Model Based Metrics

IPPM Working Group
IETF 89, London, Mar 3

Matt Mathis <mattmathis@google.com>
Al Morton <acmorton@att.com>

Outline

- Why Model Based Metrics matter (2 slides)
- Document status update & major changes
- Future work & open issues

Why Model Based Metrics matter

- IP metrics to assure TCP performance
- Original intent was for SLAs
  - ISPs are selling IP service
  - User wants to buy end-to-end application performance
  - Most of the system is out-of-scope for the ISP: rest of path, end hosts
  - Want to know if the ISP's portion is
- Use models to derive IP requirements from application targets
  - Can test IP properties in isolation
  - Any sub-path that fails any test vetoes the end-to-end performance
  - Out of scope components (e.g. host software) don't matter
    - If the IP layer passes all tests
    - and the transport and rest of end system are state of the art
    - then applications can be expected to meet the performance targets

Why Model Based Metrics matter

- The approach
  - Open loop Congestion control systems
    - Eliminate out-of-scope influences on the traffic
  - Test with traffic that resembles long RTT TCP
    - But decoupled from the details of the actual path
  - Measure the delivery statistics
    - Pass if better than required by models
    - No long term state in the system
- Key new properties
  - Vantage independence
    - Target RTT effects traffic and success criteria via models
    - Test RTT affects neither traffic nor success criteria
  - IP level tests are generally a ctionable by ISPs
  - Tests can be independently verified
  - Isolate effects of “out of scope” components
    - Mostly inside the model

Document Updates (Draft -02)

- It is now logically complete
  - No substantial missing material
  - The structure & flow are good
  - Some of the tests (section 8) need more detail
- Much of attention to terminology and consistent usage
- It defines a framework for defining “Target Diagnostic Suites”
  - The examples in the document could be fleshed out to be full metrics
- Added a (draft) new alternate statistical criteria section
  - Need to better understand the statistics
  - Further evaluation needed

Future work and open issues

- Tighten up some of the testing details
  - Tie to existing metrics and tools where possible
- More testing
- Feedback!

Backup Slides

Overview

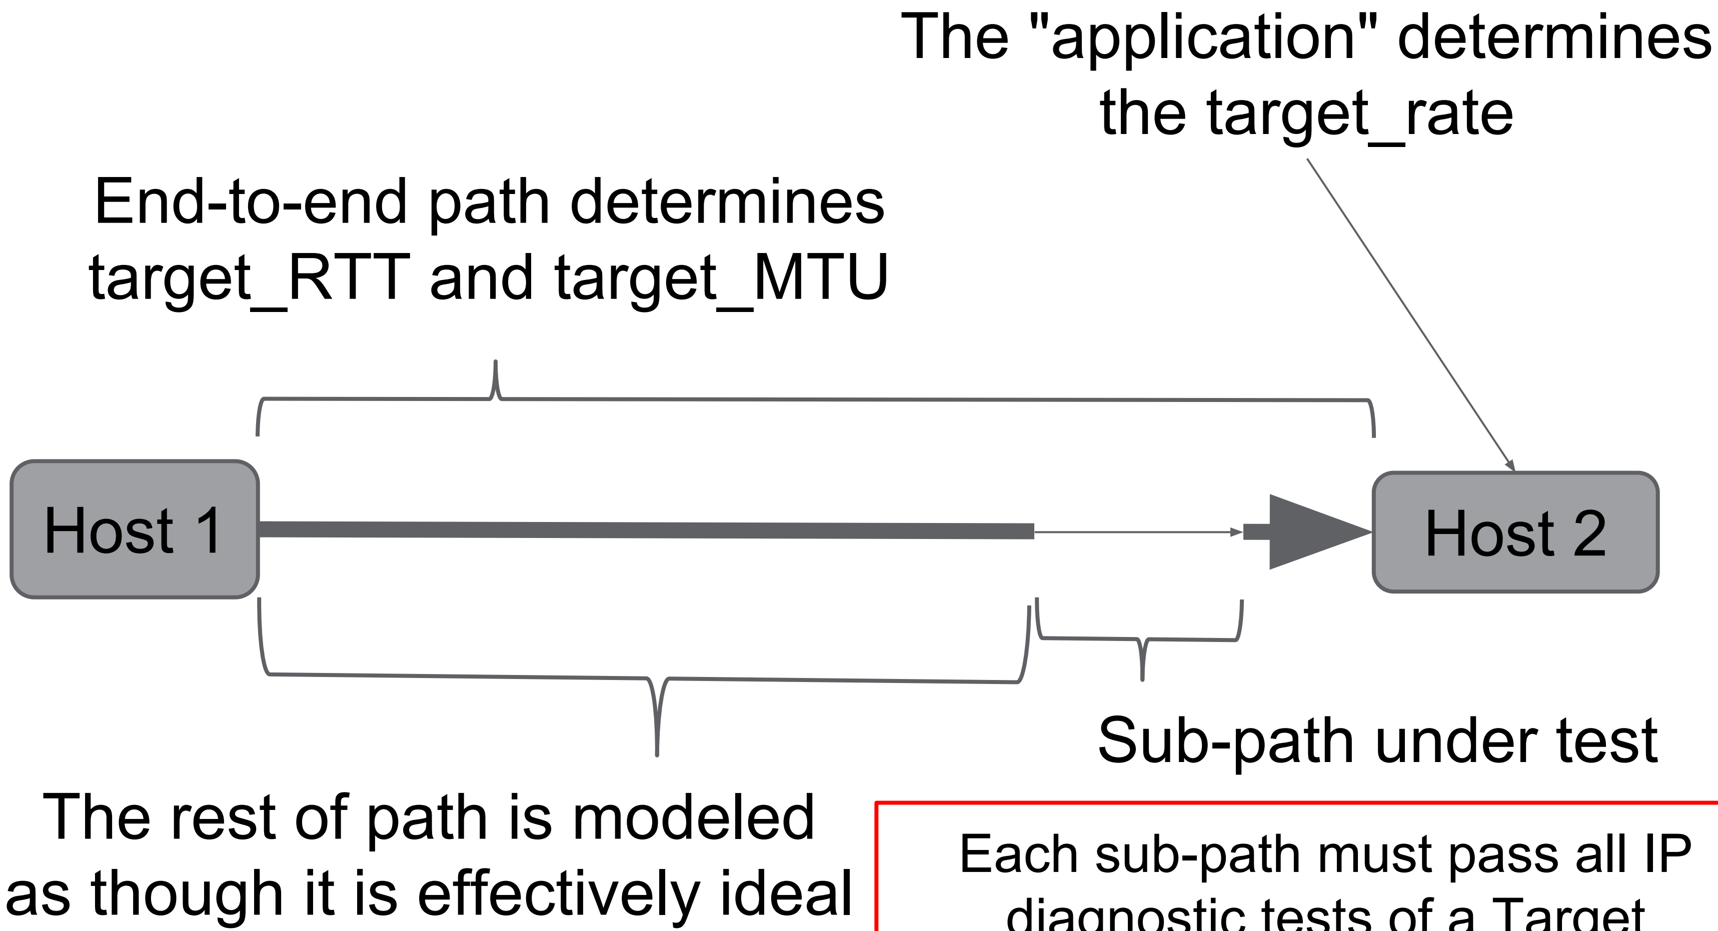

Each sub-path must pass all IP diagnostic tests of a Target Diagnostic Suite (TDS).

Overall Methodology

- Choose end-to-end Target Application (TCP) Parameters
  - Target_data_rate, target_RTT, and target_MTU
- Compute common model parameters
  - target_pipe_size - required average window size (packets)
  - target_run_length - required spacing between losses/ECN marks, etc
- Generate a Targeted Diagnostic Suite (TDS)
  - Pass/Fail/Inconclusive tests of all important IP properties
    - Average spacing between losses (run length)
    - Sufficient buffering at the dominant bottleneck
    - Sufficient tolerance for IF rate bursts
    - Appropriate treatment of standing queues (AQM, etc)

Example: HD Video at moderate range (50 mS)

- Target: 5 Mb/s (payload) rate; 50 mS RTT; 1500 Byte MTU
- Model:
  - Target_pipe_size = 22 packets
  - Target_run_length = 1452 packets
- Computed TDS:
  - Run length longer than 1452 packets (no more than 0.069% loss)
  - Tolerates 44 packet slowstart bursts (twice the actual bottleneck rate)
    - (Peak queue occupancy is expected to be 22 packets)
  - Tolerates 22 packet bursts at server interface rate
    - (Peak bottleneck queue also expected to be 22 packets)
  - Standing queue test:
    - First loss/ECN is more than 1452 packets after the onset of queueing
    - First loss/ECN is no later than 3*1452(?) packets after queueing onset
      - Precise success criteria still under evaluation

An easier (combined) test procedure

- Fold most of the TDS into a single combined test
- Send 22 packet server rate bursts every 50 mS
  - Must average <1 loss/ECN every 66 bursts (1452 packets)
  - This has the same average data rate
  - ...same stress on the primary bottleneck (although more frequent)
  - ...same or higher stress on the rest of the path
- Downside: symptoms become ambiguous
- This test may actually be too conservative
  - A path that can withstand this test is likely to meet a higher target
  - This was the motivation for "derating"

Quasi-passive under streaming content delivery

- Diagnosis as a side effect of delivering real content
  - e.g. Using RFC 4898 - TCP ESTATS MIB
- Requires non-throughput maximizing traffic
  - To avoid self inflicted congestion
  - E.g. any streaming media < target_rate
- Requires serving RTT < target_RTT
- Compute test_window = target_data_rate*serving_RTT
- Clamp serving cwnd to test_window
  - Average rate over any full RTT will be smaller than target_rate
  - All bursts will be smaller than test_window (also target_pipe_size)
  - Compute run length from actual delivery statistics