IGP Convergence and Stability: Lets have it both

Cengiz Alaettinoglu
cengiz@packetdesign.com

Why care about convergence?

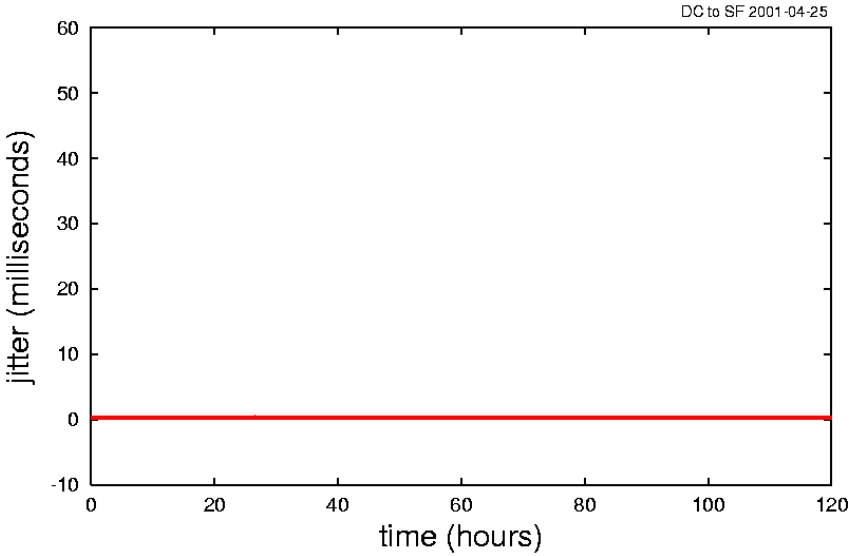

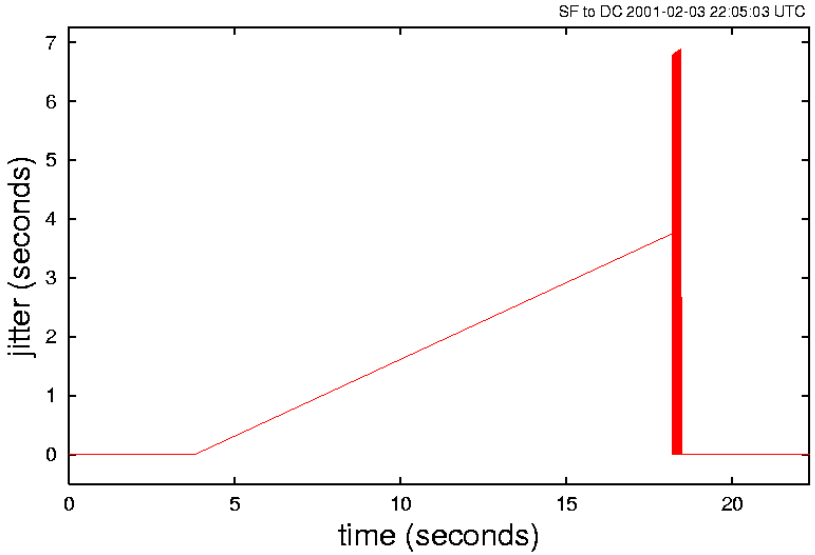

- Your IP network
  - → 69 million *happy* packets
  - → Zero packets lost
  - → 100% jitter < 700µs

- Your IP network during IGP convergence
  - → outages upto 2.5 mins
  - → massive reordering & jitter
    - routing loops result in blender events

(data from NANOG 22 talk)

Convergence vs restoration times

- Convergence time: all routers have heard the news and computed new routing tables
  - → SPF time + propagation delays + per hop scheduling delays
- Restoration time: time to first successful data packet transmission after failure
  - → SPF time + per hop scheduling delays
    - – because the link state packet is ahead of the first data packet by an SPF time

- *I am ignoring detection time and FIB install time here, vendors are way ahead in detection aspect, for FIB install see my feasible next hop talk at Atlanta IETF.*

SPF Times

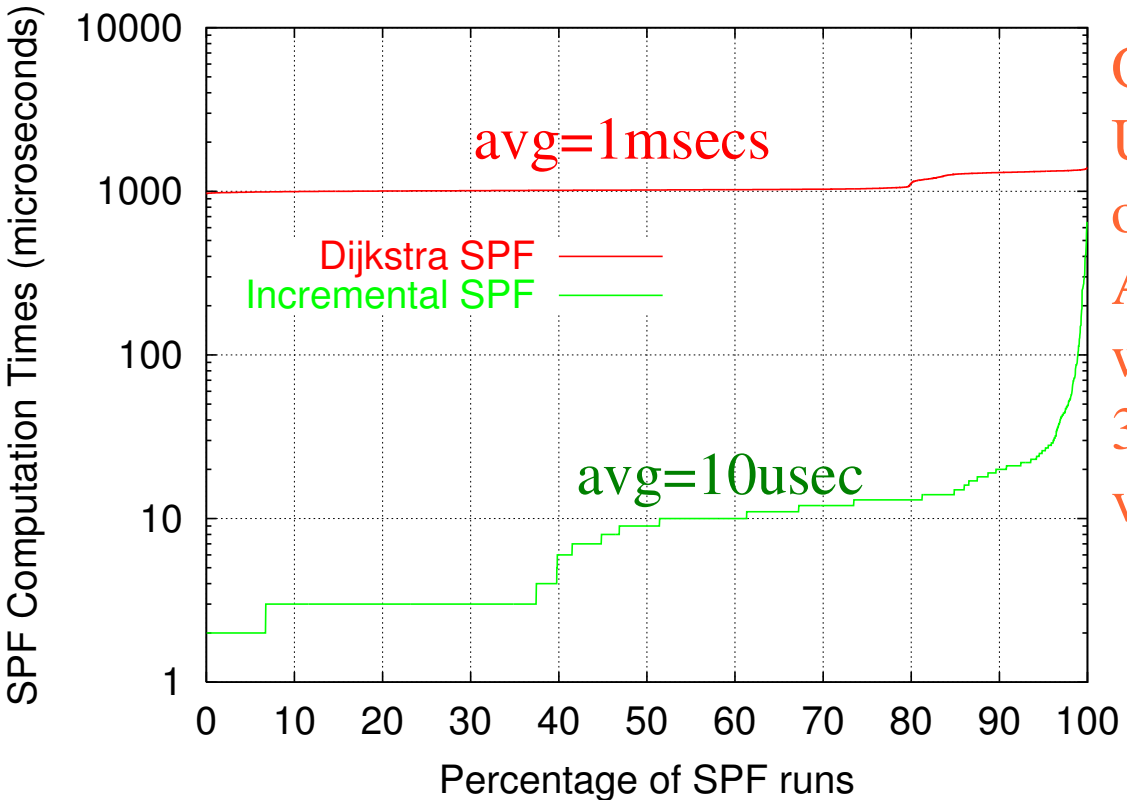

Qwest topology
Using week worth of routing events
Actual spf times were roughly 30 msecs at two vendors' routers

- Benefits of incremental algorithms
  - → scaling to number of nodes
  - → to full mesh (regular SPF goes up to seconds)
  - → less cpu intensive farther from the failure

Convergence vs restoration times: the math

- Convergence time
  - → SPF time + propagation delays + per hop scheduling delays
  - → low hundreds of milliseconds

- Restoration time
  - → SPF time + per hop scheduling delays
  - → tens of milliseconds

Why arent we there?

- → We are afraid (for good reasons) to hurt ourselves!
- → We need a defense mechanism.

Stability vs Restoration Time

- After a certain level of external instability (e.g. flaky layer 2 stuff), routing system itself starts introducing instability, ..., causing a network wide meltdowns
  - → Many ISP examples to choose from

- Defense mechanism: rate limit SPF computation
  - → This hurts convergence time
  - → and causes routing loops (NANOG 24)

We need a better defense mechanism that works and doesnt hurt convergence!

Defense mechanism: damping

- Multiple layers of defense:
  - → At link layer damp flaky links only on recovery
    - never on failure where the convergence matters
  - → Damp flaky links again at routing layer
    - dont trust the device driver writer
    - again damp good news only
  - → Damp routers who don't implement this right
    - dont trust the other vendor
    - again damp good news only, per link?
  - → Damp your SPF (rate limit)
    - only if you are spending > x% of cpu on spf

Challenges

- Understanding the IGP behavior
  - → One set of parameters does not fit all ISPs
  - → Measurement and analysis
- A solid damping implementation
  - → Simulate, emulate, and test using measured/random data
- Parameters
  - → If 10 parameters needs to be configured, it wont happen
    - adaptive parameters w/ good starting defaults
    - aggressiveness configurable
- Winning back ISPs' trust

Acknowledgments

- Graphs are from earlier talks in collaboration w/ Stephen Casner, Haobo Yu, Cha-chi Quan, and Van Jacobson.
- Our ISP partners for comments and for letting us use their topologies for analysis.
- And many in the routing community for constructive criticism and suggestions.