A Status Report
Modeling the ATM Testbed

- Very early measurements
- SIMDAQ model
- Comparison of SIMDAQ with measurements
- Next Steps

Early ATM Measurements

- 16 Port FORE switch
- 1-8 300MHz Pentium II as Global Processors
- 1-2 333 MHz PPros as supervisor RoI Processors
- 100 microsec "algorithm" time

Simple setup for ATM measurements

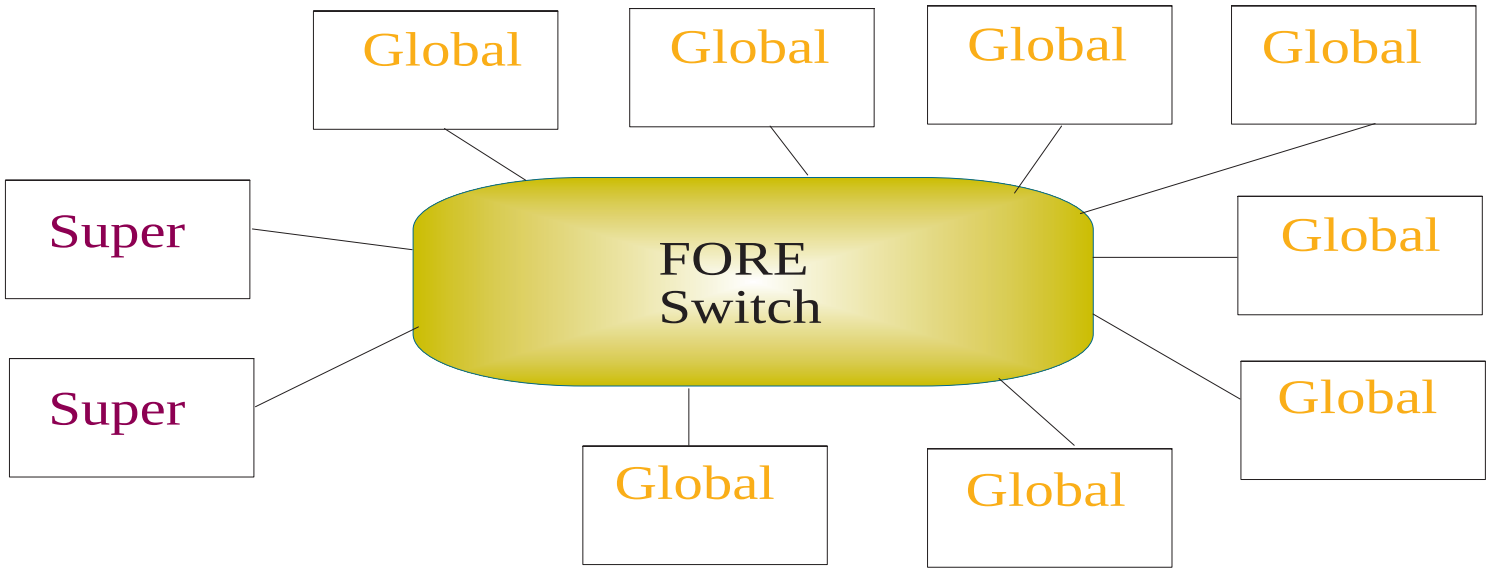

o Each Super queues up to 8 events in each Global (aka DST)
o Each Global waits 10 microsec and responds with event accept/reject
o Measure rates/latency dist. vs. #Supers/#Globals

SIMDAQ version of measurements

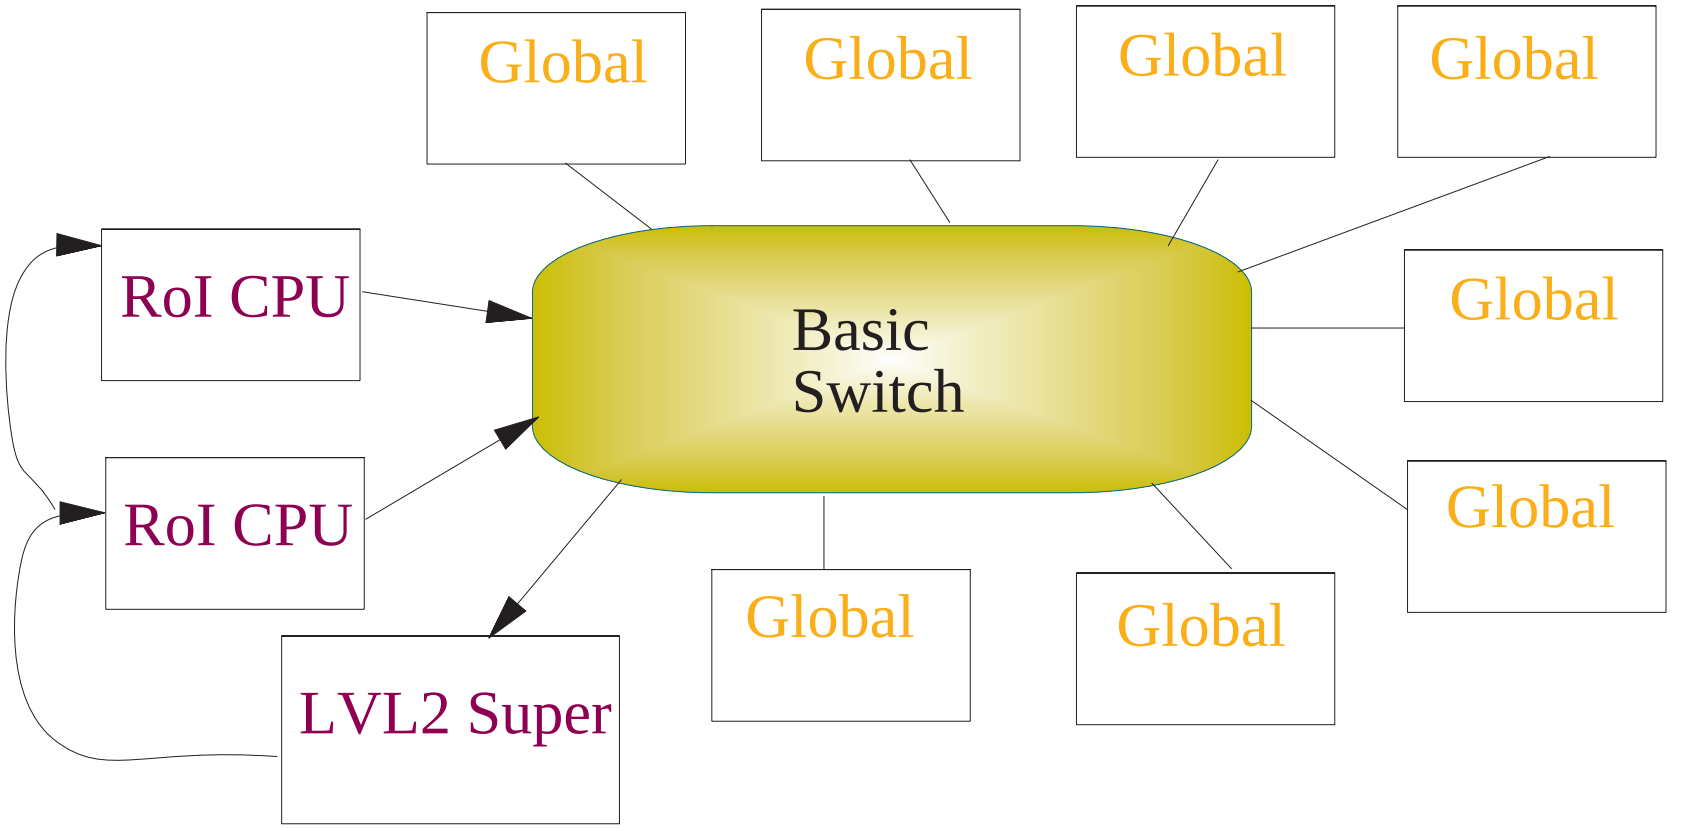

Differences

- SIMDAQ has data requests to ROB measured doesn't (data back =0 though)
- SIMDAQ has one Supervisor receiver even with two CPU Supervisor
- ATM is not modelled in SIMDAQ (yet)
- Decision Block ???
- Depth of queue for each processor set globally in SIMDAQ and independently in measurement (per Super CPU)

SIMDAQ configuration (for study offline)

After the following tune:
set scheduling times to
5 microsec (vs. 25)
take run with one Supervisor and 8 Globals and tune supervisor timings to make the rate match (tuning dependence?)

Warning!

- Take all of this with a 50 lb salt block (instead of a "grain of salt")
- Too many mysteries to count and this all suggests that none of this is right yet
- The results speak for themselves in this regard
- Need to really fix the "details" first and then understand the quirks

Rates (meas. vs. SIMDAQ)

- 8 global 1 super (tuned to) 44kHz vs. 46kHz
- 8 global 2 super 86 kHz vs. 56 kHz
- 1 global 1 super 11 kHz vs. 7 kHz
- 1 global 2 super 12 kHz vs. 7 kHz
- 8 global 1 super 39 kHz vs. 46 kHz (100 microsec. global algorithm)
- 8 global 2 super 42 kHz vs. 46 kHz (100 microsec. ...)

Latency

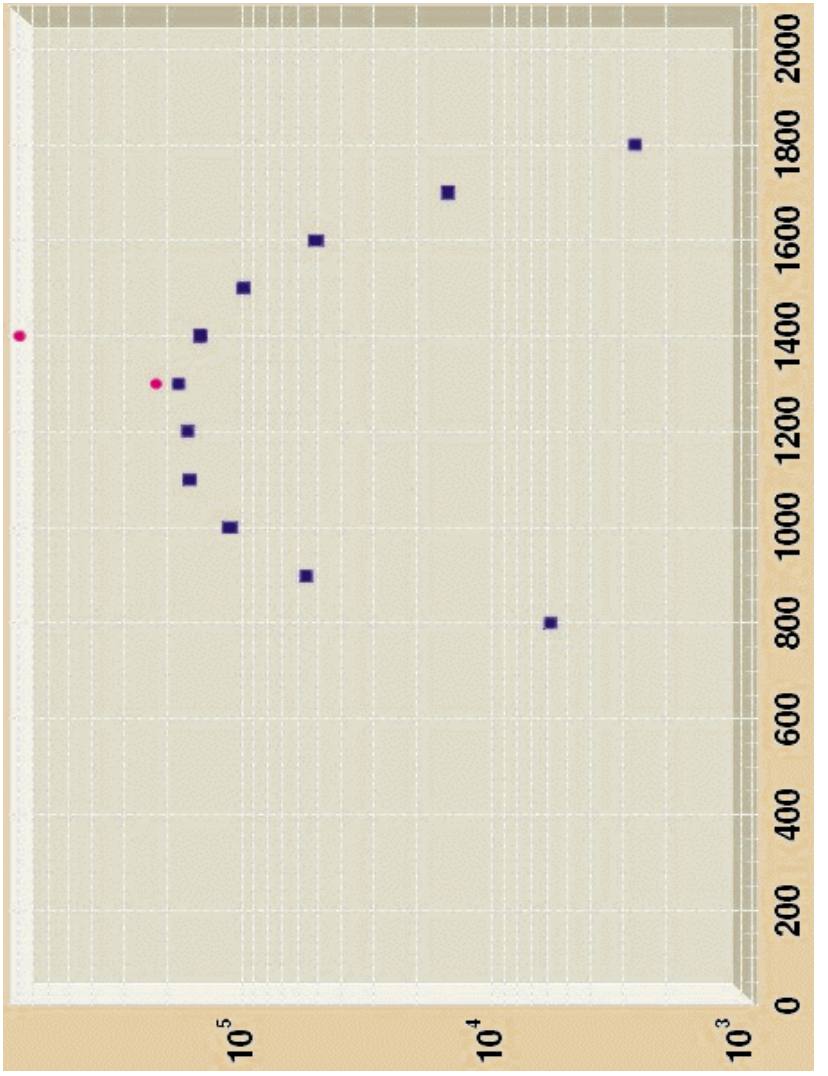
0
200
400
600
800
1000
1200
1400
1600
1800
2000
10^5
10^4
10^3

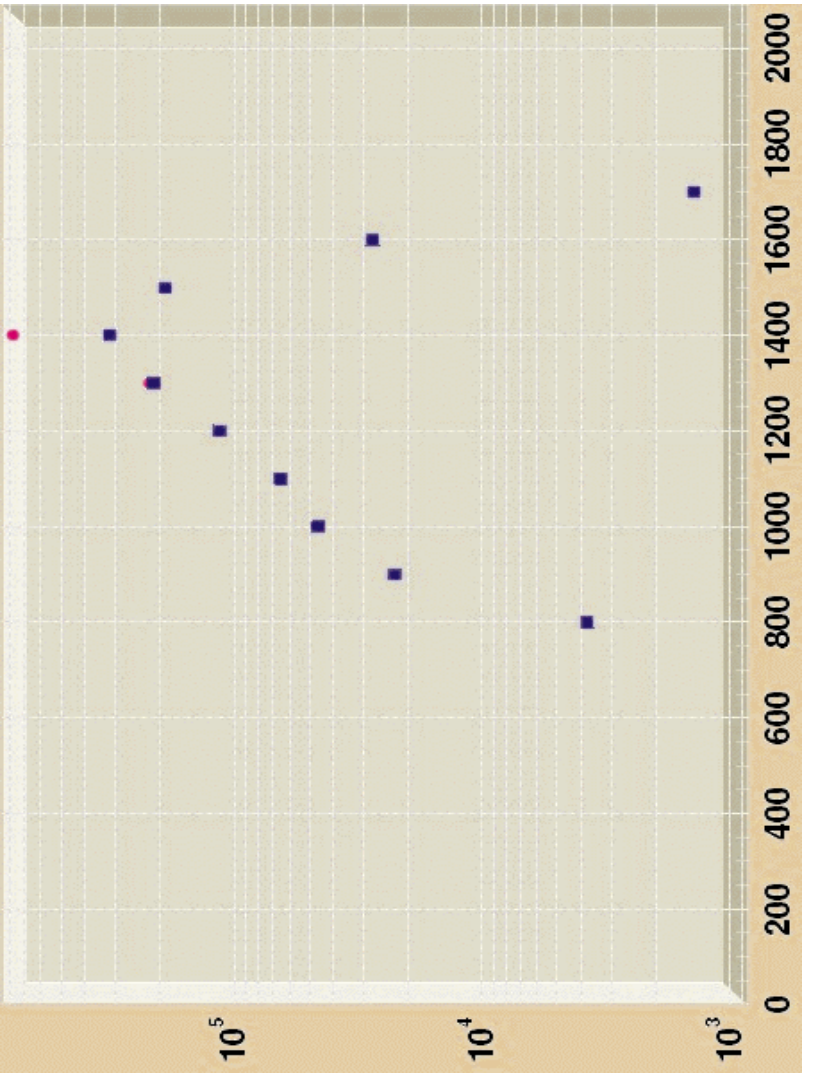
0
200
400
600
800
1000
1200
1400
1600
1800
2000
10^5
10^4
10^3