Modeling Expected and Anomalous Performance of End-to-end Application Workflows

Prasad Calyam, Ph.D.
pcalyam@osc.edu

DOE COMBINE Workshop
September 2012

Topics of Discussion

- Application Workflow “Agendas”
- Expectation Management in Agendas
- Characterization and Modeling of Agendas
- Lessons from Prior Work
- Some Larger Questions…

Application Workflow "Agendas"

- Distributed computing applications in DOE community have "inherent workflow agendas"
  - Application: Bulk file transfers from research instrumentation sites in the LHC data transfers from Tier-0 to Tier-1 and Tier-2 sites
    - E.g., Agenda: Data sharing amongst worldwide collaborators for replicating results, and refining conclusions in LHC Tier-2 site collaborations
  - Application: Multi-user remote instrumentation steering and visualization in Remote access of PNNL Confocal microscopes in GTL project
    - E.g., Agenda: Remote analytics for real-time experimentation in the ITER inter-pulse data analysis using simulation codes at remote supercomputers

User QoE needs to be assessed over "Agendas" versus just at a session flow level; weakest component affects overall user productivity

Expectation Management in DOE Agendas

What could be the expectations…

- Substantial infrastructure investments are being made, hence there are high application performance expectations from users
- Examples of user expectations could include:
  - (a) moving a Terabyte of LHC data within 6 hours between international collaborator sites
  - (b) smooth remote steering of the PNNL Confocal microscope that generates 12.5 Gbps high-definition video stream per camera to deliver “at-the-instrument” user experience for multiple geographically dispersed remote users
  - (c) a west-coast user experiencing reliable performance over long time-periods when manipulating simulation codes and their graphical user interfaces pertaining to 2 to 3 Gbytes ITER inter-pulse data being transferred and analyzed every 20 minutes at NERSC supercomputer
  - (d) ….

Characterization and Modeling of Agendas

What we could be doing…

- Salient workflow agenda flows need to be characterized and modeled on *realistic testbeds*
  - User, application and network interplays could be understood
  - Dominant factors that affect performance could be identified to reduce sample space of data
  - 'Ideal' and 'Performance bottleneck' states could be catalogued in conjunction with *user surveys*
  - Correlation analysis can be performed on observed phenomena to compare with expectations

Characterization and Modeling of Agendas (2)

What we could be doing…

- "Expectation-management" tools that allow repeatability need to be built, deployed and refined
  - They can be used to exercise, analyze and visualize if inherent workflow agendas are performing as expected or are anomalous (particularly if they are faulty)
  - They can be integrated to extend familiar and widely-adopted middleware software interfaces (e.g., Pegasus Workflow Management System, perfSONAR, …)
  - Gather real-world data, re-train models - to refine tools and make them more relevant, and keep them relevant!

Lessons from Prior Work

Case study with Videoconferencing Workflow Agendas…

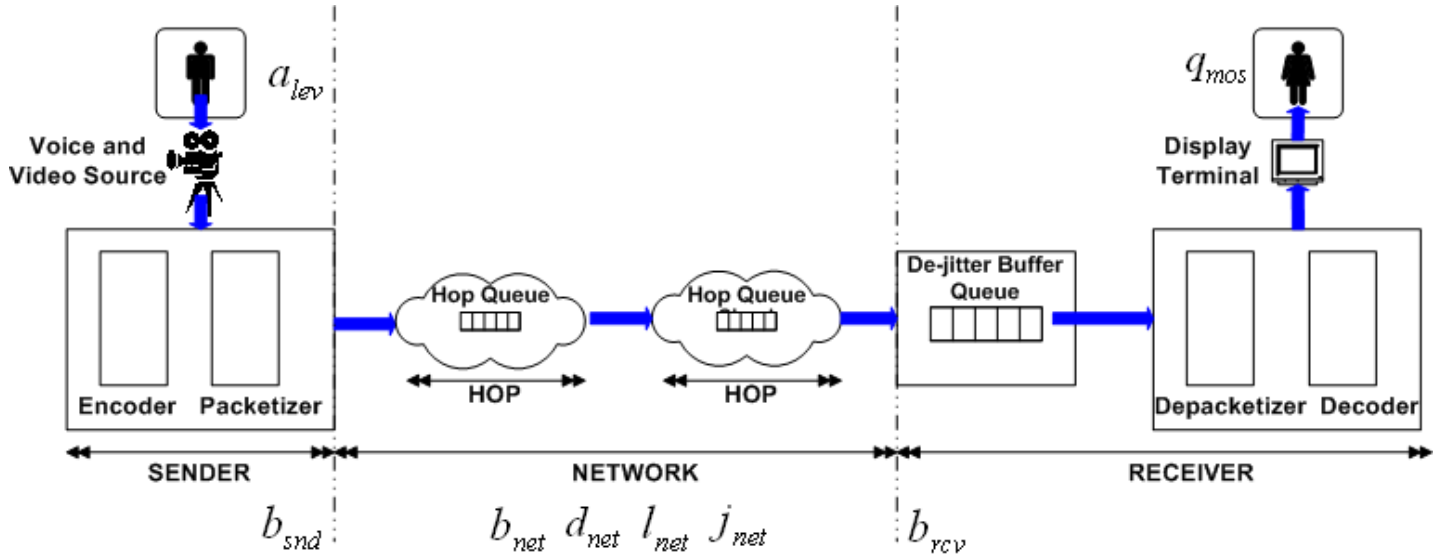

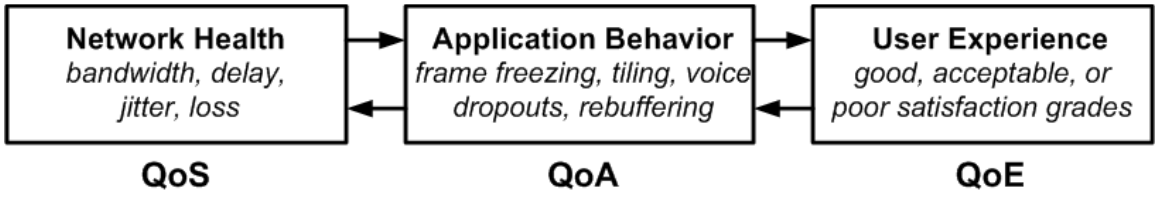

$$QoS = f(QoA, QoE)$$

$$QoE = f(QoS, QoA)$$

Video a_{lev}

- Low a_{lev} - Slow body movements and constant background; E.g. *Claire* video sequence
- High a_{lev} - Rapid body movements and/or quick scene changes; E.g. *Foreman* video sequence
- 'Listening' versus 'Talking'
  - Talking video a_{lev}(i.e., High) consumes more bandwidth than Listening video a_{lev} (i.e., Low)

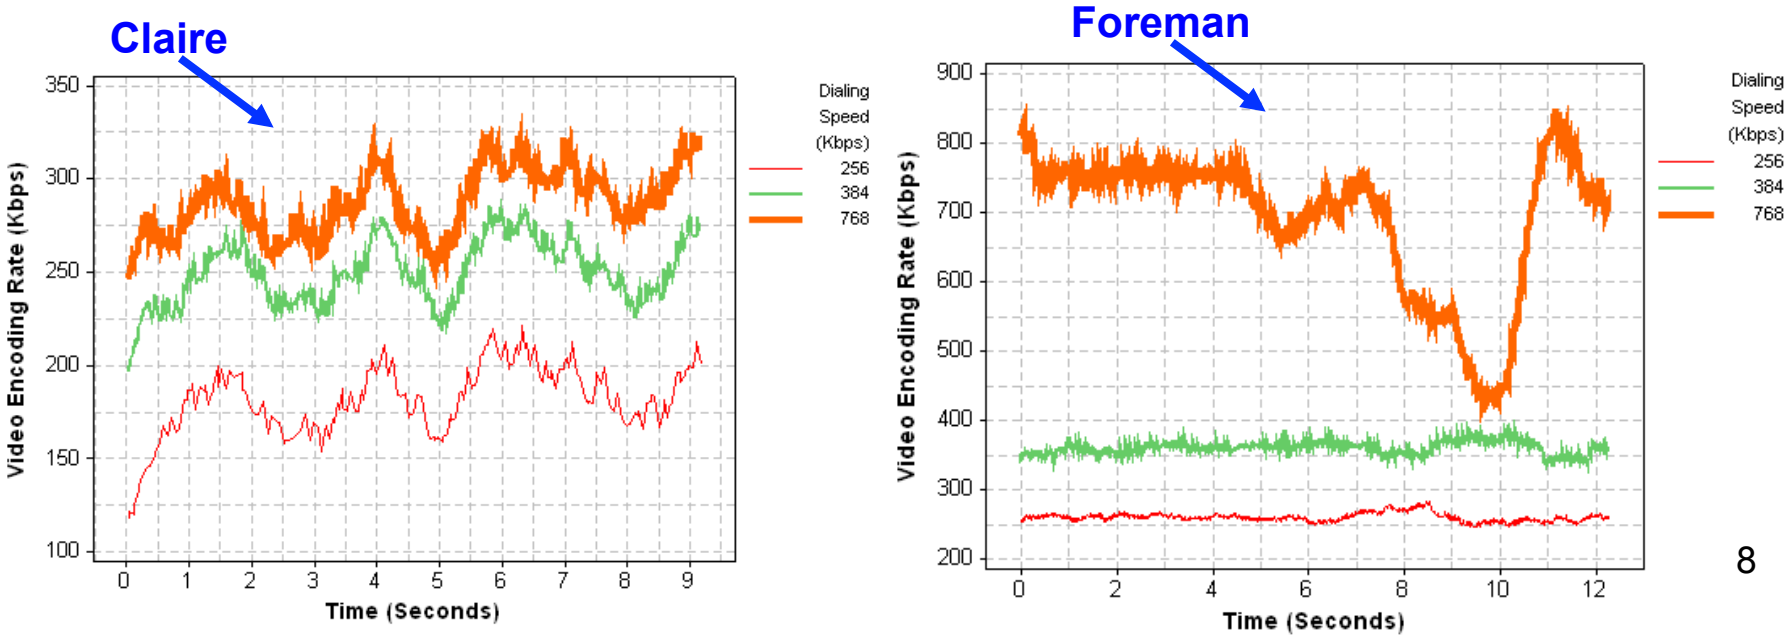

Agenda Exercising "Vperf" Tool

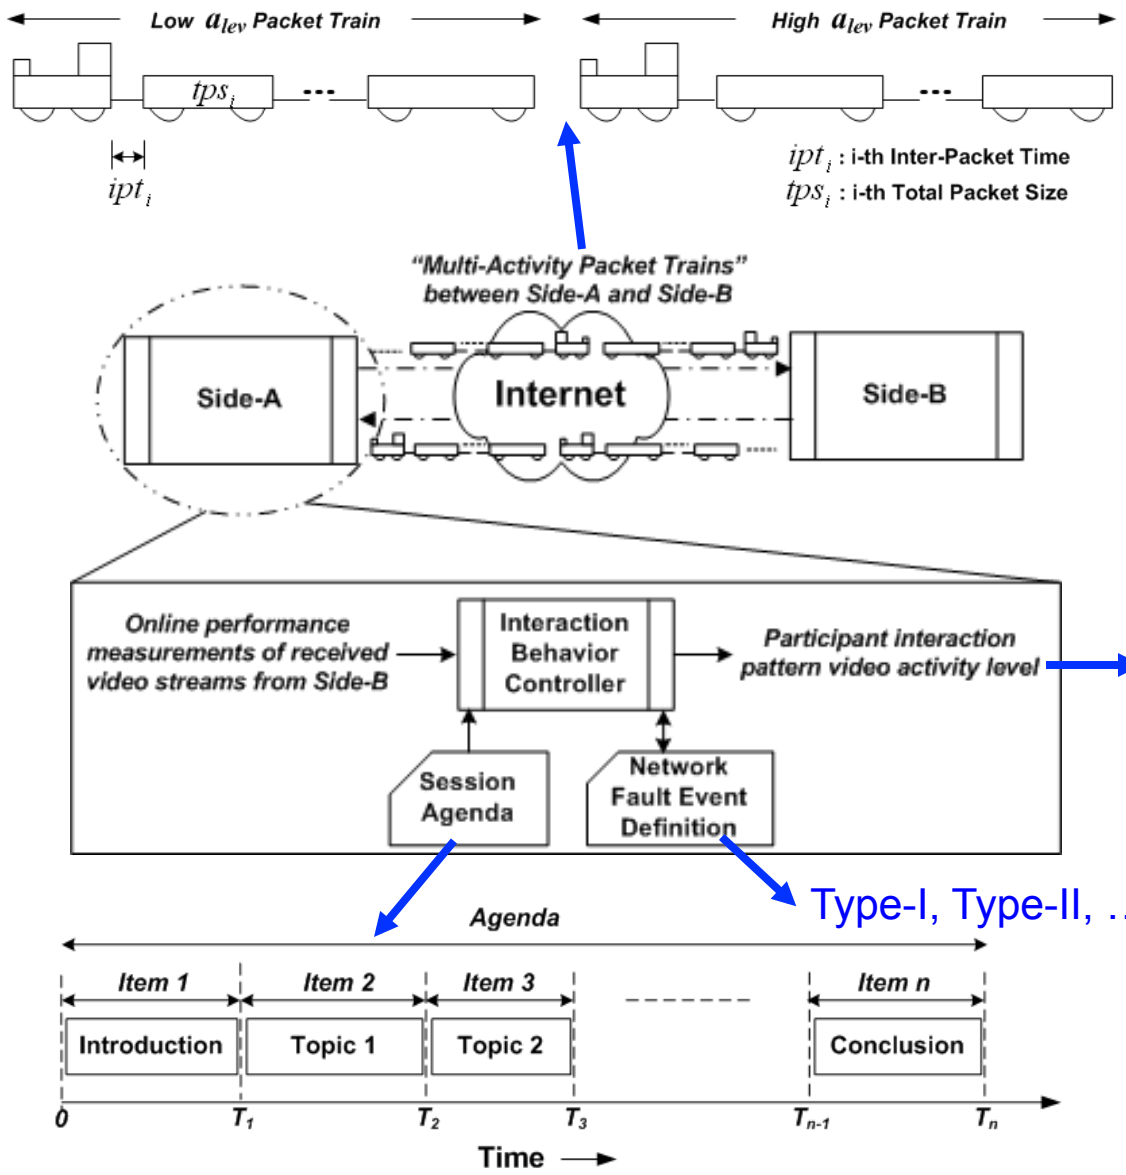

"Can you repeat what you just said?"

"This line is noisy, lets hang-up and reconnect…"

Type-I, Type-II, … Type-N fault detection

Vperf Tool Implementation

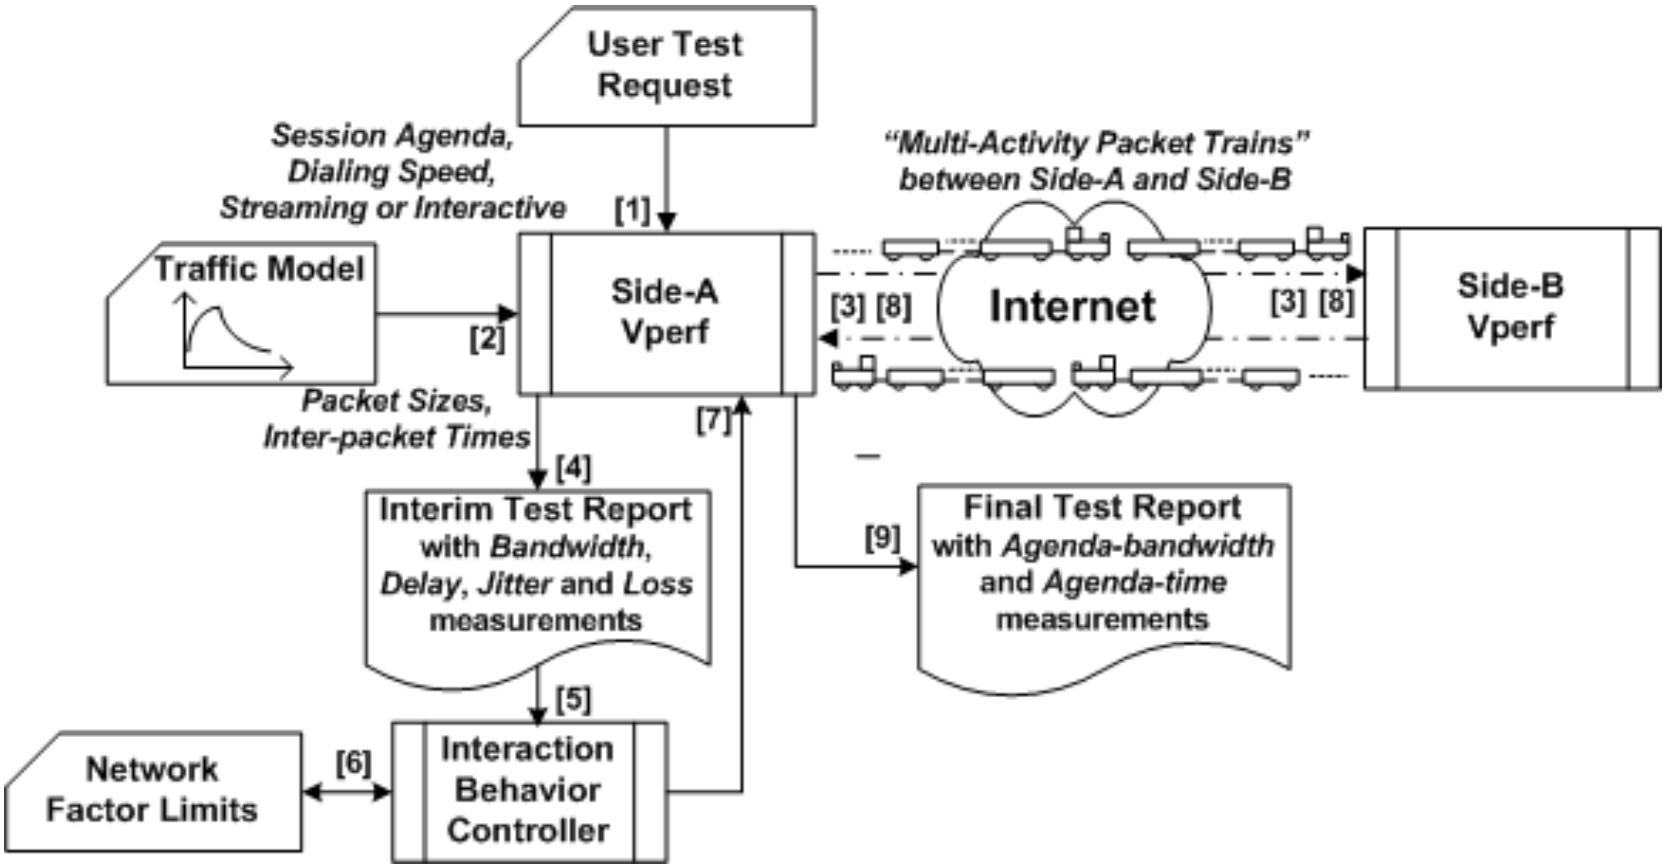

Vperf Measurements Evaluation

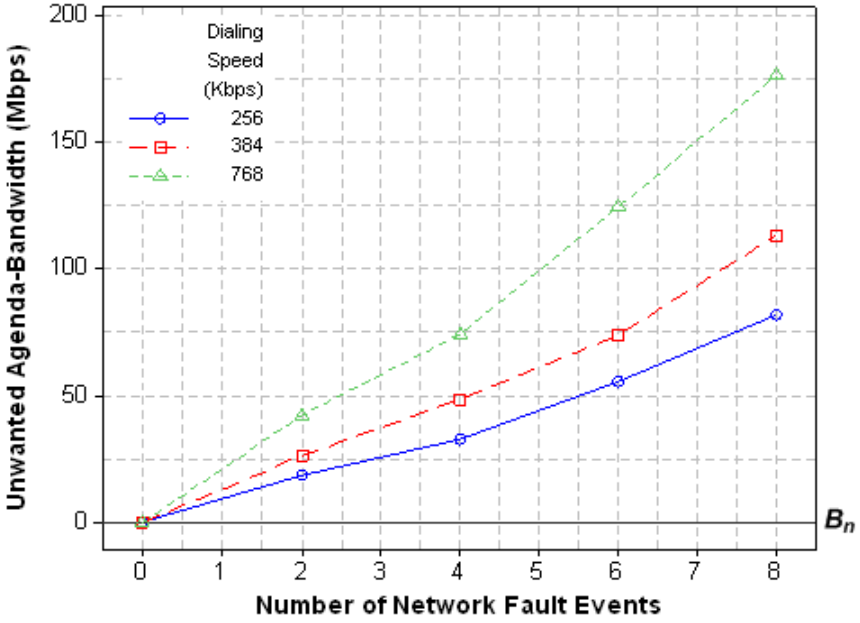

(a) Impact of Type-I Network Fault Events on Unwanted Agenda-Bandwidth

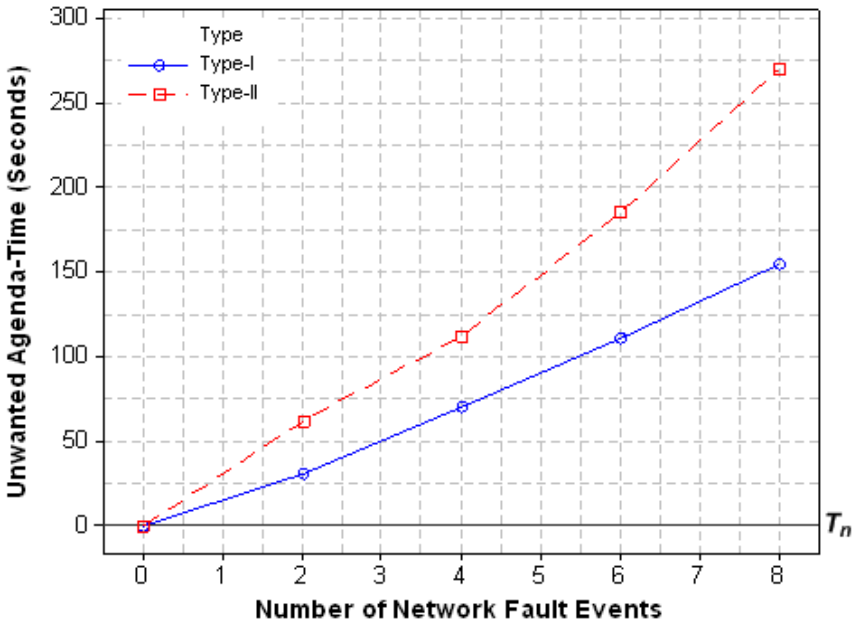

(b) Impact of Type-I and Type-II Network Fault Events on Unwanted Agenda-Time

More than one major product vendor and service provider have found use for the Vperf tool…

Some Larger Questions…

What challenges we face with salient workflow agendas…

- How to account for human behaviors when interacting with remote instruments or other humans and set “expectations”?
- How to characterize and model combination of both bandwidth-intensive and latency-sensitive traffic flows under ideal and fault conditions? How to *find users* and *conduct* related user surveys?
- How to account for diverse and ever changing end-system, network technologies and build general models?
- What metrics and performance bounds can be used in agenda-exercising tools to collect parameters for large-scale simulations?
- How to instrument real-world user applications with such tools to derive tangential benefits to users/operators in the short-term, *but also help continuously re-train models*?

Thank you for your attention!

Questions?

Voice and Video Packet Streams

$$b_{snd} = b_{voice} + b_{video} = tps_{voice}\left(\frac{b_{codec}}{ps}\right)_{voice} + tps_{video}\left(\frac{b_{codec}}{ps}\right)_{video}$$

- Total packet size (tps) – sum of payload (ps), IP/UDP/ RTP header (40 bytes), and Ethernet header (14 bytes)
- Dialing speed is $\lceil b_{video} \rceil$; $\lceil b_{voice} \rceil$ = 64 Kbps fixed for G.711 voice codec
  - Voice has fixed packet sizes (tps_{voice} ≤ 534 bytes)
  - Video packet sizes are dependent on a_{lev} in the content