



IEEE 802.1Qbp: Hash Proposal

Brad Matthews & Puneet Agarwal

June 23, 2011

V 0.2

Agenda



- Overview
- Hash Functions Evaluated
- Review Results
- Observations
- New Functions Evaluated -1

Goals



- Identify hashing strategies that provide good flow distribution for multi-hop networks in a deterministic manner

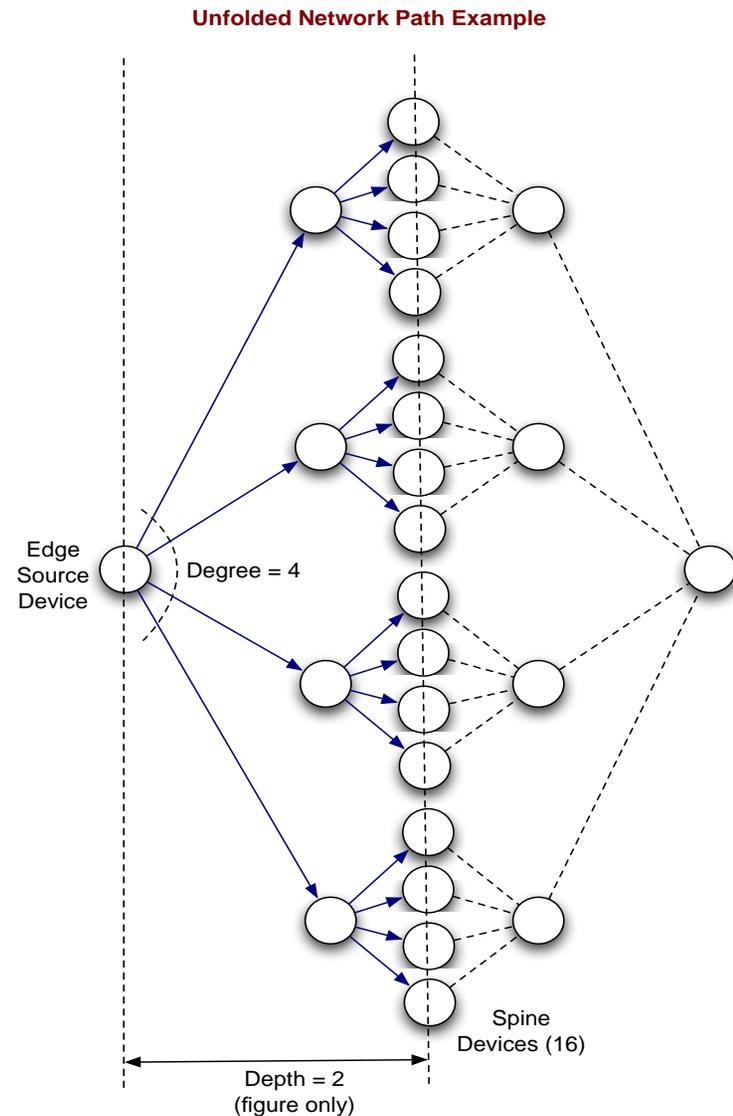
Evaluating Load Balancing Performance

Approach

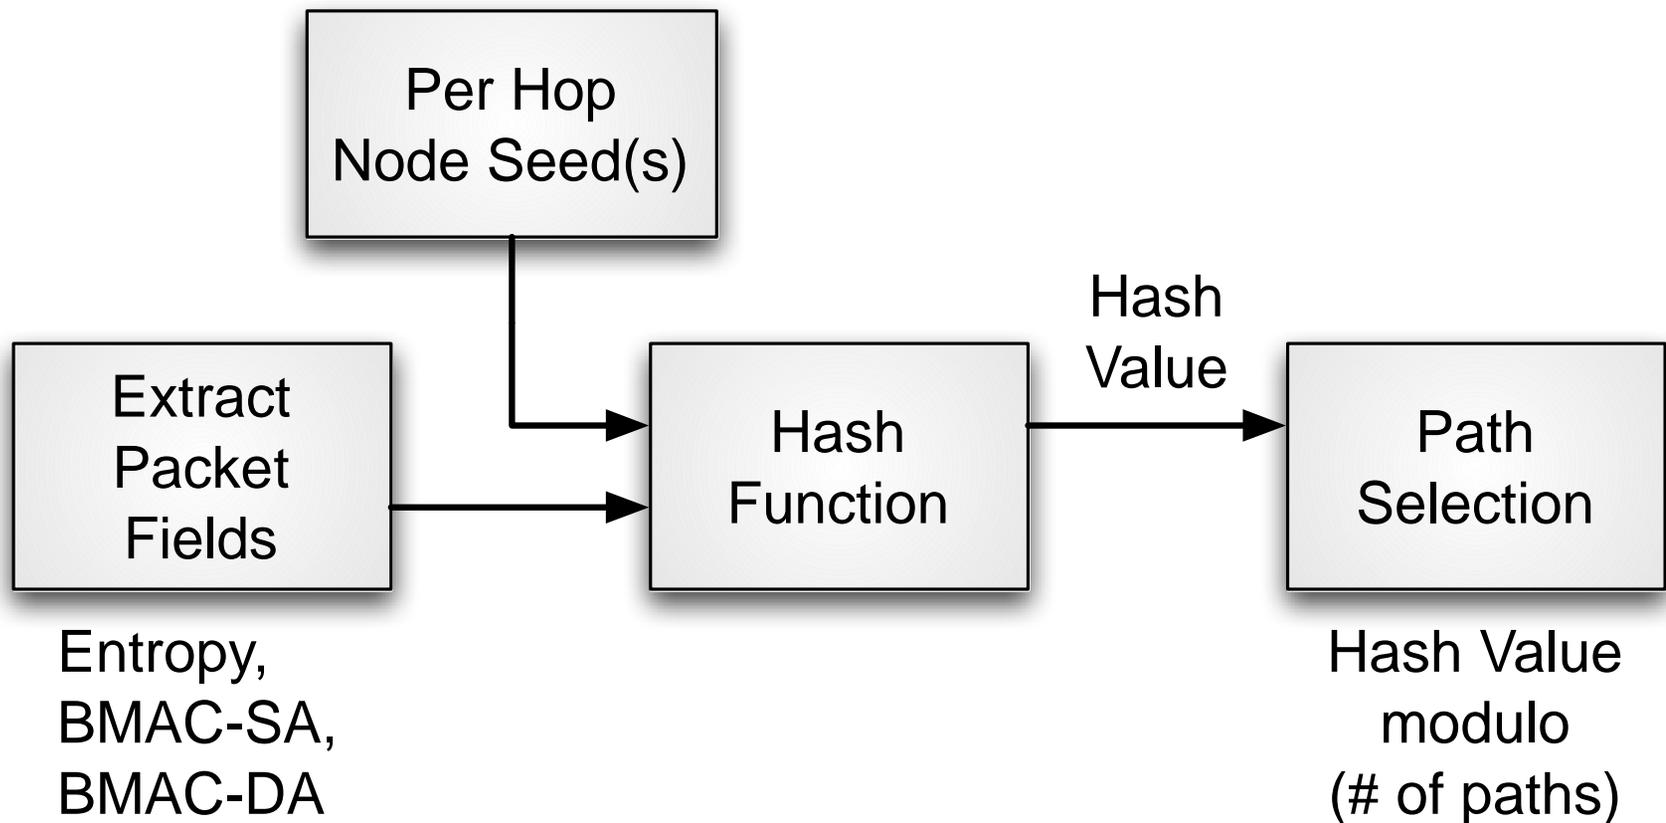
- Transmit flows from Edge source device (root node) and measure flow distribution across spine devices
- Use an N-ary tree with a degree of 4 and a depth of 3 hops

Measure

- Standard deviation of flows received at the spine devices



Path Selection Data Flow



Hash Input Fields



- Entropy (16-bit)
- Per-hop Node Seed
- BMAC SA
- BMAC DA

Agenda



- Overview
- Hash Functions Evaluated
- Review Results
- Observations
- New Functions Evaluated -1

Hash Functions Evaluated – 1



- **Baseline Hash Function**

- XOR of the following attributes:
 - 16-bit Entropy Value
 - 16-bit Node Seed (unique per hop)

Hash Functions Evaluated – 2



- **Baseline + BMAC SA + BMAC DA Hash Function**
 - XOR of the following attributes:
 - 16-bit Entropy Value
 - BMAC SA
 - BMAC DA
 - 16-bit Node Seed (unique per hop)

Hash Functions Evaluated – 3



- **CRC16-CCITT**

- CRC based on the following packet attributes:
 - 16-bit Entropy Value
 - BMAC SA
 - BMAC DA
 - 16-bit Node Seed (unique per hop)
- CRC Polynomial: $x^{16} + x^{12} + x^5 + 1$

Hash Functions Evaluated – 4



- **Baseline + Node Seed Shift**

- XOR of the following attributes:
 - 16-bit Entropy value shifted (circular) by the amount in Node Seed[3:0]
 - BMAC SA shifted (circular) by the amount in Node Seed[7:4]
 - BMAC DA shifted (circular) by the amount in Node Seed[11:8]
 - Node Seed[31:16]
- Node Seed is unique per hop

Agenda



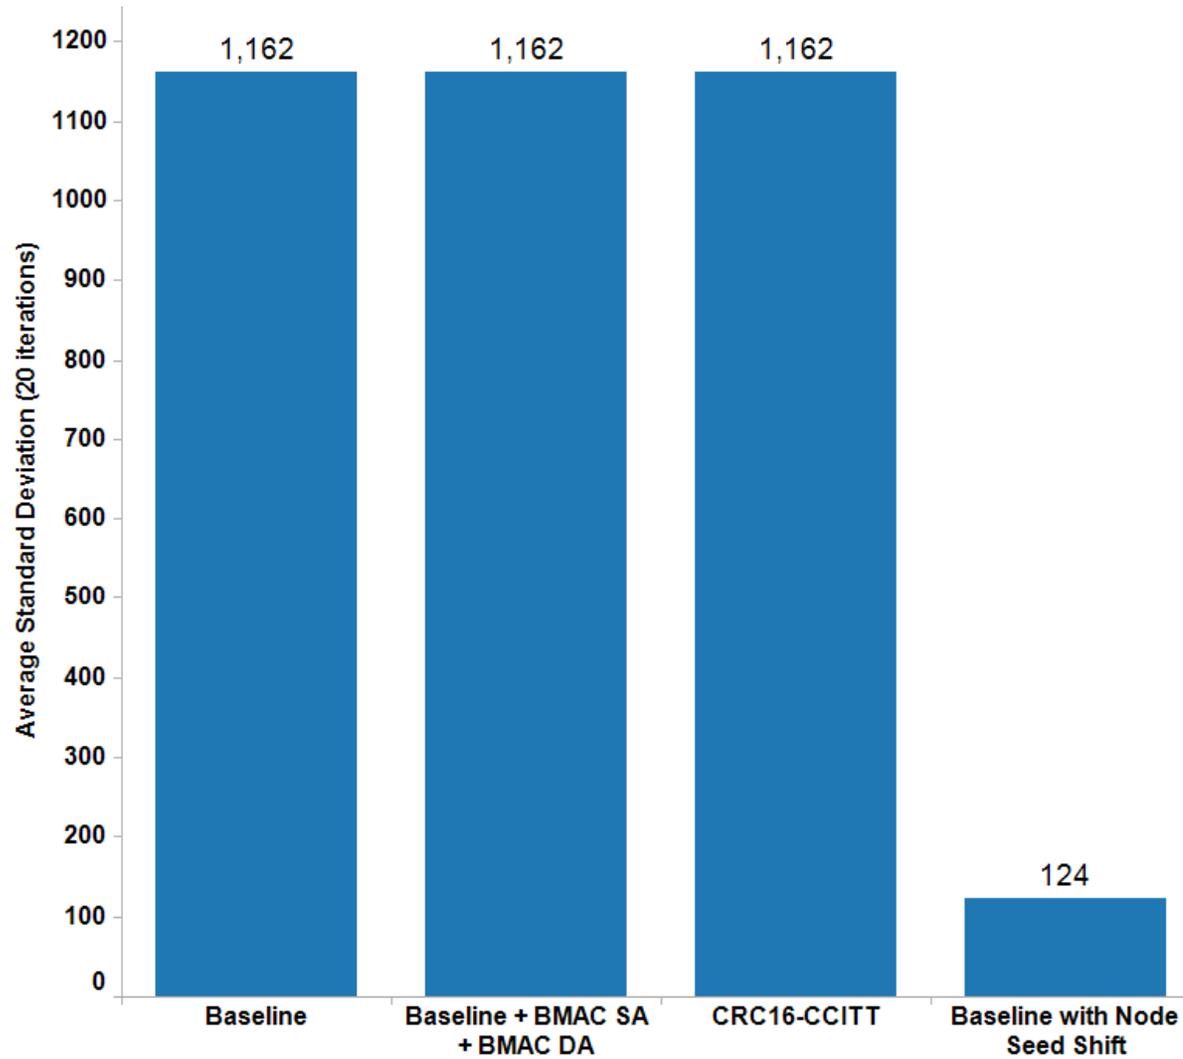
- Overview
- Hash Functions Evaluated
- Review Results
- Observations
- New Functions Evaluated -1

Test Setup



- **Topology: N-ary Tree**
 - Degree: 4
 - Depth: 3 hops
 - Spine devices: 64
- **Simulation Constraints**
 - 19,200 flows originating at edge source device
 - (300 flows) x (# of spine devices)
 - BMAC SA/DA limited to 64 unique values

Simulation Results



Agenda



- Overview
- Hash Functions Evaluated
- Review Results
- Observations
- New Functions Evaluated -1

Observations and Next Steps



- Observations

- XOR with circular shift based on a per-node seed provided good performance with low implementation cost

- Next Steps

- Look at other functions
 - FNV
 - Ideal

Agenda



- Overview
- Hash Functions Evaluated
- Review Results
- Observations
- New Functions Evaluated -1

New Hash Functions Evaluated



- Case A - Ideal
 - Random value generated for every flow at every node
 - Next Hop Selection: Random Value mod (# of Next Hops)

Hash Functions Evaluated

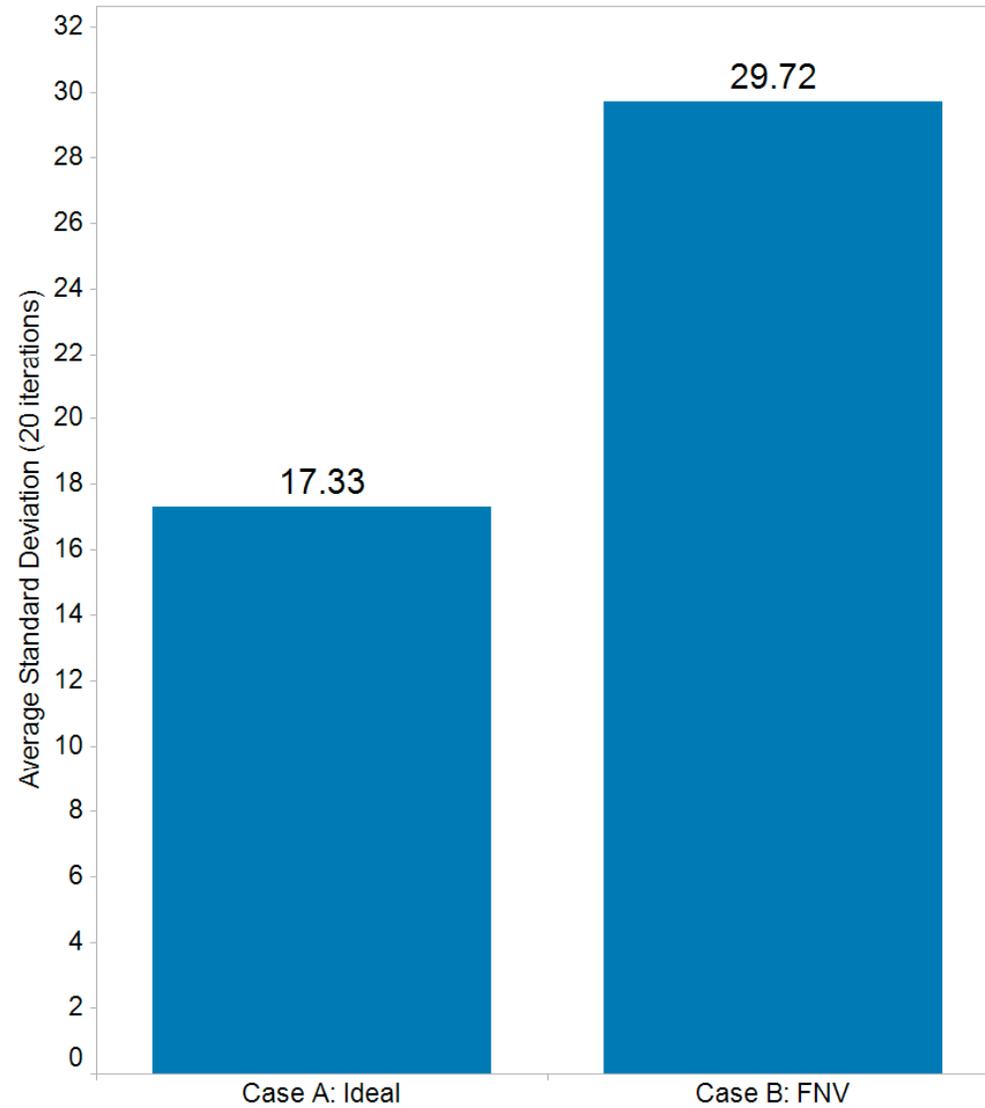


- Case B - FNV

- FNV-16

- FNV32 with 32-bit output folded using XOR of:
 - Hash Value[15:0]
 - Hash Value[31:16]
 - Offset-basis: 0x811c9dc5
 - Octets of Data:
 - Entropy (2 octets)
 - Node Seed (2 octets)
 - BMAC SA (2 octets)
 - BMAC DA (2 octets)

Simulation Results



New function Observations and Next Steps



- Observations
 - FNV shows very good behavior and is approaching very close to ideal

- Next Steps



Thank You