

Slides
rev-09

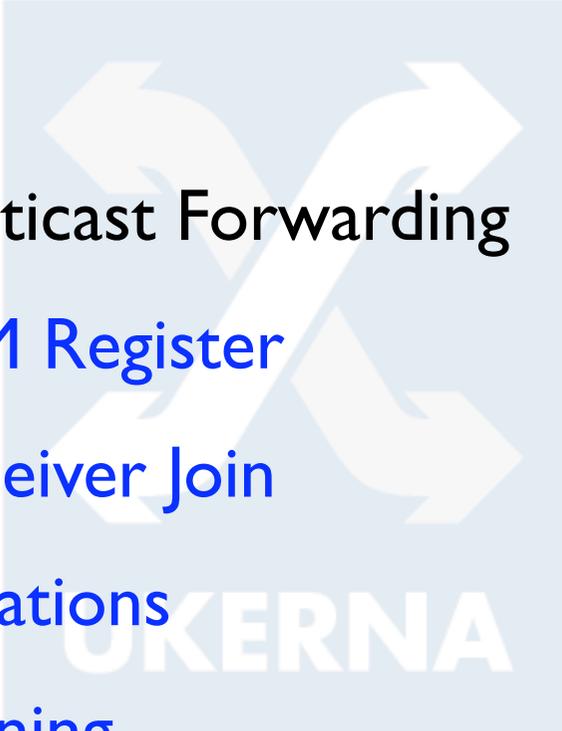
PIM Routing

Gorry Fairhurst
University of Aberdeen
gorry@erg.abdn.ac.uk

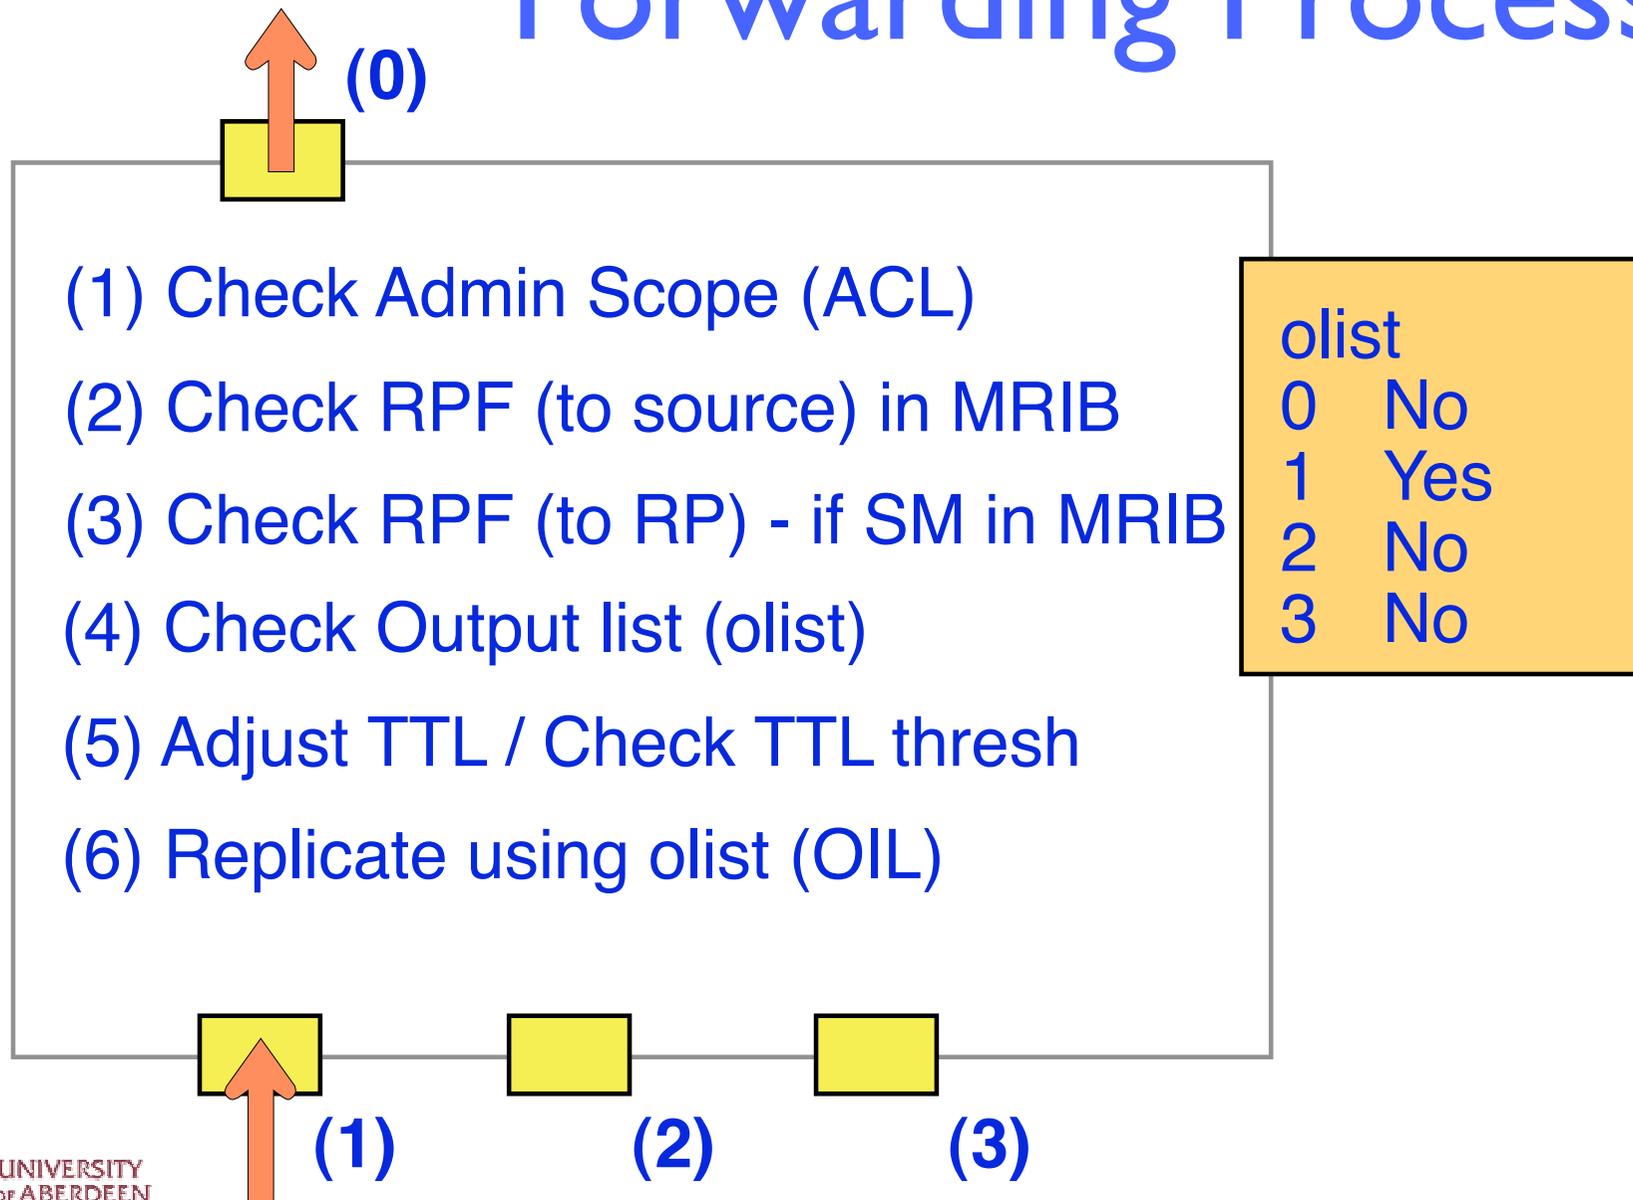
IP Multicast Workshop at Networkshop 2006

Slides (as flash, QT) at:

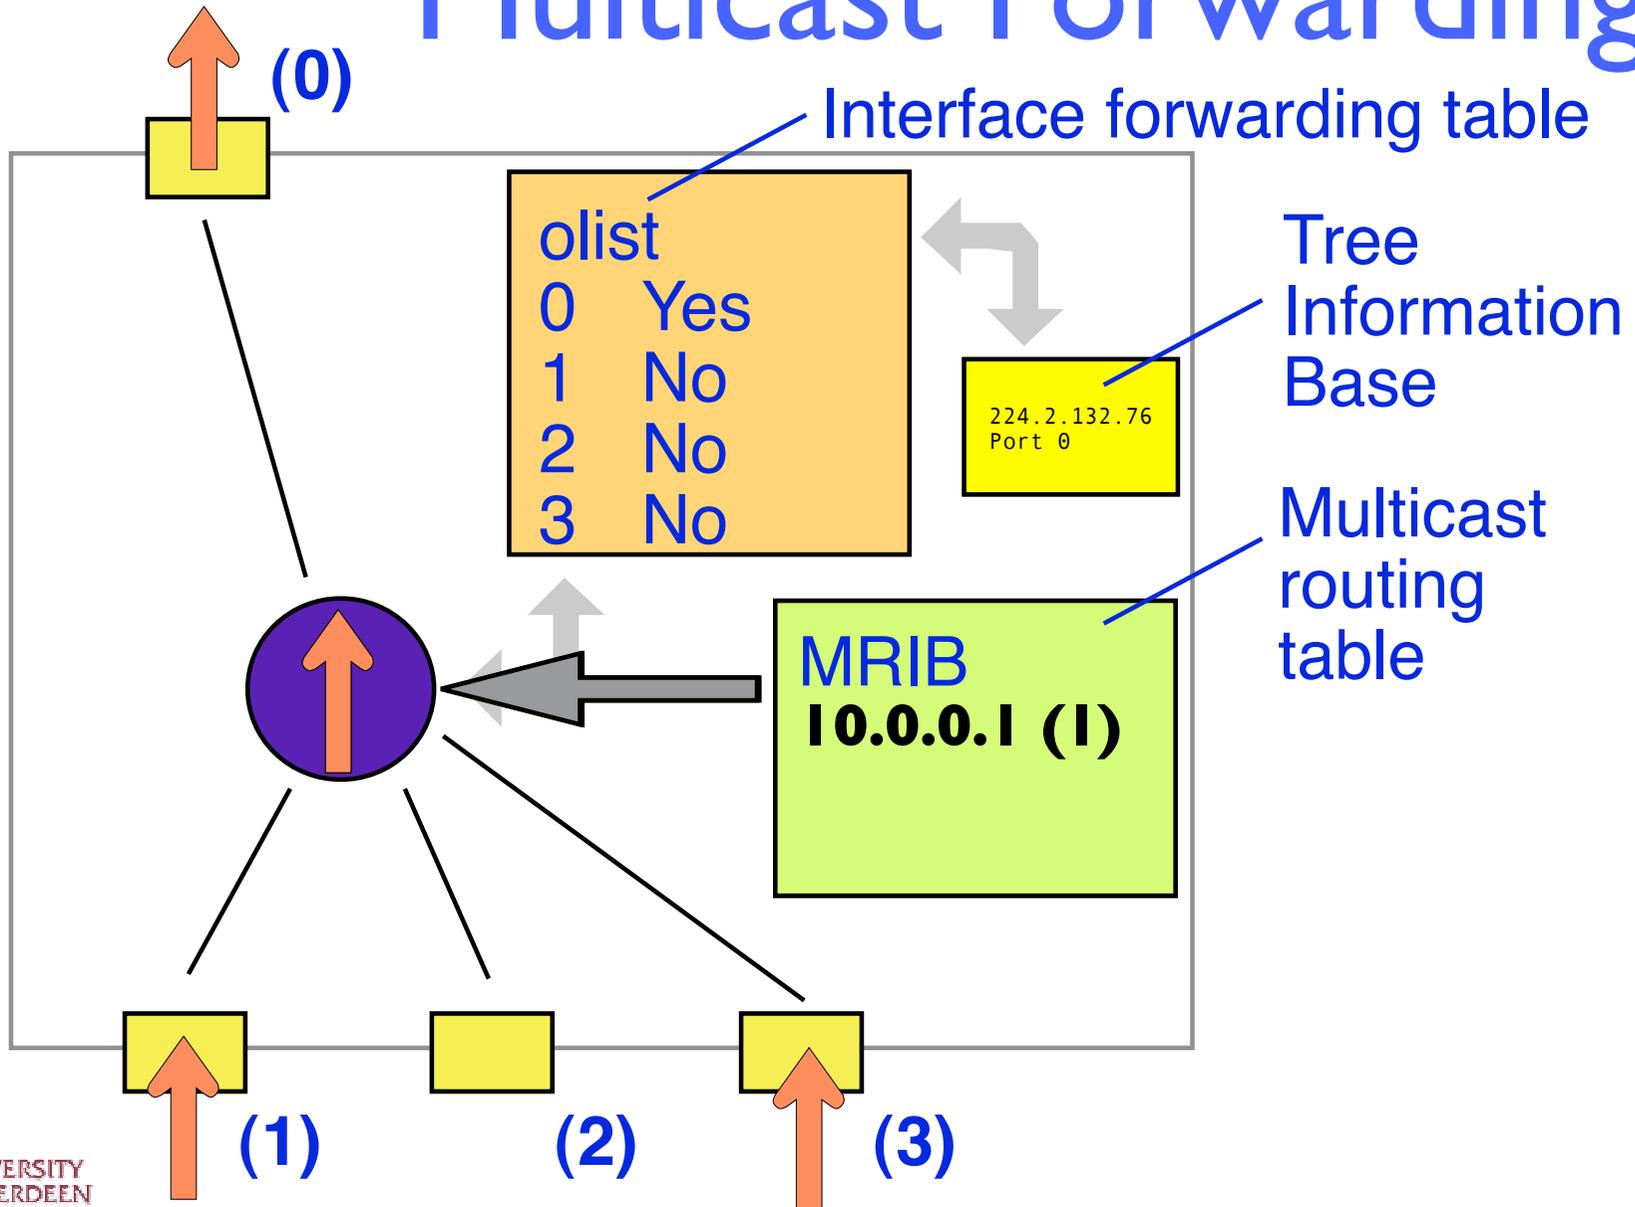
<http://www.erg.abdn.ac.uk/users/gorry/ipmulticast/docs/>

- 
- 
- Part 1: Multicast Forwarding
 - Part 2: ASM Register
 - Part 3: Receiver Join
 - Part 4: Variations
 - Part 5: Pruning

Forwarding Process

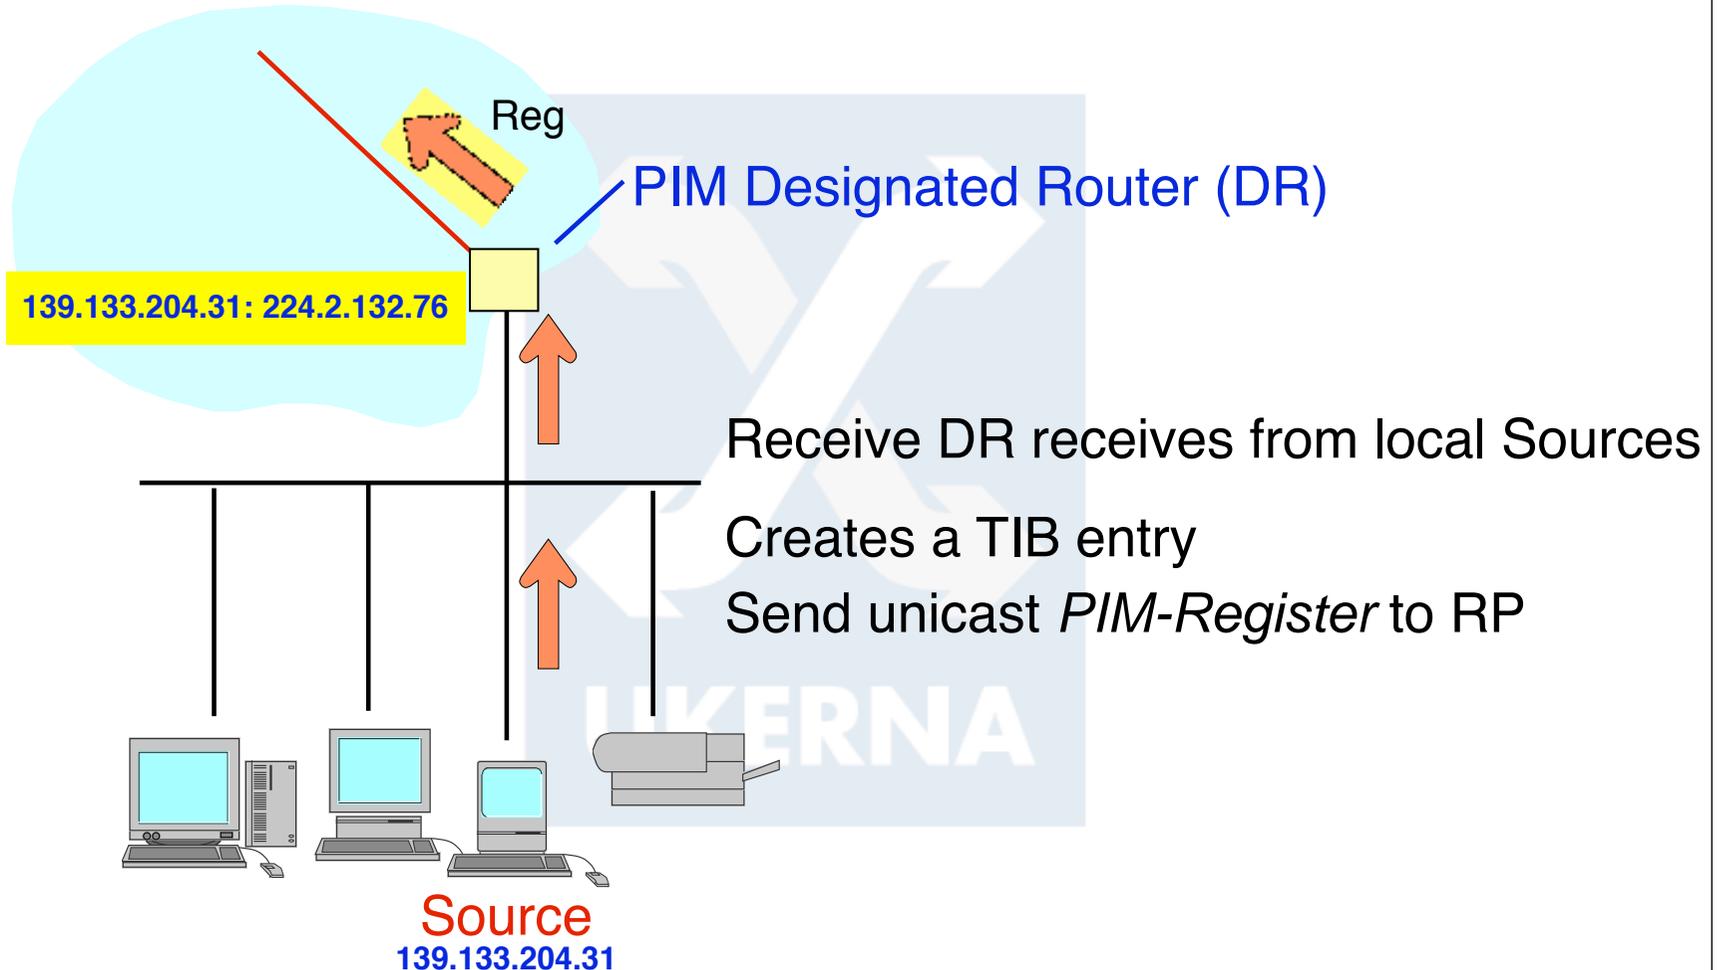


Multicast Forwarding



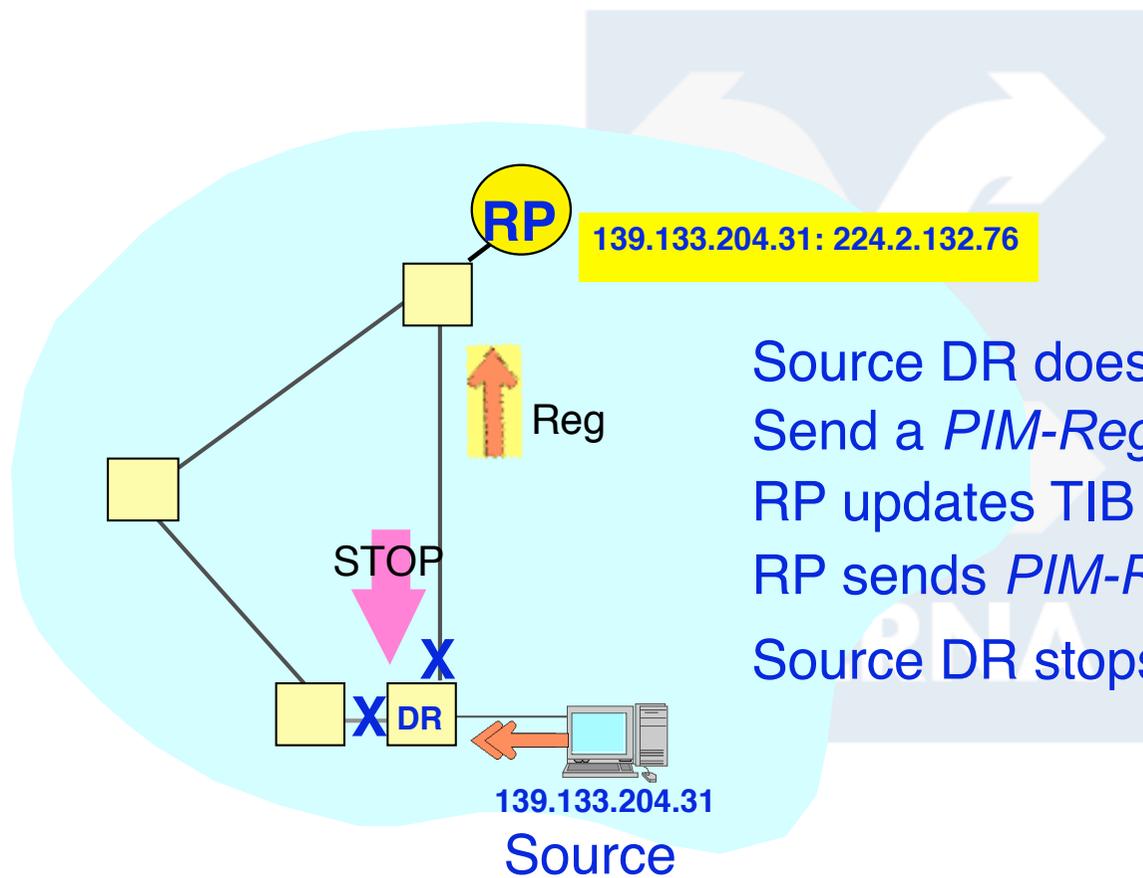
- 
- 
- Part 1: Multicast Forwarding
 - Part 2: ASM Register
 - Part 3: Receiver Join
 - Part 4: Variations
 - Part 5: Pruning

Sending to the LAN DR



PIM Designated Router (DR) may be IGMP Querier

Telling the RP about a Source

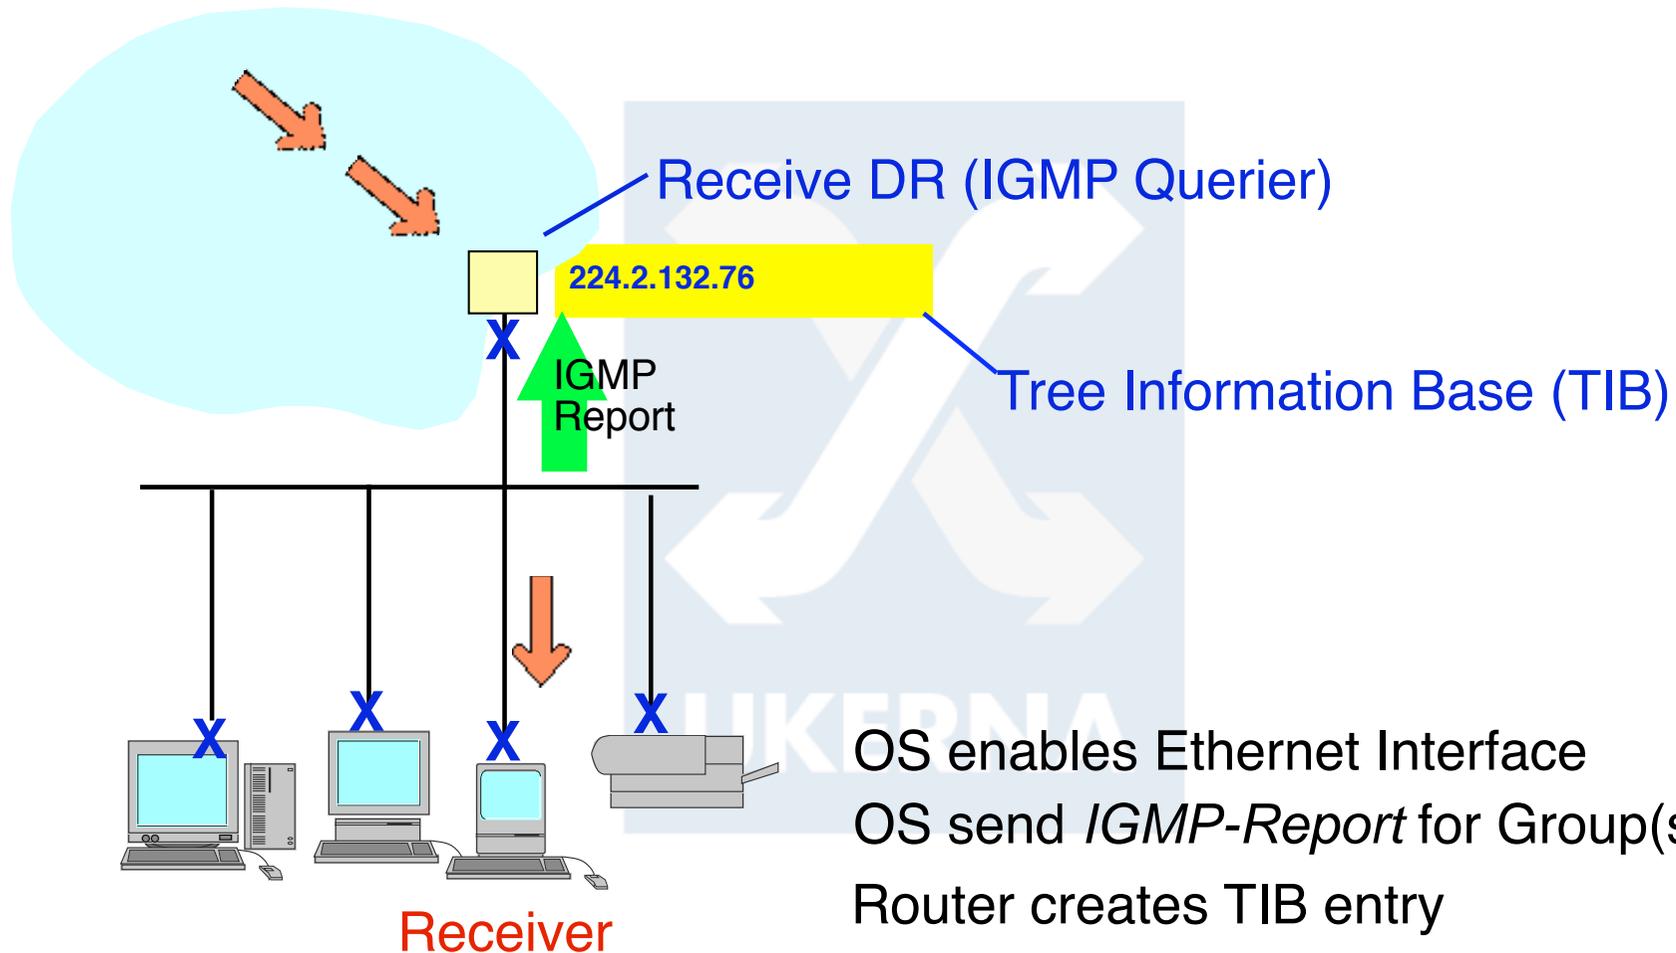


- Source DR does not forward flows
- Send a *PIM-Register* to the RP
- RP updates TIB
- RP sends *PIM-Register-STOP*
- Source DR stops registering (for a while)

DR waits for a *PIM-Join*

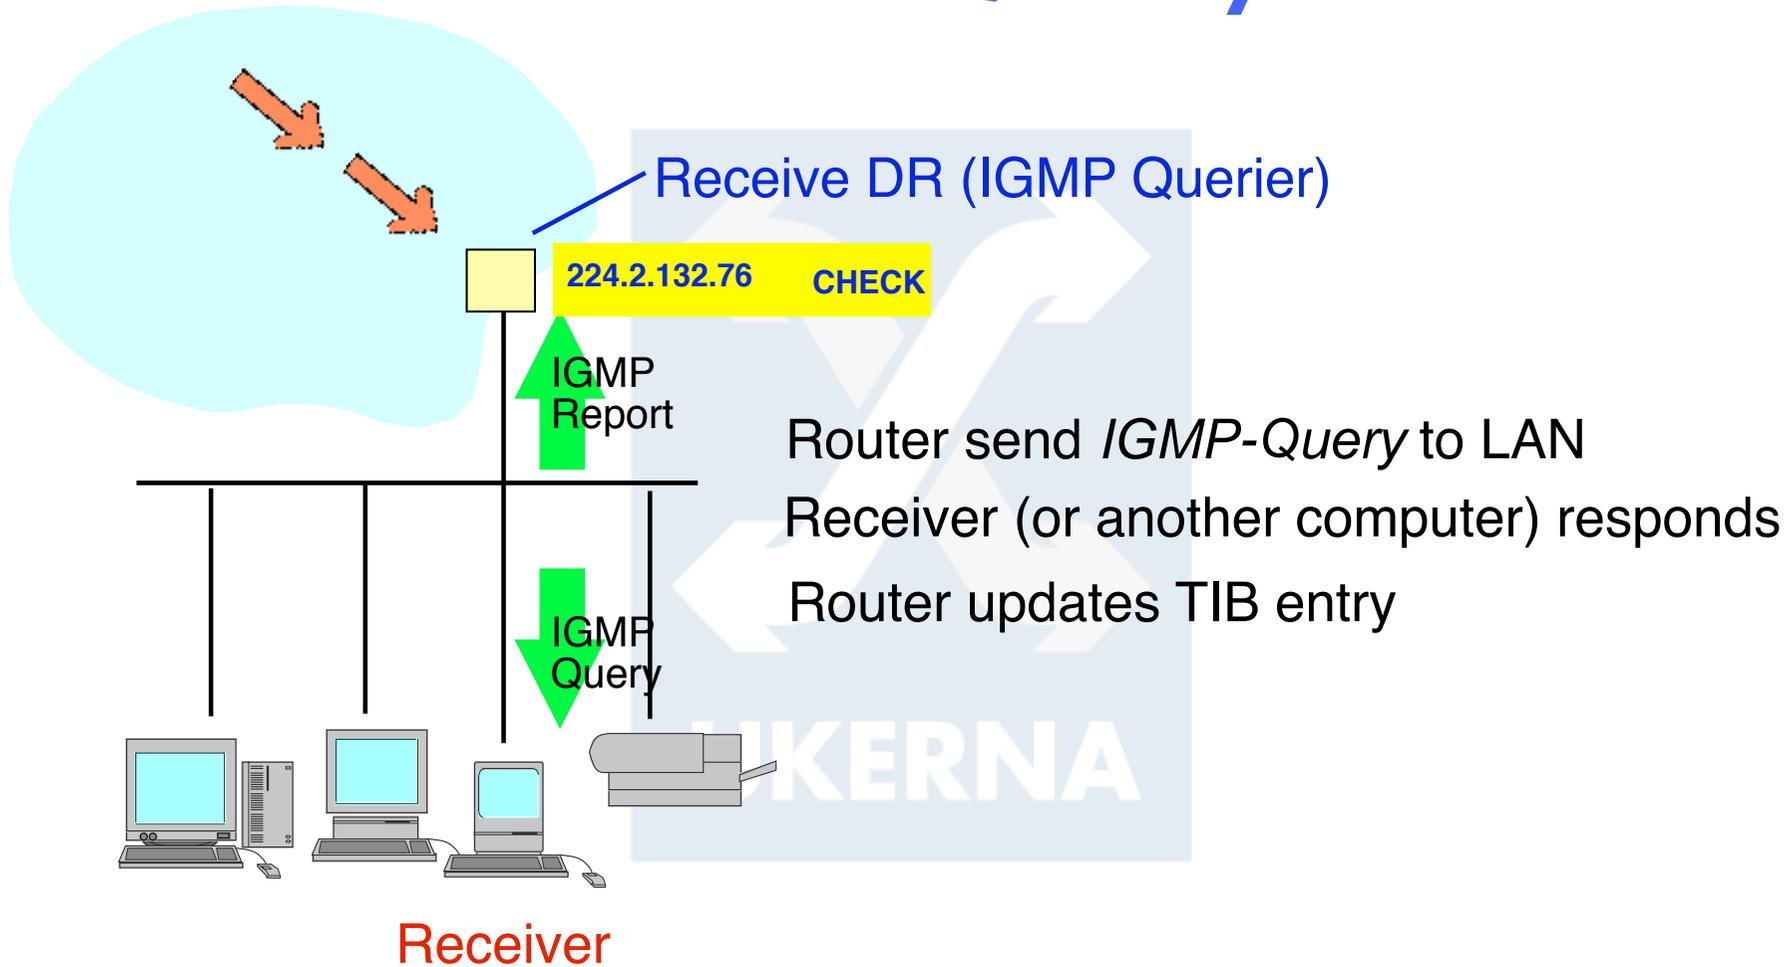
- 
- 
- Part 1: Multicast Forwarding
 - Part 2: ASM Register
 - Part 3: Receiver Join
 - Part 4: Variations
 - Part 5: Pruning

OS Sends IGMP Report



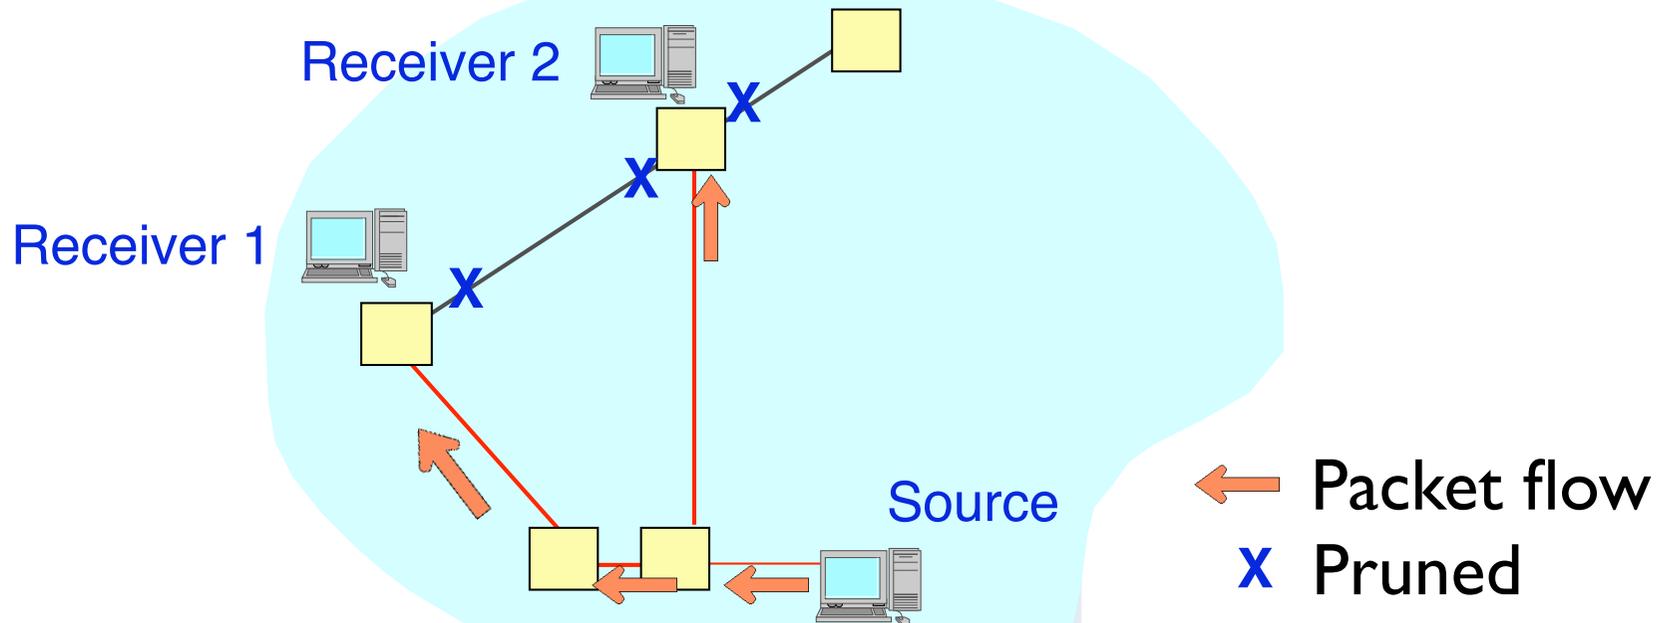
- OS enables Ethernet Interface
- OS send *IGMP-Report* for Group(s)
- Router creates TIB entry
- Router enables LAN interface
- Router forwards packets to Receiver

IGMP Query



Router continues to forward packets to Receiver

Joining / Pruning the Tree



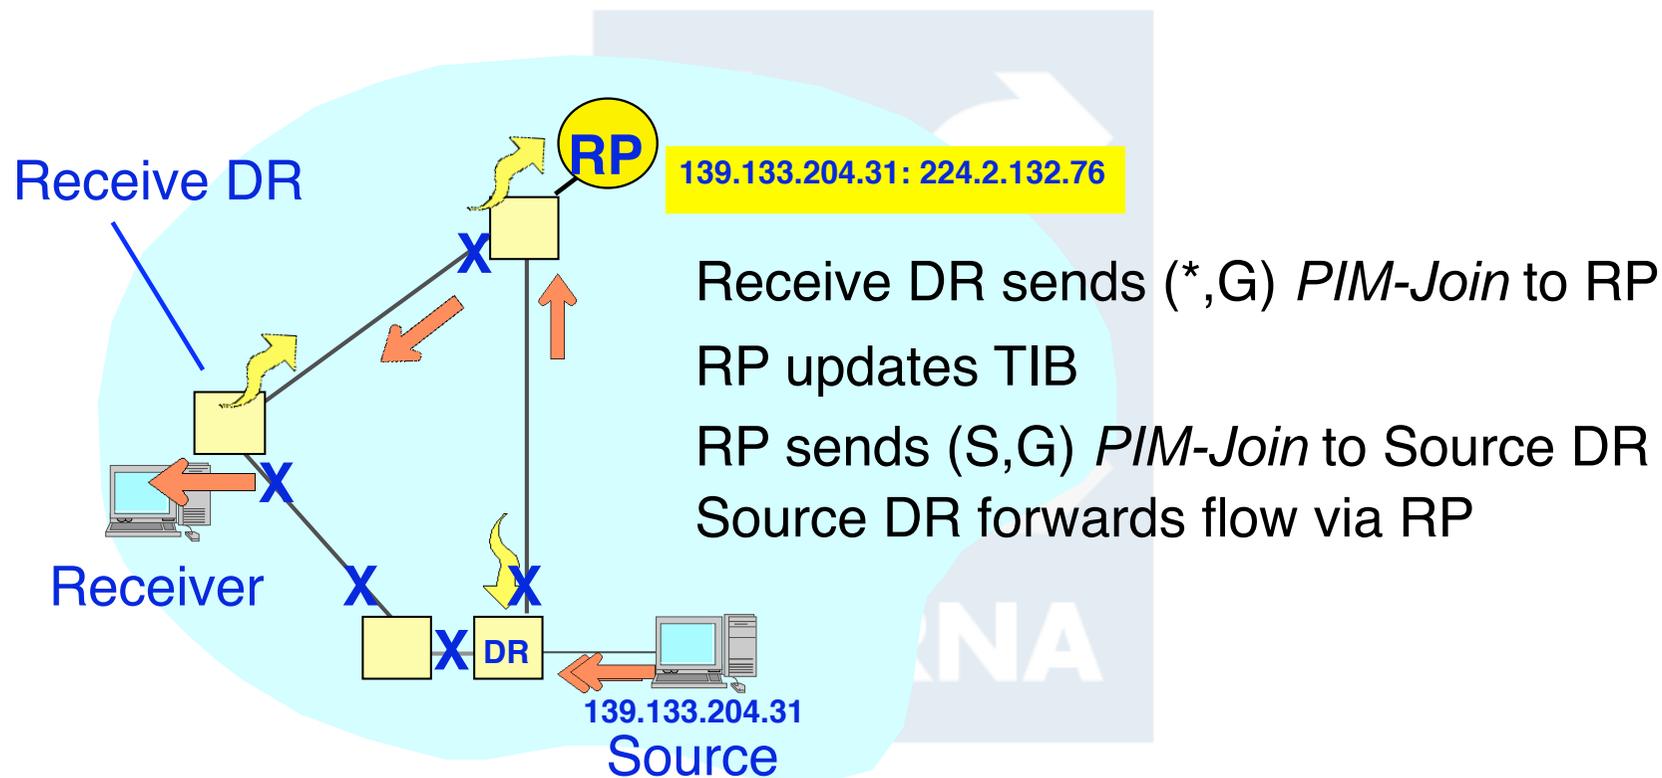
Joining

Routers do not forward flows until a *Join* is received
Routers send *Join* if they wish to receive a flow

Pruning

Routers *Prune* flows when there are no *Joins*

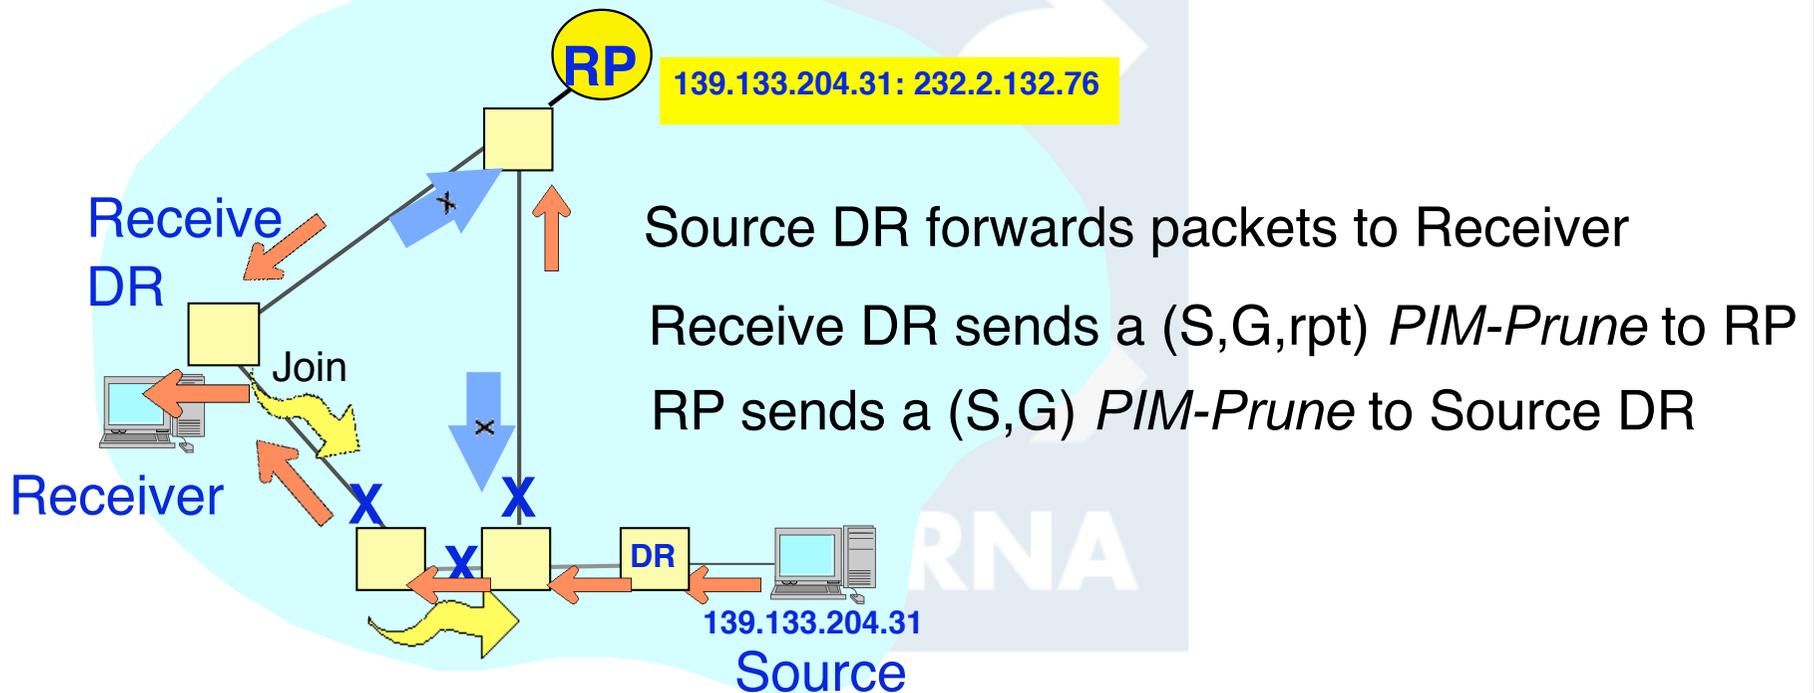
Receiving via the RP



Receiving from Source

Receive DR discovers Source address

Receive DR sends (S,G) *PIM-Join* to Source



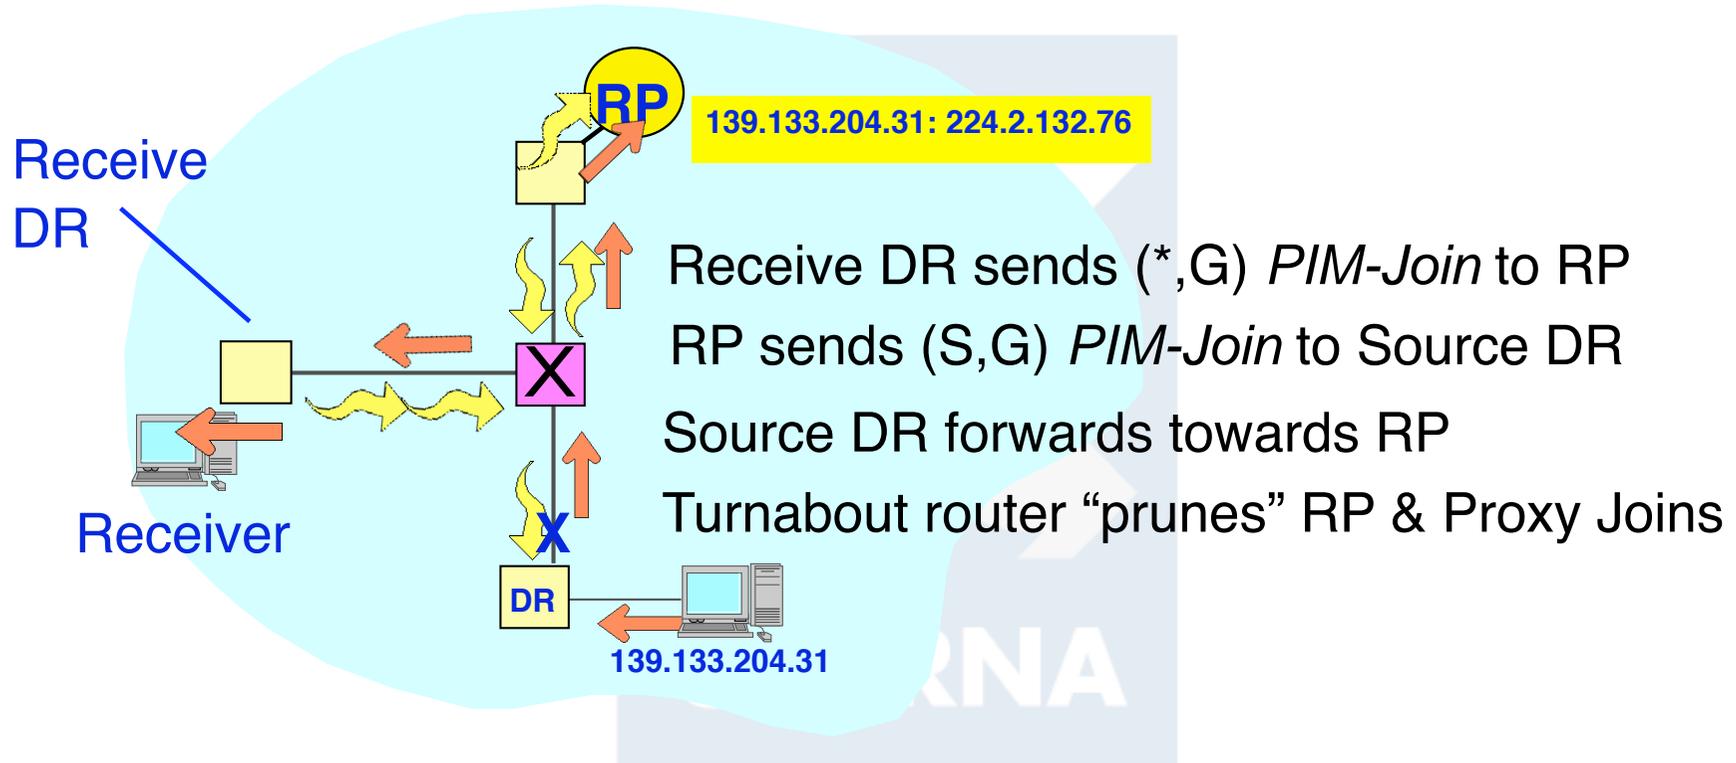
Source DR forwards packets to Receiver

Receive DR sends a (S,G,rpt) *PIM-Prune* to RP

RP sends a (S,G) *PIM-Prune* to Source DR

Packets forwarded on shortest path (SPT)

Turnabout Router



Traffic never flows up the RPT
to then flow back down the same SPT!

- 
- Part 1: Multicast Forwarding
 - Part 2: ASM Register
 - Part 3: Receiver Join
 - Part 4: Variations
 - Part 5: Pruning

PIM-SSM

- when there are few sources

Bi-Dir PIM

- when there are many sources

Sending SSM from the DR

Rules change for 232/8 (SSM!)

Source DR role changes

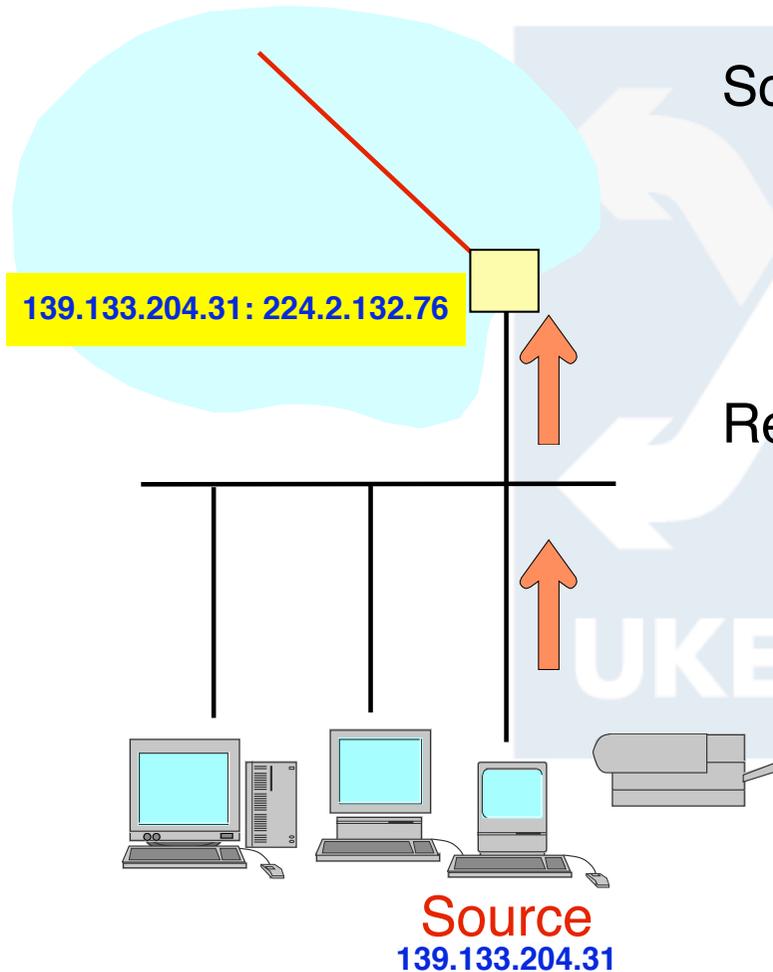
No register sent from DR

No (*,G) PIM-Join for 232/8

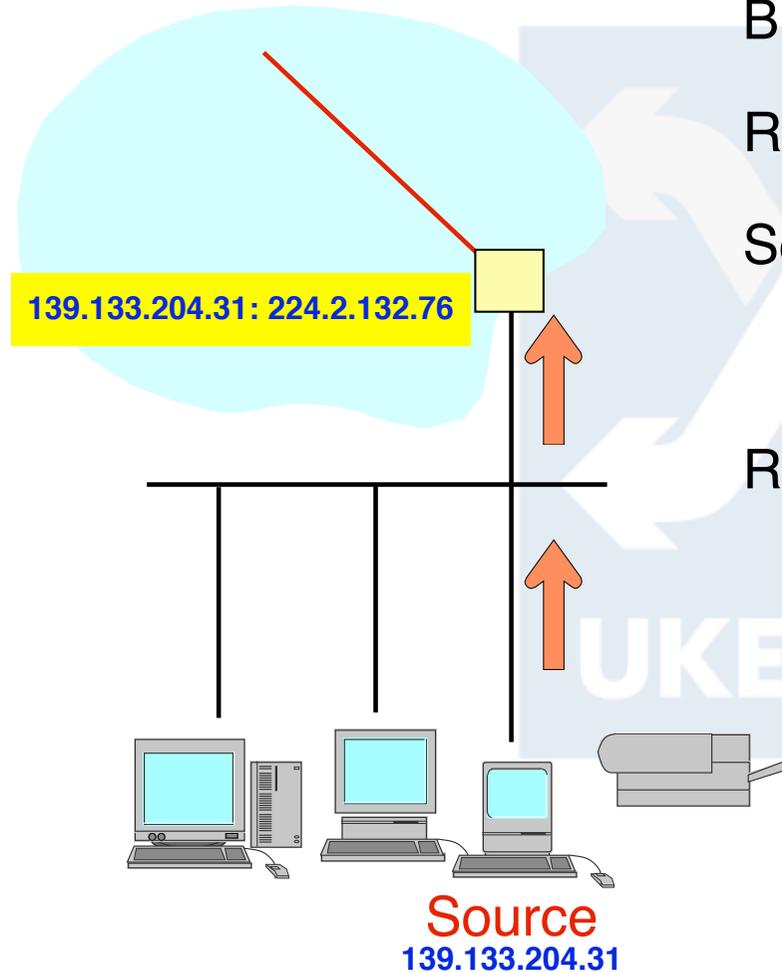
Receive DR Join changes

(S,G) Join directly to source (IGMPv3)

RPs not to handle these groups



Sending BiDir from the DR



BiDir negotiated in *PIM Hello*

Rules change

Source DR role changes

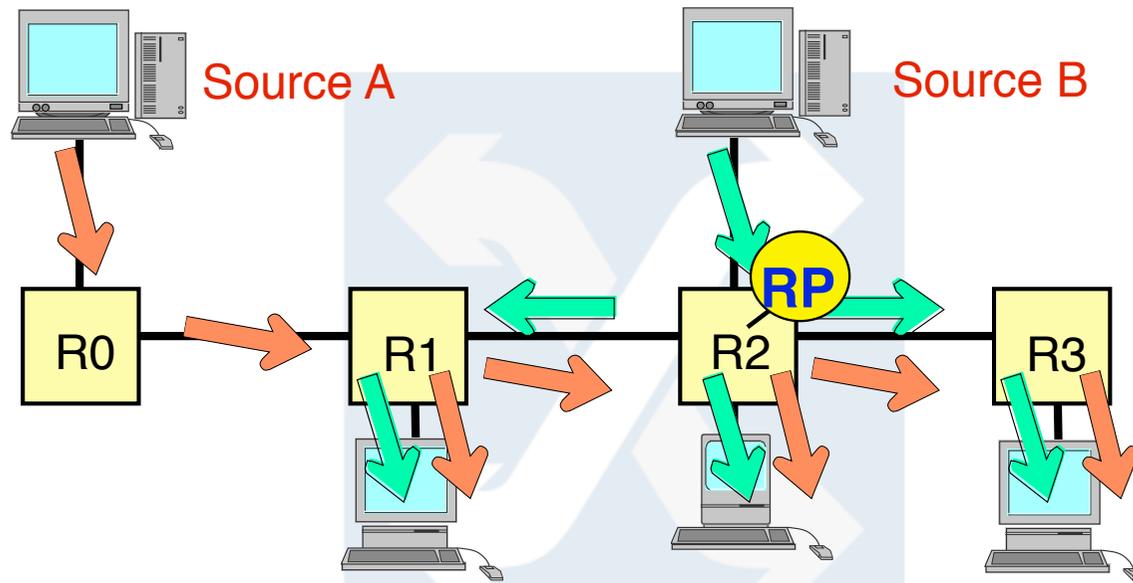
No *Register* sent from DR

Receive DR forwarding changes

Send natively to RP

Other routers can see the packets

BiDir Routing



No per-source state

Good for many sources (c.f. SSM)

No problems with bursty sources (as in SSM)

Network changes for BiDir

RP much simpler (No tunnels)

No problems with bursty sources

Can be any multicast router

Problems

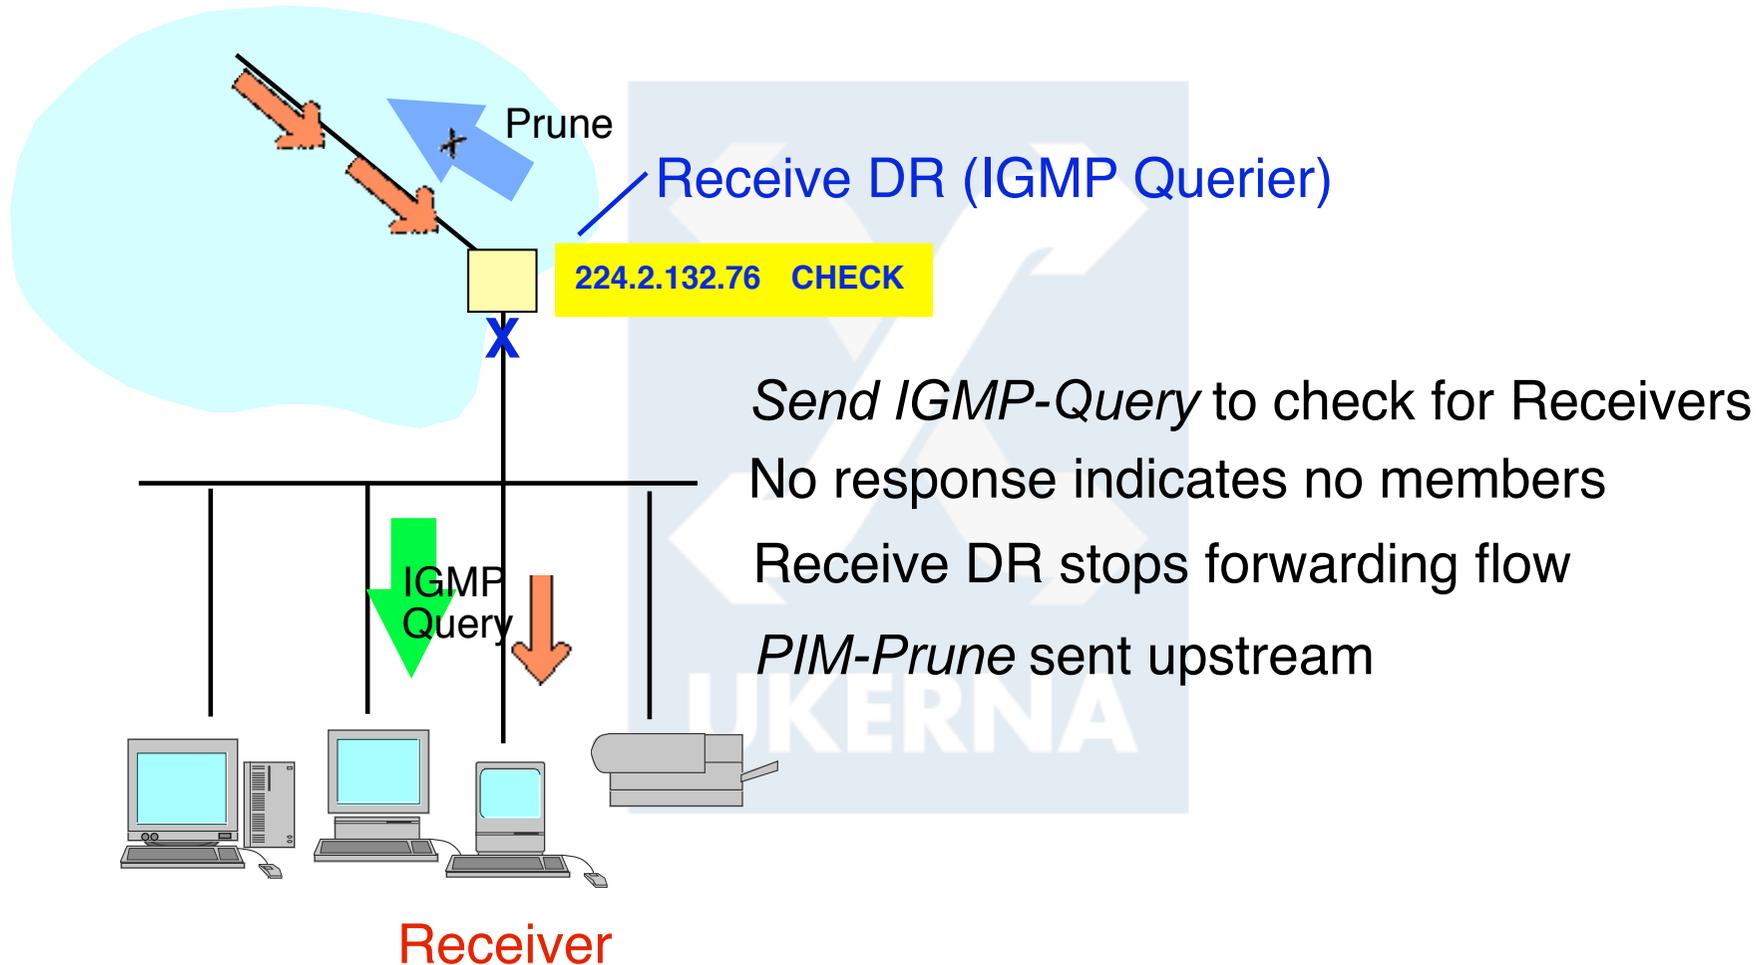
Does not directly interoperate with SM

All traffic can flow to the RP

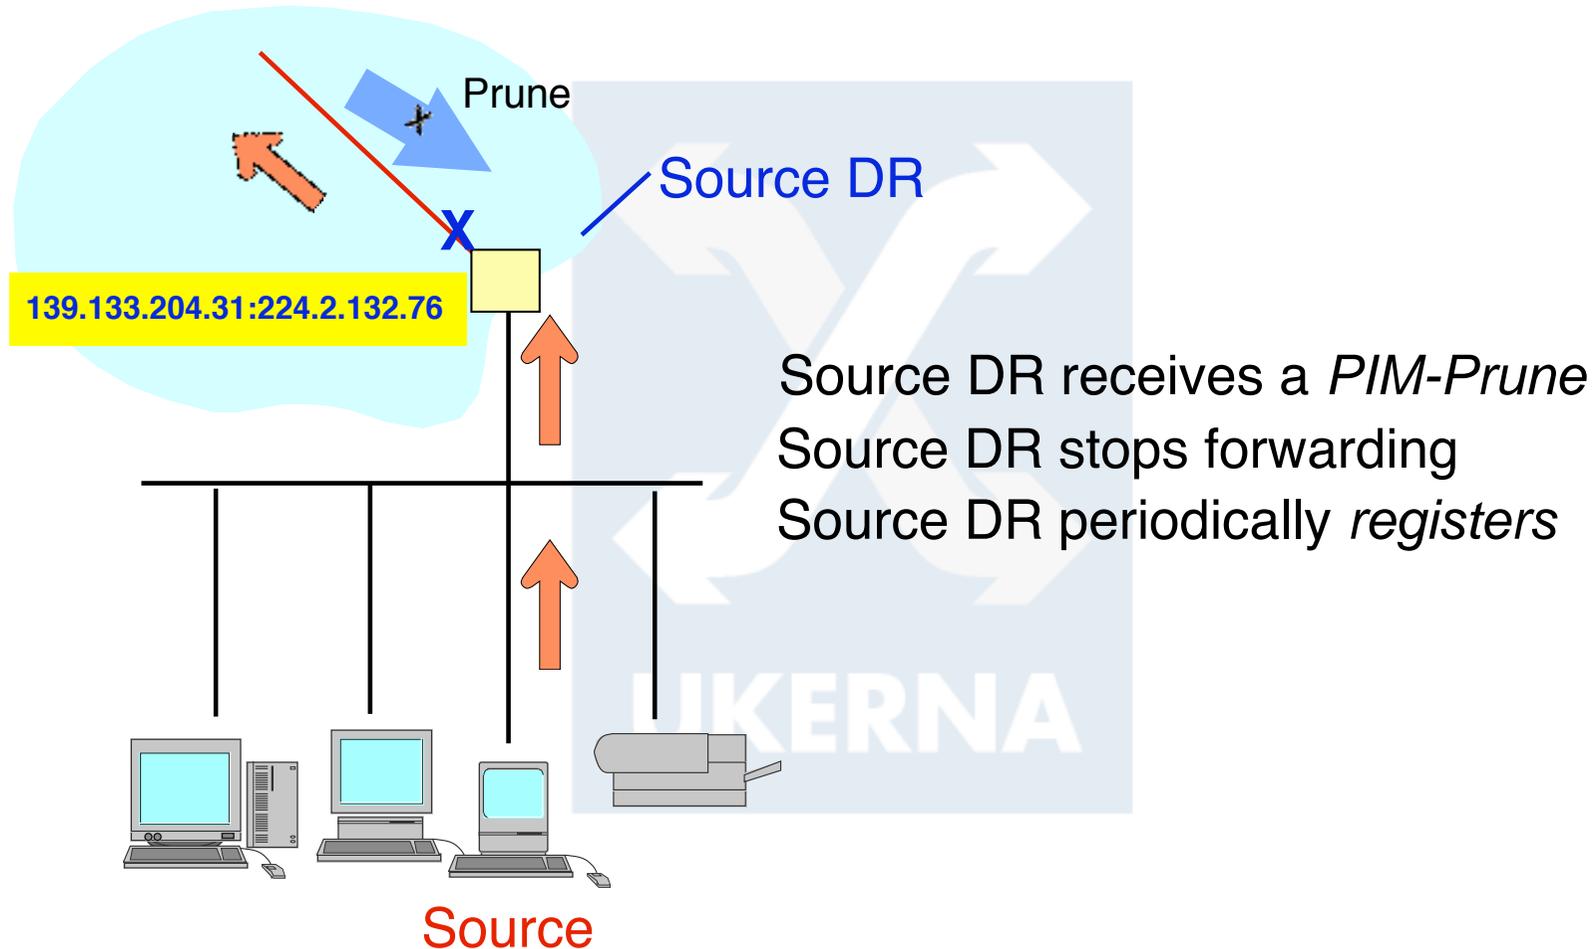
RP placement important

- 
- 
- Part 1: Multicast Forwarding
 - Part 2: ASM Register
 - Part 3: Receiver Join
 - Part 4: Variations
 - Part 5: Pruning

Leaving the Group



Stopping a Source



So Finally....

Three key things

IGMP tells a Receive Router which flows to Receive
RP's tell Receive Routers about Sources (via packets)
Join/Prune used to build trees

Multicast isn't hard

It is very different to unicast
Need to work with new "tools"

Question & Answers



Slides (as flash, QT) at:

<http://www.erg.abdn.ac.uk/users/gorry/ipmulticast/docs/>

PIM Message Types

PIM Messages

- 0: HELLO
- 1: REGISTER (includes tunnelled packet)
- 2: REGISTER STOP
- 3: JOIN/PRUNE
- 4: Bootstrap Router (BSR) - SM Only
- 5: ASSERT
- 6: GRAFT - DM Only
- 7: GRAFT-ACK - DM Only
- 8: CANDIDATE RP-ADVERTISEMENT (with BSR)
- 10: BiDir

Common PIM Flags (displayed by "mroute")

- C Connected - One or members wish to receive flow for Group G
- F Register - This DR has a local Source for Group G
- J Join - Router has sent a Join message for this flow
- P Pruned - Router does not need to forward flow for Group G
- T Tree - Router is receiving directly from the Source (S,G)
- X Turnabout - Router is ending proxy-joins to the Source (S,G)