Aalto University
School of Science
and Technology

Towards More Adaptive Voice Applications

Jörg Ott <jorg.ott@aalto.fi>

Md. Tarikul Islam, Cheng Luo
Nagasai Panchakarla, Anssi Turkulainen

13 October 2011

Adaptive Applications

- Varying radio link conditions and coverage/connectivity
  - Often less predictable changes
  - Congestion vs. errors
- Varying path characteristics in the Internet
  - Variable load
  - Route changes

Delay **Loss** **Date Rate** **MTU**

- "Fair" sharing of communication resources
  - Utilize available resources effectively, but do not overload
- Obtain sufficient application performance in spite of the above

Aalto University
School of Science
and Technology

Implementing Adaptivity: Examples

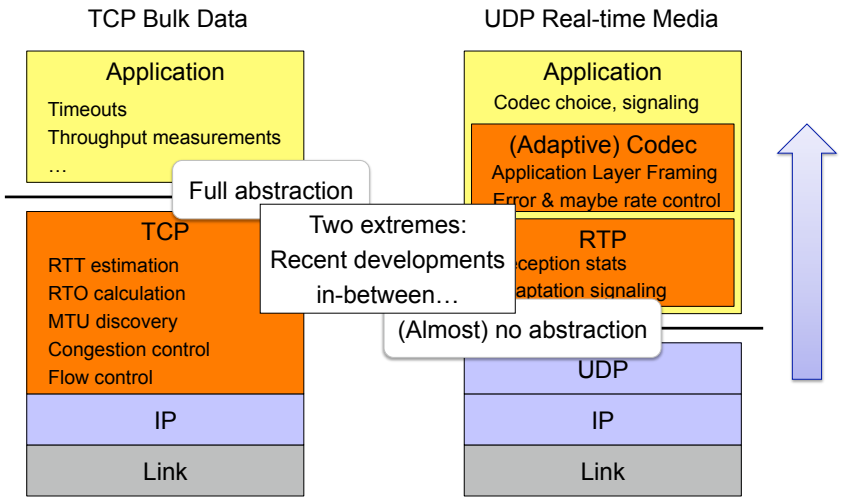

Aalto University
School of Science
and Technology

Implicit Assumptions

1. Applications are capable of adapting across a sufficiently wide range of communication characteristics

OR

2. The best effort service delivered will just be good enough for the applications to work well

Aalto University
School of Science
and Technology

Adaptive Protocols and Applications

Delay | **Loss** | **Date Rate** | **MTU**

- Limitations in the operational range
  - Minimum performance requirements needed for acceptable operation
  - Maximum they are (practically) able to utilize (mostly data rate)
  - Capability to “cancel out” over- and underperforming over time

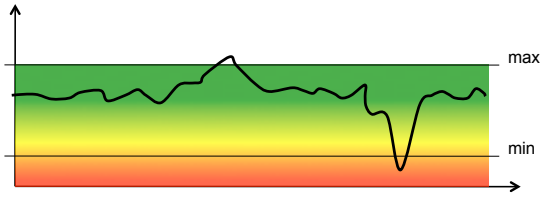

- With insufficient performance, users may get annoyed, give up, need to retry (later), …

Aalto University
School of Science
and Technology

Some Examples…

	Delay	Loss	Data rate	MTU
Bulk data	Don’t care as long as TCP does not stall or disconnect… (P2P even better)			
Interactive web	RTT<300ms	< 2%	100 kbit/s – 1 Mbit/s	1500 bytes ok
	Interactivity = f(loss, delay) needs to be sufficient			
Streaming	seconds	< 1%	100 kbit/s – 100 Mbit/s	1500 bytes ok (could be larger)
	Data rate = f(loss, delay) needs to be sufficient			
VoIP	< 200ms	< 5%	4 kbits – 100+ kbit/s	< 100s bytes

Aalto University
School of Science
and Technology

When Best Effort is Not Enough…

- Mobile Internet access and wireless networks

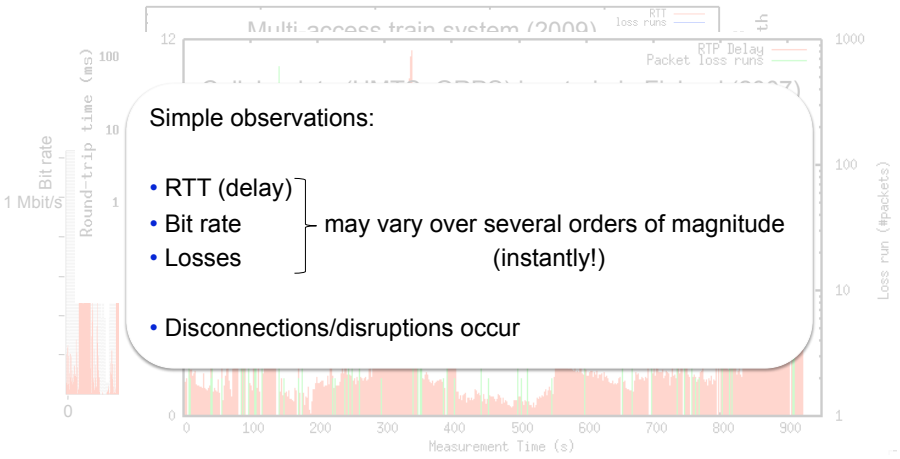

Simple observations:

- RTT (delay)
- Bit rate
- Losses

} may vary over several orders of magnitude (instantly!)

- Disconnections/disruptions occur

Aalto University
School of Science
and Technology

Adaptation continued…

Given a path and the need to send a certain amount of information…

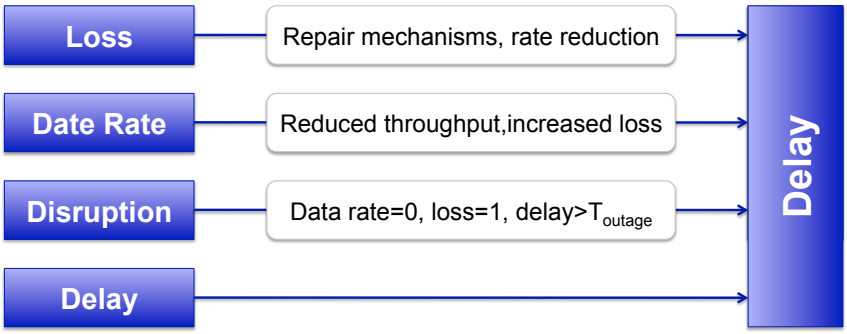

Aalto University
School of Science
and Technology

A Case Study on Asynchronous Voice

Voice Messaging

- Boring…
  - Answering machines
  - Walkie-talkies
  - Push-to-talk (over Cellular)
- …but useful in diverse scenarios…
  - Decouple sender and receiver (just like email)
- …especially in opportunistic ad-hoc networks…
  - No infrastructure
  - No stable paths or no end-to-end paths at all
- …or when facing instantaneously insufficient access links
  - Smoothen utilization of cellular infrastructure
  - Expand multiplexing in the time domain (particularly when mobile)

Aalto University
School of Science
and Technology

DT-Talkie: Asynchronous Voice

- Push-to-talk for single-hop or multi-hop opportunistic networks
  - May also use infrastructure (cellular, WLANs)
- Reliable (hop-by-hop) communication to deal with losses
  - Speech quality is not impacted, only delay is
- Delay tolerance: decoupling sender and receiver
  - Asynchronous interaction without dedicated mediator
  - Optional support via infrastructure servers for rendezvous

Aalto University
School of Science
and Technology

DTN-based Voice

- Plain and simple: record – send – forward – receive – playback
  - Based upon user-indicated (button press) statements

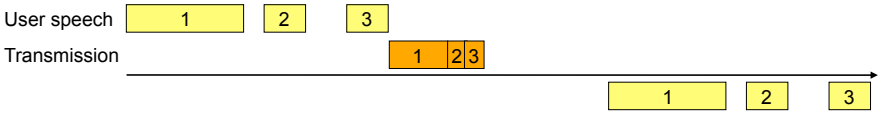

- Subtleties: message size?
  - Semantic fragmentation (Application Layer Framing)
  - Keep talkspurts together ("MTU")
  - Good connectivity and short messages: interactive communication workable

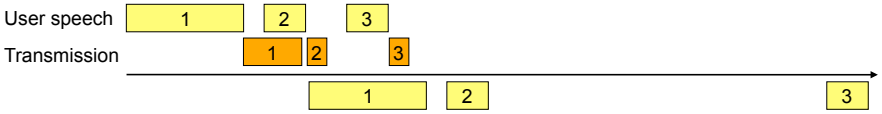

Aalto University
School of Science
and Technology

And if the network works just fine…?

- From adaptive message size…

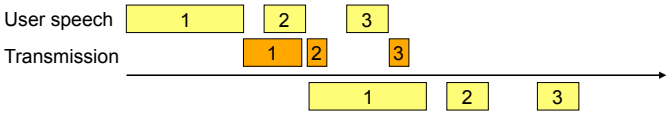

- To constant message size…

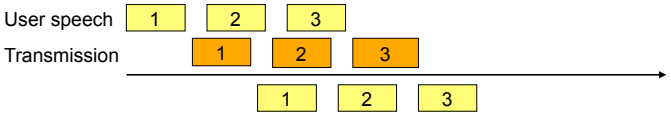

- To small packets: synchronous voice!

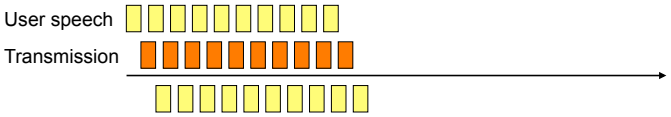

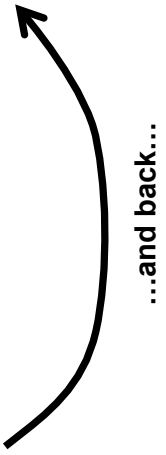

Aalto University
School of Science
and Technology

Increasing Freedom for Adaptation

Aalto University
School of Science
and Technology

A Simple SIP Application

- Vertical adaptivity for VoIP yields a certain operational range

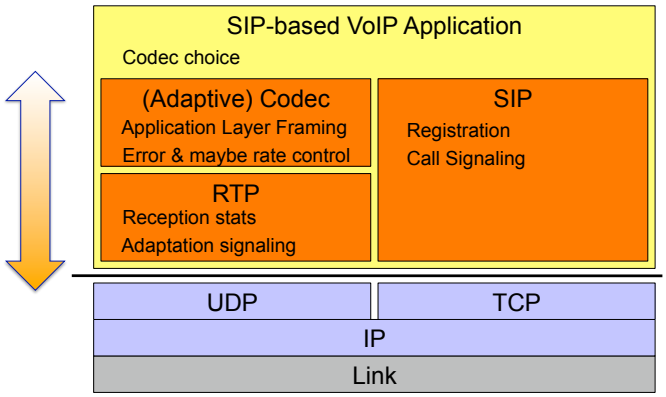

Aalto University
School of Science
and Technology

Taking a step back looking at the semantics…

- Advancing individual (vertical) adaptation per function…
  …to integrate them across different ones (horinzontal adaptation)

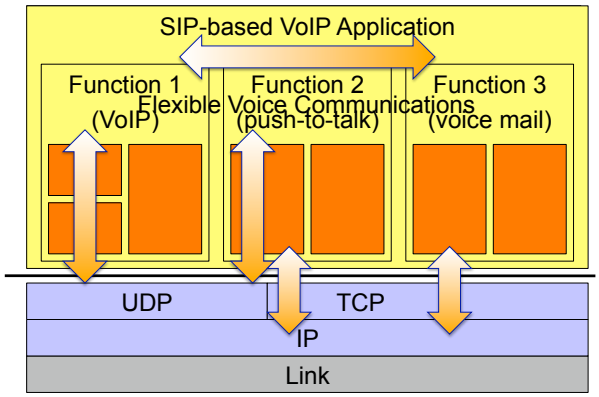

Aalto University
School of Science
and Technology

…to fully adaptive Voice

- Micro adaptation: error and rate control
- Macro adaptation: data unit size and reliability as a function of path properties and delay tolerance

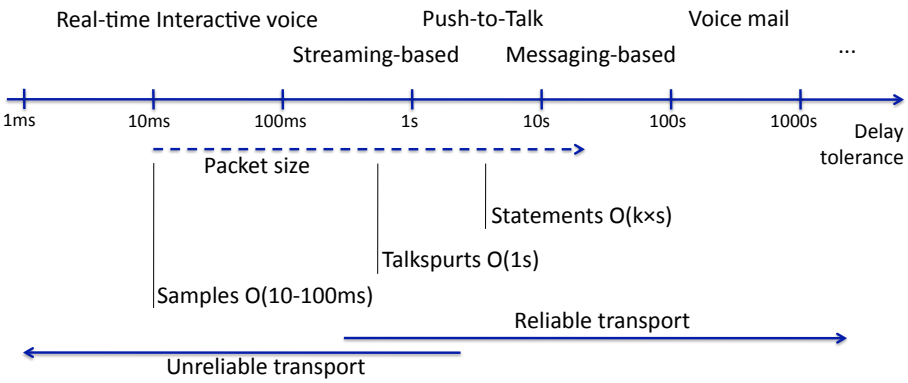

Aalto University
School of Science
and Technology

Impact on Capacity Sharing

- Don't limit competing for capacity only on a short time scale
  - Extending the altruistic behavior of TCP congestion control further

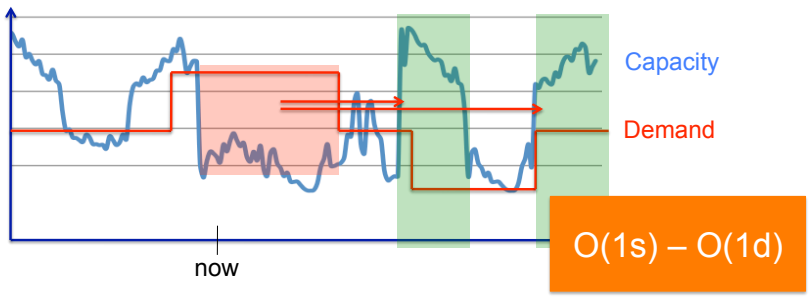

- Expend resources (energy, spectrum) when it makes sense

Aalto University
School of Science
and Technology

A Trivial Algorithm

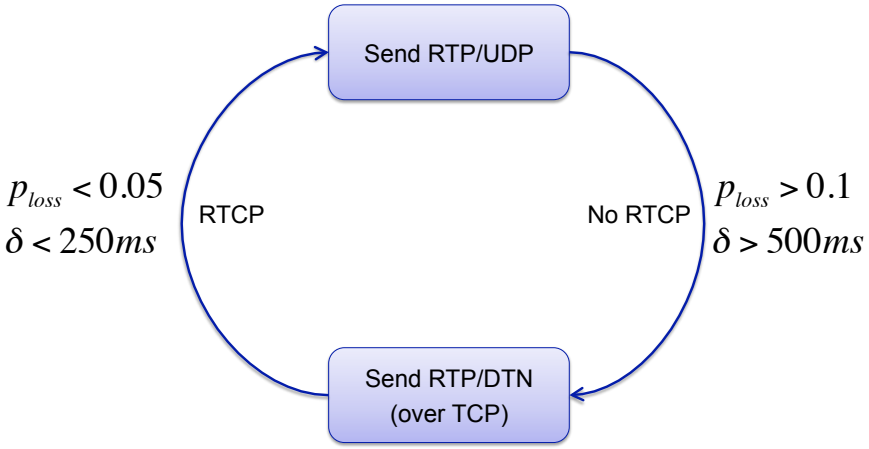

Aalto University
School of Science
and Technology

More Flexibility for Known Content

- Example: Streaming
  - Buffering does already part of the job (but still competes for capacity)
- Anticipate bottlenecks
  - E.g., Learning from history (own, others)

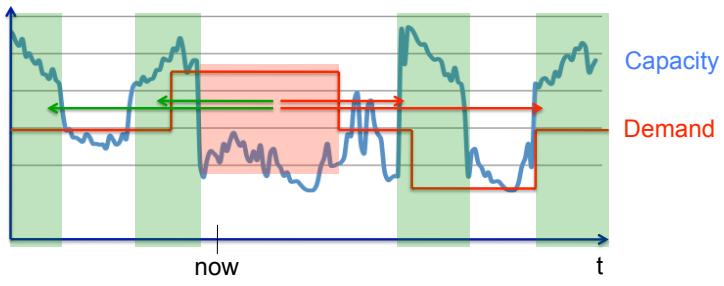

Aalto University
School of Science
and Technology

Some Random Thoughts…

- How is adaptivity supposed to work?
  - Specific vs. generic monitoring mechanisms?
  - Time-scale?
  - Relying on (predicting) future communication opportunities?
- How about fairness…?
  - Is additional delay another dimension to consider?
  - E.g., more data in return for less urgent data?
- Incentives?
- Can some common abstraction be provided?

Aalto University
School of Science
and Technology

Conclusion

- Allowing for delay tolerance may extend adaptation capabilities
- Requires looking (again) at the application semantics
- May not be as evil as it seems
- Endpoints and user interfaces matter
- Networks (or network-related services) may provide support

Aalto University
School of Science
and Technology

EC FP7 project CHIANTI
http://www.chianti-ict.org/

Finnish ICT-SHOK Future Internet project:
http://www.future-internet.fi/