Università degli Studi di Padova

Case Study Analysis: Autonomous Production Site

Armir Bujari, Claudio E. Palazzi, Daniele Ronzani
Email: *cpalazzi@math.unipd.it*

Rise of the Working Robots

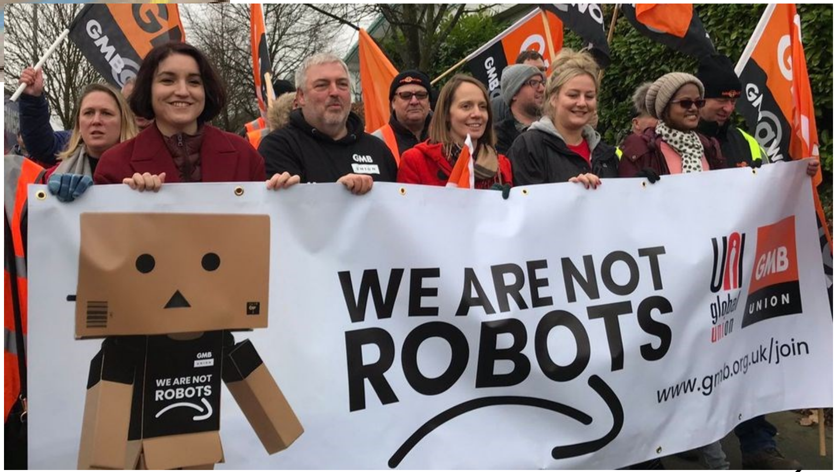

Rise of the Working Robots

- Small robots are more and more used in industry
- They have become famous thanks to Amazon's and Alibaba's warehouse (pick up) robots
- Are now used even in productions sites
  - E.g., to bring semi-finished products through their stages of production
- Operators have plans to automate almost every physical move in their facilities within the next 2/3 years
- Robotic automation can extend the capacity, hours of operation, and life of a production site
  - And do not complain about wages

https://www.youtube.com/watch?v=8gy5tYVR-28

https://www.youtube.com/watch?v=qU4YMDJNzpg

Automated Guided Vehicle (AGV)

- An AGV is a portable robot that follows markers or wires on the floor, or uses vision, magnets, or lasers for navigation

- AGVs have been used for case, pallet, bulk, or specialized container movement for decades across a wide range of industries and applications
  - **Now:** massive numbers and small devices

- **Problem statement:** **How many AGVs can be supported by current wireless network technology in a production site?**

- **Goal:** **understand the limitations of an AGV-based autonomous production site in terms of network capability**

Two (Control) Cases for AGVs

- Dynamic Traffic Control, fine grained (best performance)
  - One message sent/received by any AGV every 100ms
  - More messages allows fine remote control of AGVs, dynamic decisions about movements and task orders (no predetermined/fixed paths)
  - More messages may congest the network
    - If the network becomes congested, messages are lost or delivered with delay jeopardizing the performance of the system

- Dynamic Traffic Control, coarse grained (worst performance)
  - One message sent/received by any AGV every 500ms
  - Less messages only allows for static routes of AGVs (virtual/painted rails)
  - Less messages maintains the network uncongested

Experimental Scenario

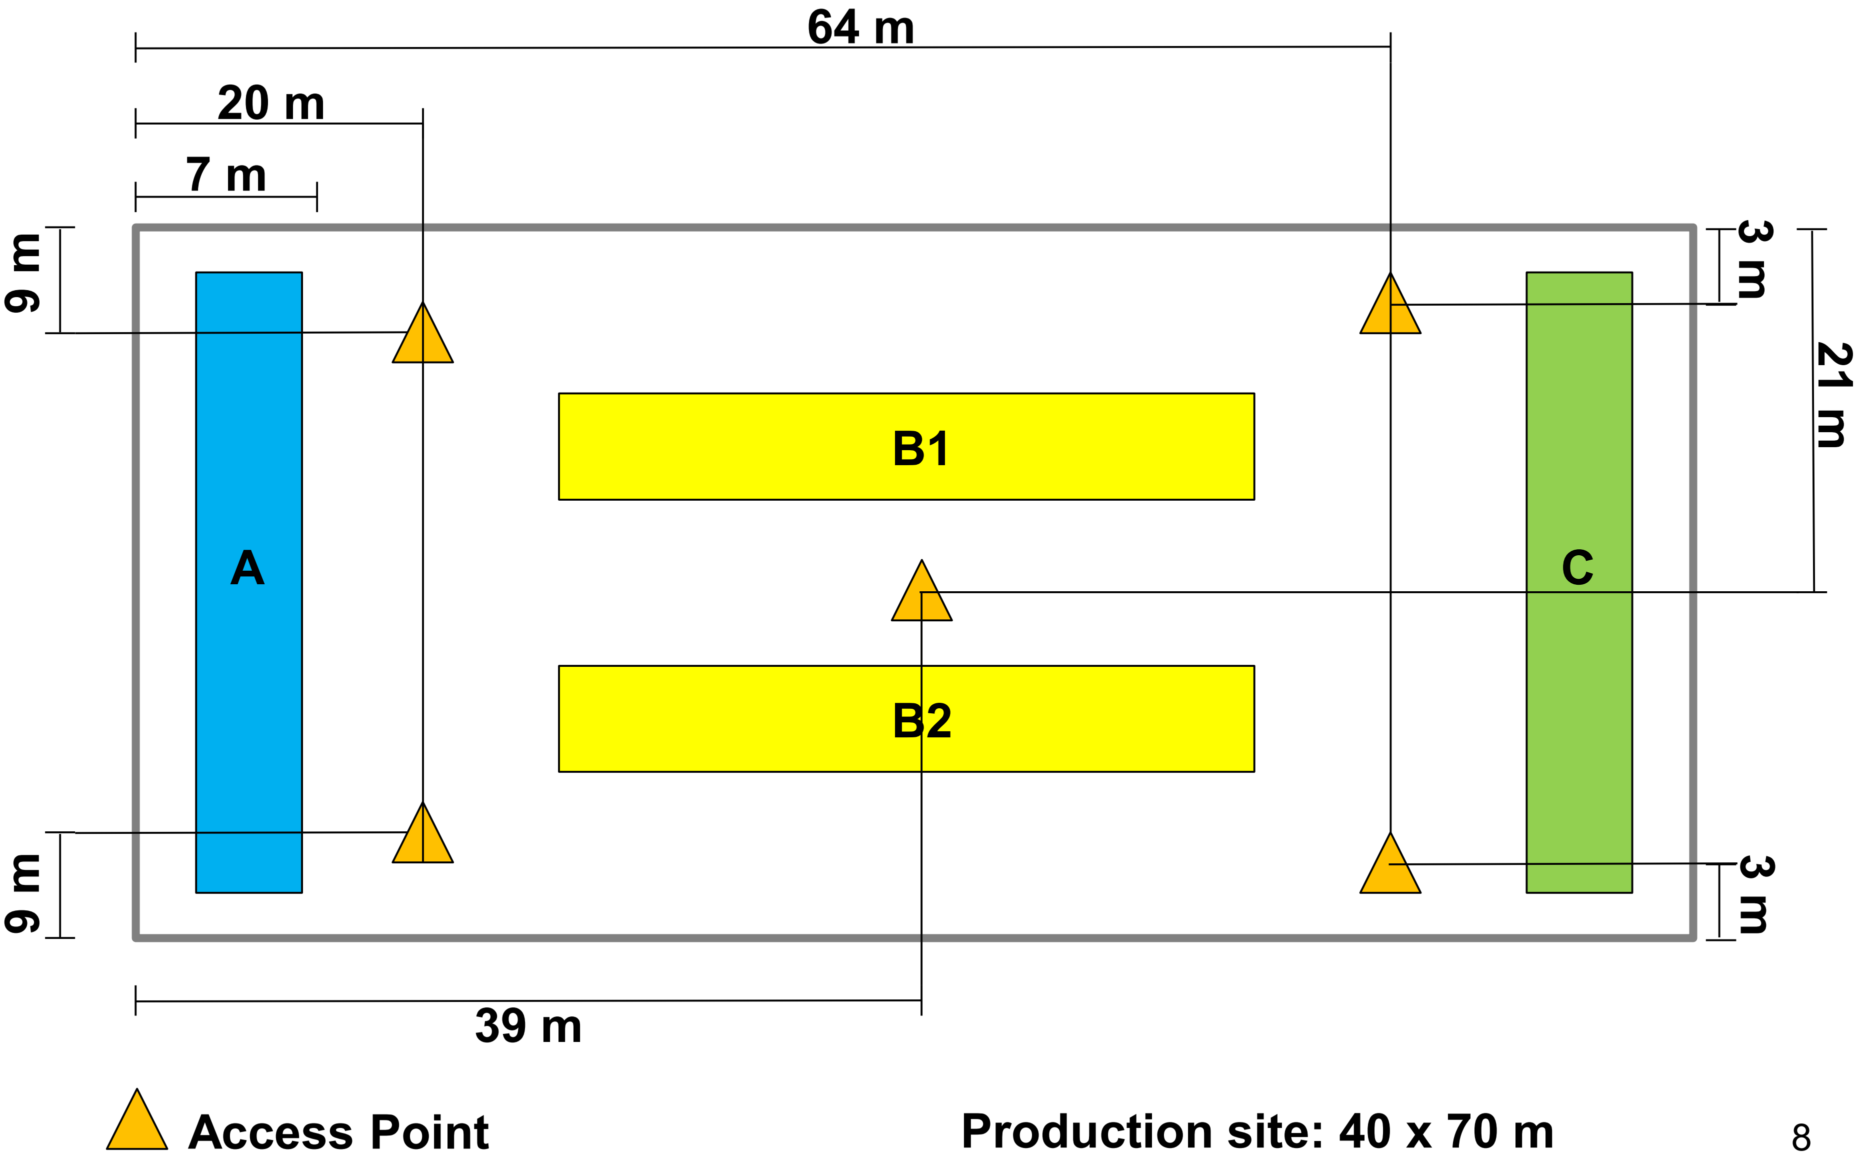

Production site: 40 x 70 m

Rover Movements - Partial Scenario

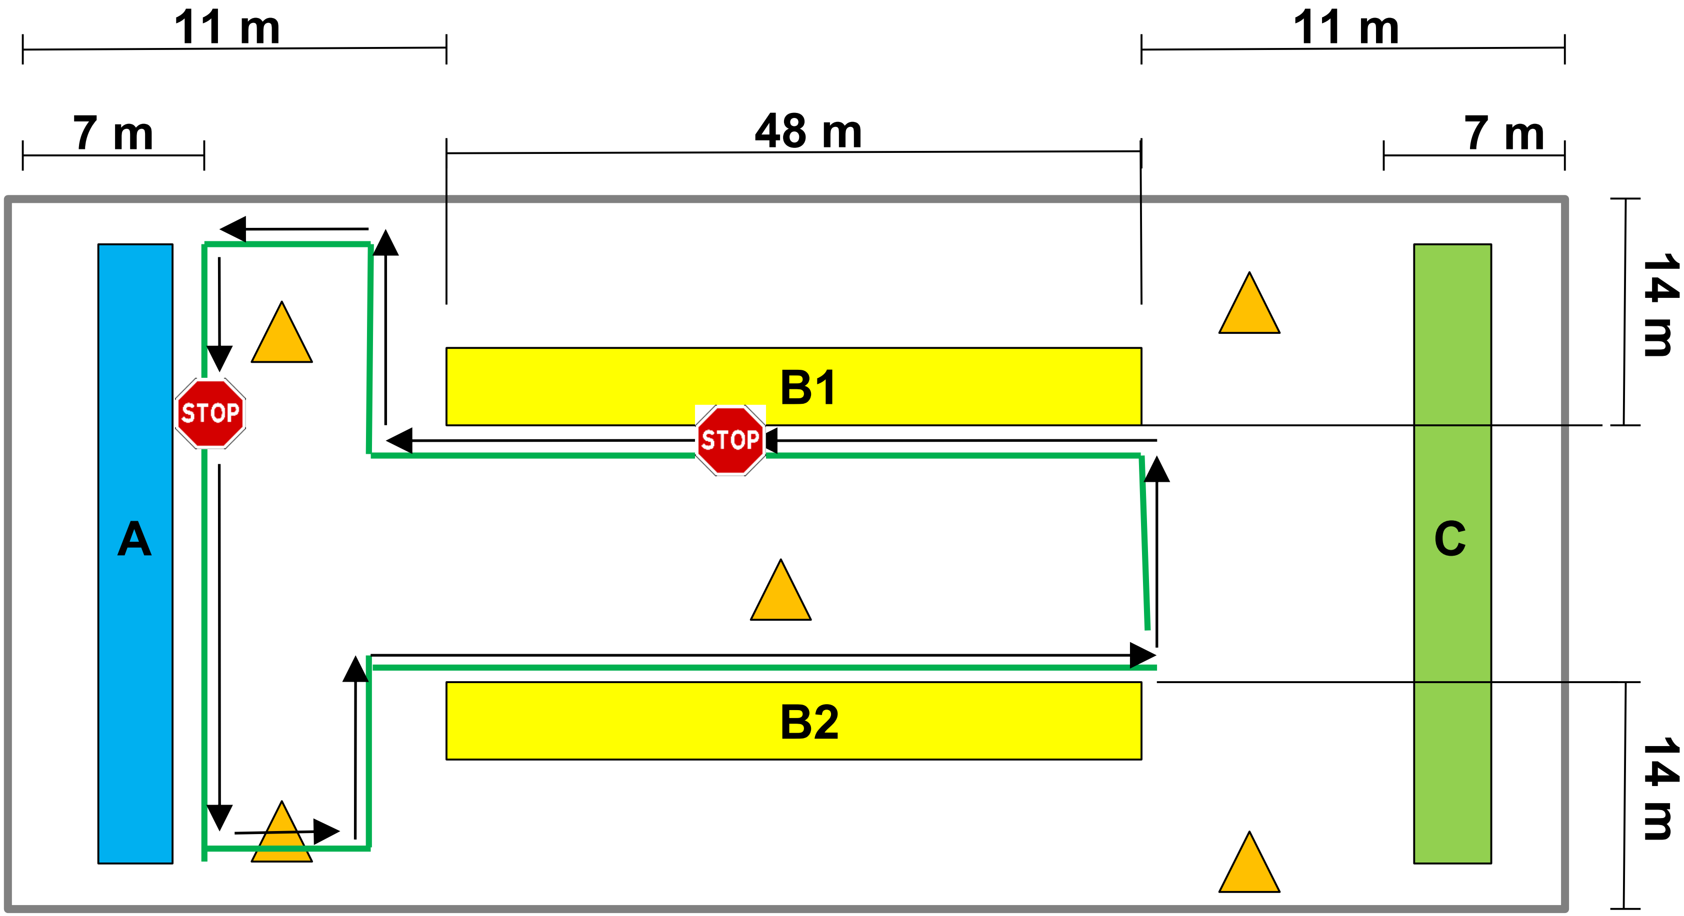

Rover Speed: 1 m/s

64 rovers moving between A and B1/B2

In essence, the area denoted by **A** has 20 potential stop positions while **B1** and **B2** have 10 each, uniformly distributed on their length. Each rover is assigned two stops (one from A and B1 or one from A and B2). The rover moves at constant speed from one stop to the other.

Rover Movements

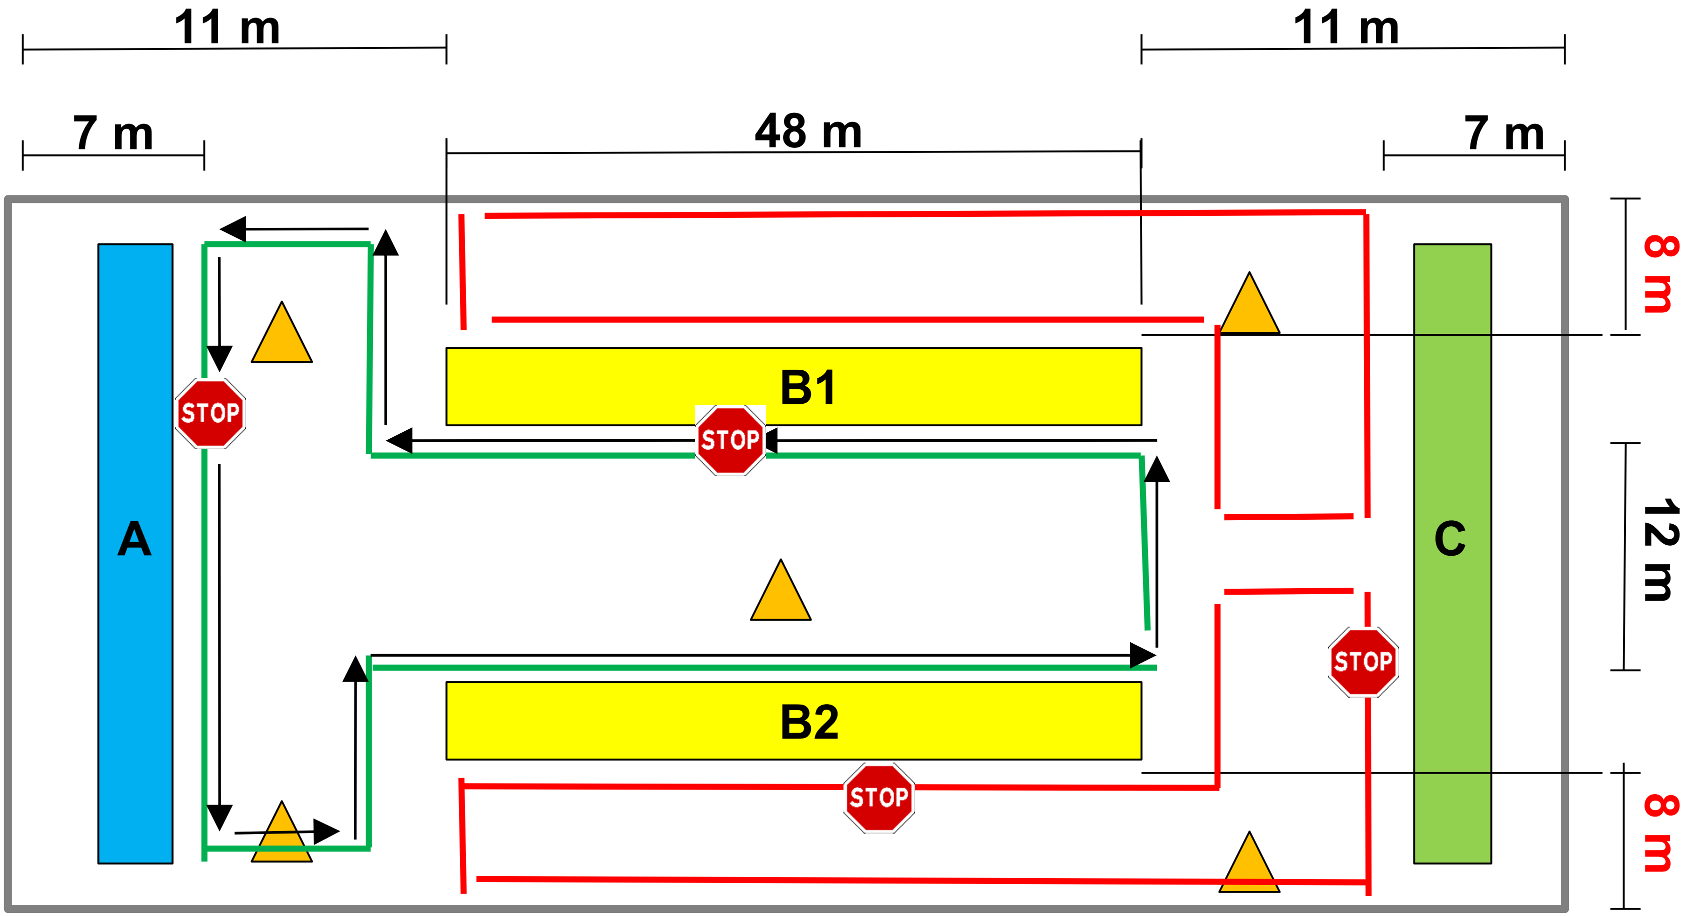

Rover Speed: 1 m/s

Other 13+13 rovers moving between B1/B2 and C

Areas **B1** and **B2** have 10 potential stop positions each while C has 20. Stops are uniformly distributed on each area. Each rover, upper and bottom area, is assigned two stops from B1 and C and B2 and C, respectively.

Scenario – Network Traffic

TransportControl - Target Network Load

Number of bytes sent through network: single com cycle | single Target | IPv4					
	UDP payload	UDP header	IPv4 header	WLAN overhead (~50%)	Total
Target > Server	45 bytes	8 bytes	20 bytes	37 bytes	110 bytes
Server > Target	25 bytes	8 bytes	20 bytes	27 bytes	80 bytes
Total	70 bytes	16 bytes	40 bytes	63 bytes	**189 bytes**

Number of bytes sent through network: single com cycle | single Target | IPv6					
	UDP payload	UDP header	IPv6 header	WLAN overhead (~50%)	Total
Target > Server	45 bytes	8 bytes	40 bytes	47 bytes	140 bytes
Server > Target	25 bytes	8 bytes	40 bytes	37 bytes	110 bytes
Total	70 bytes	16 bytes	80 bytes	83 bytes	**249 bytes**

Minimum data rate for 100 Targets				
	1 Target	100 Targets	500 ms lifecycle*	100 ms lifecycle*
IPv4	189 bytes	18900 bytes	302 kbit/s	1512 kbit/s
IPv6	249 bytes	24900 bytes	398 kbit/s	**1992 kbit/s**

2 cases

*server lifecyle interval affects the traffic control performance:

100 ms	dynamic traffic control, best performance
500 ms	dynamic traffic control, worst performance
> 500 ms	block signaling only, not recommended

Experimental Configuration

- Network Simulator 3 (NS3)
- IEEE 802.11n (classic Wi-Fi)
- Each AP uses a different channel
  - 802.11n (5 GHz) is less prone to inter-channel interference w.r.t 802.11g
- Messages sent to/by each rover every 50 ms, <u>100 ms</u> or 500 ms

Network Performance Metrics

- Packet loss
  - Measured at the application layer (end-to-end)
  - The *acceptable* packet loss rate depends on the criticality of the exchanged message content
- Message delay
  - Must be strictly lower than the operational cycle, otherwise the exchanged operational data might not be valid anymore
  - Measured delay corresponds to the time added by the last wireless tier

Mobile Scenario

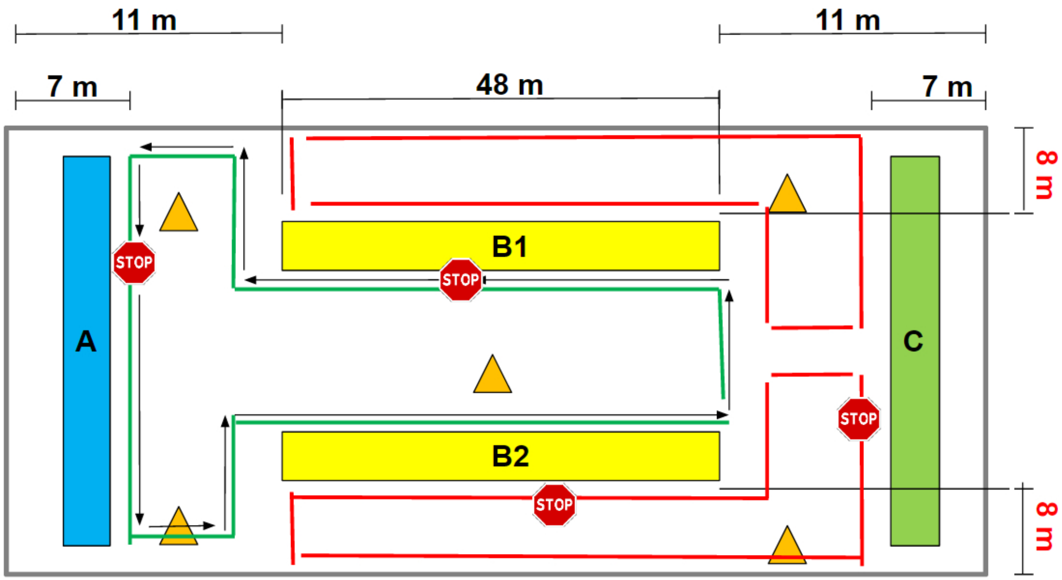

Mobile Scenario

- Scenario modeled according to the specifications
  - Rovers move at 1 m/s and occasionally stop
- Studied different deployment strategies
  - Uniform - AP0/AP1/AP2 (30/30/30)
  - AP0/AP1/AP2 (13/13/64)
  - AP0/AP1/AP2 (19/19/40)
- Network performance metrics
  - Packet loss at the app. layer
  - End-to-end delay at the app. layer

Packet loss – 100 ms duty cycle

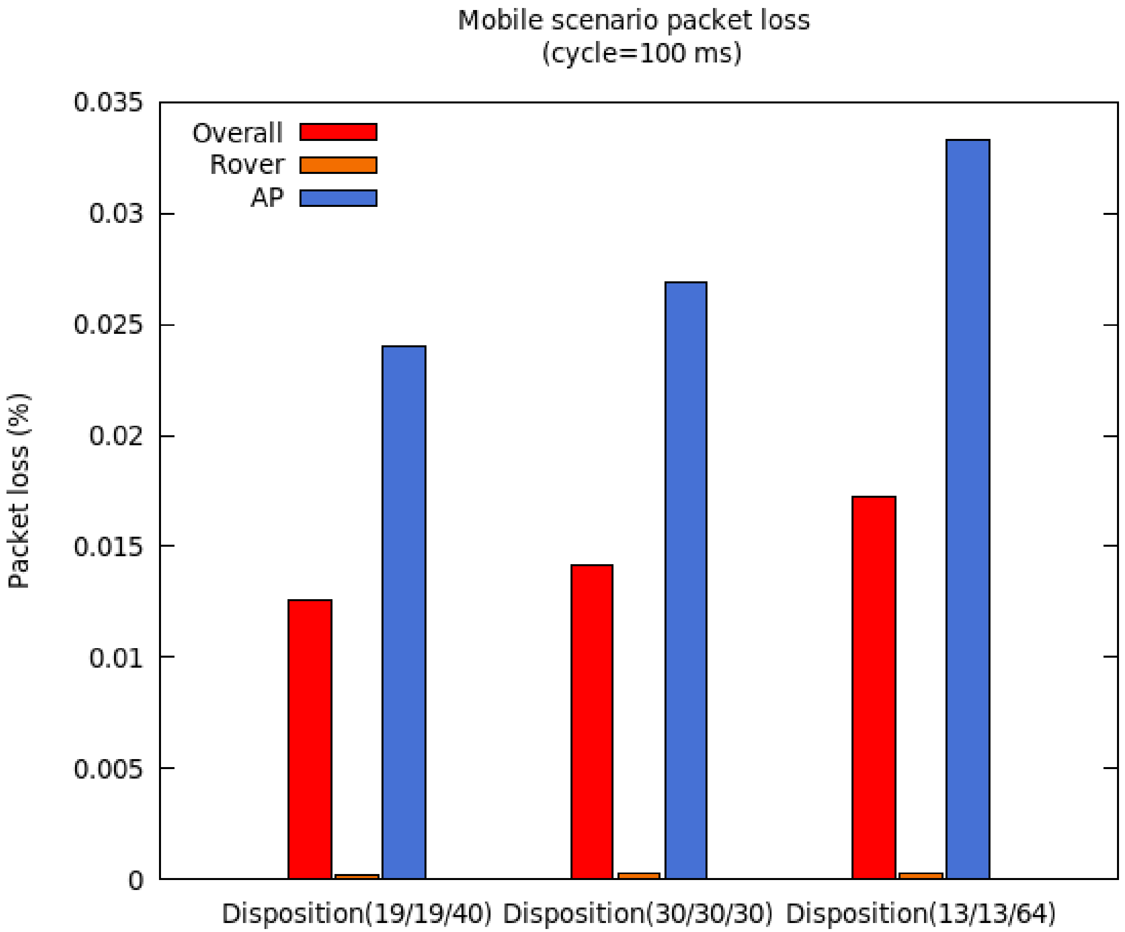

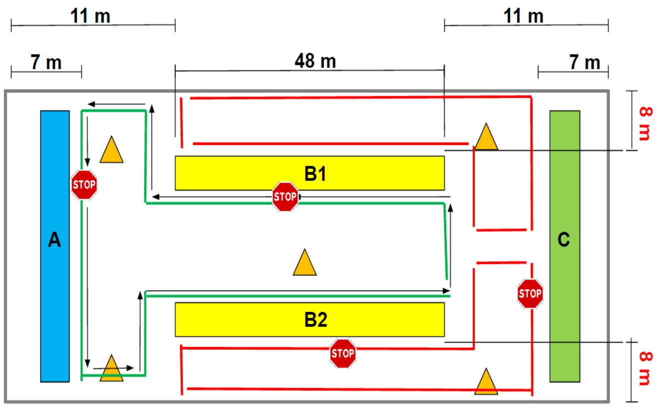

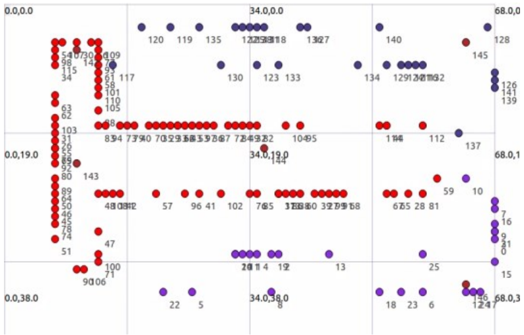

Basically, in all configurations there are very few packet losses (less than 0.035%)

E2E delay AP per Deployment – 100 ms duty cycle

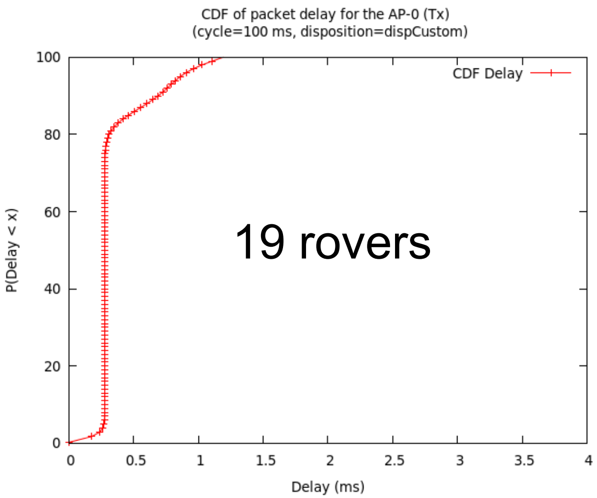

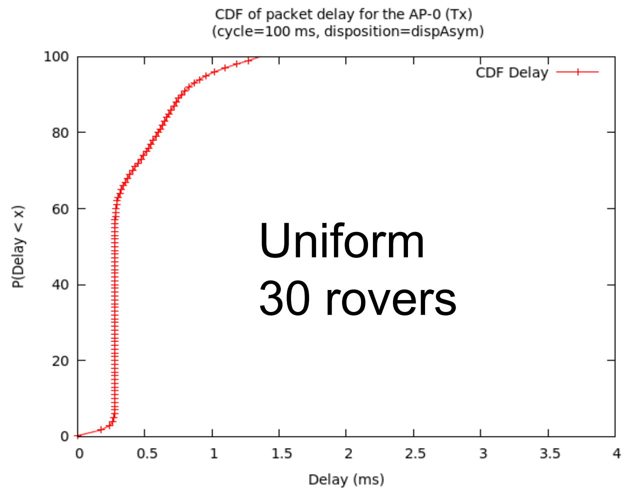

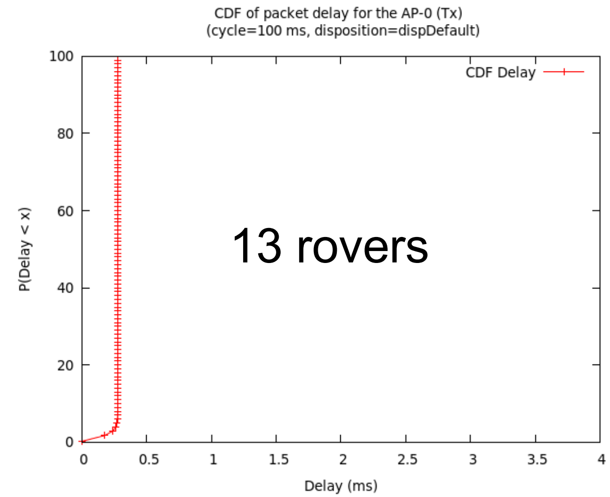

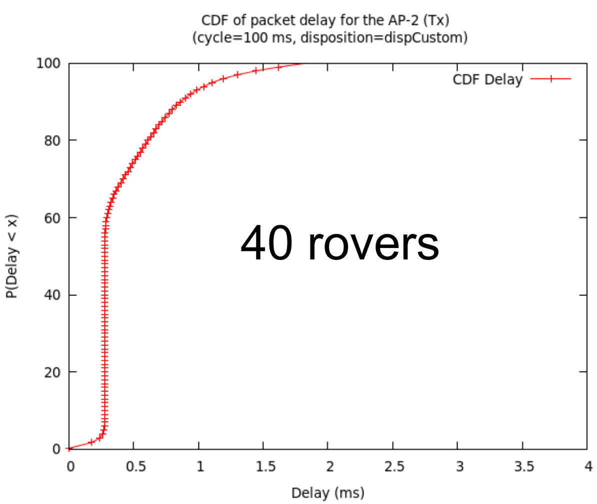

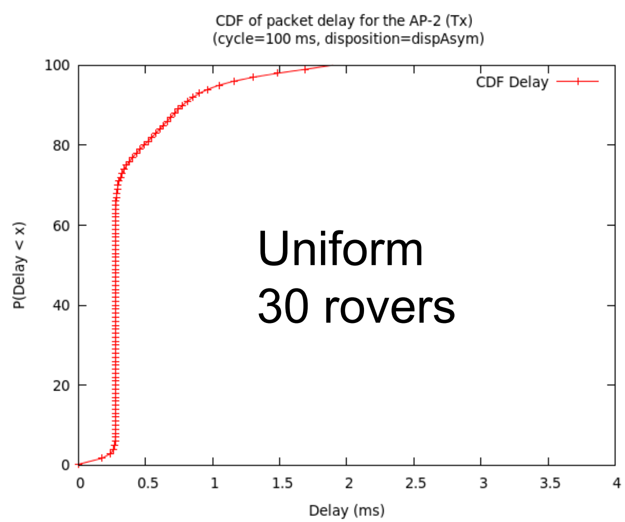

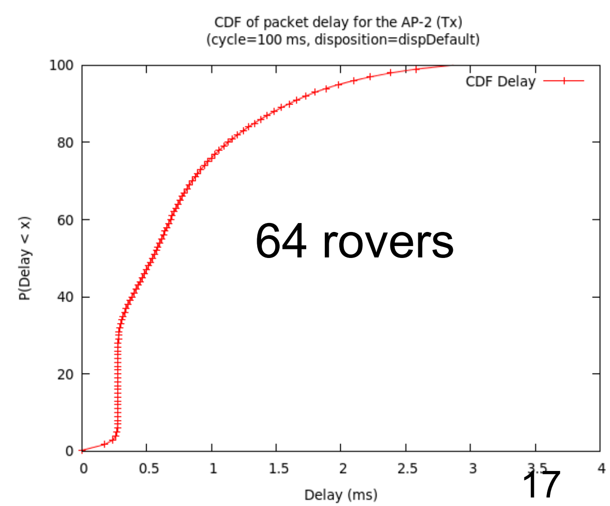

E2E delay Rover per Deployment – 100 ms duty cycle

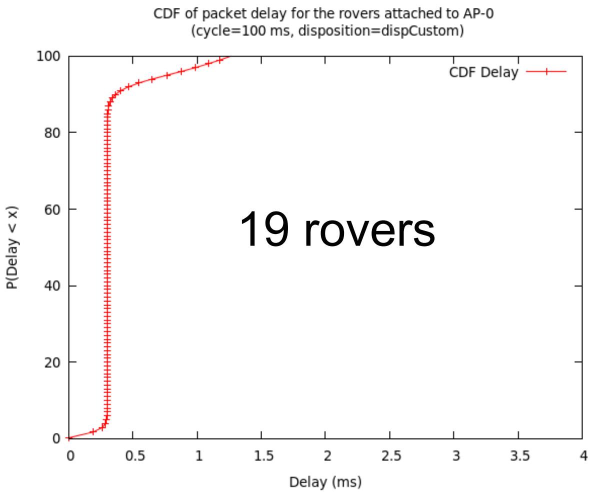

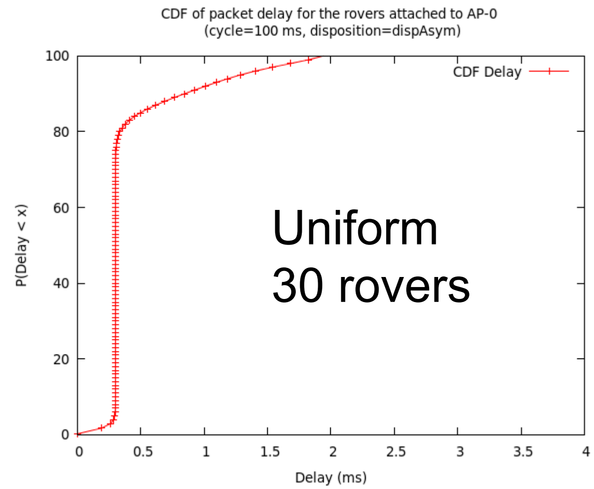

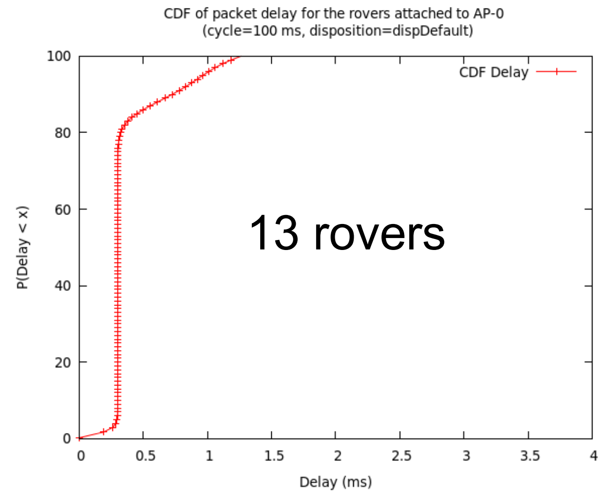

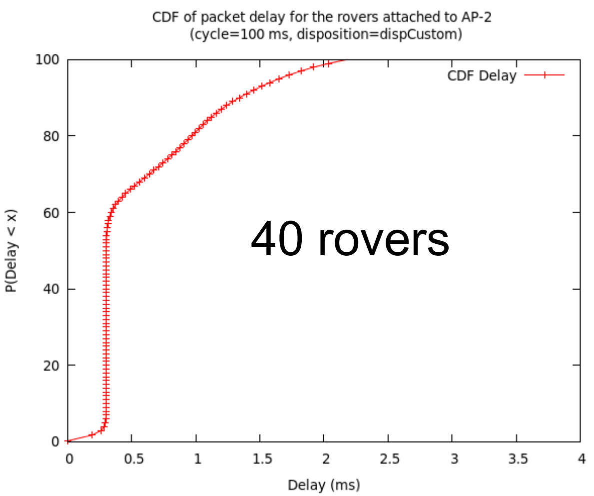

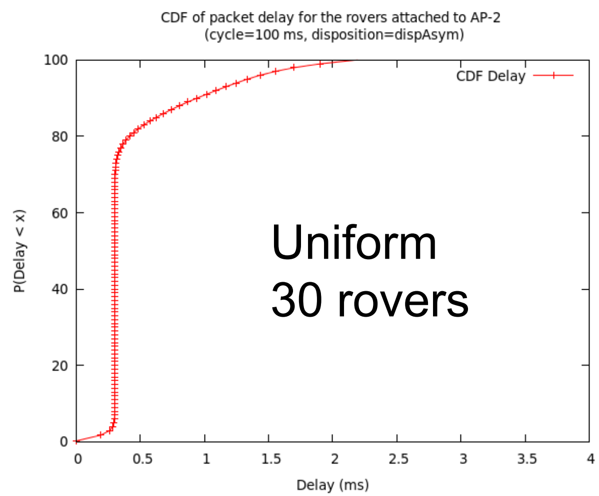

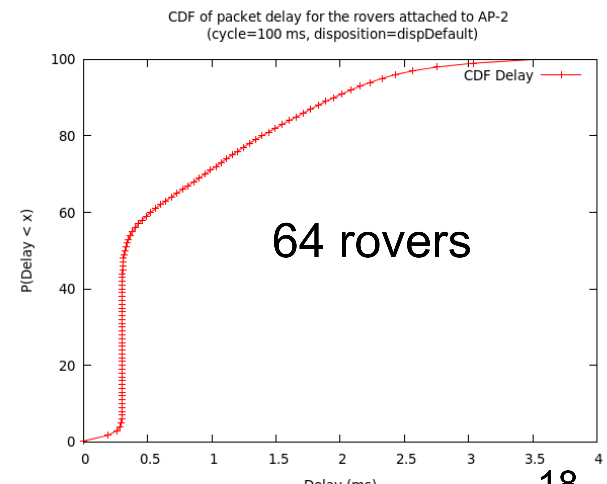

Summary of Results

Considering the per-packet delay at the Access Point (packets traveling toward the rovers), we had less than 4 ms of added delays between the AP and any rover. Also, there is no queuing up; so the system can sustain the network traffic.

Considering the per-packet delay at the rovers (packets traveling toward the AP), in all configurations we had less than 4 ms of added delays between any rover and the AP. Also, there is no queuing up; so the system can sustain the network traffic

In summary there seems to be no problem in handling the considered scenario with a packet interdeparting time of 100 ms.

We have considered even the case with 500 ms of interdeparting time. Results are obviously even better; there is no need to report them here.

Static Scenario

To test the limit of the system we have also considered the case with a single AP and many nodes (rovers) connected to it.

We considered various configurations with different number of nodes (not moving).

Bandwidth consumption is much higher than throughput due to channel contention mechanisms.

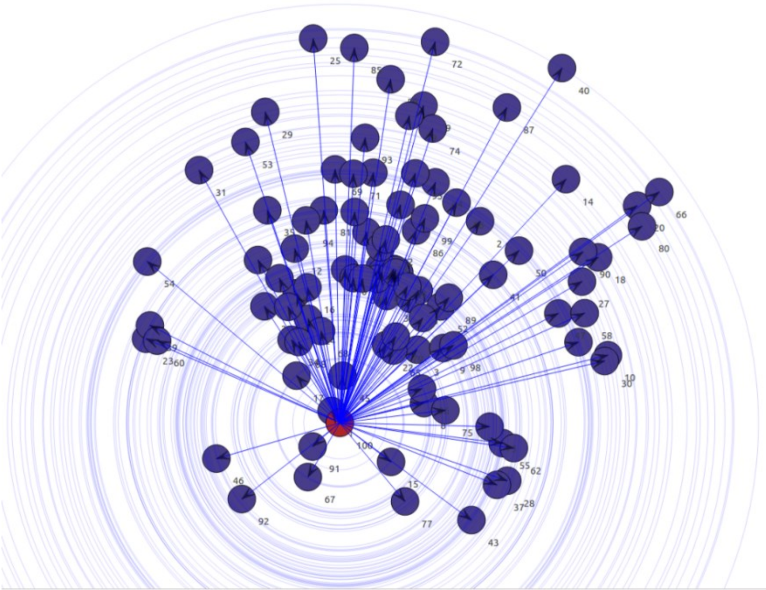

Packet loss – 100 ms duty cycle

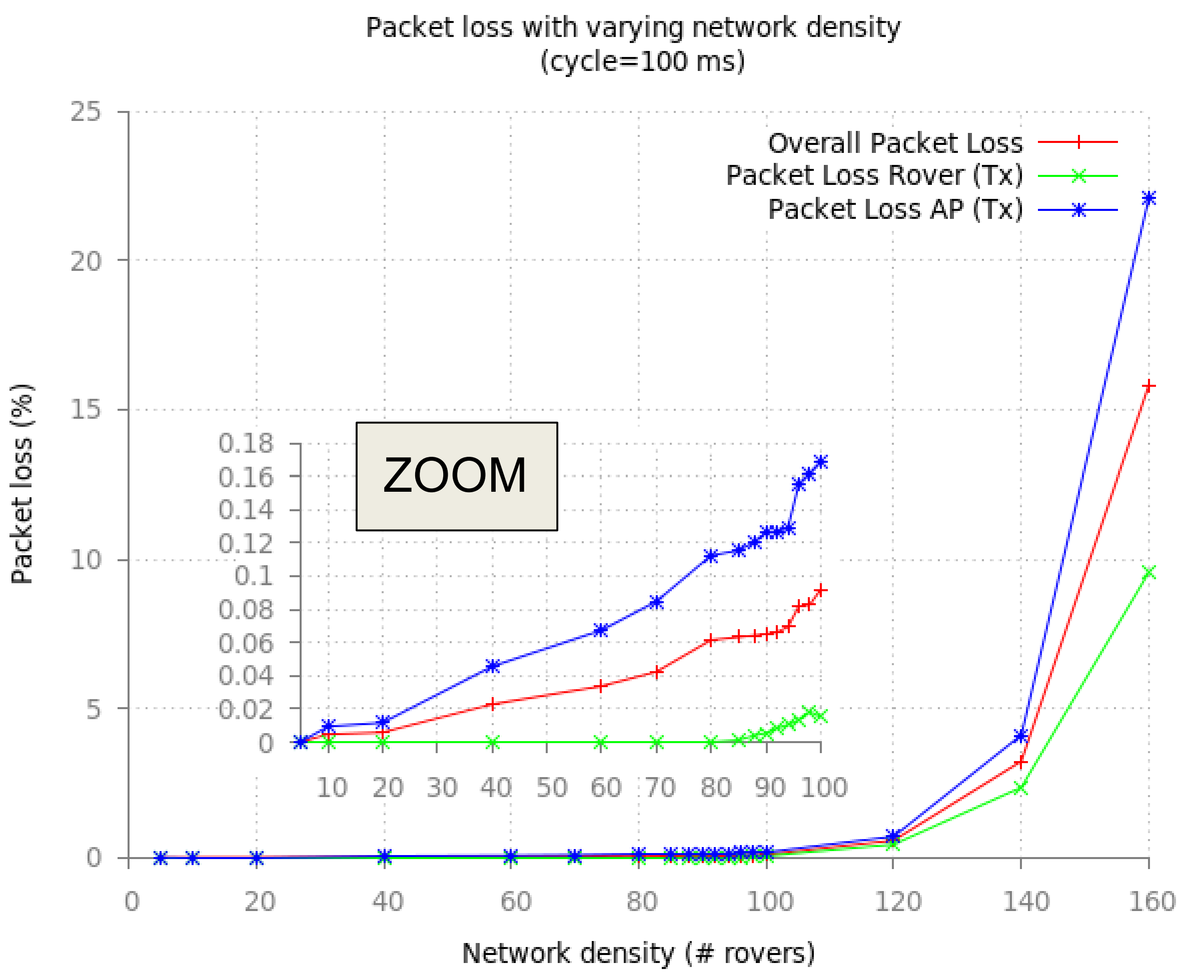

Beyond 120 nodes we see an increase of the packet loss

E2E delay AP - 100 ms

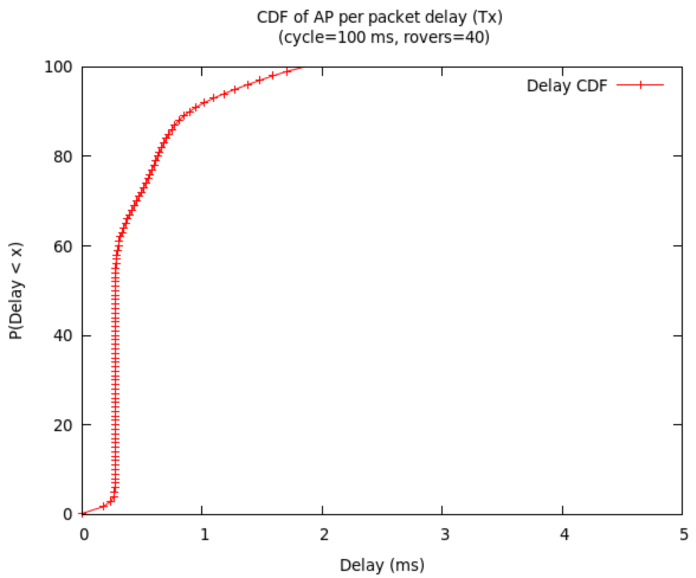

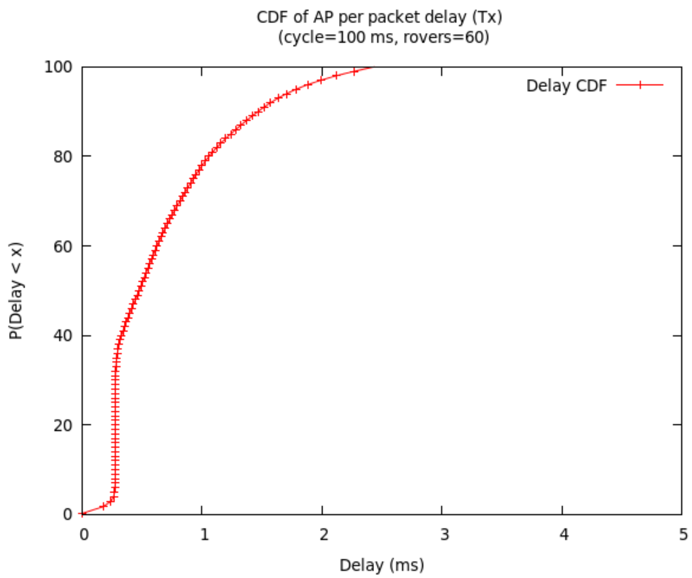

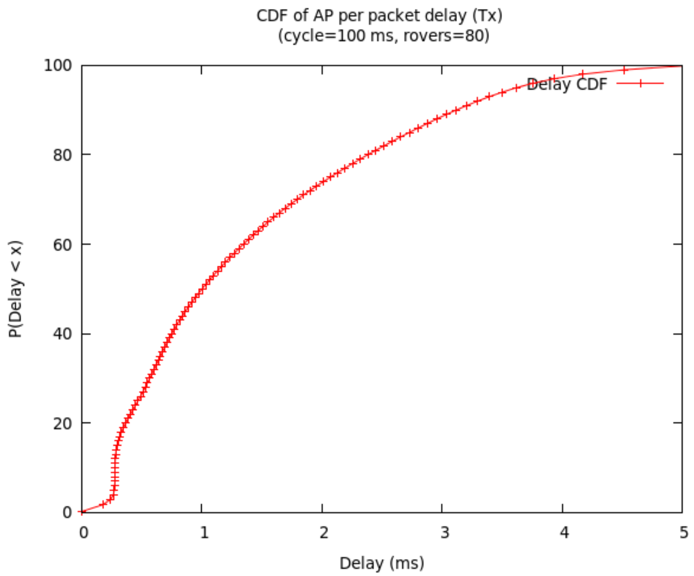

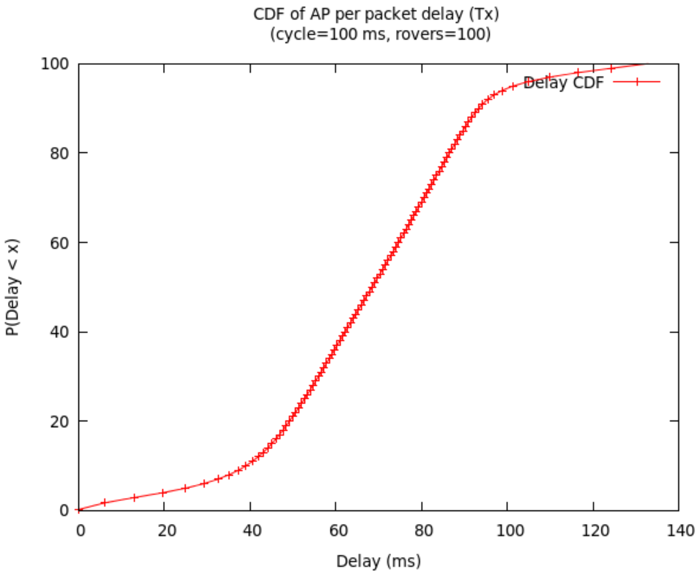

E2E delay Rover –100 ms

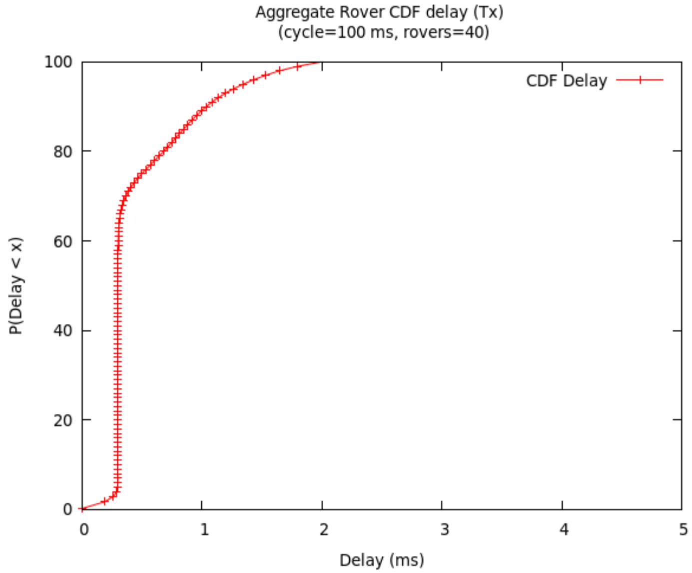

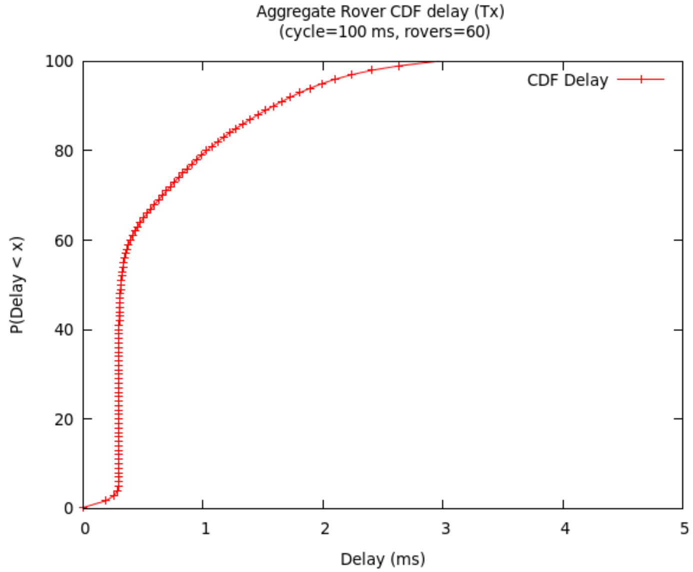

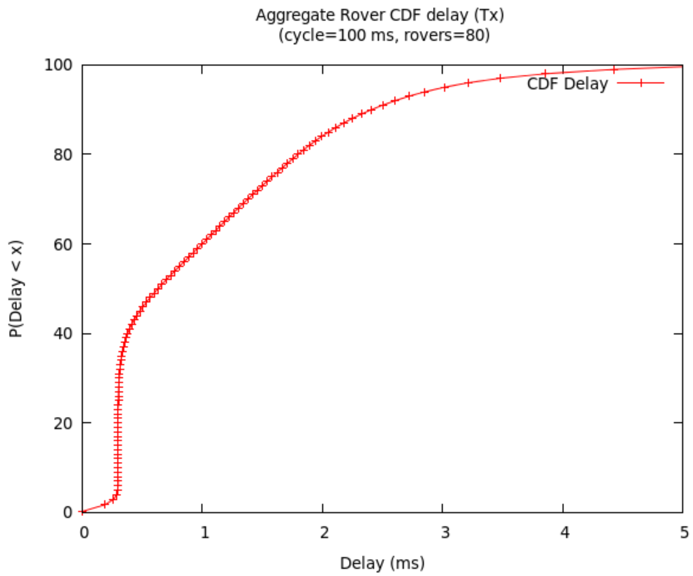

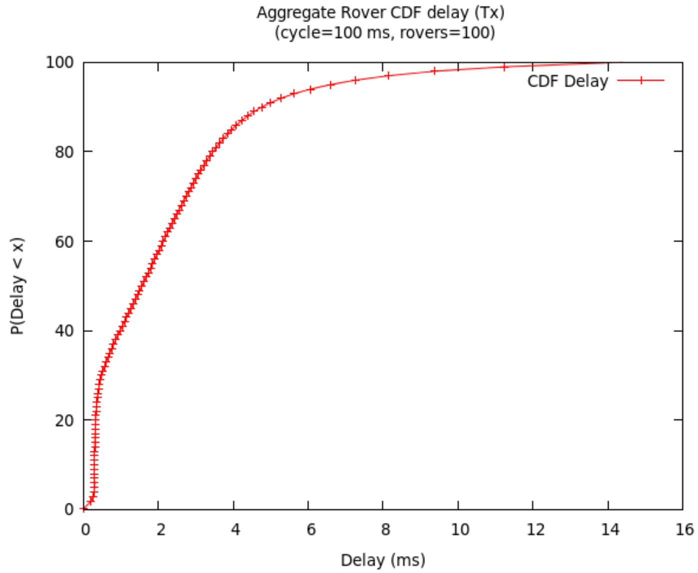

How far can we push ? 50 ms duty cycle

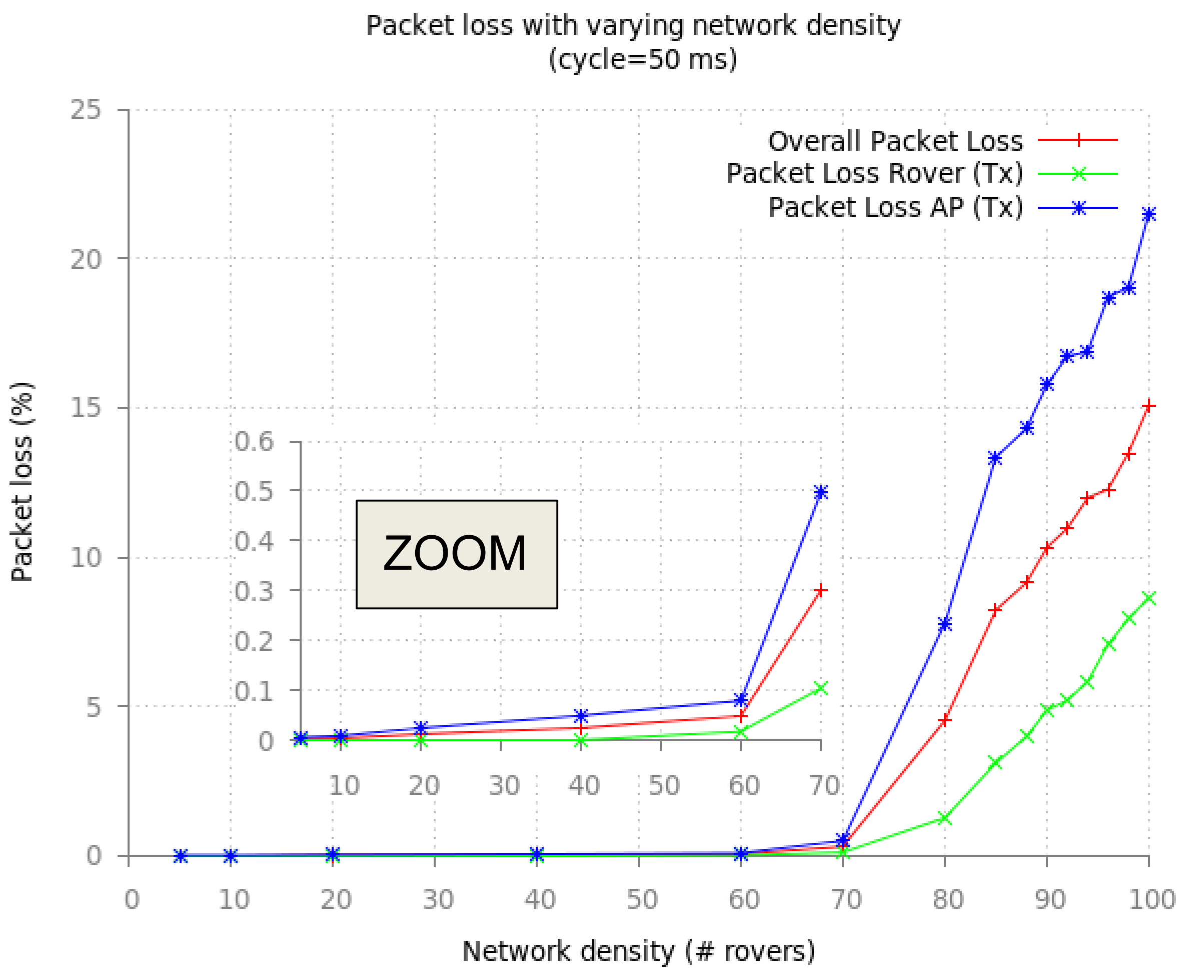

With 50ms of interdeparting time we have a packet loss increase after 70 rovers

E2E delay AP – 50 ms

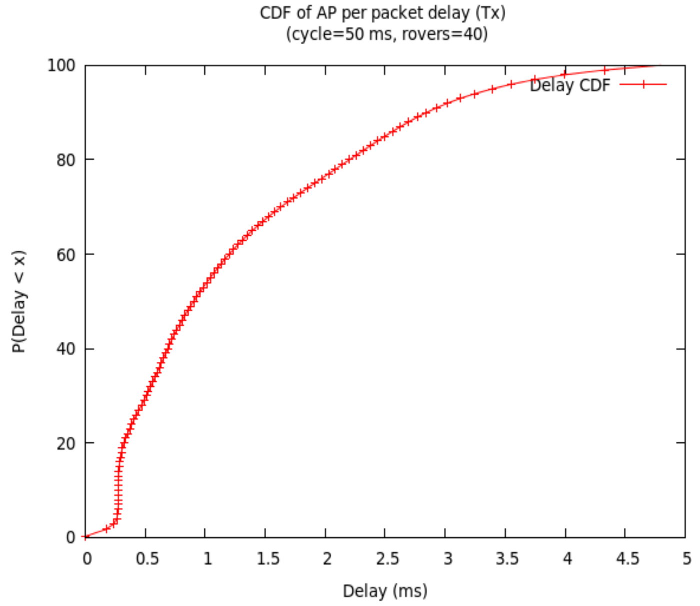

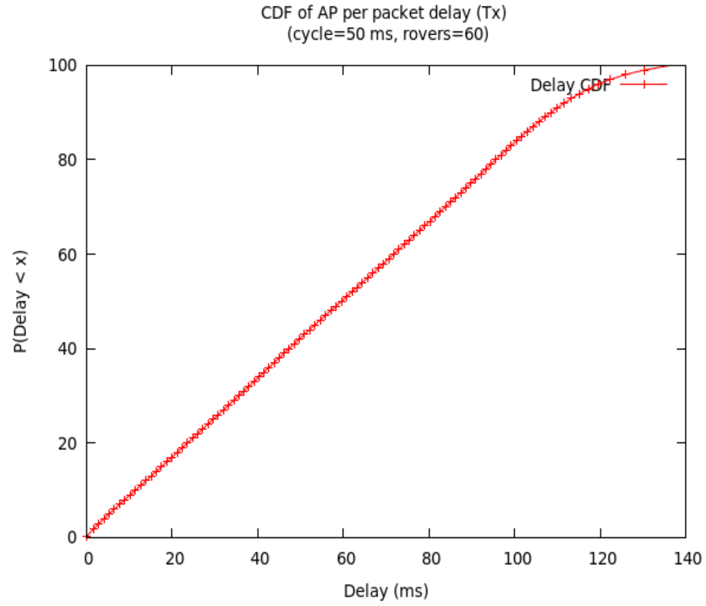

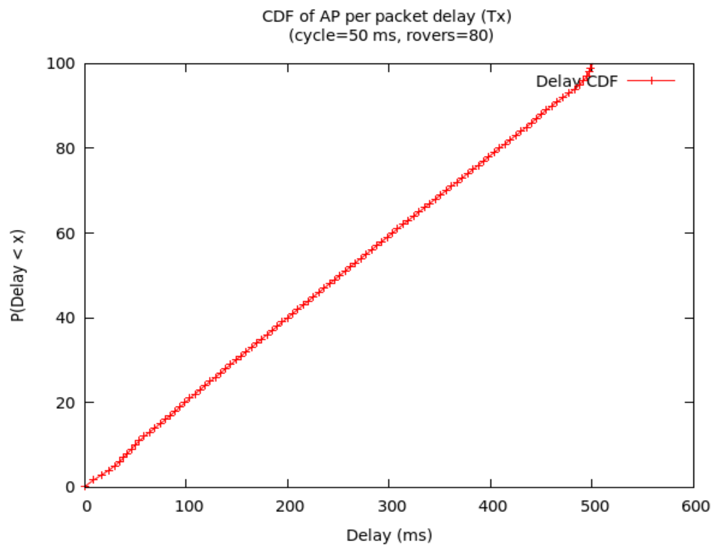

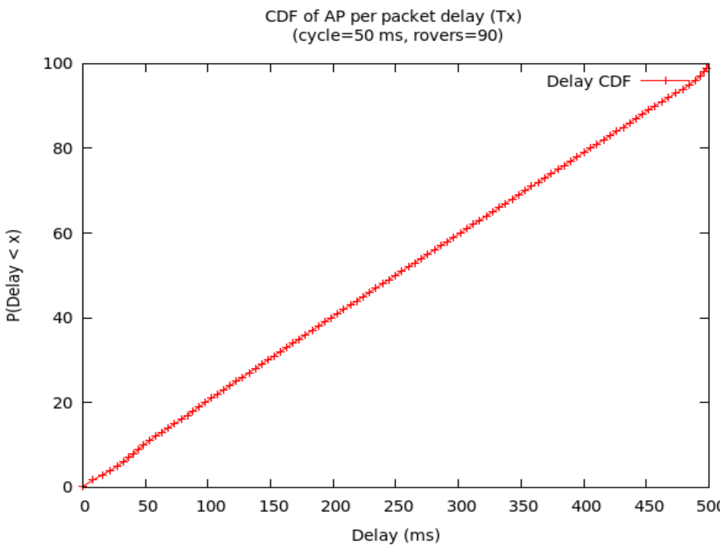

E2E delay Rover – 50 ms

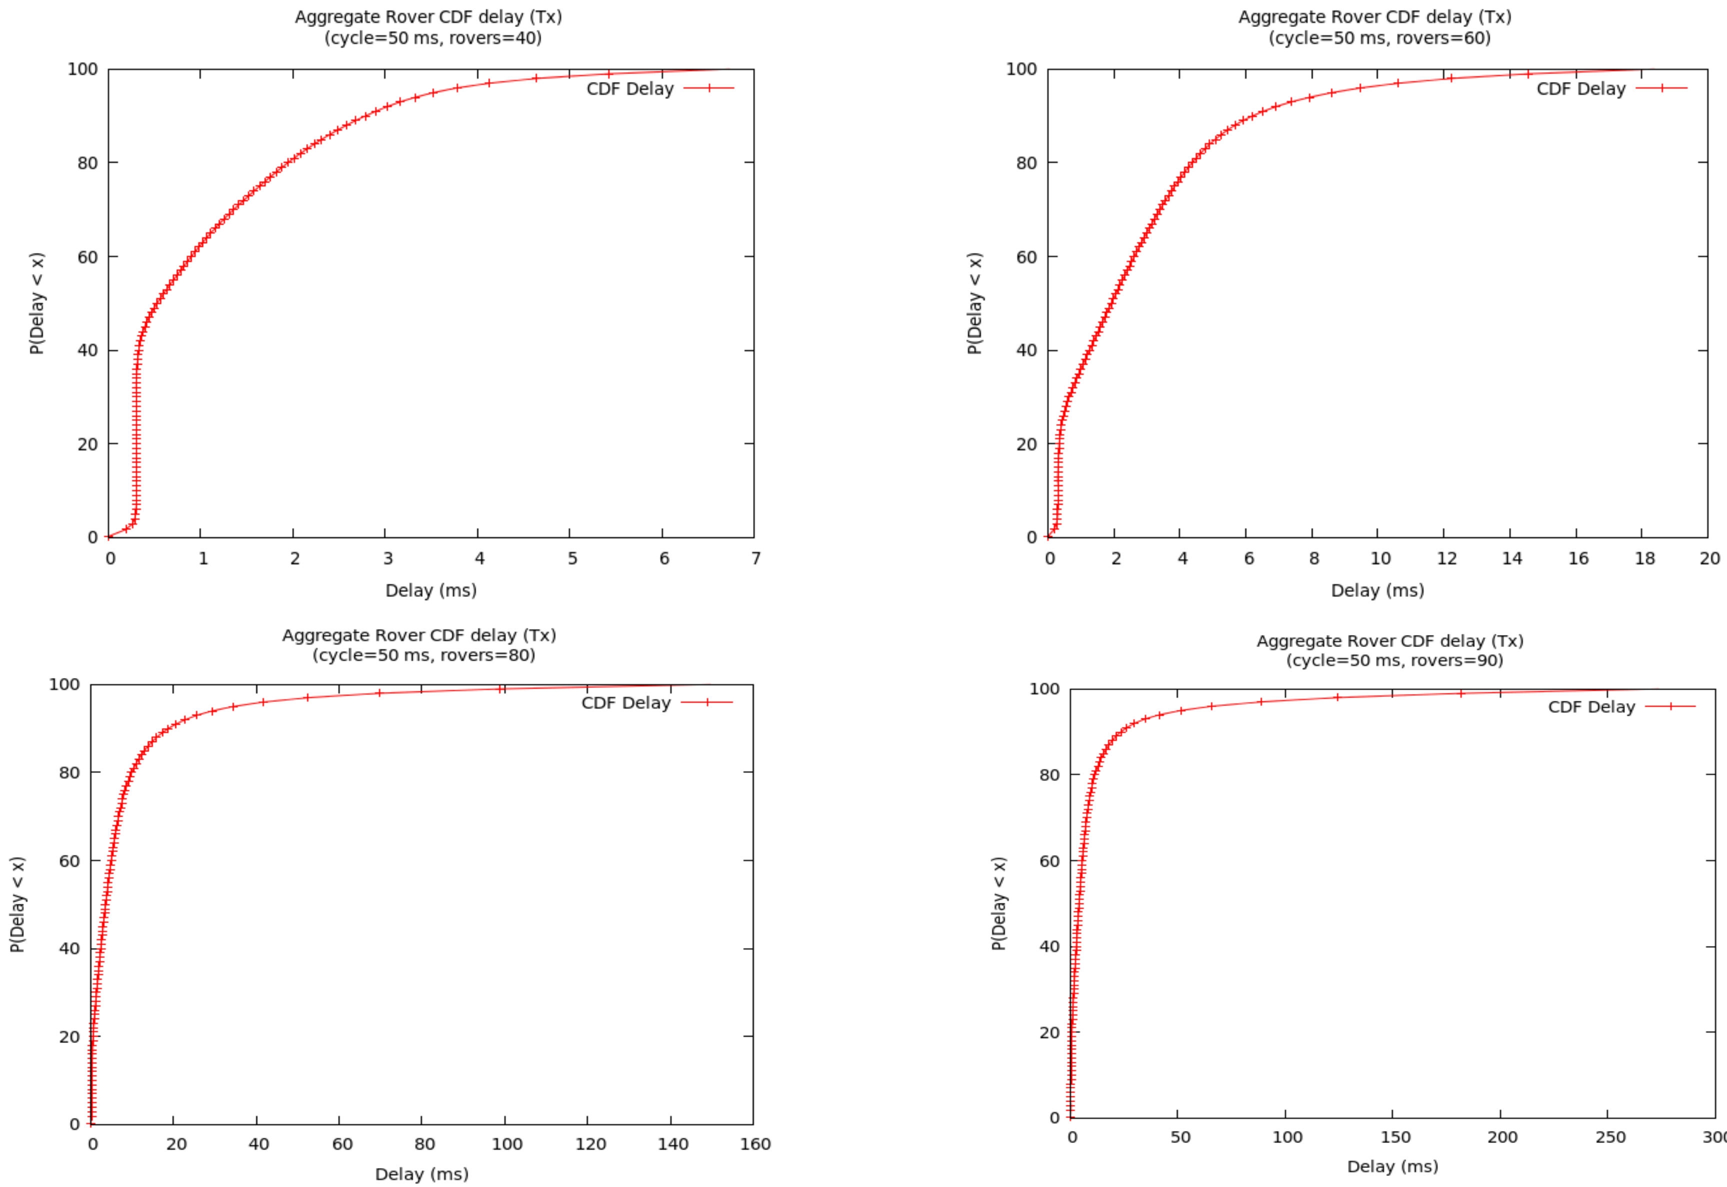

Adaptive Rate Manager

Recent APs use improved techniques such as, for instance, Adaptive Rate Managers. With them (present in the actually purchased APs) results improve.

In the following, we have considered a modified configuration exploiting an adaptive rate manager (Minstrel).

Adaptive Rate Manager in Short

- A table of acknowledgement probability estimates is maintained per neighbour per (physical layer) rate
- The ratio of transmission attempts to acknowledgements received is maintained using an exponential weighted moving average to smooth the probability estimation
- On a frequent basis, the table is scanned to find an approximation to the best performing rate and retry chain, and that is used for transmission for the next interval
- With a moderate frequency, frames are selected to probe presently unused rates
    - feedback from those probe frames maintains the probability estimates for unused rates so that can be chosen if needed

Minstrel Rate Manager – 100 ms duty cycle

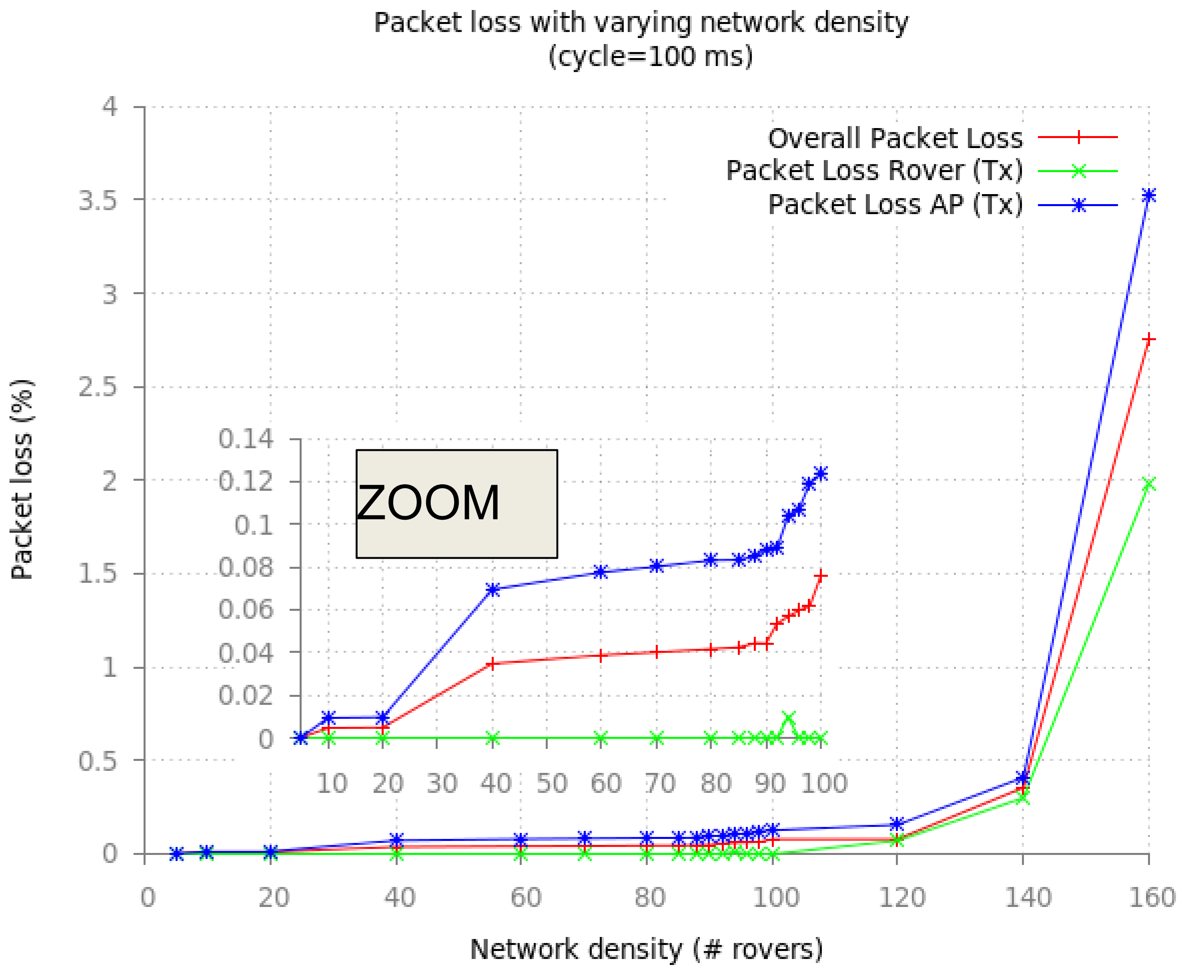

Generally, with an adaptive rate manager performance improves. As evidenced in the chart, the system can support even 140 rovers with less than 1% of packet loss.

E2E delay AP – 100 ms

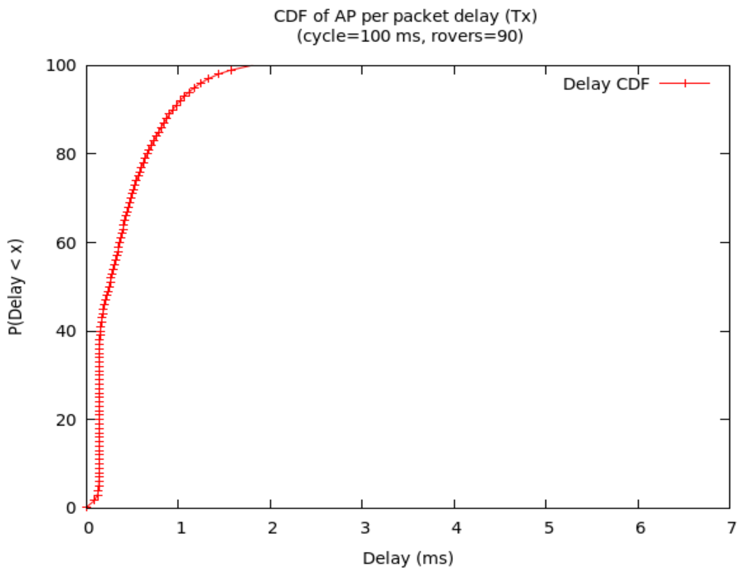

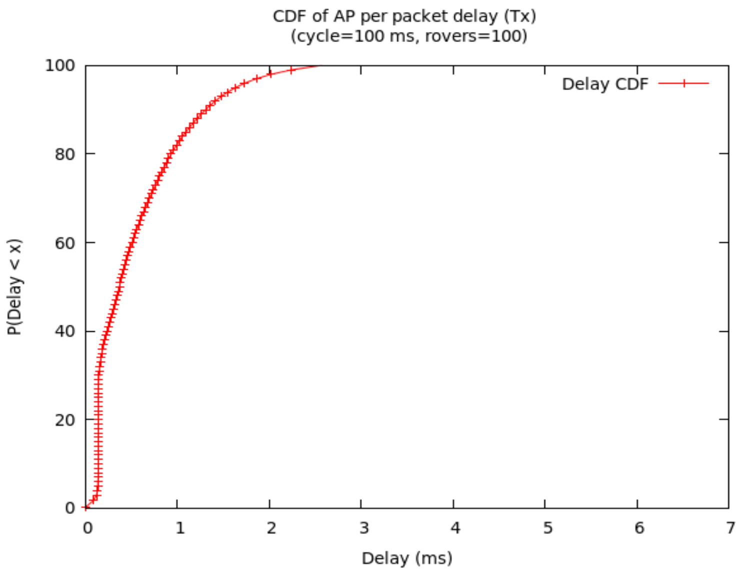

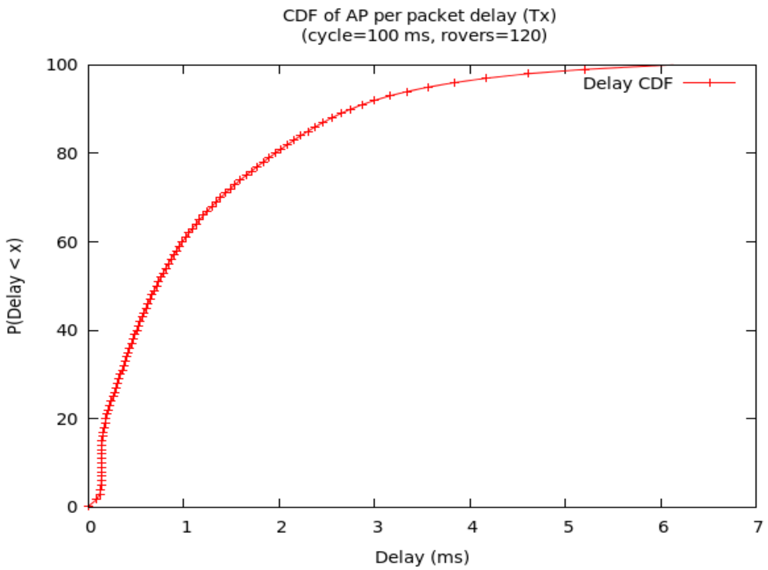

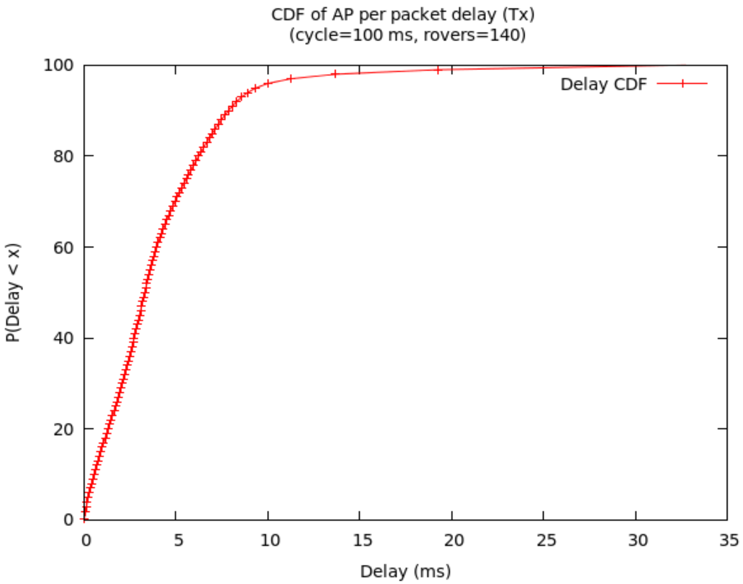

E2E delay Rover – 100 ms

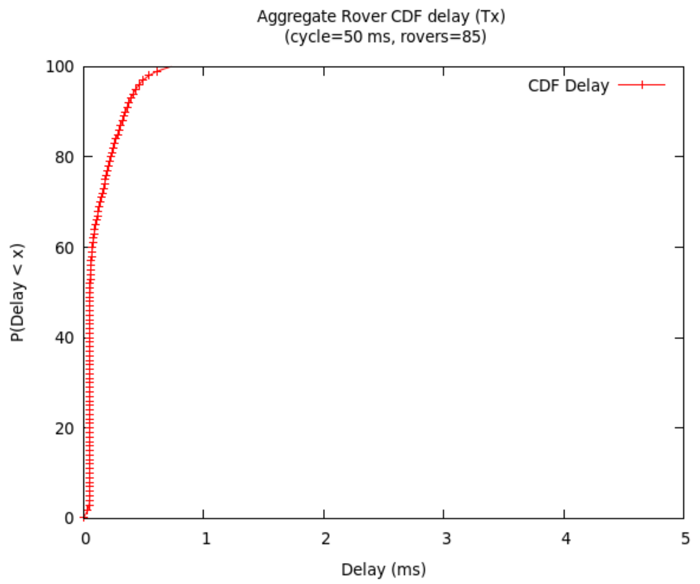

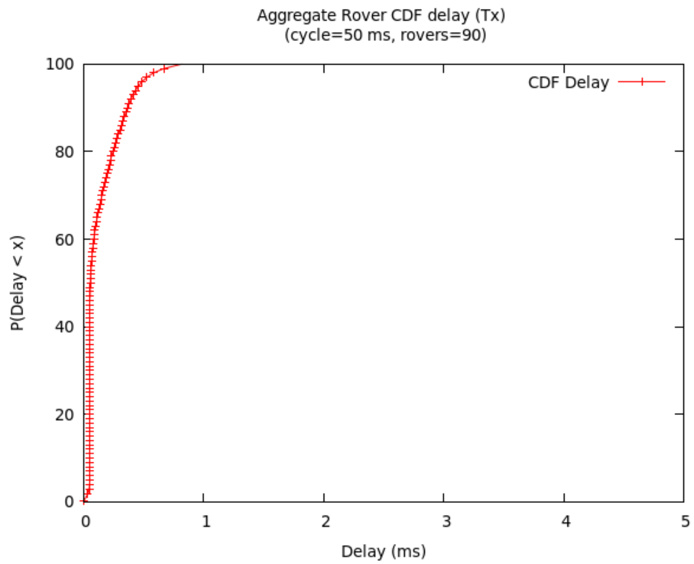

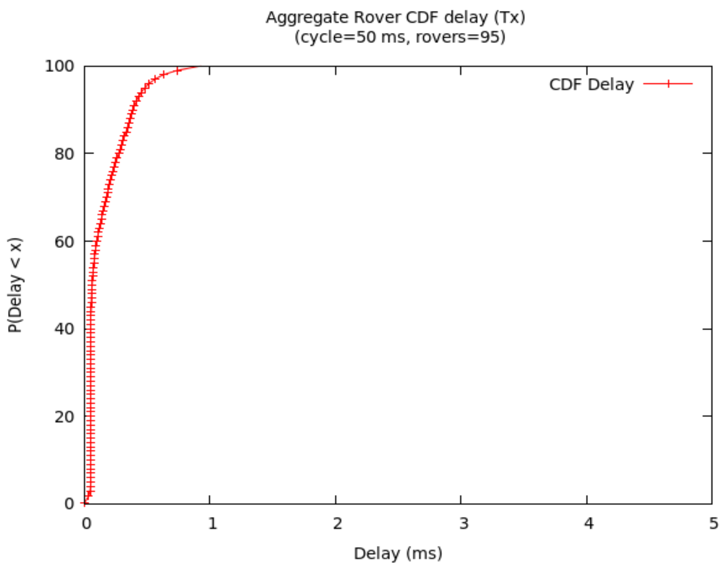

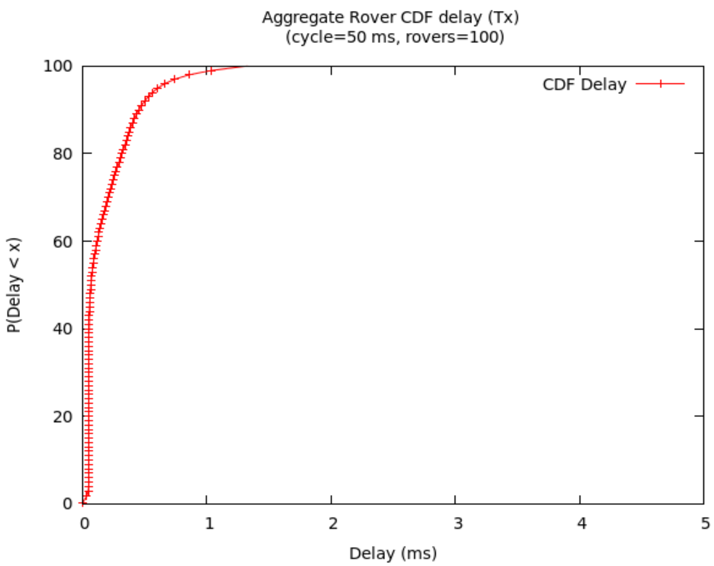

Conclusion - Operational upper bounds

100 ms duty cycle	Constant Rate	Minstrel
Packet Loss	120 rovers (< 1%)	140 rovers (< 1%)
Delay	80 rovers (< 5 ms)	120 rovers (< 6 ms)
Recommended Max	80 rovers	120 rovers

50 ms duty cycle	Constant Rate	Minstrel
Packet Loss	70 rovers (< 1%)	100 rovers (< 1%)
Delay	40 rovers (< 5 ms)	65 rovers (< 5 ms)
Recommended Max	40 rovers	65 rovers

For greater number of AGVs, the AGVs must relay on less communication with a central station, thus giving up on some flexibility or increasing the level of movement decisions that can be taken autonomously by the AGVs