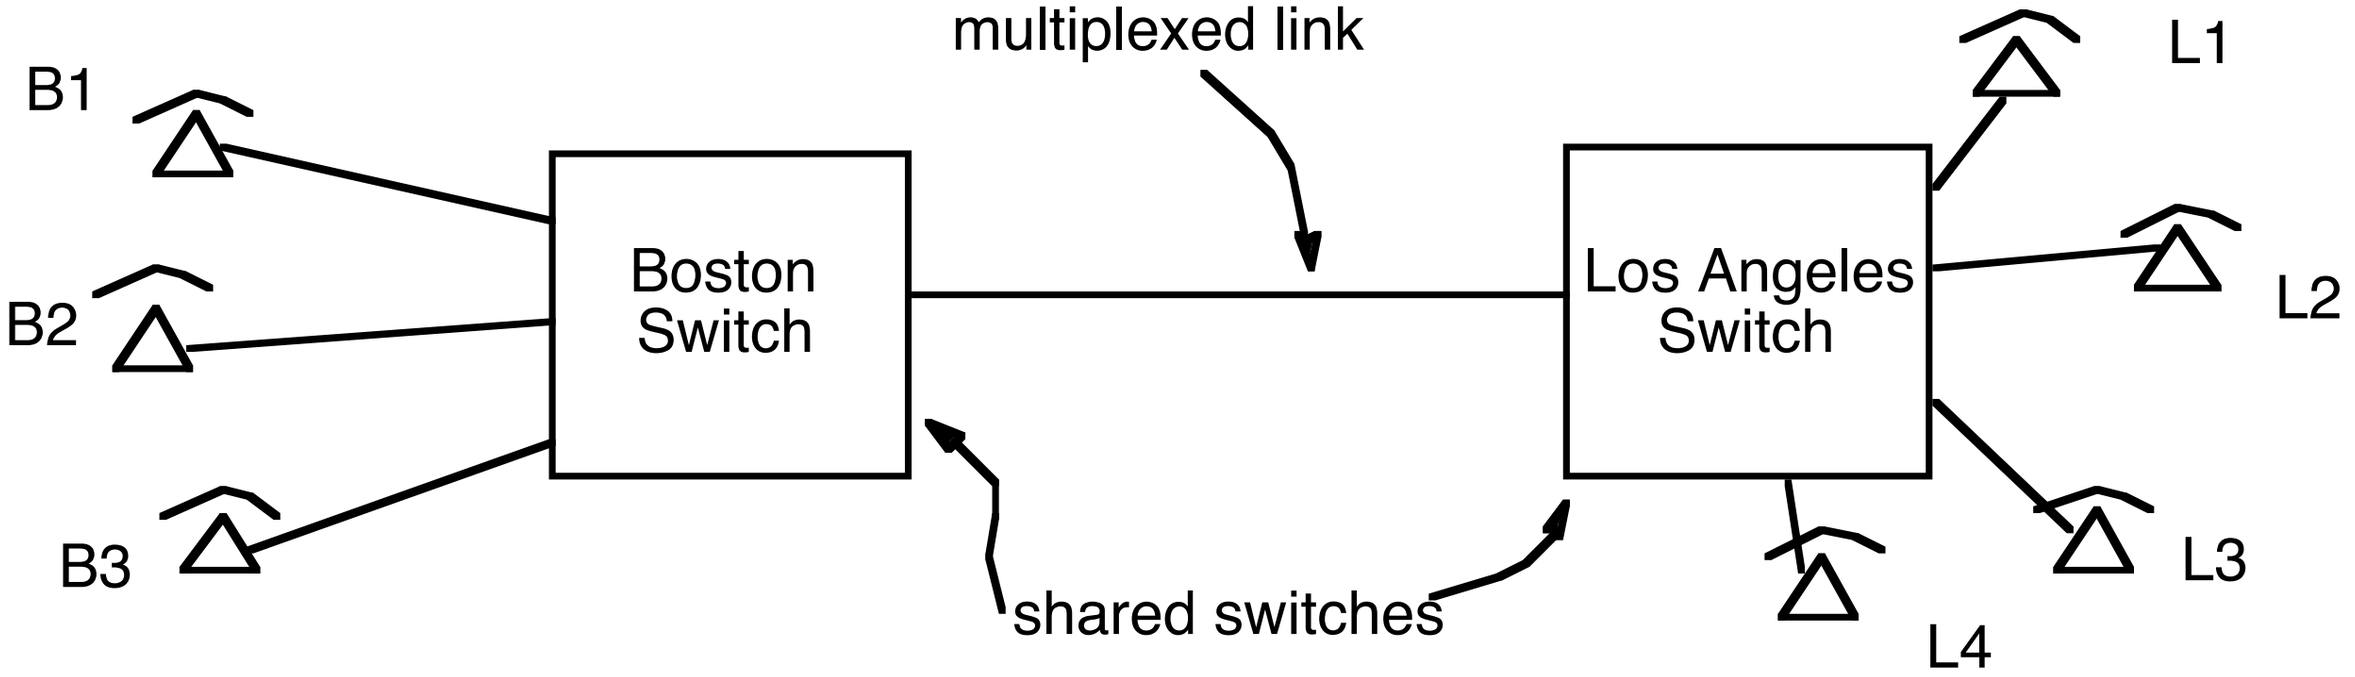

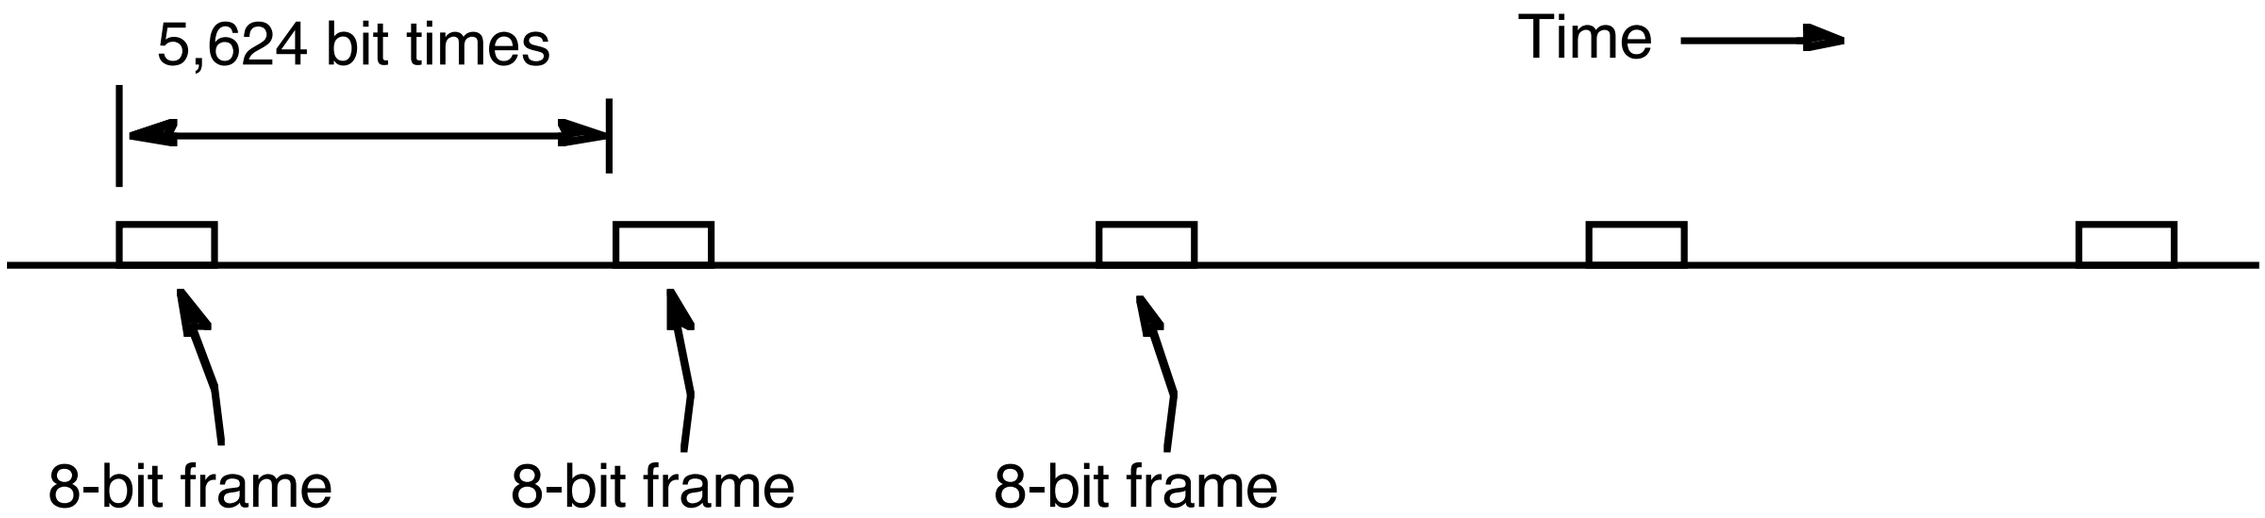

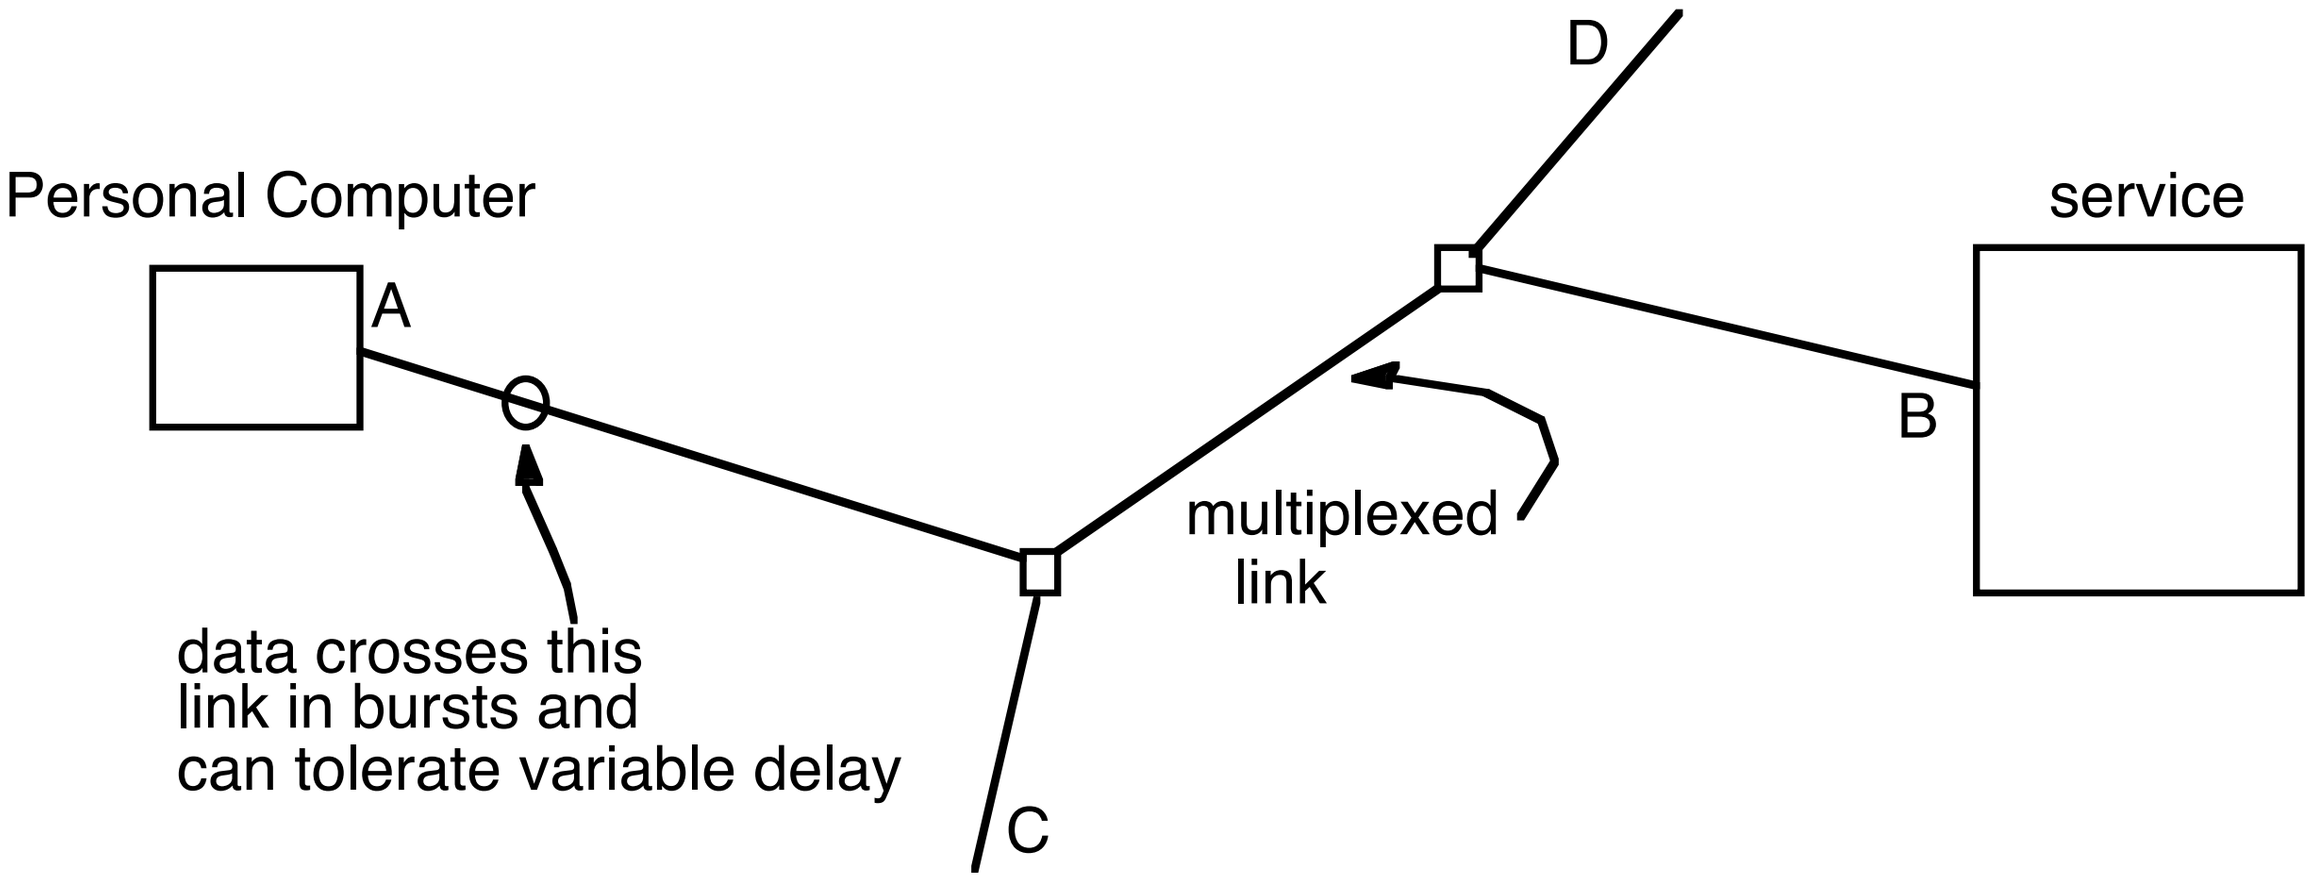

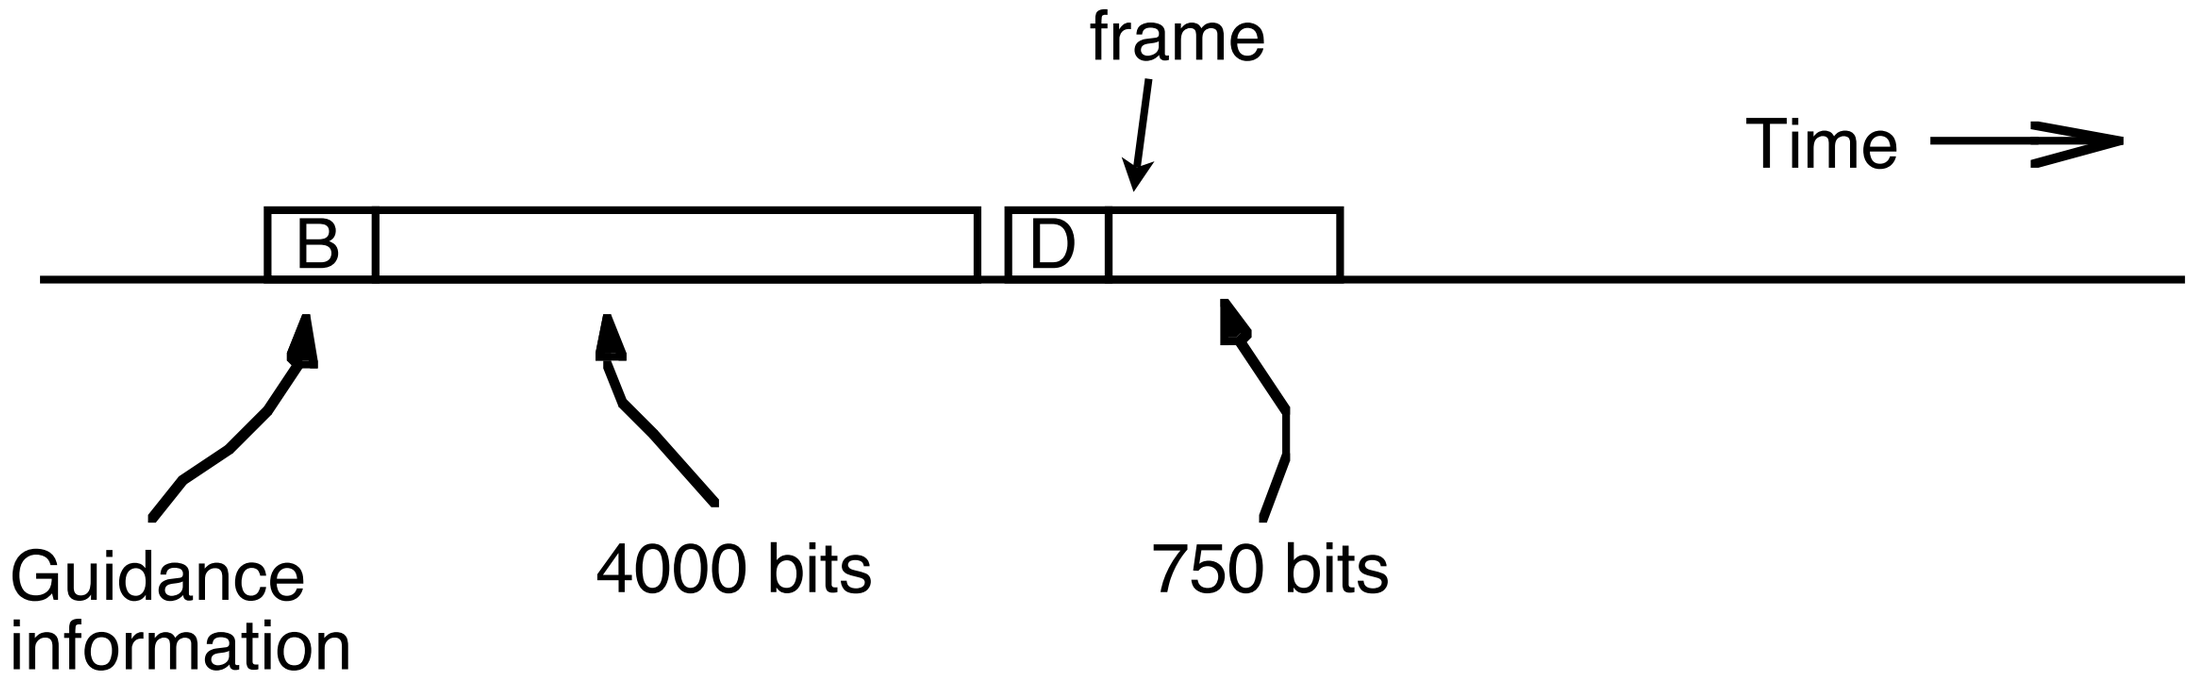

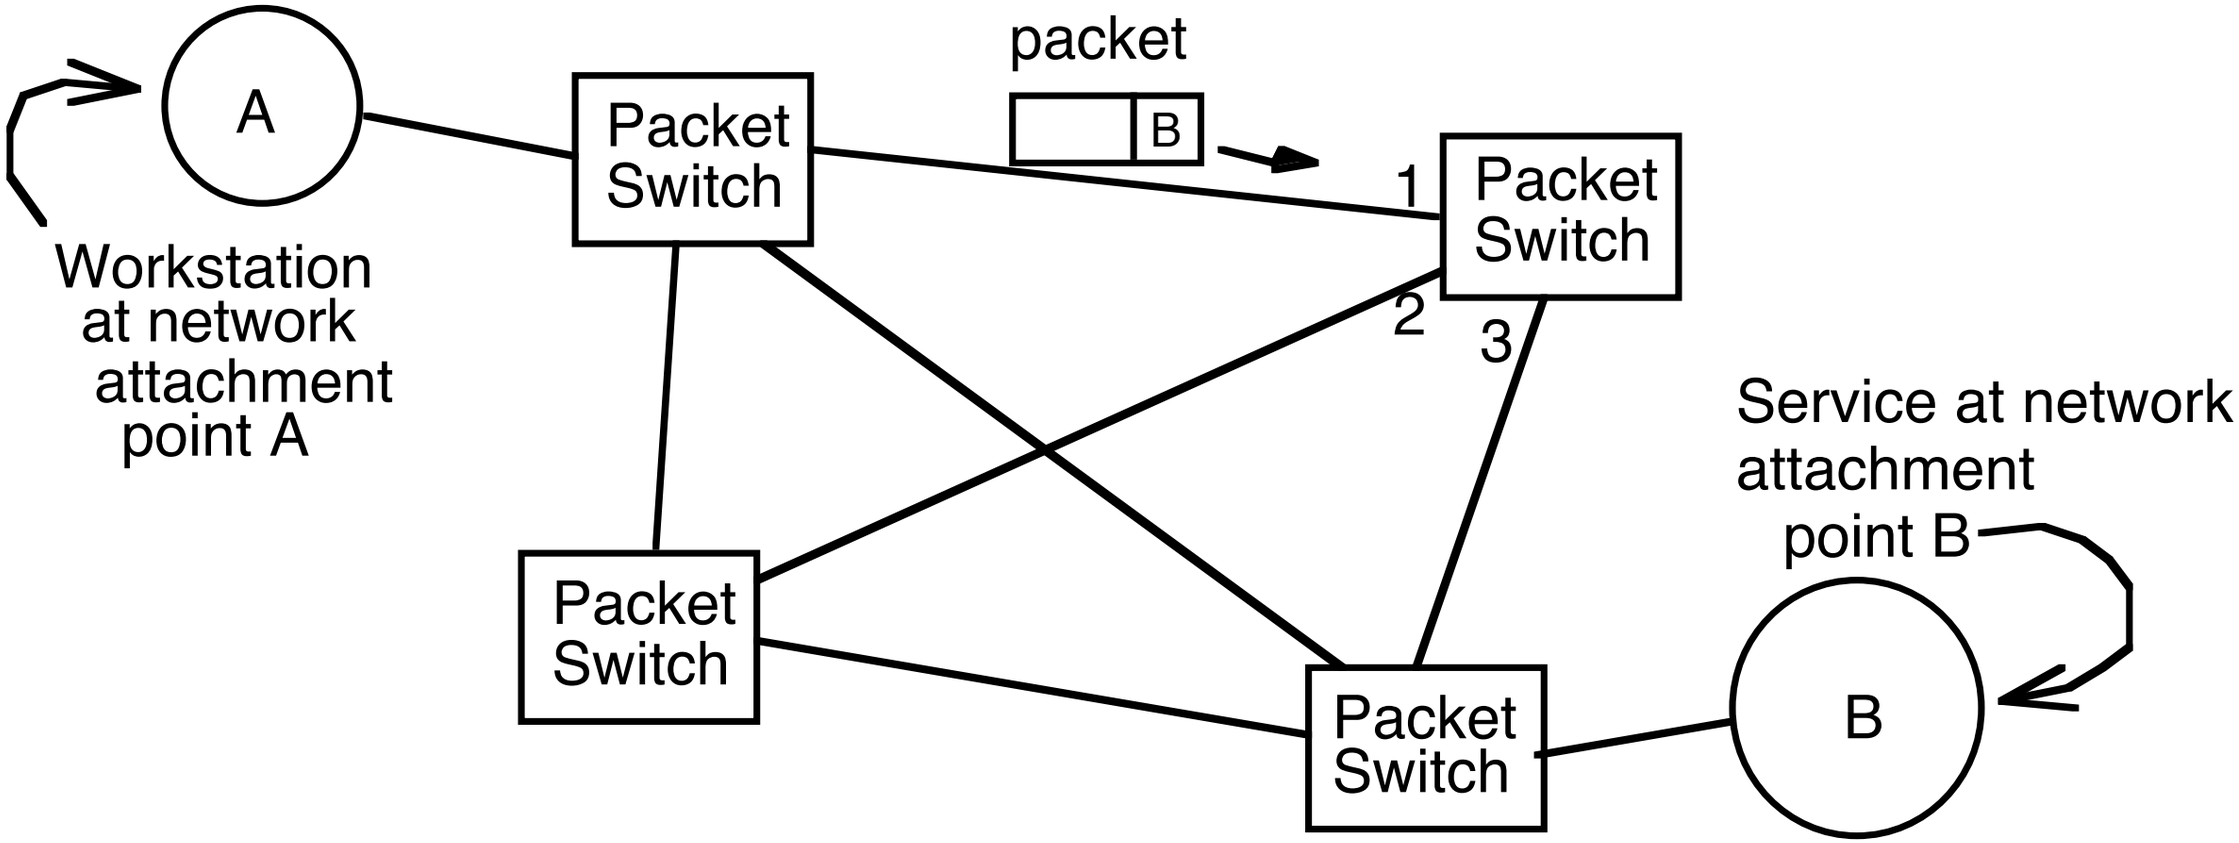

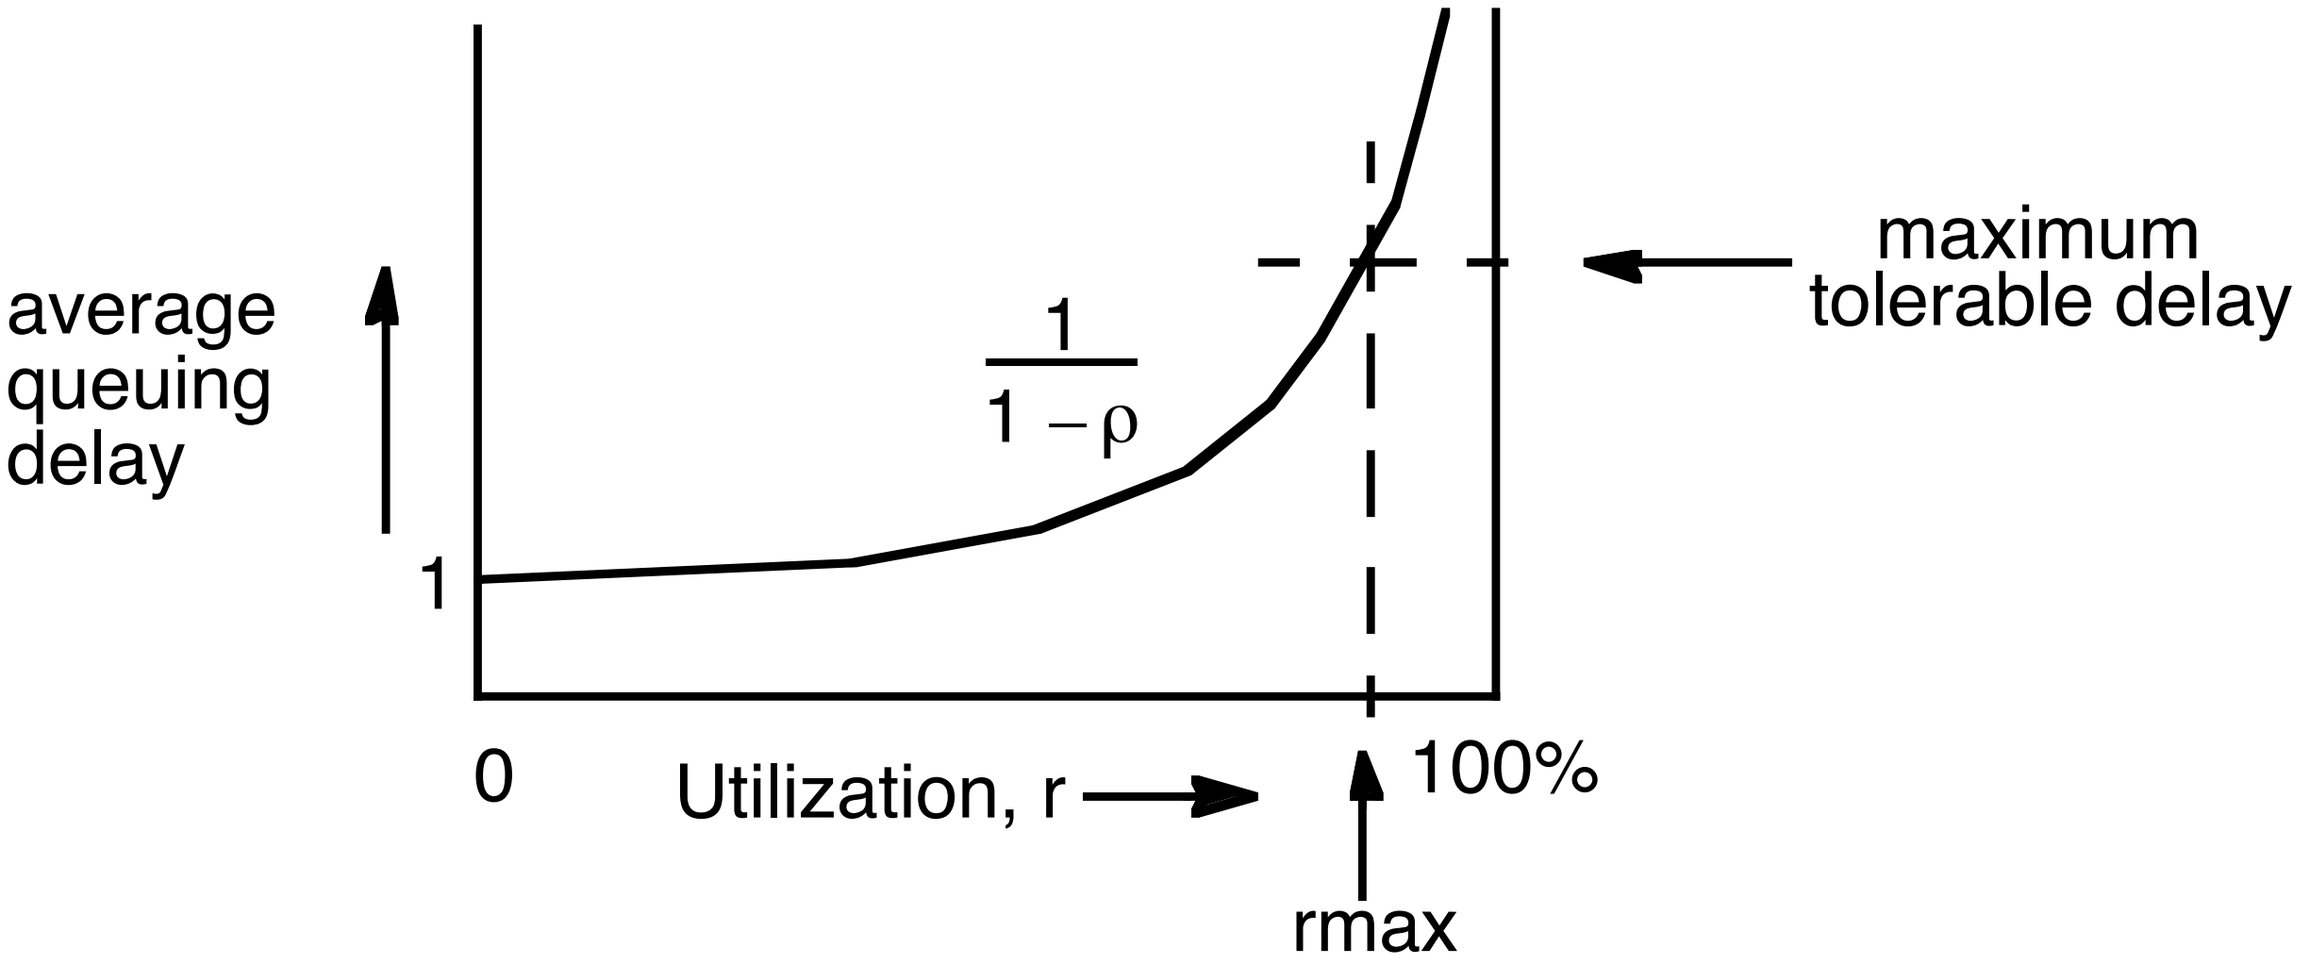

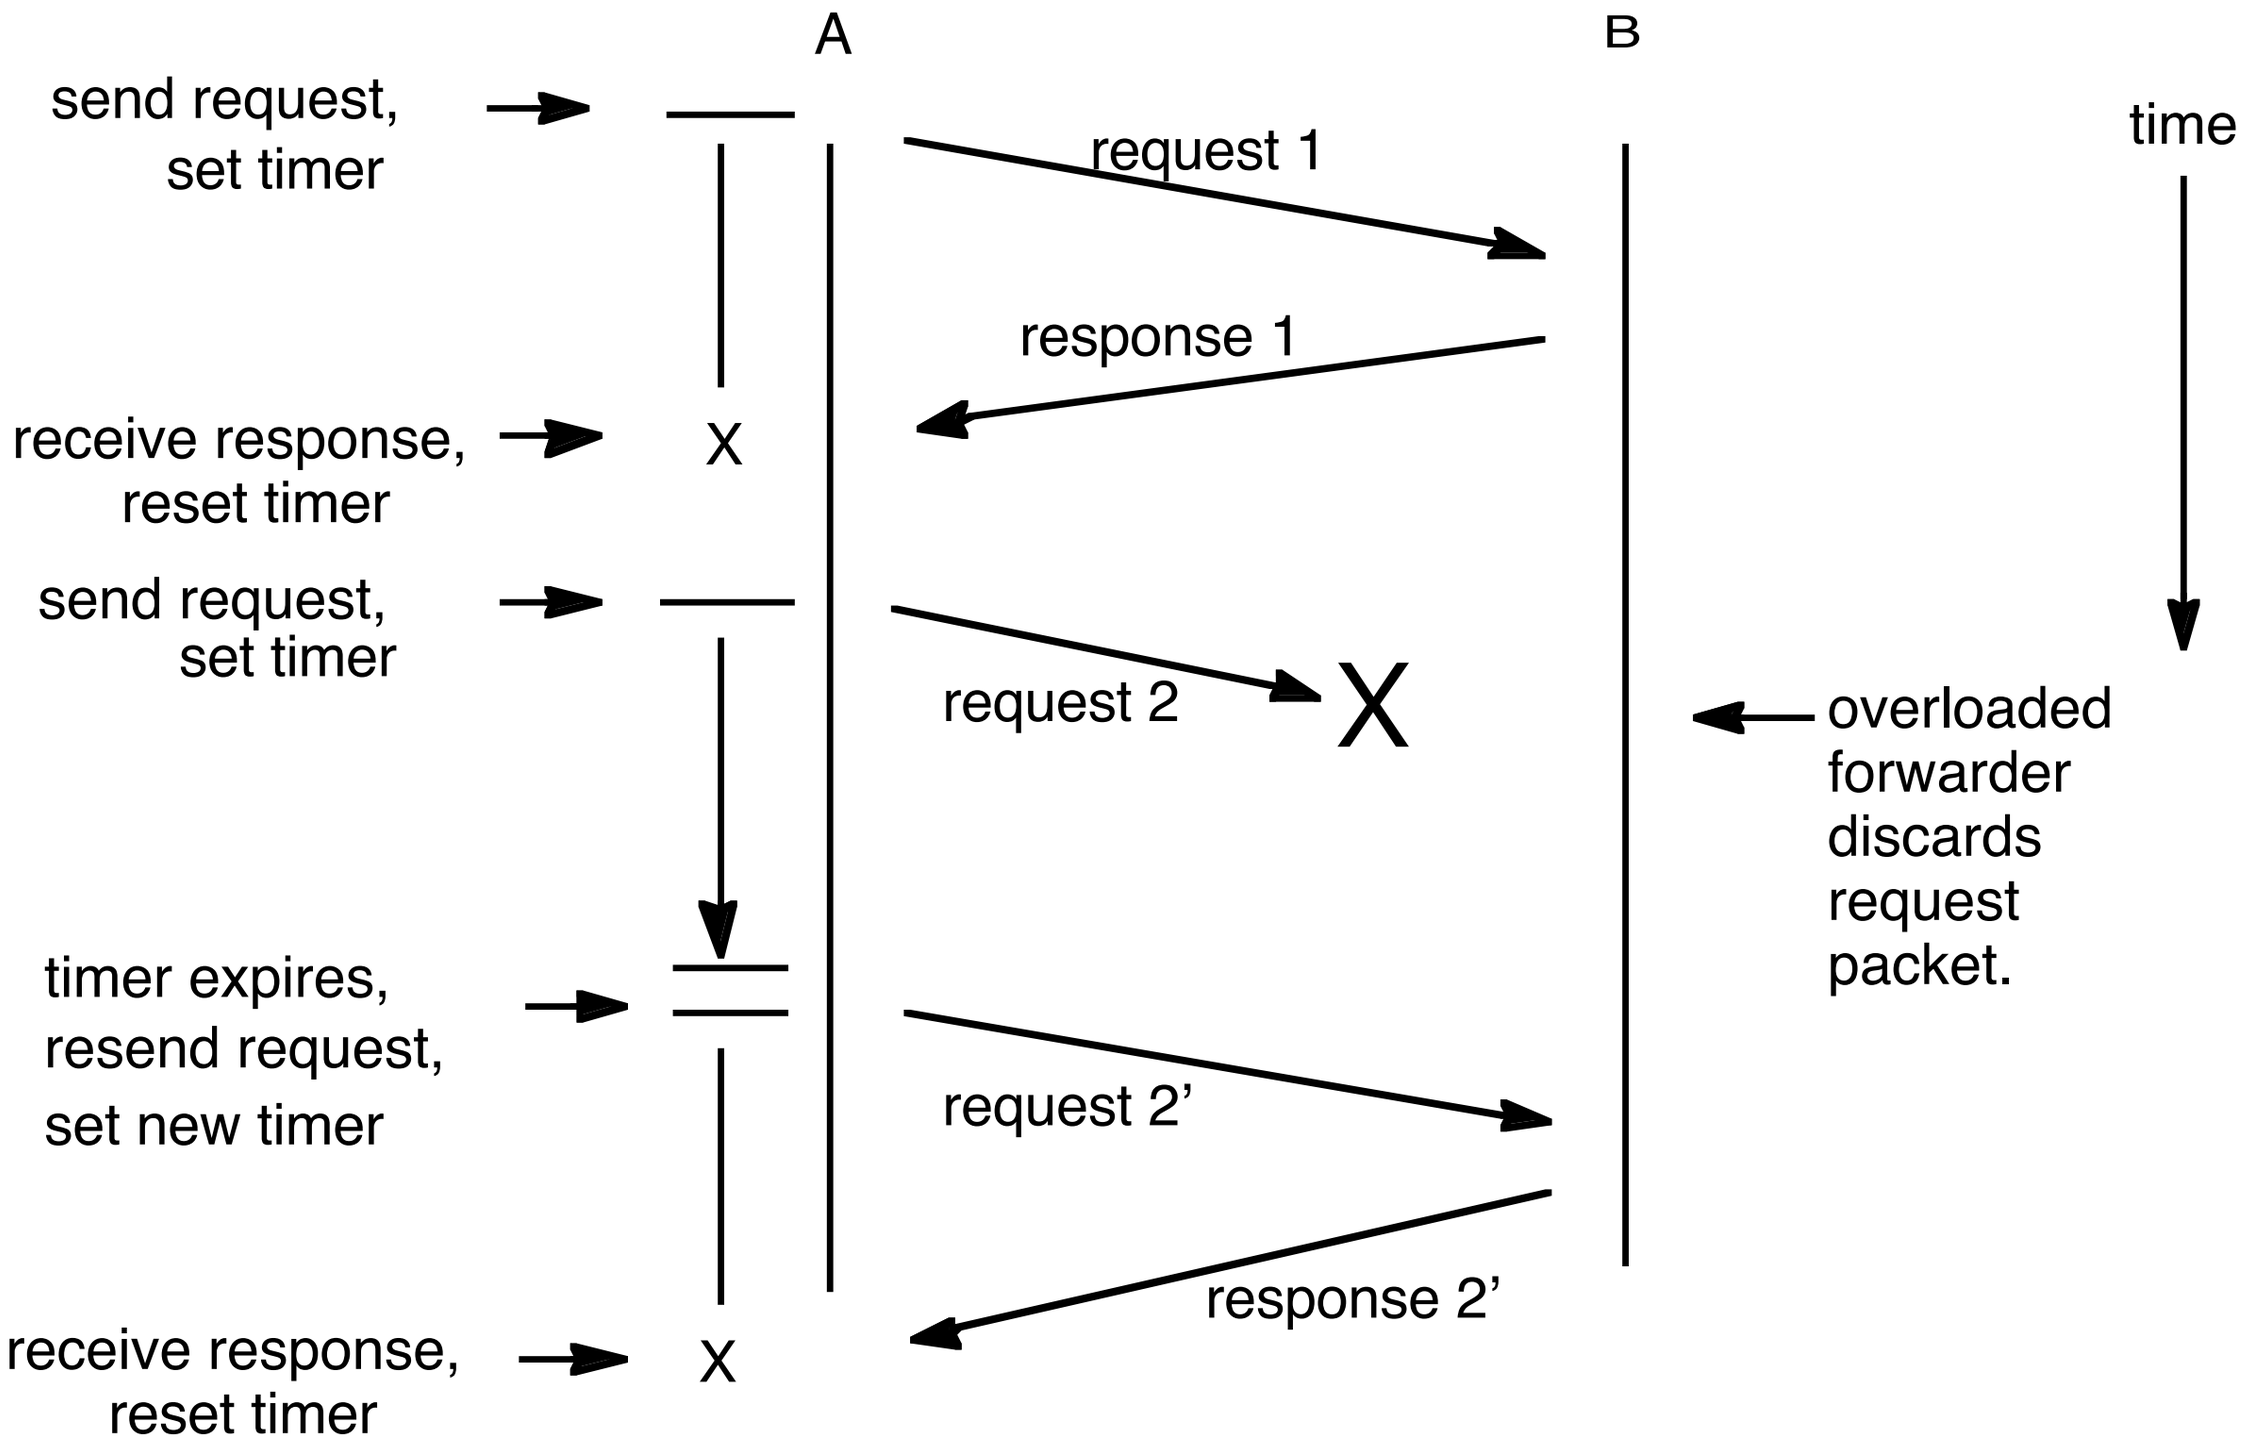

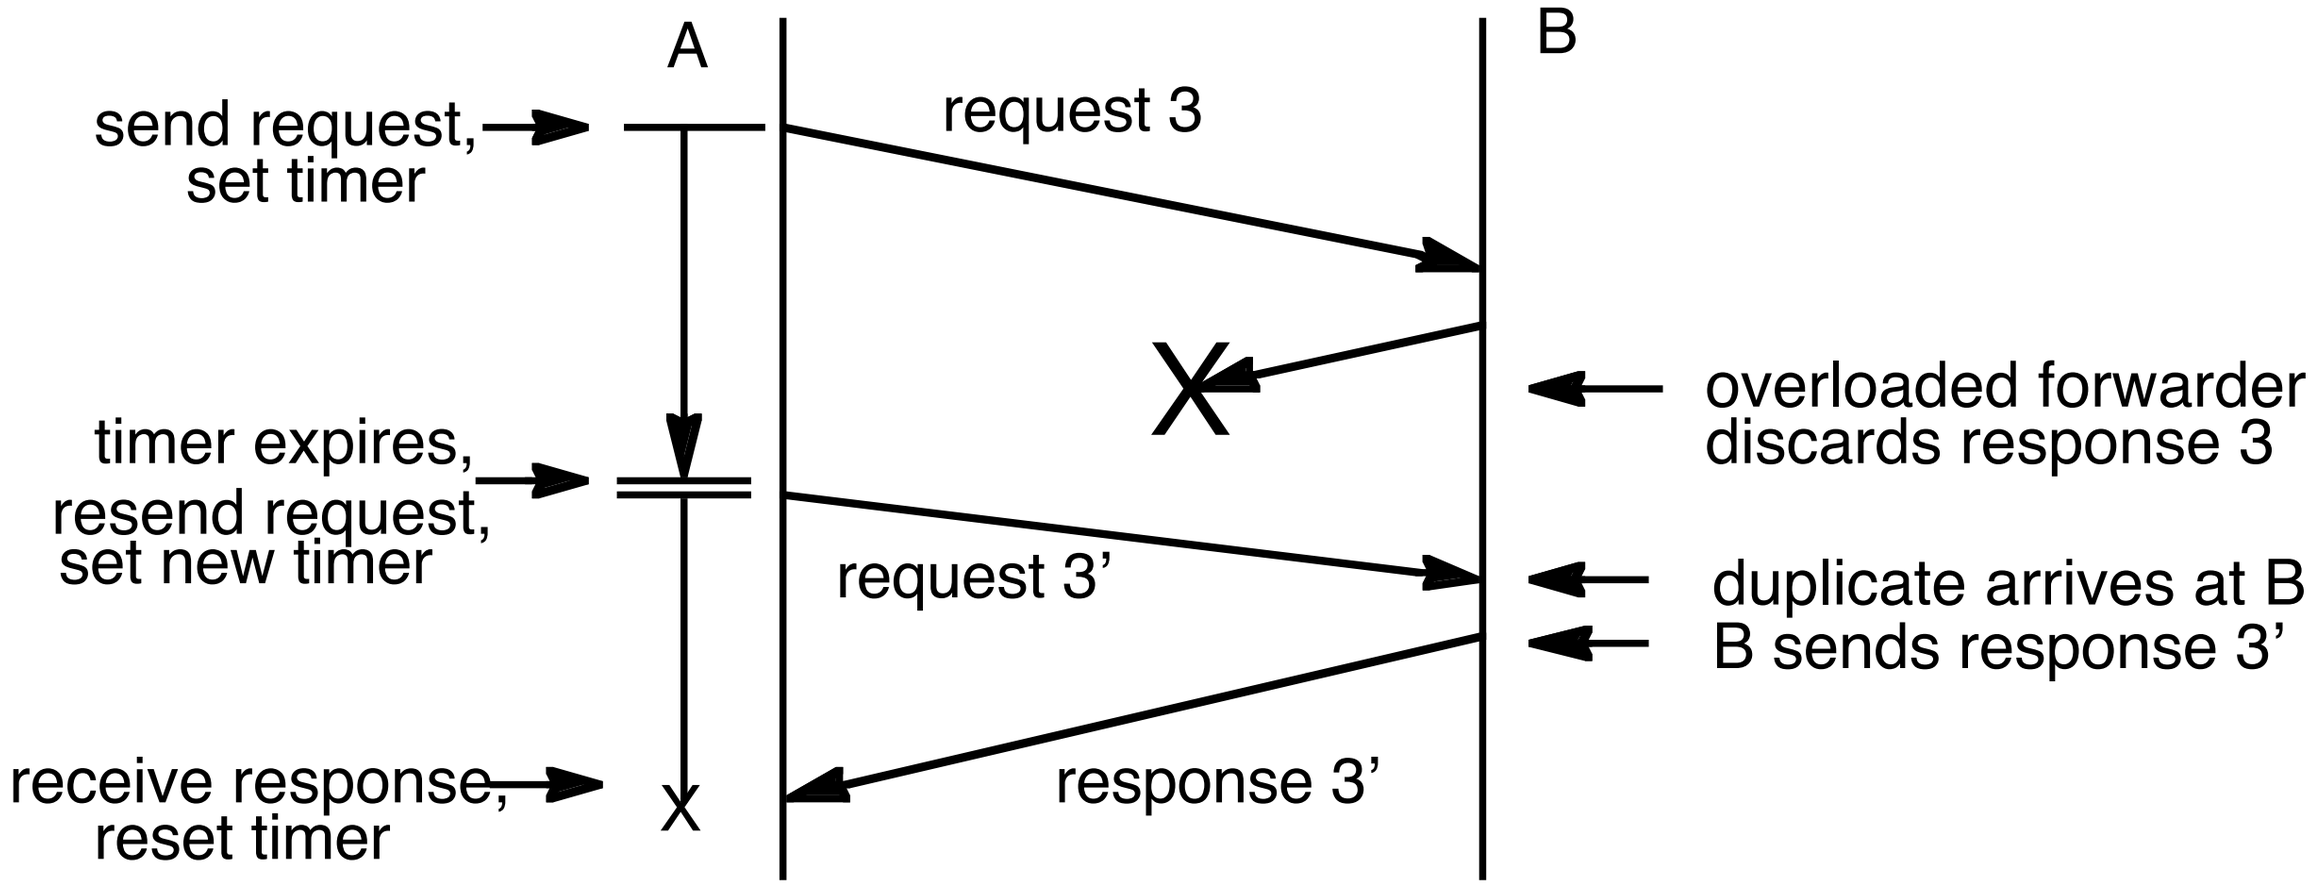

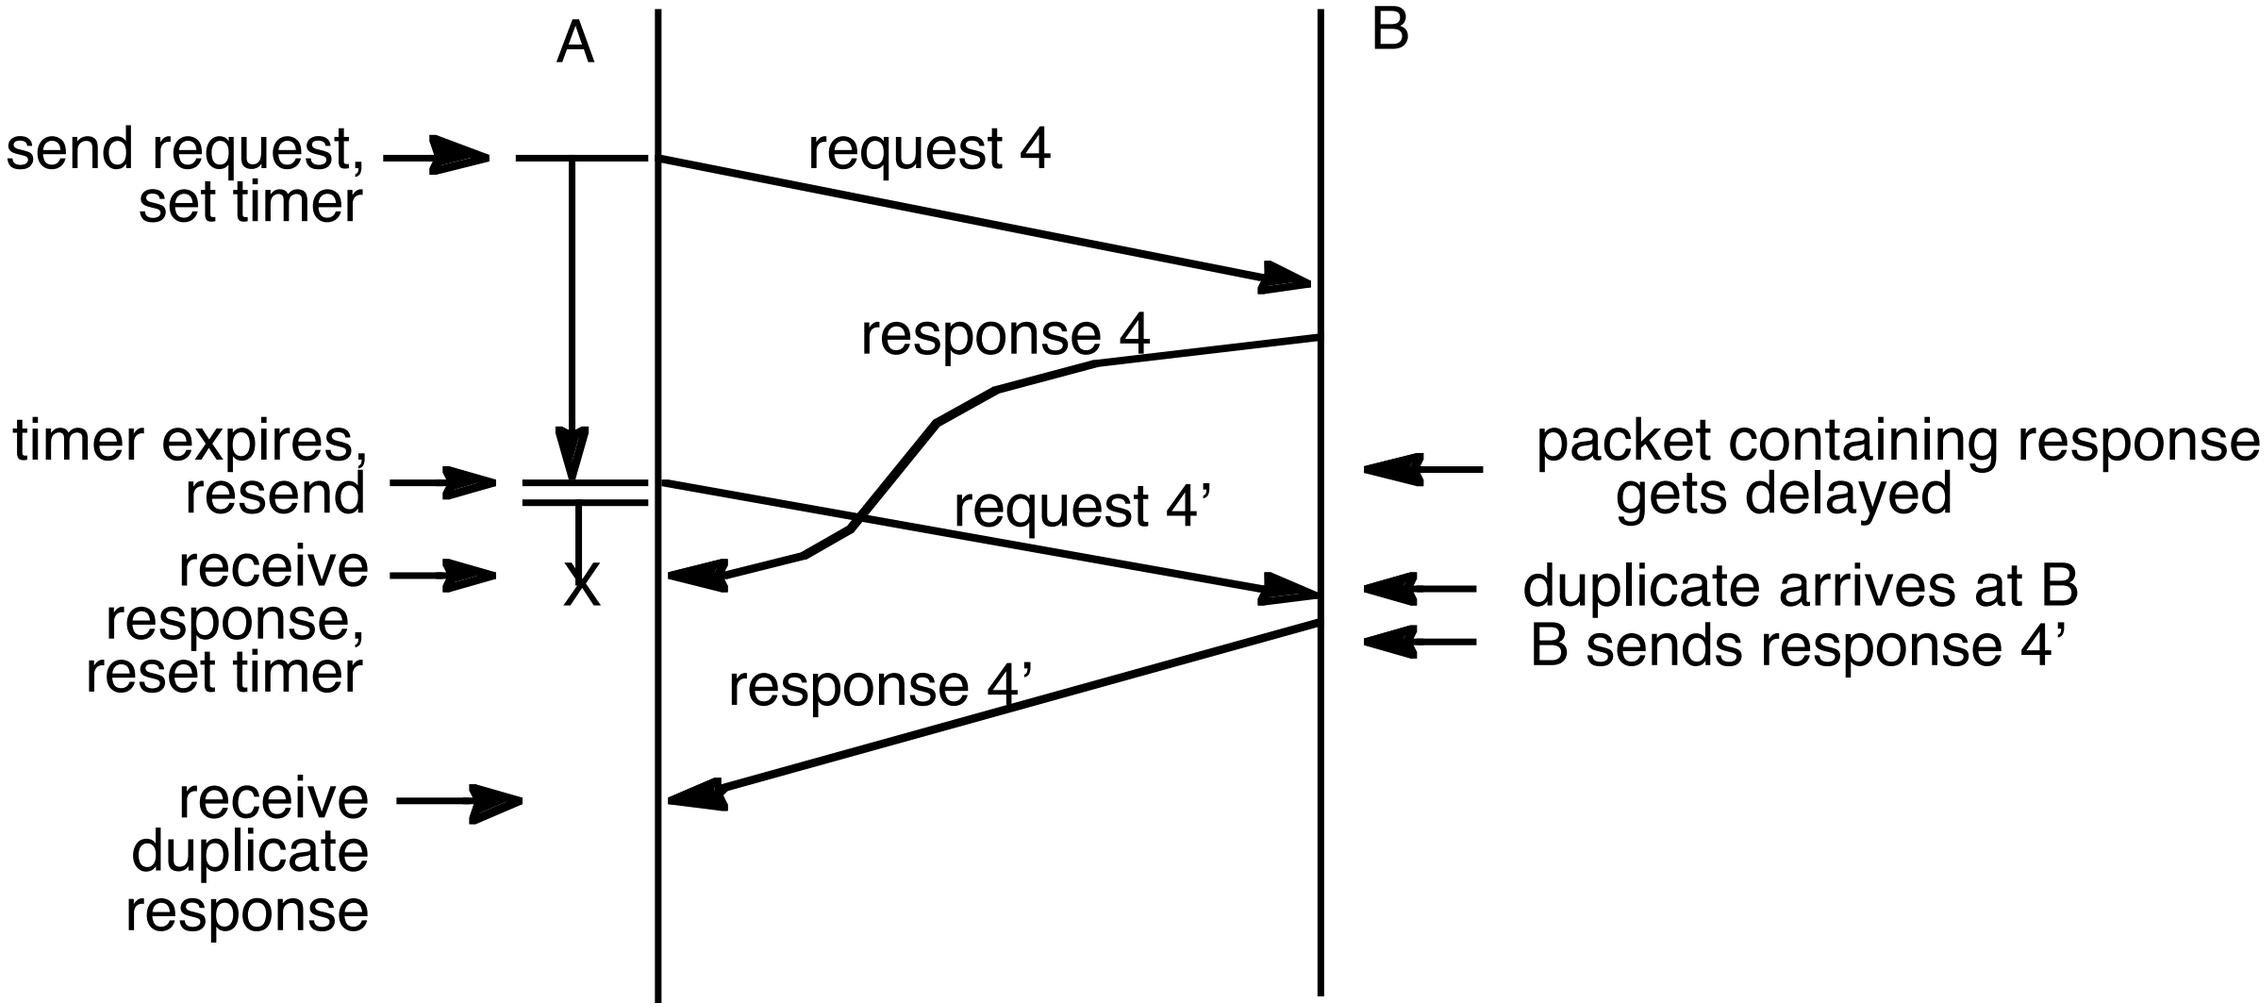

Network Type	Application characteristics: Continuous stream (e.g., interactive voice)	Application characteristics: Bursts of data (most computer-to-computer data)	Response to load variations
isochronous (e.g., telephone network)	good match	wastes capacity	(hard-edged) either accepts or blocks call
asynchronous (e.g., Internet)	variable latency upsets application	good match	(gradual) 1 variable delay 2 discards data 3 rate adaptation

Networks encounter a vast range of

- Data rates
- Propagation, transmission, queuing, and processing delays.
- Loads
- Numbers of users

Networks traverse hostile environments

- Noise damages data
- Links stop working

Best-effort networks have

- Variable delays
- Variable transmission rates
- Discarded packets
- Duplicate packets
- Maximum packet length
- Reordered delivery

result ← FIRE (#, *target*, *action*)

procedure FIRE (*nmiss*, *where*, *react*)

...

return *result*

Client stub

Service stub

Prepare request message. Send to service Wait for response.

request:

```
proc: FIRE
args: 3
type: integer
value: 2
type: string
value: "Lucifer"
type: procedure
value: EVADE
```

Receive request message. Call requested procedure. Prepare response message. Send to client.

response:

```
acknowledgment
type: string
value: "destroyed"
```

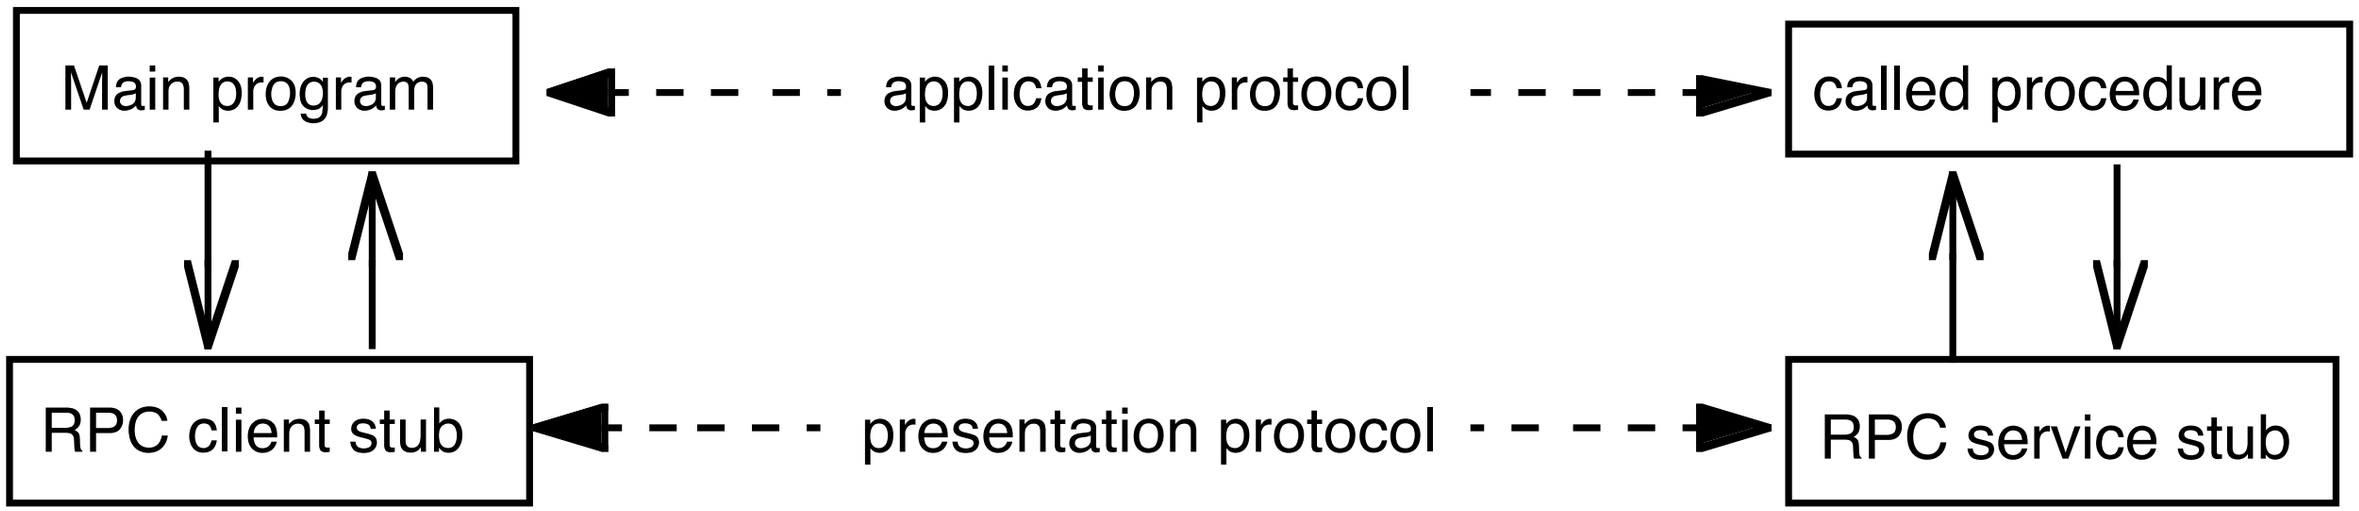

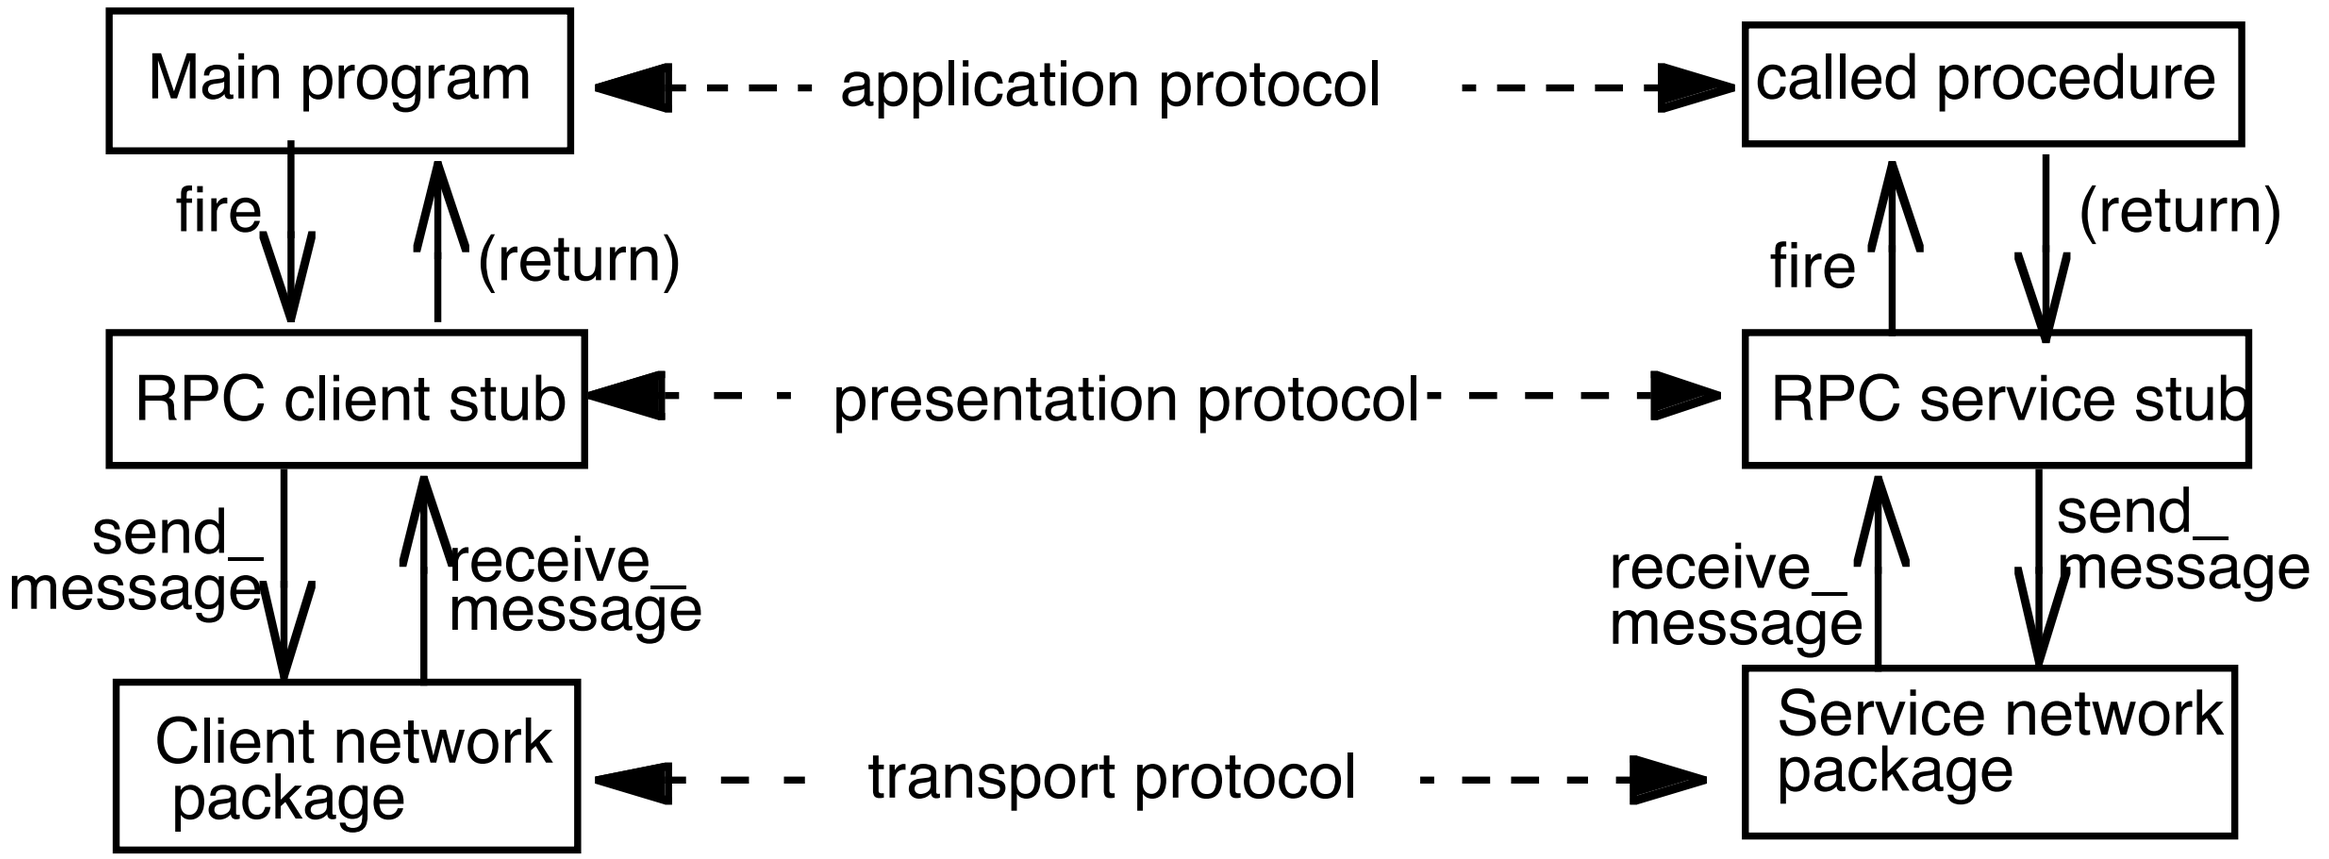

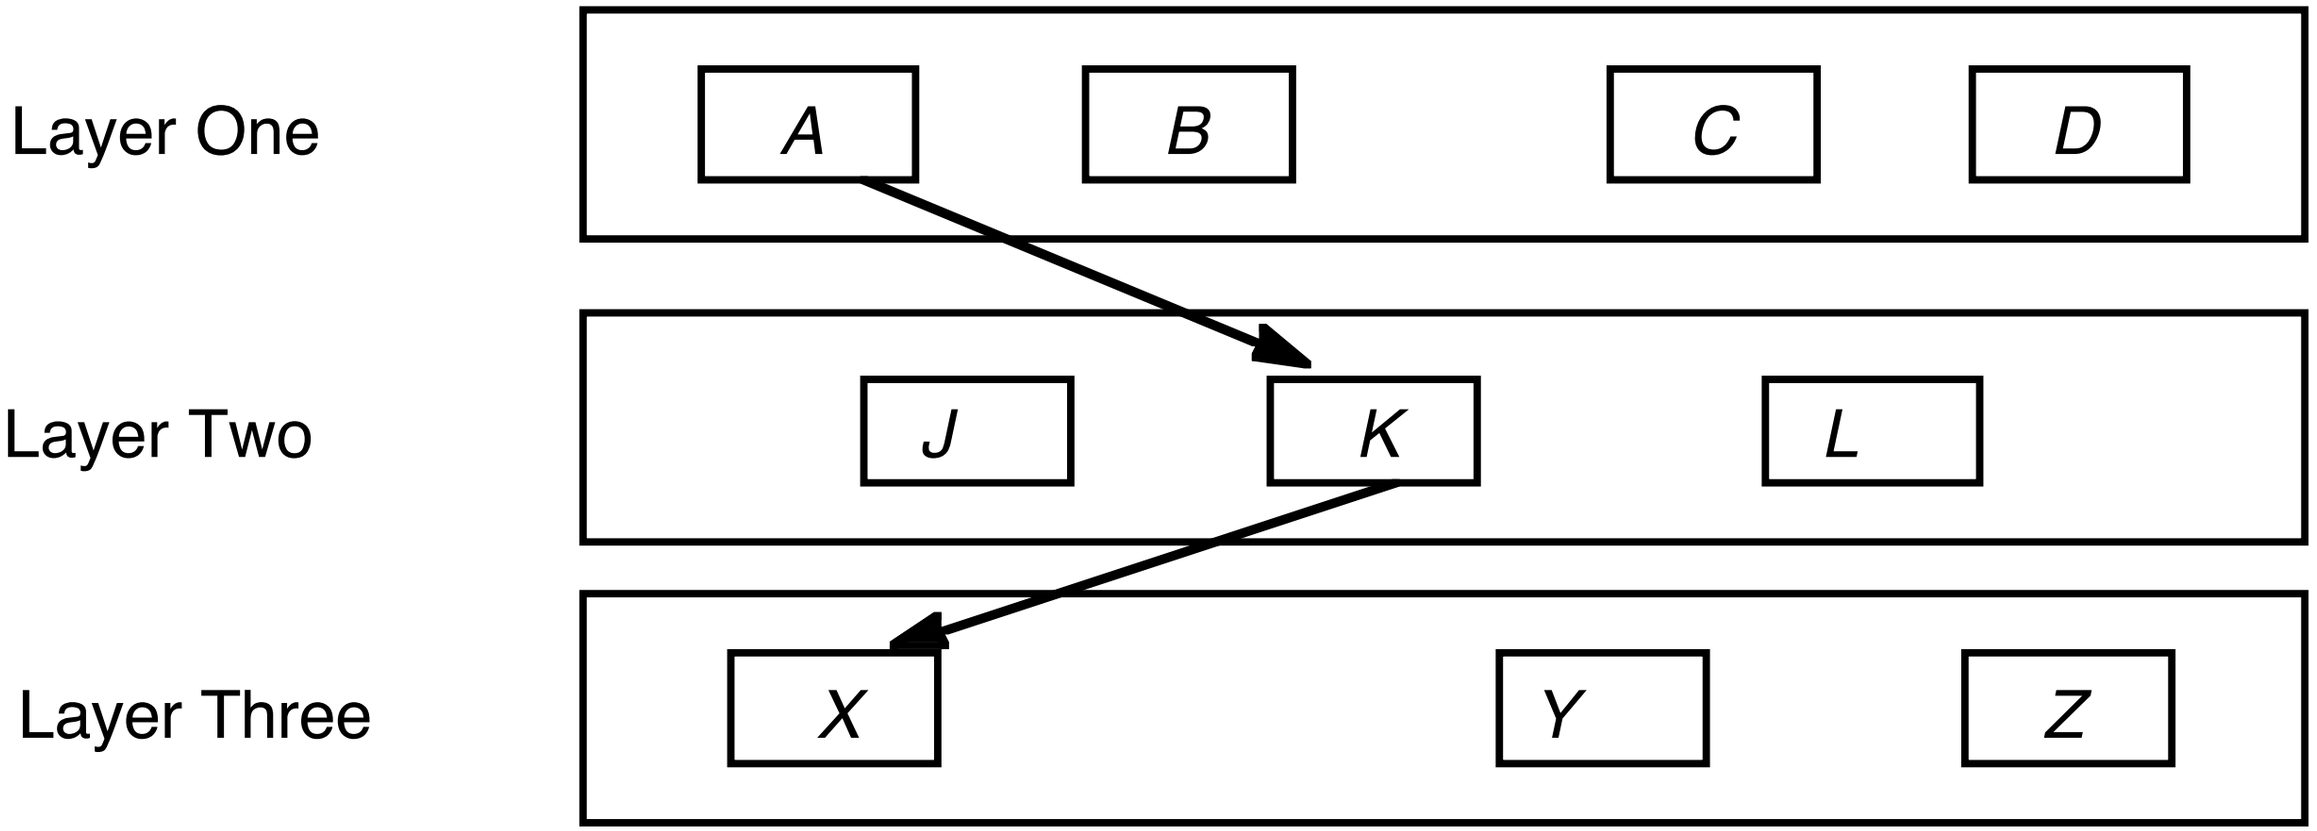

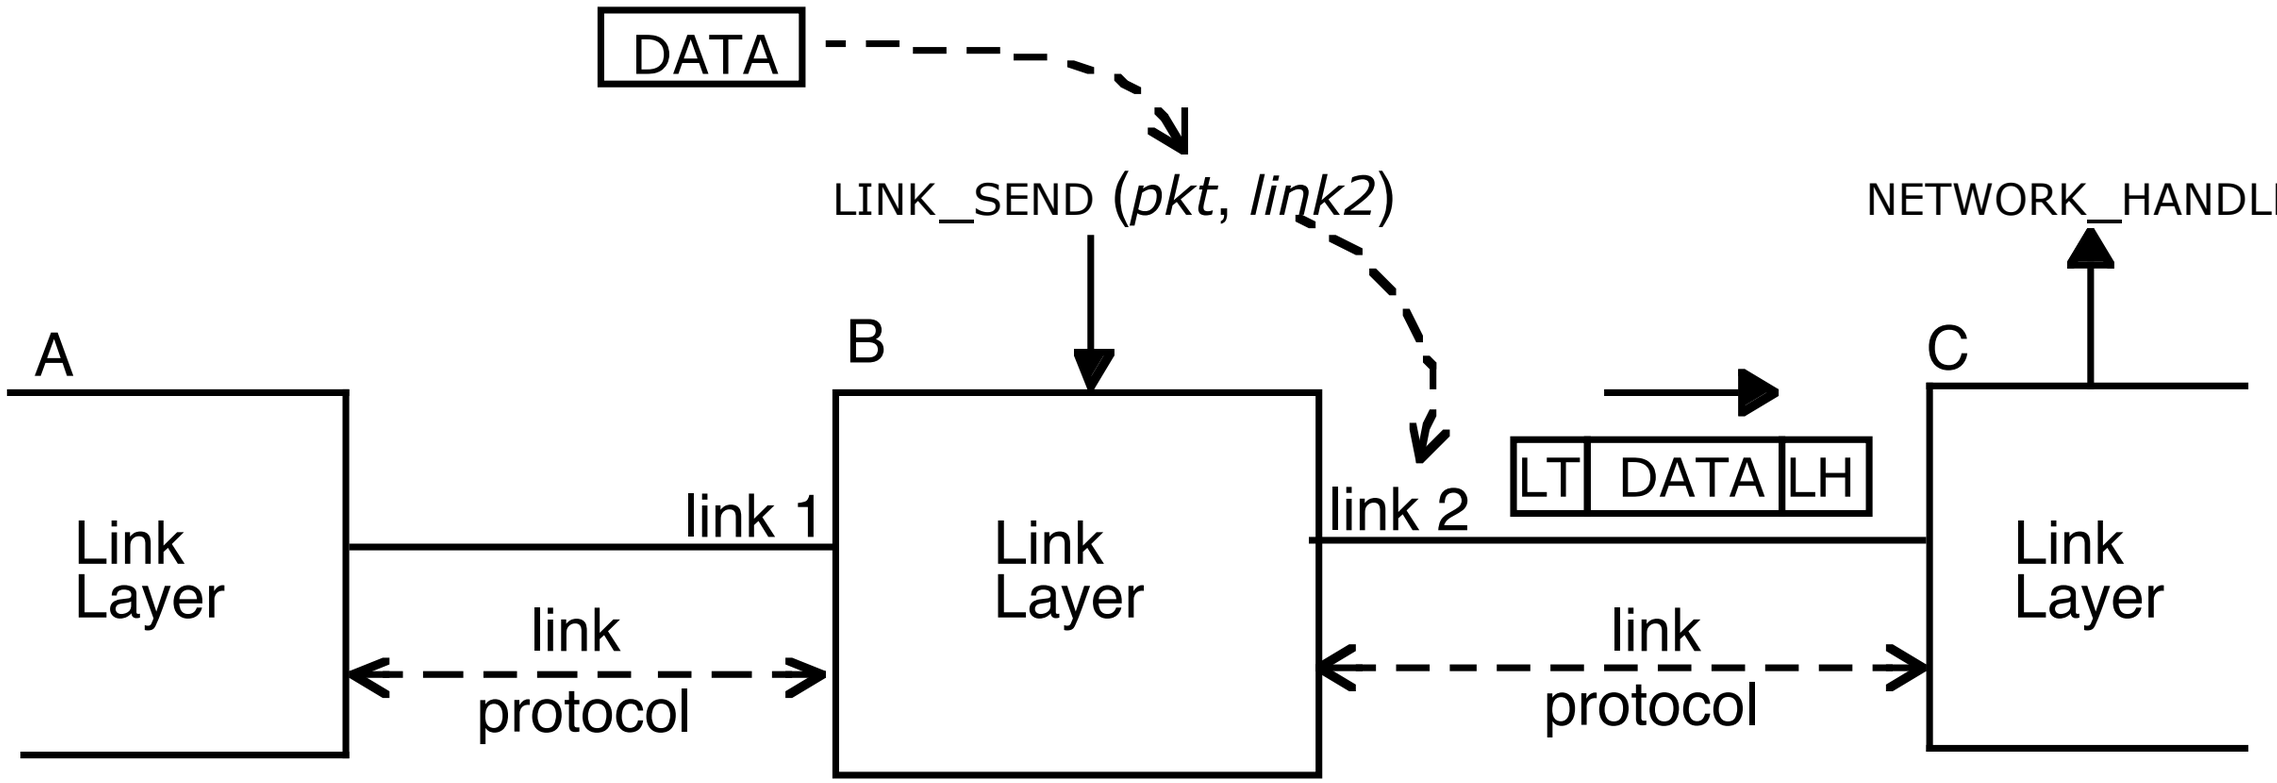

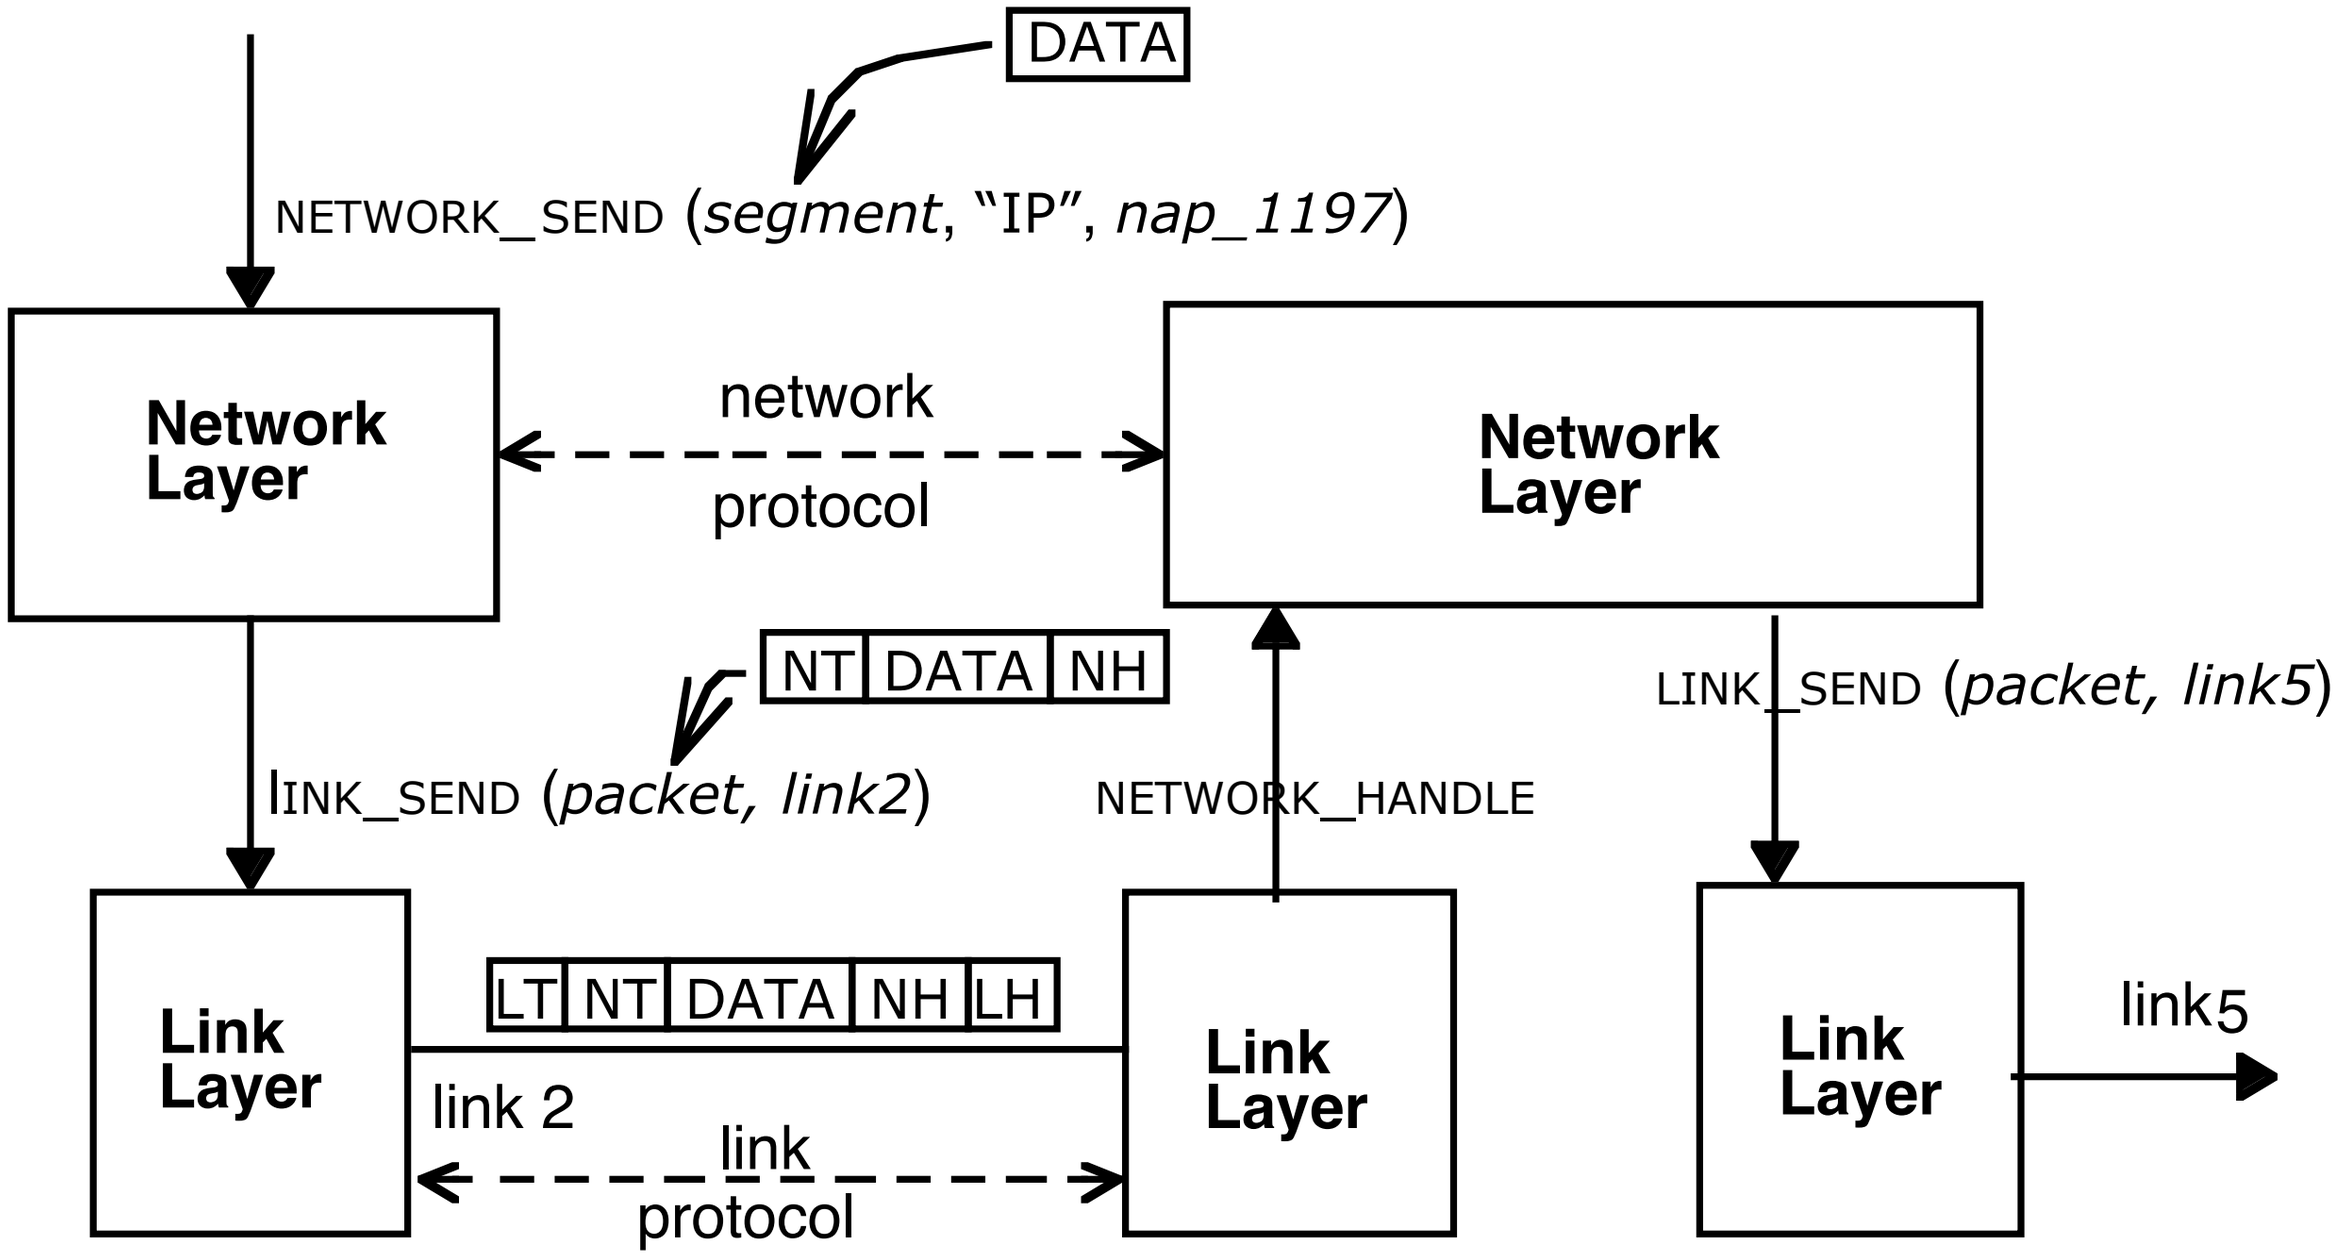

FIRE (7, "Lucifer", *evade*)

DATA

FIRE (7, "Lucifer", *evade*)

End-to-End Layer (RPC)

end-to-end protocol

End-to-End Layer (RPC)

ET | DATA | EH

Network Layer

Network Layer

Network Layer

NT | ET | DATA | EH | NH

Link Layer

Link Layer

Link Layer

Link Layer

LT | NT | ET | DATA | EH | NH | LH

The end-to-end argument

The application knows best.

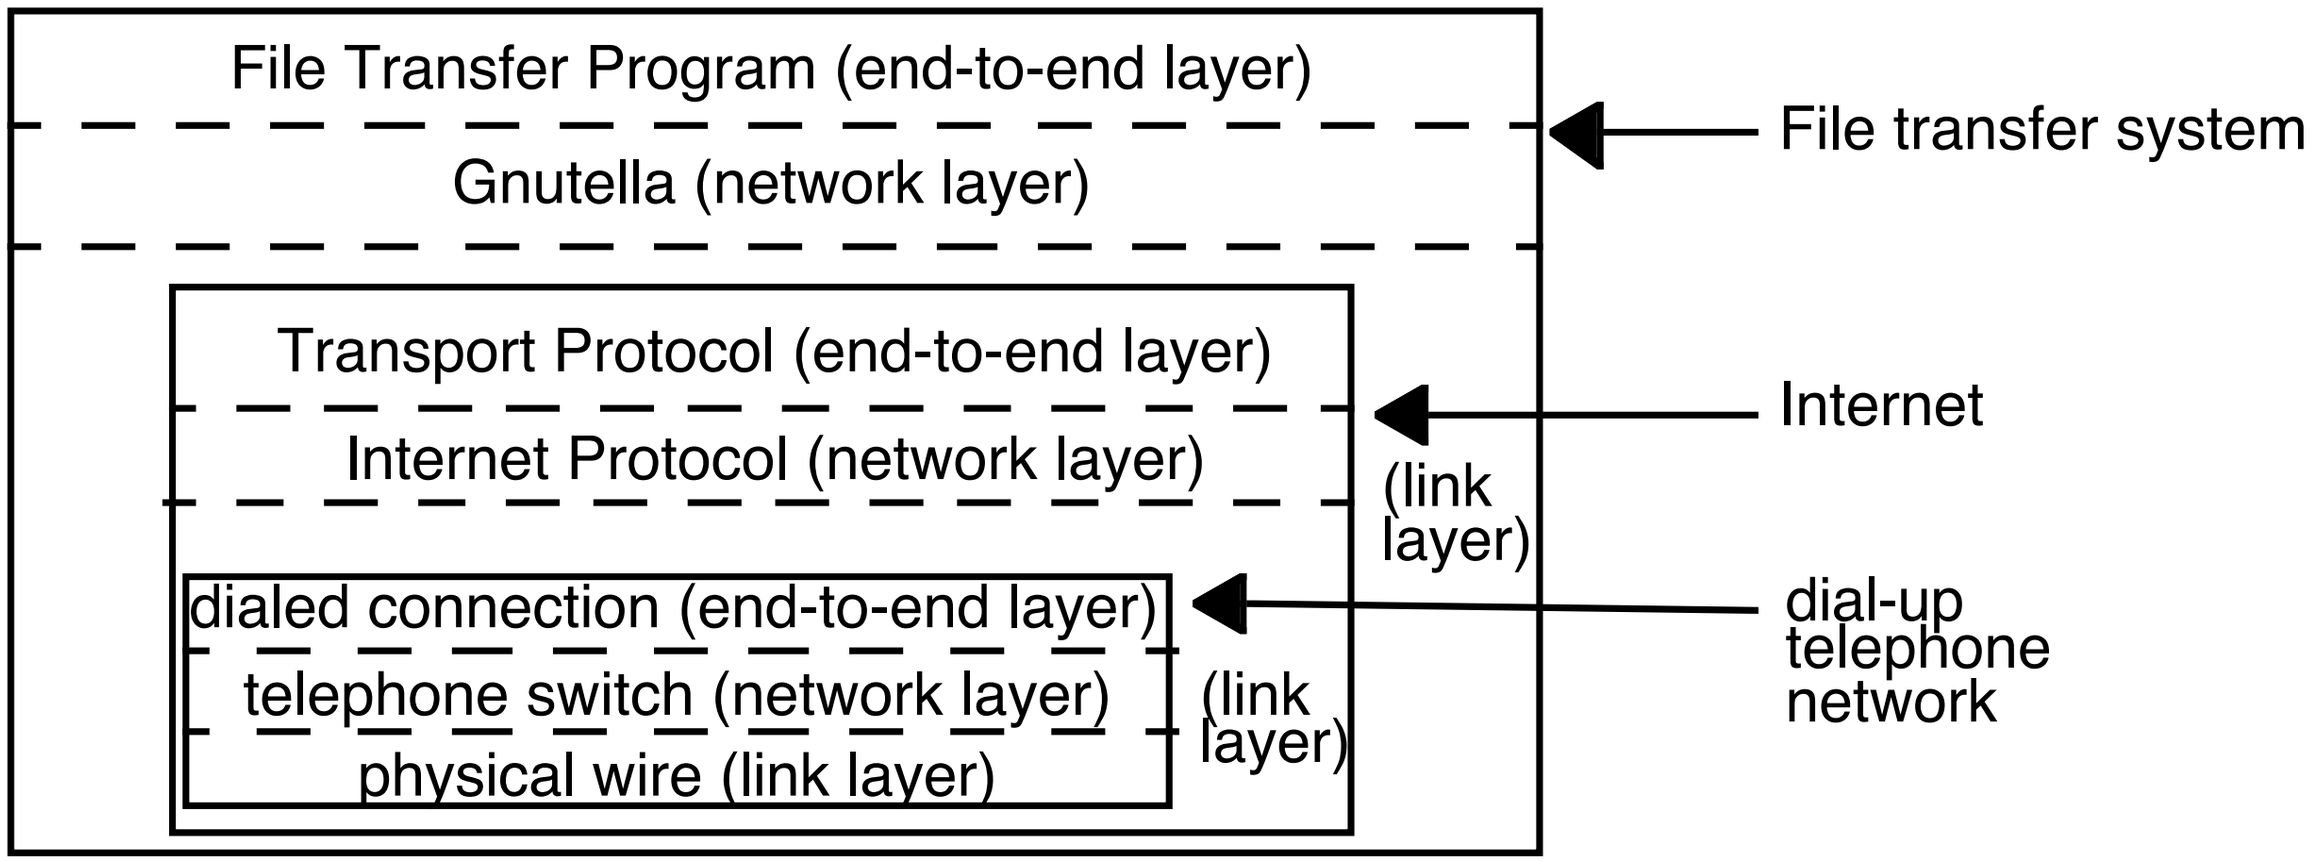

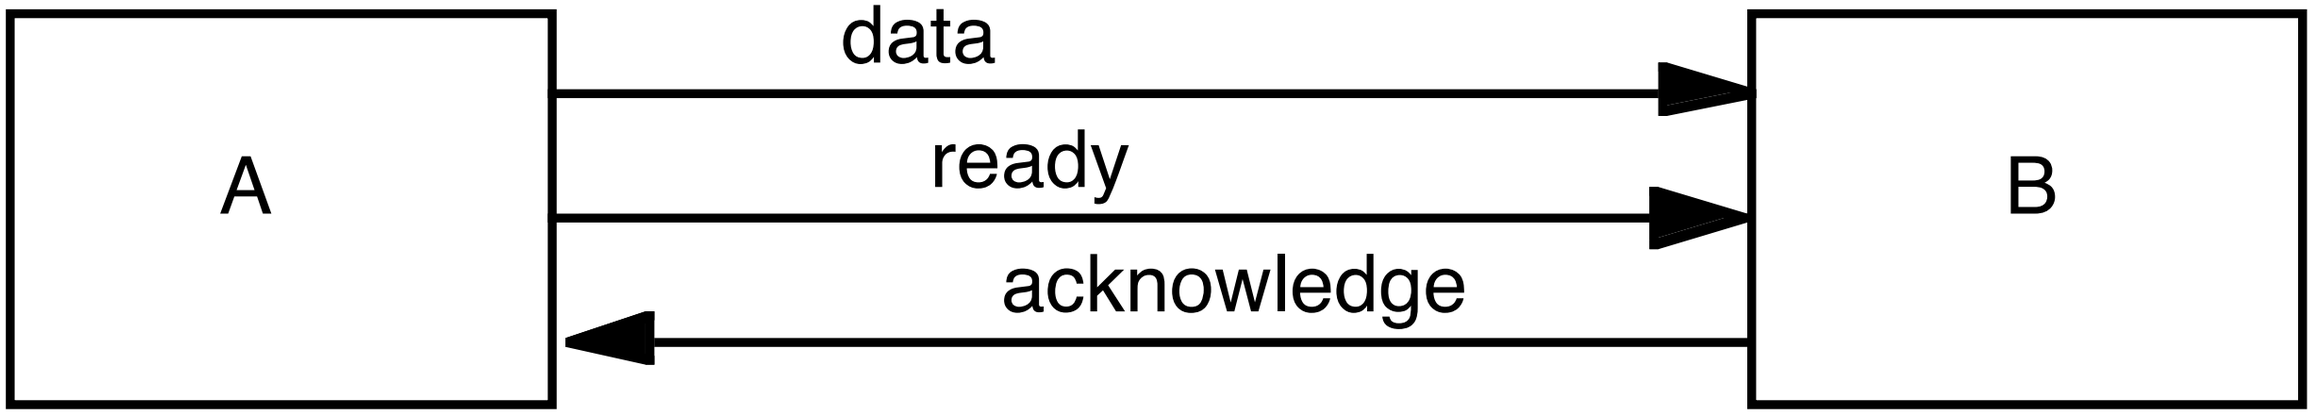

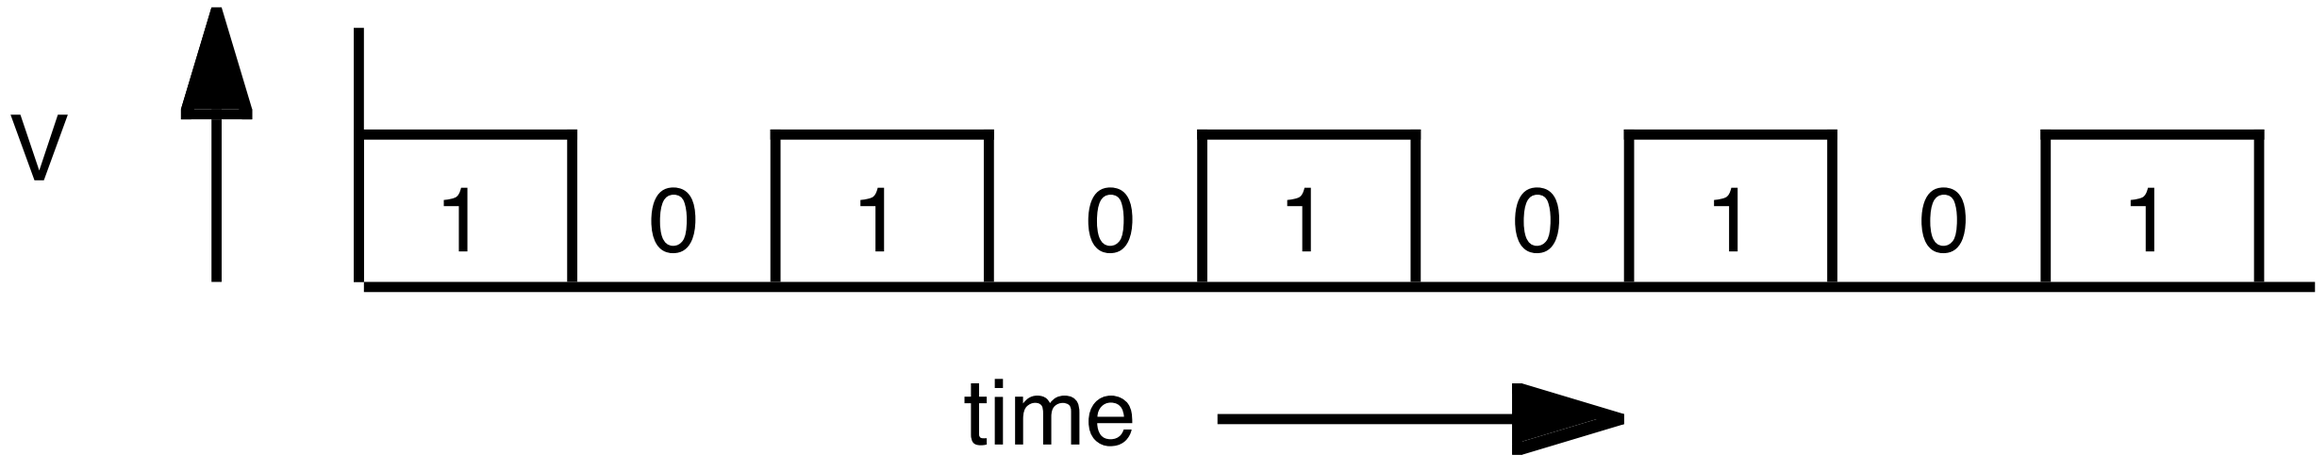

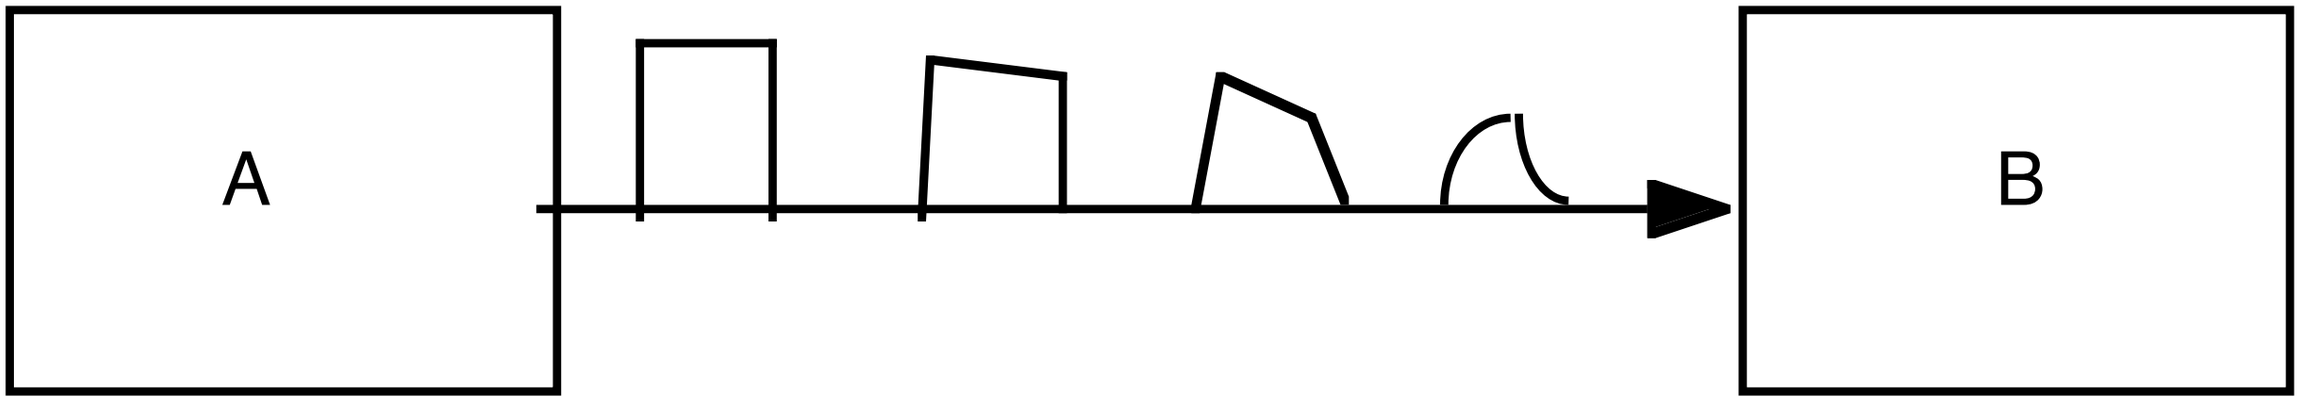

```
procedure FRAME_TO_BIT (frame_data, length)
    ones_in_a_row = 0
    for i from 1 to length do                  // First send frame contents
        SEND_BIT (frame_data[i]);
        if frame_data[i] = 1 then
            ones_in_a_row ← ones_in_a_row + 1;
            if ones_in_a_row = 6 then
                SEND_BIT (0);                  // Stuff a zero so that data doesn't
                ones_in_a_row ← 0;             // look like a framing marker
    else
        ones_in_a_row ← 0;
for i from 1 to 7 do                           // Now send framing marker.
    SEND_BIT (1)
```

```
procedure BIT_TO_FRAME (rcvd_bit)
    ones_in_a_row integer initially 0
    if ones_in_a_row < 6 then
        bits_in_frame ← bits_in_frame + 1
        frame_data[bits_in_frame] ← rcvd_bit
        if rcvd_bit = 1 then ones_in_a_row ← ones_in_a_row + 1
        else ones_in_a_row ← 0
    else                          // This may be a seventh one-bit in a row, check it out.
        if rcvd_bit = 0 then
            ones_in_a_row ← 0                    // Stuffed bit, don't use it.
        else                                     // This is the end-of-frame marker
            LINK_RECEIVE (frame_data, (bits_in_frame - 6), link_id)
            bits_in_frame ← 0
            ones_in_a_row ← 0
```

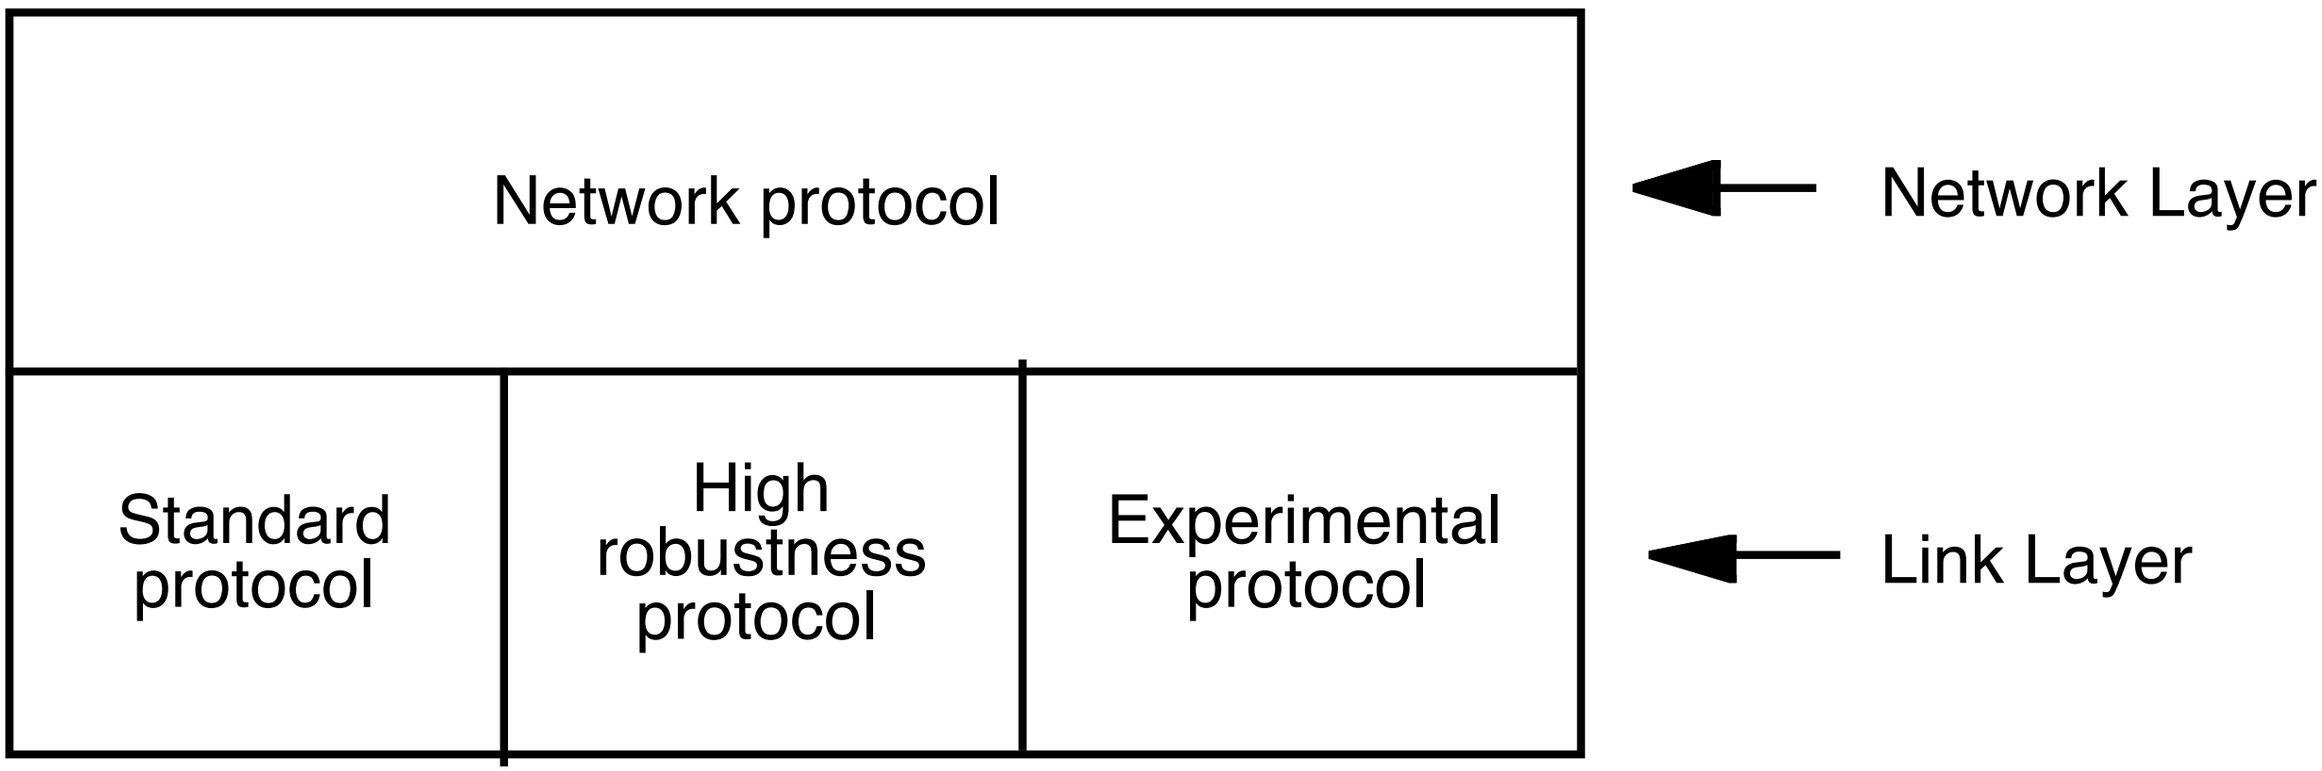

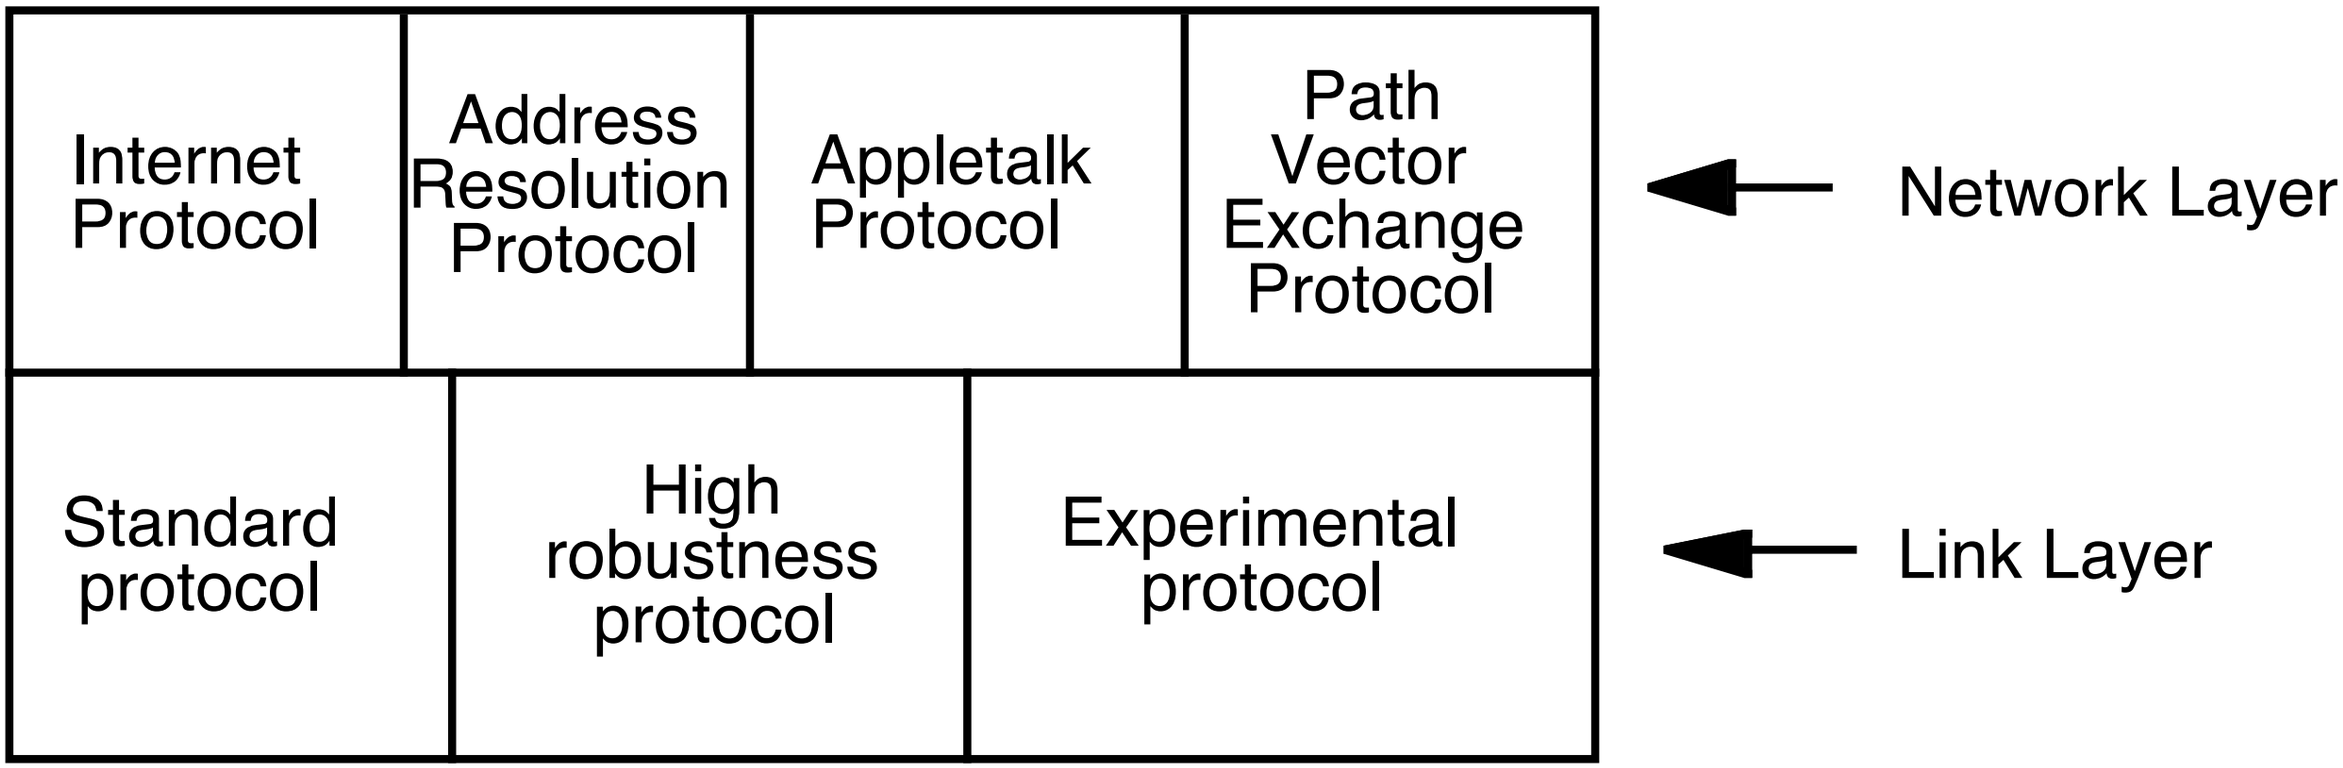

```
structure frame
    structure checked_contents
        bit_string net_protocol                    // multiplexing parameter
        bit_string payload                         // payload data
    bit_string checksum

procedure LINK_SEND (data_buffer, link_identifier, link_protocol, network_protocol)
    frame instance outgoing_frame
    outgoing_frame.checked_contents.payload ← data_buffer
    outgoing_frame.checked_contents.net_protocol ← data_buffer.network_protocol
    frame_length ← LENGTH (data_buffer) + header_length
    outgoing_frame.checksum ← CHECKSUM (frame.checked_contents, frame_length)
    sendproc ← link_protocol[that_link.protocol]          // Select link protocol.
    sendproc (outgoing_frame, frame_length, link_identifier)   // Send frame.
```

```
procedure LINK_RECEIVE (received_frame, length, link_id)
    frame instance received_frame
    if CHECKSUM (received_frame.checked_contents, length) =
                                                    received_frame.checksum
        then                                        // Pass good packets up to next layer.
        good_frame_count ← good_frame_count + 1;
        GIVE_TO_NETWORK_HANDLER (received_frame.checked_contents.payload,
                                 received_frame.checked_contents.net_protocol);
    else bad_frame_count ← bad_frame_count + 1     // Just count damaged frame.

// Each network layer protocol handler must call SET_HANDLER before the first packet
// for that protocol arrives…

procedure SET_HANDLER (handler_procedure, handler_protocol)
    net_handler[handler_protocol] ← handler_procedure

procedure GIVE_TO_NETWORK_HANDLER (received_packet, network_protocol)
    handler ← net_handler[network_protocol]
    if (handler ≠ NULL) call handler(received_packet, network_protocol)
    else unexpected_protocol_count ← unexpected_protocol_count + 1
```

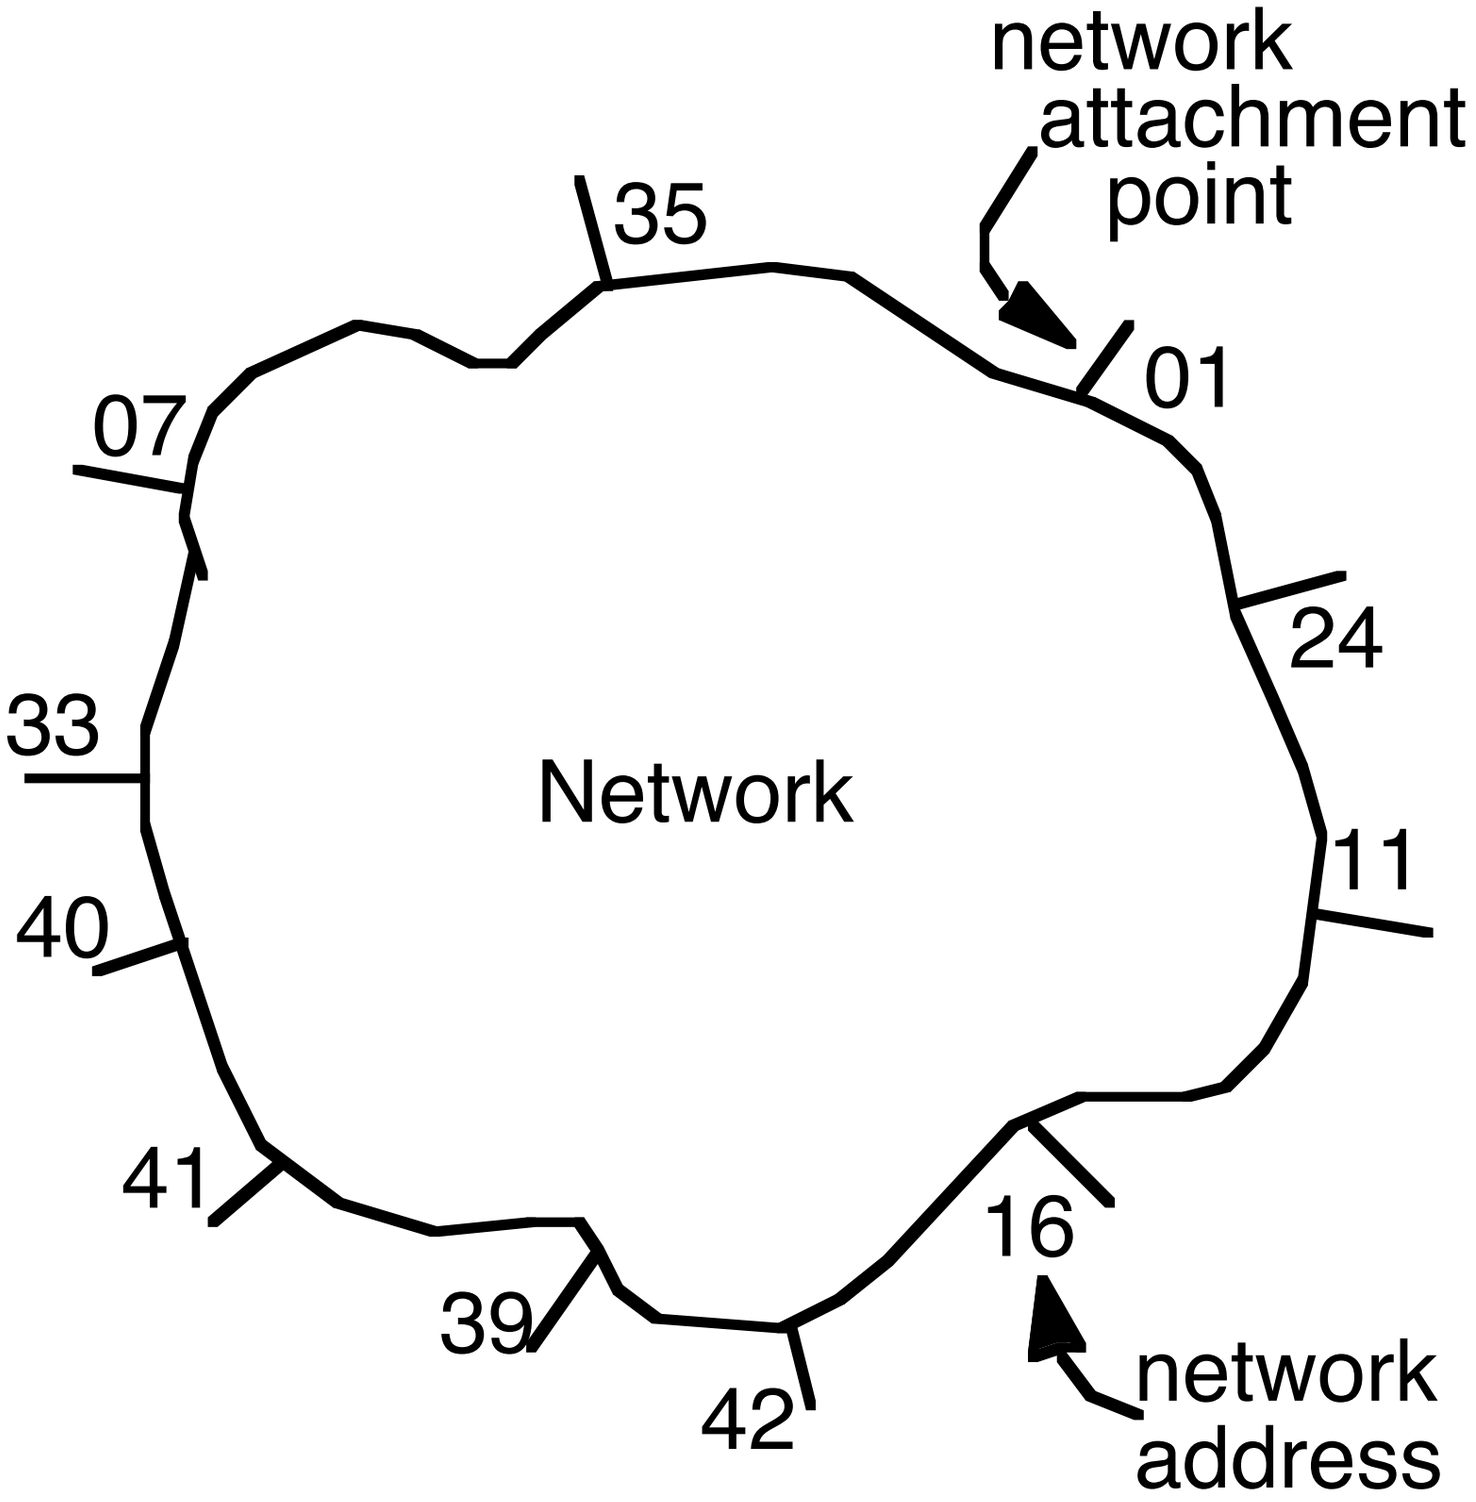

```
    structure packet
        bit_string source
        bit_string destination
        bit_string end_protocol
        bit_string payload

procedure NETWORK_SEND (segment_buffer, destination,
                                        network_protocol, end_protocol)
    packet instance outgoing_packet
    outgoing_packet.payload ← segment_buffer
    outgoing_packet.end_protocol ← end_protocol
    outgoing_packet.source ← MY_NETWORK_ADDRESS
    outgoing_packet.destination ← destination
    NETWORK_HANDLE (outgoing_packet, net_protocol)
```

```
procedure NETWORK_HANDLE (net_packet, net_protocol)
    packet instance net_packet
    if net_packet.destination ≠ MY_NETWORK_ADDRESS then
        next_hop ← LOOKUP (net_packet.destination, forwarding_table)
        LINK_SEND (net_packet, next_hop, link_protocol, net_protocol)
else
    GIVE_TO_END_LAYER (net_packet.payload,
                       net_packet.end_protocol, net_packet.source)
```

Segment presented to the network layer →

DATA

Packet presented to the link layer →

source & destination	end protocol	DATA

Frame appearing on the link →

frame mark	network protocol	source & destination	end protocol	DATA	check sum	frame mark

Example →

1111111	IP	41 —> 24	RPC	“Fire”	97142 55316	1111111

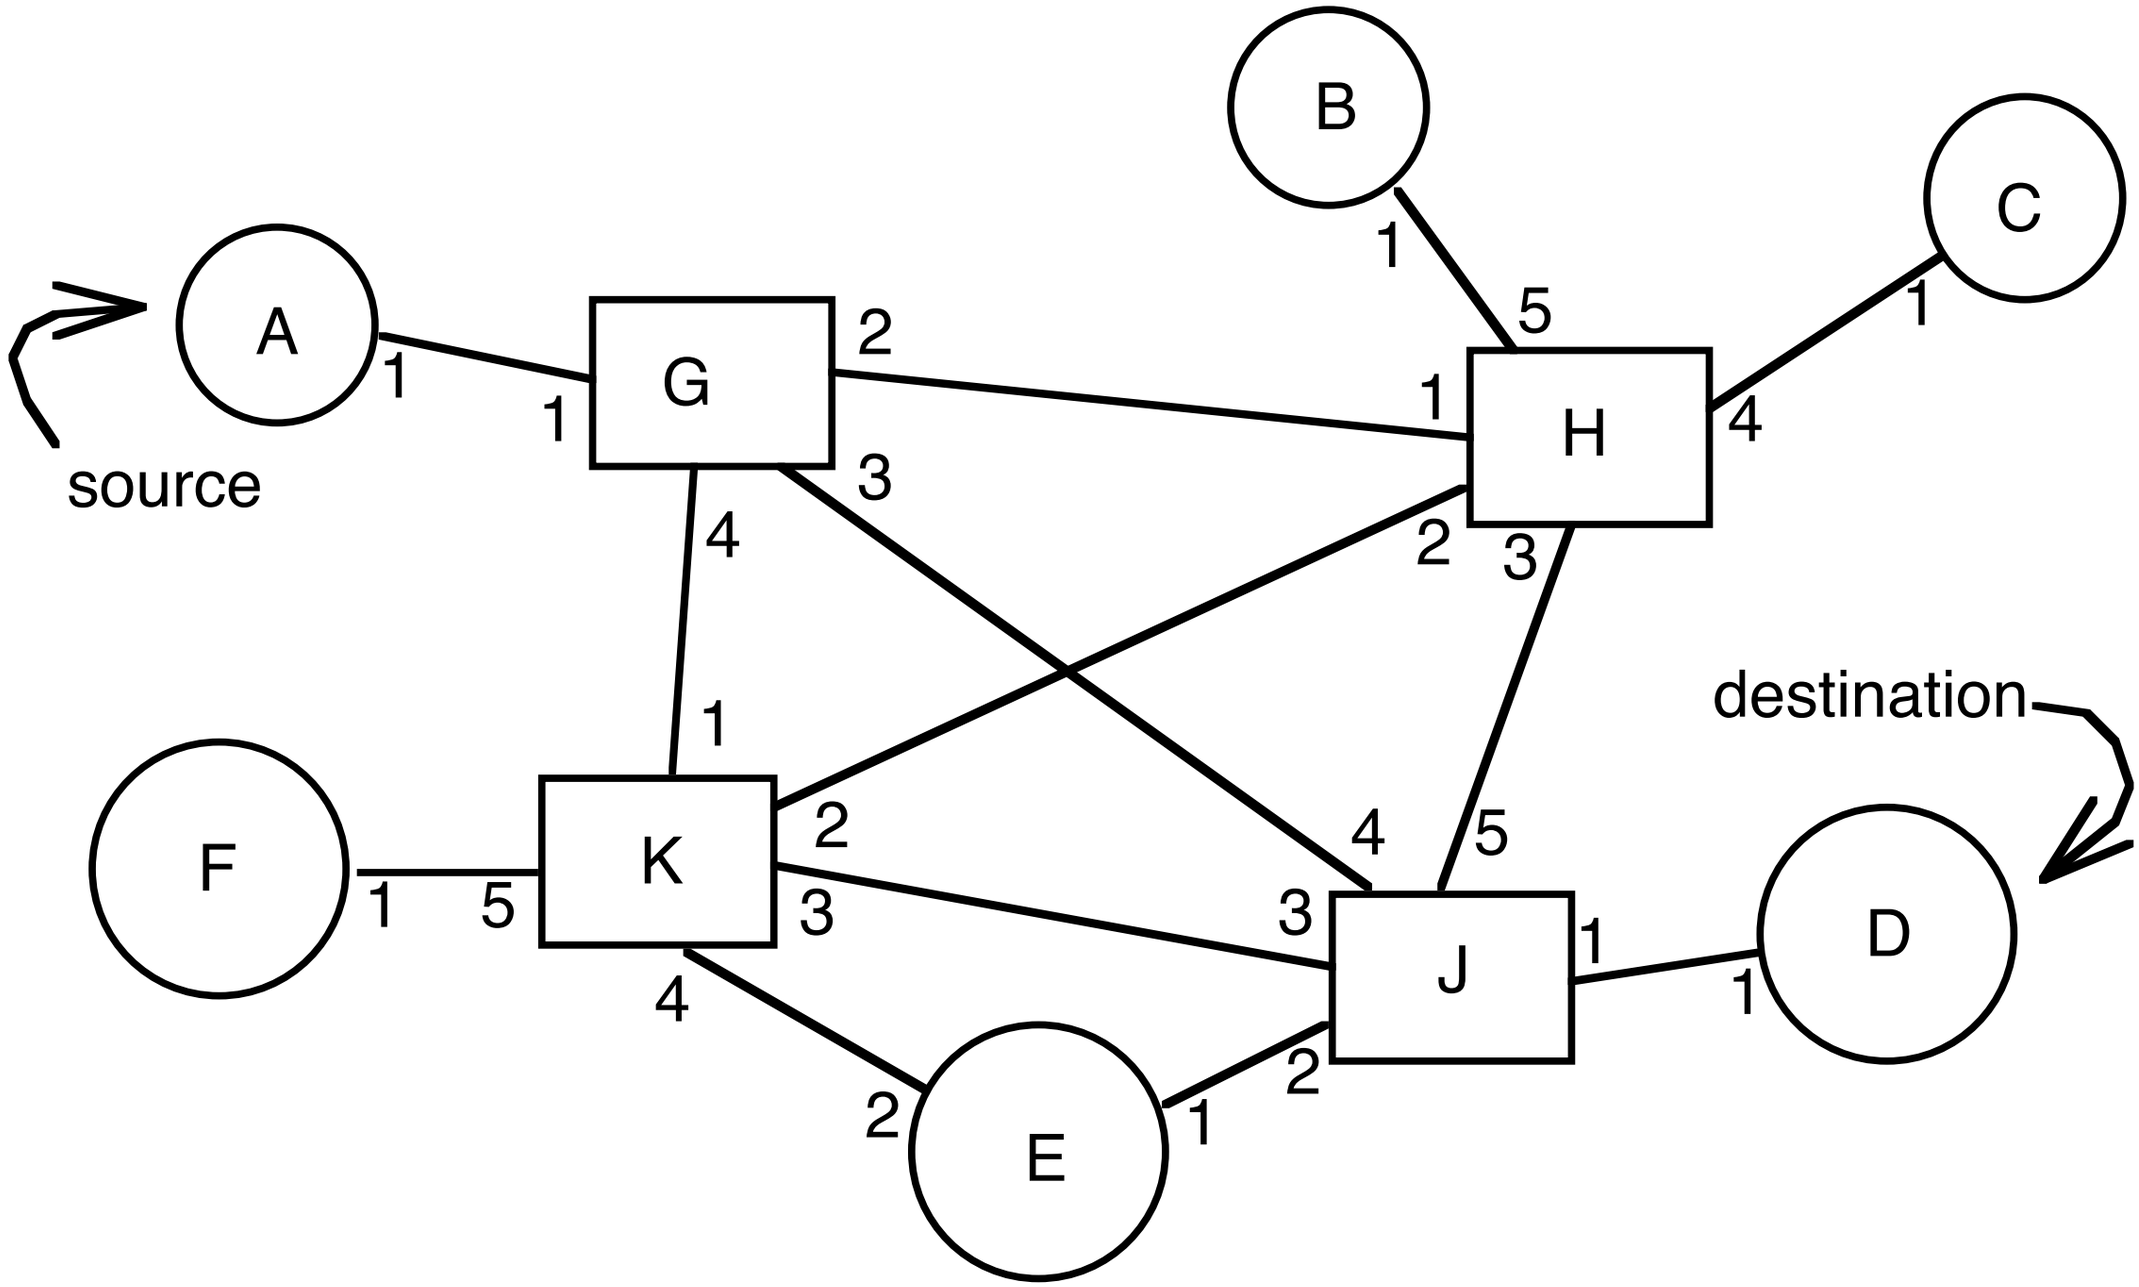

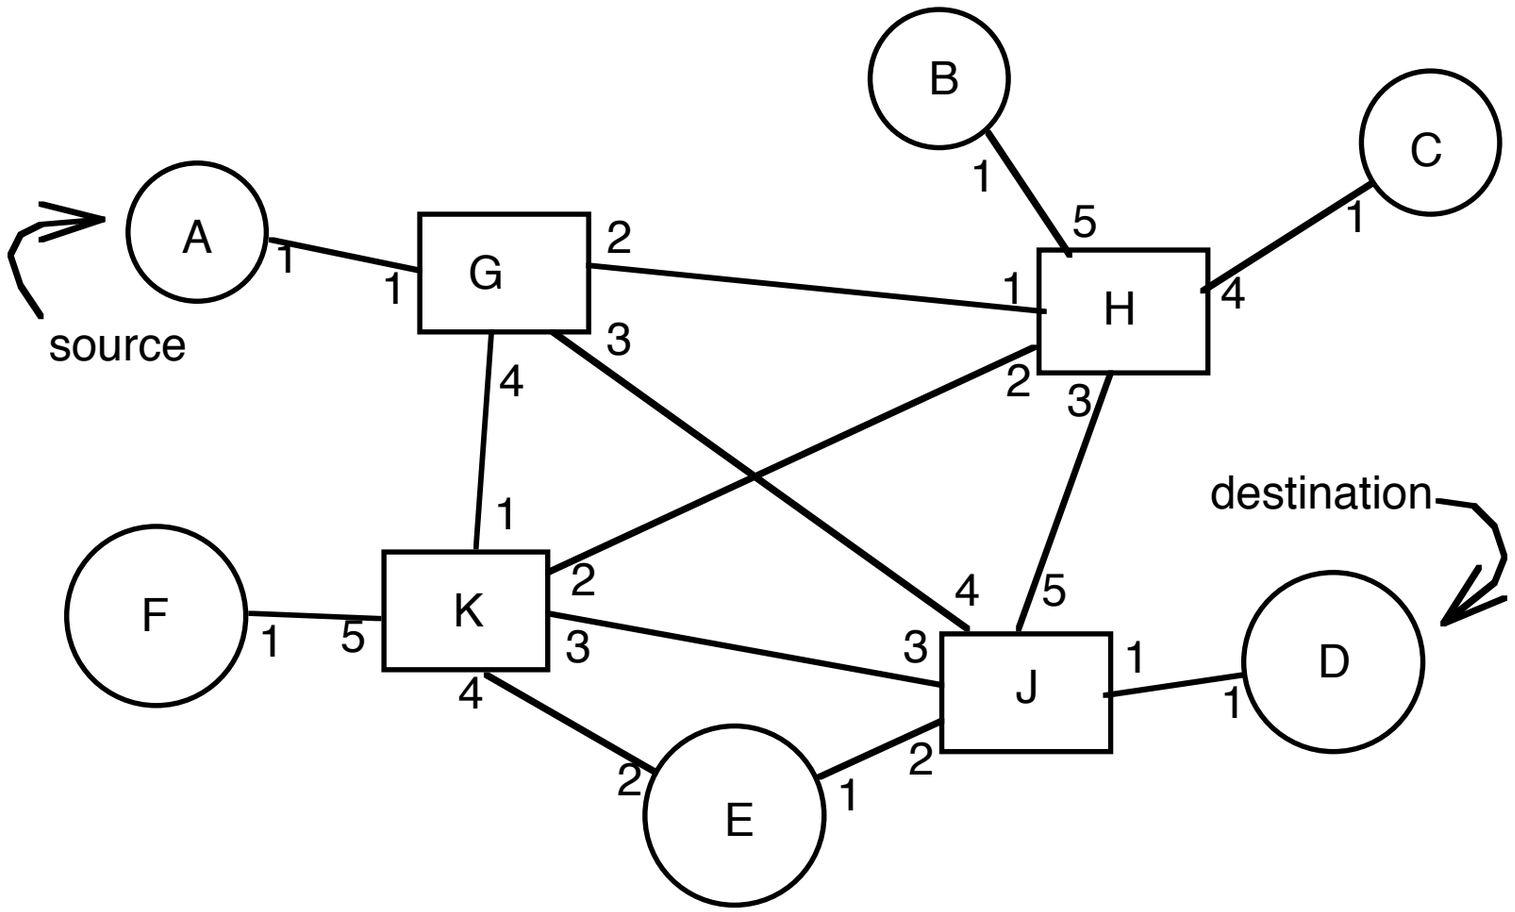

destination	link
A	end-layer
all other	1

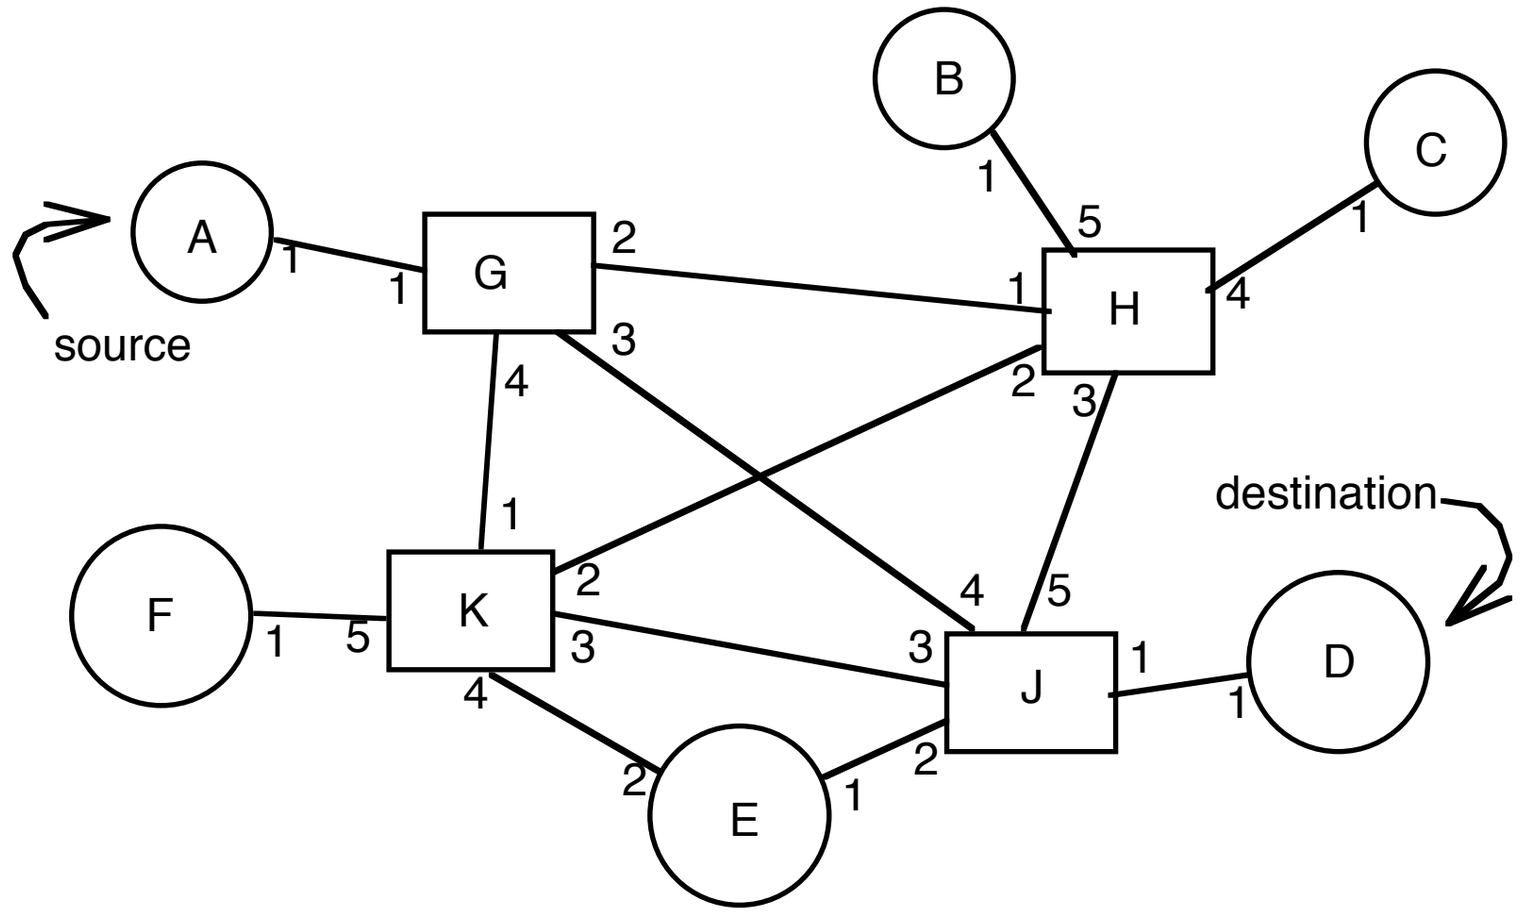

destination	link
A	1
B	2
C	2
D	3
E	4
F	4
G	end-layer
H	2
J	3
K	4

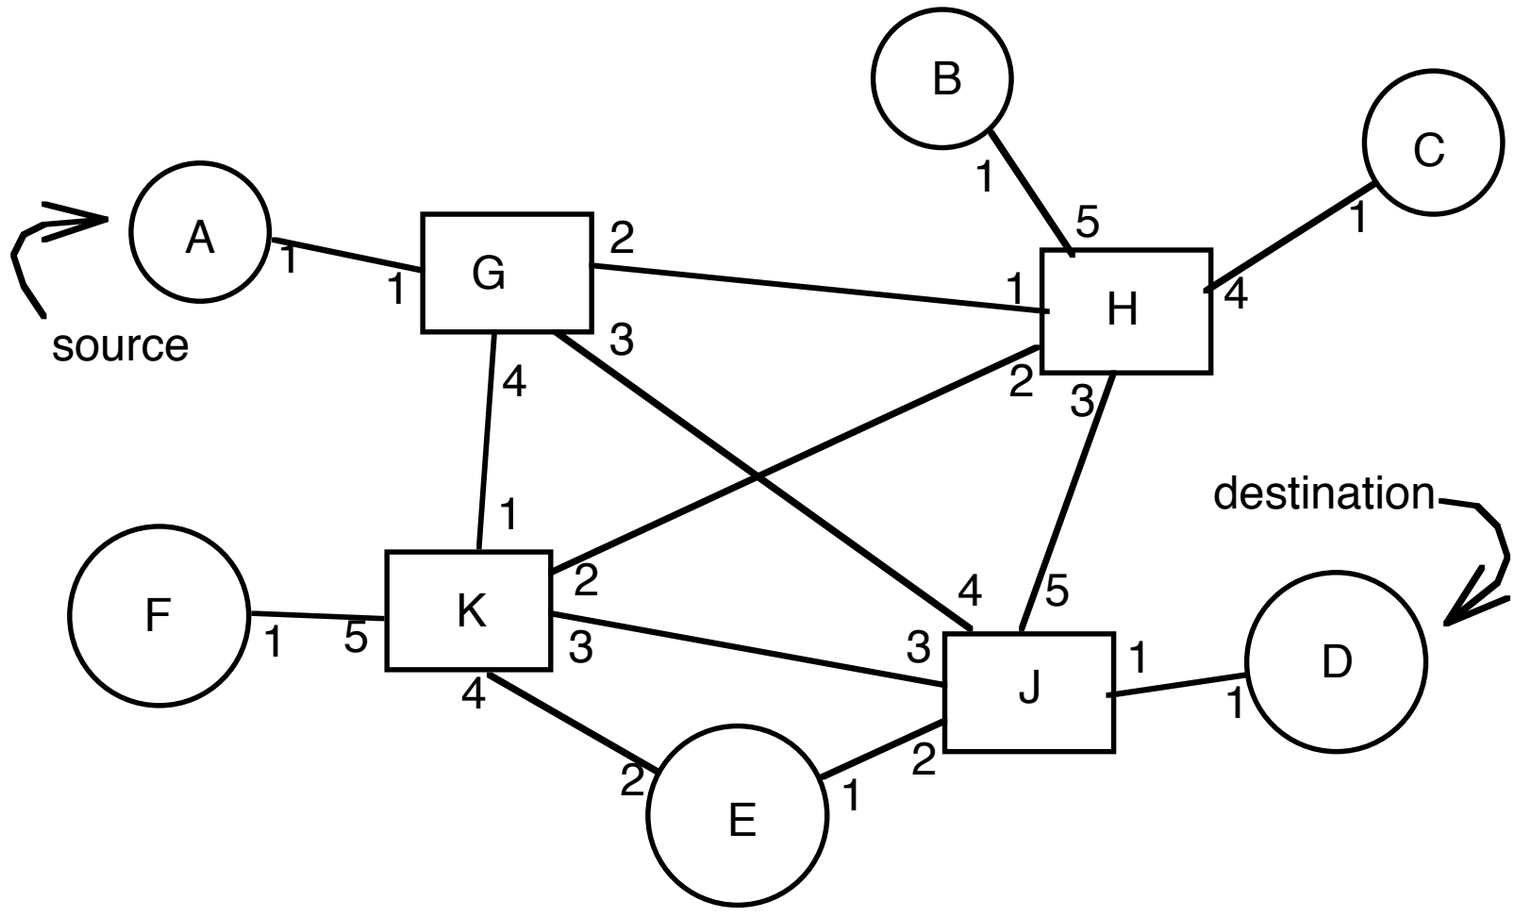

to	path
G	< >

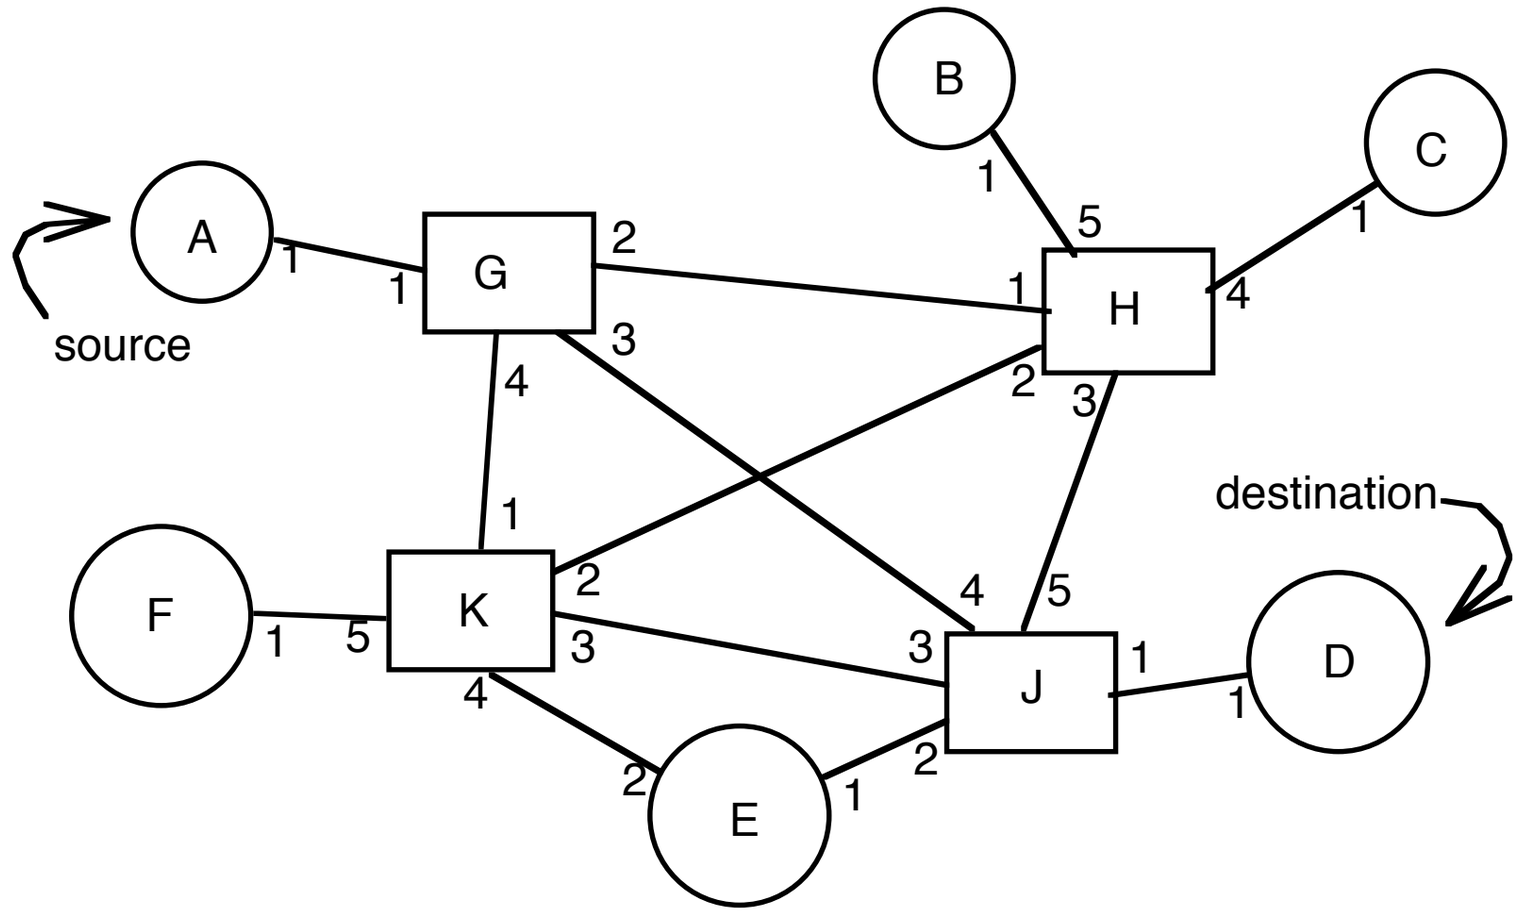

From A,
via link 1

to	path
A	< >

From H,
via link 2:

to	path
H	< >

From J,
via link 3:

to	path
J	< >

From K,
via link 4:

to	path
K	< >

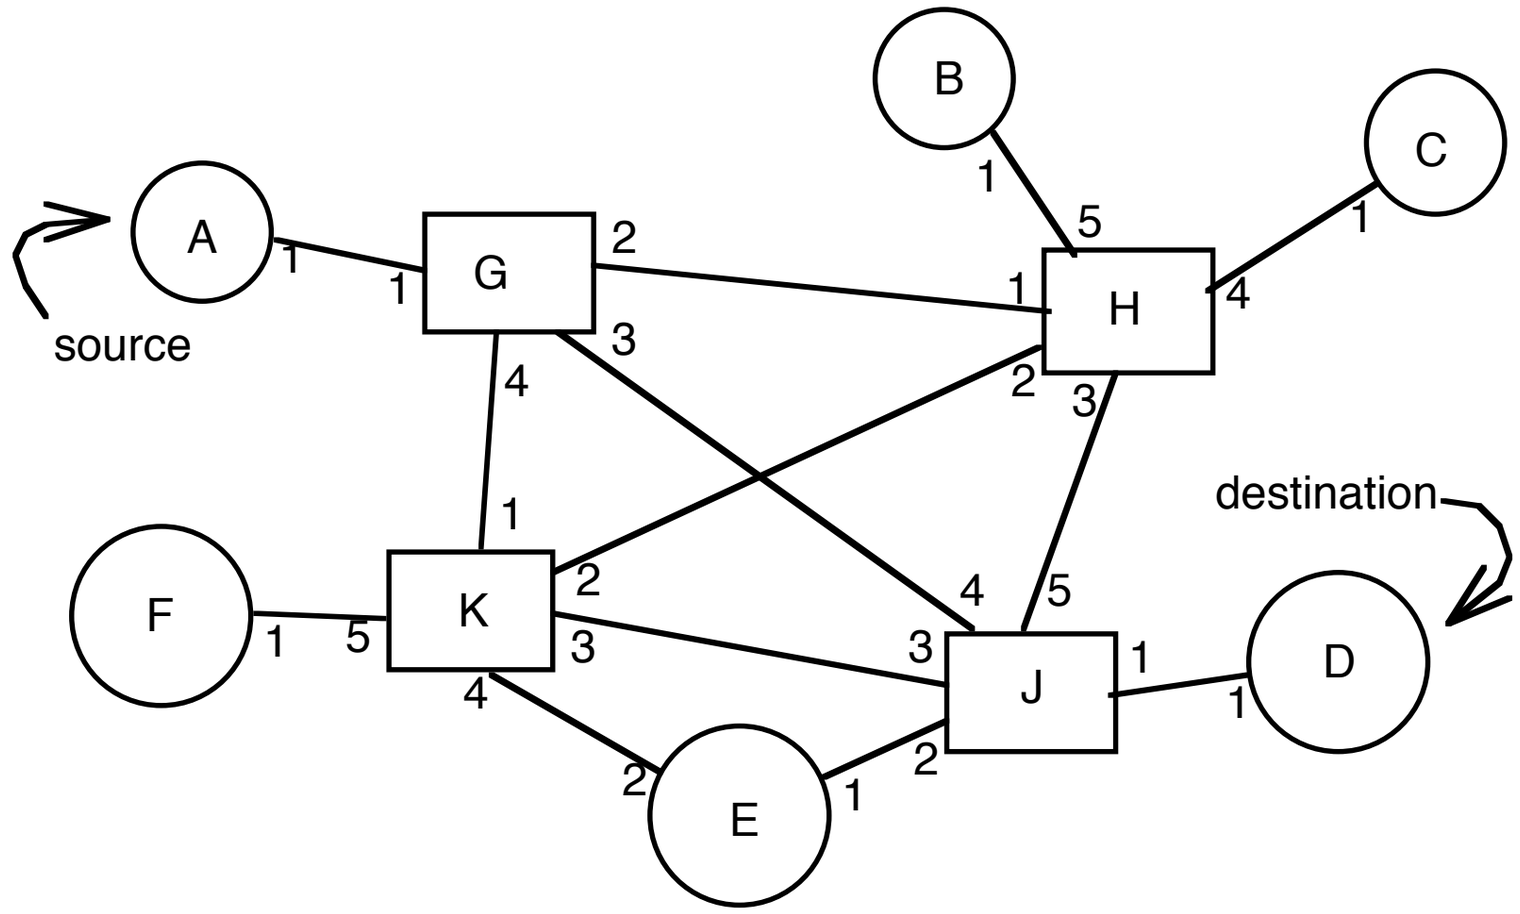

path vector

to	path
A	<A>
G	< >
H	<H>
J	<J>
K	<K>

forwarding table

to	link
A	1
G	end-layer
H	2
J	3
K	4

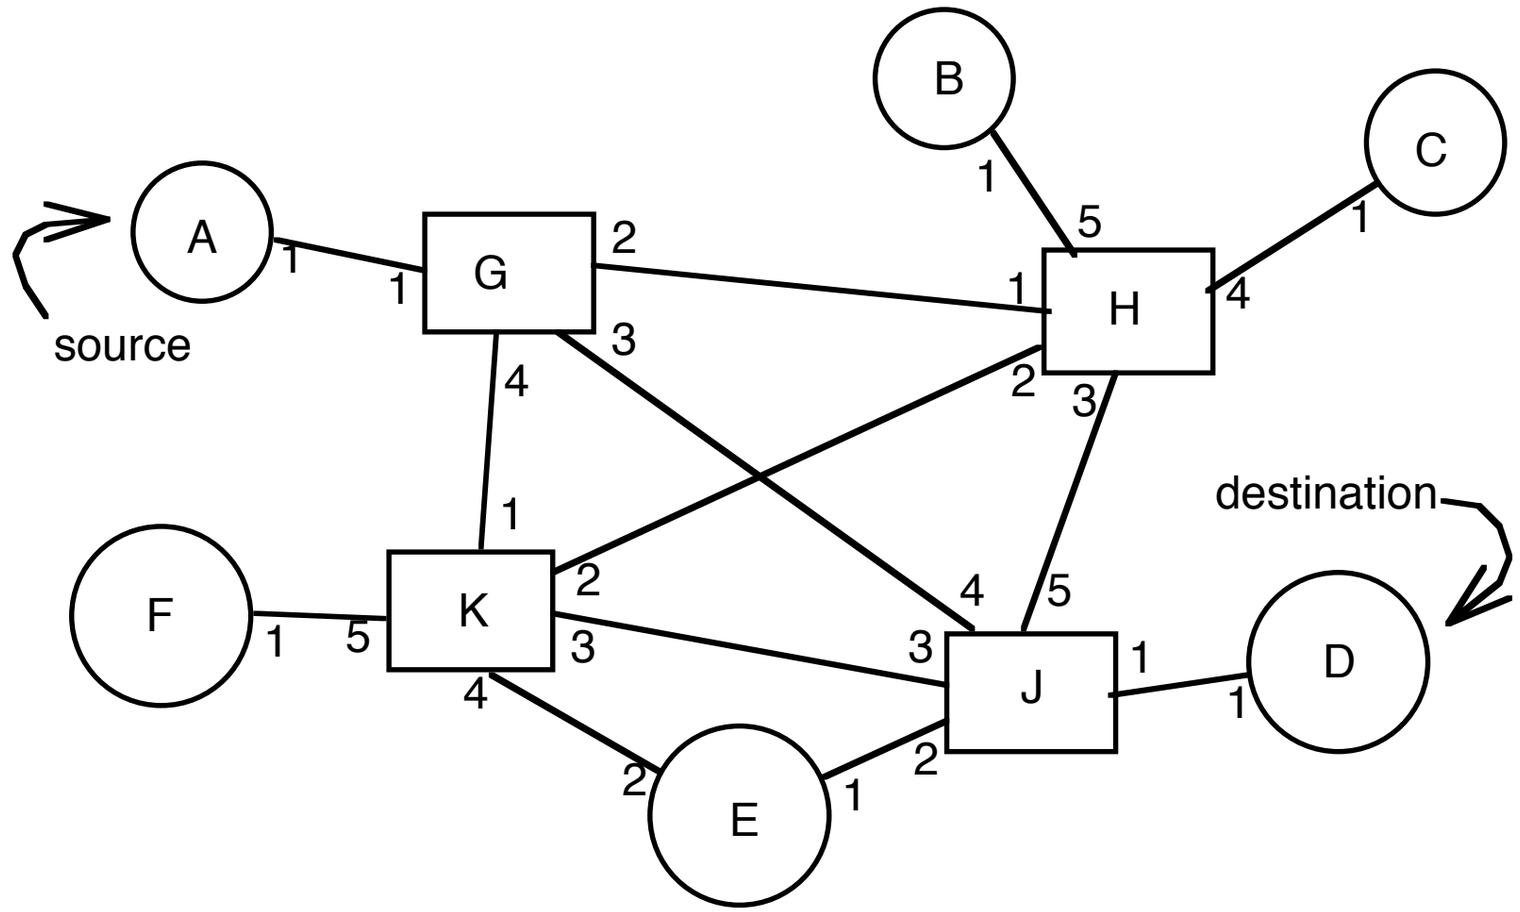

From A,
via link 1

to	path
A	< >
G	<G>

From H,
via link 2:

to	path
B	<B>
C	<C>
G	<G>
H	< >
J	<J>
K	<K>

From J,
via link 3:

to	path
D	<D>
E	<E>
G	<G>
H	<H>
J	< >
K	<K>

From K,
via link 4:

to	path
E	<E>
F	<F>
G	<G>
H	<H>
J	<J>
K	< >

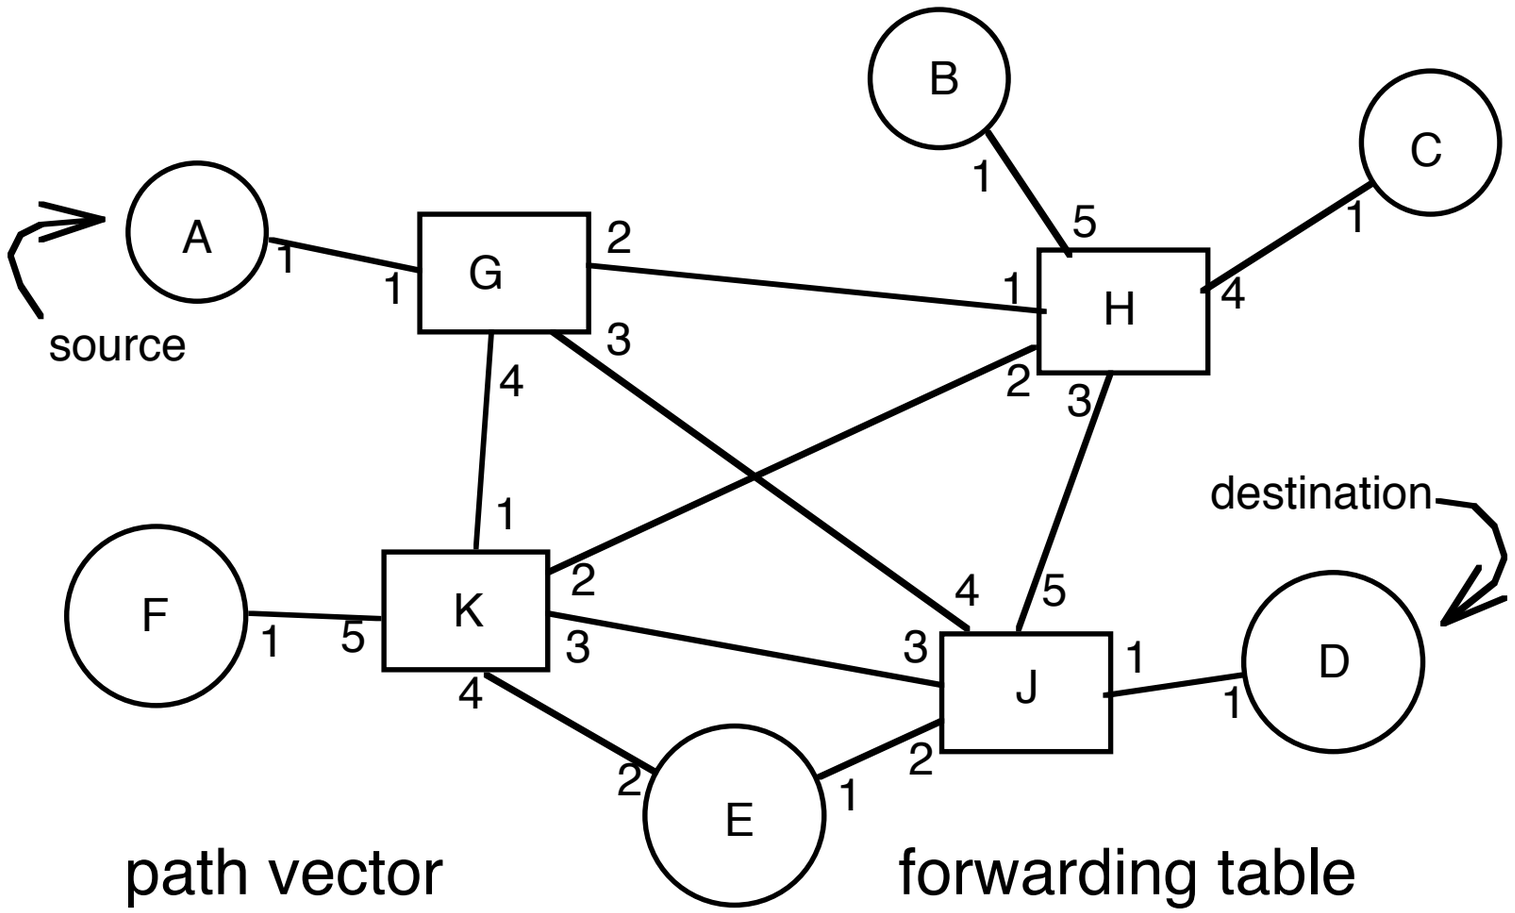

path vector

to	path
A	<A>
B	<H, B>
C	<H, C>
D	<J, D>
E	<J, E>
F	<K, F>
G	< >
H	<H>
J	<J>
K	<K>

forwarding table

to	link
A	1
B	2
C	2
D	3
E	3
F	4
G	end-layer
H	2
J	3
K	4

```
// Maintain routing and forwarding tables.

vector associative array          // vector[d_addr] contains path to destination d_addr
neighbor_vector instance of vector   // A path vector received from some neighbor
my_vector  instance of vector     // My current path vector.
addr associative array            // addr[j] is the address of the network attachment
                                  // point at the other end of link j.
                                  // my_addr is address of my network attachment point.
                                  // A path is a parsable list of addresses, e.g. {a,b,c,d}

procedure main()                  // Initialize, then start advertising.
    SET_TYPE_HANDLER (HANDLE_ADVERTISEMENT, exchange_protocol)
    clear my_vector;              // Listen for advertisements
    do occasionally               // and advertise my paths
        for each j in link_ids do // to all of my neighbors.
            status ← SEND_PATH_VECTOR (j, my_addr, my_vector, exch_protocol)
            if status ≠ 0 then    // If the link was down,
                clear new_vector  // forget about any paths
                FLUSH_AND_REBUILD (j)   // that start with that link.
```

```
procedure HANDLE_ADVERTISEMENT (advt, link_id)      // Called when an advt arrives.
    addr[link_id] ← GET_SOURCE (advt)               // Extract neighbor's address
    neighbor_vector ← GET_PATH_VECTOR (advt)        //    and path vector.
    for each neighbor_vector.d_addr do              // Look for better paths.
        new_path ←{addr[link_id], neighbor_vector[d_addr]}      // Build potential path.
        if my_addr is not in new_path then          // Skip it if I'm in it.
            if my_vector[d_addr] = NULL) then       // Is it a new destination?
                my_vector[d_addr] ← new_path        // Yes, add this one.
            else                                    // Not new; if better, use it.
                my_vector[d_addr] ← SELECT_PATH (new_path, my_vector[d_addr])
FLUSH_AND_REBUILD (link_id)
```

```
procedure SELECT_PATH (new, old)          // Decide if new path is better than old one.
    if first_hop(new) = first_hop(old) then return new   // Update any path we were
                                                         // already using.
    else if length(new) ≥ length(old) then return old    // We know a shorter path, keep
    else return new                                      // OK, the new one looks better.

procedure FLUSH_AND_REBUILD (link_id)     // Flush out stale paths from this neighbor.
    for each my_vector,d_addr
        if first_hop(my_vector[d_addr]) = addr[link_id] and new_vector[d_addr] = NULL
            then
            delete my_vector[d_addr]      // Delete paths that are not still advertised.
    REBUILD_FORWARDING_TABLE (my_vector, addr)          // Pass info to forwarder.
```

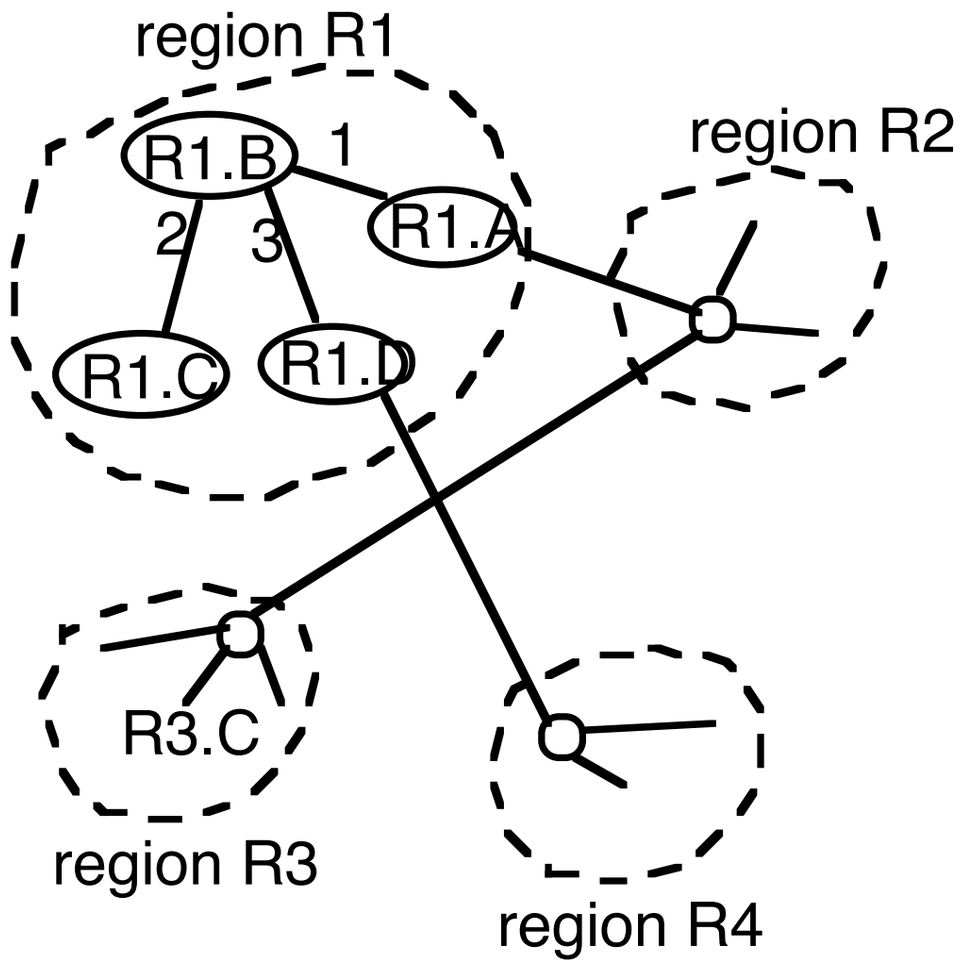

forwarding table in R1.B

region forwarding section

to	link
R1	local
R2	1
R3	1
R4	3

local forwarding section

to	link
R1.A	1
R1.B	end-layer
R1.C	2
R1.D	3

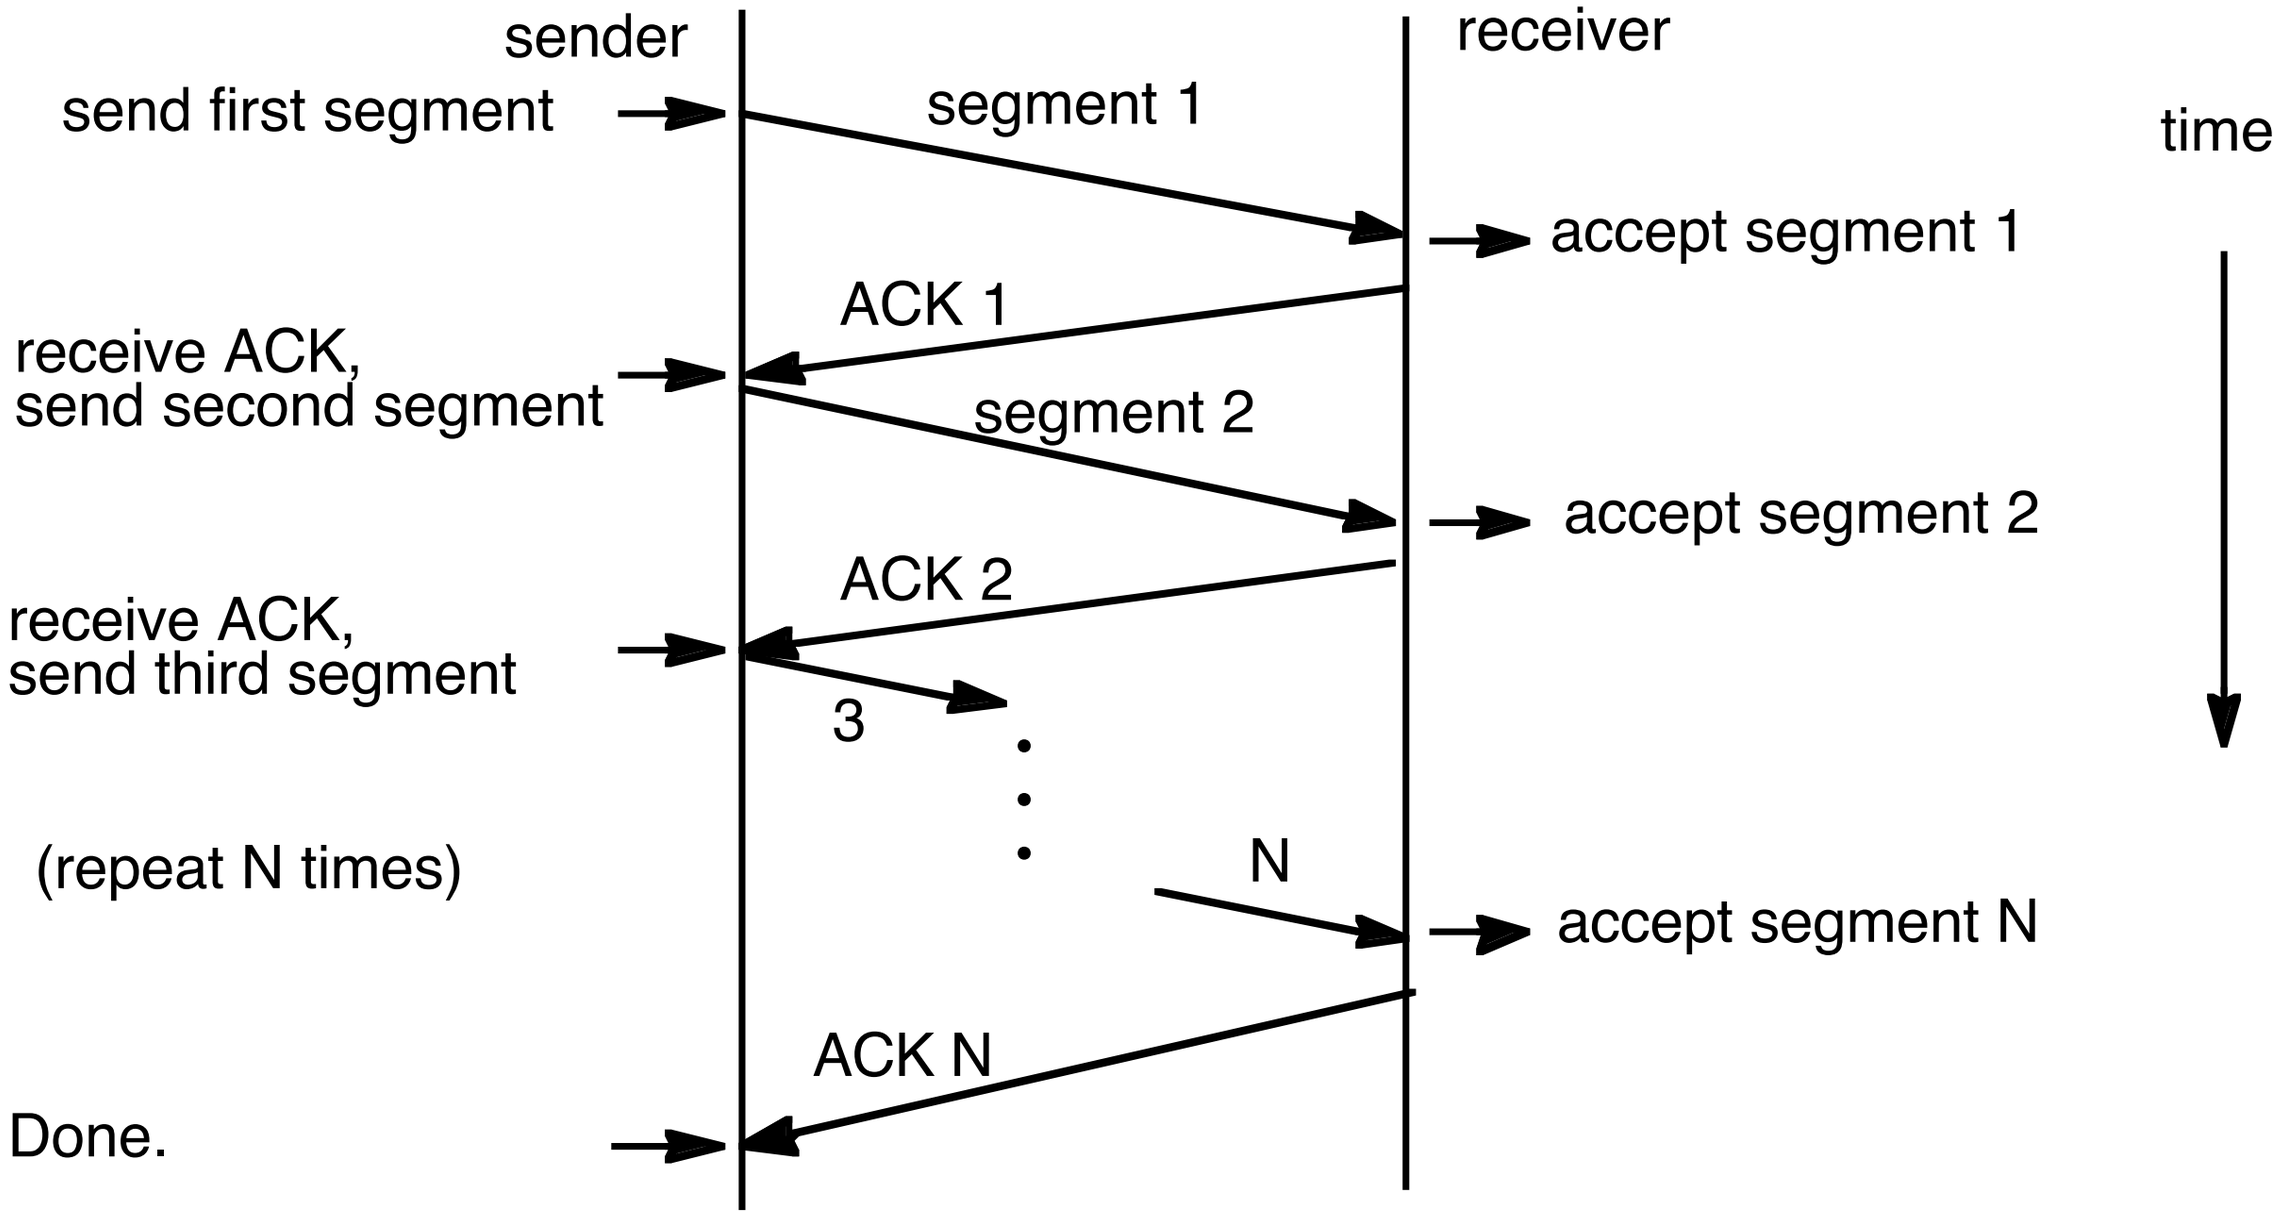

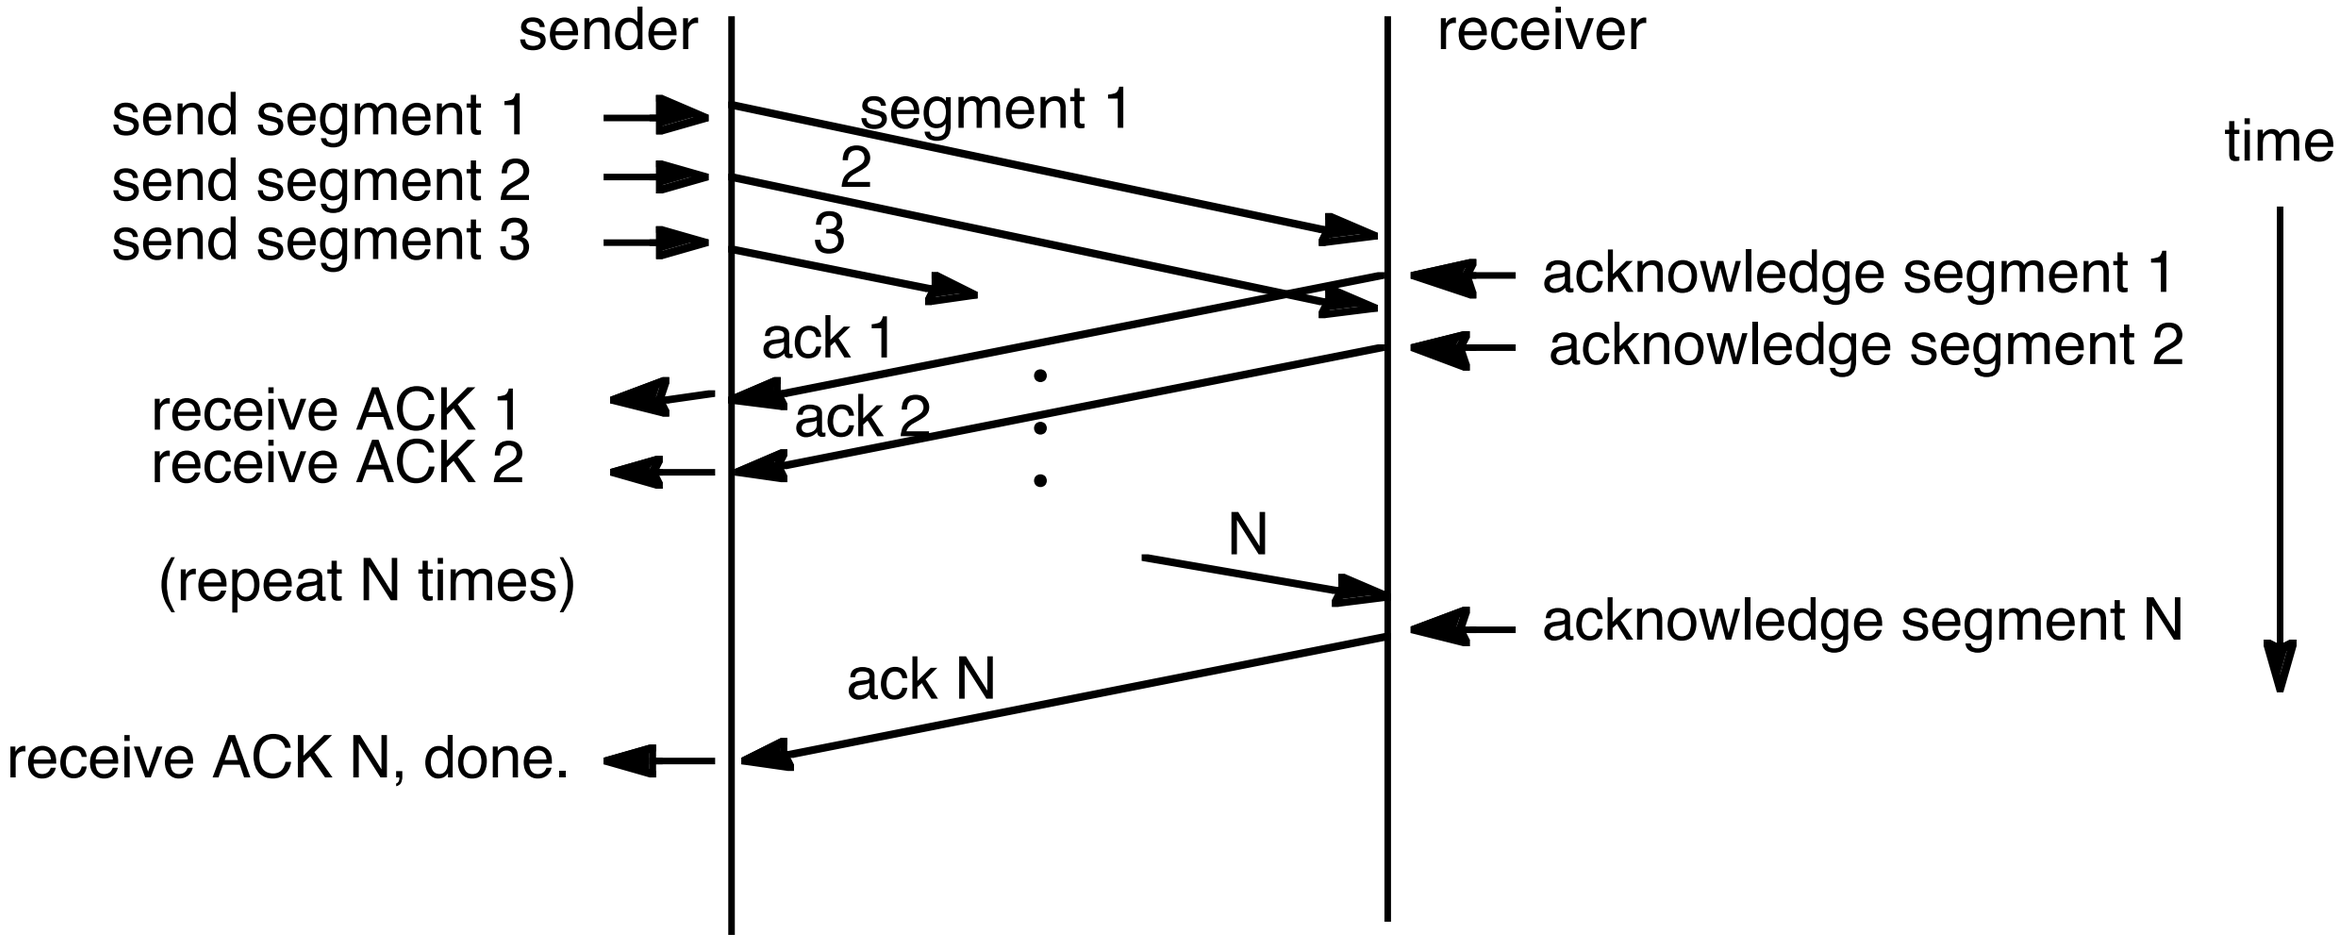

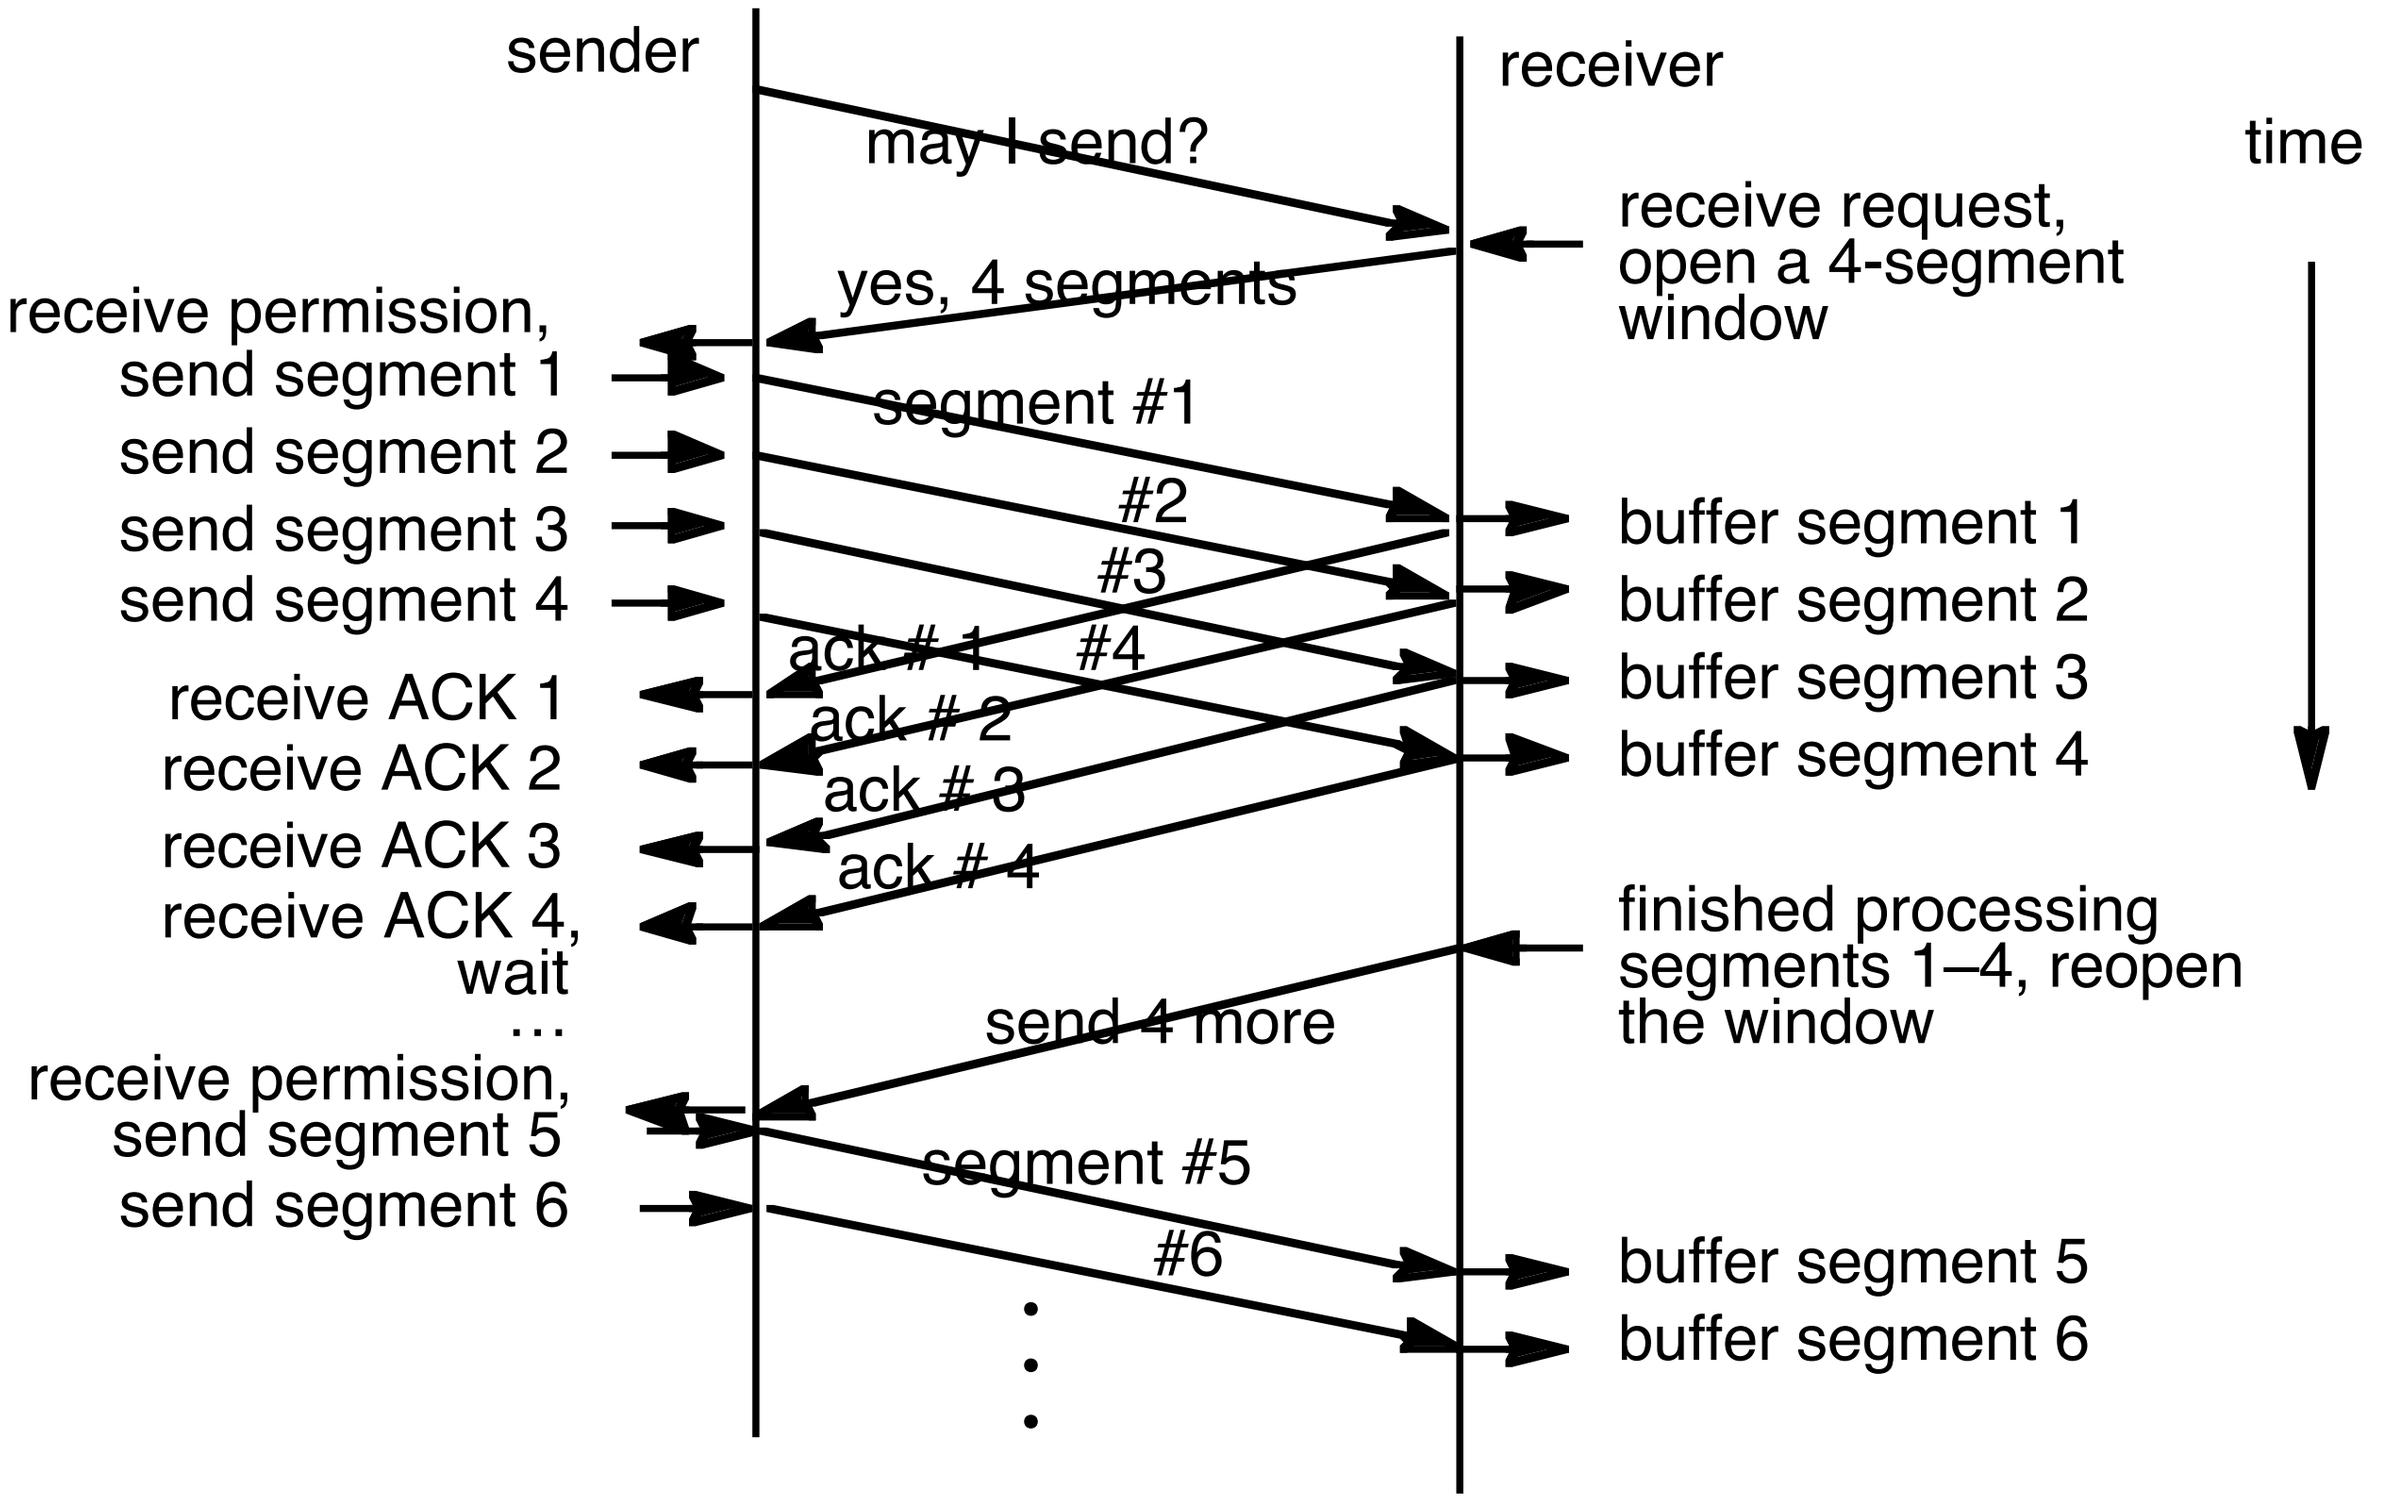

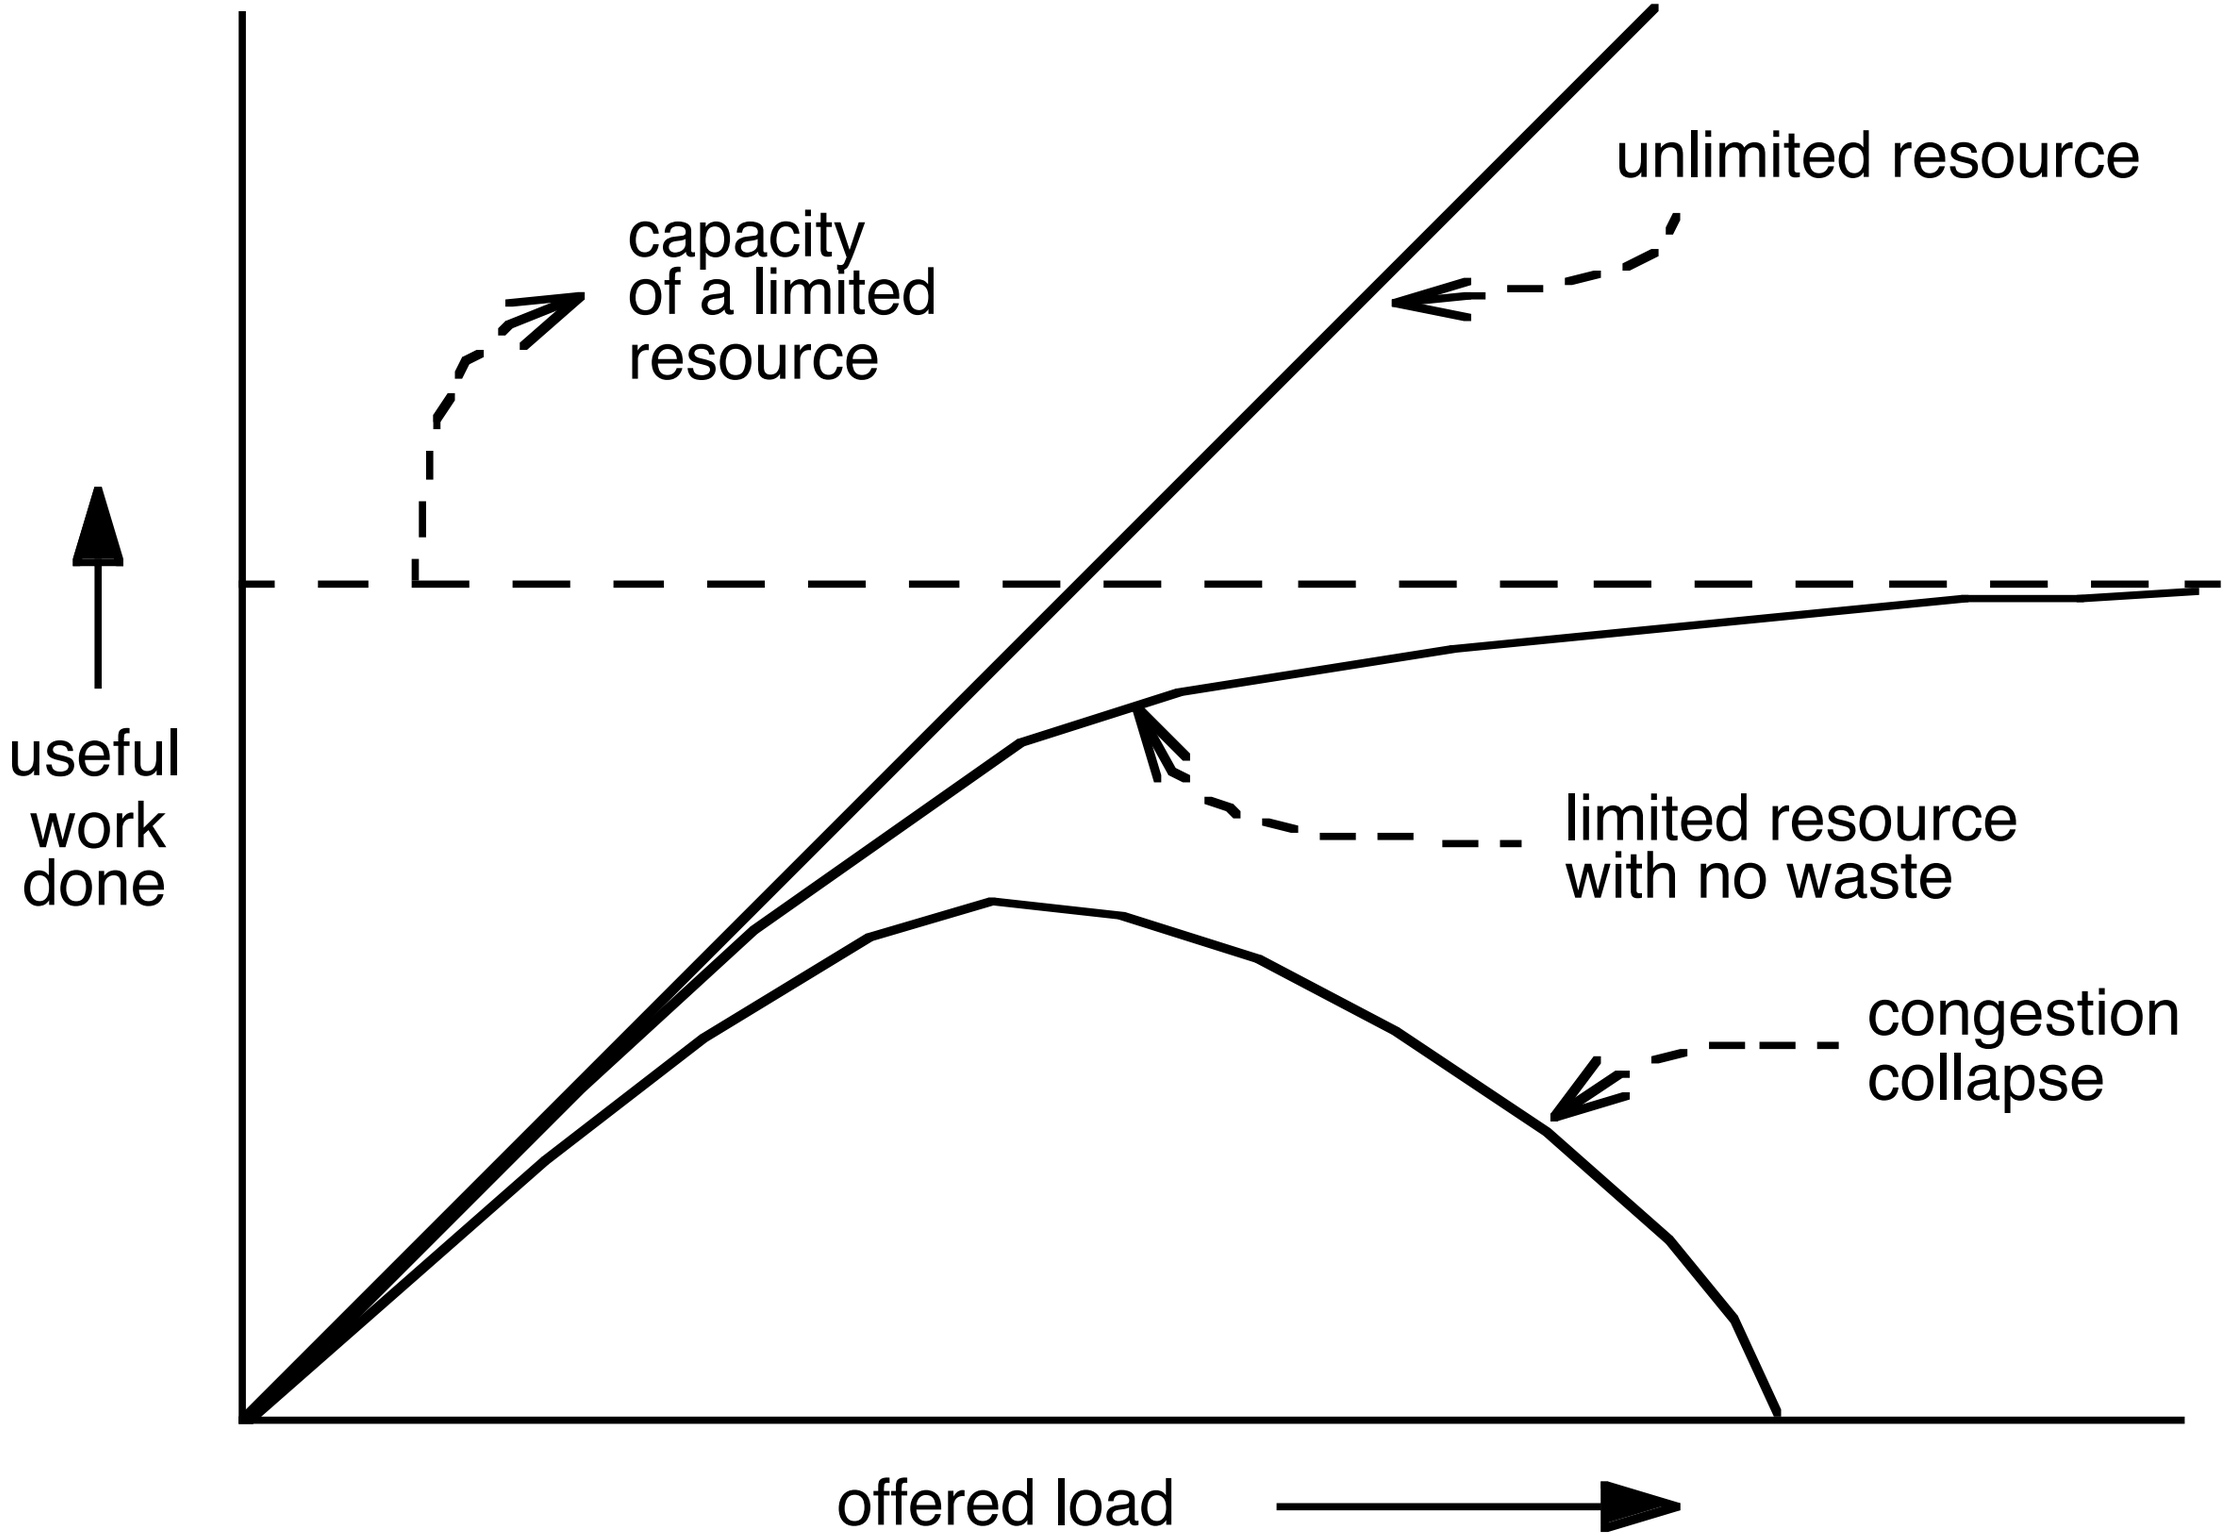

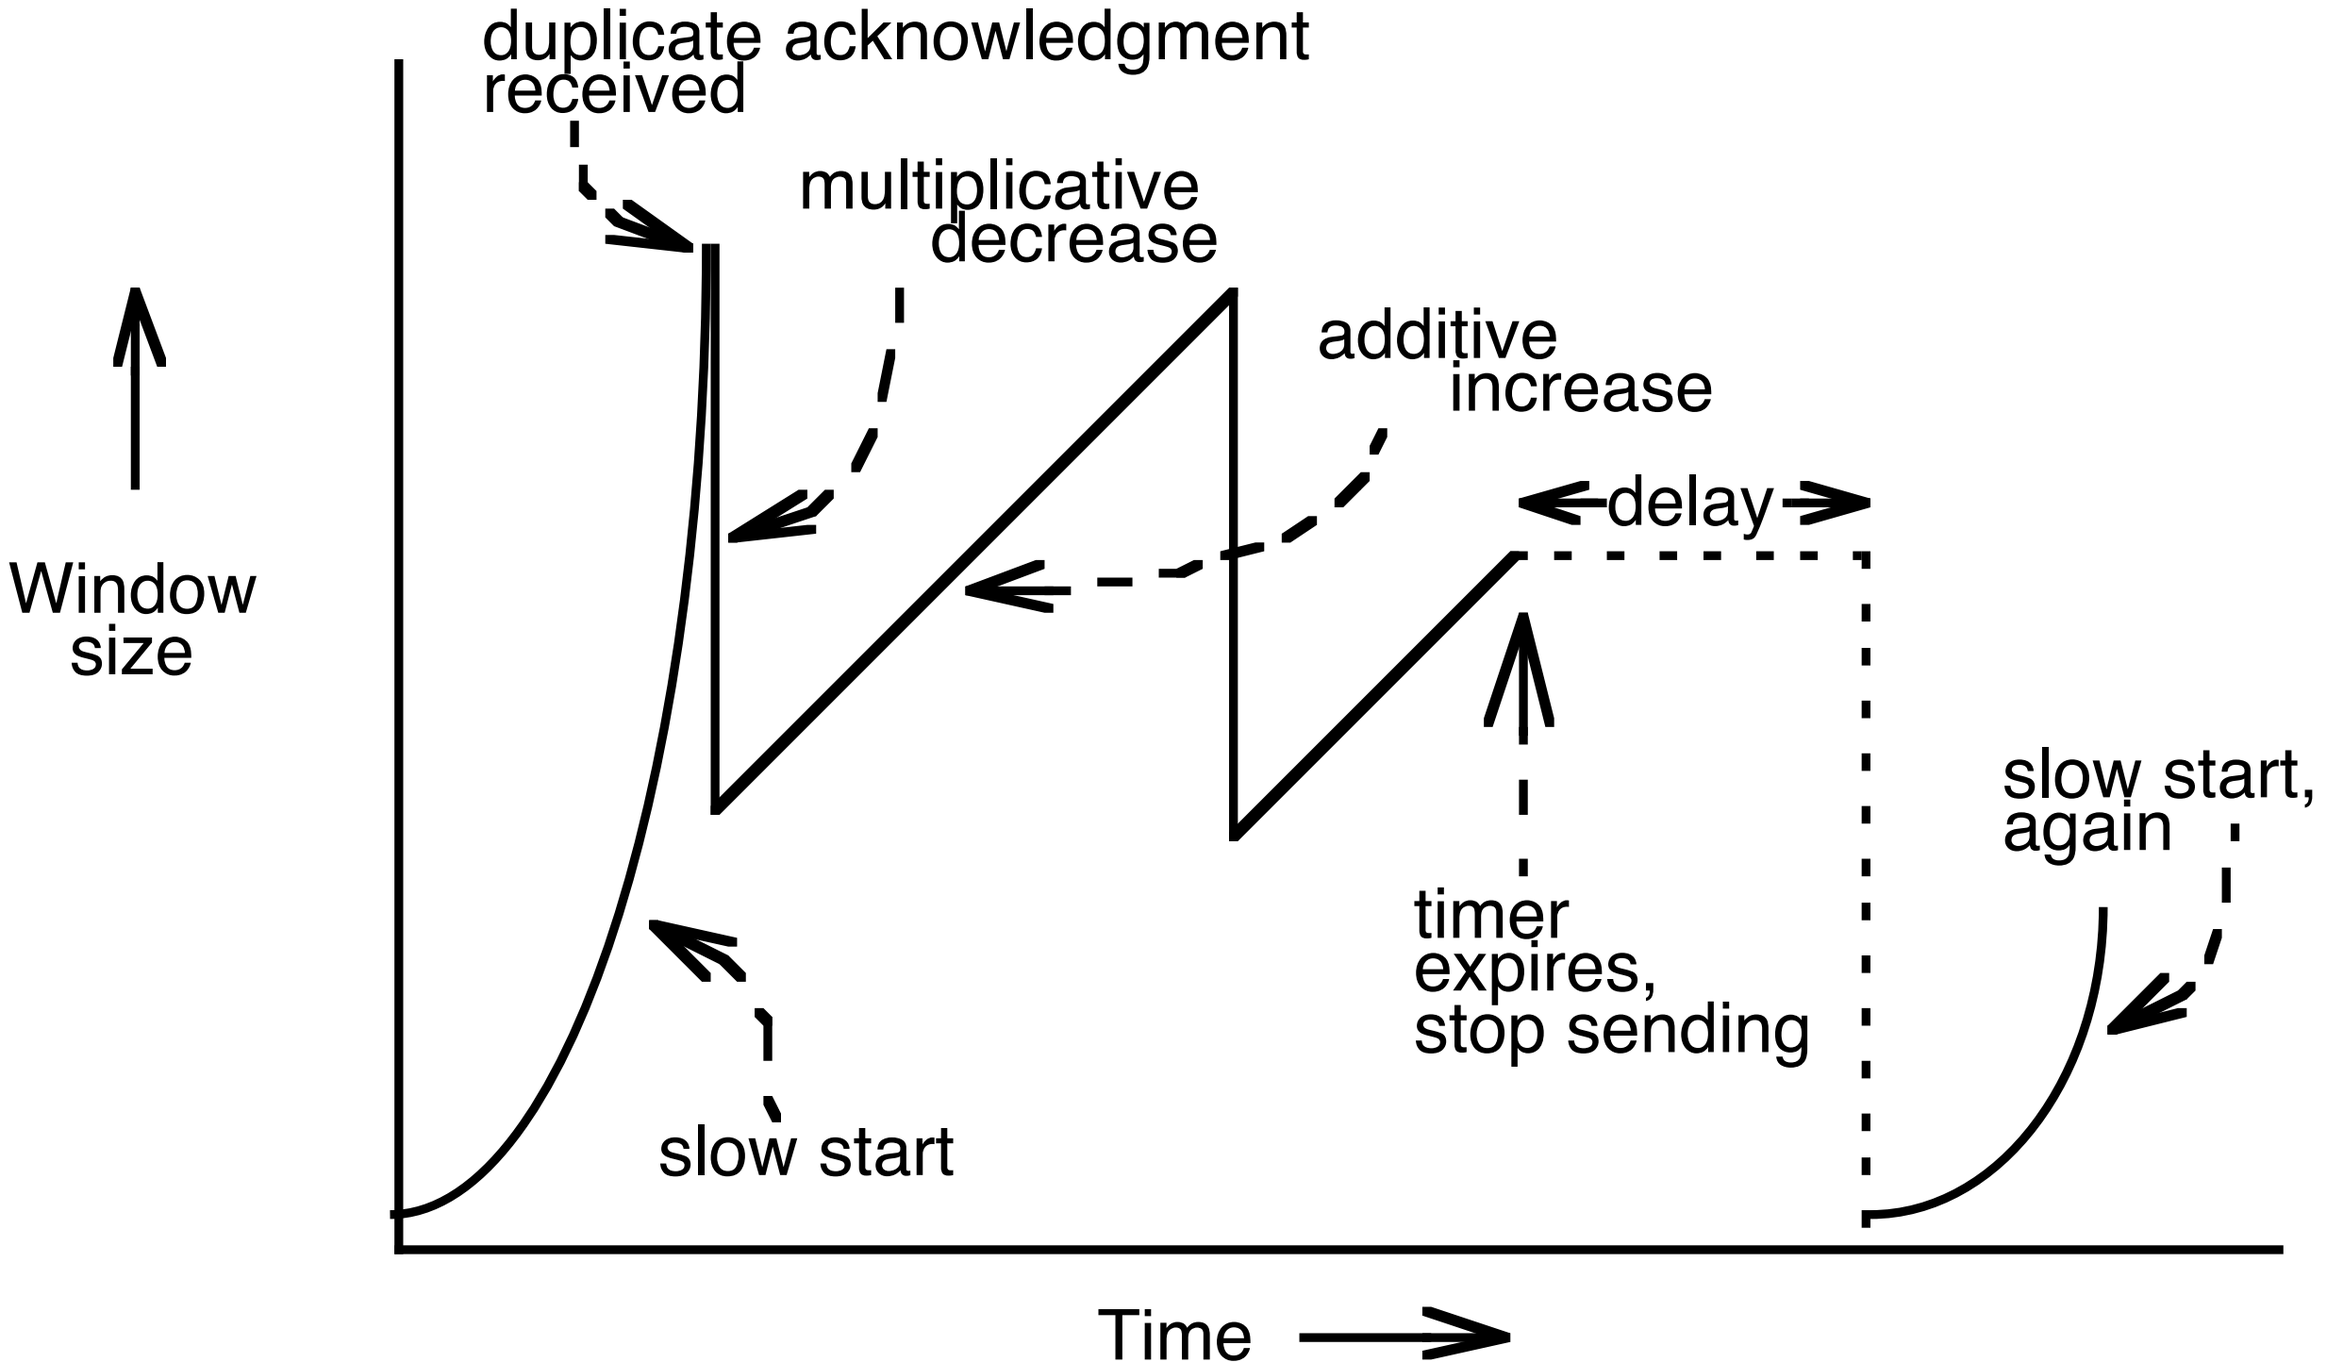

leader	destination	source	type	data	checksum
64 bits	48 bits	48 bits	16 bits	368 to 12,000 bits	32 bits

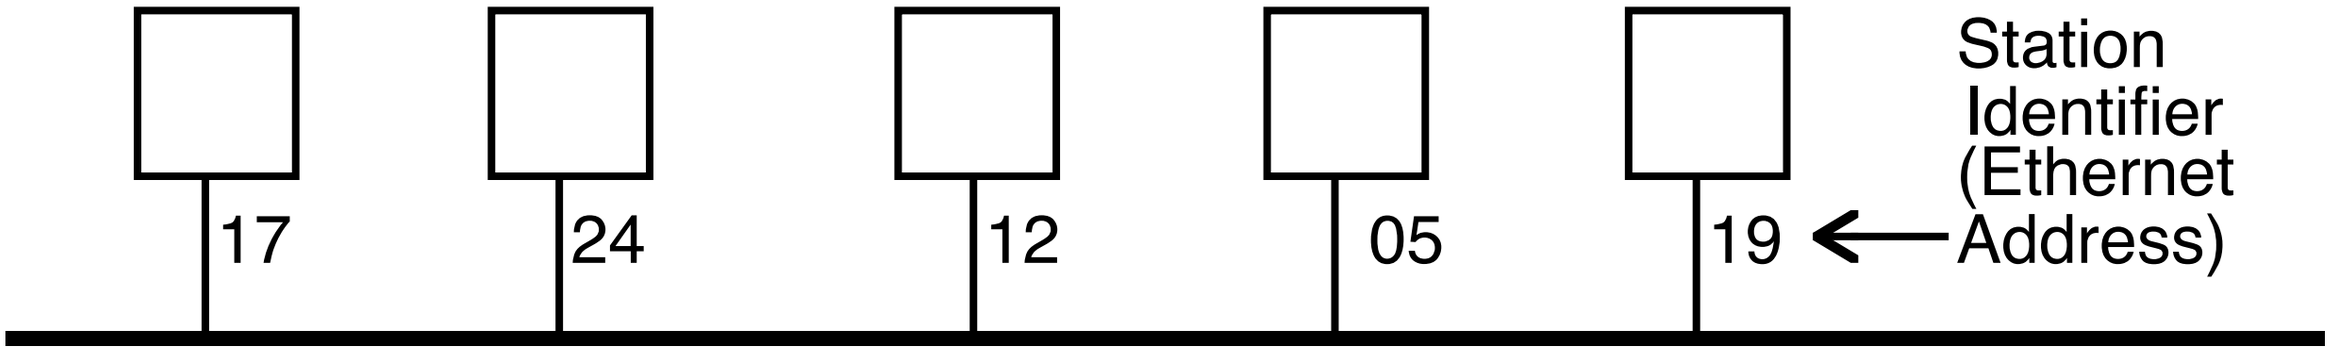

```
procedure ETHERNET_HANDLE (net_packet, length)
    destination ← net_packet.target_id
    if destination = my_station_id or destination = BROADCAST_ID then
        GIVE_TO_END_LAYER (net_packet.data,
                           net_packet.end_protocol,
                           net_packet.source_id)
    else
        ignore packet
```

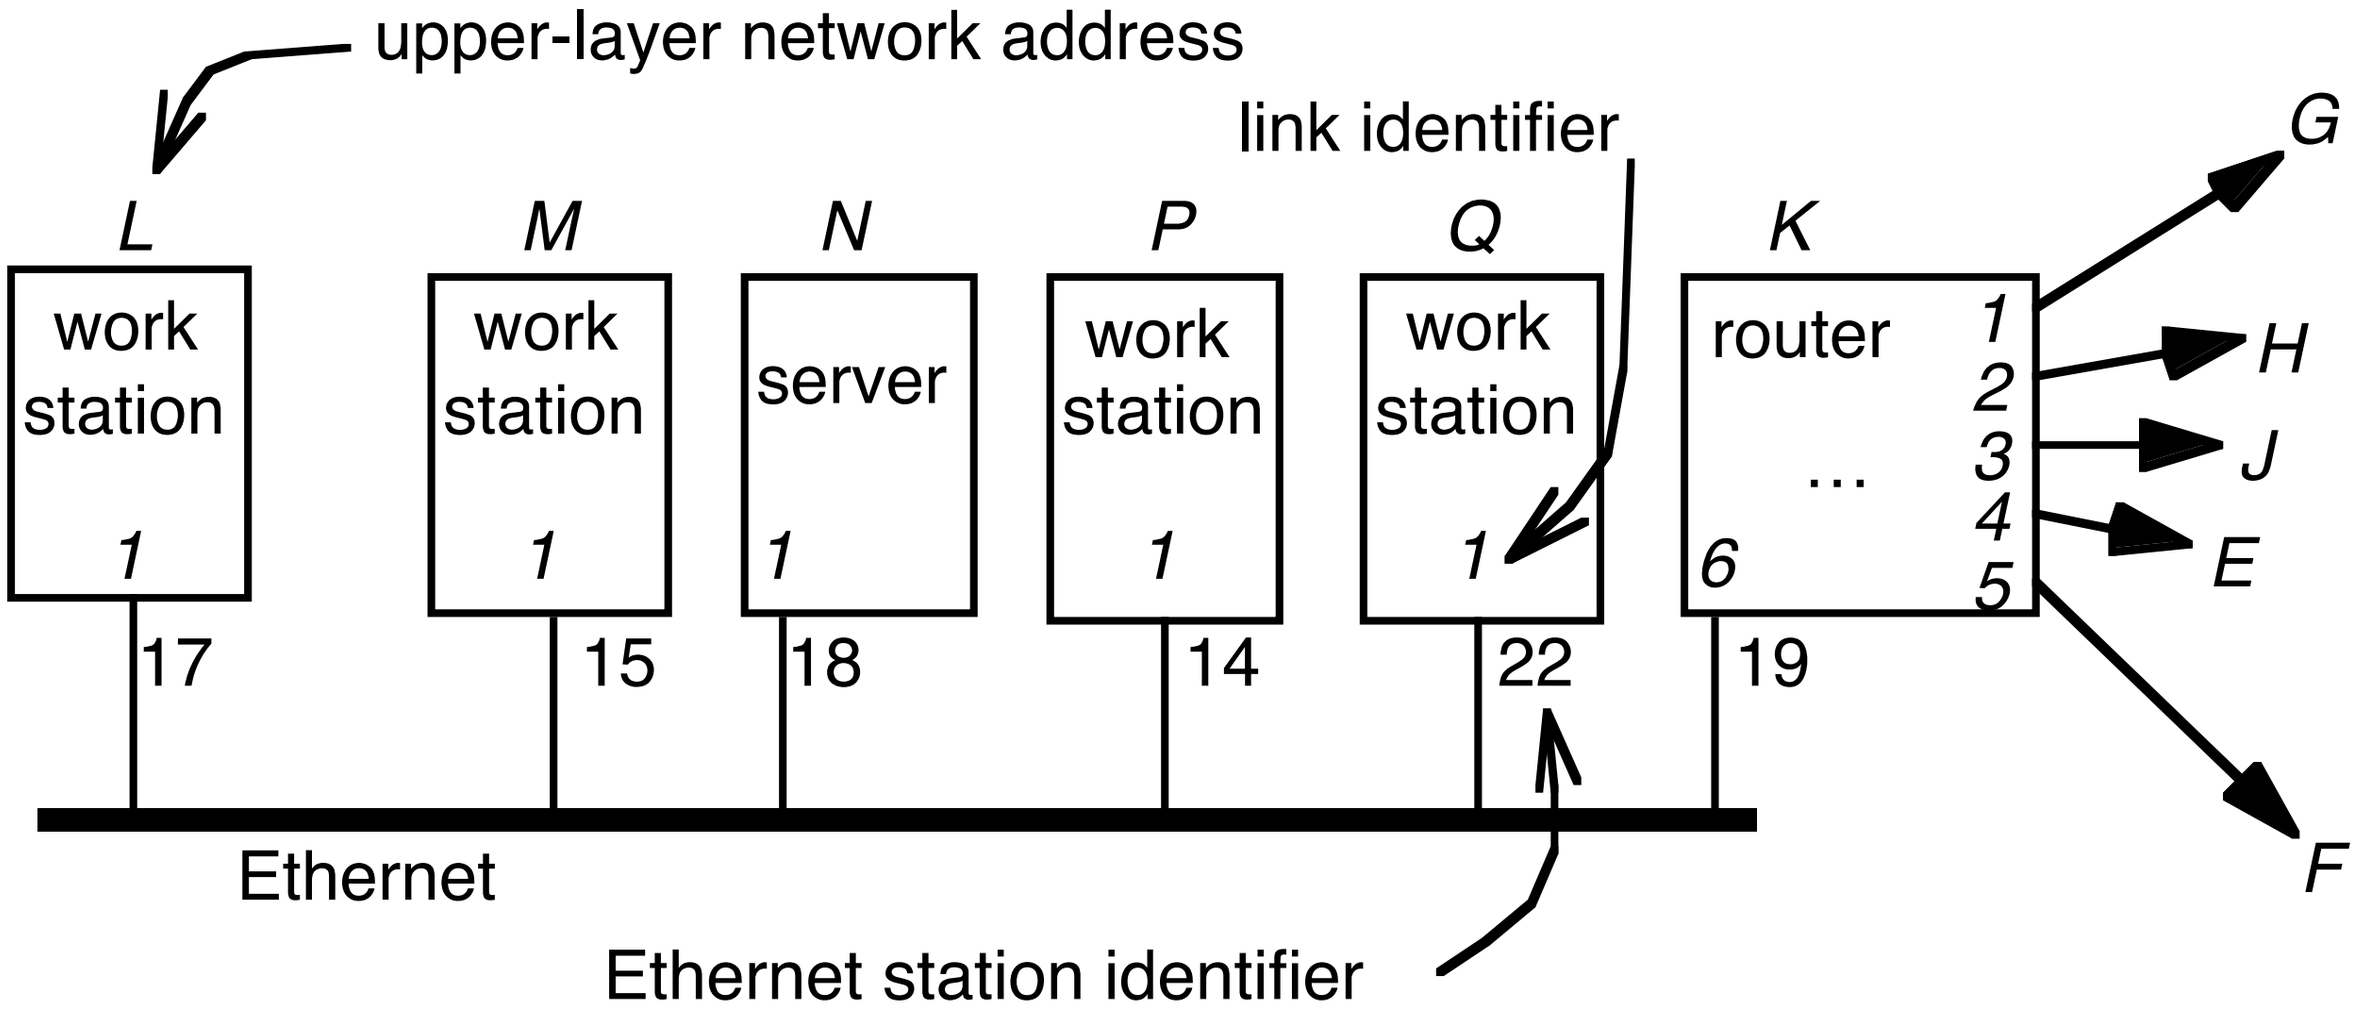

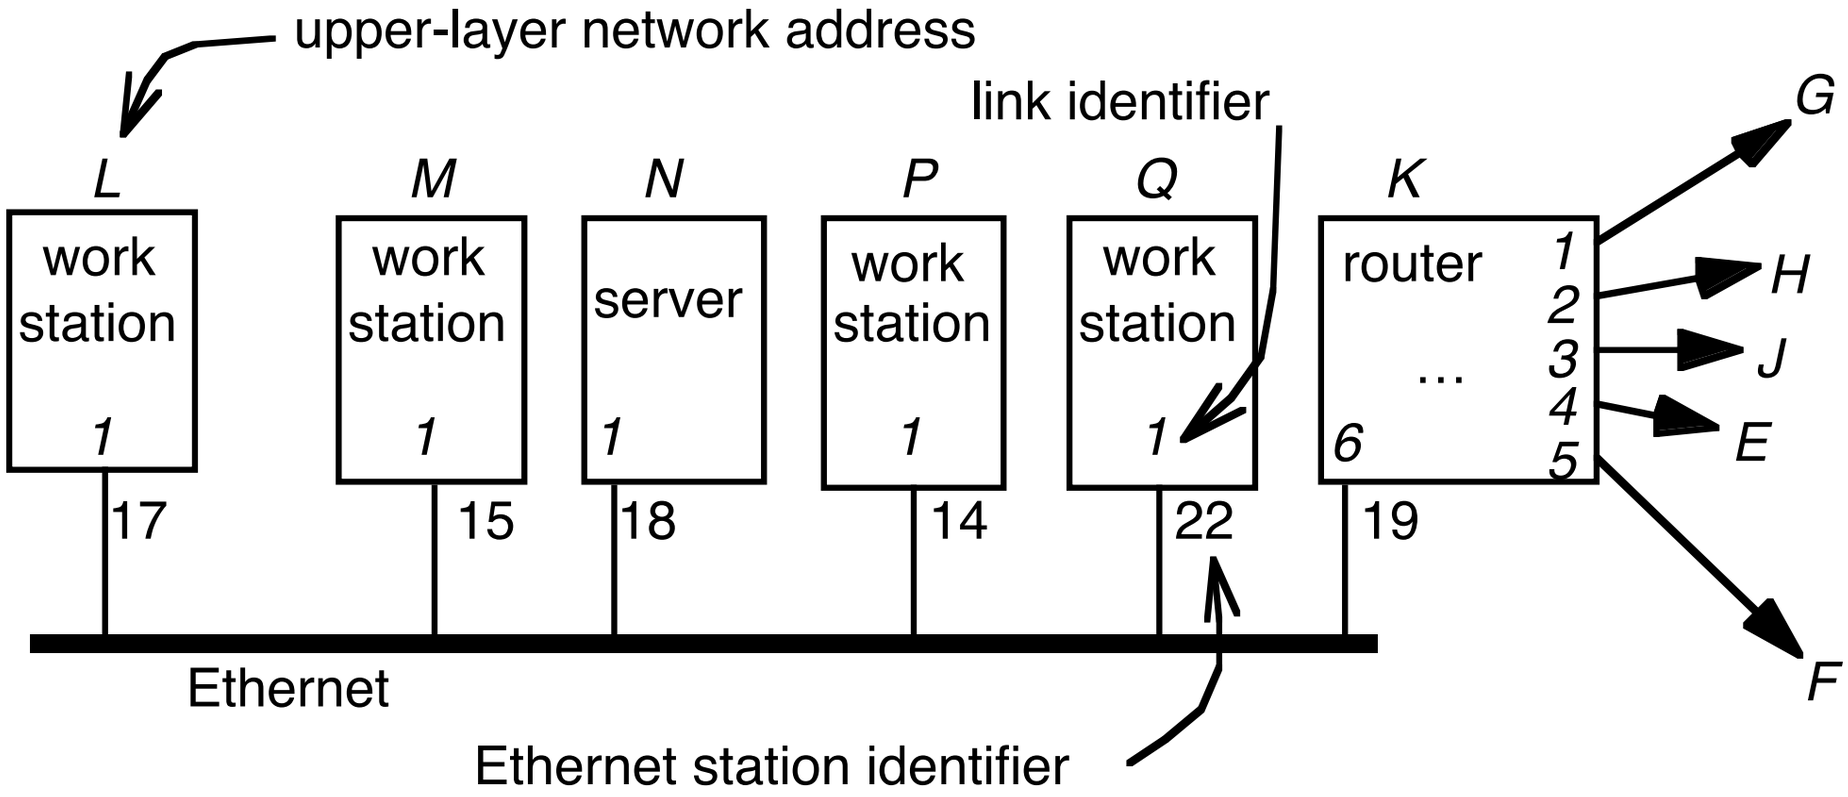

internet address	Ethernet/ station
M	enet/15
N	enet/18
P	enet/14
Q	enet/22
K	enet/19
E	enet/19

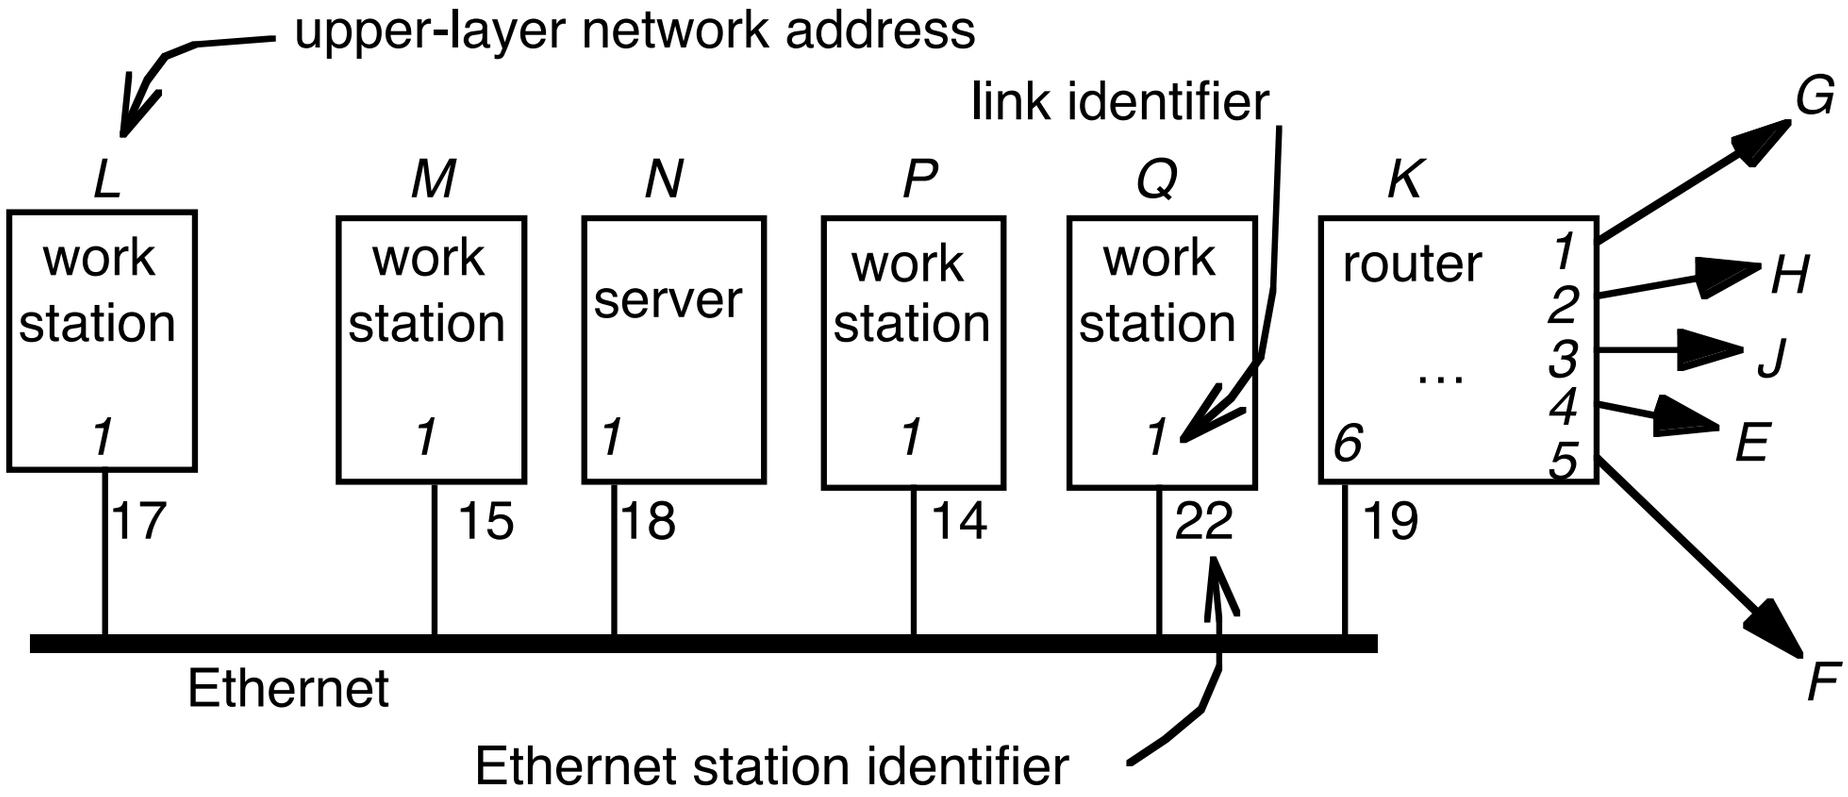

internet address	Ethernet/ station
M	enet/15

internet address	Ethernet/ station
M	enet/15
E	enet/19