H.O. #6

Winter 98-99

- Generic Attribute Registration Protocol (GARP)
- GARP Multicast Registration Protocol (GMRP)
- Recommended Reading:
  - IEEE 802.1D Chapter 10
  - IEEE 802.1D Chapter 12, Sections 12.1 through 12.9

Issues on Multicast Handling

- A ***Broadcast*** packet is sent from one station to ***all*** stations in the LAN; everybody gets it.
- A ***Multicast*** packet is sent from one station to a ***group*** of stations, not necessarily all stations.
- "Basic" bridges treat multicast packets the same as broadcast packets: the packet is sent out in all forwarding ports, except the one where it was received.
- Problem: unnecessary transmissions to LAN segments where nobody is interested in the packet.
- Issue: bridges need to know which segments contain stations interested in each specific multicast.

Example

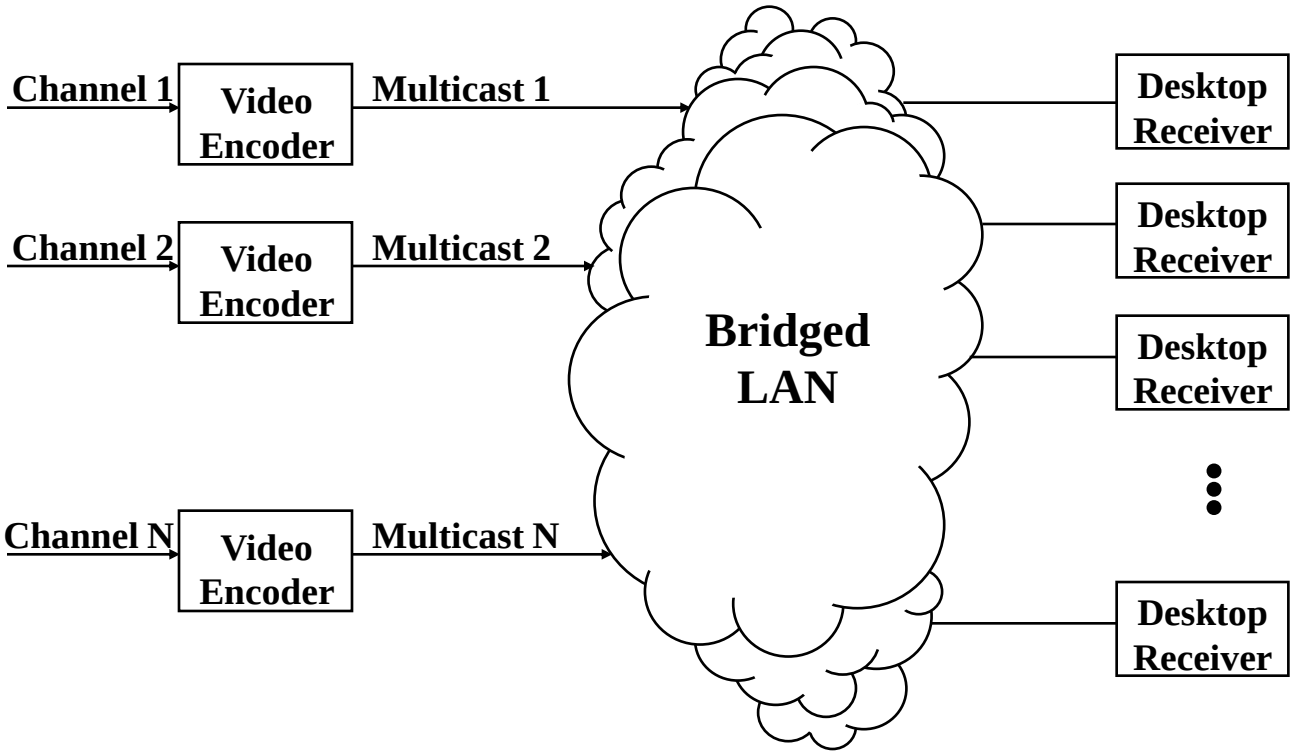

- Each encoder sends its stream to a multicast address
- Decoders "tune in" a stream by receiving from that multicast address
- If multicasts go everywhere, the network runs out of bandwidth quickly (imagine 10 Mb/s Ethernets and 4 Mb/s MPEG-2 streams)

Proprietary Solutions

- Main issue: the bridges need to know who is interested in each multicast - but until now, this information only exists at the next layer, in a protocol-specific fashion.
- 3Com: bridges know about the IP multicast protocol, and “snoop” it.
- Cisco: bridges talk to their routers and learn from the routers the location of the groups, using a proprietary protocol.
- Need to define something ***standard*** at layer 2.

GMRP
GARP Multicast Registration Protocol

- A mechanism that allows bridges and end stations:
  - to dynamically register (and subsequently de-register) group membership information with the MAC bridges attached to the same LAN segment
  - to disseminate that information across all bridges in the bridged LAN that support Extended Filtering Services
- Operation of GMRP relies on the services provided by the Generic Attribute Registration Protocol (GARP)

Type of Information (1)

- Information registered, de-registered and disseminated via GMRP:
  - a) group membership information
    - indicates that one or more GMRP participant that are members of a particular group (or groups) exist
    - carries the group MAC address(es) associated with the group(s)

  $\Rightarrow$ results in the creation or updating of *group registration entries* in the filtering database to indicate the port(s) on which members of the group(s) have been registered

Type of Information (2)

- B) group service requirement information
  - indicates that one or more GMRP participants require *Forward All Groups* or *Forward Unregistered Groups* to be the default group filtering behavior

Open Host Group Concept

- Any GMRP participants that wish to receive frames transmitted to a particular group or groups register their intention to do so by requesting membership to the group(s) concerned
- Any MAC service user that wishes to send frames to a particular group can do so *from any point of attachment* in the bridged LAN
- MAC service users that are sources of MAC frames for the group do not have to register as members of the group themselves unless they also wish to receive frames sent to the group by other sources

Result of Group Membership Information Registration and Propagation

- Frames sent to a particular group can be received on all LAN segments to which registered GMRP participants are attached
- Bridges filter frames on ports which have not had group registration entries created by GMRP
  - frames are not transmitted on LAN segments which neither have registered GMRP participants, nor are in the path through the *active topology* between the sources of the frames and the registered members

Example of an Active Topology

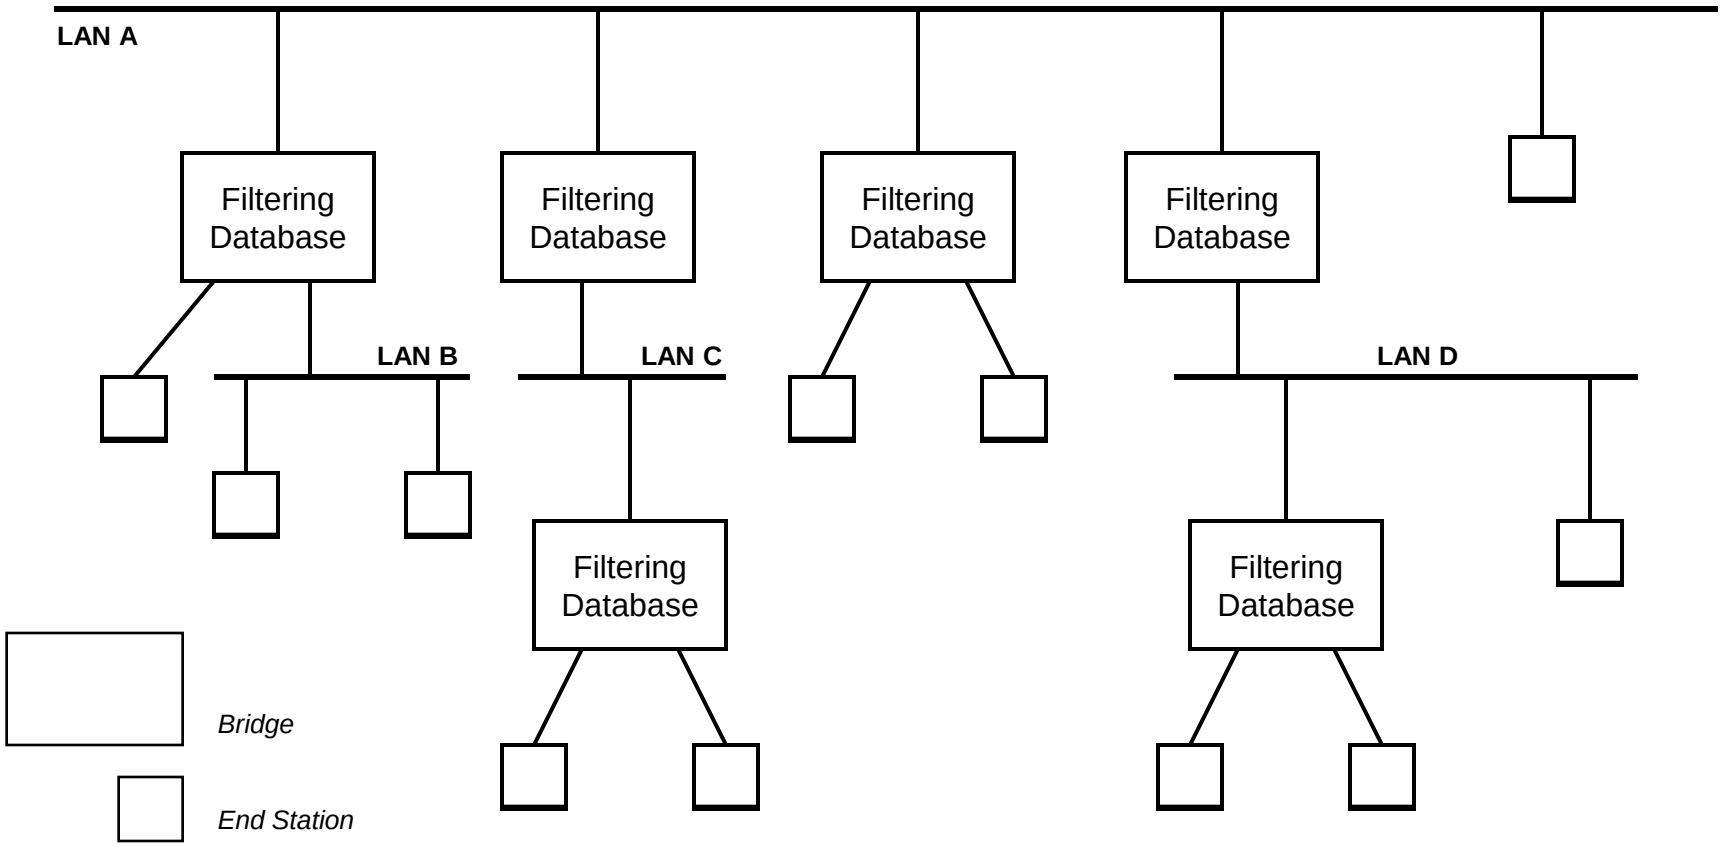

Active Topology with Group Members

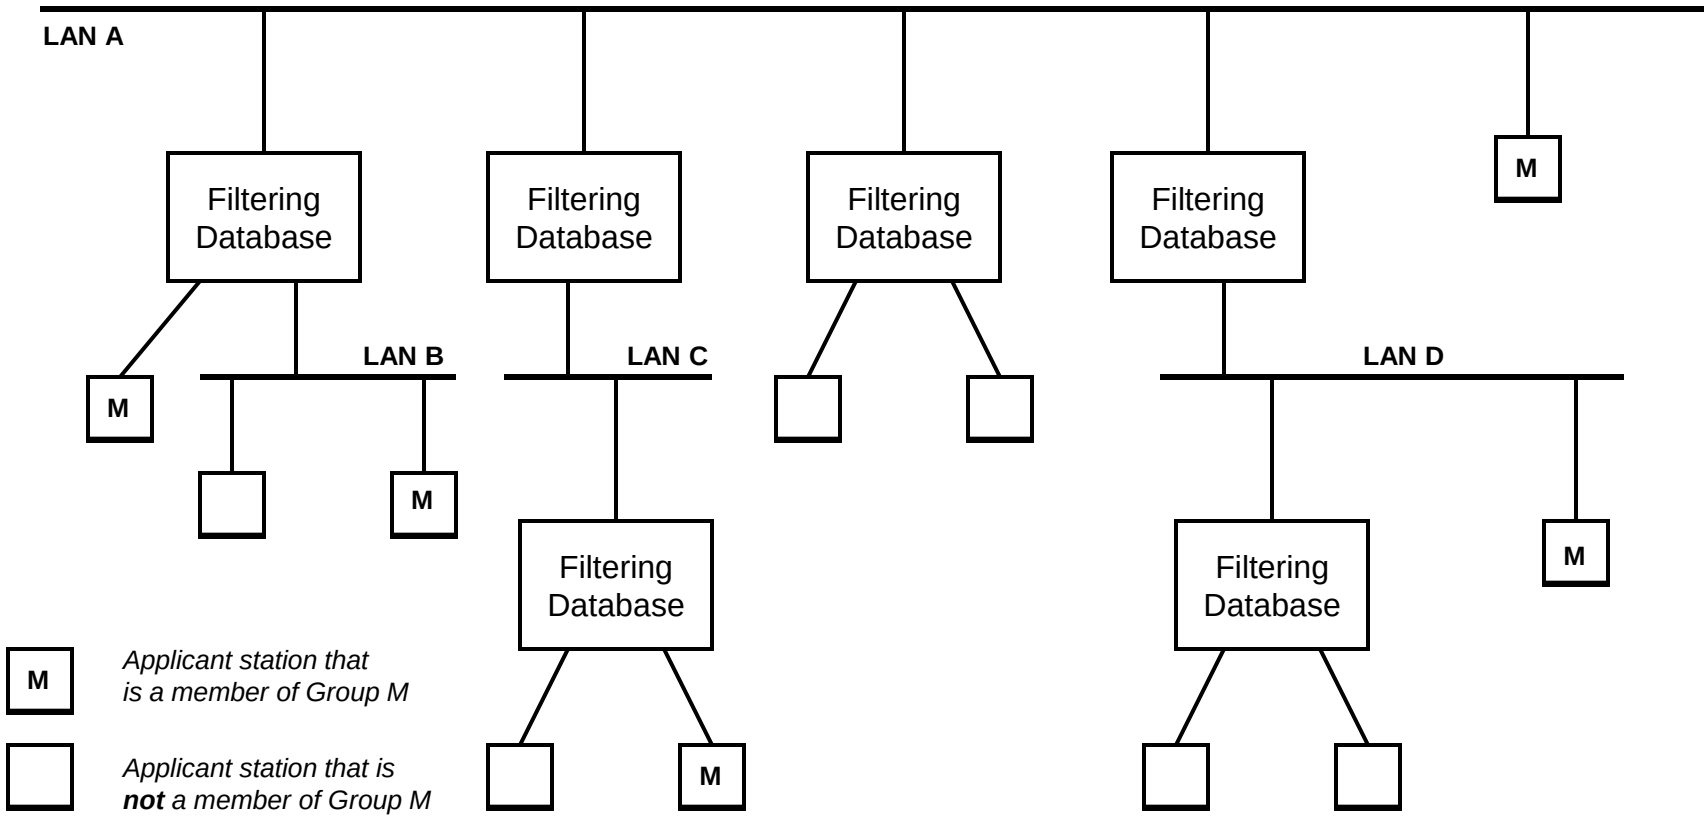

Example 1

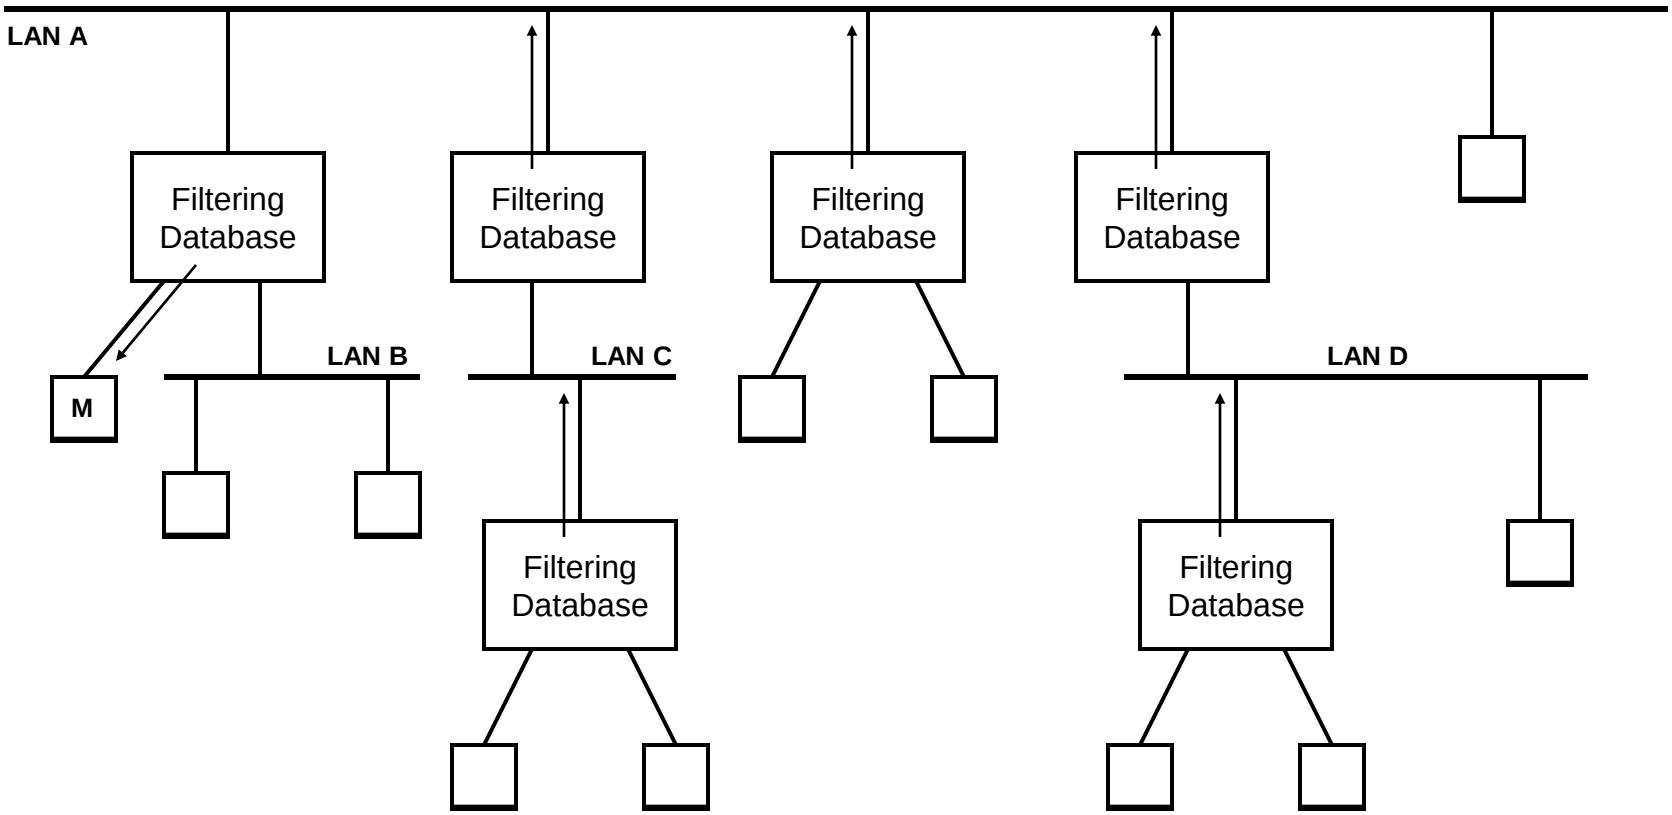
LAN A
Filtering Database
Filtering Database
Filtering Database
Filtering Database
LAN B
LAN C
LAN D
M
Filtering Database
Filtering Database

Example 2

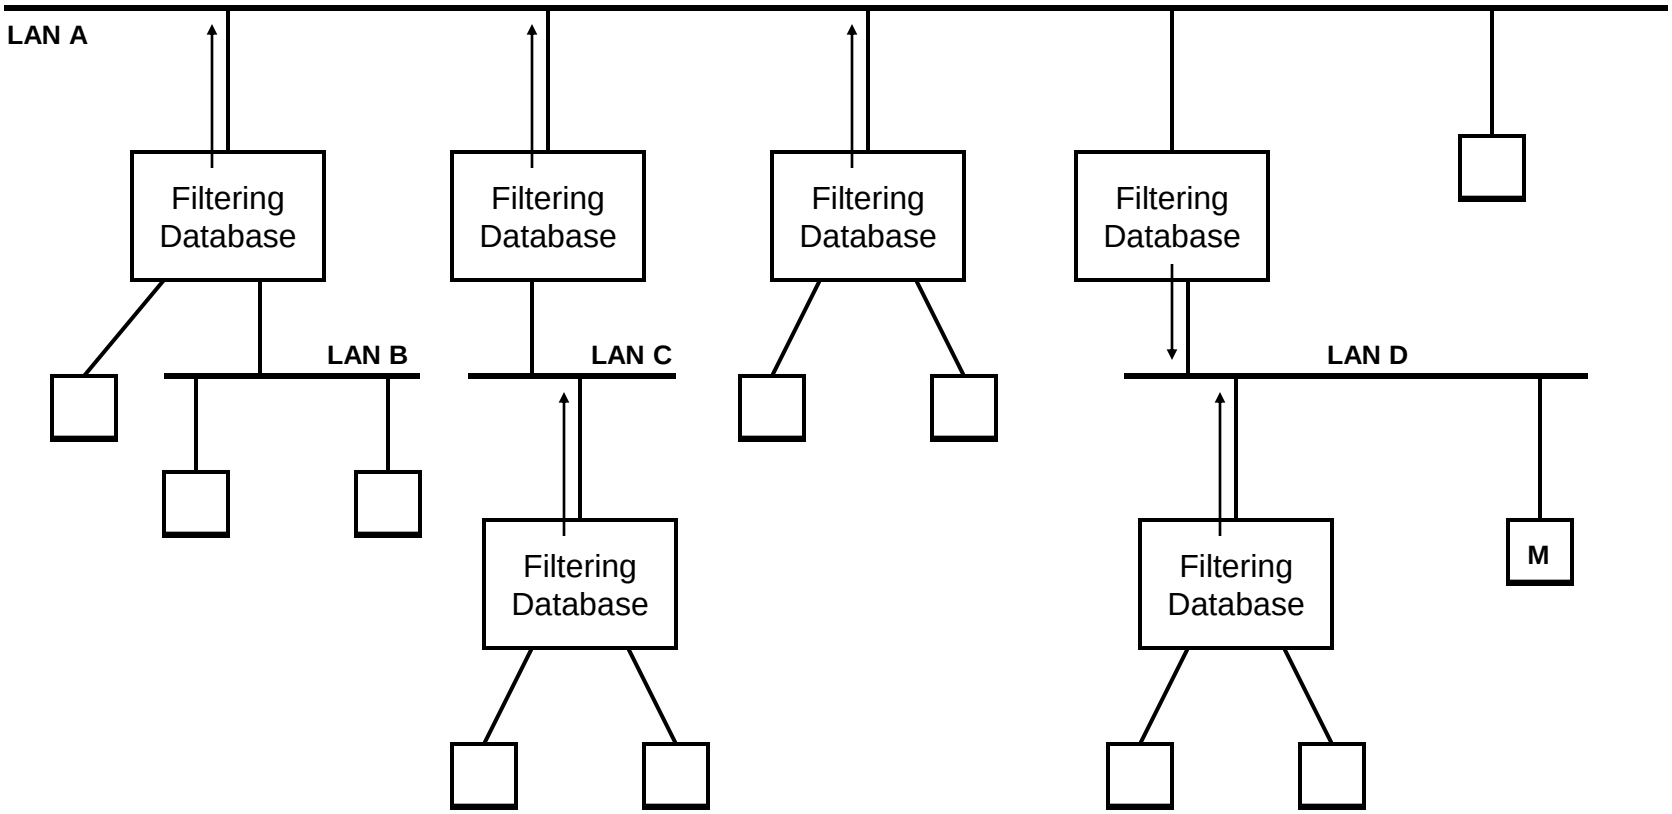
LAN A
Filtering Database
Filtering Database
Filtering Database
Filtering Database
LAN B
LAN C
LAN D
Filtering Database
Filtering Database
M

Example 3

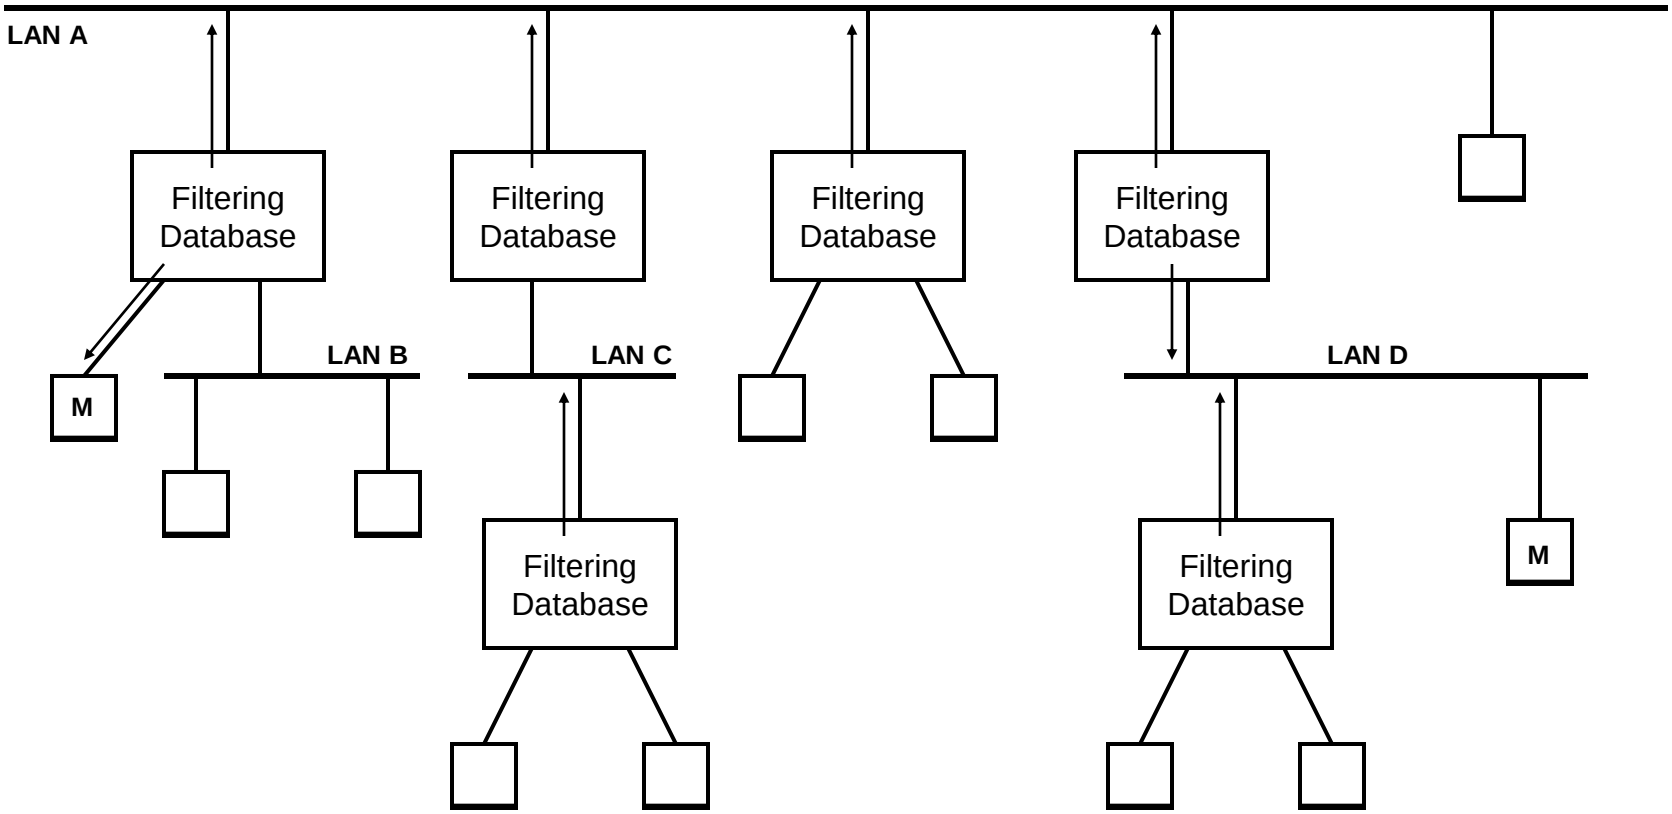
LAN A
Filtering Database
Filtering Database
Filtering Database
Filtering Database
LAN B
LAN C
LAN D
M
Filtering Database
Filtering Database
M

Group Registration Entries in Filtering Databases Resulting in a Directed Graph

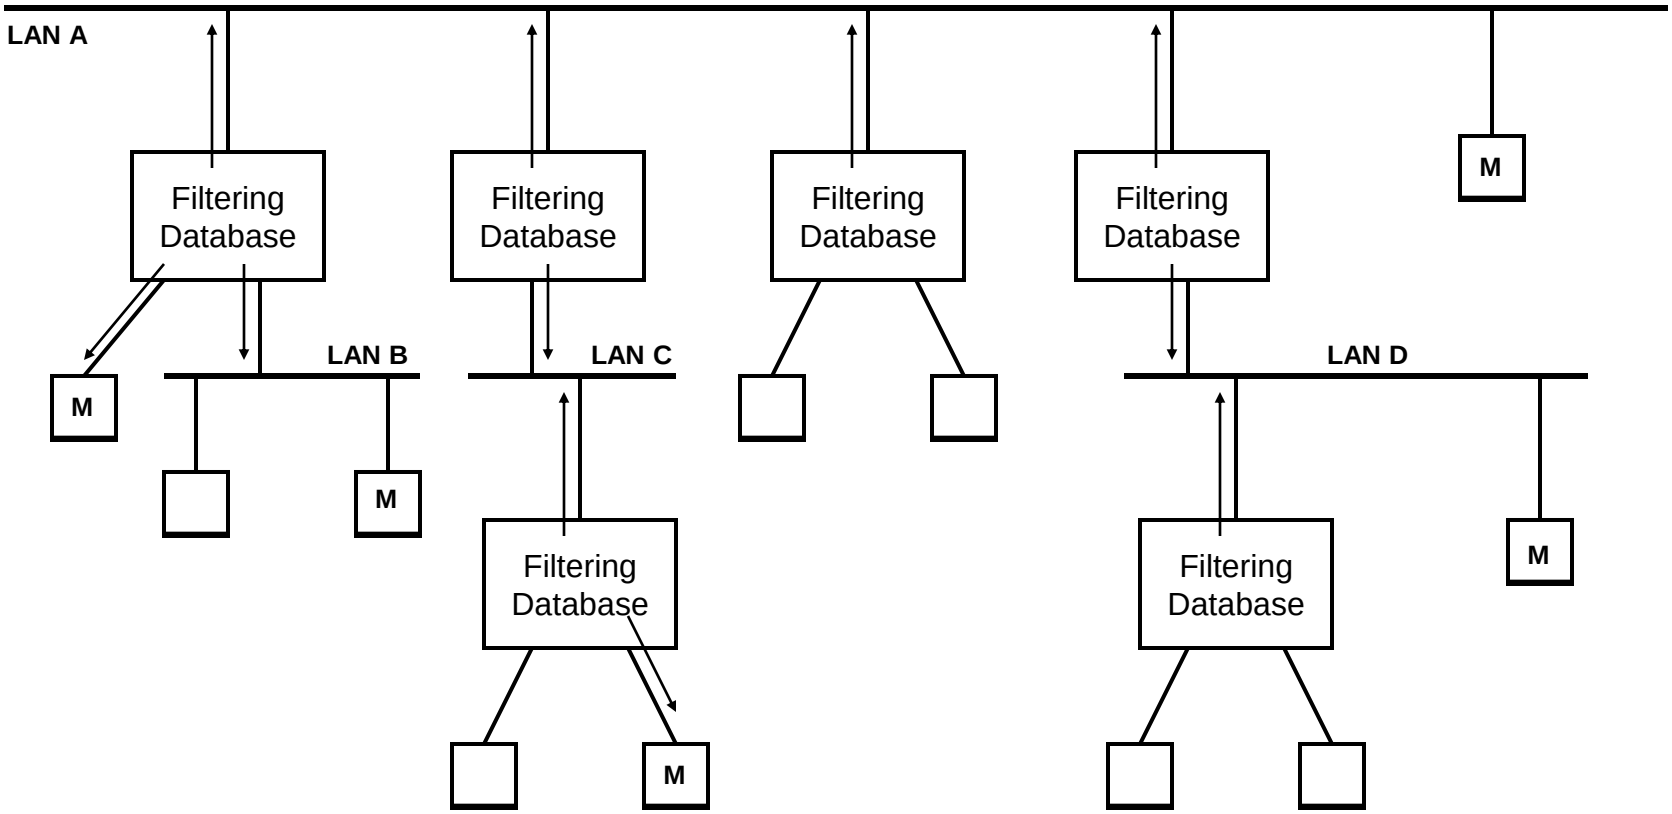

Source Pruning

- End stations use Group Membership information registered via GMRP to keep track of the set of Groups for which active members exist
- End stations may suppress the transmission of frames for which there are no valid recipients
- Avoids unnecessary flooding of traffic on the LAN if there are no members that wish to receive such traffic

Use of Forward All Groups

- To ensure that regions of the Bridged LAN that contain legacy devices can receive all multicast frames
  - Can use static configuration to explicitly disallow certain multicast traffic
- To allow successful operation of devices that require promiscuous reception
  - Routers
  - Network monitors

Use of Forward Unregistered Groups

- Group addresses which do not have dynamic filtering database entries are forwarded
- Group addresses which have dynamic filtering entries are filtered or forwarded based on the dynamic entry
- Useful in circumstances where GMRP-aware devices distinguish between legacy multicast addresses for which they do not register and “new” multicast addresses for which they do register
- Must insure GMRP-aware end stations do not register for legacy multicast addresses

Use of Filter Unregistered Groups

- Group addresses which do not have dynamic filtering database entries are filtered
- Group addresses which have dynamic filtering entries are filtered or forwarded based on the dynamic entry
- Intended for operation with GMRP-aware end stations only

GARP

- Objective:
  - registration and dissemination of information of any generic attribute over a bridged LAN
  - end stations can issue/revoke declarations for the attribute values
- Attributes are opaque to GARP
  - it is up to the GARP Application to define and interpret the generic attribute
  - Example: GMRP is a GARP application, where group membership is the attribute

GARP: Basic Notions (1)

- Simply fully-distributed many-to-many protocol
- Scalable (collapse the number of messages)
- Soft state (no acks, no confirmation, needs garbage collection)
- Protocol participants communicate their current state instead of just the triggering events

GARP: Basic Notions (2)

- Resilient against loss of a single packet
- Resilient against failure of GARP participants
- Operates in homogeneous and heterogeneous bridged LANs
- Small overhead (in terms of bandwidth)

GARP: Basic Notions (3)

- A participant that wishes to make a declaration sends Join messages
- If an applicant sees two Join messages, it does not need to send a Join message itself in order to participate in the declaration
- A participant that wishes to withdraw a declaration needs to send a single Leave message
- A periodic garbage collection is done using the LeaveAll procedure

Elements in a Participant

- There is a Registrar and an Applicant per attribute value that the participant knows of
- Registrar
  - records attribute registration declared by the other participants in the same segment
- Applicant
  - ensures that this participant's declarations are registered by other participant's registrars
  - ensures that when a participant de-registers an attribute, other participants that wish to maintain registration have a chance to re-declare (rejoin) it
  - i.e.. it looks after the interest of all would be participants

GARP Architecture

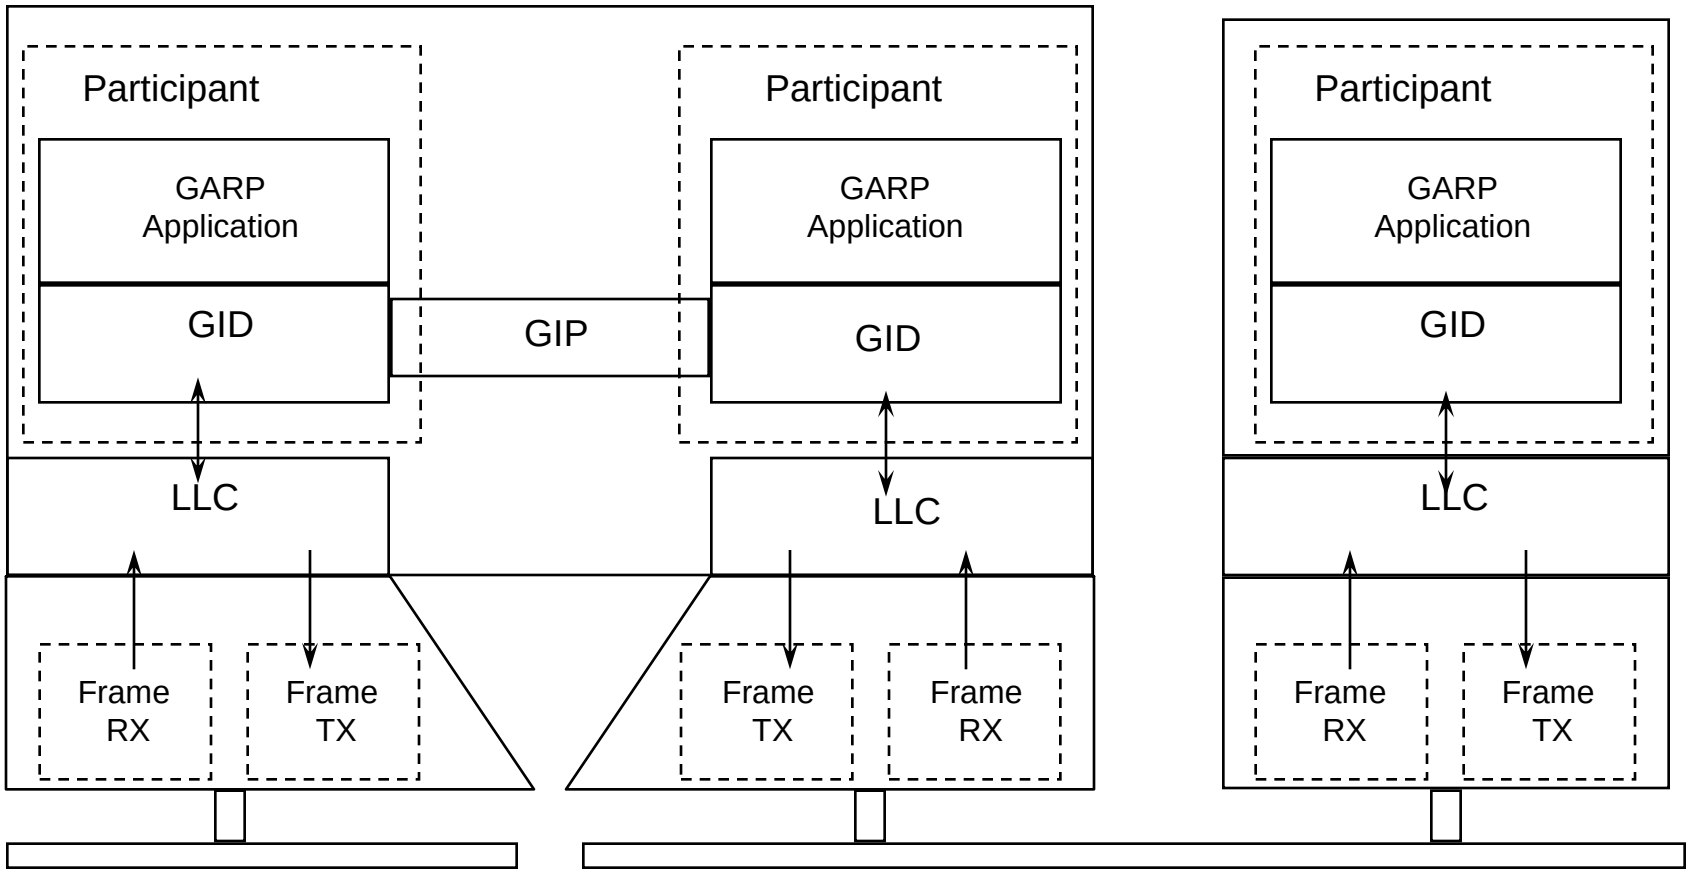

GID and GIP

- GID: GARP Information Distribution
  - distributes attribute information between participants on the same segment
  - contains the Registrar and the Applicant
- GIP: GARP Information Propagation
  - propagates attribute information between participants in the same bridge
    - registrations are propagated automatically
    - de-registrations are only propagated if this was the last registered port (excluding the one you propagate it to)

A Registration Example

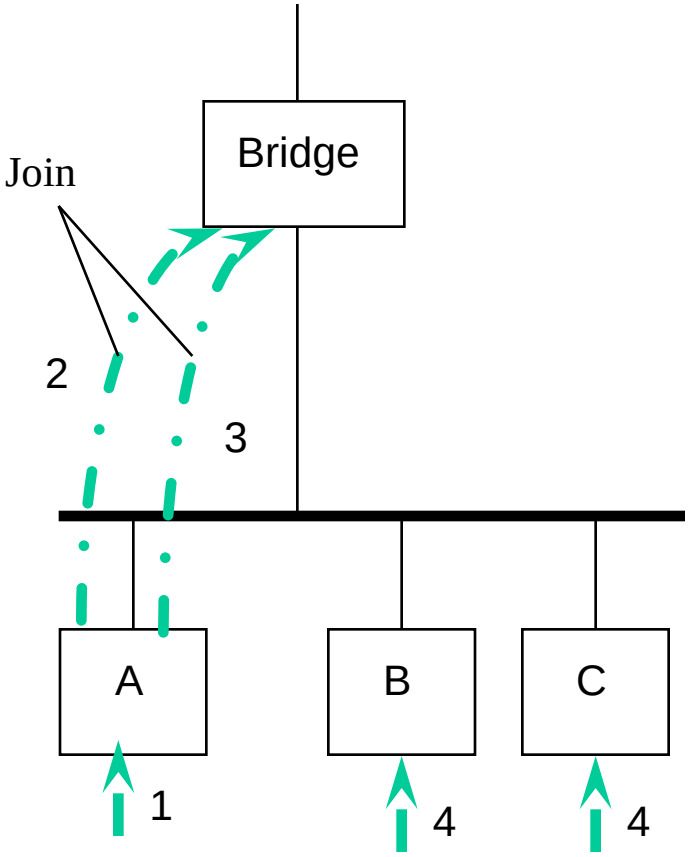

1.- station A decides to join group G
2.- after a JoinTime, A sends a Join(G). The bridge starts forwarding frames destined to group G
3.- after another JoinTime, A sends a second Join(G). This is done to be resilient to simple packet losses
4.- users B and C decide to join group G, but their Registrars have already recorded A's registration, so they don't send more messages. This is done for scalability and overhead reduction

A De-registration Example (1)

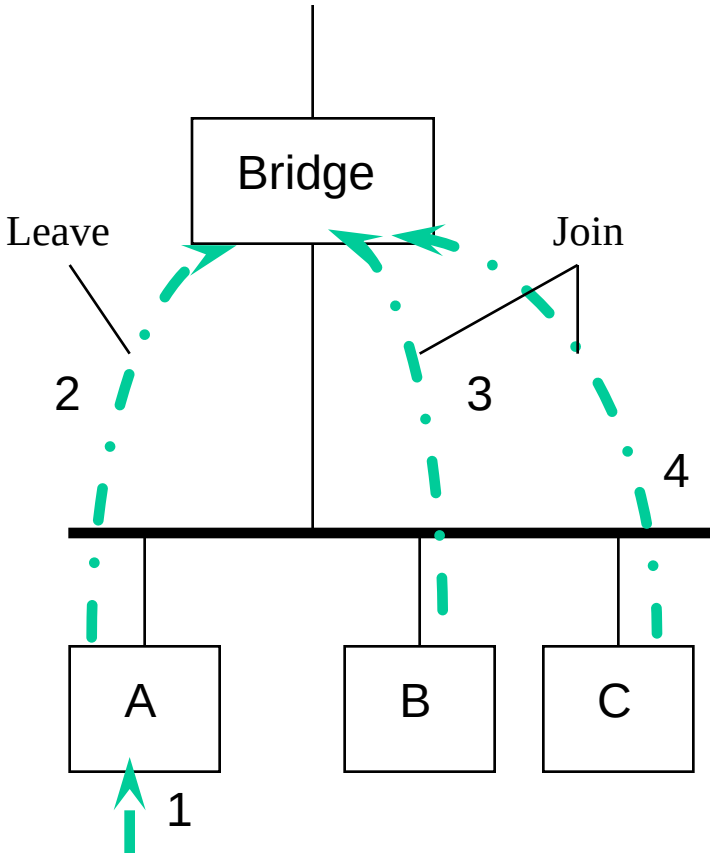

1.- the user at station A decides to leave group G, however B and C are still interested in receiving group G
2.- after a JoinTime, A sends a Leave(G)
3.- after another JoinTime, B sends a Join(G)
4.- after another JoinTime, C sends a second Join(G). Note that the second Join can be sent by another participant

Note that the bridge has never stopped forwarding frames destined to group G

A De-registration Example (2)

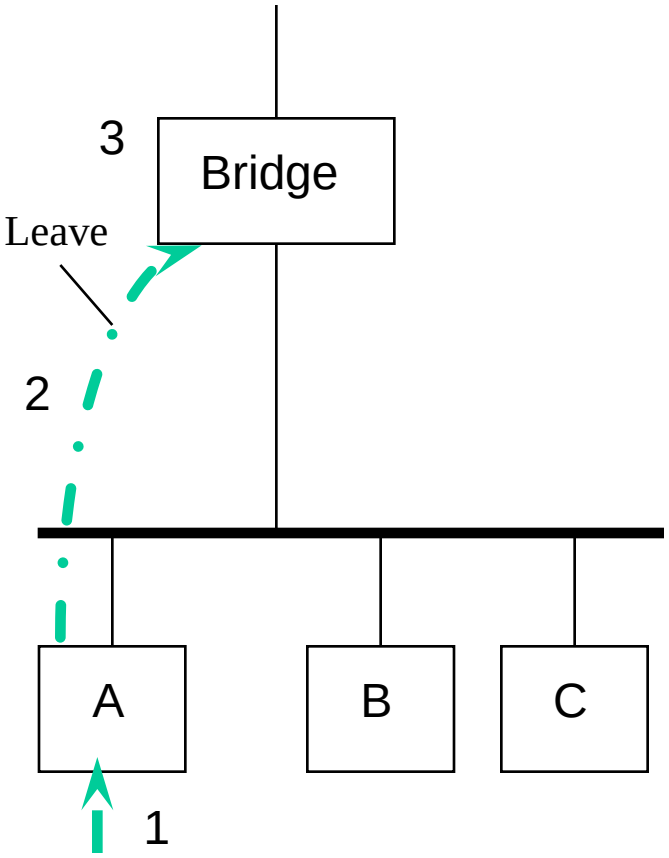

1.- the user at station A decides to leave group G. B and C left the group before
2.- after a JoinTime, A sends a Leave(G)
3.- after a LeaveTime without receiving a Join the bridge deregisters group G from this port and stops relaying frames for group G

LeaveAll Procedure

- periodically there is a clean up procedure
- a LeaveAll message is issued, which is like sending a Leave message to all groups that are registered
- if stations are active, they should reregister the groups that they are interested in. This might create a message implosion
- this happens every 10 s

Join/Leave Messages

- Join/Join All:
  - sent by an station/proxy that wants to join some (all) group(s). It sends two join messages unless it hears another join message from another station, then send just one.
- Leave All:
  - sent by the bridge after an inactivity period. It is a “are-you-alive?” message. If stations are alive, they should join all they groups they are subscribed
- Leave:
  - used to leave a group, if sent by an station, or to dismiss all group members, if sent by a bridge.
  - If a station hears of a Leave message, then it should join the group again, unless it hears of two other Join messages

Registrar State Machine (1)

- IN: someone registered this attribute value on this segment.
- Empty (MT): all declarations for this attribute value on this segment have been withdrawn.
- Leaving (LV): I had registered this attribute value, but now I am timing out the registration. If I don't see a (re)declaration before the LeaveTimer expires I will become MT.

Registrar State Machine (2)

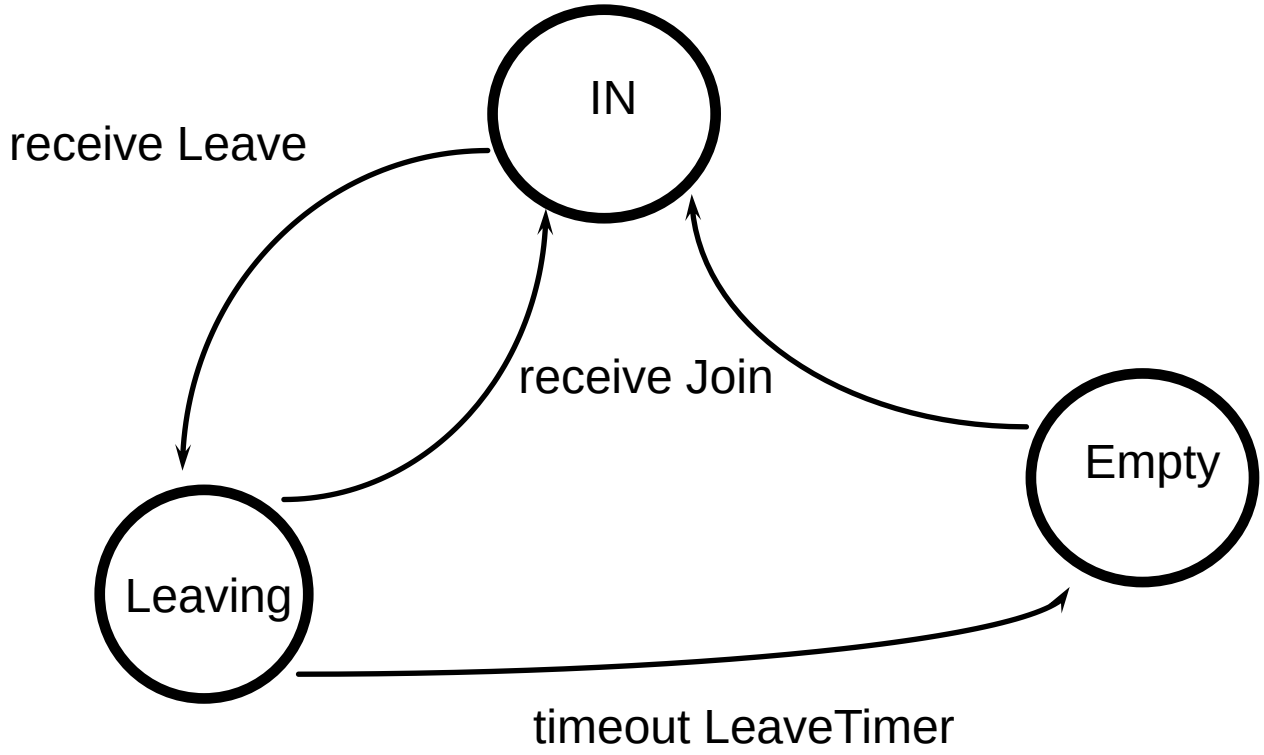

The Registrar does not send messages, but its state is visible through the messages sent by the Applicant: JoinIn and JoinEmpty.

Applicant State (1)

- Very Anxious (V): No JoinIn message has been received and no Join message has been sent. The applicant has no reason to believe that the other Registrars have registered this attribute value (Join_Msgs=0).
- Anxious (A): if no messages have been lost other Registrars will have registered this attribute value (Join_Msgs=1); a second Join needs to be sent.
- Quiet (Q): no need to send more messages (Join_Msgs=2).
- Leaving(L): need to send a Leave message

Applicant State (2)

- Active member (A): attempts to maintain the registration of an attribute value
- Passive member (P): attempts to maintain the registration of an attribute value, but still hasn't sent a Join message to declare it
- Observer (O): tracks the attribute state, but does not wish to make a declaration

Applicant State(3)

- The combined state is:

	Very Anxious	Anxious	Quiet	Leaving
Active Member	VA	AA	QA	LA
Passive Member	VP	AP	QP	--
Observer	VO	AO	QO	LO

- Note that there is no LP (Leaving Passive), since a Passive Member can directly transition to Observer

Applicant State Machine

- If an Observer (O) wants to become a Member, it first needs to become a Passive Member (P). If it were Quiet, it can just become Passive and Quiet without sending messages. Otherwise, it will become a Very Anxious Passive Member.
- If a Passive Member sends a Join, it becomes an Active Member.
- If an Active Member receives a Leave or LeaveAll, it becomes a Passive Member. It will then become Very Anxious if it still wants to be a Member.
- Full state machine is in table 12-3 in the standard.

Applicant-Only Participants

- Simplification of the GARP Participant, where it only wants to make declarations - it does not include the Registrar.
  - Example: a station that only wants to receive multicast frames.
  - This station will not send, so it does not need to implement source pruning (it doesn't care about the state of the group registration).
- For the generated messages, the Registrar is assumed in the IN state.
- No need to support the LO state or a LeaveAll timer (and messages).

Simple Applicant Participants

- Further simplification of the Applicant Only Participant.
- Simplest compatible GARP Participant
- The Passive Member and Observer states are removed.
- Result: initial Join and final Leave are always sent.
- If you have lots of these, there may be a lot of additional traffic; in this case, the use of the Applicant Only Participant is recommended.

GARP Application

- Defines:
  - Attribute types, set of values and their semantics
  - structure and encoding of the GARP PDU
  - Group MAC address used
  - Application Identifier

GARP PDU

Header Structure

1	2	3	4	5	6
Protocol Identifier		Application ID		No. of Messages	

Message Structure

1	2	3	4	5	6
Message Length		Msg. Type	Attrib. Type	No. of Messages	
Message Header				msg. body (may be null)	

Msg. Body Structure

1	…	…	…	msg. length - 4
Operator 1		…	Operator N	

Operator Structure

1	2	3	…	Operator Length
Operator Length	Operator Type	Operand (opaque to GARP)		

GARP Application Addresses

Assignment	Value
GMRP address	01-80-C2-00-00-20
VMRP (for VLANs)	01-80-C2-00-00-21
reserved	from 01-80-C2-00-00-22 to 01-80-C2-00-00-2F

Timers (1)

- Join Timer (jt): delays the transmission of any message so that there is no message implosion or redundancy
- Leave Timer (lt): delays the de-registration of an attribute so that other would be participants can (re)join
- Leave All Period (lap): clean-up/ are-you-alive procedure. Forces all participants to reregister all attributes

Timers (2)

- LeaveTime > 2×JoinTime
- LeaveAllTime > LeaveTime

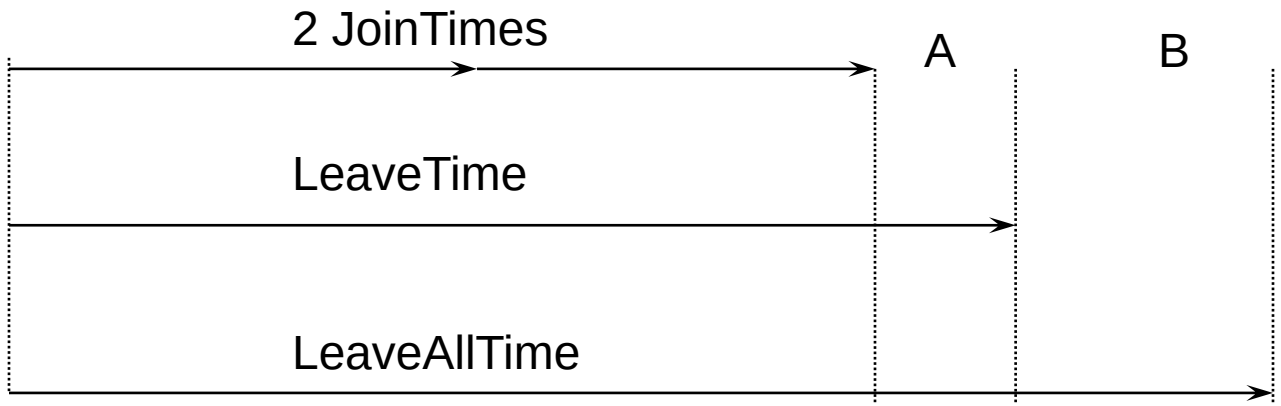

Timers (3)

- All timeouts are uniformly distributed between 0 and JoinTime
- Recommended values

Parameter	Value (milliseconds)
JoinTime	200
LeaveTime	600
LeaveAllTime	10000

Additional Characteristics

- Network administrator can disable/force the registration/application of particular attributes in certain ports
  - e.g. if a IP multicast router is attached to a port he/she can set the port of the bridge it is attached to to be in forwarding mode B (Forward all unregistered addresses)