



Adapting PnP to RMAP and Other Recent Changes

Peter Mendham, Albert Ferrer Florit,
Stuart Mills, Steve Parkes

Space Technology Centre
University of Dundee



Overview

- Principles
- Simplifications and additions
- “Ownership” of resources and RMW
- Notification mechanism
- Service specification



Principles of PnP

- Central goal of PnP is interoperability at the network level
 - Promote hardware and software reuse
 - Create more potential for off-the-shelf components
 - Permit network discovery and verification
- PnP should provide support for features defined in the SpaceWire standard
- If it is optional in the SpaceWire standard it should be optional in PnP
- It should be possible to detect what features hardware supports
- There should be methods of working with existing hardware



Principles of PnP using RMAP

- Try to keep things as simple as possible
- Use the features of RMAP appropriately
 - Verified/acknowledged writes
 - Read-modify-write
 - Error codes
- Separate read-only and read-write parameters
 - Separates interrogation/configuration use cases
 - Less likely to perform accidental writes
- Group related parameters together into structures/data-types
 - Create a logical arrangement of parameters
 - Support common use cases



Simplifications

- Original proposal had “data-types”
- Packet containing a data-type always contains the entire data-type
- Data-types have become “structures”
 - We can view these as structures in memory
 - Or as larger data-types
 - Name change is purely to distinguish the change to RMAP
- RMAP can read/write parts of a structure
- Create larger structures than in original proposal
- Remove anything that is not strictly necessary



Additions

- Since September (ISC2007)
- Time-code handling support
- Time-code generation support
- Protocol ID counters
 - Number of unrecognised protocol IDs
 - Number of extended protocol IDs
- Document restructure
 - Refer to, rather than repeat, detail from the RMAP standard



Two Issues in the Current Proposal

- “Ownership” mechanism
- Notification mechanism
- Features in RMAP can help address these issues
- Will examine each in turn



“Ownership” and RMW - Introduction

- In a multi-host system, hosts may compete for authority over parameters
- Solved by permitting a host to “own” a device
- Notification table entries (slots) may also be owned by hosts
- For this to work, there must be an atomic method for setting the ownership ID
 - Otherwise there will be race for ownership



“Ownership” and RMW – Original Method

- Original method used conditional write behaviour built into hardware
- Used read, write and reset commands
- A value can only be written if it is currently in its reset state
- Used as follows (pseudo-code):

```
// Acquire ownership
do {
    write(id_location, my_id);
    test_id = read(id_location);
} while (test_id != my_id);
// Do things
...
// Release Ownership
reset(id_location);
```



“Ownership” and RMW – Naïve RMAP Port

- Ported directly to RMAP
- No reset command, use a write of the reset value
- Behaviour the same:
 - If a parameter is in the reset state it can be written with any value
 - If a parameter is not in the reset state, only the reset value can be written

```
// Acquire ownership
do {
    write(id_location, my_id);
    test_id = read(id_location);
} while (test_id != my_id);
// Do things
...
// Release Ownership
write(id_location, reset_value);
```

Reset became a write





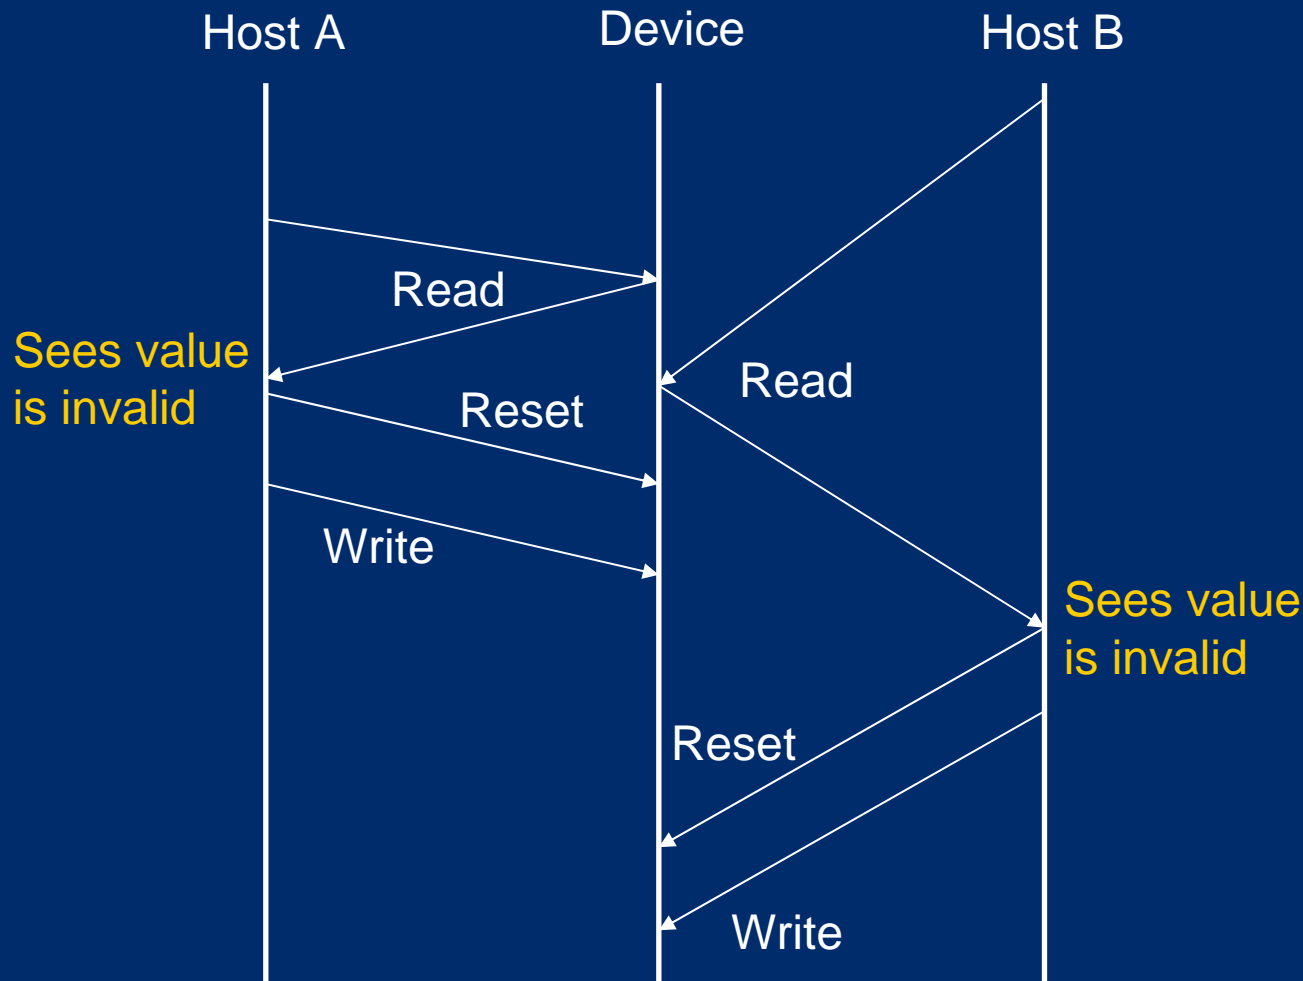
“Ownership” and RMW – Remaining Problem

- The reset value is invalid and therefore “special”
- Any other value is not
- If a host owns the parameter and then dies there is a race condition for re-ownership
- If a device is disconnected and re-connected without being reset there is a race condition
- Let’s examine what happens once it is known that the owner of a device is no longer valid
 - The parameter has a value, but this value is not the reset value
 - And yet the value is invalid and must be treated as “special”



Race Condition Explanation

- Upon seeing value is invalid, each host tries for ownership





“Ownership” and RMW – Solution using RMAP

- Implement RMW as conditional write
- Mask value is used for verification
- Write will only complete if the current value matches the mask value
- “Special” value is mask value
- Any value can be made “special”
- Solves race condition problem
- Is more consistent with RMAP



“Ownership” and RMW – Example Usage

- Read-Modify-Write command takes mask value and data value, returns read value
- If the read returns the reset value, which should also be used for the mask, the write succeeded

```
// Acquire ownership
do {
    test_id = rmw(invalid_value, my_id);
} while (test_id != invalid_value);

// Do things
...

// Release Ownership
write(id_location, reset_value);
```

Annotations in the code block:


- Mask value**: Points to the `invalid_value` argument in the `rmw` function call.
- Data value**: Points to the `my_id` argument in the `rmw` function call.
- Read value**: Points to the `test_id` variable being assigned the result of the `rmw` function call.



Notification Mechanism – Current Situation

- Routers to sends messages to hosts when ports connect or disconnect
- Currently uses acknowledgements and retries
 - Router expects an acknowledgement
 - If it does not get one after a timeout, it will retry
- If no acknowledgement, retries infinitely
 - This would occur if there was a problem with a host
- Each slot in the notification table is owned
- No way to identify if owner is still a valid/present host
 - If hosts repeatedly become unavailable, slots will get used up
- **Note: None of this works with current hardware**

Notification Mechanism - Proposal



- Timeout table and event table
- Timeout table entry has two timeouts
 - Reply timeout
 - Re-read timeout
- If a timeout table entry is non-zero it is being used
- Can use RMW for access



Notification Mechanism – Proposed Process

- Set timeouts in empty (zero-valued) slot
 - Reply timeout
 - Re-read timeout
- Read corresponding slot in event table
- Device will reply when:
 - Event occurs (connect/disconnect)
 - Reply timeout elapses
- Host should then read again
- If any events happened in the mean time, reply will be immediate
- If no read before re-read timeout, timeouts cleared (slot free)



Notification Mechanism – Proposal Overview

- No infinite retries
- No problems if host becomes unavailable
- Host can keep track of timeout and know if
 - Router is no longer active
 - Packet was lost
- Slight increase in network traffic
- Long timeout reduces network traffic but means longer before detection of lost packets
- On most SpaceWire networks this shouldn't be a problem
- Would be a problem if there was a high likelihood of lost packets



Service Identification – The Problem

- Need to know:
 - What a device is (i.e. device type)
 - How to talk to it (available services)
- The essential parts of an electronic datasheet
- Device type is already present
- Want a simple way of detailing the services a device provides and to what level they are supported
- “Service” defined as
 - A given protocol (from an ID)
 - Providing a given service
 - With a level of support



Service Identification – Proposed Solution

- Propose a simple binary table
- Lists protocols with standard IDs
- Protocols are responsible for defining standard services (if applicable)
- Both protocols and services supply:
 - Version number
 - Support bitmap (defined by protocol/service)
- For example:
 - RMAP protocol
 - Services are standardised address spaces
 - Support bitmap equivalent to RMAP conformance tables



Summary

- Current proposal applies RMAP
- Uses features of RMAP to simplify the proposal
- Conditional write should be swapped for standard RMW due to race condition
- Alternative notification mechanism should be considered
- Service identification should be discussed and considered



Backup Slides

“Leading Zero” Issue

- In order to discover networks, there must be a consistent method of addressing nodes and routers
- Address routers as follows:



Configuration
Port



- Routers configuration port “sees” everything from LA to EOP
- If a similar packet is sent to a node, the node will “see” everything from 0 to EOP
- Network discovery packets arriving at node will have an extra leading zero

Handling the Leading Zero at Nodes

- For nodes that support network discovery
 - The leading zero indicates that the packet refers to configuration/discovery of the node (the configuration port)
 - This is in line with the SpaceWire standard (it is not explicitly included, or prevented, but it is conceptually consistent)
- For nodes that do not support network discovery (e.g. legacy nodes)
 - If the node supports the PID specification, the packet will not conform and will therefore be discarded
 - If the node does not support the PID specification behaviour is undefined
 - If the node supports a leading zero (i.e. a configuration port) but not network discovery it will check the PID, which will be a PID it does not handle, it will then discard the packet



Standardising the Leading Zero

- The PnP proposal standardises operation through the configuration port for both routers and nodes
- No additional behaviour is required of other nodes beyond what is already standardised
- Does handling of the behaviour of a leading zero need adding to a standard?
 - Is this really normative?
 - Or is it just informative?
- Note: the leading zero is *not* an anomaly, it is part of the valid operation of PnP, which consistently uses a configuration port



Whoa, There!

Standardising RMAP Address Spaces

- RMAP is a very useful protocol
- It may be useful to standardise some RMAP address spaces at a WG level
- Each address space also would need to specify a set of operating practices
 - Fixed, or how to determine maximum packet sizes
 - RMW implementation (if supported)
 - May need to define further semantics of operation (like notification under PnP)
- Question for the WG: how should these address spaces be standardised and identified?