Topology and Protection issues

Italo Busi, Alcatel (Presenter)

Anoop Ghanwani, Lantern

Jeanne De Jaegher, Alcatel

Pietro Grandi, Alcatel

Vittorio Mascolo, Alcatel

Draft 03 - XX January 2002

Table of Contents

Introduction

Scope of the Presentation

- Describe the Alladin proposal for the topology discovery
  - Some minor changes have been made since the last meeting
- Provide motivation on why the Alladin topology discovery better suits for the RPR requirements

- Describe the Alladin proposal for the protection switching
  - Some changes have been made since the last meeting
- Provide motivation on why the Alladin protection better suits for the RPR requirements

Topology Discovery Protocol

Protocol Requirements – 1

- It shall be scalable to 255 nodes on the ring
- It shall determine the topology of the ring in terms of the list of nodes, their order and their interconnections
- It shall validate the topology of the ring
- It shall ensure that all nodes on the ring have the same topology map
- It shall automatically update the topology after a station insertion/removal
- It shall be tolerant to message losses
- It shall cause minimal overhead

Protocol Requirements – 2

- It shall operate without any master node in the ring
- It shall be able to work independently on the presence or absence of a network management system
- It shall provide a mean for station to exchange some pieces of information
  - Standard information – e.g. for auto-negotiating the options supported by the nodes along the ring or to distribute the weights
  - Private information (i.e. proprietary vendor specific information)
- The topology image built by the protocol will be used by other RPR control protocols (e.g. the protection switching protocol and the fairness algorithm)

Topology Changes

- The ring topology consists of list of stations, their order and their interconnections
  - The topology changes only because of a node insertion or removal
- It is outside the scope of the topology discovery protocol to discover the dynamic status of the links
  - The protection protocol has been designed for such a scope
- Separating the topology discovery and the protection protocols is a good architectural design
  - They work independently each other – no redundancy of information
  - They can be developed separately
  - They can be tested separately
  - They can be managed separately by the service providers

Basic principles
Initiating a topology exchange procedure

- In order to cause minimal overhead the topology information should be exchanged only when needed
- Topology change events happen very rarely – there is no need to continuously update the topology of the network
- Adjacent nodes periodically check (e.g. each second) if they have a consistent topology image
  - The topology exchange is initiated only when two adjacent nodes detects a misalignment between their topology databases or when a node is inserted or removed from the ring

Neighbor_Hello Messages

- Periodically exchanged between two adjacent nodes on each ringlet
  - The ringlet_id field allow to validate the ringlet configuration (e.g. mis-cabling errors)
  - The ring_image_version, that summarizes the topology information of the node, is used to check that two adjacent nodes have the same topology image
  - The neighbor MAC address is derived from the SA of the message

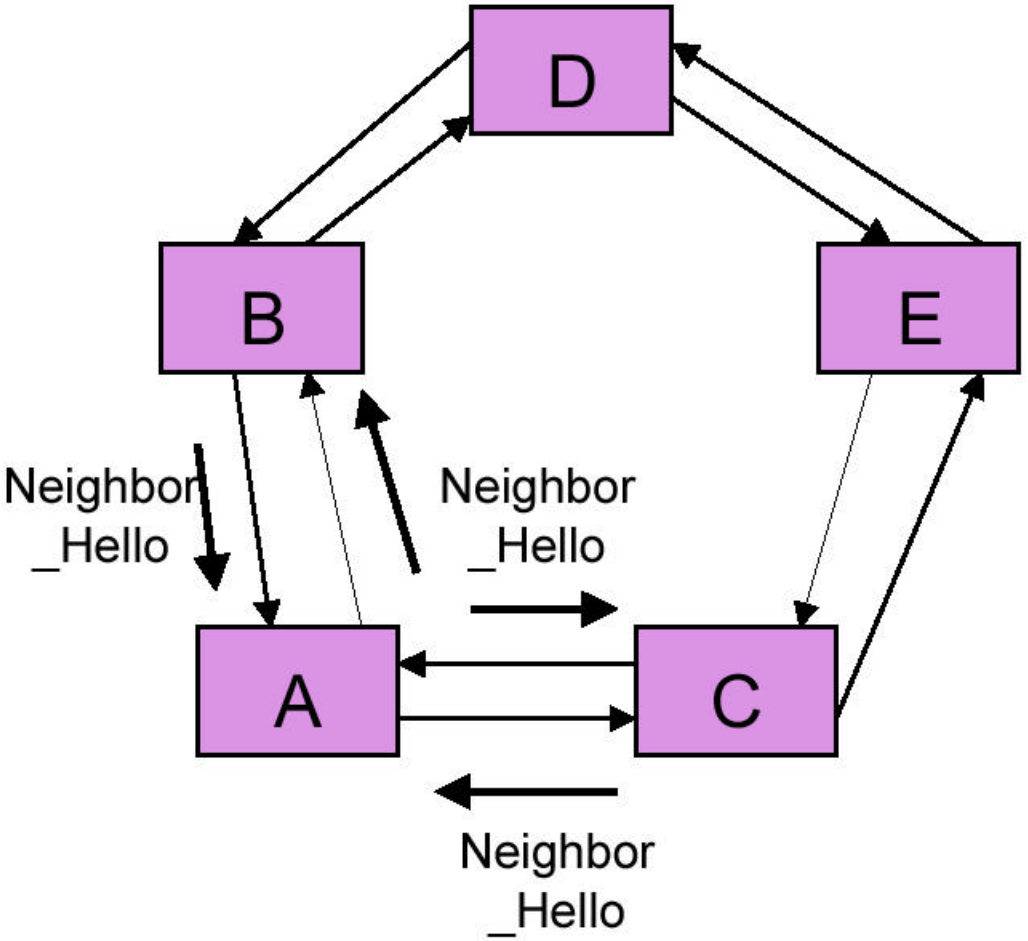

Topology_Status Messages

- Broadcast messages on the ring to report the local status of the node
  - MAC addresses of the neighbors
  - The station image version that is useful to recognize the most updated information
  - The options supported by the node (for auto-negotiation purposes)
  - The weight of the node
  - Private data for vendor specific proprietary extensions
- They are sent by any node that detects
  - A neighbor change
  - A validation failure

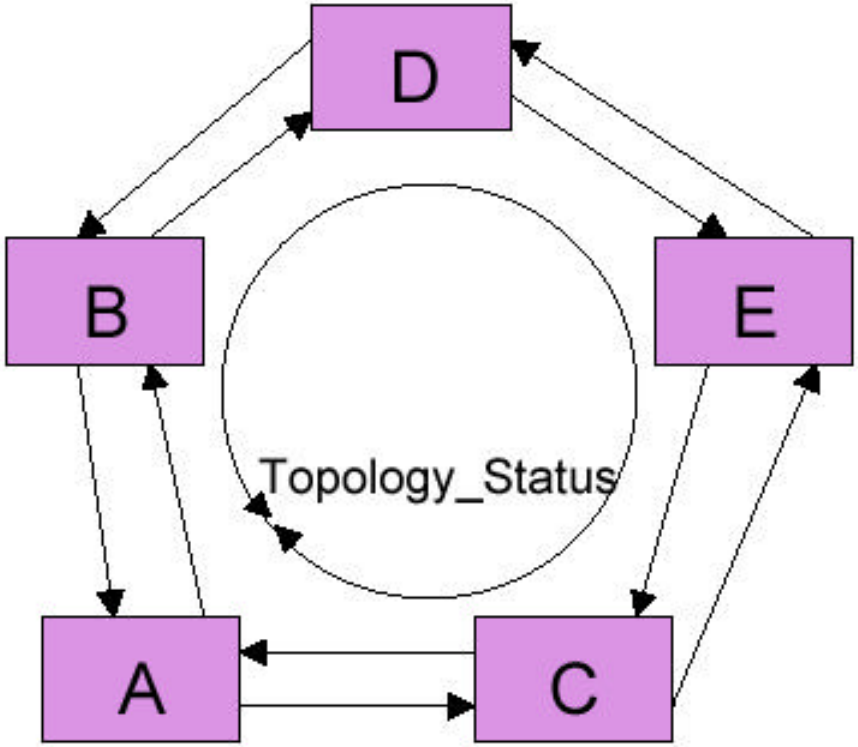

Basic Principles
Node initialization

- Local topology image is initialized containing only the local node, no links and unknown neighbors – the station image version is set to 0 (new station)
- A Topology_Status is broadcasted on all the ringlets
- The Neighbor_Hello messages are periodically sent to all the neighbors
- Any node in the ring that receives a Topology_Status with a station image version equal to 0, understands that there is a new station on the ring and broadcasts a Topology_Status on all its ringlets
  - By listening to these Topology_Status messages the new node discovers the ring topology
- When the node receive a Neighbor_Hello message from a neighbor node, it discover its neighbor, updates the local topology map and broadcasts a Topology_Status message on all the ringlets
  - By listening to these Topology_Status messages from the new node, all the other nodes in the ring discover where the new node has been inserted

Protection Protocol

Protocol Requirements – 1

- Protection switching should be performed in less than 50 ms
- The protocol shall scale to up 255 nodes
- It shall support a quick dissemination of protection states along the ring
- It shall be tolerant to message losses
- It shall operate without any master node in the ring
- It shall be able to work independently on the presence or absence of a network management system
- It shall work independently from the topology discovery protocol
  - It only assumes that the ring topology has been discovered by the topology discovery protocol

Protocol Requirements – 2

- It shall be independent on the protection type (steering or wrapping)
- It shall cause minimal overhead
- Protection is triggered by both Layer 1 or Layer 2 events
  - Keep_Alive messages are periodically sent (e.g. each ms) to detect Layer 2 failures
- It should guarantee the highest possible availability
  - Islands should be formed in case of multiple faults
- The protocol is always initiated by the nodes that detect a fault or receive an operator command
- The ring shall react quickly also when protection messages are lost
  - It is very beneficial to continuously send protection messages

Protection Triggers

- Layer 1 triggers
  - Signal Fail
  - Signal Degrade (optional)
- Operator Commands
  - Manual Switch
  - Force Switch
- Layer 2 triggers
  - Loss of Keep-Alive messages
    - It is handled as a Signal Fail from the protection protocol point of view

Protection Events

- Protection events have to be processed according to their priority
  - FS – Forced Switch – operator generated (highest priority)
  - SF – Signal Fail – automatically generated
  - SD – Signal Degrade – automatically generated
  - MS – Manual Switch – operator generated
  - WTR – Wait Time to Restore – automatically generated
  - NR – No Request present (lowest priority)
- Automatically generated events are subject to the hold-off timer
- Only events with the same priority can co-exist on the ring
  - The only exception are the SF and FS that can co-exist
  - The low priority events are withdrawn as long as the higher priority events are present

Protection Notifications

- Protection notifications have to be emitted as soon as a failure is detected
  - Protection notifications have to be continuously sent to be tolerant to message losses and node restarts
  - Protection notifications have to be sent also to signal that there are no requests on the ring

- It seems useful to use the Keep_Alive messages to notify the protection status of the ring
  - The MAC address field represent the node that has the request
  - The flags field represent the direction of the request (same or opposite ringlet)
  - The upstream_link_status field represent the kind of request

Protection Notification Propagation

- When there are no requests, Keep_Alive messages are sent when the emission timer expires with all the fields set to 0
- As soon as a local request is detected, Keep_Alive messages are asynchronously sent on all the ringlets without waiting for the timer expiration
- As soon as an incoming request is received by a neighbor, its priority is checked with the local request
  - If the local request is a SF or a FS, nothing is done
  - If the local request has higher or equal priority than the incoming one, nothing is done
  - If the local request has lower priority than the incoming one, a Keep_Alive message is asynchronously sent on the same ringlet without waiting for the timer expiration

Ring with no Faults

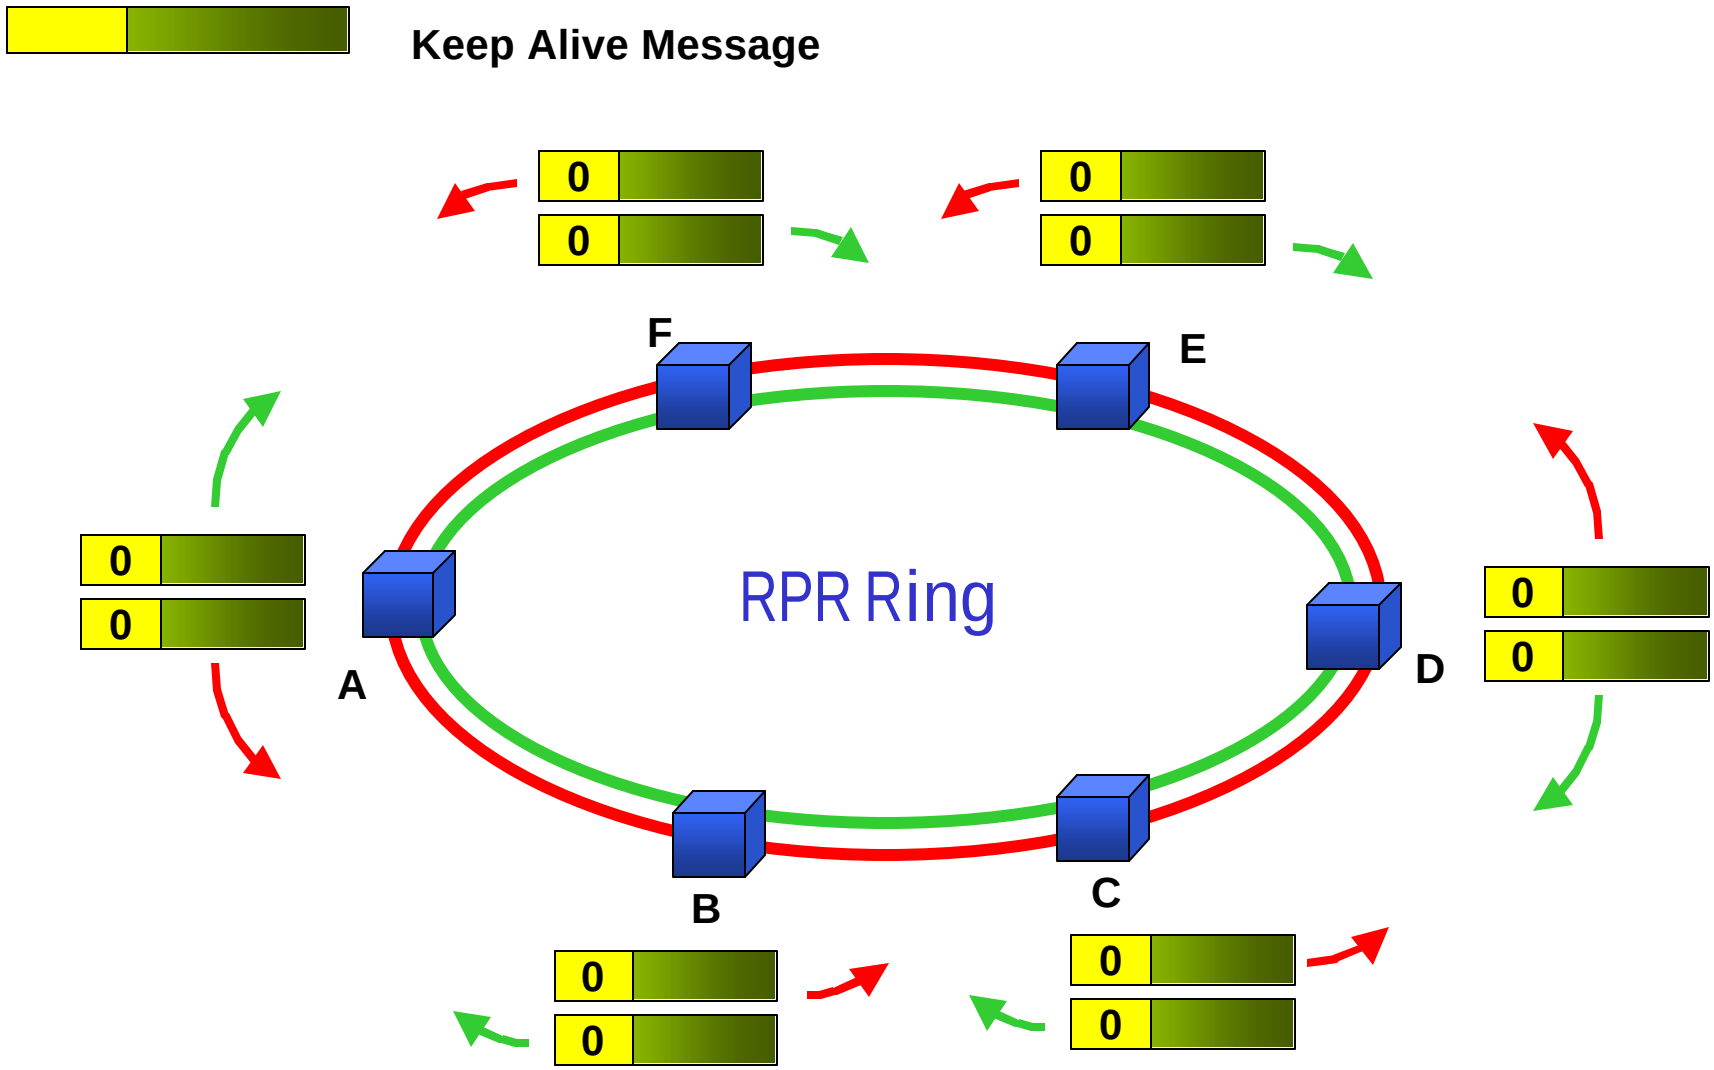

Ring with one span failure (SF)

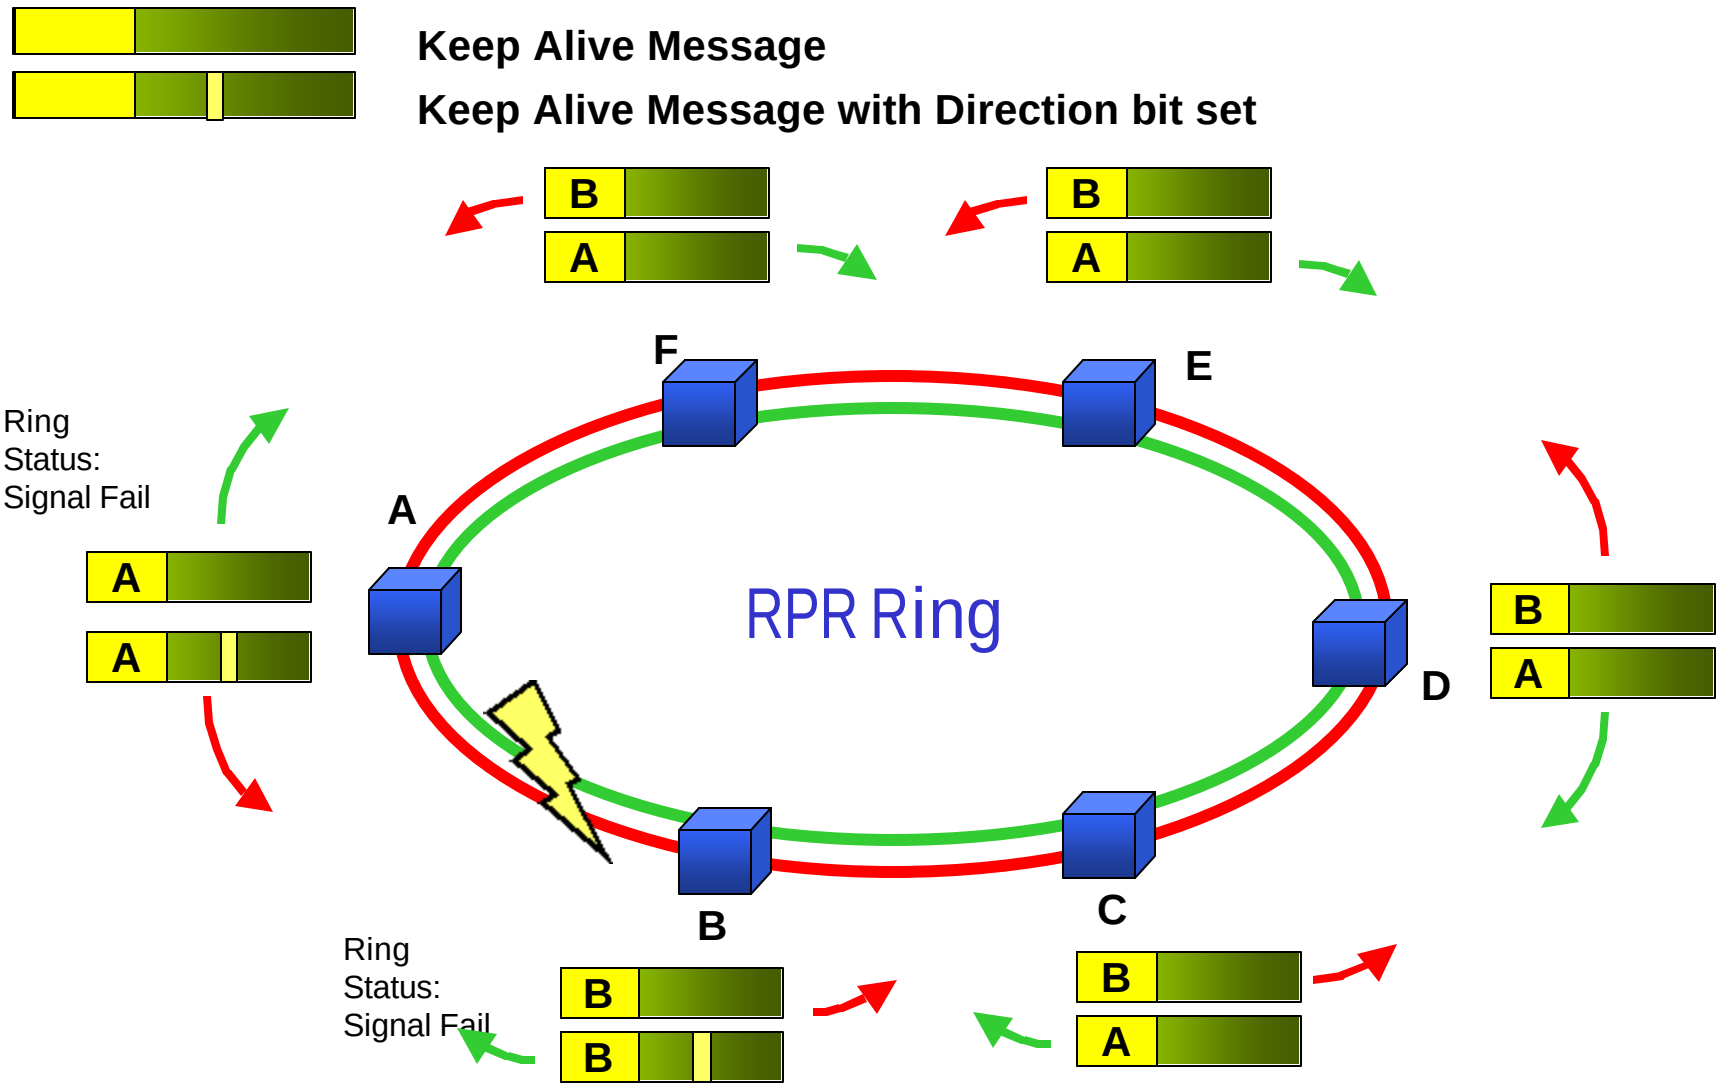

Conclusion

Conclusions – 1

- It is important to decouple the topology discovery from the protection protocols
  - The two protocols presented are able to work independently each other
- The topology does not change because of a link failure
  - Link failures are detected and propagated only by the protection protocol
- The topology exchange is performed only when needed
  - This allow to minimize the overhead – topology changes are very rare events

Conclusions – 2

- The protection notification mechanism proposed is derived from the well-established mechanism in SONET/SDH rings

- The protection protocol is responsible to notify protection requests along the ring
  - A mechanism based on continuously sending information among the rings, like in SONET/SDH networks, is proposed
- The protection protocol notifies only the needed requests
  - Low priority requests are not notified
- The protection protocol is responsible to detect layer 2 failures
  - The Keep_Alive messages have been designed for this purpose