

IBDM

IBNL

A Generic IB Netlist Format

Revisions Table:

Revision	Date	Updated By	Modification High Level
1.0	29/9/2004	Eitan Zahavi	Initial revision
1.1	5/7/2005	Eitan Zahavi	Remove IBADM from the definition. Real effect only on target dir for the package: \$prefix/lib/ibdm1.0

Overview

IBDM topology file provide means to describe the IB fabric using a set of predefined systems. A system definition is provided in a single file in IBNL format that describes the internal InfiniBand connectivity of the system in terms of boards and devices. When IBDM starts it parses all the available system definition files before it handles the topology file. The files are located in the following directory relative to the installation prefix: <prefix>/lib/ ibdm1.0/ibnl. Any new system that is described using the IBNL file format and has a corresponding file in this directory is supported by IBDM.

This appendix describes the IBNL file format used to define an arbitrary IB system internal IB connectivity. It outlines the main concepts used by the file, provides details for how to obtain such a file, a formal definition of the file syntax in BNF like format (YACC/Bison readable) and concludes with real life example.

1. Main Concepts:

The described file format for defining generic system IB connectivity uses the following terminology:

NODE – The instantiation of an IB device may it be a switch or a channel adapter.

SUB-SYSTEM – The instantiation of a board or module in a system.

SYSTEM – A definition of a board or a module.


TOPSYSTEM – The top most system defined in the given file.

SUB-SYSTEM MODIFIER – Many systems allow multiple variations of their sub-system such as a 12X version of a leaf board in a MTS9600. The modifier is a suffix to the board name and should follow the format “: <any name>”. The IBNL format supports assigning multiple names to the same board.

2. File Format:

The IBNL file is line sensitive as it requires different sections to start on a new line. The file is broken into several SYSTEM (optional) and one TOPSYSTEM sections. Each has one or more names. The following diagram describes the file in more details:

Figure 1 - IBNL File Sections



Connections are defined inside the system level only and thus might be of two types:

1. Between any node or sub-system to other node or sub-system
2. From any node or sub-system to a connector of the system.

Note: The top system can define ports that might be redundant and not connected to any lower level board connector. In these cases the ports are simply omitted from the resulting system. This feature enables defining the front panel ports of a 12X/4X such that if a 12X leaf is selected only the 12X front panel ports are used.

Another important note is that port width and speed provided at the lowest level have precedence over definitions provided at upper levels of the hierarchy.

3. Writing a System IBNL File

The following procedure should be followed in order to provide a new system IBNL:

- a) Name the file after the system name: <any sys name (no spaces)>.ibl
- b) Define a SYSTEM section for each board included in the system
- c) The port names of the boards are just strings, we use the simple format of P<N> where N is just a serial number but you can pick any name you want. Just make sure it is unique.
- d) When different flavors of the boards exists – like when you have a 4x and 12x option for a board – name the optional boards with a “modifier postfix”. An example for such system that supports a LEAF board of 12x and 4x types would be to define two SYSTEMS:

```
SYSTEM LEAF,LEAF:4x,LEAF:4X
```

```
...
```

```
SYSTEM LEAF:12x,LEAF:12X
```

```
...
```

Note that the instantiations of the LEAF boards in another board or the top system need not specify the postfix and are being decided upon only when the topology file is being parsed. The “SYSTEM LEAF” with no postfix will be used by default. To continue the example here is how the LEAFs are instantiated in the top system:

```
TOPSYSTEM MyIbSystem
```

```
LEAF leaf1
```

```
...
```

```
LEAF leaf2
```

```
...
```

The actual 4x or 12x version of the LEAF board can then be specified in the topology file CFG section to select the right combination of optional boards in the system. An example will be:

```
MyIbSystem N001 CFG: leaf2=12x
```

In this case leaf1 will be 4x as no special modifier is defined for it (and LEAF is by default a 4x leaf). Leaf2 will be 12x as defined in the CFG section,

- e) Special considerations should be made to design the top system section when several optional board types exist. Top system section might include multiple definitions for front panel plugs like P1-4x and P1-12x (yes these are just names that should follow the writing on the front or back panels). As the different flavors of the boards are not required to define the same names for their ports – including a 12x port might leave some of the top level connections dangling (not connected) and thus the front panel ports of the other flavor will be removed from the final system definition.

As an example let us consider a simple board LEAF with 3 4x port flavor and one 12x port flavors. We recommend connecting it in the top level using the following scheme:

```
SYSTEM LEAF:4x
```

```
NODE ... U1
```

```
  1 -4x-> 4XP1
```

```
  2 -4x-> 4XP2
```

```
  3 -4x-> 4XP3
```

```
.....
```

```
SYSTEM LEAF:12x
```

```
NODE ... U1
```

```
  1 -12x-> 12XP1
```

```
.....
```

```
TOPSYSTEM myIbSystem
```

```
SUBSYSTEM LEAF leaf1
```

```
  4XP1 -> L1/P1
```

```
  4XP2 -> L1/P2
```

```
  4XP3 -> L1/P3
```

```
 12XP1 -> L1/P1-12x
```

- f) Place the file in the <prefix>/lib/IBDM/ibdm1.0/ibnl directory
- g) To check the file syntax you have two options:
- Using the ibnlparse utility:
Usage: ibnlparse <ib netlist file> <sys type>
Where the system type is the top system name.
 - By parsing it in context of a topology file using the utility ibdmtr (trace a direct route through the fabric):
 - Define a topology file with your new system
 - Set some CFG options if applicable
 - Invoke: ibdmtr -t <topo file> -s <any node>/U1 -p 1 -d 0,1

4. Formal definition in YACC Syntax:

```
INT ::= ([1-9][0-9]*|0) ;

WIDTH ::= (4x|8x|12x) ;

SPEED ::= (2.5G|5G|10G) ;

NODETYPE ::= (SW|CA|HCA) ;

NAME ::= ([\[\]\*\*/A-Za-z0-9_.:%@\~]+) ;

NL:
  LINE
  | NL LINE;

ONL:
  | NL;

ibnl: ONL systems topsystem;

systems:
  | systems system ;

sub_inst_attributes:
  | sub_inst_attributes sub_inst_attribute NL;

sub_inst_attribute:
  NAME '=' NAME '=' NAME
  | NAME '=' NAME '=' INT
  | NAME '=' NAME ;

topsystem:
  TOPSYSTEM system_names NL sub_inst_attributes insts ;

system:
  SYSTEM system_names NL insts ;

system_names:
  system_name
  | system_names ',' system_name ;

system_name:
  NAME ;

insts:
  | insts node
  | insts subsystem ;

node:
  node_header NL node_connections ;

node_header:
  NODE NODETYPE INT NAME NAME ;
```

```

node_connections:
  | node_connections node_connection NL ;

node_connection:
  node_to_node_link
  | node_to_port_link ;

node_to_node_link:
  INT '-' WIDTH '-' SPEED '-'>' NAME
  | INT '-' WIDTH '-'>' NAME INT
  | INT '-' SPEED '-'>' NAME INT
  | INT '-'>' NAME INT ;

node_to_port_link:
  INT '-' WIDTH '-' SPEED '-'>' NAME
  | INT '-' WIDTH '-'>' NAME
  | INT '-' SPEED '-'>' NAME
  | INT '-'>' NAME ;

subsystem:
  subsystem_header NL subsystem_connections ;

subsystem_header:
  SUBSYSTEM NAME NAME ;

subsystem_connections:
  | subsystem_connections subsystem_connection NL ;

subsystem_connection:
  subsystem_to_subsystem_link
  | subsystem_to_port_link ;

subsystem_to_subsystem_link:
  NAME '-' WIDTH '-' SPEED '-'>' NAME NAME
  | NAME '-' WIDTH '-'>' NAME NAME
  | NAME '-' SPEED '-'>' NAME NAME
  | NAME '-'>' NAME NAME ;

subsystem_to_port_link:
  NAME '-' WIDTH '-' SPEED '-'>' NAME
  | NAME '-' WIDTH '-'>' NAME
  | NAME '-' SPEED '-'>' NAME
  | NAME '-'>' NAME ;

```

5. Example IBNL for MTS14400

We provide here a partial (one spine and 2 leafs) IBNL for MTS14400

SYSTEM LEAF,LEAF:4x,LEAF:4X

NODE SW 24 MT47396 U1

- 1 -> P13
- 2 -> P14
- 3 -> P15
- 4 -> P16
- 5 -> P17
- 6 -> P18
- 7 -> P19
- 8 -> P20
- 9 -> P21
- 10 -> P22
- 11 -> P23
- 12 -> P24
- 13 -> P1
- 14 -> P2
- 15 -> P3
- 16 -> P4
- 17 -> P5
- 18 -> P6
- 19 -> P7
- 20 -> P8
- 21 -> P9
- 22 -> P10
- 23 -> P11
- 24 -> P12

SYSTEM LEAF:12x,LEAF:12X

NODE SW 24 MT47396 U1

- 1 -12x-> P13
- 4 -12x-> P14
- 7 -12x-> P15
- 10 -12x-> P16
- 13 -> P1
- 14 -> P2
- 15 -> P3
- 16 -> P4
- 17 -> P5
- 18 -> P6
- 19 -> P7
- 20 -> P8
- 21 -> P9
- 22 -> P10
- 23 -> P11
- 24 -> P12

SYSTEM SPINE,SPINE:4x,SPINE:4X

NODE SW 24 MT47396 U1

- 21 -> P63
- 22 -> P30
- 23 -> P18
- 24 -> P6
- 1 -> P27
- 2 -> P15
- 3 -> P3
- 4 -> P31
- 5 -> P19
- 6 -> P7
- 7 -> P34
- 8 -> P22
- 9 -> P10

10 -> P43
11 -> P55
12 -> P67
13 -> P46
14 -> P58
15 -> P70
16 -> P42
17 -> P54
18 -> P66
20 -> P51
19 -> P39

NODE SW 24 MT47396 U2

10 -> P44
11 -> P56
12 -> P68
13 -> P47
14 -> P59
15 -> P71
16 -> P41
17 -> P53
18 -> P65
20 -> P50
19 -> P38
21 -> P62
22 -> P29
23 -> P17
24 -> P5
1 -> P26
2 -> P14
3 -> P2
4 -> P32
5 -> P20
6 -> P8
7 -> P35
8 -> P23
9 -> P11

NODE SW 24 MT47396 U3

9 -> P12
10 -> P45
11 -> P57
12 -> P69
13 -> P48
14 -> P60
15 -> P72
16 -> P40
17 -> P52
18 -> P64
20 -> P49
19 -> P37
21 -> P61
22 -> P28
23 -> P16
24 -> P4
1 -> P25
2 -> P13
3 -> P1
4 -> P33
5 -> P21
6 -> P9
7 -> P36
8 -> P24

TOPSYSTEM MTS14400,Rhino

spine1/U1=isBmaAgent

SUBSYSTEM LEAF leaf1

P1 -> spine1 P1

P2 -> spine1 P2
P3 -> spine1 P3
P4 -> spine1 P4
P5 -> spine1 P5
P6 -> spine1 P6
P7 -> spine2 P7
P8 -> spine2 P8
P9 -> spine2 P9
P10 -> spine2 P10
P11 -> spine2 P11
P12 -> spine2 P12
P13 -> L1/P1
P14 -> L1/P2
P15 -> L1/P3
P16 -> L1/P4
P17 -> L1/P5
P18 -> L1/P6
P19 -> L1/P7
P20 -> L1/P8
P21 -> L1/P9
P22 -> L1/P10
P23 -> L1/P11
P24 -> L1/P12

SUBSYSTEM LEAF leaf2

P1 -> spine1 P13
P2 -> spine1 P14
P3 -> spine1 P15
P4 -> spine1 P16
P5 -> spine1 P17
P6 -> spine1 P18
P7 -> spine2 P19
P8 -> spine2 P20
P9 -> spine2 P21
P10 -> spine2 P22
P11 -> spine2 P23
P12 -> spine2 P24
P13 -> L2/P1
P14 -> L2/P2
P15 -> L2/P3
P16 -> L2/P4
P17 -> L2/P5
P18 -> L2/P6
P19 -> L2/P7
P20 -> L2/P8
P21 -> L2/P9
P22 -> L2/P10
P23 -> L2/P11
P24 -> L2/P12

SUBSYSTEM SPINE spine1

P30 -> leaf3 P6
P18 -> leaf2 P6
P6 -> leaf1 P6
P27 -> leaf3 P3
P15 -> leaf2 P3
P3 -> leaf1 P3
P31 -> leaf10 P6
P19 -> leaf11 P6
P7 -> leaf12 P6
P34 -> leaf10 P3
P22 -> leaf11 P3
P10 -> leaf12 P3
P43 -> leaf9 P6
P55 -> leaf8 P6
P67 -> leaf7 P6
P46 -> leaf9 P3
P58 -> leaf8 P3
P70 -> leaf7 P3
P42 -> leaf4 P6

P54 -> leaf5 P6
P66 -> leaf6 P6
P39 -> leaf4 P3
P51 -> leaf5 P3
P63 -> leaf6 P3
P29 -> leaf3 P5
P17 -> leaf2 P5
P5 -> leaf1 P5
P26 -> leaf3 P2
P14 -> leaf2 P2
P2 -> leaf1 P2
P32 -> leaf10 P5
P20 -> leaf11 P5
P8 -> leaf12 P5
P35 -> leaf10 P2
P23 -> leaf11 P2
P11 -> leaf12 P2
P44 -> leaf9 P5
P56 -> leaf8 P5
P68 -> leaf7 P5
P47 -> leaf9 P2
P59 -> leaf8 P2
P71 -> leaf7 P2
P41 -> leaf4 P5
P53 -> leaf5 P5
P65 -> leaf6 P5
P38 -> leaf4 P2
P50 -> leaf5 P2
P62 -> leaf6 P2
P28 -> leaf3 P4
P16 -> leaf2 P4
P4 -> leaf1 P4
P25 -> leaf3 P1
P13 -> leaf2 P1
P1 -> leaf1 P1
P33 -> leaf10 P4
P21 -> leaf11 P4
P9 -> leaf12 P4
P36 -> leaf10 P1
P24 -> leaf11 P1
P12 -> leaf12 P1
P45 -> leaf9 P4
P57 -> leaf8 P4
P69 -> leaf7 P4
P48 -> leaf9 P1
P60 -> leaf8 P1
P72 -> leaf7 P1
P40 -> leaf4 P4
P52 -> leaf5 P4
P64 -> leaf6 P4
P37 -> leaf4 P1
P49 -> leaf5 P1
P61 -> leaf6 P1