Routing Indices For Peer-to-Peer Systems

A Critical Review

By Alireza Mortezaei

WATERLOO ENGINEERING

engineering.uwaterloo.ca

Outline

- Summary
- Technical Content
  - Strengths
  - Weaknesses
- Thoughts and Recommendations
- Final Evaluation

Background

- P2P systems facilitates content sharing
- No global knowledge over the network
- Routing is a major open problem in P2P networks

Taxonomy of Solutions

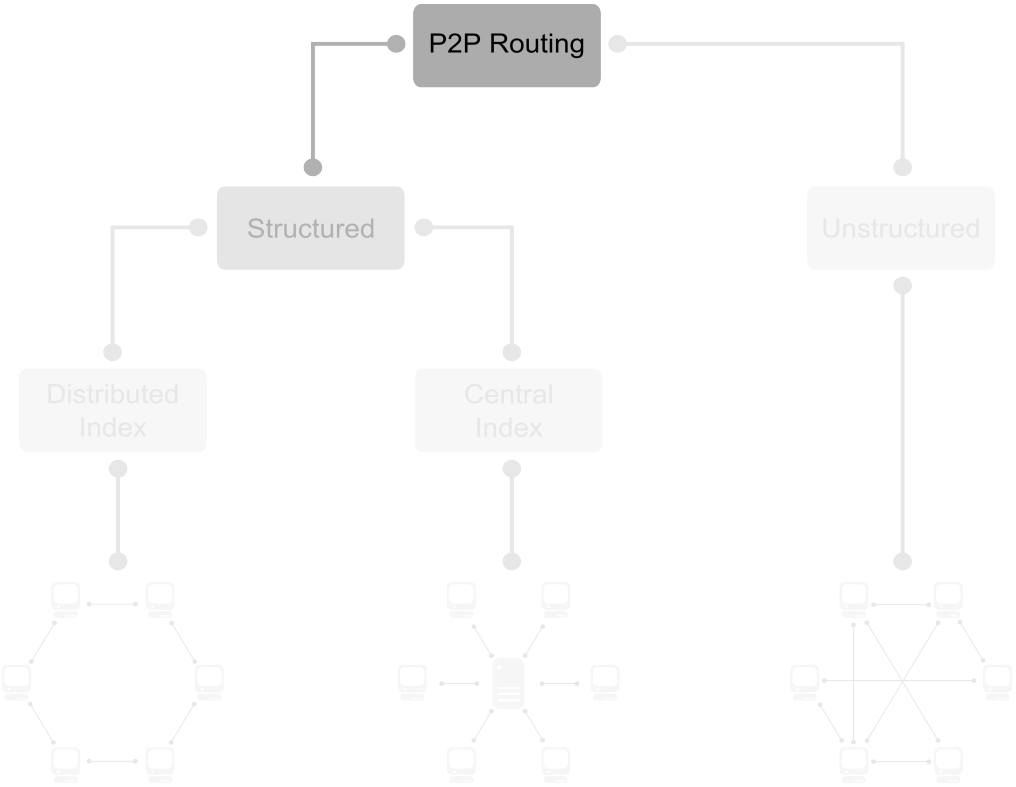

WATERLOO ENGINEERING

The Nature of Query

- Conjunction of keywords
- Forwarded until satisfied
- **Stop condition**: certain number of documents
- **Keywords:** subject topics e.g.

> "*Find me 100 documents on **database** and **networks***"

General Idea

- Direction versus location
- Use hints to guide the query through the network
- Basis of routing:

> "**Choose the neighbor that is more likely to yield better result**"

General Idea Cont'

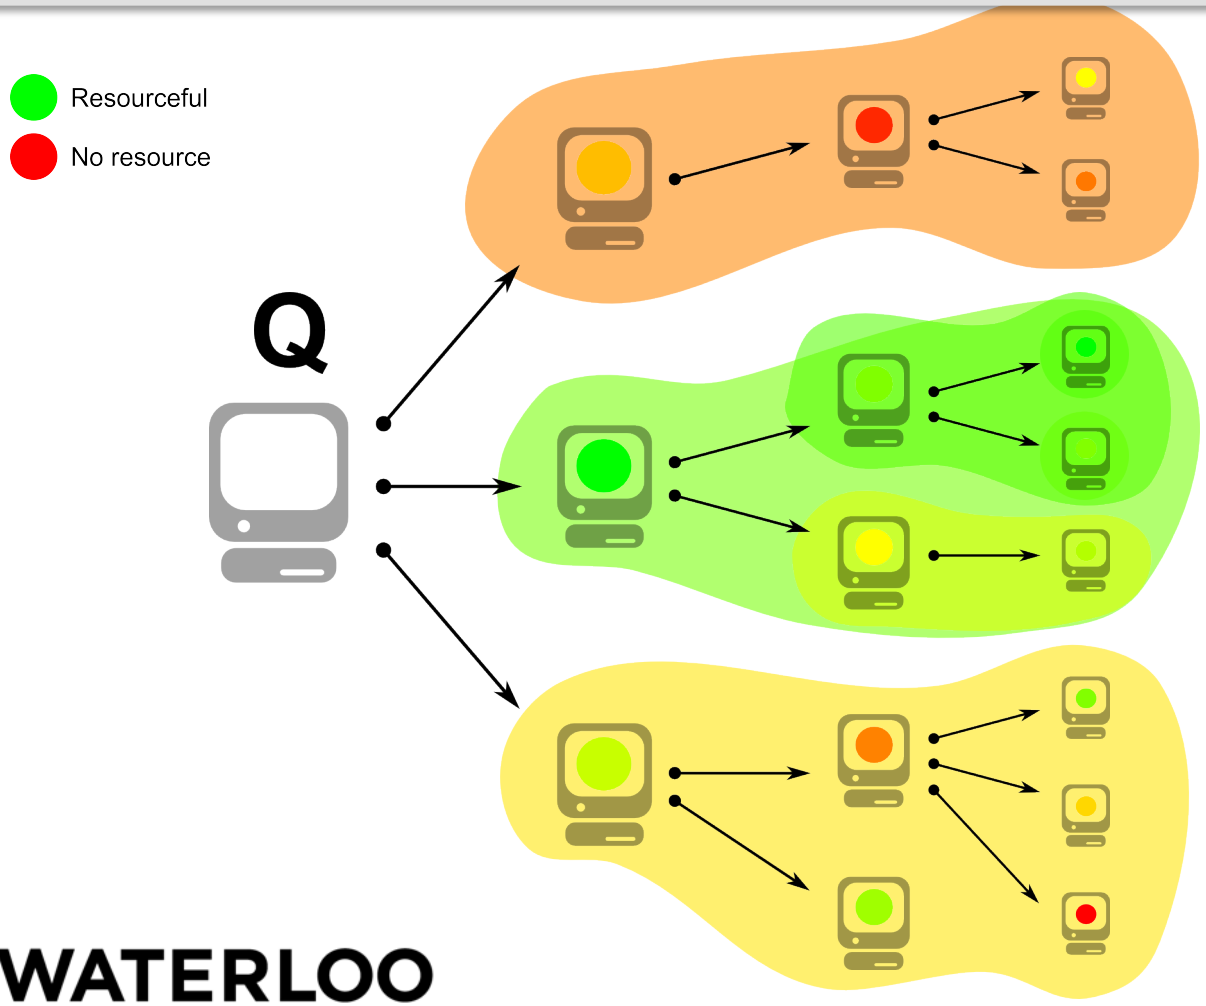

WATERLOO
ENGINEERING

Basic Model

- Compound Routing Indices (CRI):
  - Local indices
  - Routing indices
- Route table functionality
- Summarize RIs by coarsely categorizing topics:
  - Compresses RI table
  - Information loss due to summarization

Basic Model Cont'

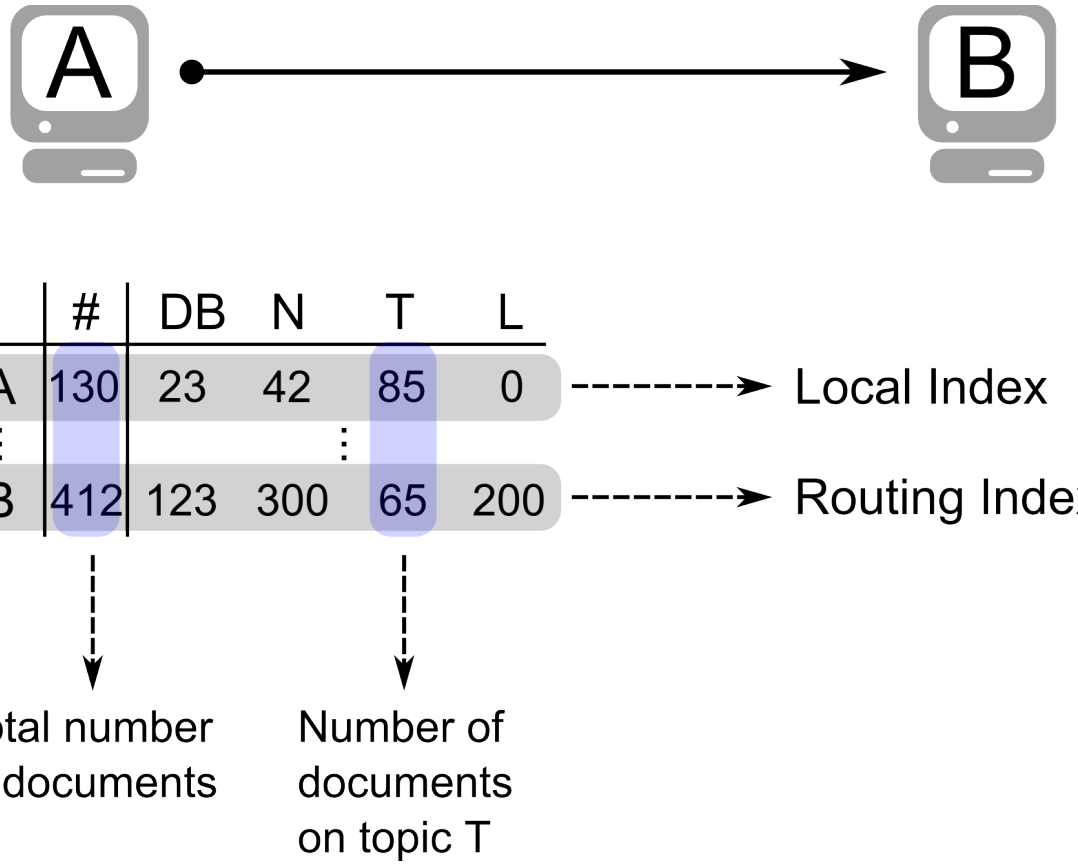

Basic Model Cont’

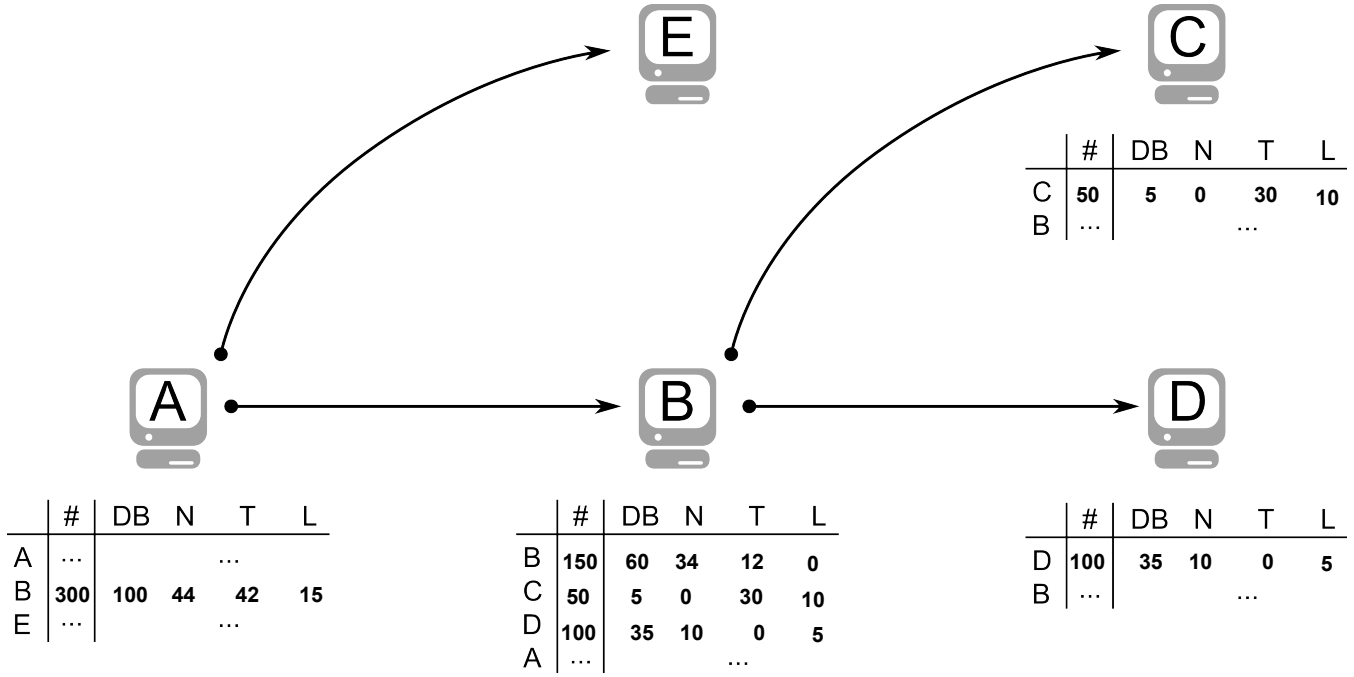

Estimation Model

- Estimate the number of documents along the path as a measure of “goodness”
- Prioritize neighbors based on their estimated goodness
- GlOSS estimator:

$$NumOfDocs \times \prod_{i} \frac{CRI(s_i)}{NumOfDocs}$$

Estimation Model Cont'

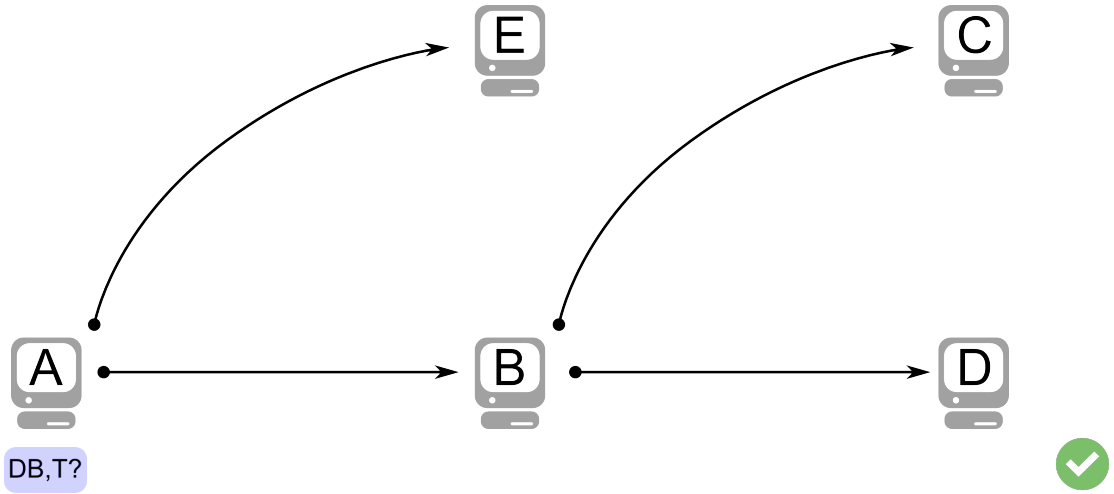

	#	DB	N	T	L
A	…			…	
B	500	500	04	20	16
D	200	56	00	08	520

$$G(b) = 500 \times \frac{500}{500} \times \frac{30}{50} \times \frac{24}{300} = 8$$

$$G(d) = 200 \times \frac{56}{200} \times \frac{8}{200} = 0$$

WATERLOO ENGINEERING

Creation and Maintenance

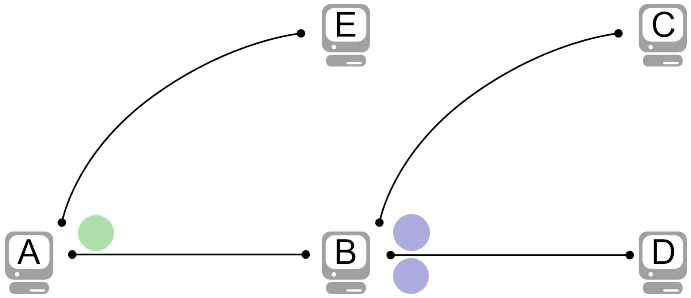

A's table

	#	DB	N	T	L
A	400	240	120	5	45
E	200	50	0	28	120
B	300	100	44	42	15

B's table

	#	DB	N	T	L
B	150	60	34	12	0
C	50	5	0	30	10
D	100	35	10	0	5
A	600	290	120	33	165

700 340 164 47 165 Update for E

850 395 464 42 170 Update for C

800 355 154 75 175 Update for D

WATERLOO ENGINEERING

Variations of the Design

- Compound Routing Indices
- Hop-count Routing Indices
- Exponentially Aggregated Routing Indices

Hop-count RI

- Store aggregated RIs for a maximum of **m** hops. **m** is also called *horizon*
- New goodness: The hop that yield the most result with the least number of hops
- Assumptions:
  - Uniform distribution
  - Regular tree with constant fan-out
- Estimator:

$$goodness_{hc}(Neighbor_i, Q) = \sum_{j=0..h} \frac{goodness(N_i[j], Q)}{F^j}$$

Hop-count RI Cont'

	1 Hop					2 Hop				
	#	DB	N	T	L	#	DB	N	T	L
B	150	60	34	12	0	230	70	140	0	23
C	50	5	0	30	10	100	35	8	45	50

$$G(b) = \frac{[150 \times \frac{34}{150} \times \frac{12}{150}]}{1} + \frac{[230 \times \frac{140}{230} \times \frac{0}{230}]}{2} = 2.72$$

$$G(c) = \frac{[50 \times \frac{0}{50} \times \frac{30}{50}]}{1} + \frac{[35 \times \frac{8}{35} \times \frac{45}{35}]}{2} = 140$$

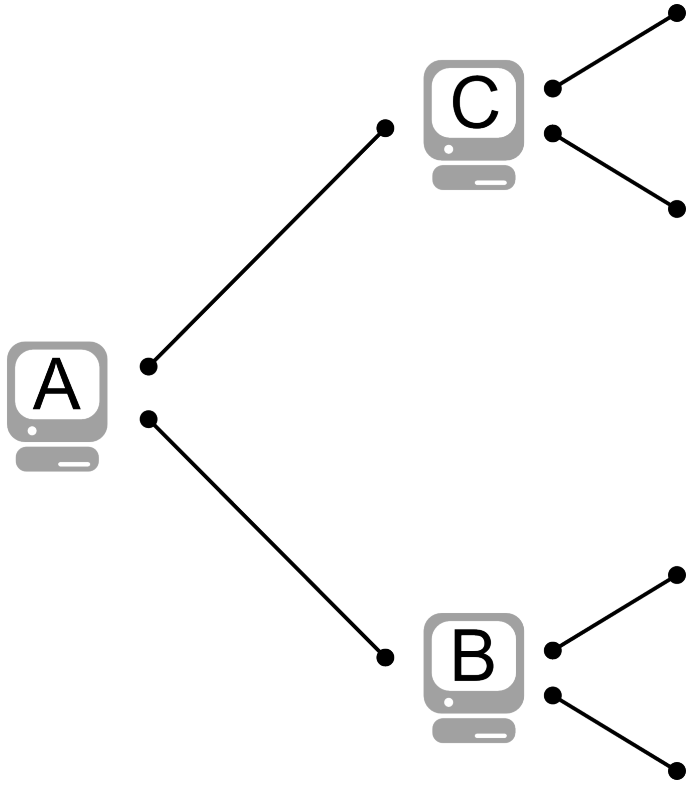

WATERLOO ENGINEERING

Exponentially Aggregated RI

- Compresses HRI tables into one table
- Simply apply the regular-tree cost to HRI model
- Is not bounded by the *horizon*
- Effects of updates dissipates through out the path
- Trades storage for accuracy

Cycles in the Network

Updates can loop indefinitely throughout the network

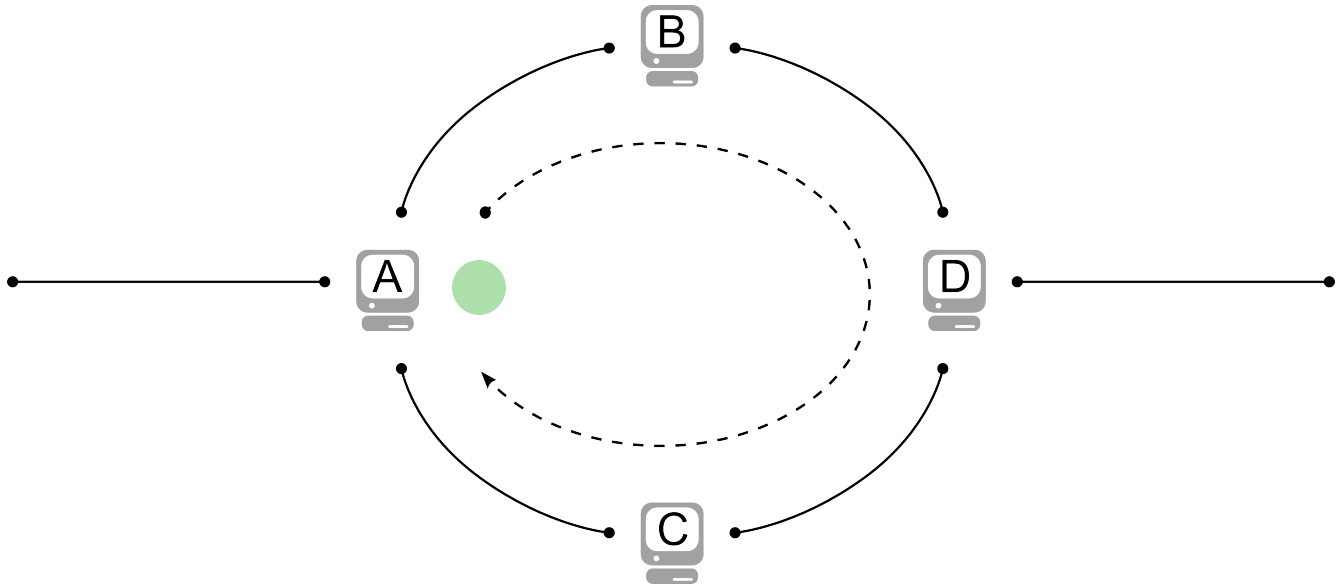

Dealing With Cycles

- No solution
- Cycle avoidance solution
- Cycle detection and recovery solution

No Solution?

Simply no changes are made to any of the algorithms

- Pros:
  Works for Hop-count and Exponential RI
- Cons:
  Can loop indefinitely for CRI

Cycle Avoidance

Do not allow update connections that results in cycles

- Pros:
  Consistency is maintained since there are no loops
- Cons:
  Sub-optimal update model

Cycle Detection & Recovery

Include a message identifier in the update in order to detect cycles

- Pros:
  Consistency of the system is maintained
- Cons:
  Race condition, ignores hop-count

Strengths

- Major: Great experiment setup and high gain
- Minor: High peer autonomy
- Minor: DV inspired update propagation

Weaknesses

- Major: Starvation
- Major: keywords, compression and narrowed application domain
- Major: Race condition regarding cycle recovery solution

Starvation

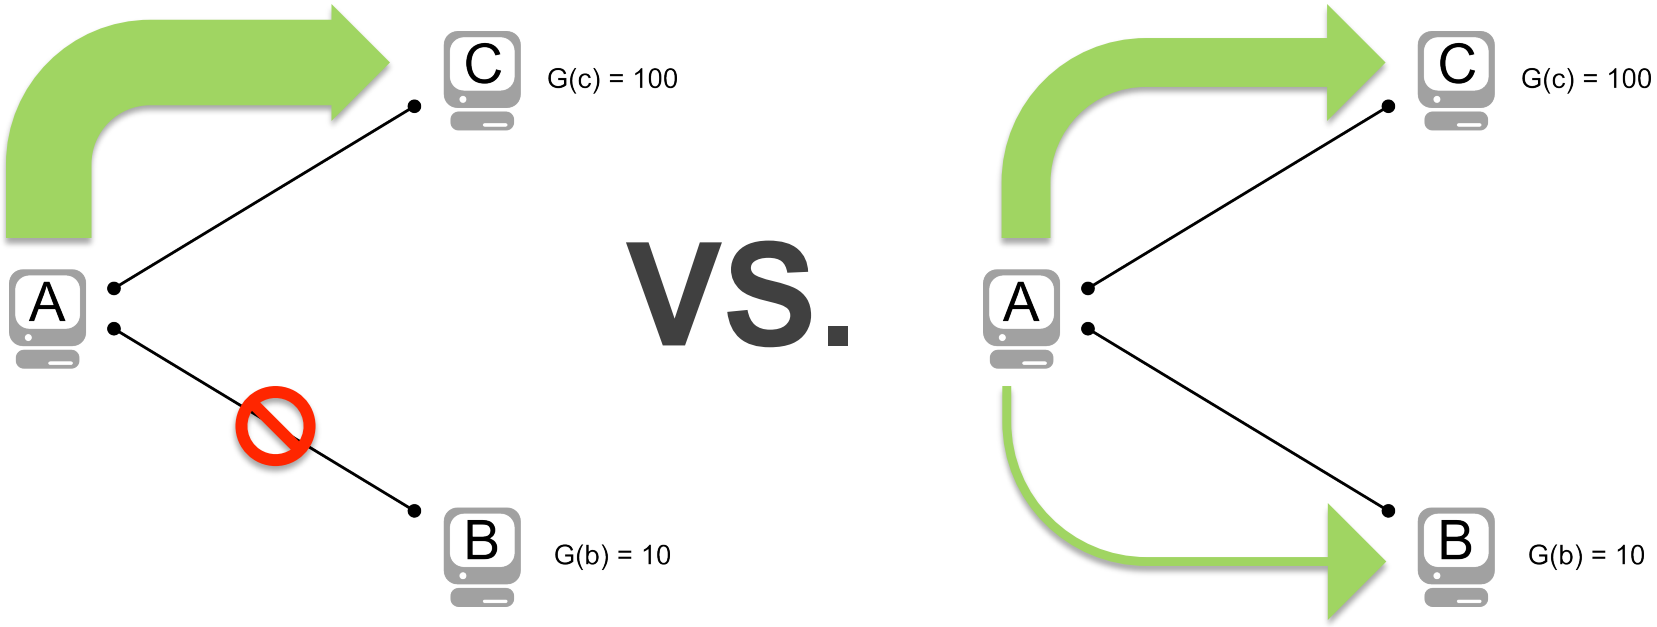

WATERLOO ENGINEERING

Race Condition

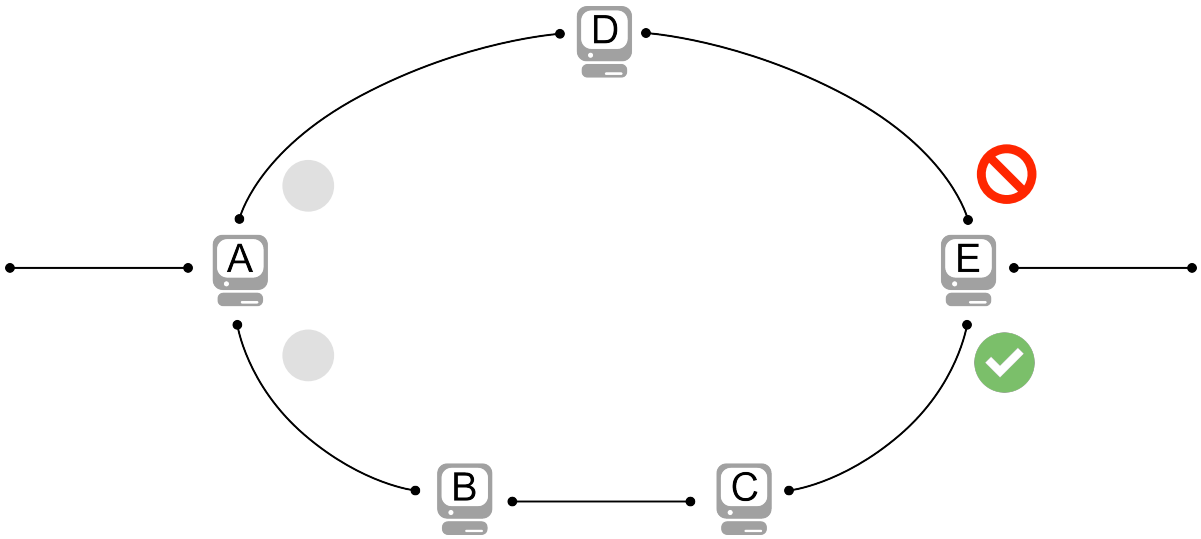

Future Extension

- Probabilistic Routing Model:

  Instead of blindly choosing the neighbor with the highest goodness, distribute query among paths based on the combination of goodness and query size.

Final Evaluation

- Main issues:
  - Radical assumptions regarding certain design decisions
  - Narrowed application space
- Final score: **2.75**

WATERLOO ENGINEERING

Questions?

Thank You!

WATERLOO
ENGINEERING