Pathological Behaviors for RLM and RLC

Arnaud Legout/Ernst W. Biersack

Eurecom

Arnaud.Legout@eurecom.fr

Outline

- Motivation.
- RLM reminder.
- RLM pathological behaviors.
- RLC reminder.
- RLC pathological behaviors.
- Conclusion.

Motivation

- RLM is the first and most popular receiver-driven layered multicast congestion control protocol:
  - Only a few studies about RLM exist that essentially show it performs reasonably well.
- RLC is a TCP-like version of RLM:
  - Not aware of any studies about RLC.
- We present simple scenarios where RLM and RLC exhibit fundamental pathological behaviors:
  - Fundamental: the problems are inherent to the protocol itself.
  - Pathological: we observe undesirable behaviors that significantly reduce the performance of RLM/RLC.

RLM Reminder

- Receiver-driven cumulative layered multicast Congestion Control (CC) protocol for video dissemination:
  - The video stream can be organized and striped in **cumulative** layers.
  - Multicast capable network.
  - The source sends each layer on a different multicast group.
- All the protocol machinery is at the receiver side (receiver-driven).
- Timers:
  - Join timer T_j: Periodicity of the join experiments.
  - Detection timer T_d: Estimation of the time to decide if a join experiment has succeeded.

RLM reminder

- Bandwidth inference mechanism:
  - Make a join experiment every T_j (multiplicatively increased when a join experiment fails, reduced (relaxed) every T_d). Add the layer if the join experiment succeeds, i.e. no loss during a T_d after the join experiment has started.
  - In case of loss observed wait a T_d in the hysteresis state. Drop a layer if at the end of the hystereris period there is more than 25% loss rate (i.e. congestion). Only **one** layer dropped per T_d.
- Shared learning: In case of join experiment, send a message to the whole group. Precludes a join experiment at a higher layer while there is an experiment for a lower layer. Receivers learn from failed join experiment of the other receivers.

RLM Simulations: Convergence

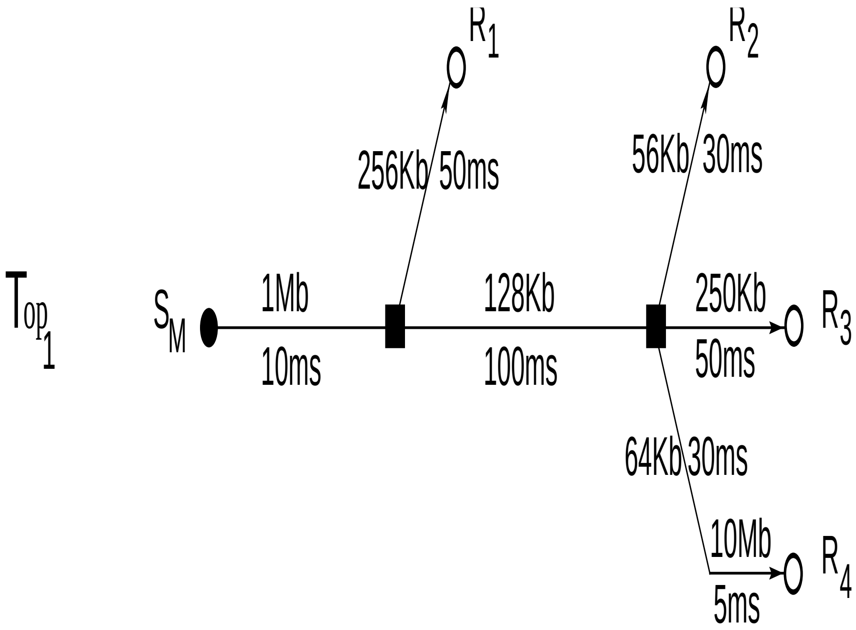

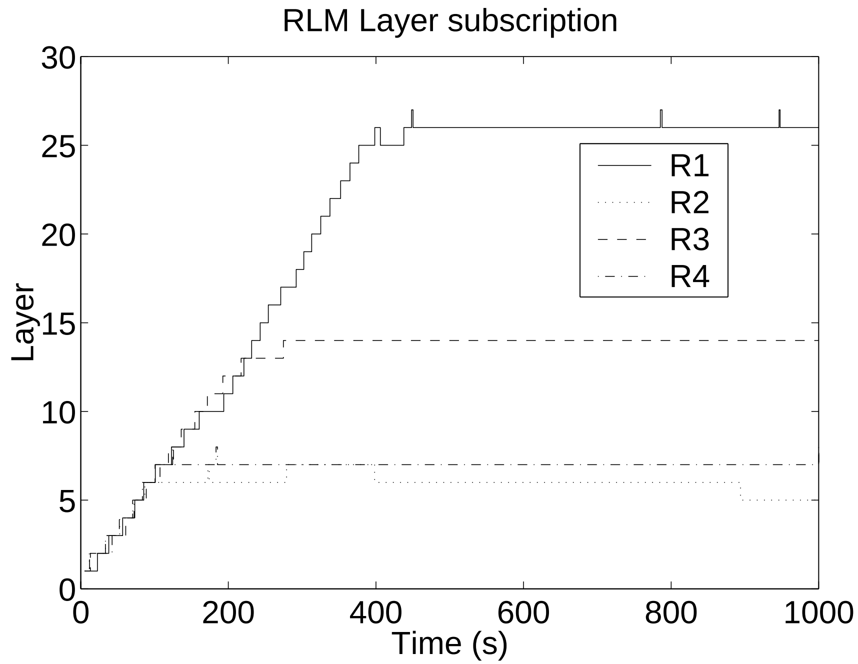

- Evaluation of the speed, stability, and accuracy of the RLM convergence in the context of a large heterogeneity of delay and bandwidth
- 10Kbit/s per layer (tough test).

- Very slow convergence (Minimum join timer T_j set to 5 seconds).
- 3.2% mean loss rate (25% loss threshold).
- Low number a join experiments.

RLM Simulations: Scalability

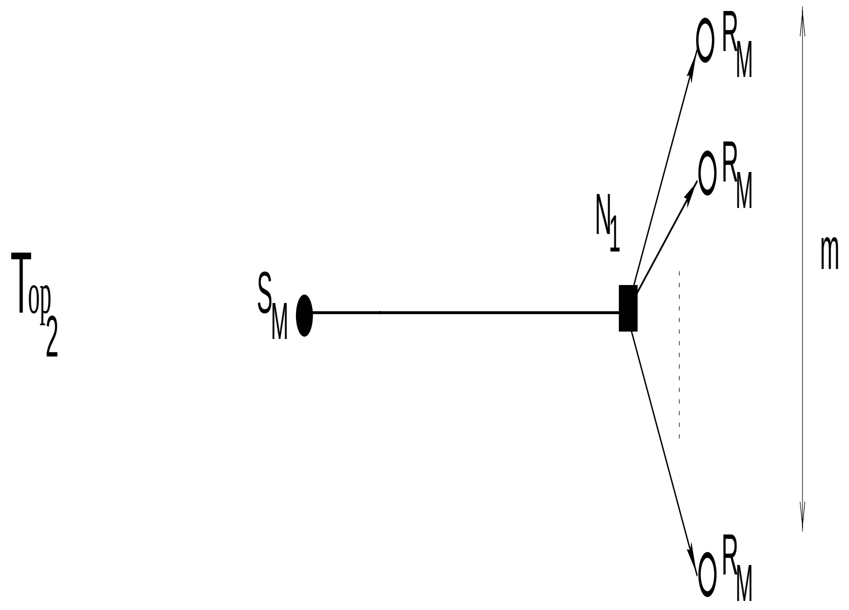

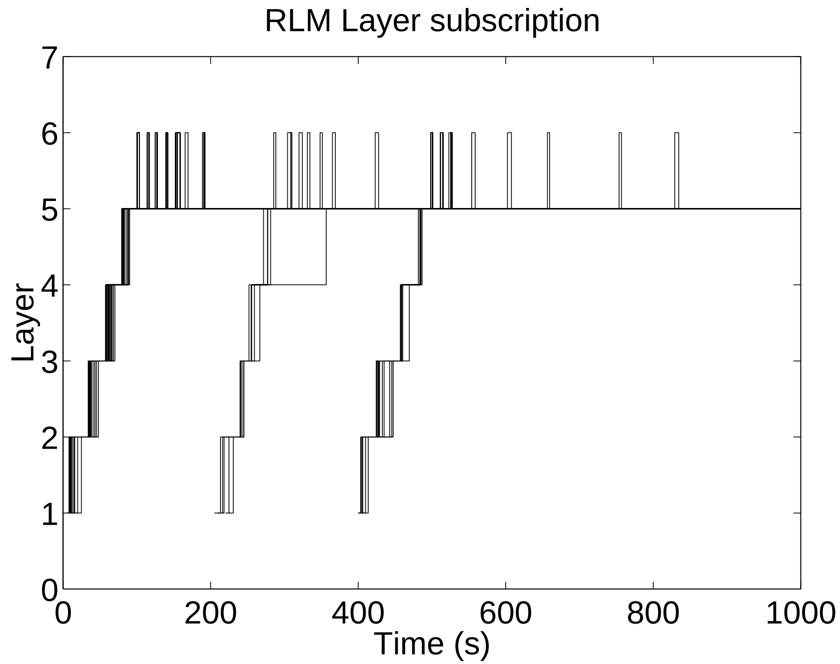

- Evaluation of the RLM scalability with the number of receivers and with late joins
- 50Kbit/s per layer.
- 20+5+5 receivers.
- Receiver synchronization due to the shared learning (precludes joining an upper layer while there is a join experiment for a lower layer).

RLM Simulations: Dynamics

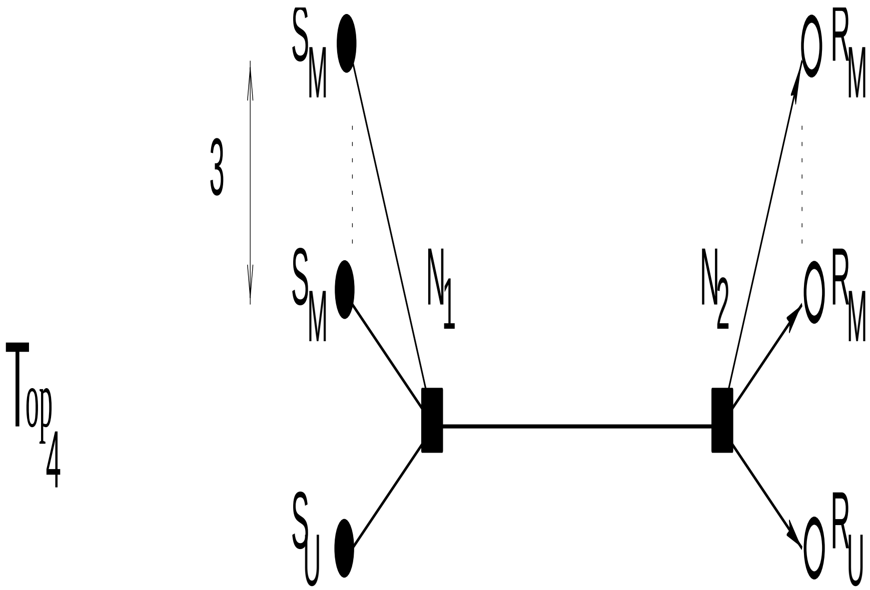

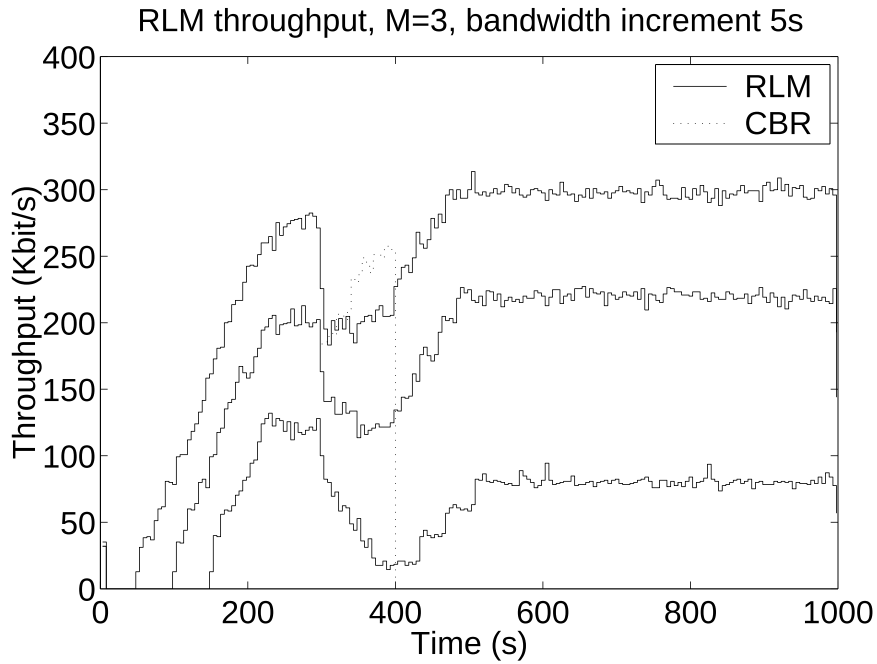

- 3 RLM + 1 CBR. Evaluation of the scalability of RLM with the number of session, RLM adaptation to heavy congestion.
- 20 Kbit/s per layer.
- Slow convergence (Min T_j=5 s).
- High unfairness.

RLM Simulations: Dynamics

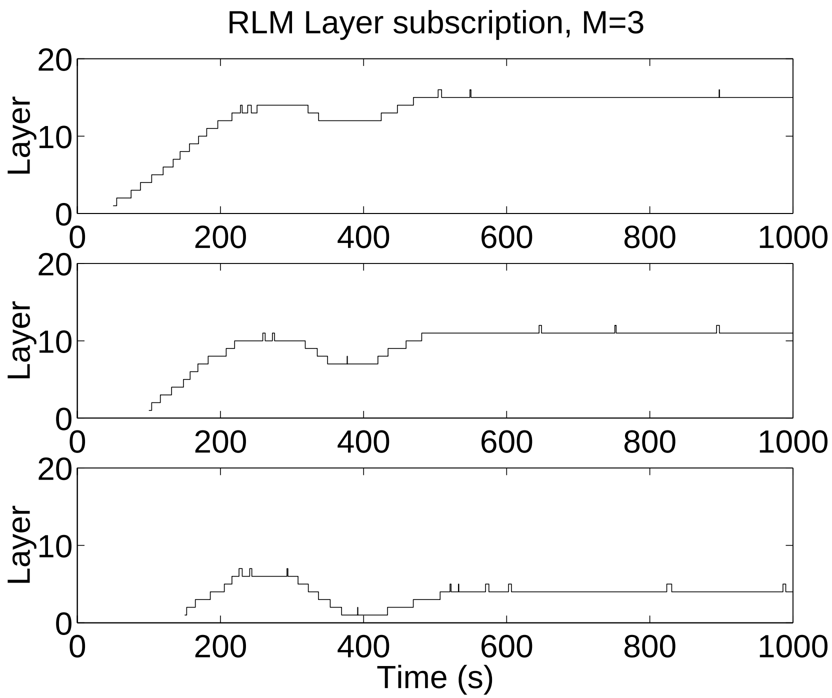

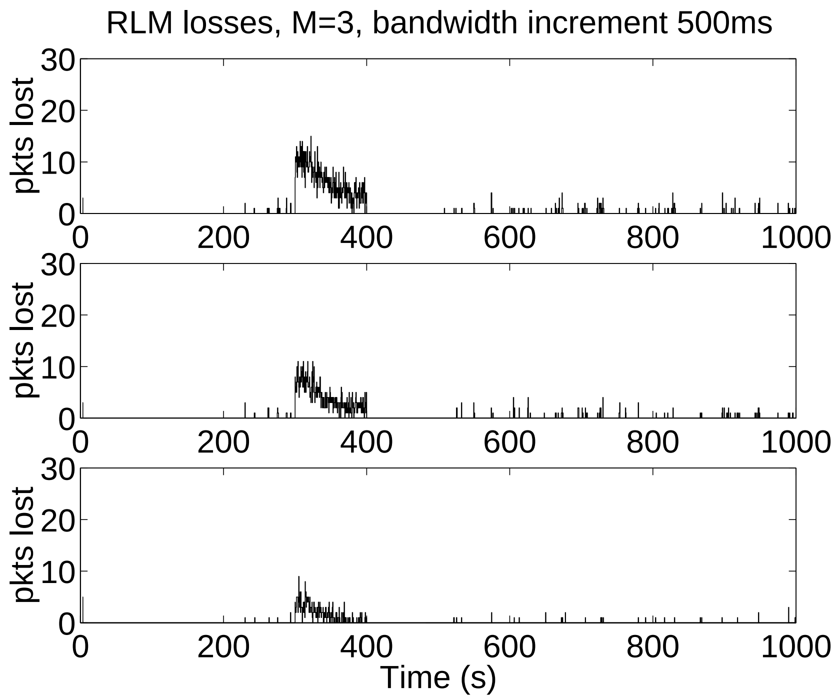

- The process of dropping layers is very conservative (one layer dropped per detection timer).
- High number of losses in case of congestion: 2.3% mean loss rate.

RLM Simulations: RLM and TCP

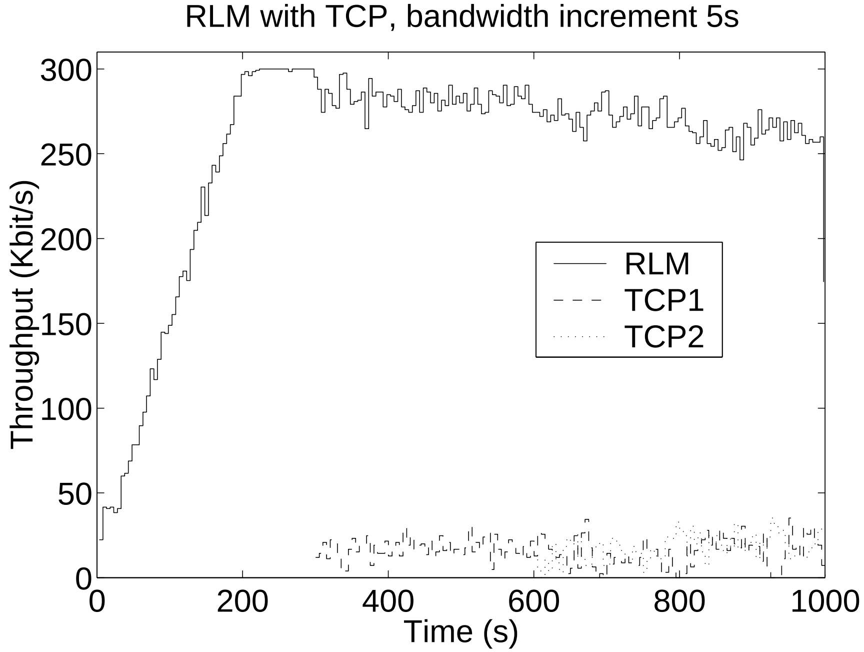

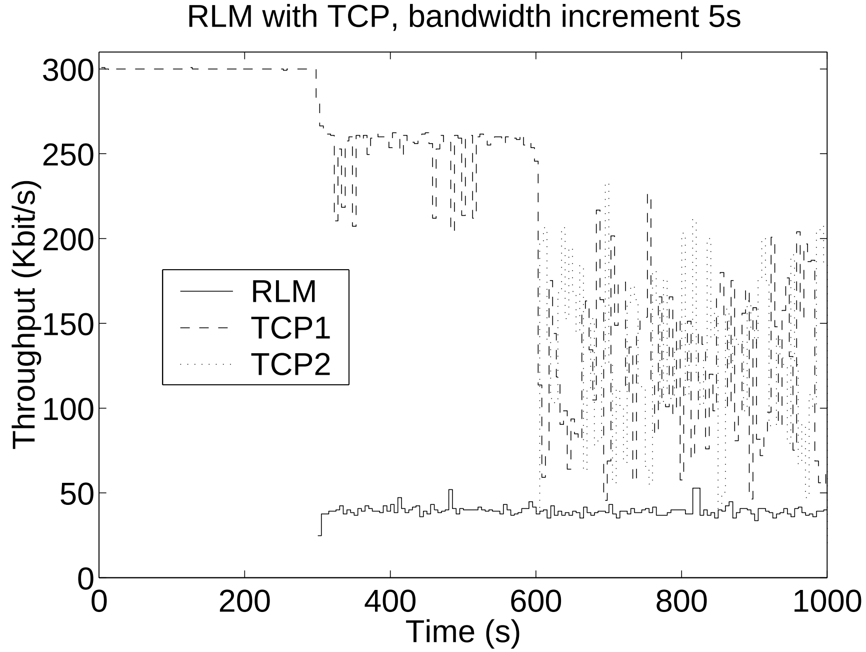

- 1 RLM + 2 TCP. RLM starts first.
- RLM gets all the available bandwidth. RLM needs to experience high losses to drop a layer: loss threshold is 25%. TCP cannot grab bandwidth.
- 20 Kbit/s per layer

- 1 RLM + 2 TCP. RLM starts after TCP1.
- RLM is unable to grab bandwidth. A join experiment succeeds only when there is no loss during a detection timer period. TCP produces at least a loss per cycle.

RLM Pathological Behaviors

- Minimum join timer (slow convergence):
  - Tradeoff between speed of convergence and periodic congestion due to the join experiments.
- High loss threshold/Hysteresis state (high loss rate, starves TCP when RLM starts first):
  - Tradeoff between a conservative and a reactive behavior in case of losses.
- Shared learning (receiver synchronization), Conservative join experiments (TCP starves RLM when TCP starts first):
  - Foundations of RLM.
- Conservative layer drop process (transient periods of high congestion):
  - Necessary to avoid cascade drops, very hard to tune.

RLC Reminder

- Receiver-driven cumulative layered multicast CC protocol:
  - Data that can be organized and striped in **cumulative** layers.
  - Multicast capable network.
  - The source sends each layer on a different multicast group.
  - Layers exponentially distributed.
- Bandwidth inference: **Periodic bursts** (double the throughput for a short **fixed** period of time) followed by an idle period.
- Synchronization points (SP) on each layer, spaced proportionally to the bandwidth of the corresponding layer (exponentially), and always located at the end of a burst.

RLC Reminder

- Mechanism:
  - Add a layer at a SP if no losses are experienced during the burst preceding that SP.
  - Drop a layer on congestion (a loss), one layer drop per **deaf period** (fixed value). TCP-like behavior: exponential decrease in case of loss.

RLC Simulations: Convergence

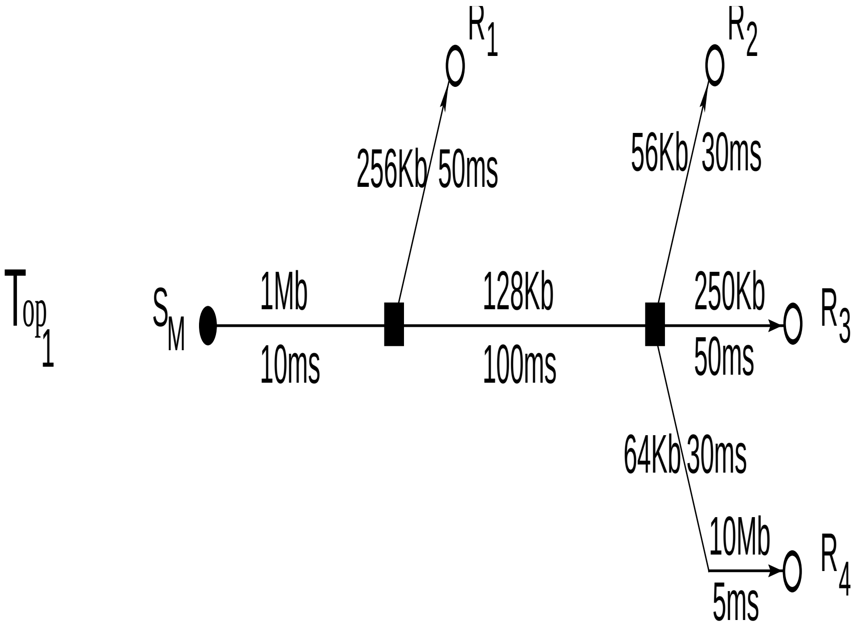

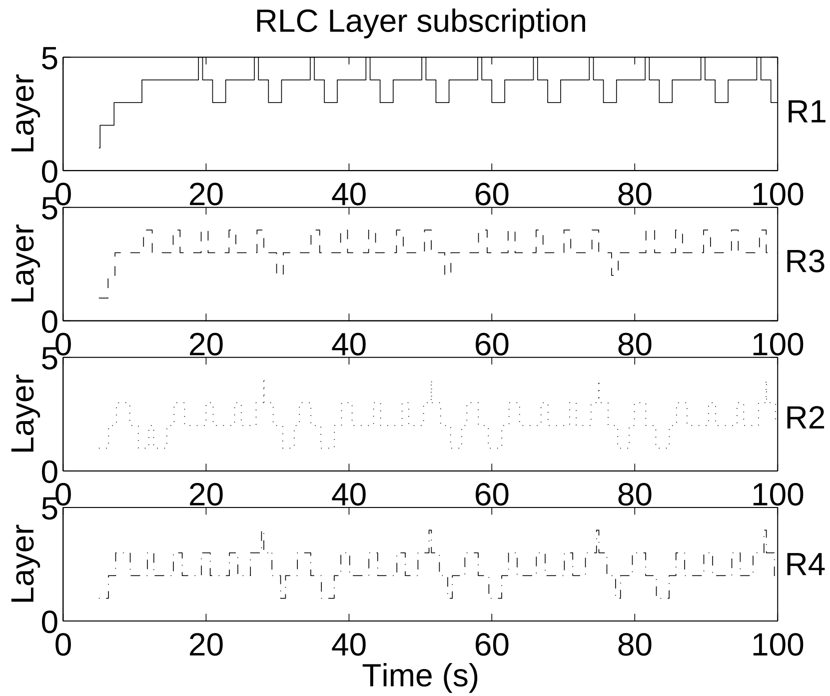

- Evaluation of the speed, stability, and accuracy of RLC convergence in the context of a large heterogeneity of delay and bandwidth.
- Base layer: 32 Kbit/s.
- The periodic bursts do not succeed to make the bottleneck queue overflow (erroneous bandwidth inference).
- Mean loss rate: 13%.

RLC Simulations: Scalability

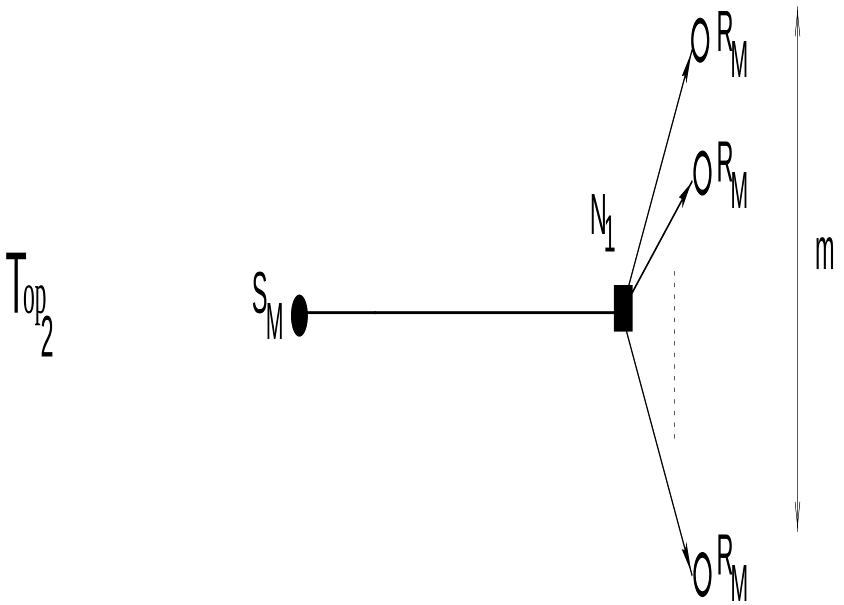

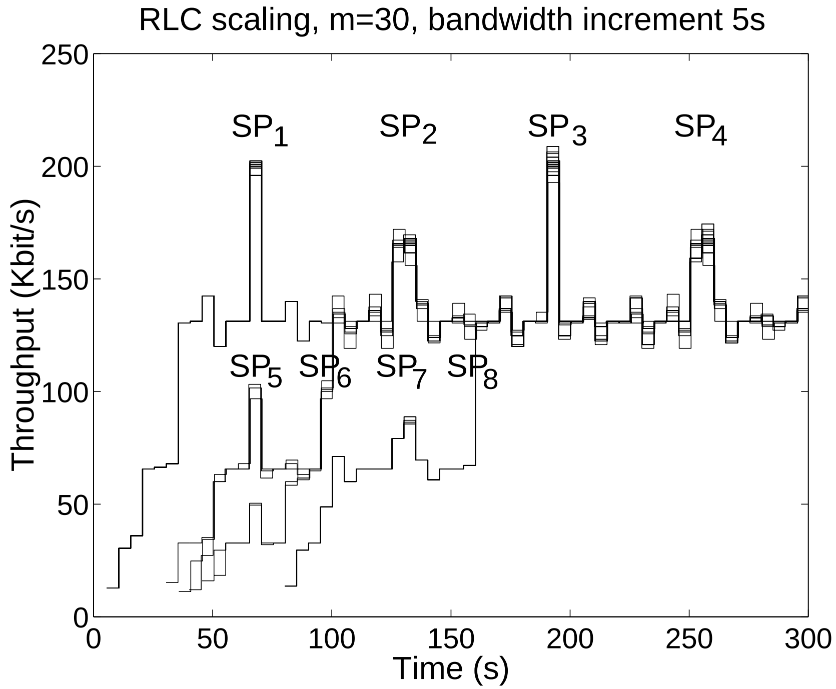

- Evaluation of the RLC scalability with the number of receivers and with late joins.
- 20+5+5 receivers.
- At SP1 a burst creates congestion.
- As SP5 and SP1 are synchronized, the late receivers must wait until SP6 to add a layer.

RLC Simulations: RLC and TCP

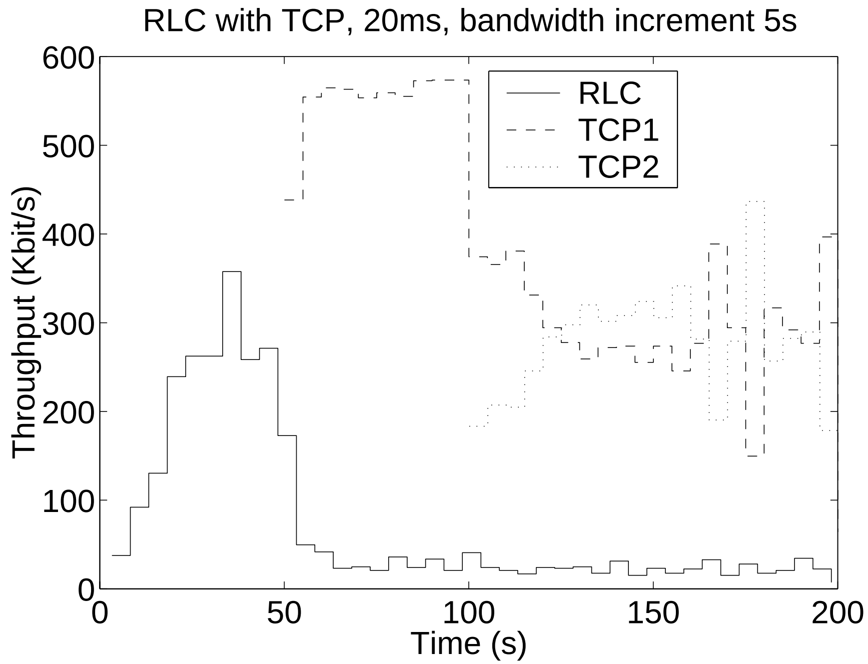

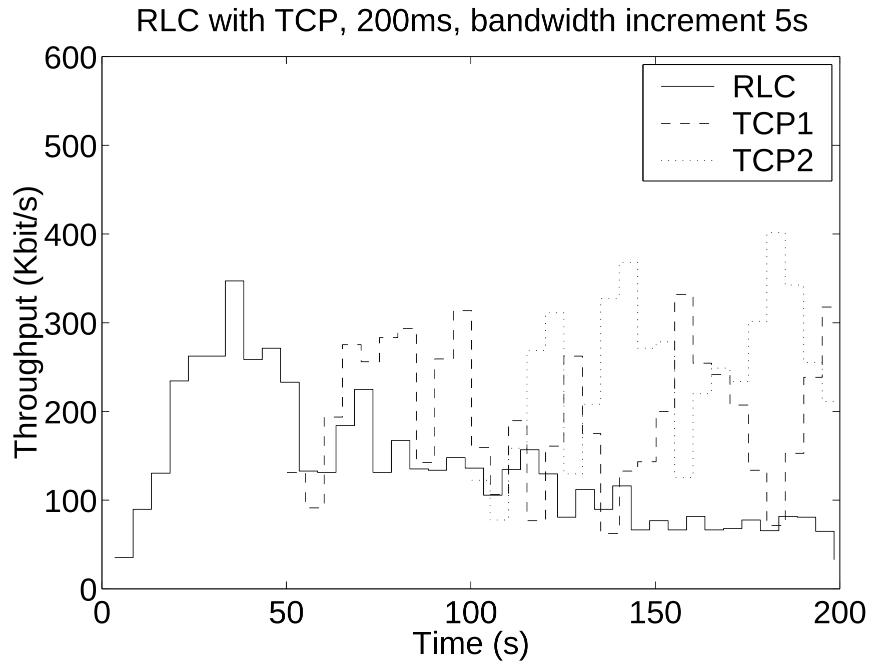

- 20 ms bottleneck link delay.
- RLC shares unfairly the bandwidth with TCP.
- A small RTT leads to a small TCP cycle (frequent periodic losses).

- 200 ms bottleneck link delay.
- RLC shares fairly the bandwidth with TCP.
- A large RTT leads to a large TCP cycle (sparse periodic losses).

RLC Pathological Behaviors

- The bandwidth inference mechanism based on periodic bursts does no succeed (does not make the queue overflow):
  - We need to know how long the bursts should persist to make the queue overflow. Comes close to a new bandwidth inference mechanism.
- Synchronization points as distributed in RLC significantly slow down the convergence of the RLC receivers:
  - Open problem!
- TCP-like: responsive to losses but independent of the RTT (unfairness with TCP):
  - Open problem!

Conclusion

- RLM and RLC exhibit fundamental pathological behaviors:
  - RLM:
    - Slow convergence.
    - Sustained loss rate.
    - Receiver synchronization.
    - Conservative/aggressive with TCP.
  - RLC:
    - Poor bandwidth inference mechanism.
    - Slow convergence.
    - TCP-like but independent of the RTT.

Thanks!

Questions?