt-specifier = cat( unit-address, interrupt-specifier )
19                 this-node = parent-node
20             else \ this is the root node, we're done
21                 return( unit-interrupt-specifier )
22             then
23
24     else      \ not "interrupt-controller"
25     if present( "interrupt-map", this-node ) \ we have a mapping to perform
26         if present( "interrupt-map-mask", this-node )
27             mask = valueof( "interrupt-map-mask" )
28             unit-interrupt-specifier = unit-interrupt-specifier & mask
29         then
30             init-decode-cells( "interrupt-map", this-node )
31             found? = false
32             begin
33                 child-specifier = decode-cells( sizeof( unit-interrupt-specifier ) )
34                 parent-node = decode-cells( 1 )
35                 if present( "#address-cells", parent-node )
36                     #cells = valueof( "#address-cells", parent-node )
37                 else
38                     #cells = 0
39                 then
40                     #cells = #cells + valueof( "#interrupt-cells", parent-node )
41                     if child-specifier == unit-interrupt-specifier
42                         found? = true
43                     else
44                         dummy = decode-cells( #cells )
45                     then
46                         until found?
47                         interrupt-specifier = decode-cells( #cells )
48                         if present( "#address-cells", parent-node )
49                             unit-address = valueof( "reg"[0], this-node )
50                         else
51                             unit-address = NULL
52                         then
53                             unit-interrupt-specifier = cat( unit-address, interrupt-specifier )
54                             this-node = parent-node
55
56     else      \ no "interrupt-map" table
57     if present( "interrupt-parent", this-node )
58         this-node = valueof( "interrupt-parent", this-node )
59
60     else      \ no "interrupt-parent" property
61         this-node = device-tree-parent( this-node )
62     then
63
64     again

```