fast efficient and robust networking FERN

Fast, Efficient, and Robust Multicast in Wireless Mesh Networks

Ian Chakeres
Chivukula Koundinya
Pankaj Aggarwal

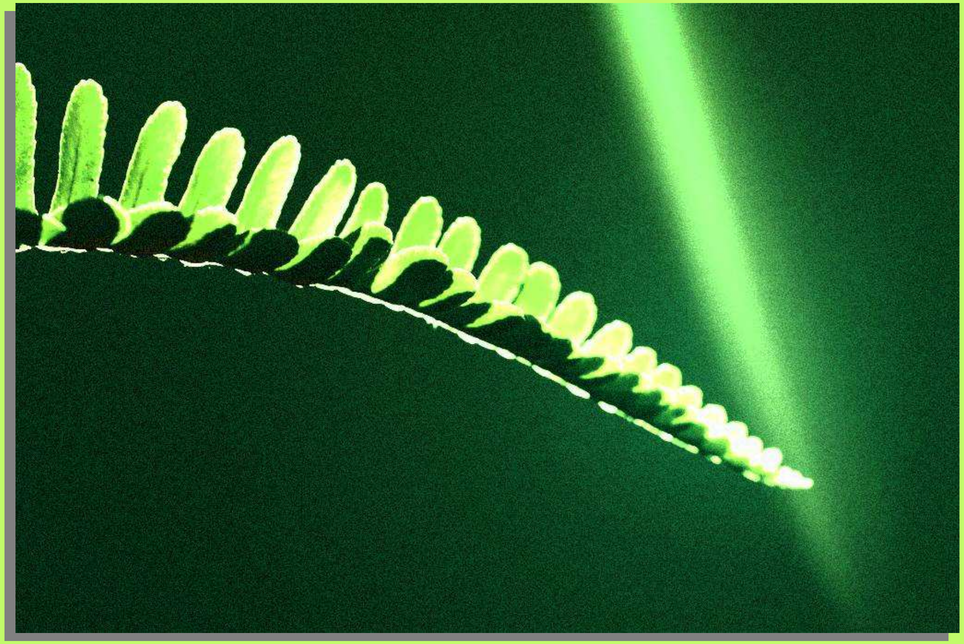

Outline

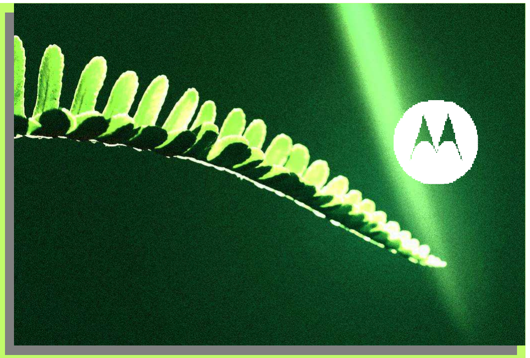

- IEEE 802.11s mesh multicast
- FERM algorithms analyzed
- Experimental evaluation
  - Load
  - Loss
- Conclusions and future work

Motivation

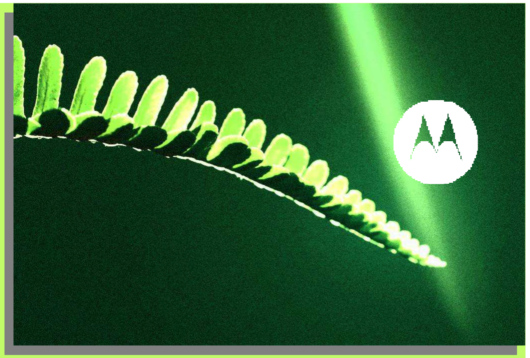

Multimedia applications are often challenged by wireless networks' low bandwidth, long medium access times, and high loss rates.

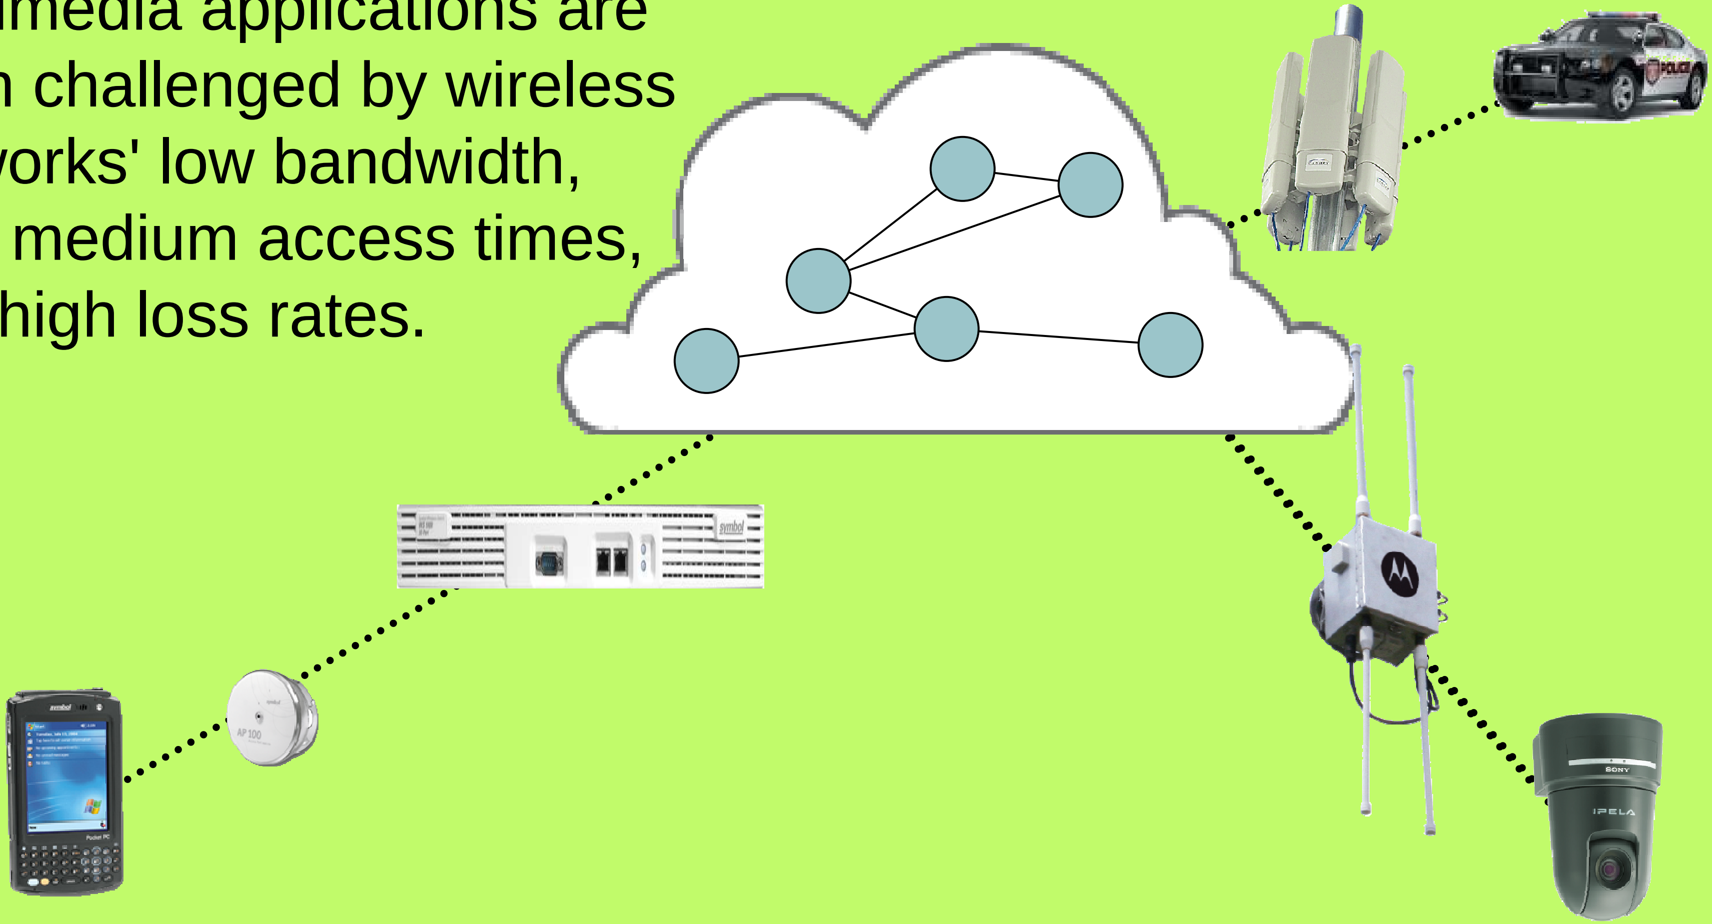

Fast Efficient and Robust Networking FERN

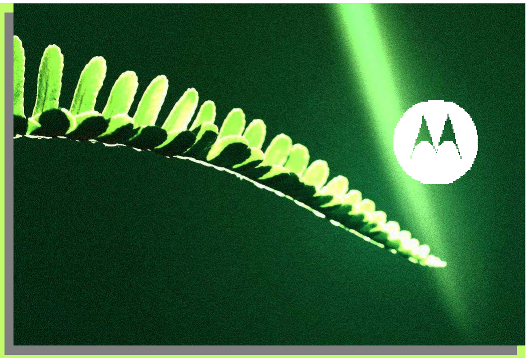

Network devices working individually or cooperatively to intelligently improve the end-user application experience

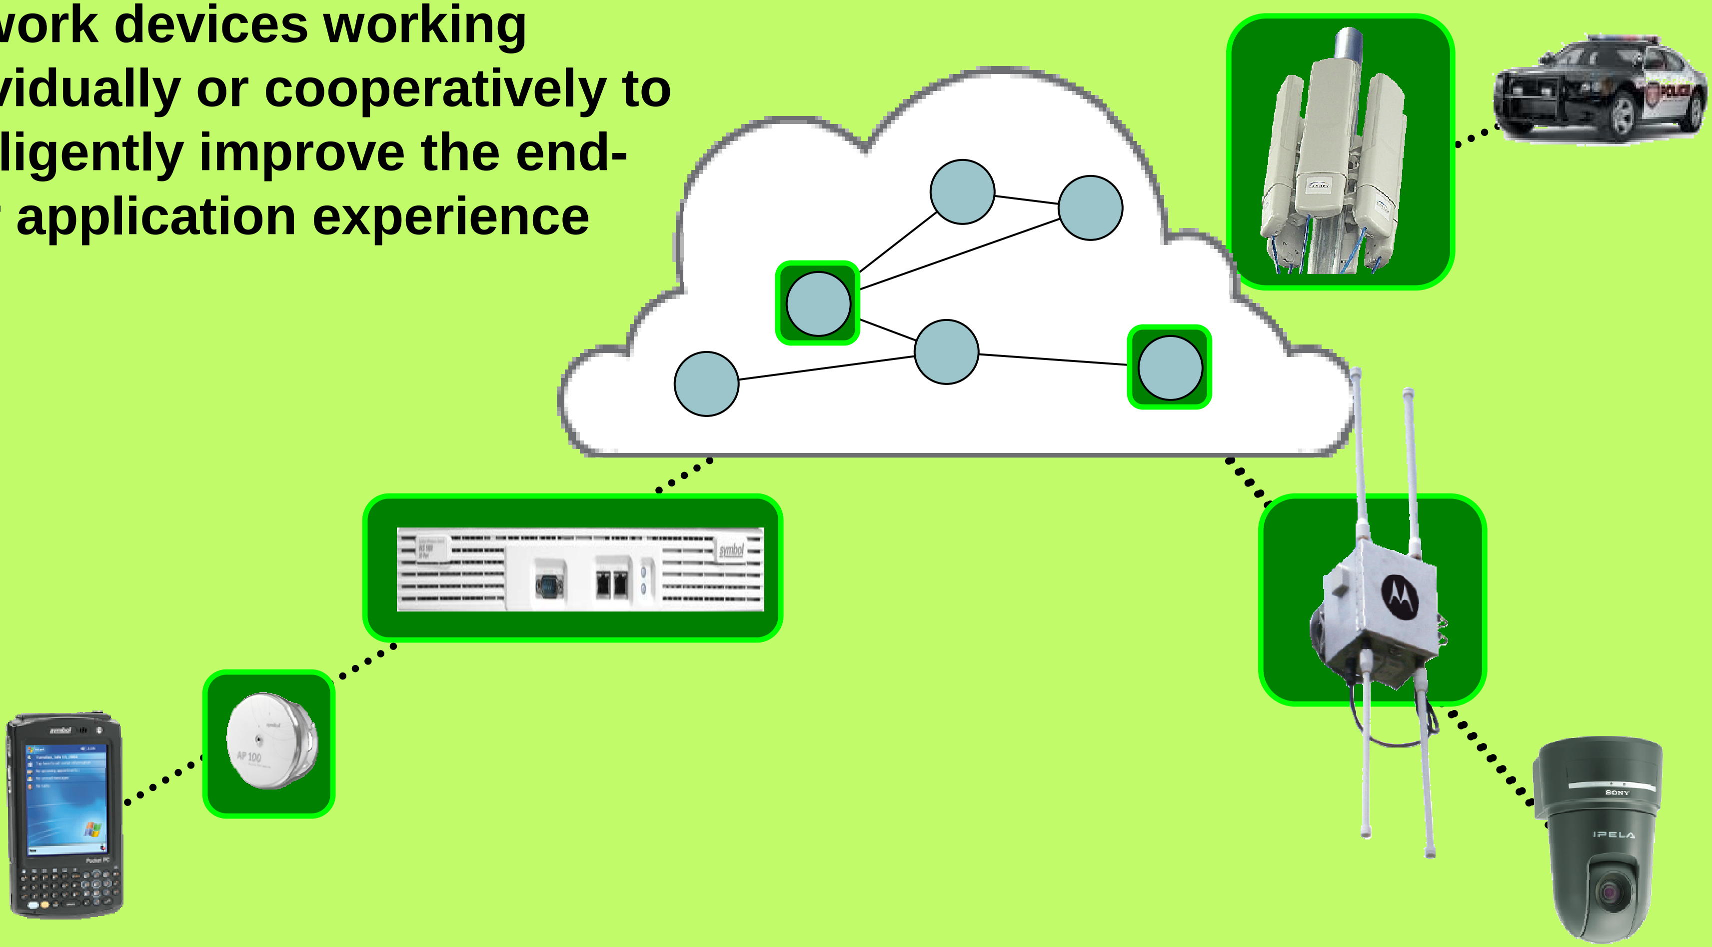

Fast Efficient and Robust Networking FERN

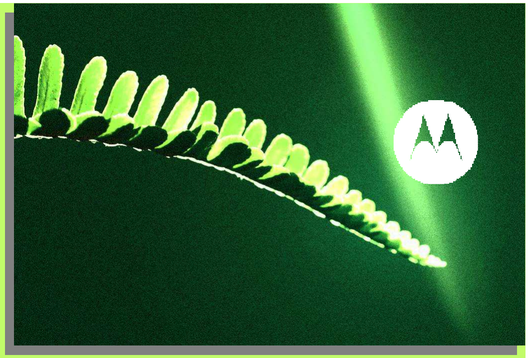

Multicast

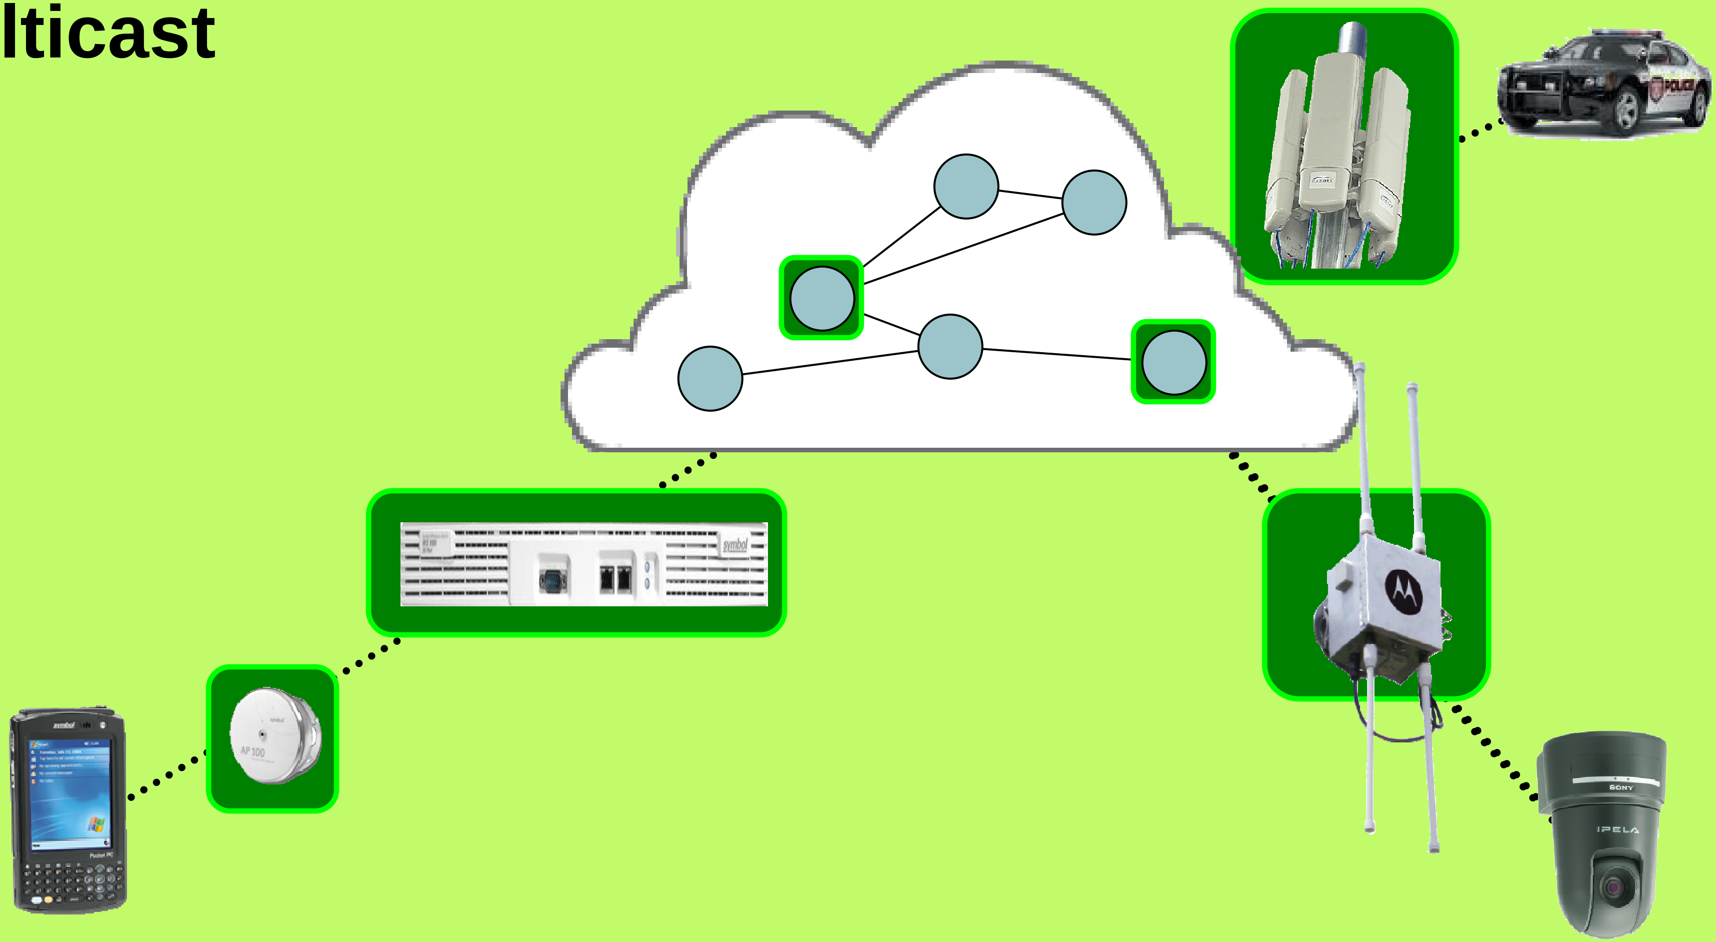

Fast Efficient and Robust Networking FERN

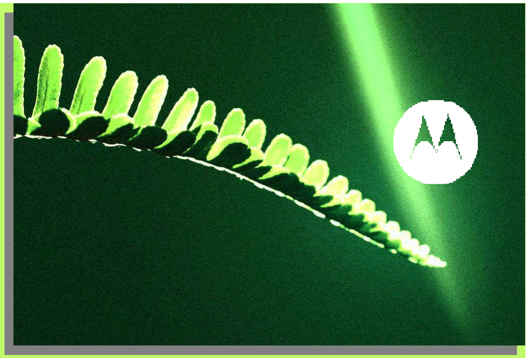

Multicast

Slow & lossy

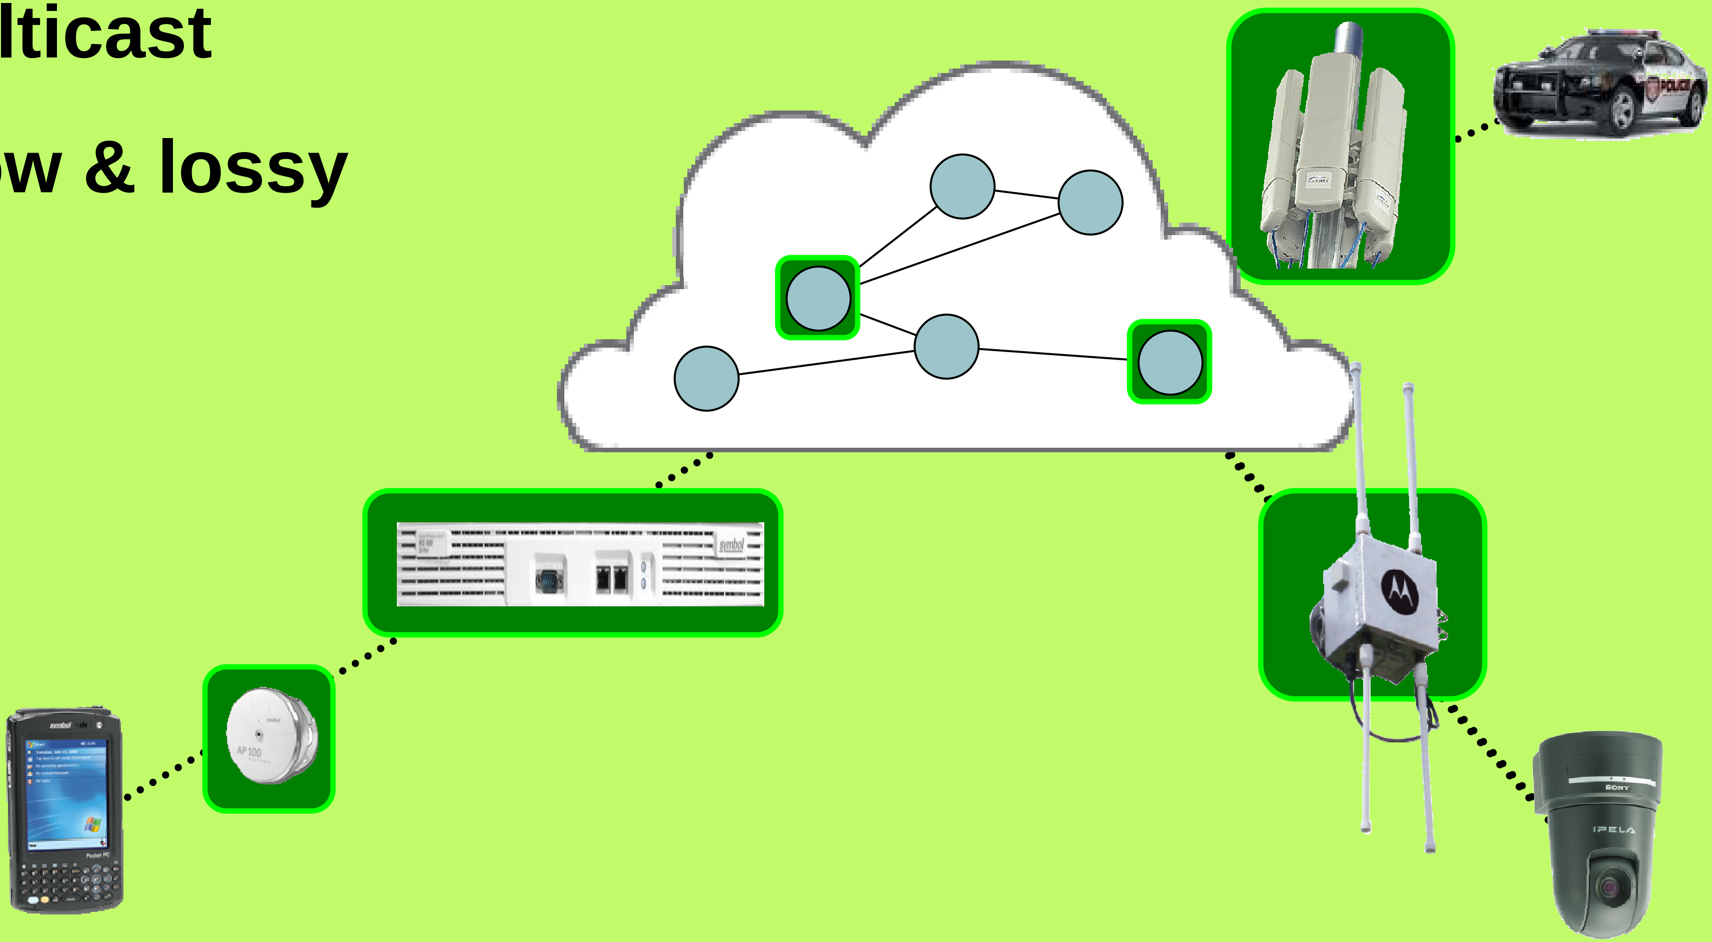

Multicast introduction

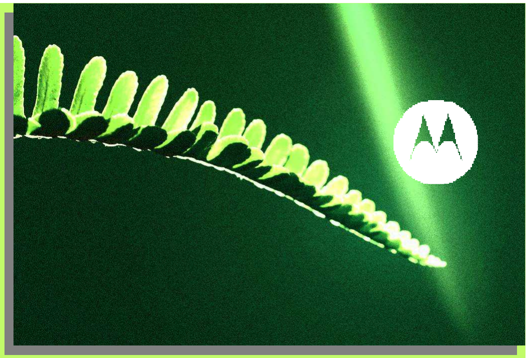

- **Multicast in multihop IEEE 802.11s mesh networks**
- **Single-hop transmission**
  - **IEEE 802.11 basic service set**
    - **1 or 2mbps - IEEE 802.11none/b/g/n**
    - **6mbps - IEEE 802.11a**
  - **No robustness = no ack,**
    no retransmission, no backoff
- **Multihop dissemination**
  - **Classical flooding**
  - **Leaf routers don’t retransmit**
- **We call this algorithm Default Broadcast (DB)**

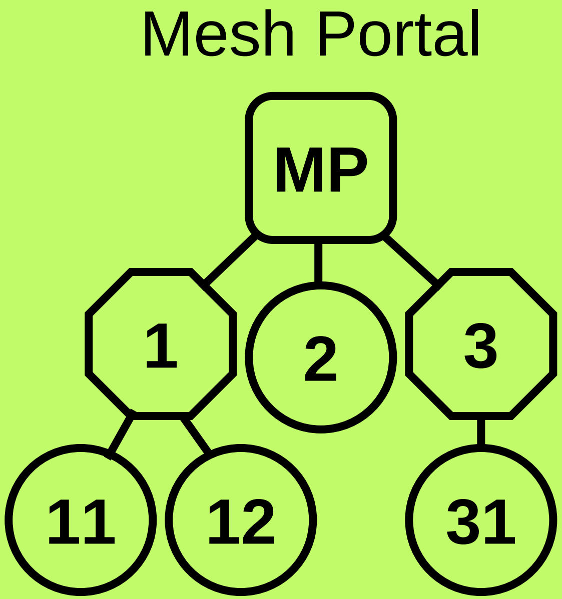

Problems with DB

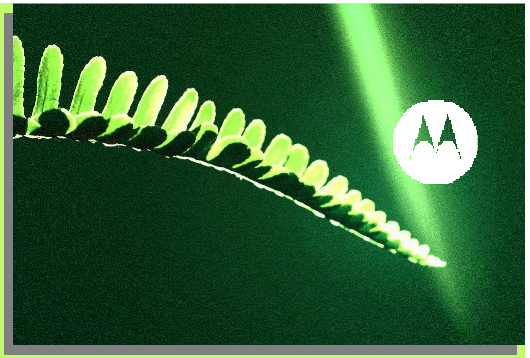

- **Slow**
  - **IEEE 802.11 basic service set rate**
- **Not robust to loss**
  - **Does nothing to counter losses**

FERM algorithms surveyed

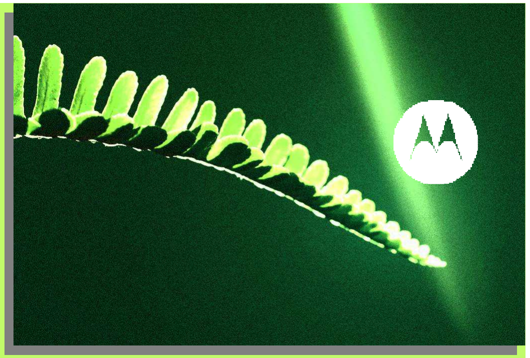

- **Survey**
  - **Default broadcast (DB)**
  - **Fast broadcast (FB)**
  - **Selective unicast (SU)**
  - **Multiple unicast (MU)**
  - **Ack-oriented in-mesh multicast (AM)**

Fast broadcast (FB)

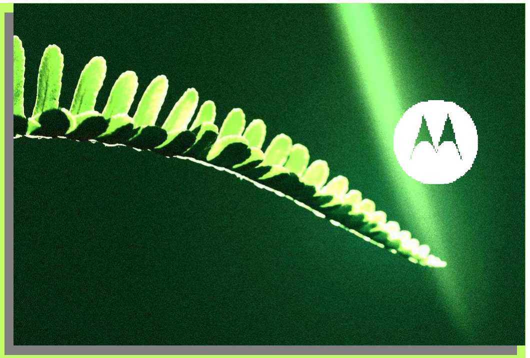

- **Use "fastest" data rate, e.g. 36 mbps**
- **Requires knowledge about downstream neighbors**
- **Intermediate routers re-transmit packets**
- **Faster than DB, still no robustness**

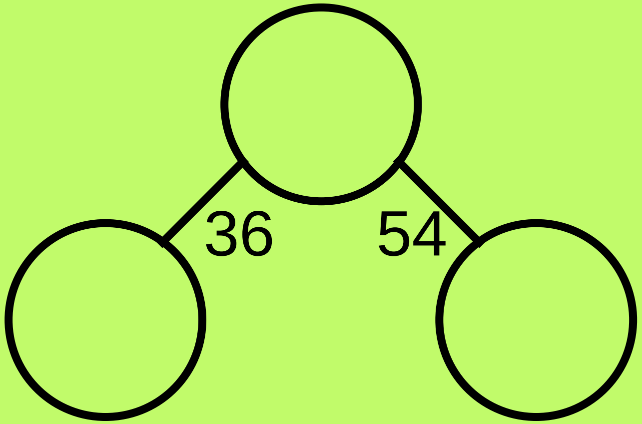

Selective unicast (SU)

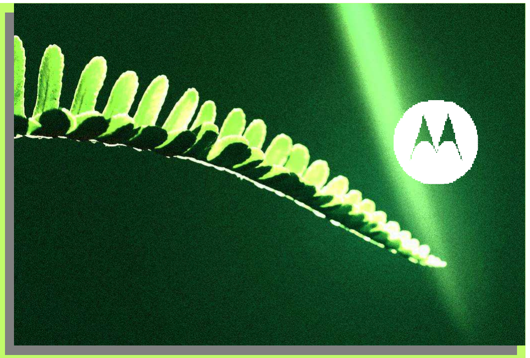

- **Use “fastest” data rate**
- **Unicast to one downstream neighbor**
  - **Use unicast reliability, L2_ACK & retransmissions**
  - **We choose the downstream router with the highest link metric (poorest link)**
- **Other receivers promiscuously process the packet**
- **Fast with some robustness**

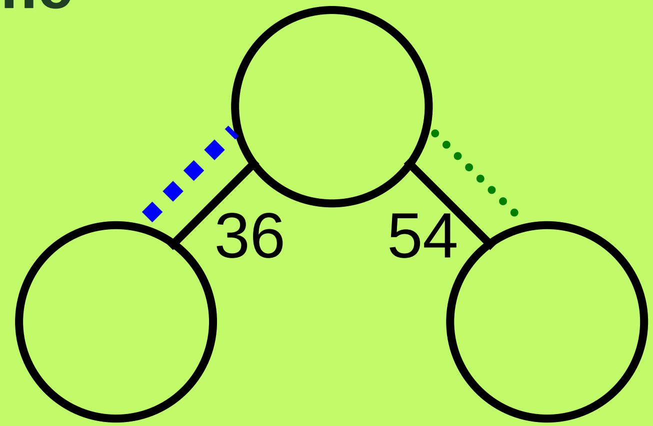

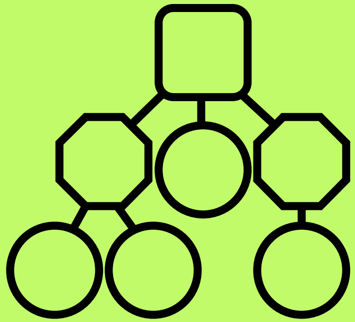

Multiple unicast (MU)

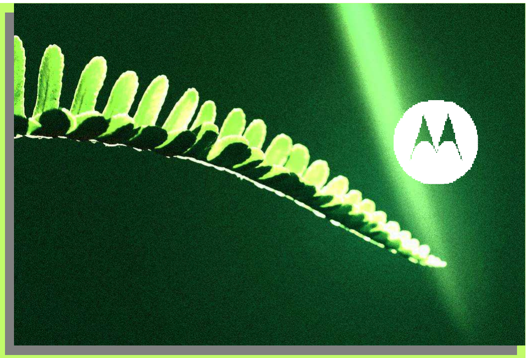

- **Transmit each multicast frame as multiple unicast transmissions**
  - **Use fastest data rate for each transmission**
  - **Maximizes reliability**
  - **L2_ACK, retransmission, & backoff**
- **For dense or low loss scenarios MU is expensive**
- **Fast & highly robust, but resource intensive**

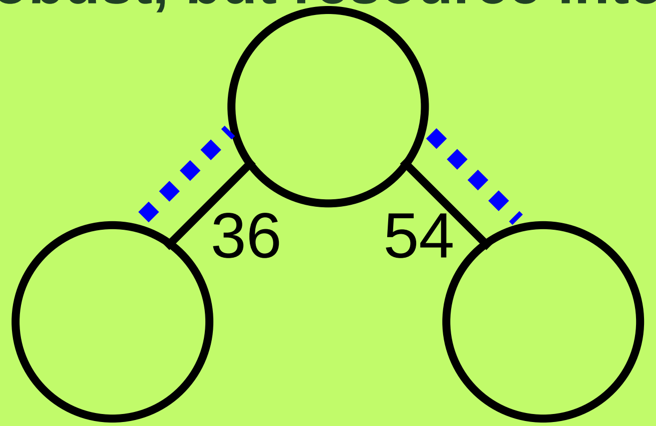

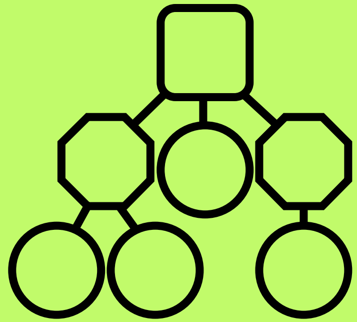

Ack-oriented in-mesh multicast (AM)

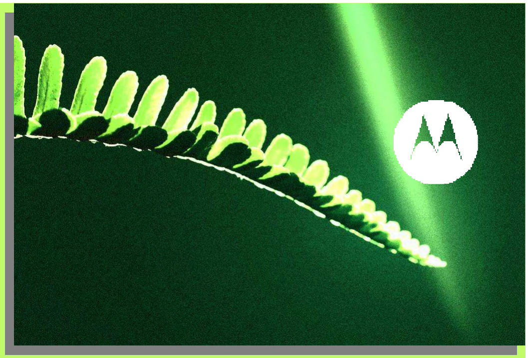

- **Phase 1 - Distribute data packet - flood like FB**
- **Phase 2 - Leaf routers unicast PACKs**

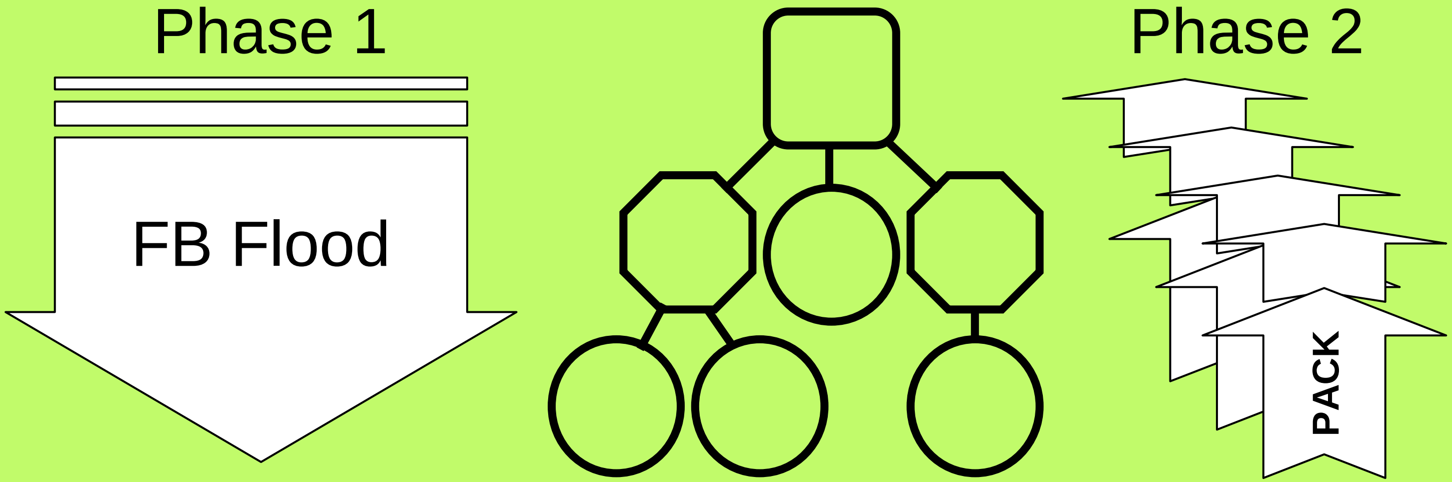

- **Repeat - If PACKs are not received from each leaf router**

Ack-oriented in-mesh multicast (AM)

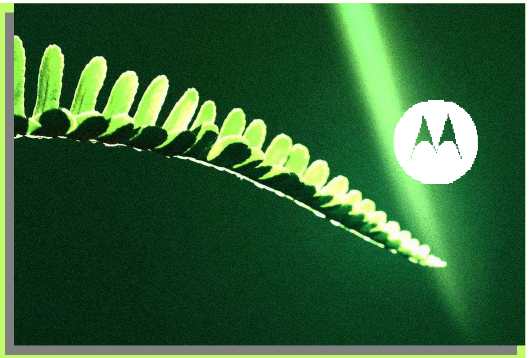

- Source (MP) buffers packet and multicasts it
- Phase 1 - like FB
- Use “fastest” data rate like FB
- Phase 2 - upon receipt of data, leaf routers send unicast PACKs (packet acknowledgements)
- MP maintains list of leaf routers from which it received PACK
- If PACKs are not received from all leaf routers, MP re-floods packet

- **Fast & robust**
- **Resource expensive?**
- **Several optimizations are possible**

Simple analysis of algorithms without loss

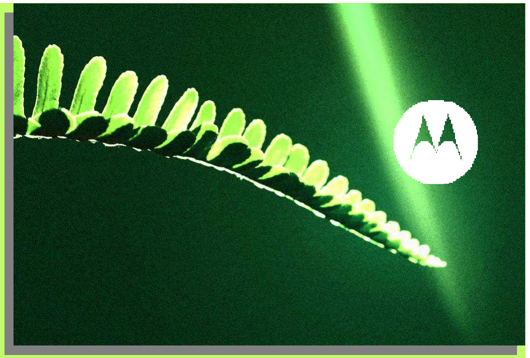

- IEEE 802.11s used for unicast routing and multicast tree building
- MP
- Intermediate routers
- Leaf routers
- Leaf routers do not re-transmit

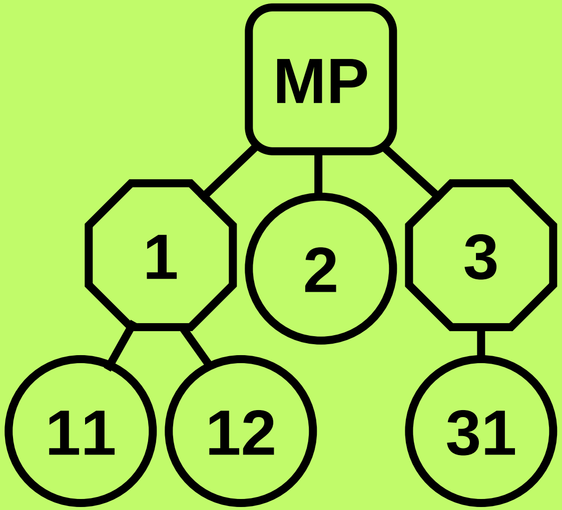

- All links are IEEE 802.11a
- All links are capable of 36mbps
  Engineered mesh network deployment
- DB example
  3 multicast transmissions at 6mbps

Simple analysis of algorithms without loss

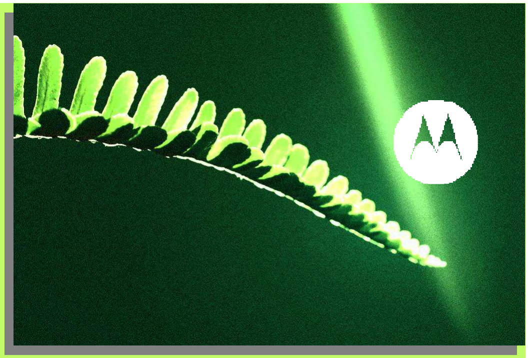

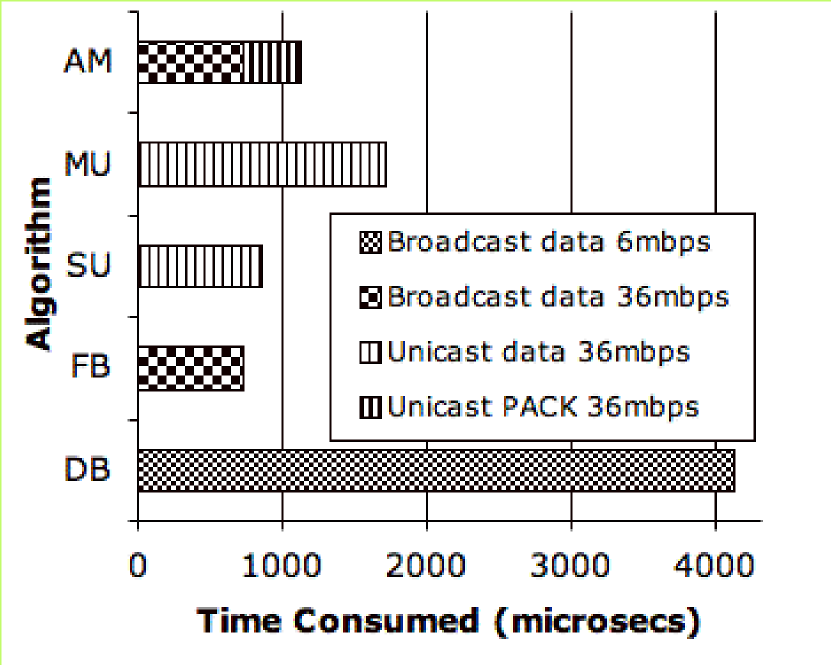

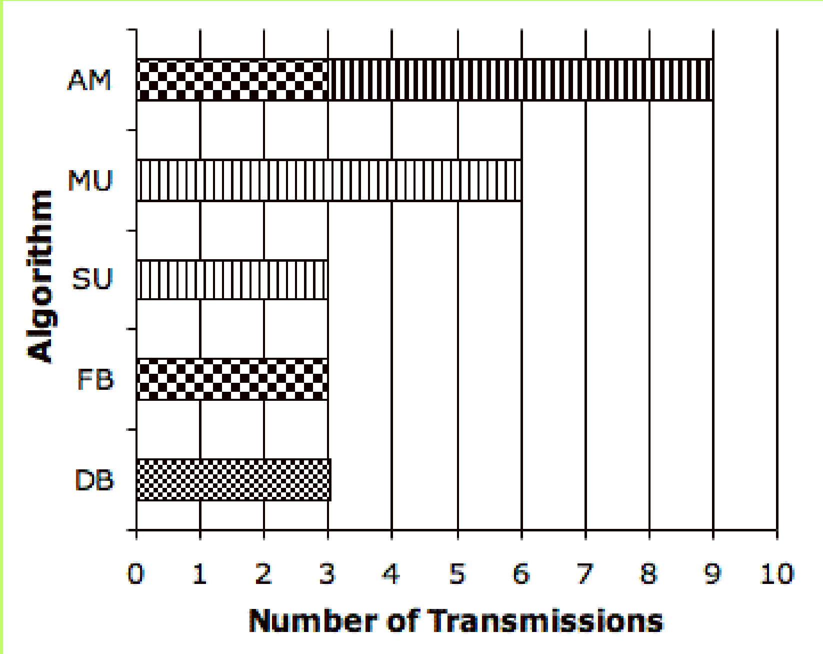

Simple analysis of algorithms with one loss

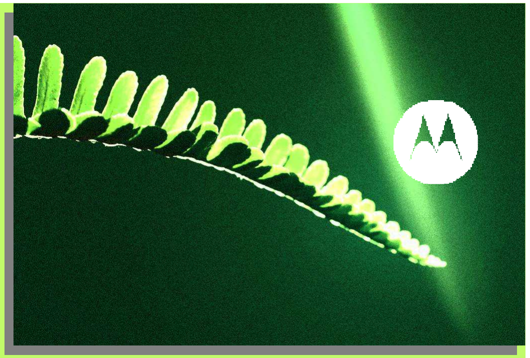

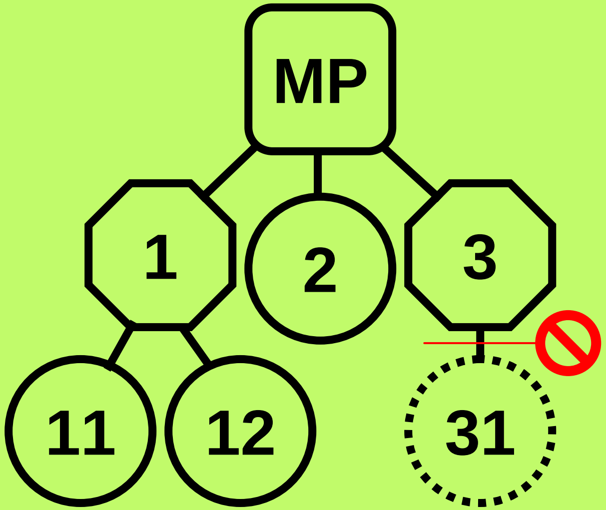

Simple analysis of algorithms with one loss

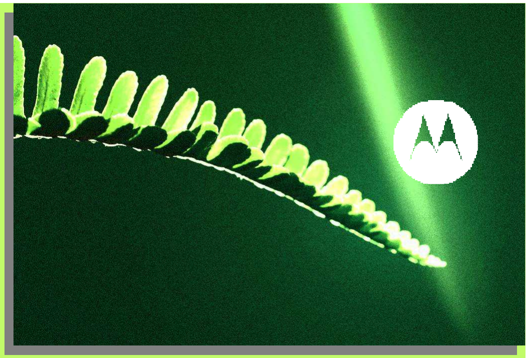

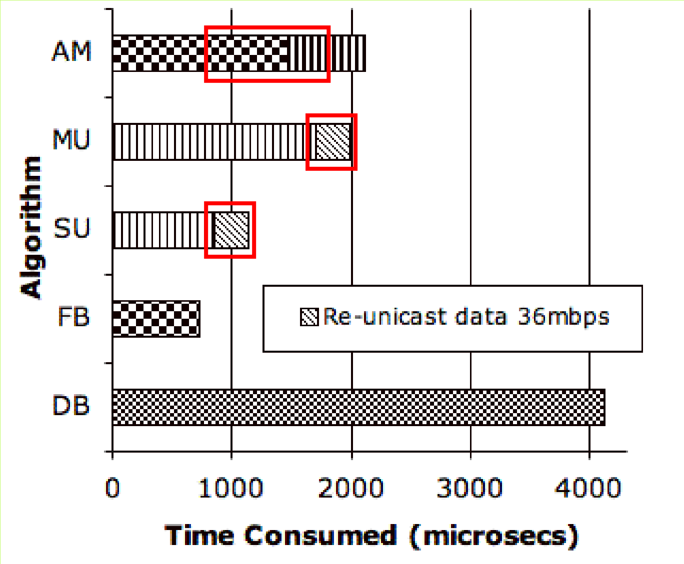

-1
-1

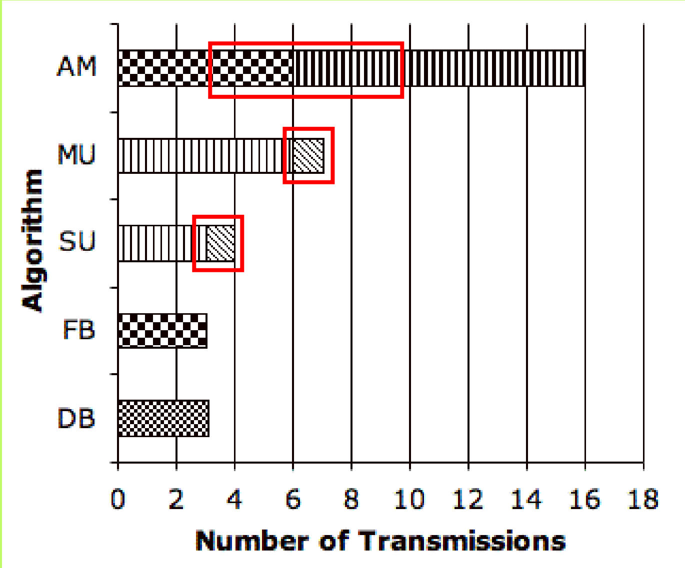

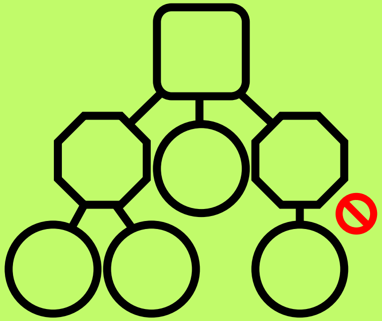

Cascading losses

- A single loss can result in a portion of the network not receiving a multicast packet
- Example

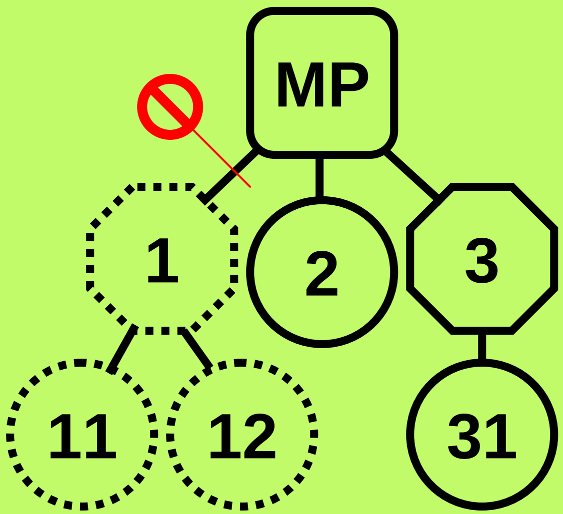

Simulation

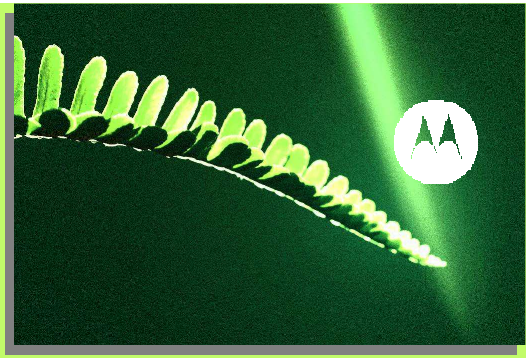

- Increasing load without loss
- Increasing loss with fixed load
- **OPNET™ 11.5**
- **Fixed L2 data rates**
- **Fixed packet reception error rates**
- **CBR traffic, 1000-byte packets**
- **No RTS/CTS, *RTS/CTS also examined**
- **Walls to create multihop network**
- **Simulation traffic runs for 30 sec after reaching steady state, *other times tested**
- **Five runs with different seeds averaged**

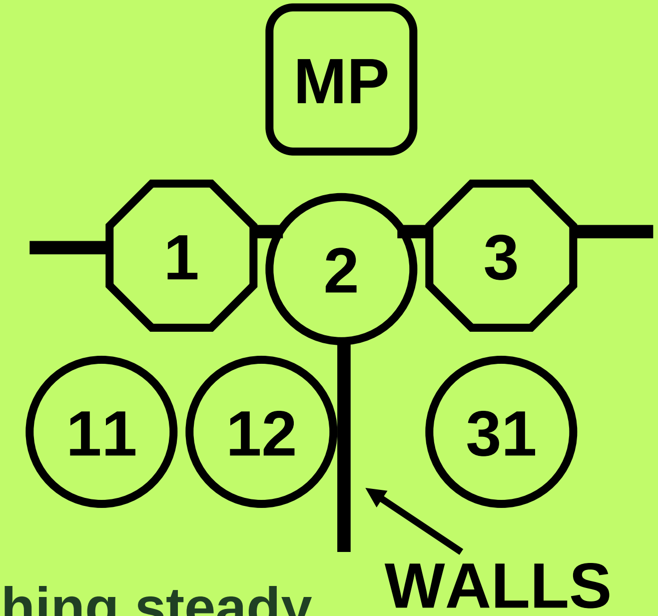

Metrics

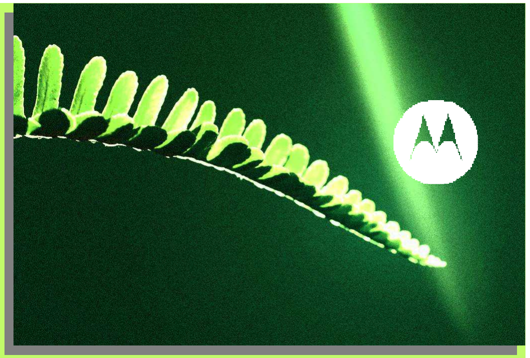

- **Packet delivery ratio (PDR)**

  The ratio of multicast data packets received divided by the number of multicast packets generated at the source. This metric indicates the goodput of the multicast protocol as well as its robustness to errors.

- **End-to-end delay (delay)**

  Time between when the source generates a packet and the time a receiver receives it. This metric indicates the amount of time it takes to deliver a data packet.

- **Wireless resources utilized**

  The number of transmissions and the amount of wireless channel time these transmissions consume.

The metrics above are averaged at each receiver, then the minimum (minimum receiver average), maximum (maximum receiver average) and average of all receivers are displayed in the graphs.

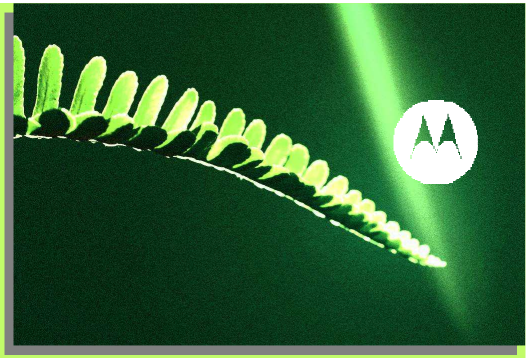

Load results without loss

- Data points and min/max bars represent range of average performance at each of the different receiving routers
- Algorithms perform relative to their efficiency until their congestion point

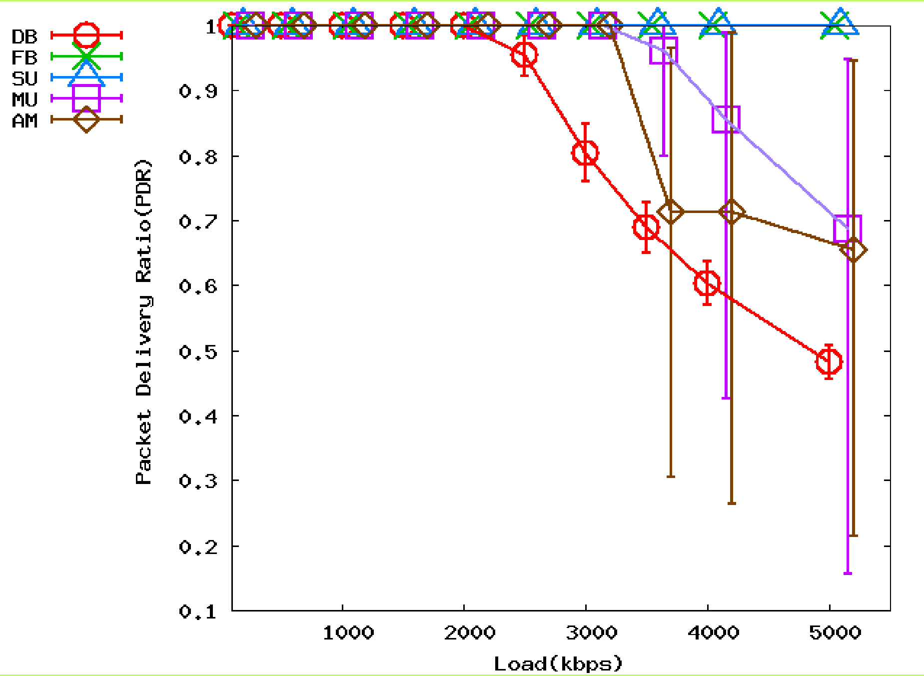

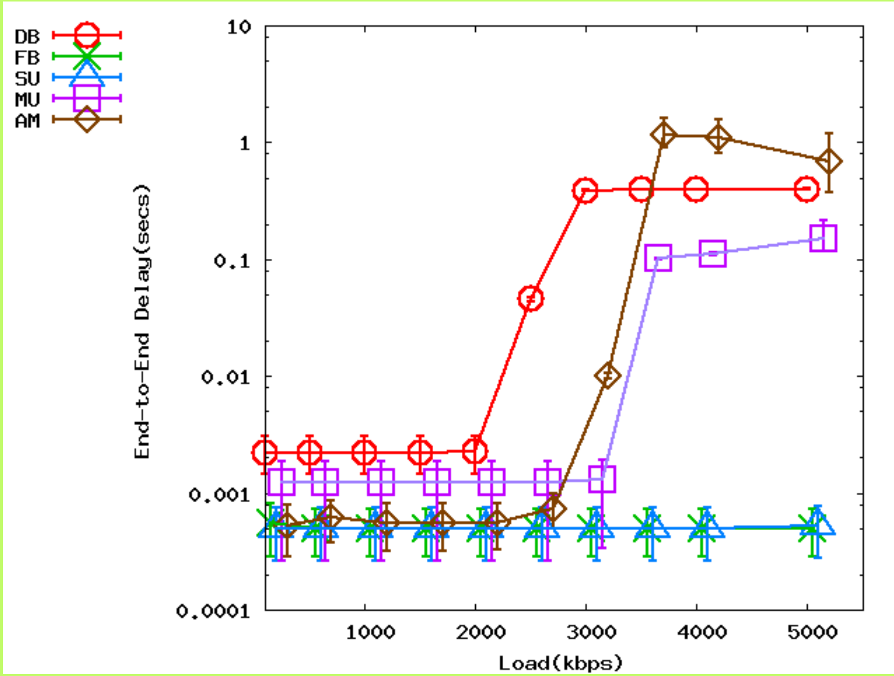

Load results without loss

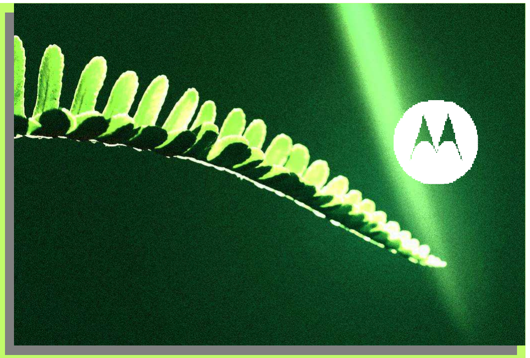

- DB consumes the most medium resource time and therefore congests first
- AM congest before MU even though MU utilizes more medium resource time because AM has several more transmissions. Medium access delays due to more competition cause AM to congest
- FB and SU do much better than others, lowest number of transmissions, highest number of potential receivers

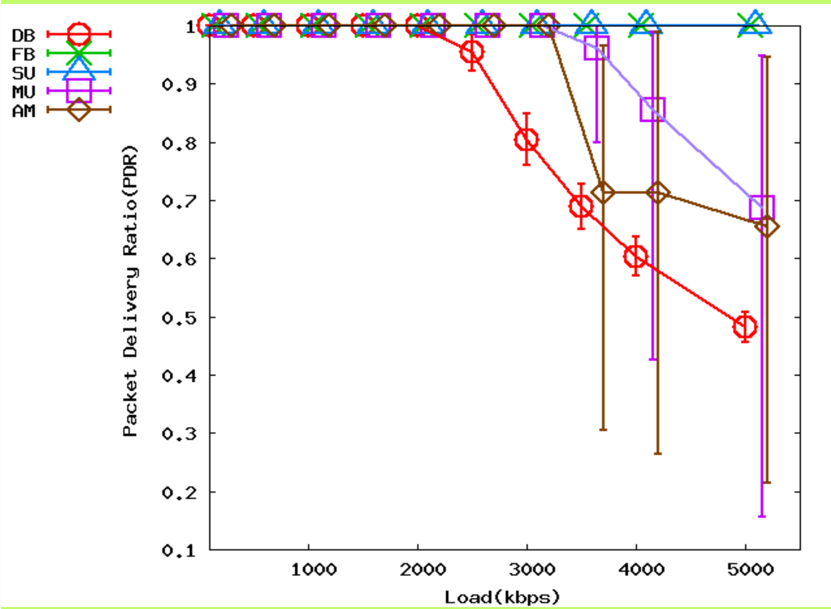

Algo	Congestion Point
AM	3 mbps
MU	3.5 mbps
SU	8 mbps
FB	10 mbps
DB	2.5 mbps

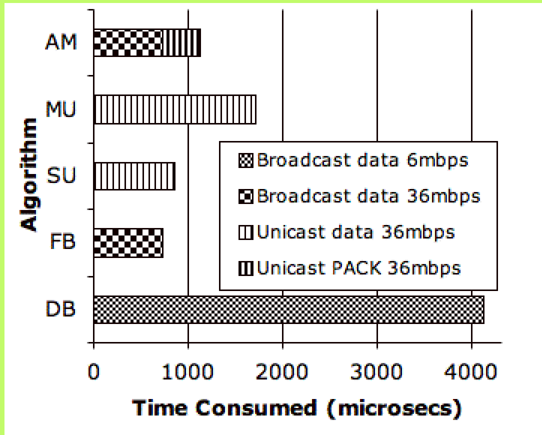

Results with loss below congestion point

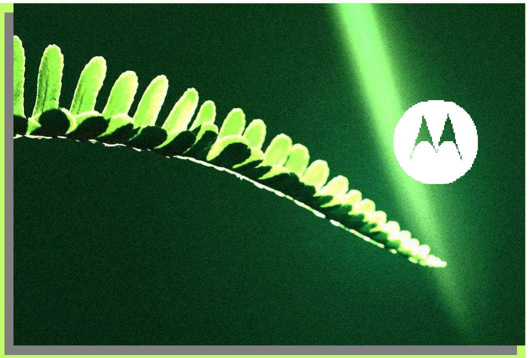

- DB and FB degrade linearly with the loss rate
- SU tends to degrade more slowly with increasing error, since it recovers from errors to the chosen receiver
- AM degrades even less than SU with increasing error
- MU does not experience any error, since each receiver performs retransmissions for robustness
- For DB, FB, SU, and MU all the packets are delivered quickly.
- For AM the packet delay is significant. The large delay value is due to the re-flood transmission timer; one second

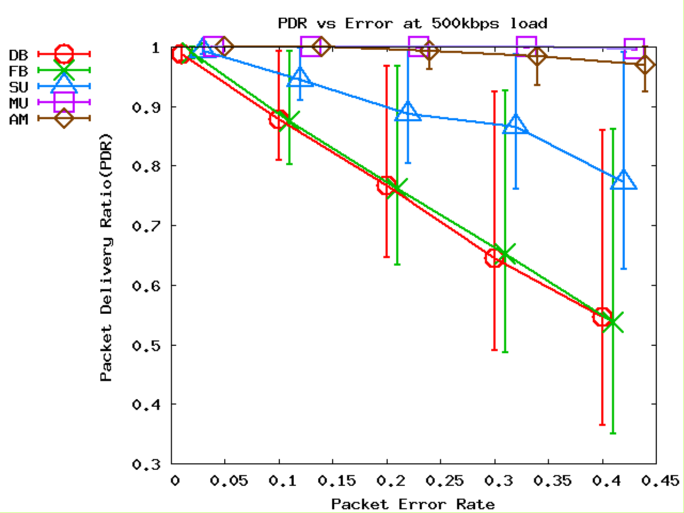

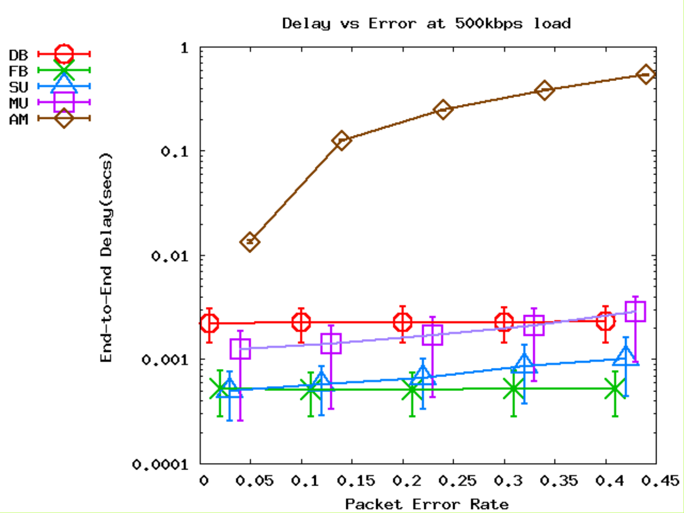

Robust mechanisms resource usage with 30% loss

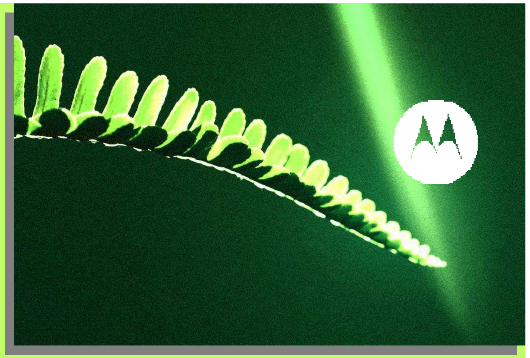

- Robust mechanisms can quickly consume resources, doubling or tripling resource usage
- Robust mechanisms result in less loss and lower delay when resources are available

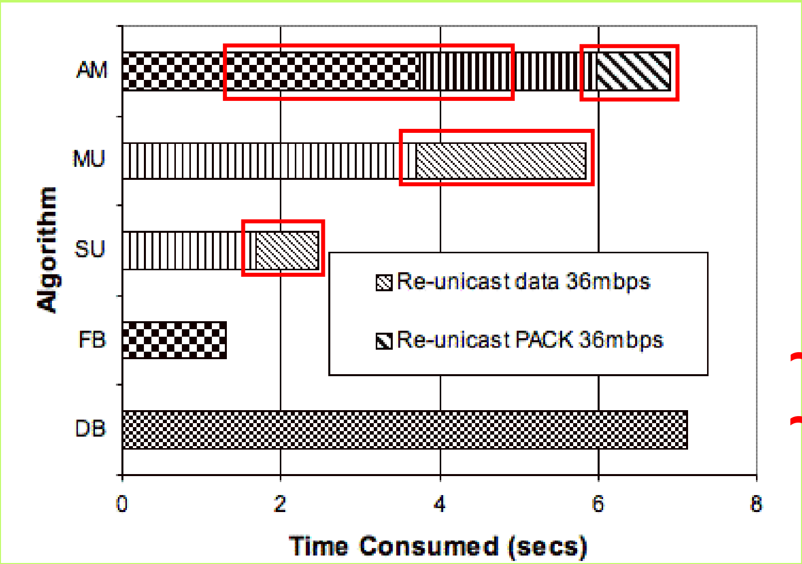

~30%
~30%

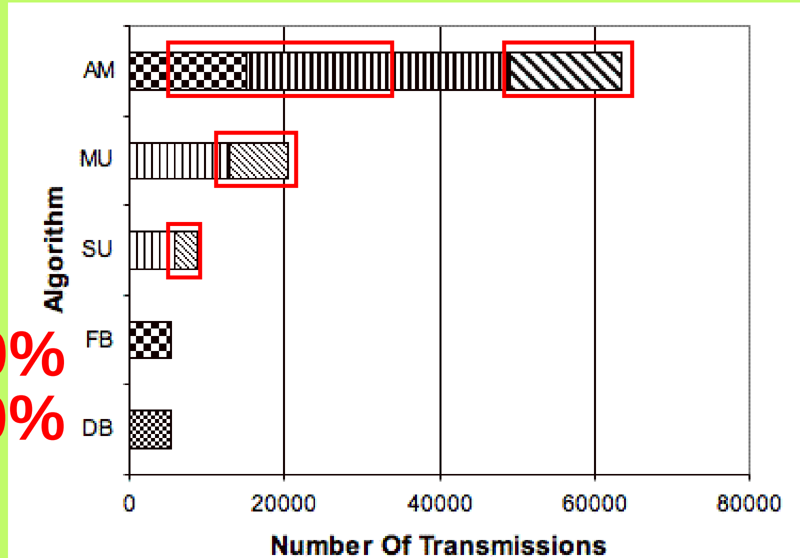

Effect of RTS/CTS

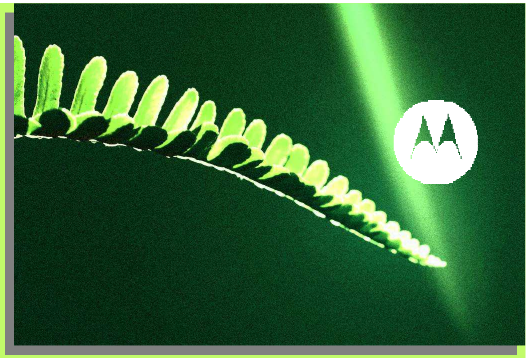

- **RTS/CTS for unicast can add significant overhead**
- **When RTS/CTS is utilized the congestion point occurs at a much lower load for all algorithms utilizing unicast**

Algorithm	**Congestion Point no RTS/CTS**	**Congestion Point With RTS/CTS**
DB	2.5 mbps	2.5 mbps
FB	10 mbps	10 mbps
SU	**8 mbps**	**4 mbps**
MU	**3.5 mbps**	**1.5 mbps**
AM	3 mbps	3 mbps

Multicast conclusions

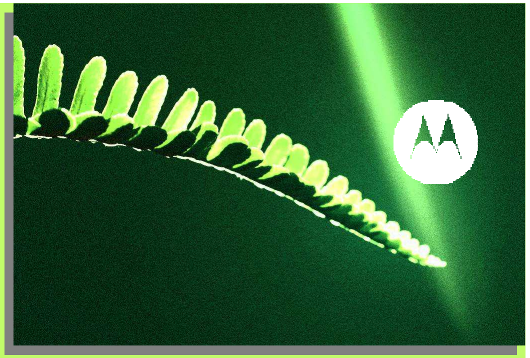

- Robustness mechanisms have desirable effects when their mechanisms don't induce congestion
- Large return on using higher data rates
- In-mesh robustness measures can be expensive
- Traffic load, receiver loss statistics, traffic type, expected/predicted number of redundant receptions, number of downstream devices
- Which technique(s)? When? Policing & enforcement

Summary

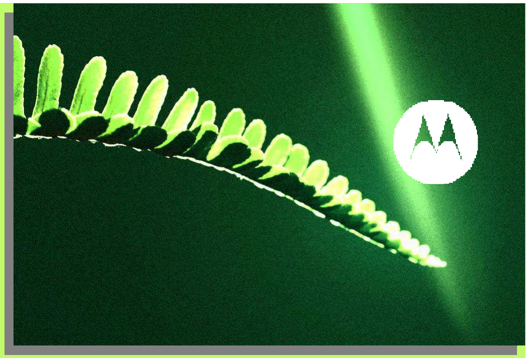

- **Robustness mechanisms have desirable effects when their mechanisms don't induce congestion**
  - **In this region, robustness can be traded for goodput**
- **AM robustness measures are expensive**
  - **MP buffers packets, messaging, loss behavior, delay**
  - **Occupy many resources**
- **In low loss scenarios, FB almost always does better than robust algorithms**
- **In high loss scenarios, SU can provide increased robustness without much increase in resource utilization**
- **Packet delivery delay after congestion is dominated by queuing delay**
- **AM delay is long in comparison to other robustness approaches**

Future Work

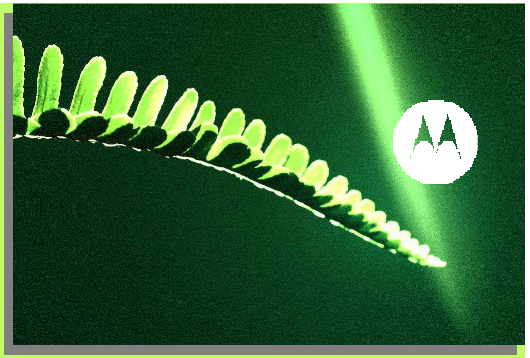

- **Heterogeneous links**
- **Multiple traffic classes**
- **Multiple traffic types (unicast/multicast)**
- **NACK-oriented robustness schemes**
- **In-mesh FEC**

fast efficient and robust networking FERN

Questions

Ian Chakeres - ian.chakeres@motorola.com
Chivukula Koundinya - koundyc@gmail.com
Pankaj Aggrawal - pankajaggrawal@motorola.com