Real-Time Performance Guarantees over Wireless LANs

Olaf Meyer

CIS 640, Spring ‘98

Organization

- Wireless LAN technologies
- Wired RETHER
- Wireless RETHER
- Mobility related issues
- Conclusion

References

- Chitra Venkatramani and Tzi-cker Chiueh, *Design, Implementation and Evaluation of a Software-based Real-Time Ethernet Protocol*, Proceedings of ACM SIGCOMM '1995
- Prashant Pradhan and Tzi-Cker Chiueh, *Real-Time Performance Guarantees over Wired/Wireless LANs*, SUNY at Stony Brook (to appear)
- *Wireless LAN Technologies*, Jean Tourrilhes

CSMA/CD (wired Ethernet)

- check the medium (wire) to assess if it is free
- immediately send data if the line is idle
- each transmitter can detect a collision
- on a collision the transmitter retries transmitting (with exponential back off)

Wireless LANs using CSMA/CA

- transmitter listens to other signals to assess the availability of the channel
- a transmitter cannot detect collisions (its transmission masks all other signals)
- when channel is busy wait until it becomes idle, then start *contention*

Avoiding packet loss on the air

Reasons for packet loss

- high BER on the air
- collision of packets

How do we avoid packet loss?

- positive acknowledgement
- MAC level retransmissions
- Fragmentation

Hidden node problem

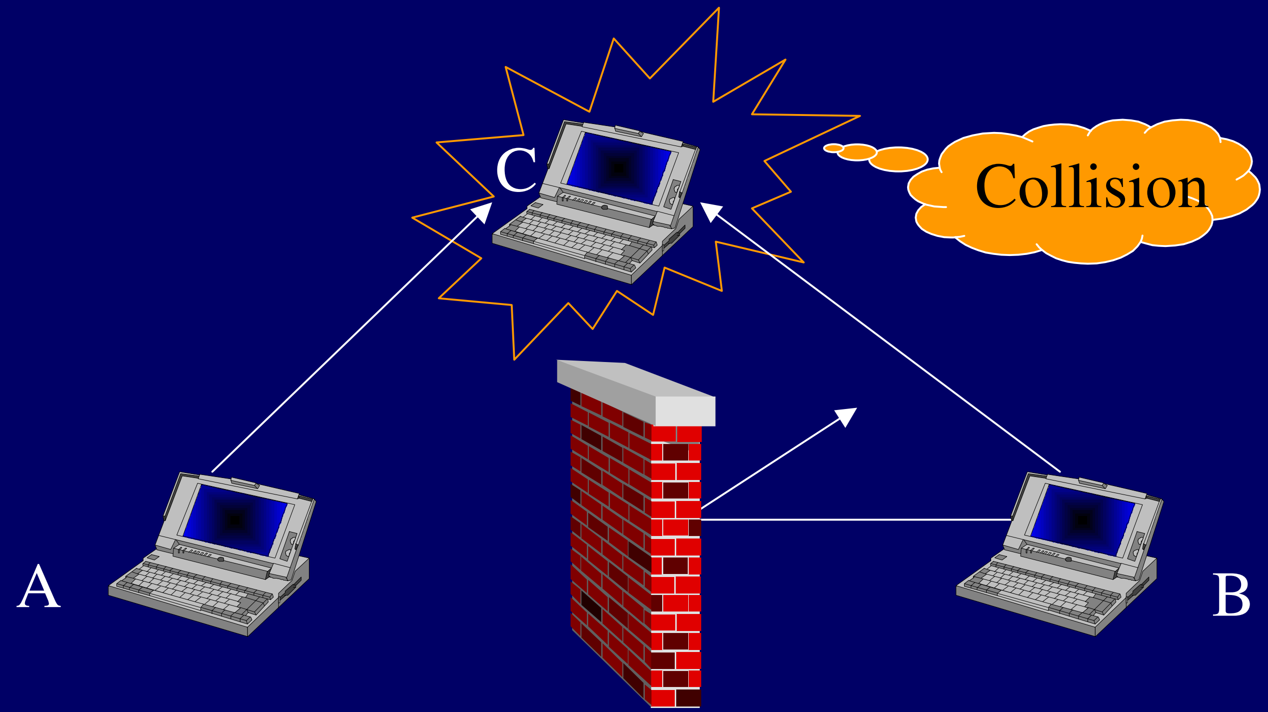

Channel Access Mechanisms

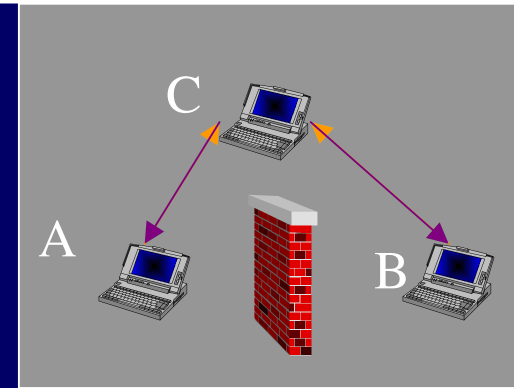

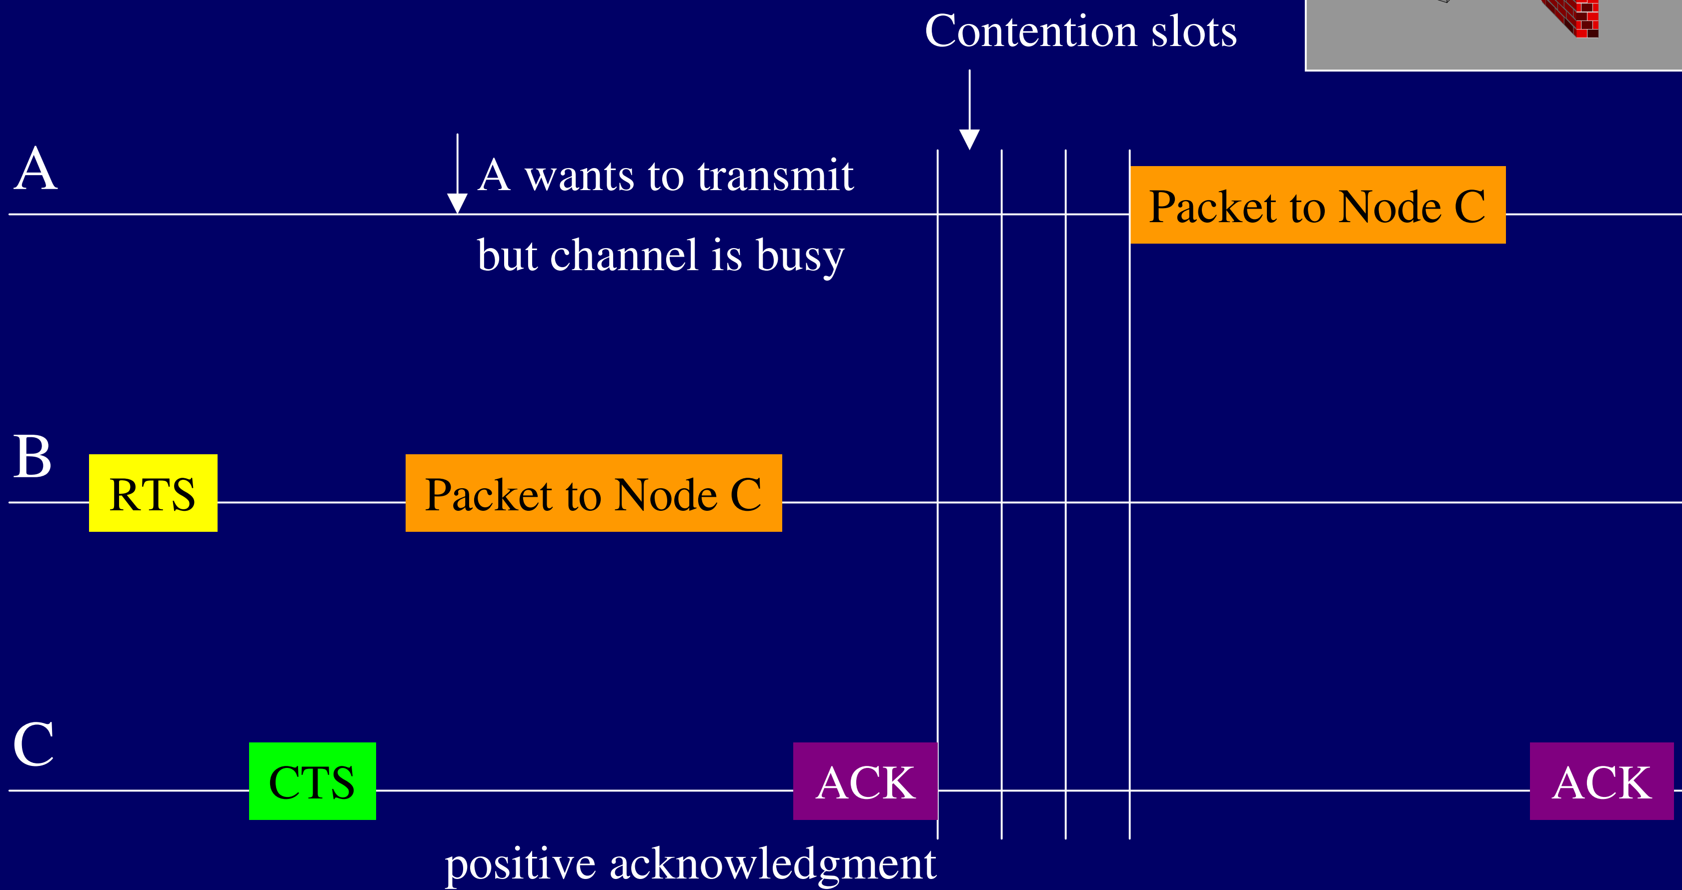

Throughput considerations

Wireless LAN protocols have a high overhead

- radio technology
  - large synchronization fields
  - large slots in the contention window
- additional features
  - larger MAC headers (NWID, encryption)
  - new management packets (AP registration)
- trade-offs to increase the reliability

throughput usually between 50 and 70% of the signaling rate

QoS on LANs

CSMA/CD and CSMA/CA do not provide any bandwidth or latency guarantees

How can we provide bandwidth guarantees and deterministic, periodic network access?

- Time Division Multiplexing (TDM)
- Token passing

The RETHER Protocol

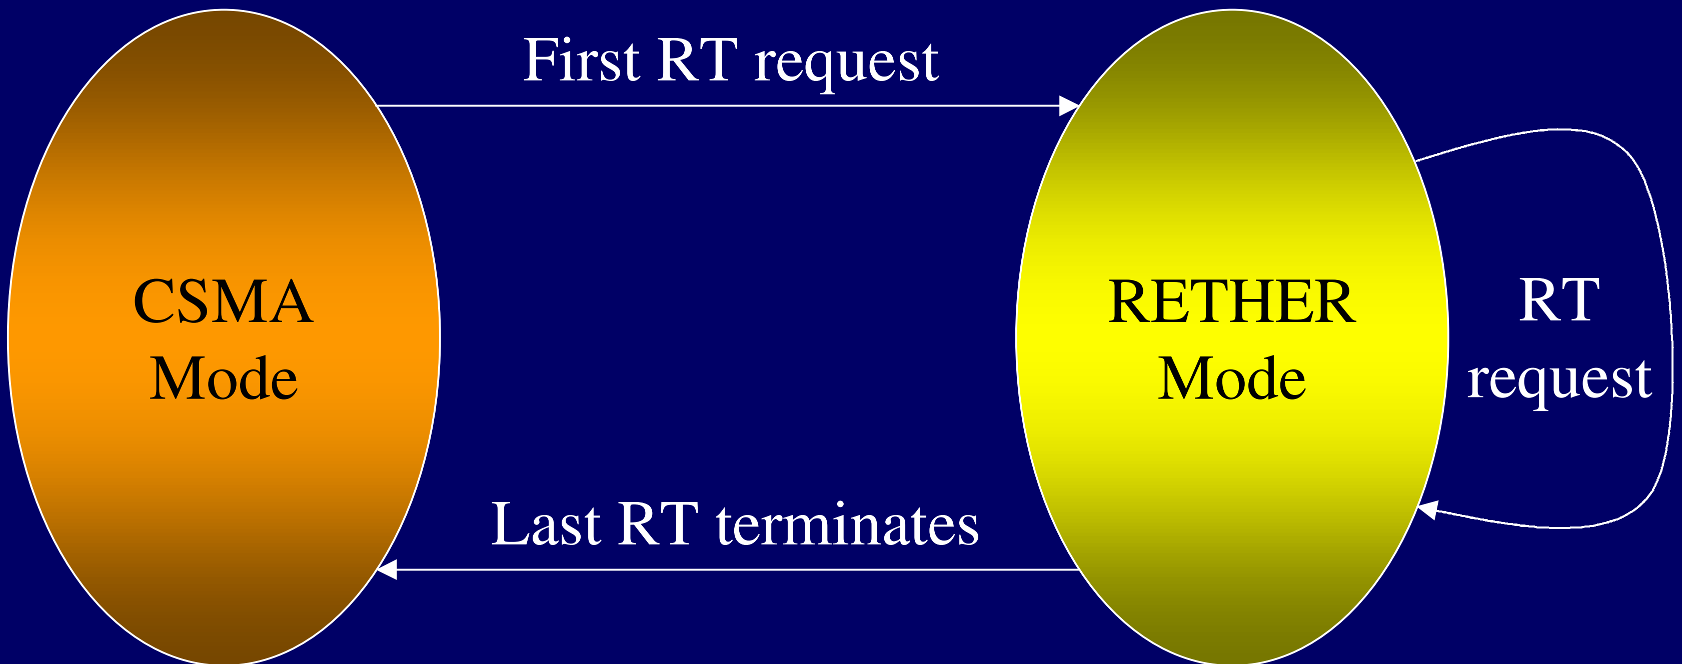

Different Types of Data

- *RT* flow
  - created when a node requests a certain amount of bandwidth per TRT
- *NRT* data
  - all non real-time data belongs to this category (e.g. NFS, X)
  - a certain amount of the TRT is reserved for *NRT* data

Admission control

RETHER	*wired*	*wireless*
Admission control	distributed	centralized
Requests	• *CSMA mode* broadcasted • *RETHER mode* waits for NRT	sent to the base station
Token	carries state	stateless

Overhead at the Base Station

% of Real-Time Reservation	*Interrupt Overhead as % of Cycle Time (66 ms)*
1.2%	2.64%
24%	2.37%
96%	1.47%

- Network interrupts do not seriously degrade the performance of the base station

Performance guarantees

The worst case network access latency for *NRT* nodes (ignoring token passing overhead) is:

$$NRT_{latency} = \frac{TRT \cdot (N-1)}{\left\lfloor \frac{(1-k) \cdot TRT}{MTU / B} \right\rfloor}$$

worst case latency for the *N-th* node

nodes that can send *MTU* of *NRT* data

- *TRT* is the token rotation time
- *N* is the number of nodes
- *B* is the bandwidth of the network
- *k* is the percentage of *RT* data

Multi-Segment RETHER

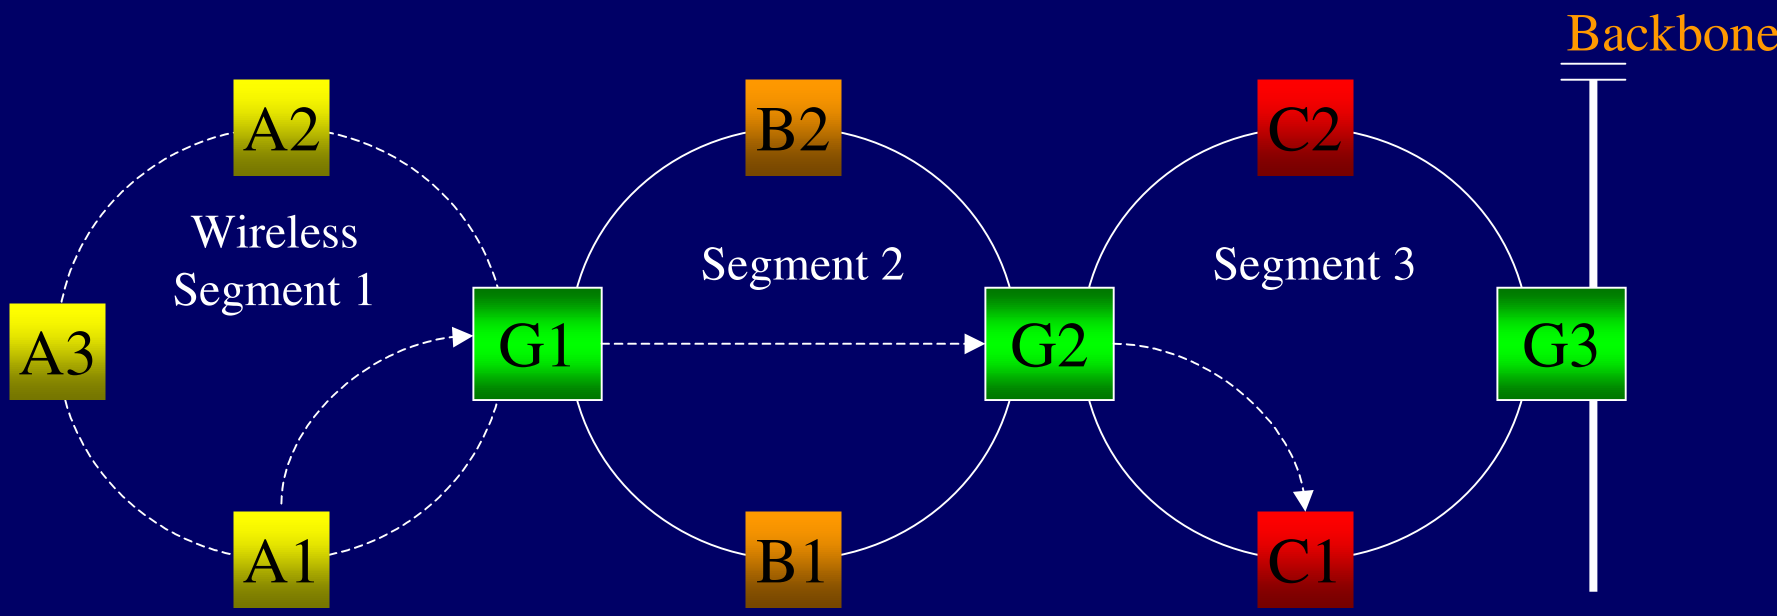

- A1 sends a connection request to C1
- C1 then initiates the connection setup

  send and *receive* sub-connections are established at the gateways

Multi-Segment RETHER (cont.)

- Admission control is performed independently on each segment
- latency guarantees are harder to determine
- extra buffering at the routers is needed

Roaming between Subnets

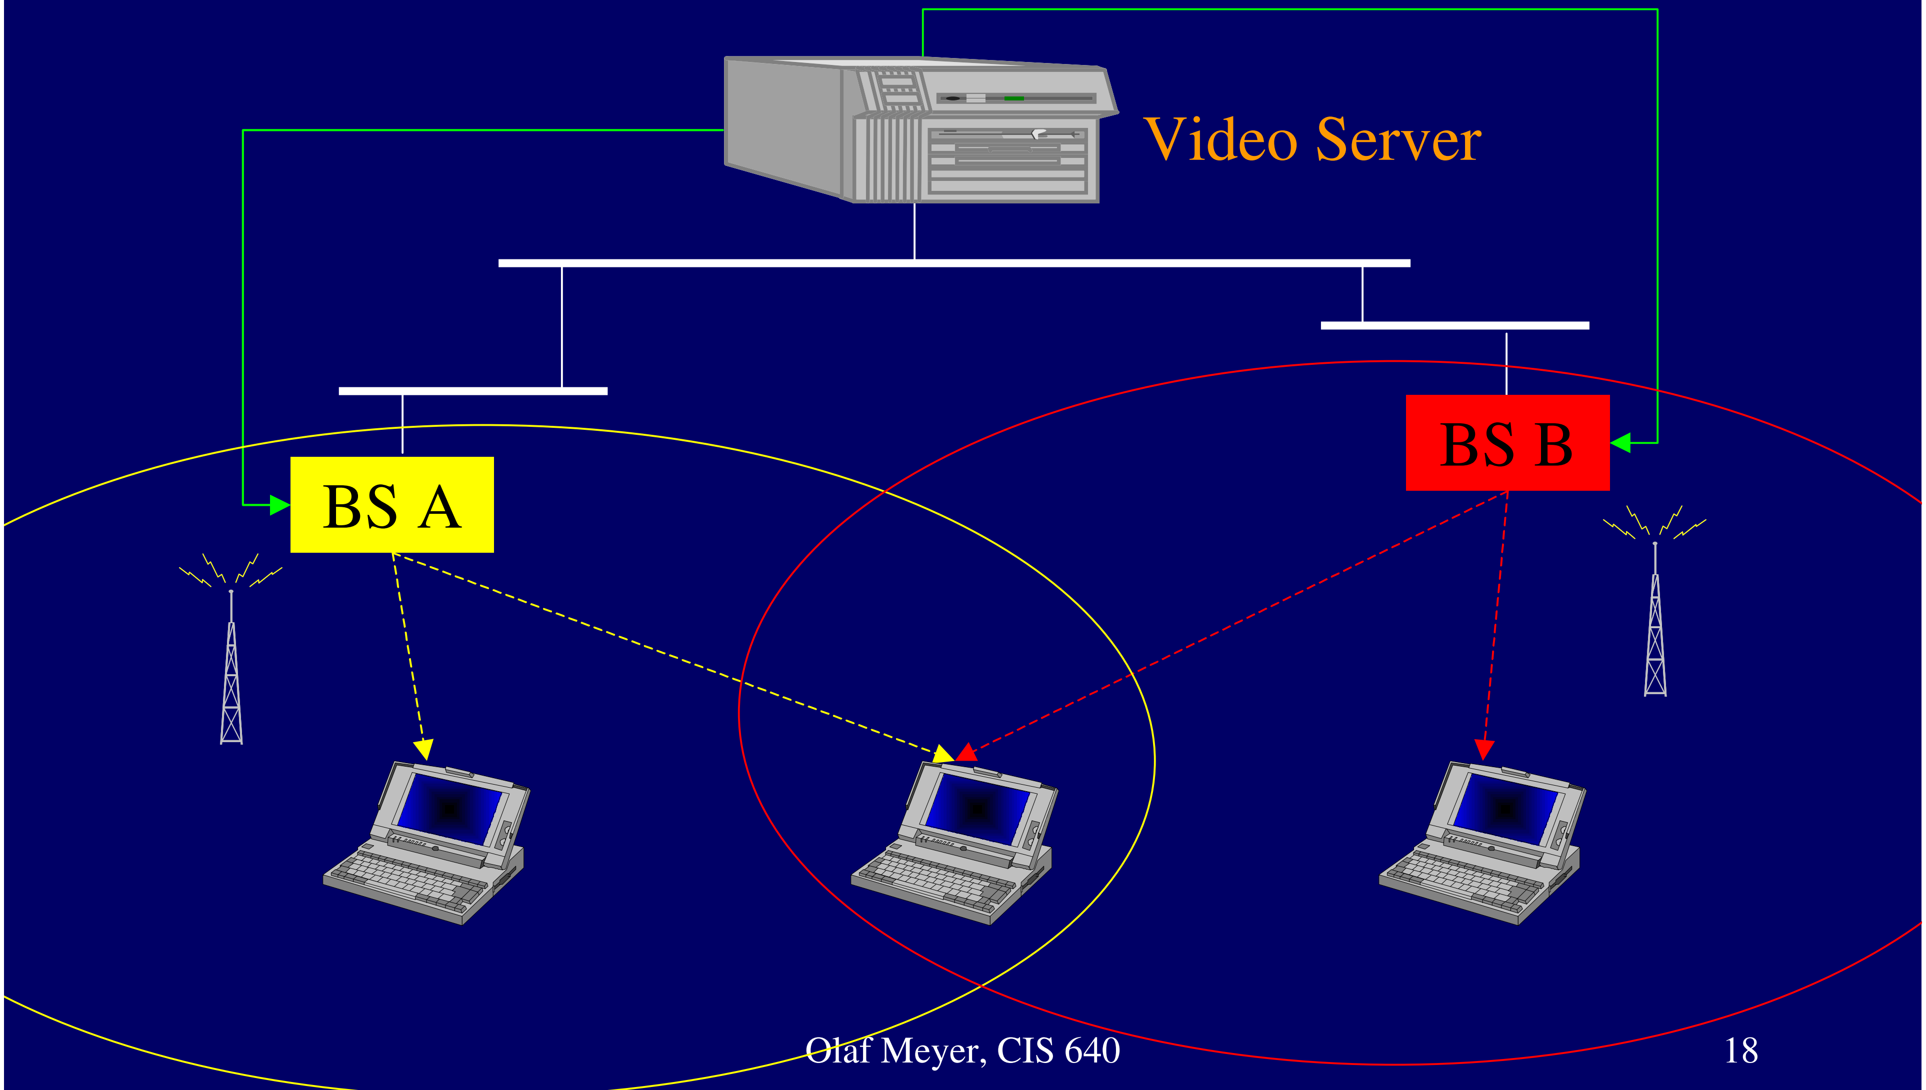

Performance Considerations

The total cell switching time is given by

$$\underbrace{H \cdot P_{beacon} + NRT_{prioritized}}_{\text{switching decision}} + \underbrace{NRT_{latency}}_{\text{connection request}} + C_{setup}$$

- H is the number of consistent 'wins' of signal strength
- P_{beacon} is the period between beacon signals
- $NRT_{prioritized} = k \cdot TRT$

For the experimental setup we get a time of 1.82 sec

(2 Mbps, MTU 1500 bytes, 66 msec TRT, H=3,

P_{beacon} = 500msec, C = 23.72 msec, two-segment connection)

Conclusions

- RETHER provides an end-to-end real-time link-layer solution across LANs connected by wired and wireless segments
- Protocol fully implemented in software using off-the-shelf hardware
- SBVS (Stony Brook Video Server) demonstrates

$$NRT_{latency} = \frac{TRT \cdot (N-1)}{\frac{(1-k) \cdot TRT}{MTU / B}}$$

$$H \cdot P_{beacon} + NRT_{prioritized} + NRT_{latency} + C$$

$$\underbrace{H \cdot P_{beacon} + NRT_{prioritized}}_{\text{switching decision}} + \underbrace{NRT_{latency}}_{\text{connection request}} + C_{setup}$$