Experimental Congestion Control & the IETF/IRTF

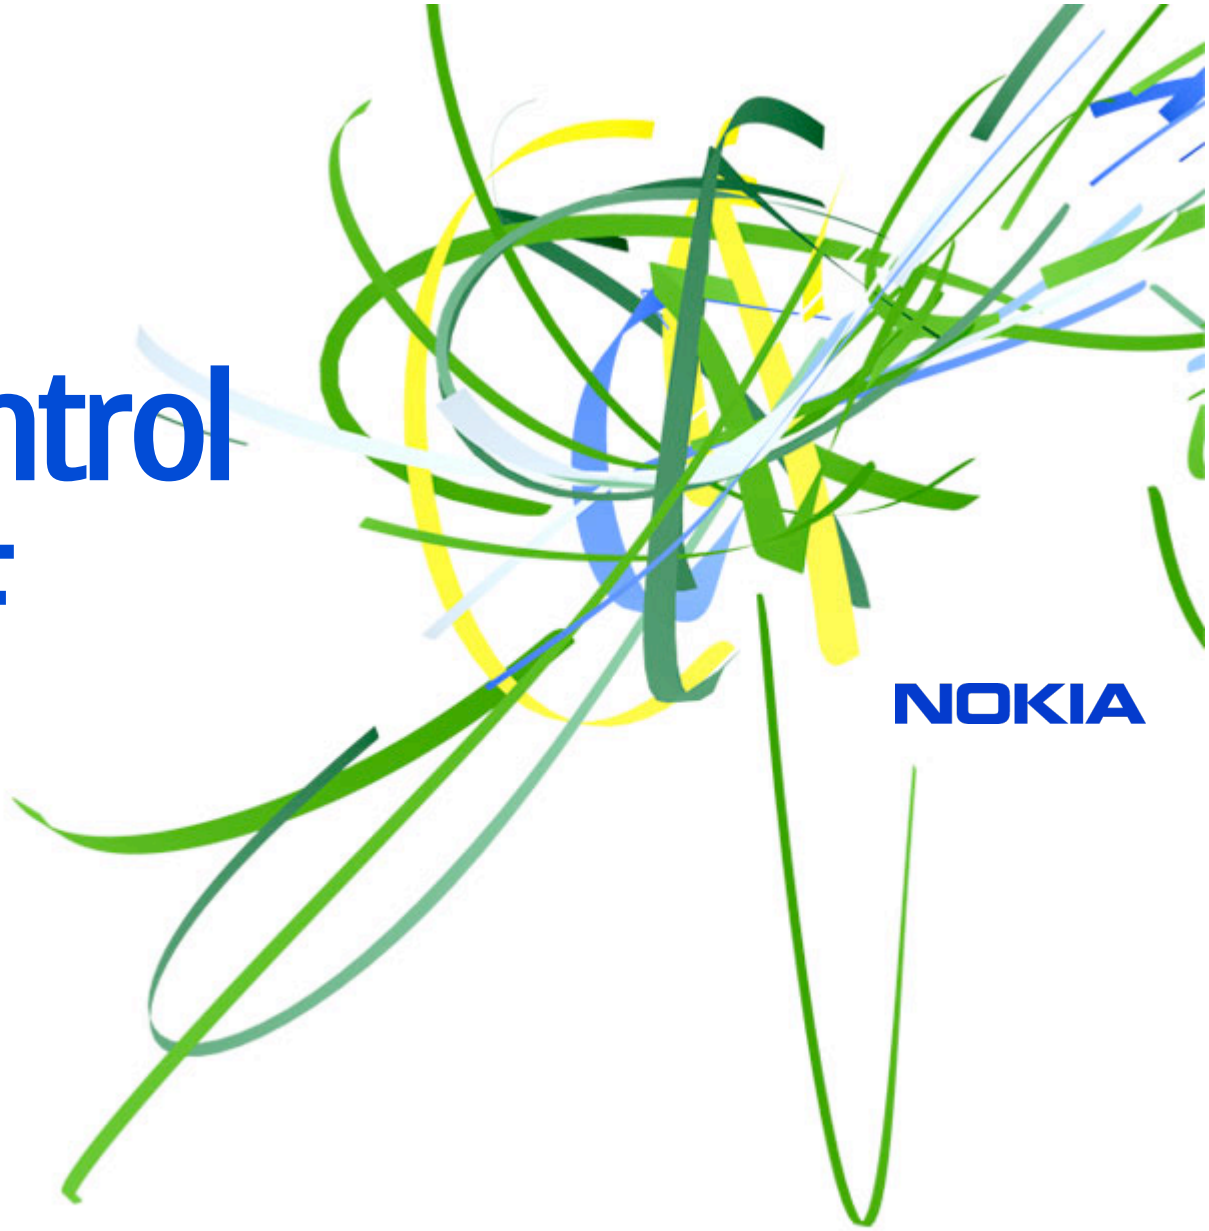

Lars Eggert

IRTF ICCRG Meeting

USC/ISI, Marina del Rey, CA, USA

2007-2-12

Where are we?

- lots and lots of promising congestion control research
  - for fat paths, but also other scenarios
  - some schemes useful on an Internet-wide scale
- potential for benefit is usually demonstrated
  - papers, etc.
- potential for bad interactions is less well investigated
  - because it's hard & boring :-)
- metrics & scenarios for comparing schemes are unclear
  - which TCP variant is "the best" and what does that mean?

Where do we want to go?

- we'd all like to evolve TCP forward
  - TCP = Internet-wide congestion control standard
  - safe in all environments, performs OK in many
- standard ≈ agreed-upon social contract for CC
  - "how we all use the shared resource we communicate over"
- safe ≈ prevents congestion collapse, some fairness
- an "evolved TCP" needs to be a safe standard
  - not safe ➔ Internet melts down
  - not standard ➔ interactions between different CC (safe?)
potential for arms race
hard enough to get one variant right

Why is there an issue?

- interest in new CC features for major TCP stacks
  - some new CC has already leaked out onto the Internet
  - some stacks move beyond RFC mechanisms
- we don't know what major stacks do anymore
  - insufficient documentation, insufficient review
- is this safe? what is safe?
  - optimistic view: "well, the Internet hasn't melted yet"
  - pessimistic view: "but we don't know if it will stay this way"
- the IETF is the originator and maintainer of TCP
  - we want to provide the venue for evolving it

IETF/IRTF involvement

- encourage the proposers and implementers of new CC to participate in the IETF/IRTF

- goals
  1. document current stack behavior

     "we'd like you to know, this is what our stack does"

  2. proposals for eventual standardization

     "we think this may eventually become a recommended mechanism, and would like people to experiment with it..."

Class 1 – Document current stack behavior

- goal: documentation to inform the community
  - subsets of RFCs implemented or ignored & why
  - which additional mechanisms implemented & why
  - lessons learned
- existing examples
  - deployed TCP reactions to ICMP soft errors
  - FreeBSD: SYN cookie extensions
- future examples?
  - Linux: delayed-ACK suppression during slow-start
  - Vista: impact of enabling ECN, window-scaling, etc.
- vessel: Internet Drafts intended for Informational RFC, published out of the TSV area

Class 2 – "Experimental" specifications

- goal: mechanisms that may eventually progress onto the standards track
  - "we think this may eventually become a recommended mechanism, and would like people to experiment with it…"
    - "…on the global Internet"
    - "…in scenarios that are restricted in the following ways…"
- vessel: Internet Drafts intended for Experimental RFC
  - technical specification to guide implementers
  - discussion & data in preparation of community consensus
  - Sally's BCP draft has some guidelines
    - draft-floyd-tsvwg-cc-alt-00, soon-to-be draft-ietf-tsvwg-cc-alt-00

Proposed approach, phase 1

- work split between IETF & IRTF
- bring individual Internet Draft to ICCRG first
  - IETF will redirect
  - RFC Editor may want to do similarly
- ICCRG reviews draft & existing body of work
  - “is this safe for limited, experimental use?”
  - on the Internet, or in restricted environments
- after ICCRG consensus, send draft & review to TSV area
- if adopted, publish Experimental RFC out of the TSV area

Proposed approach, phase 2

- assume we have a number of such Experimental RFCs
- we'd eventually like to move one (several?) towards STD
  - "the IETF recommends you implement this"
- need to gather experience with them
- need to evaluate them
  - related IRTF TMRG draft: draft-irtf-tmrg-metrics-06
- how? there is research left to be done
  - the IETF is not a research organization - but the IRTF is
- ICCRG coordinates this effort
  - results feed into a follow-on TSV area effort

Comments?

Discussion points

- what does it mean to “review a protocol?”
  - criteria for “safe”, “good”, etc.
  - pass/fail safety criteria as an initial step?
  - hard test cases vs. soft guidelines?
- list of things proposers need to bring to the RG
  - credible attempt at a technical spec & paper citations with data
  - use TMRG scenarios & metrics?

- RG review is iterative process
- organization of the structure and process of RG reviews
- H-TCP is there, CUBIC and C-TCP coming soon
- clarify the benefit of the process to the proposer