Distance Halving

Continuous Graphs

Matthias Ohst

Ferienakademie im Sarntal 2008
FAU Erlangen-Nürnberg, TU München, Uni Stuttgart

September 30, 2008

Overview

Overview

1 Introduction

2 Continuous Graphs
- The Distance Halving Graph
- From Continuous Graphs to Discrete Graphs

3 Insertion of Peers and the Principle of Multiple Choice
- The Principle of Multiple Choice
- Two Lemmas Concerning the Principle of Multiple Choice
- Insertion of Peers

4 Routing in the Distance Halving Network
- Simple Algorithm
- Congestion Optimized Algorithm

5 Conclusion

Motivation

- Goal: constant degree and logarithmic diameter (degree minimized network)
- *Viceroy*: complex network

The Distance Halving Network

- 2003: Moni Naor, Udi Wieder

The Distance Halving Network

- 2003: Moni Naor, Udi Wieder
- Put great emphasis on the principle of *continuous graphs*
  - Actually used in networks *CAN* and *Chord*
  - Formalized first by Naor and Wieder

The Distance Halving Network

- 2003: Moni Naor, Udi Wieder
- Put great emphasis on the principle of *continuous graphs*
  - Actually used in networks *CAN* and *Chord*
  - Formalized first by Naor and Wieder
- Graph: Pair (V, E) with $E \subseteq V \times V$
  - Discrete Graph: finite set V
  - Continuous Graph: infinite set V

Overview

The Distance Halving Graph

- Vertex set: $V = [0,1) \subseteq \mathbb{R}$
- Edge set: $E \subseteq V \times V$

The Distance Halving Graph

- Vertex set: $V = [0, 1) \subseteq \mathbb{R}$
- Edge set: $E \subseteq V \times V$
- Four types of edges ($x \in [0, 1)$):
  - Left edges: $(x, \frac{x}{2})$
  - Right edges: $(x, \frac{1}{2} + \frac{x}{2})$
  - Backward left edges: $(\frac{x}{2}, x)$
  - Backward right edges: $(\frac{1}{2} + \frac{x}{2}, x)$

The Distance Halving Graph

- Vertex set: $V = [0, 1) \subseteq \mathbb{R}$
- Edge set: $E \subseteq V \times V$
- Four types of edges ($x \in [0, 1)$):
  - Left edges: $(x, \frac{x}{2})$
  - Right edges: $(x, \frac{1}{2} + \frac{x}{2})$
  - Backward left edges: $(\frac{x}{2}, x)$
  - Backward right edges: $(\frac{1}{2} + \frac{x}{2}, x)$
- Two edges (x_1, y_1), (x_2, y_2):
  - Both left edges or both right edges: $|y_1 - y_2| = \frac{|x_1 - x_2|}{2}$
  - Hence the name: *Distance Halving*
  - Conversely both backward left edges or both backward right edges: $|y_1 - y_2| = 2|x_1 - x_2|$

The Distance Halving Graph

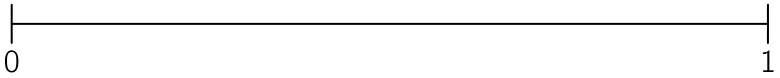

The Distance Halving Graph

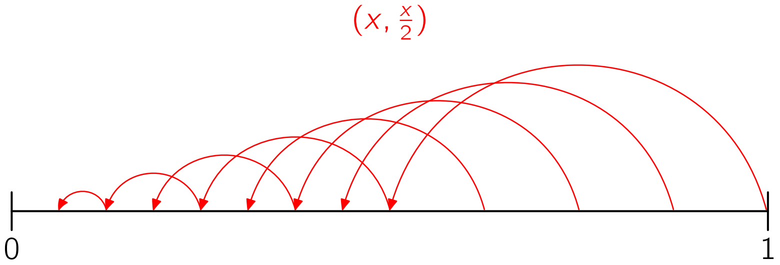

The Distance Halving Graph

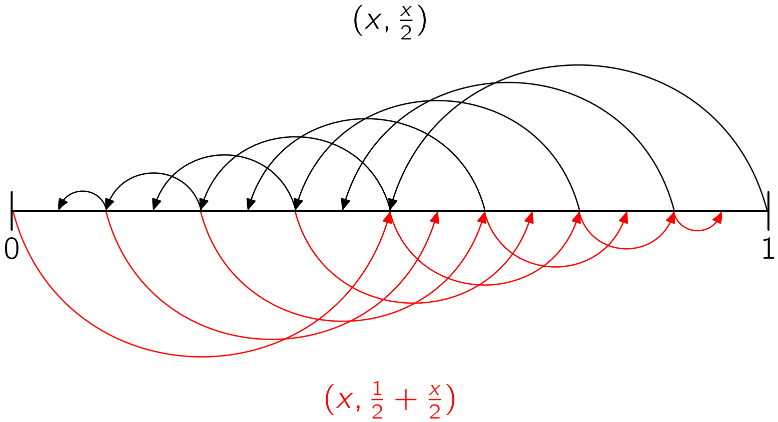

The Distance Halving Graph

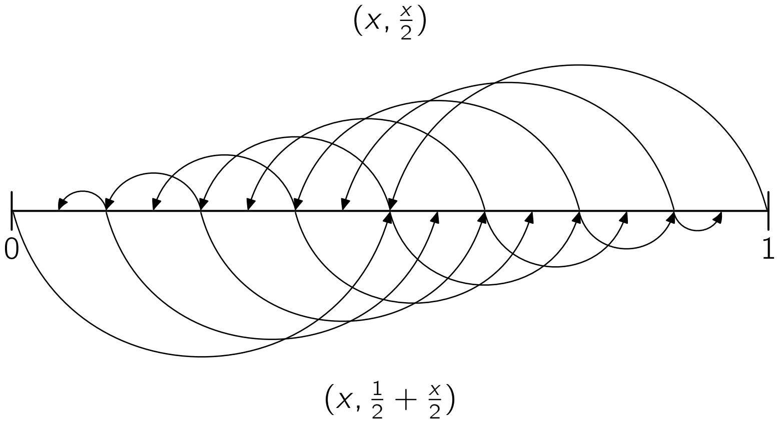

From Continuous Graphs to Discrete Graphs

- Continuous graphs: not directly useable because of the infinite number of vertices
- Partitioning the infinite vertex set V into finite many intervals (vertices of the discrete graph), called *segments*
- In our case: vertices (resp., segments) correspond to the peers in the network

From Continuous Graphs to Discrete Graphs

- Continuous graphs: not directly useable because of the infinite number of vertices
- Partitioning the infinite vertex set V into finite many intervals (vertices of the discrete graph), called *segments*
- In our case: vertices (resp., segments) correspond to the peers in the network
- Simplest case: peers will be placed randomly in the interval $[0, 1)$
- Peers: responsible for data from their posititión up to the position of their successor in the interval $[0, 1)$
- Actually a modified positioning method is used in the Distance Halving network

From Continuous Graphs to Discrete Graphs (cont.)

- Positions of the n peers: $x_1, \dots, x_n$ in ascending order, i.e. $x_i < x_j$ for $i < j$
- The peer x_i, $1 \leq i \leq n$, is assigned the segment $s(x_i) = [x_i, x_{i+1})$

From Continuous Graphs to Discrete Graphs (cont.)

- Positions of the n peers: $x_1, \ldots, x_n$ in ascending order, i.e. $x_i < x_j$ for $i < j$
- The peer x_i, $1 \leq i \leq n$, is assigned the segment $s(x_i) = [x_i, x_{i+1})$
- There is an edge between two segments $s(x_i)$ and $s(x_j)$ iff points $u \in s(x_i)$ and $v \in s(x_j)$ exist such that (u, v) is an edge in the continuous graph
- In addition there are edges between adjacent segments (existence of a ring structure)

From Continuous Graphs to Discrete Graphs (cont.)

- Positions of the n peers: $x_1, \ldots, x_n$ in ascending order, i.e. $x_i < x_j$ for $i < j$
- The peer x_i, $1 \leq i \leq n$, is assigned the segment $s(x_i) = [x_i, x_{i+1})$
- There is an edge between two segments $s(x_i)$ and $s(x_j)$ iff points $u \in s(x_i)$ and $v \in s(x_j)$ exist such that (u, v) is an edge in the continuous graph
- In addition there are edges between adjacent segments (existence of a ring structure)
- Everey path in the continuous graph can be mapped to a path in the discrete graph
- Discretization of the graph described above $\rightsquigarrow$ *Distance Halving* network

Example for Discretization

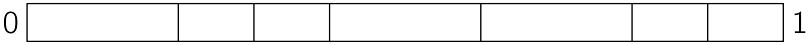

Example for Discretization

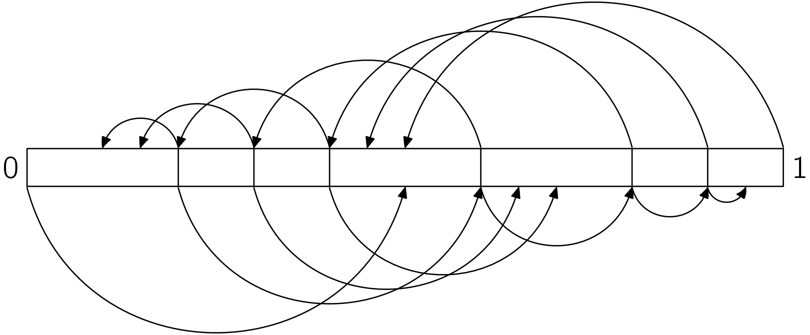

Example for Discretization

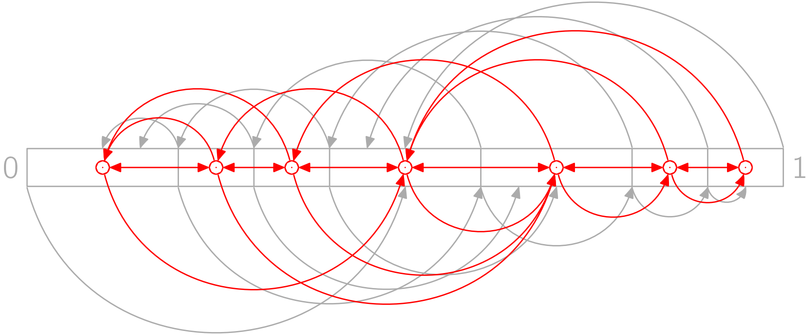

Example for Discretization

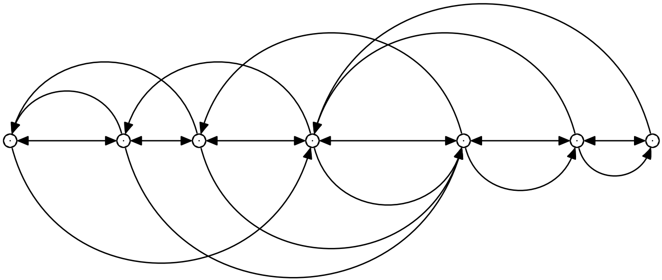

Degree of the Distance Halving Network

- The degree of the Distance Halving network is constant if the ratio of the biggest to the smallest interval is constant
- The edges of a segment map to an interval I which is for every type of edge at most twice as big as the segment itself

Degree of the Distance Halving Network

- The degree of the Distance Halving network is constant if the ratio of the biggest to the smallest interval is constant
- The edges of a segment map to an interval I which is for every type of edge at most twice as big as the segment itself
- Let $\rho = \max_{1 \leq i,j \leq n} \frac{|s(x_i)|}{|s(x_j)|}$ be the ratio of the maximal segment size to the minimal segment size
- The interval I can only intersect with at most $2\rho + 1$ segments

Degree of the Distance Halving Network

- The degree of the Distance Halving network is constant if the ratio of the biggest to the smallest interval is constant
- The edges of a segment map to an interval I which is for every type of edge at most twice as big as the segment itself
- Let $\rho = \max_{1 \leq i,j \leq n} \frac{|s(x_i)|}{|s(x_j)|}$ be the ratio of the maximal segment size to the minimal segment size
- The interval I can only intersect with at most $2\rho + 1$ segments
- A constant ratio of $\rho = 4$ can be achieved by the *principle of multiple choice*
- Increase of degree by a factor of nine by the discretization and hence a constant degree for the Distance Halving network

Overview

The Principle of Multiple Choice

- Instead of choosing a random position in the $[0, 1)$ ring at insertion, every peer looks first at $k = c \log n$ random positions $y_1, \ldots, y_k \in [0, 1)$
- For every position y_i the size $a(y_i)$ of the segment $s(x_*)$ which surrounds the point y_i is measured

The Principle of Multiple Choice

- Instead of choosing a random position in the $[0, 1)$ ring at insertion, every peer looks first at $k = c \log n$ random positions $y_1, \ldots, y_k \in [0, 1)$
- For every position y_i the size $a(y_i)$ of the segment $s(x_*)$ which surrounds the point y_i is measured
- The biggest of the segments found is chosen and the new peer is placed in the middle of that segment
- Always a relatively big segment is chosen, which implies that the distances are relatively uniformly

Example for the Principle of Multiple Choice

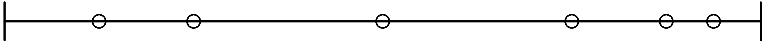

Example for the Principle of Multiple Choice

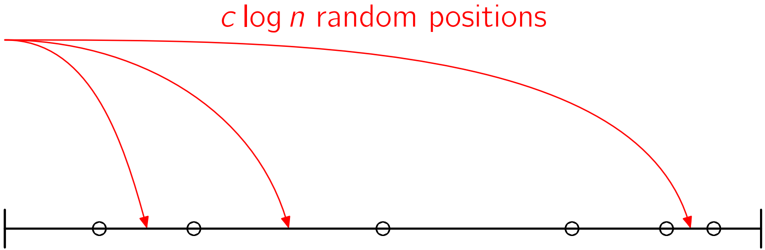

Example for the Principle of Multiple Choice

$c \log n$ random positions

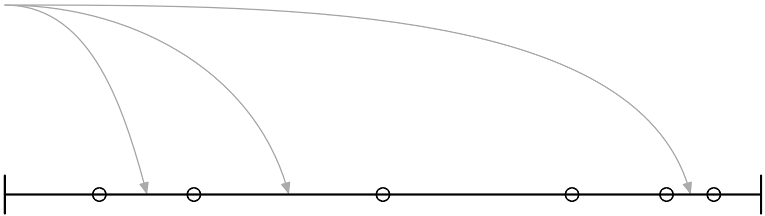

biggest found segment

Example for the Principle of Multiple Choice

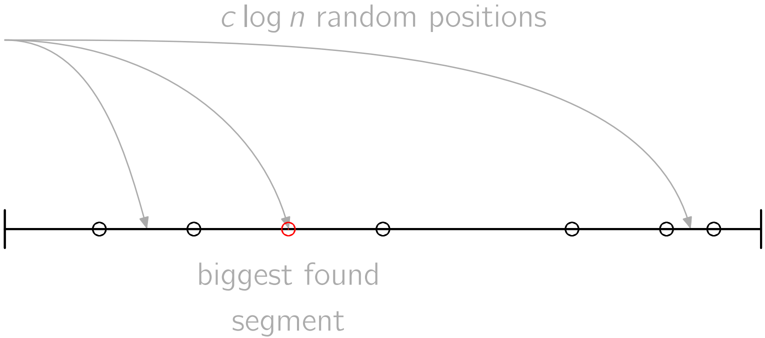

Example for the Principle of Multiple Choice

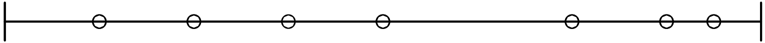

Proof of Lemma 1 (First Part)

Lemma

If $n = 2^k$, $k \in \mathbb{N}$, peers are inserted in the $[0, 1)$ ring using the principle of multiple choice, with high probability only segments of sizes $\frac{1}{2n}$, $\frac{1}{n}$ and $\frac{2}{n}$ are left.

Proof of Lemma 1 (First Part)

Lemma

If $n = 2^k$, $k \in \mathbb{N}$, *peers are inserted in the* $[0, 1)$ *ring using the principle of multiple choice, with high probability only segments of sizes* $\frac{1}{2n}$, $\frac{1}{n}$ *and* $\frac{2}{n}$ *are left.*

Proof (first part):

- Segment sizes: powers of two
- It remains to show:
  - there are no segments of size less than $\frac{1}{2n}$
  - there are no segments of size greater than $\frac{2}{n}$

Proof of Lemma 2 (first part)

Lemma

Let the biggest segment have the size $\frac{g}{n}$ (g may depend on n). Then after insertion of $\frac{2n}{g}$ peers all segments are smaller than $\frac{g}{2n}$.

Proof of Lemma 2 (first part)

Lemma

Let the biggest segment have the size $\frac{g}{n}$ (g may depend on n). Then after insertion of $\frac{2n}{g}$ peers all segments are smaller than $\frac{g}{2n}$.

Proof:

- Consider a segment of size $\frac{g}{n}$
- If $c \log n$ possible positions are examined during the insertion of every peer and $\frac{2n}{g}$ peers are inserted, the expected number of hits X in such an interval is $E[X] = \frac{g}{n} \cdot \frac{2n}{g} \cdot c \log n = 2c \log n$

Proof of Lemma 2 (second part)

Proof (second part):

- With the Chernoff bound we get for $0 \leq \delta \leq 1$:
  $\Pr[X \leq (1-\delta)E[X]] \leq n^{-\delta^2 c}$
- $\delta^2 c \geq 2$: all these intervals are hit at least $2(1-\delta)c \log n$ times
- $2(1-\delta) \geq 1$: every interval of minimum length $\frac{g}{n}$ will be divided with high probability

□

Proof of Lemma 1 (second part)

Proof (second part):

- If one applies the previous lemma for $g = \frac{n}{2}, \frac{n}{4}, \ldots, 4$, then with high probability no interval of size $\frac{g}{n}$ exists
- The number of used peers is $4 + 8 + \cdots + \frac{n}{4} + \frac{n}{2} \leq n$

Proof of Lemma 1 (second part)

Proof (second part):

- If one applies the previous lemma for $g = \frac{n}{2}, \frac{n}{4}, \dots, 4$, then with high probability no interval of size $\frac{g}{n}$ exists
- The number of used peers is $4 + 8 + \cdots + \frac{n}{4} + \frac{n}{2} \leq n$
- After the last round there are no segments bigger than $\frac{2}{n}$
- Since here only $\mathcal{O}(\log n)$ events have to arrive, the statement holds with high probability

Proof of Lemma 1 (third part)

Proof (third part):

- It remains to show: no segments smaller than $\frac{1}{2n}$ arise
- The total length of all segments of size $\frac{1}{2n}$ is at most $\frac{n}{2}$ before insertion

Proof of Lemma 1 (third part)

Proof (third part):

- It remains to show: no segments smaller than $\frac{1}{2n}$ arise
- The total length of all segments of size $\frac{1}{2n}$ is at most $\frac{n}{2}$ before insertion
- The probability that only such segments are chosen by $c \log n$ tests is at most $2^{-c \log n} = n^{-c}$
- For $c > 1$ a segment of size $\frac{1}{2n}$ is farther divided only with polynomially low probability

□

Insertion of Peers

- Needed: Approximation value of the number n of peers in the network
- Estimation achieved by the distance of neighbors

Insertion of Peers

- Needed: Approximation value of the number n of peers in the network
- Estimation achieved by the distance of neighbors
- Estimation in the Distance Halving network:
  exact except for a factor of 4
  - Biggest segment size: $\frac{2}{n}$
  - Smallest segment size: $\frac{1}{2n}$

Insertion of Peers

- At insertion the $c \log n$ segments that have to be checked are localized by a search
- For this $\mathcal{O}(\log n)$ steps are needed as we will see shortly

Insertion of Peers

- At insertion the $c \log n$ segments that have to be checked are localized by a search
- For this $\mathcal{O}(\log n)$ steps are needed as we will see shortly
- After the biggest segment was chosen:
  - The peer will be embedded in the ring structure
  - Then it establishes the further connections to the other peers with the help of the adjacent peers on the ring
- Accordingly the other neighbors in the network update, too

Overview

Routing in the Distance Halving Network

- Goal: routing algorithm which only needs $\mathcal{O}(\log n)$ steps and distributes congestion uniformly

Routing in the Distance Halving Network

- Goal: routing algorithm which only needs $\mathcal{O}(\log n)$ steps and distributes congestion uniformly
- First: Simplified version which distributes congestion not uniformly

The Simplified Version

leftRouting(*src*, *dest*)
  if *src* and *dest* adjacent **then**
    send message from *src* to *dest*
  else
    newSrc ← leftPointer(*src*)
    newDest ← leftPointer(*dest*)
    send message from *src* to *newSrc*
    leftRouting(*newSrc*, *newDest*)
    send message from *newDest* to *dest*

Example for the Simplified Version

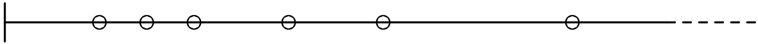

Example for the Simplified Version

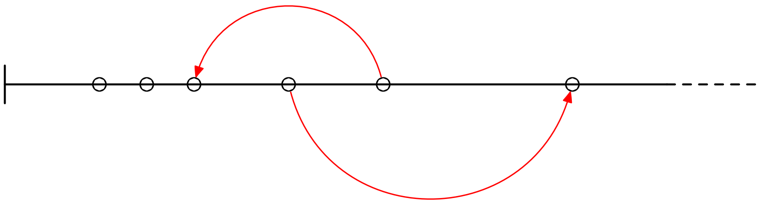

Example for the Simplified Version

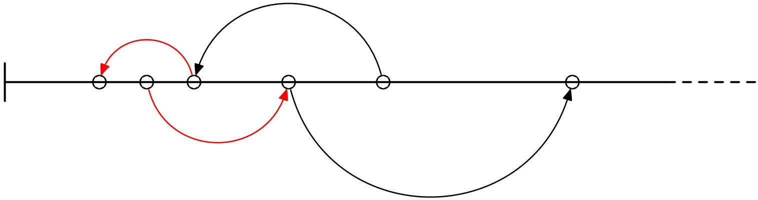

Example for the Simplified Version

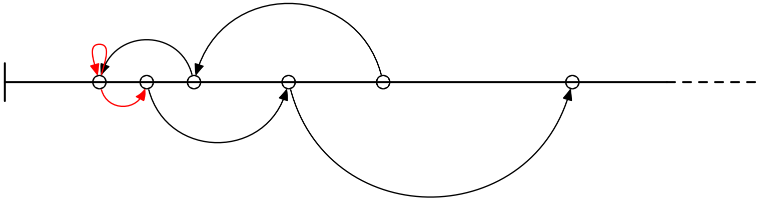

Example for the Simplified Version

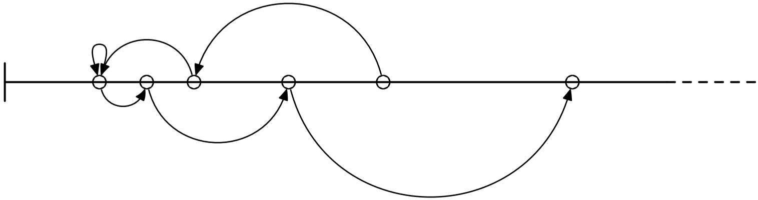

Routing in the Distance Halving Network

- This algorithm: only left edges
- The source peer calculates two intermediate stations and reduces routing to half the distance
- This continues until source and destination nodes are adjacent

Routing in the Distance Halving Network

- This algorithm: only left edges
- The source peer calculates two intermediate stations and reduces routing to half the distance
- This continues until source and destination nodes are adjacent
- The calculation of intermediate stations is done by the source node
- The intermediate stations must be told which path the message has to be carried on

The Simplified Version Using Right Edges

rightRouting(*src*, *dest*)
  if *src* and *dest* adjacent **then**
    send message from *src* to *dest*
  else
    newSrc $\leftarrow$ rightPointer(*src*)
    newDest $\leftarrow$ rightPointer(*dest*)
    send message from *src* to *newSrc*
    rightRouting(*newSrc*, *newDest*)
    send message from *newDest* to *dest*

Routing in the Distance Halving Network

- In both algorithms the distance between source and destination is halved every recursion step and every recursion step needs two steps
- Since all interval sizes differ only by a factor of $\rho = 4$, the routing algorithm needs at most $1 + \log n$ recursions to deliver a message

Routing in the Distance Halving Network

- In both algorithms the distance between source and destination is halved every recursion step and every recursion step needs two steps
- Since all interval sizes differ only by a factor of $\rho = 4$, the routing algorithm needs at most $1 + \log n$ recursions to deliver a message

Lemma

With high probability the routing in the Distance Halving network needs at most $2 \log n + 3$ messages and steps.

Routing in the Distance Halving Network

- Left and right edges can be exchanged arbitrarily in these algorithms
  $\rightsquigarrow$ possibility to decide orientation (pairwise) by coin toss

Routing in the Distance Halving Network

- Left and right edges can be exchanged arbitrarily in these algorithms $\rightsquigarrow$ possibility to decide orientation (pairwise) by coin toss
- First two algorithms: tending to send traffic into the outermost left or right corner
- This algorithm: good distribution of congestion
- One can show that congestion is very low

Congestion Optimized Algorithm

```
randomRouting(src, dest)
  if src and dest adjacent then
    send message from src to dest
  else
    if coin shows number then
      newSrc  ← leftPointer(src)
      newDest ← leftPointer(dest)
    else
      newSrc  ← rightPointer(src)
      newDest ← rightPointer(dest)
    send message from src to newSrc
    randomRouting(newSrc, newDest)
    send message from newDest to dest
```

Example for the Algorithm

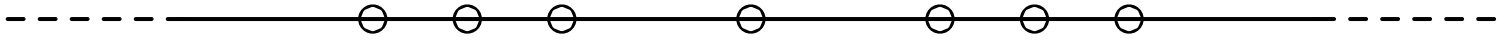

Example for the Algorithm

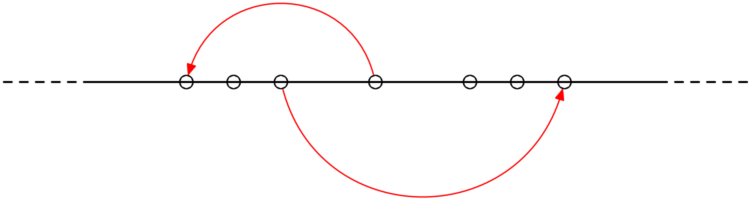

Example for the Algorithm

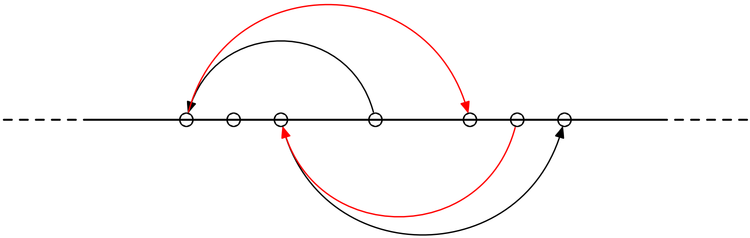

Example for the Algorithm

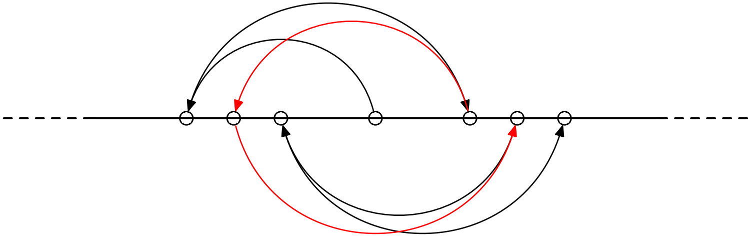

Example for the Algorithm

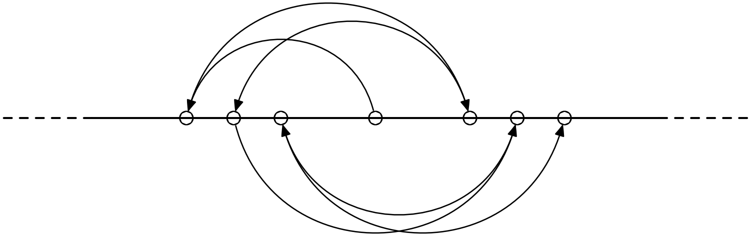

Overview

Conclusion

- Distance Halving network: degree minimized network (constant degree and logarithmic diameter)
- Elegant and simple alternative to the complex *Butterfly* graph based *Viceroy* network

The *Chernoff* bound

Theorem (Chernoff bound)

Let $X_1, \ldots, X_n$ *be independent Bernoulli experiments with probability* $\Pr[X_i = 1] = p$ *and* $X = \sum_{i=1}^{n} X_i$. *Then, for* $\delta \geq 0$,

$$\Pr[X \geq (1+\delta)pn] \leq \mathrm{e}^{-\frac{1}{3}\min\{\delta,\delta^2\}pn} .$$

Furthermore, if $0 \leq \delta \leq 1$,

$$\Pr[X \leq (1-\delta)pn] \leq \mathrm{e}^{-\frac{1}{2}\delta^2 pn} .$$

Return