Query-Flood DoS Attacks In Gnutella

Neil Daswani and Hector Garcia-Molina
Stanford University
Department of Computer Science

Problem & Approach

- Problem
  - Gnutella: multiplicative query broadcast
  - **Application-layer** denial-of-service

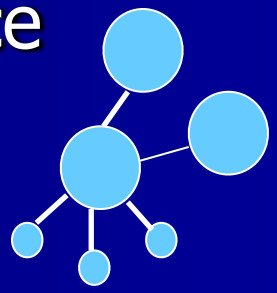

- Approach
  - Load balancing / provide fairness

How does Gnutella Work?

- Super-nodes
- Messages
  - **Ping / Pong**
  - **Query / QueryHit**
  - **Push**
- Already Seen
- Time To Live
- File X-fer: HTTP

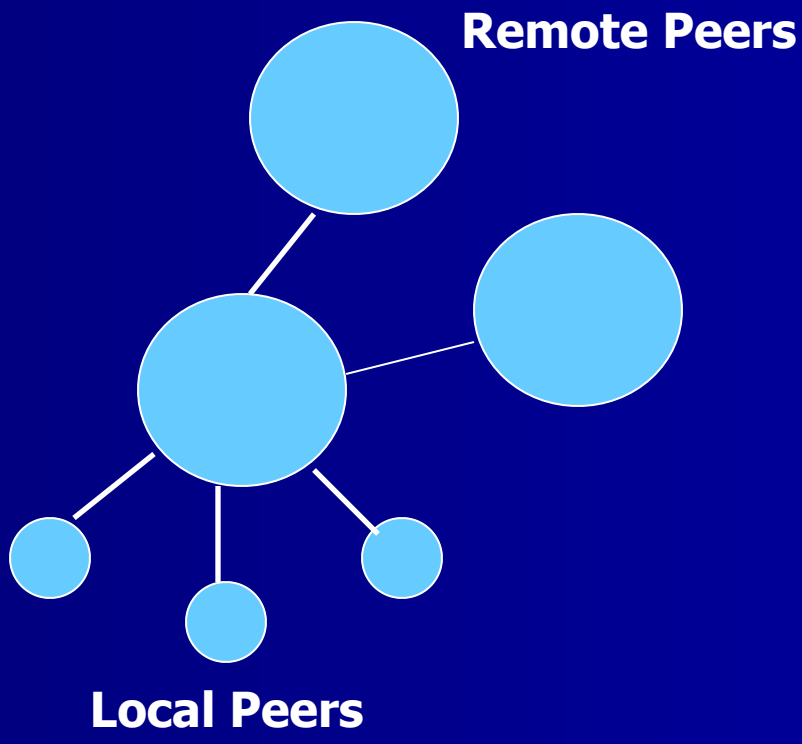

Questions

- Which queries to drop?
- Traffic management policies?
- Effect of topology?
- How is “damage” distributed?

=> Need Traffic Model & Metrics

Gnutella Traffic Model

- Discrete-event
- Only super-nodes explicitly modeled
- Only queries are modeled
- q=(origin,ttl)
- Max capacity:
  C = 6 queries / time unit

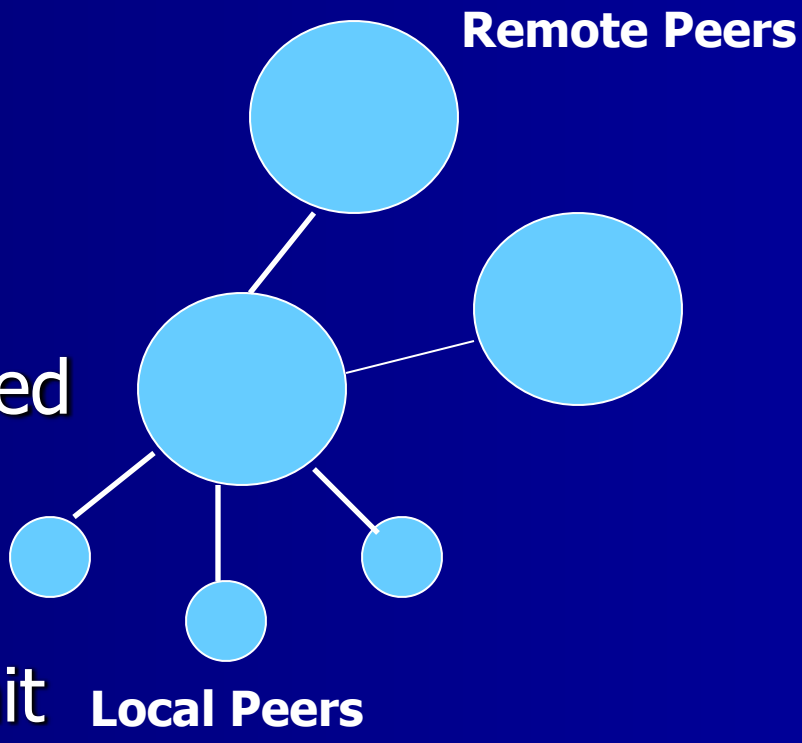

Gnutella Traffic Model

- Local Work ={q1,q2,q3}

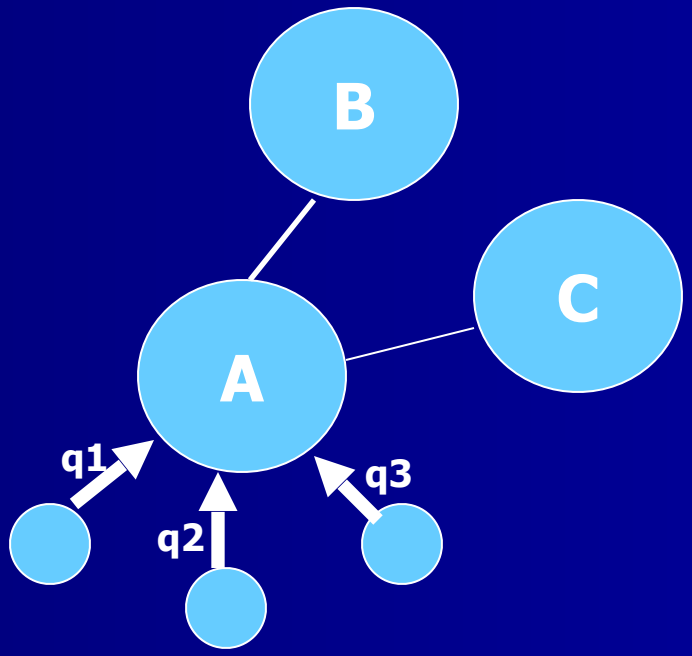

Gnutella Traffic Model

- Local Work ={q1,q2,q3}
- Remote Work = {q4,…,q9} ∪ {q10,q11}

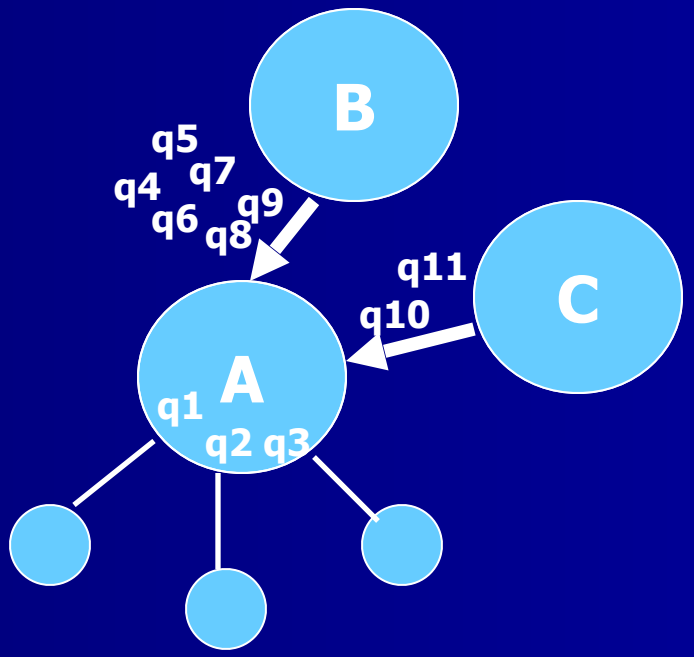

Gnutella Traffic Model

- Local vs. Remote Work:
  - Reservation Ratio (ρ)
- Remote Work:
  - How many? (IAS)
  - Which ones? (DS)

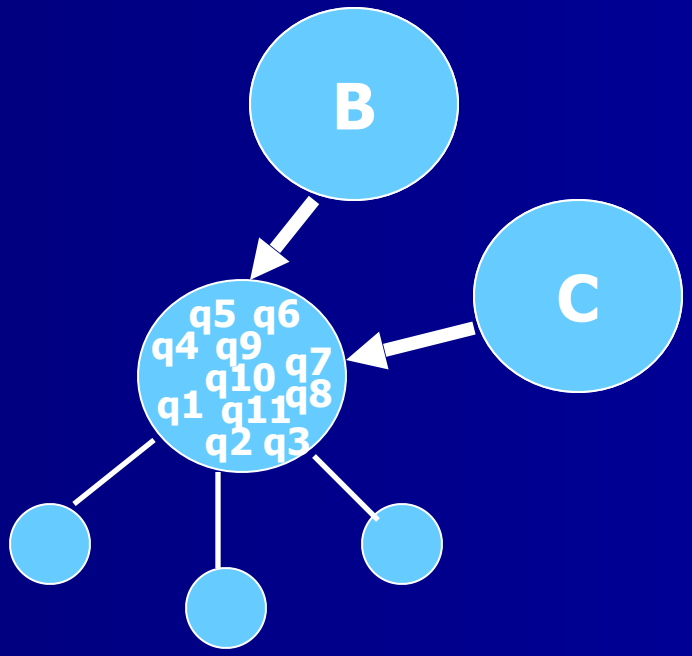

Gnutella Traffic Model

- Local Work ={q1,q2,q3}
- Remote Work = {q4,…,q9} ∪ {q10,q11}
- Local Work Accepted = {q1}
- Remote Work Accepted =
  - $I_{B,A}(1)$={q4,q5}
  - $I_{C,A}(1)$={q10,q11}

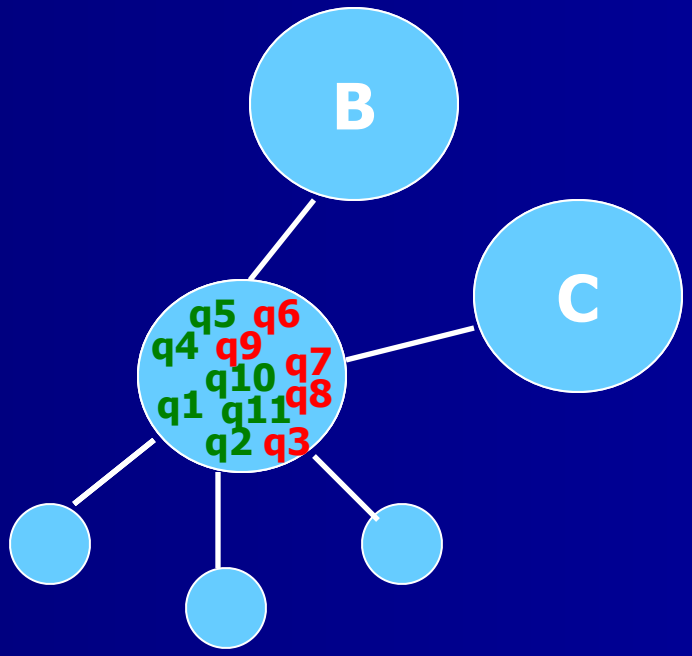

Gnutella Traffic Model

- Local Work ={q1,q2,q3}
- Remote Work = {q4,…,q9} ∪ {q10,q11}
- Local Work Accepted = {q1}
- Remote Work Accepted =
  - $I_{B,A}(1)$={q4,q5}
  - $I_{C,A}(1)$={q10,q11}
- Work Broadcasted = {q1,q4,q5,q10,q11}

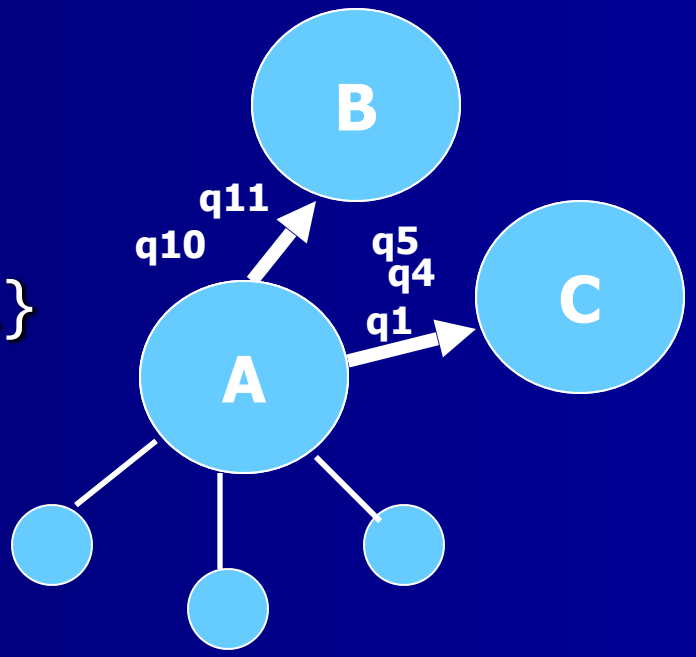

Reservation Ratio (ρ)

- Only used in high load situations.
- Max ρC queries from local peers.
- Max (1-ρ)C queries from remote peers.
- If ρ=1/3 and C=6,
  ρC=(1/3)(6)=2 Local

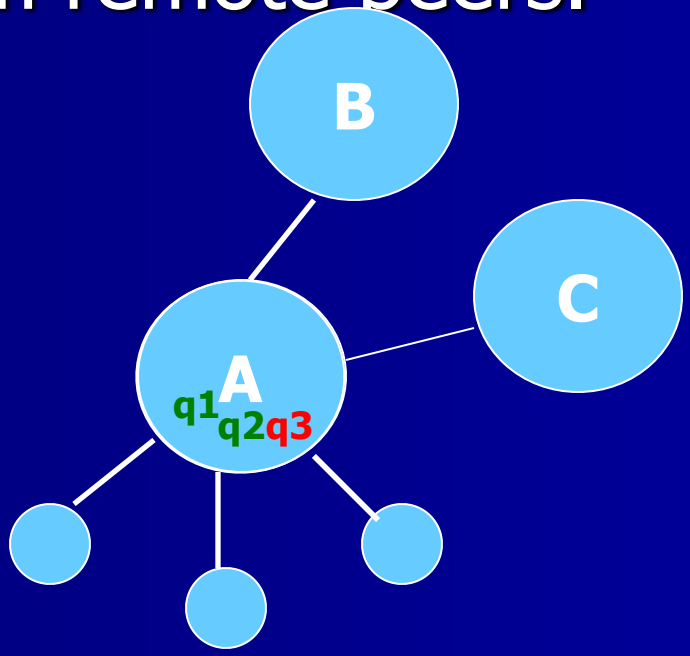

Incoming Alloc. Strategy

- $(1-\rho)C=(4/6)(6)=4$ Remote
- IAS Possibilities:
  - Fractional:
    2 from B
    2 from C
  - Weighted:
    3 from B
    1 from C

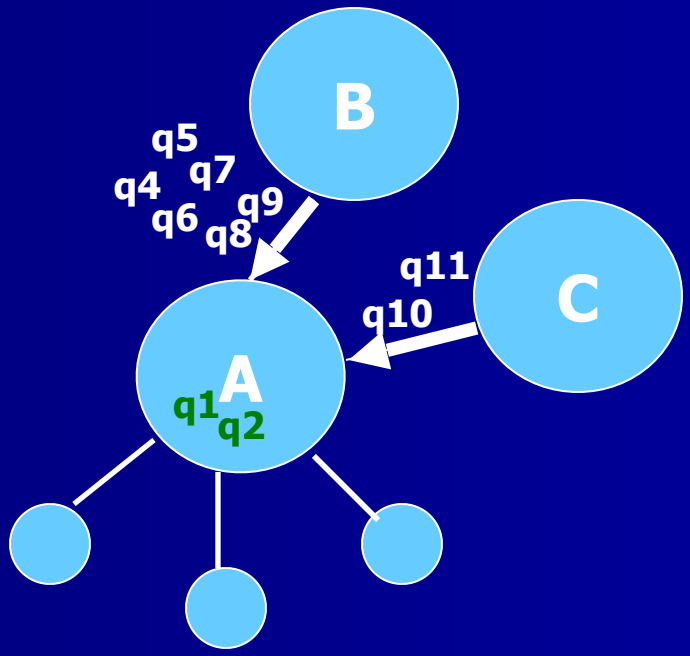

Drop Strategy

Which queries to drop?

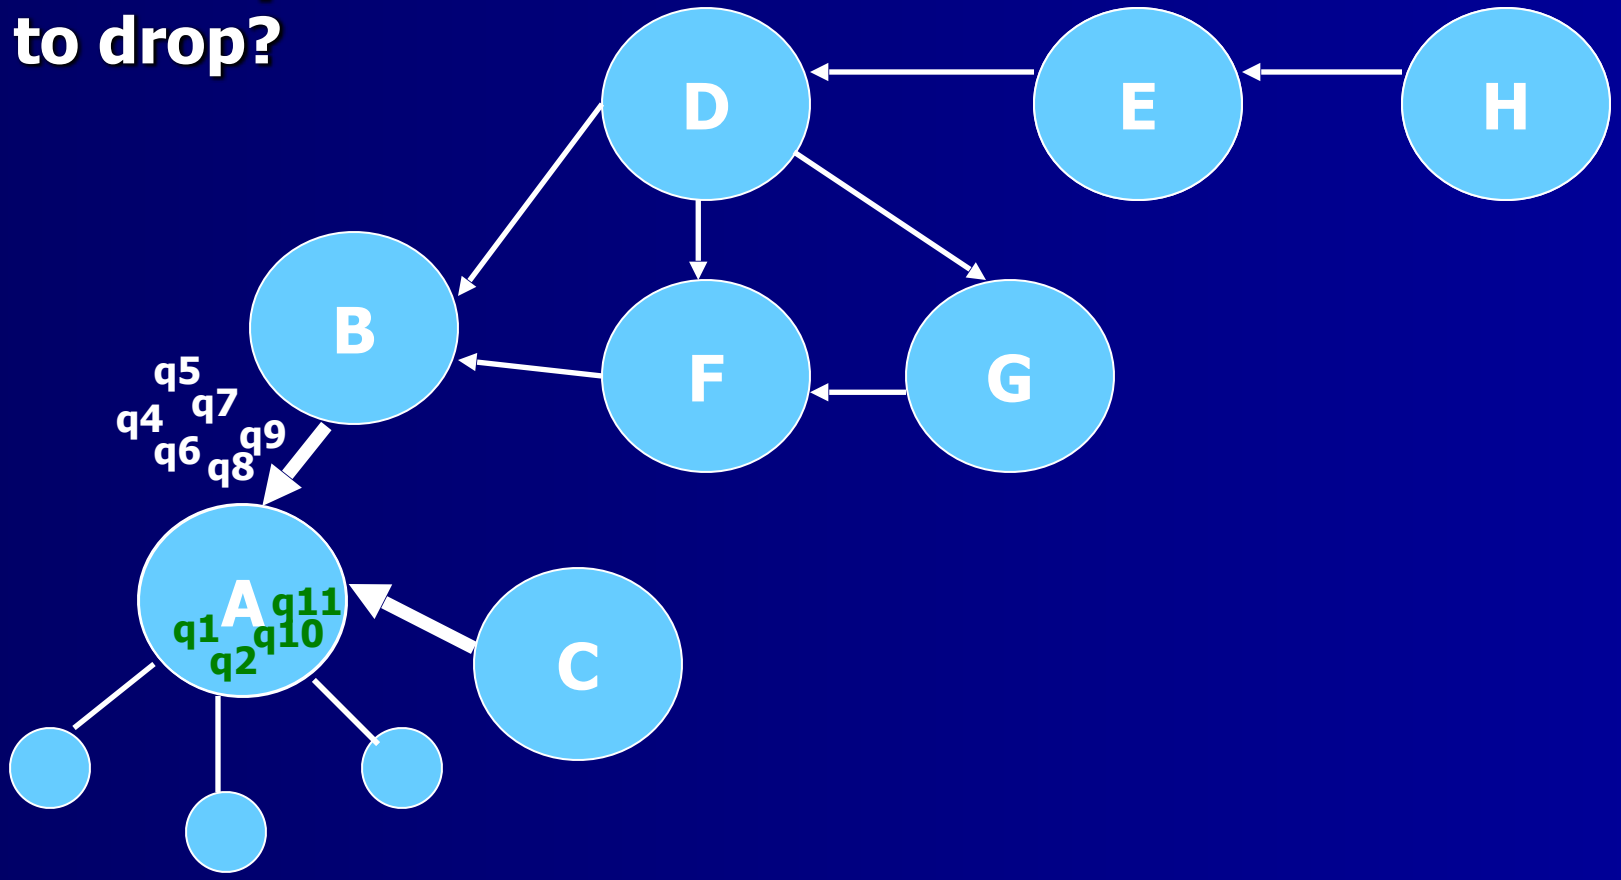

Drop Strategy

Which queries to drop?

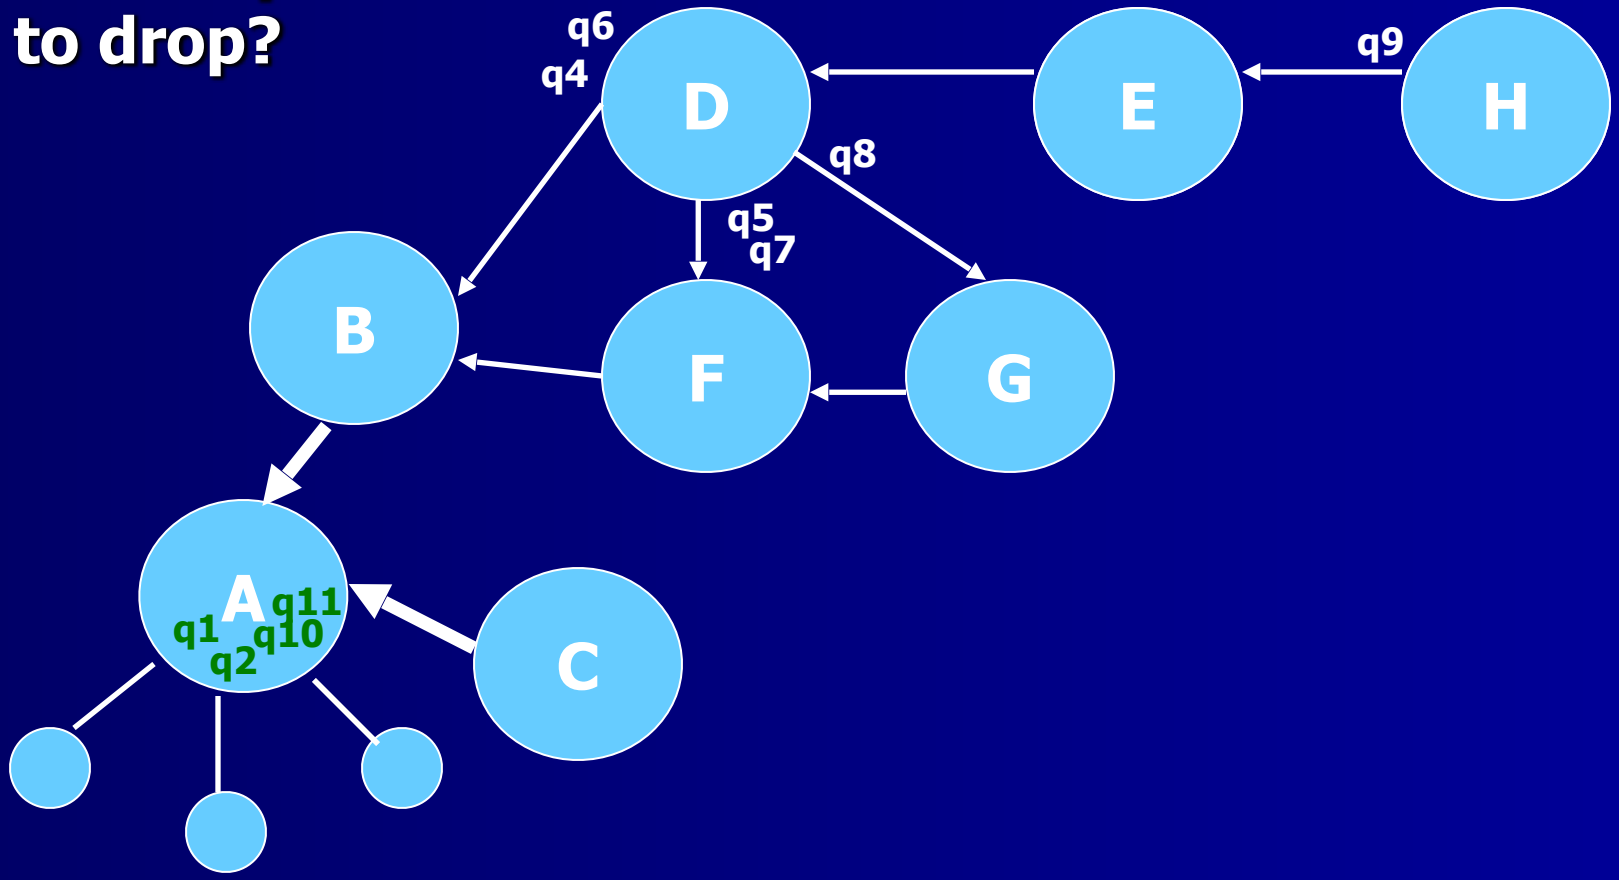

Drop Strategy

Equal

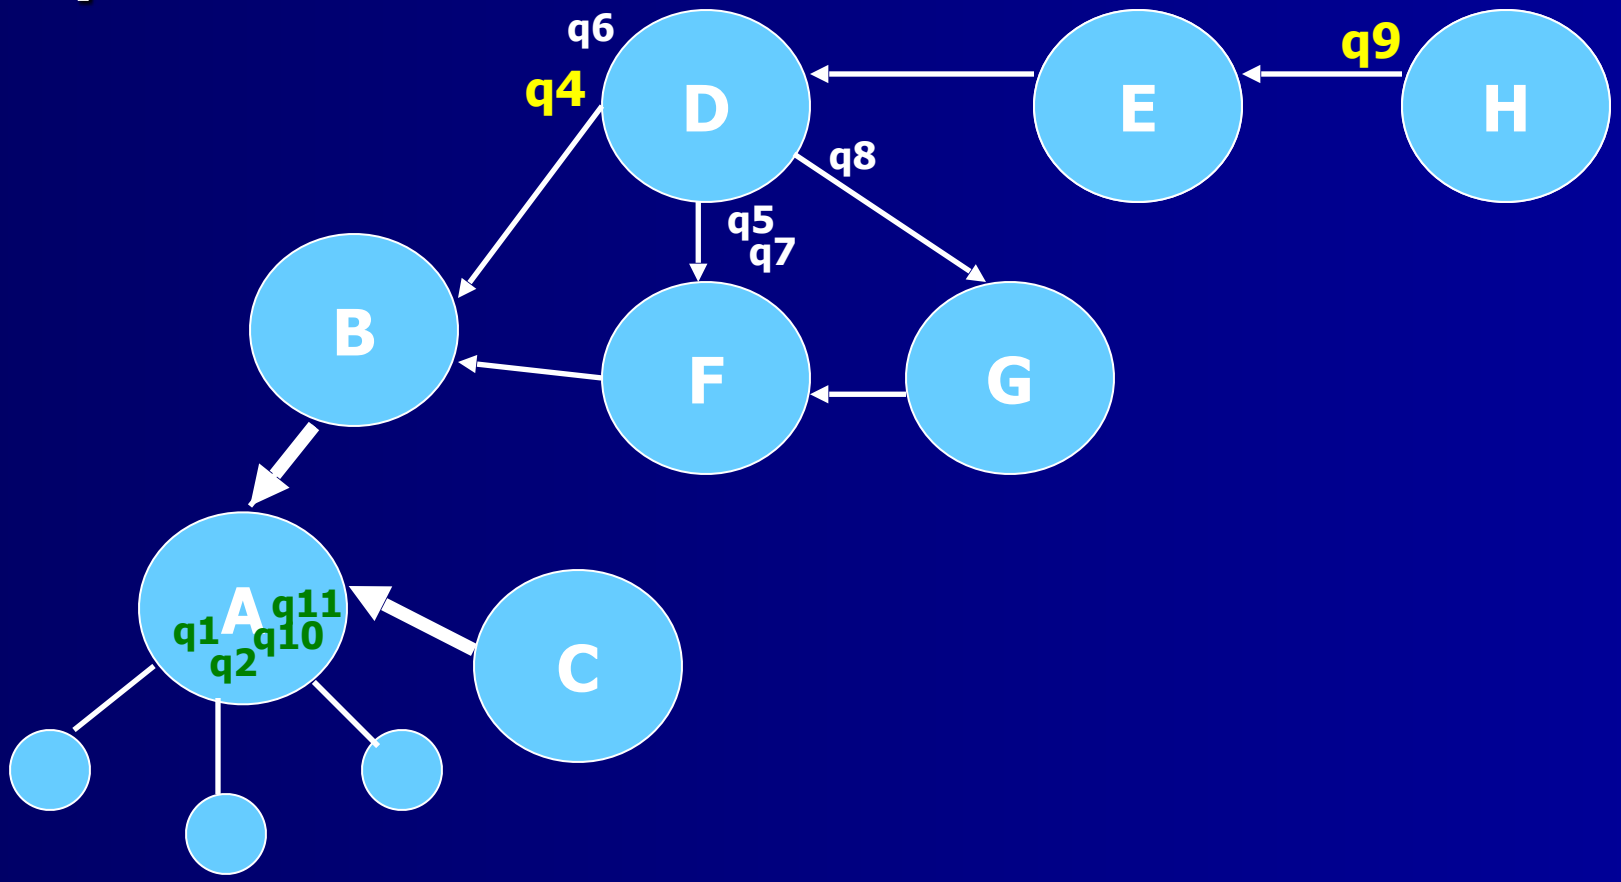

Drop Strategy

Proportional

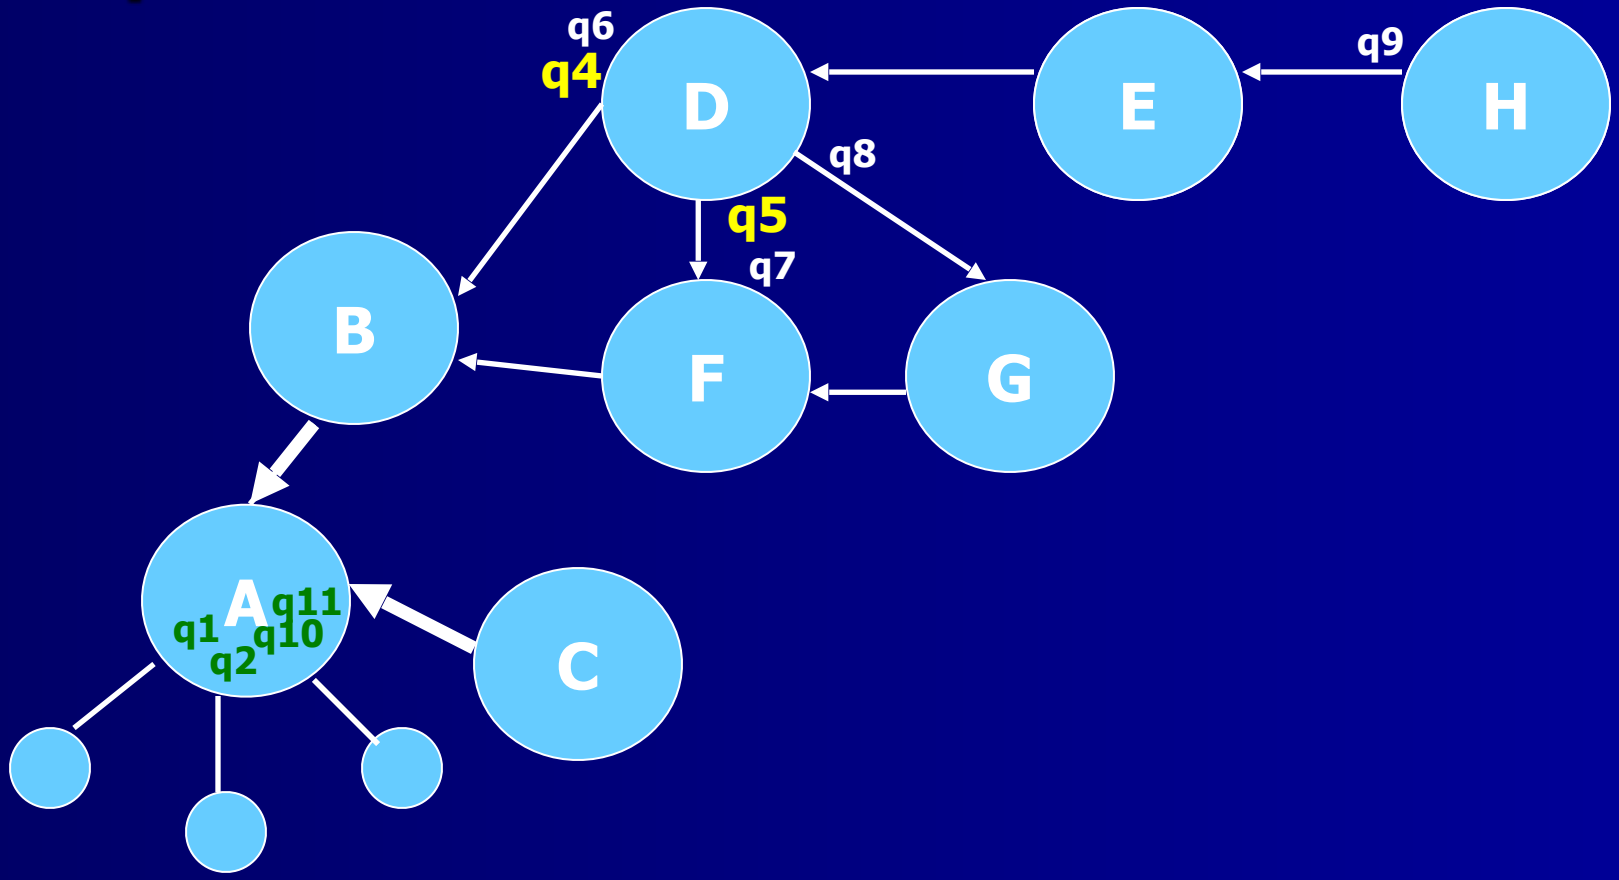

Drop Strategy

PreferHighTTL

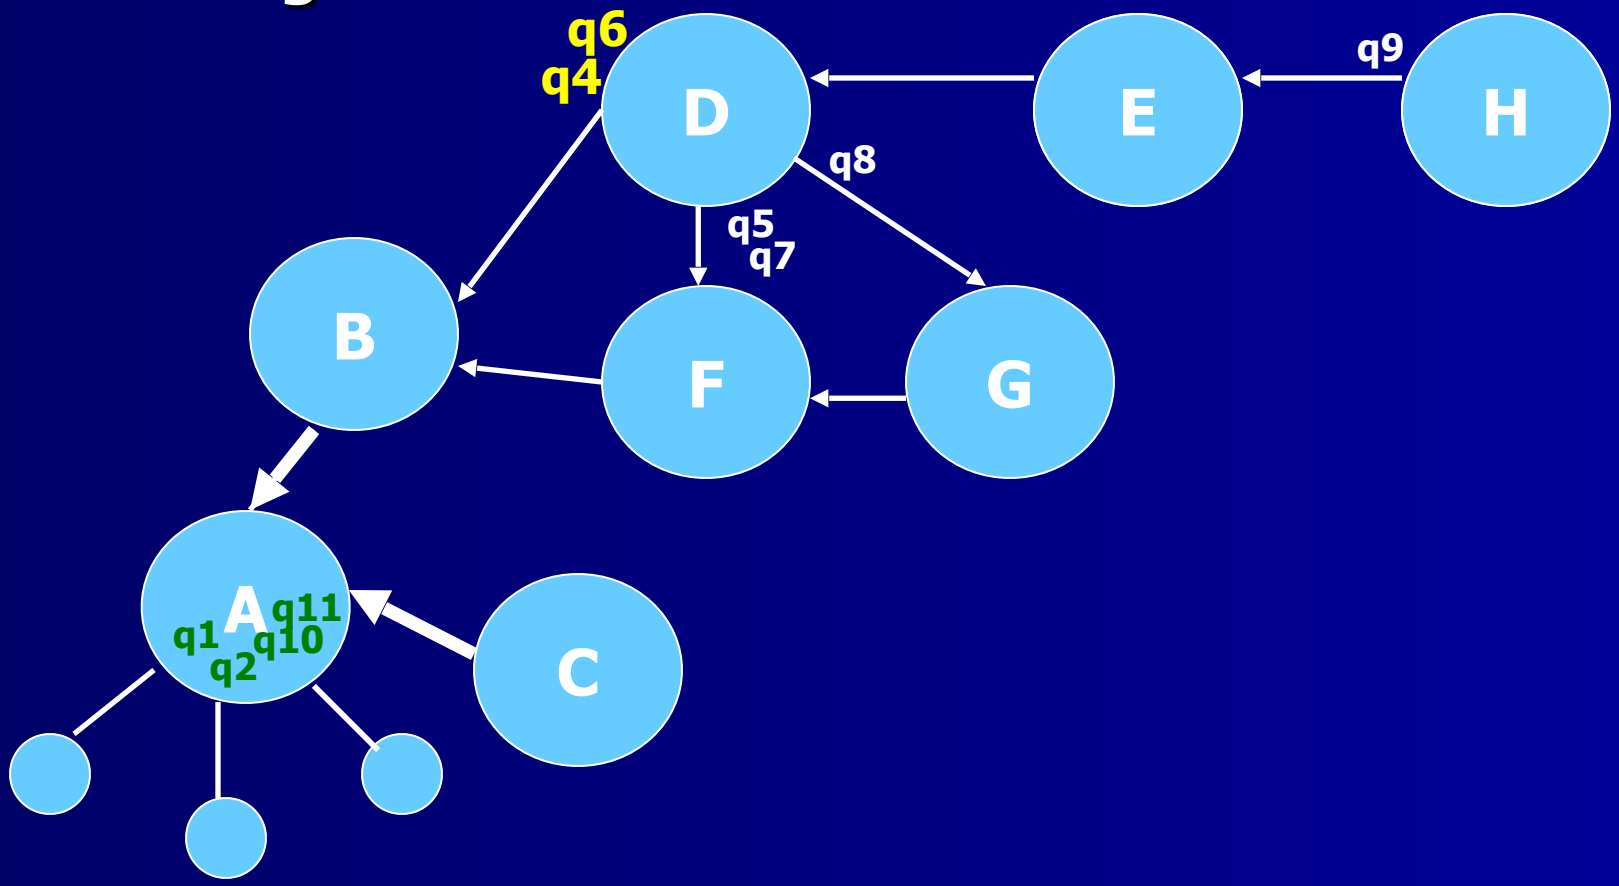

Drop Strategy

PreferLowTTL

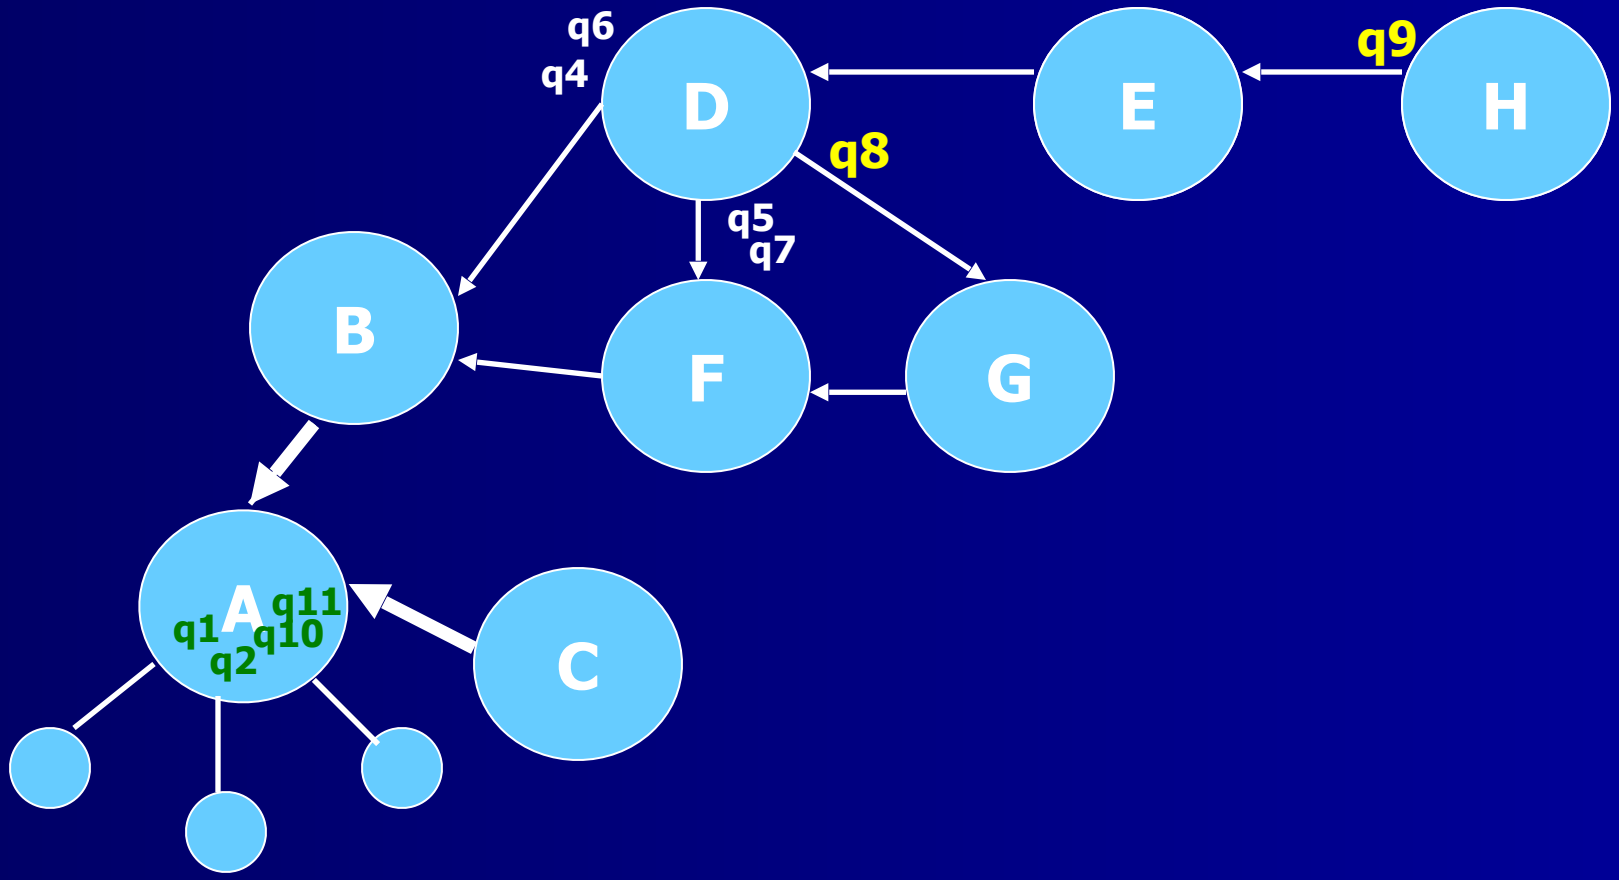

Good & Malicious Nodes

- Good nodes: $\rho = \widehat{\rho}$

K_3; τ=2

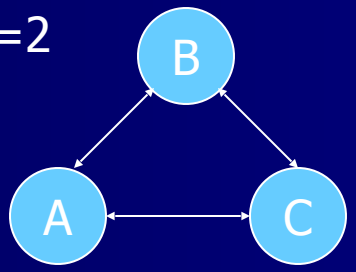

$\widehat{\rho} = 1/3$

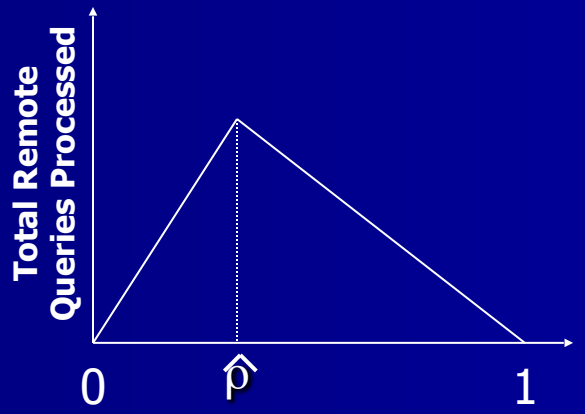

- In general, for symmetric networks:

  $\widehat{\rho} = 1 / (D(\tau) + 1)$

- Malicious nodes: $\rho_m = 1$

Damage

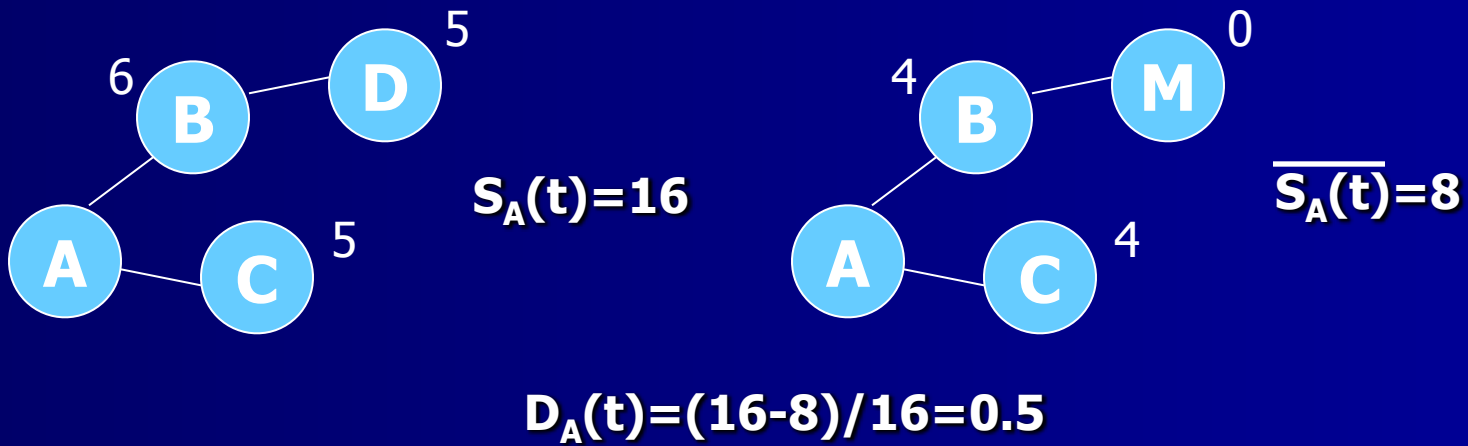

- Service Guarantee: $S_j(t)$, $\overline{S_j(t)}$
- Damage for node j (at time t):
  $D_j(t) = (S_j(t) - \overline{S_j(t)}) / S_j(t)$
- Cumulative Network Damage:
  D(t) = "bad" queries / "total" queries

Simulations

- Various Representative Topologies: K_{14}, C_{14}, G_{16}, L_{14}, P_{16}, S_{14}, W_{14}

- All IAS/DS described earlier
- Single malicious node / various placements
- Fundamental effects / trade-offs
- C=10,000; $\rho = \rho$; $\rho_m = 1$; τ=7; t=100

Results/Observations

	Fractional				Weighted			
Top(Loc)	Prop	Equal	PfHgTTL	PfLwTTL	Prop	Equal	PfHgTTL	PfLwTTL
Complete	0.143	0.143	0.143	0.143	0.545	0.545	0.545	0.545
Cycle	0.388	0.314	0.312	0.533	0.527	0.459	0.387	0.695
Grid (Ctr)	0.273	0.227	0.274	0.292	0.454	0.363	0.422	0.569
Grid (Co)	0.225	0.170	0.187	0.286	0.371	0.270	0.247	0.570
Grid (Ed)	0.282	0.191	0.208	0.378	0.412	0.306	0.294	0.553
Line (Ctr)	0.324	0.248	0.330	0.515	0.428	0.306	0.398	0.609
Line (Ed)	0.175	0.148	0.143	0.275	0.219	0.184	0.165	0.346
Pwr (H)	0.272	0.262	0.284	0.324	0.539	0.505	0.484	0.612
Pwr (L)	0.201	0.169	0.193	0.267	0.443	0.367	0.386	0.534
Star (Ce)	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
Star (Ed)	0.142	0.143	0.142	0.143	0.526	0.506	0.542	0.545
Whl (Ce)	0.386	0.386	0.386	0.386	0.726	0.751	0.717	0.751
Whl (Ed)	0.335	0.337	0.354	0.388	0.505	0.444	0.510	0.573

Results/Observations

- IAS/DS vs. Damage
  - Which IAS/DS minimizes damage?
  - Depends upon topology?
- Topology vs. Damage
  - Some topologies better than others?
  - Some nodes particularly vulnerable to attack?
- Damage Distribution
  - How is damage distributed?
  - Flood vs. Structural damage

IAS/DS vs. Damage

- O1: Fractional IAS +
  Equal or PreferHighTTL DS optimal

	Fractional				Weighted			
Top(Loc)	Prop	**Equal**	**PfHgTTL**	PfLwTTL	Prop	Equal	PfHgTTL	PfLwTTL
Complete	0.143	**0.143**	**0.143**	0.143	0.545	0.545	0.545	0.545
Cycle	0.388	0.314	**0.312**	0.533	0.527	0.459	0.387	0.695
Grid (Ctr)	0.273	**0.227**	0.274	0.292	0.454	0.363	0.422	0.569
Grid (Co)	0.225	**0.170**	0.187	0.286	0.371	0.270	0.247	0.570
Grid (Ed)	0.282	**0.191**	0.208	0.378	0.412	0.306	0.294	0.553
Line (Ctr)	0.324	**0.248**	0.330	0.515	0.428	0.306	0.398	0.609
Line (Ed)	0.175	0.148	0.143	0.275	0.219	0.184	0.165	0.346
Pwr (H)	0.272	**0.262**	0.284	0.324	0.539	0.505	0.484	0.612
Pwr (L)	0.201	**0.169**	0.193	0.267	0.443	0.367	0.386	0.534
Star (Ce)	1.000	**1.000**	**1.000**	1.000	1.000	1.000	1.000	1.000
Star (Ed)	0.142	0.143	**0.142**	0.143	0.526	0.506	0.542	0.545
Whl (Ce)	0.386	**0.386**	**0.386**	0.386	0.726	0.751	0.717	0.751
Whl (Ed)	0.335	**0.337**	0.354	0.388	0.505	0.444	0.510	0.573

IAS/DS vs. Damage

- O2: Weighted/Prop always worse than Fractional/Equal by about 2x or more

Topology (Location)	Fractional/ Equal	Weighted/ Proportional	Damage Reduction
Complete	0.143	0.545	3.8
Cycle	0.314	0.527	1.7
Grid (C)	0.227	0.454	2.0
Line (C)	0.248	0.428	1.7
Power (H)	0.262	0.539	2.1
Wheel (C)	0.386	0.726	1.9

IAS/DS vs. Damage

- O3: PreferLowTTL incurs (at least as much or) more damage than other DSs

	Fractional				Weighted			
Top(Loc)	Prop	Equal	PfHgTTL	**PfLwTTL**	Prop	Equal	PfHgTTL	**PfLwTTL**
Complete	0.143	0.143	0.143	**0.143**	0.545	0.545	0.545	**0.545**
Cycle	0.388	0.314	0.312	**0.533**	0.527	0.459	0.387	**0.695**
Grid (Ctr)	0.273	0.227	0.274	**0.292**	0.454	0.363	0.422	**0.569**
Grid (Co)	0.225	**0.170**	0.187	**0.286**	0.371	0.270	0.247	**0.570**
Grid (Ed)	0.282	**0.191**	0.208	**0.378**	0.412	0.306	0.294	**0.553**
Line (Ctr)	0.324	0.248	0.330	**0.515**	0.428	0.306	0.398	**0.609**
Line (Ed)	0.175	0.148	0.143	**0.275**	0.219	0.184	0.165	**0.346**
Pwr (H)	0.272	0.262	0.284	**0.324**	0.539	0.505	0.484	**0.612**
Pwr (L)	0.201	0.169	0.193	**0.267**	0.443	0.367	0.386	**0.534**
Star (Ed)	0.142	0.143	0.142	**0.143**	0.526	0.506	0.542	**0.545**
Whl (Ce)	0.386	0.386	0.386	**0.386**	0.726	0.751	0.717	**0.751**
Whl (Ed)	0.335	0.337	0.354	**0.388**	0.505	0.444	0.510	**0.573**

Results/Observations

- IAS/DS vs. Damage
  - Which IAS/DS minimizes damage?
  - Depends upon topology?
- Topology vs. Damage
  - Some topologies better than others?
  - Some nodes particularly vulnerable to attack?
- Damage Distribution
  - How is damage distributed?
  - Flood vs. Structural damage

Topology vs. Damage

- O4: Complete topology (K) under Frac/Eq IAS/DS least prone to damage & insensitive to malicious node position.

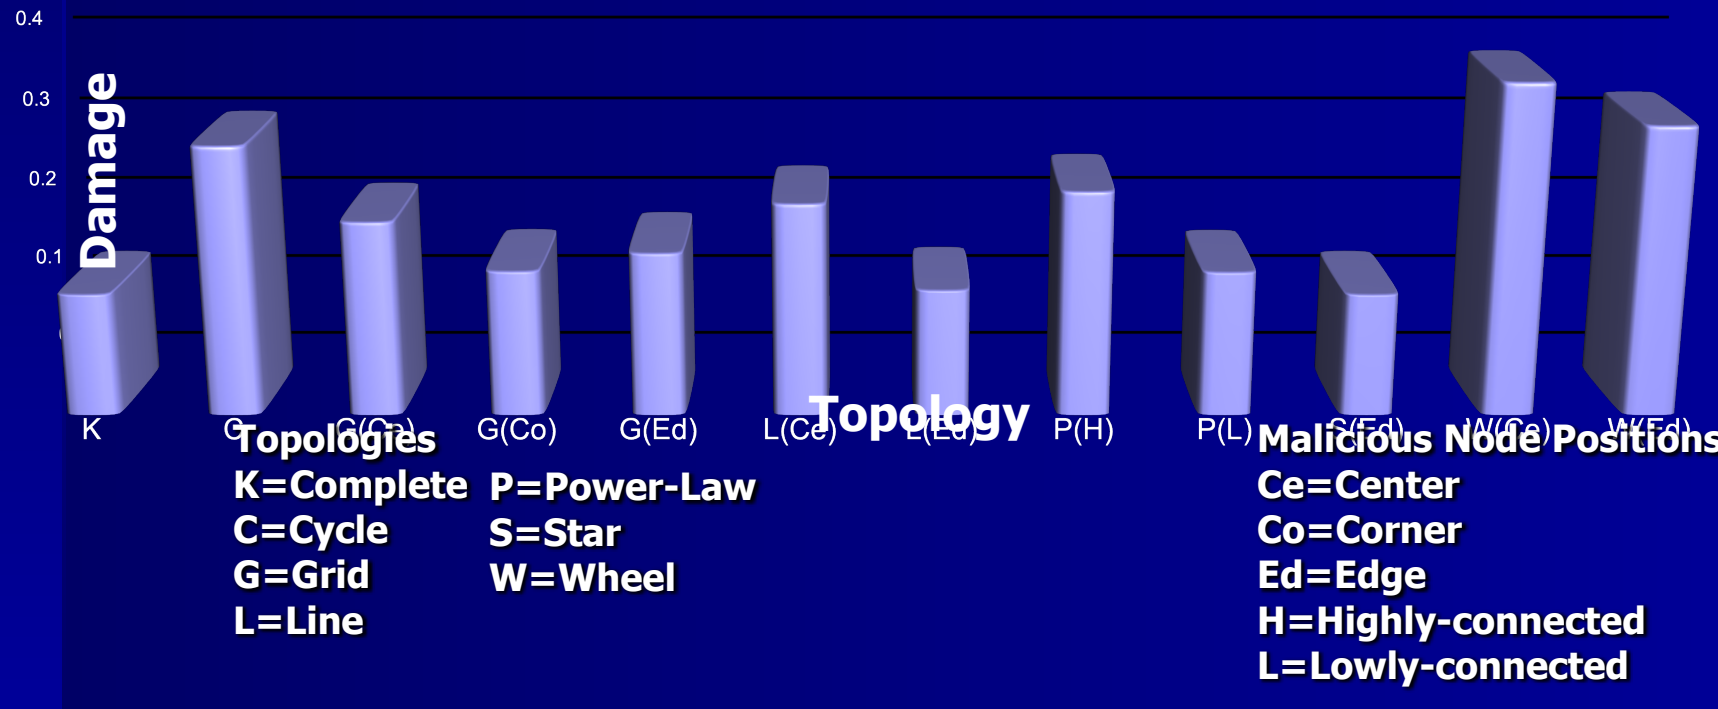

Topology vs. Damage

- O5: Good topology is not enough. Must use good policies too.

	Fractional				Weighted			
Top(Loc)	Prop	**Equal**	PfHgTTL	PfLwTTL	Prop	Equal	PfHgTTL	**PfLwTTL**
Complete	0.143	**0.143**	0.143	0.143	0.545	0.545	0.545	**0.545**
Cycle	0.388	**0.314**	0.312	0.533	0.527	0.459	0.387	**0.695**
Grid (Ctr)	0.273	**0.227**	0.274	0.292	0.454	0.363	0.422	**0.569**
Line (Ctr)	0.324	**0.248**	0.330	0.515	0.428	0.306	0.398	**0.609**
Pwr (H)	0.272	**0.262**	0.284	0.324	0.539	0.505	0.484	**0.612**
Whl (Ce)	0.386	**0.386**	0.386	0.386	0.726	0.751	0.717	**0.751**

Results/Observations

- IAS/DS vs. Damage
  - Which IAS/DS minimizes damage?
  - Depends upon topology?
- Topology vs. Damage
  - Some topologies better than others?
  - Some nodes particularly vulnerable to attack?
- Damage Distribution
  - How is damage distributed?
  - Flood vs. Structural damage

Damage Distribution (Cycle)

- O6: Nodes should distance themselves from untrusted nodes.

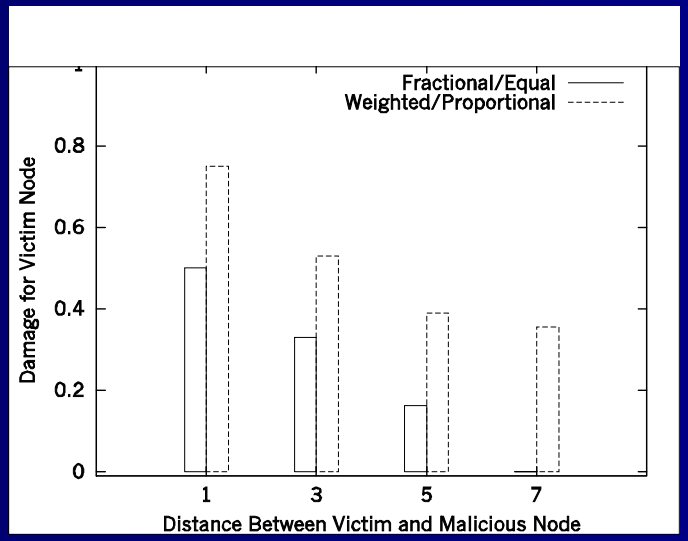

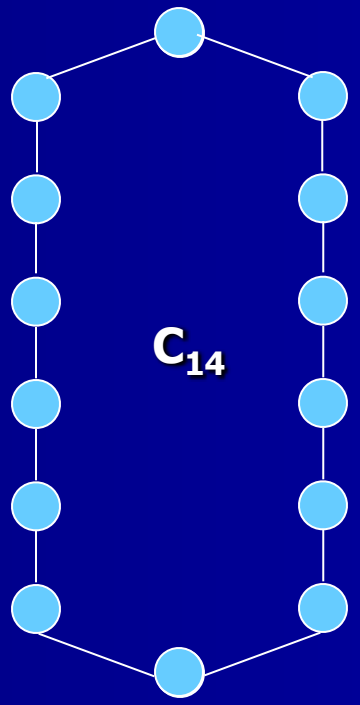

- Damage decreases as distance from malicious node increases.

Damage Distribution (Cycle)

- O7: Disconnect protocols must be used to prevent "structural" damage.

Weighted/Proportional

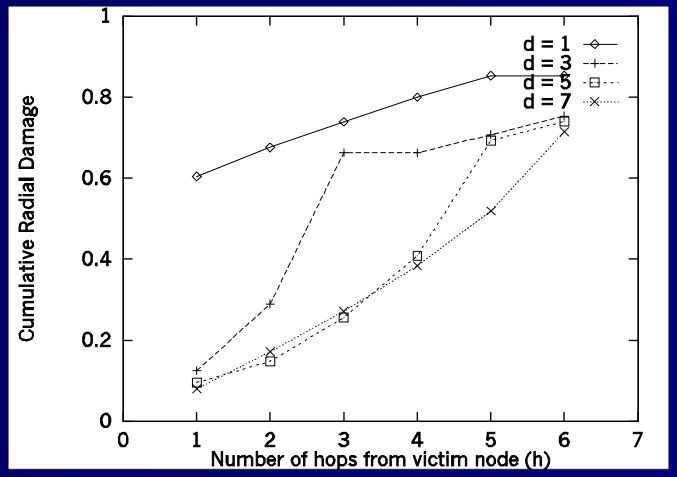

Fractional/Equal

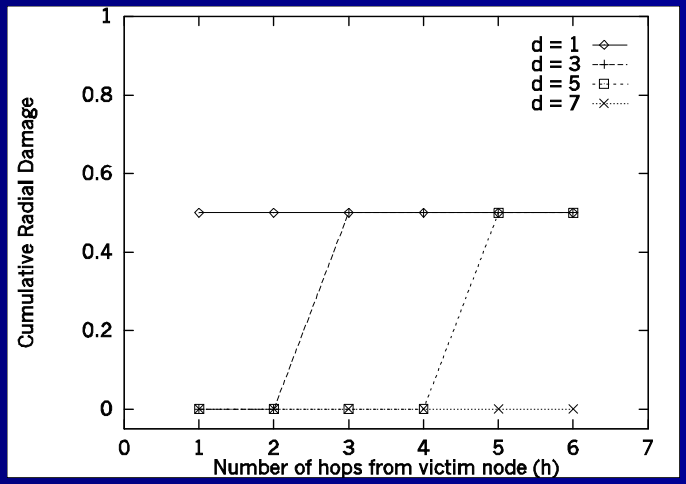

- Fractional/Equal IAS/DS minimizes "flood" damage in cycle topology.

Conclusion

- Defined model & metrics; Evaluation
- 7 observations:
  1. Fractional IAS + Equal or PreferHighTTL DS optimal
  2. Weighted IAS always worse than Fractional IAS by ~ 2x
  3. PreferLowTTL incurs more damage than other DSs (or at least as much)
  4. Complete topology (K) under Frac/Eq IAS/DS least prone to damage & insensitive to malicious node position.
  5. Good topology is not enough. Must use good policy too.
  6. Nodes should distance themselves from untrusted nodes.
  7. Disconnect protocols must be used to prevent "structural" damage.

Q & A

- Paper & slides available at: http://www.stanford.edu/~daswani