

Efficiency of navigation in indexed networks

Petter Holme

Royal Institute of Technology (Sweden), University of New Mexico (USA)


October 5, 2007, ECCS satellite workshop

<http://www.csc.kth.se/~pholme/>



ROYAL INSTITUTE
OF TECHNOLOGY


a graph



ROYAL INSTITUTE
OF TECHNOLOGY




source & target




ROYAL INSTITUTE
OF TECHNOLOGY

myopia




ROYAL INSTITUTE
OF TECHNOLOGY

navigation




ROYAL INSTITUTE
OF TECHNOLOGY

navigation




ROYAL INSTITUTE
OF TECHNOLOGY

navigation




ROYAL INSTITUTE
OF TECHNOLOGY

navigation



ROYAL INSTITUTE
OF TECHNOLOGY

navigation



ROYAL INSTITUTE
OF TECHNOLOGY

the problem: summary

- a (static) graph G of N vertices
- each vertex is associated with an index $1, \dots, N$
- packets are created with source s and destination t
- packets have memory, but no previous knowledge of the graph
- packets can observe the neighborhood of their present location
- how can we choose the indices, and a packet navigation algorithm, so that the average effective distance, or *transit time* τ , is minimized?



the problem: summary

- a (static) graph G of N vertices
- each vertex is associated with an index $1, \dots, N$
- packets are created with source s and destination t
- packets have memory, but no previous knowledge of the graph
- packets can observe the neighborhood of their present location
- how can we choose the indices, and a packet navigation algorithm, so that the average effective distance, or *transit time* τ , is minimized?



the problem: summary

- a (static) graph G of N vertices
- each vertex is associated with an index $1, \dots, N$
- packets are created with source s and destination t
- packets have memory, but no previous knowledge of the graph
- packets can observe the neighborhood of their present location
- how can we choose the indices, and a packet navigation algorithm, so that the average effective distance, or *transit time* τ , is minimized?



the problem: summary

- a (static) graph G of N vertices
- each vertex is associated with an index $1, \dots, N$
- packets are created with source s and destination t
- packets have memory, but no previous knowledge of the graph
- packets can observe the neighborhood of their present location
- how can we choose the indices, and a packet navigation algorithm, so that the average effective distance, or *transit time* τ , is minimized?



the problem: summary

- a (static) graph G of N vertices
- each vertex is associated with an index $1, \dots, N$
- packets are created with source s and destination t
- packets have memory, but no previous knowledge of the graph
- packets can observe the neighborhood of their present location
- how can we choose the indices, and a packet navigation algorithm, so that the average effective distance, or *transit time* τ , is minimized?



the problem: summary

- a (static) graph G of N vertices
- each vertex is associated with an index $1, \dots, N$
- packets are created with source s and destination t
- packets have memory, but no previous knowledge of the graph
- packets can observe the neighborhood of their present location
- how can we choose the indices, and a packet navigation algorithm, so that the average effective distance, or *transit time* τ , is minimized?




the problem: summary

- a (static) graph G of N vertices
- each vertex is associated with an index $1, \dots, N$
- packets are created with source s and destination t
- packets have memory, but no previous knowledge of the graph
- packets can observe the neighborhood of their present location
- how can we choose the indices, and a packet navigation algorithm, so that the average effective distance, or *transit time* τ , is minimized?




a tree




ROYAL INSTITUTE
OF TECHNOLOGY

search tree



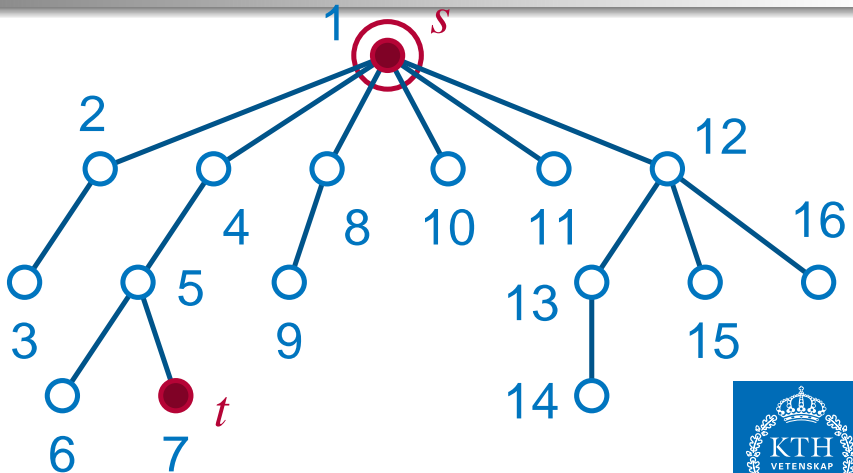
ROYAL INSTITUTE
OF TECHNOLOGY

search tree




ROYAL INSTITUTE
OF TECHNOLOGY

search tree




ROYAL INSTITUTE
OF TECHNOLOGY

search tree




ROYAL INSTITUTE
OF TECHNOLOGY

search tree




ROYAL INSTITUTE
OF TECHNOLOGY

search tree




ROYAL INSTITUTE
OF TECHNOLOGY

search tree




ROYAL INSTITUTE
OF TECHNOLOGY

accurate search down (ASD) navigation




accurate search down (ASD) navigation




ROYAL INSTITUTE
OF TECHNOLOGY

accurate search down (ASD) navigation




ROYAL INSTITUTE
OF TECHNOLOGY

accurate search down (ASD) navigation




ROYAL INSTITUTE
OF TECHNOLOGY

accurate search down (ASD) navigation




ROYAL INSTITUTE
OF TECHNOLOGY

accurate search down (ASD) navigation




ROYAL INSTITUTE
OF TECHNOLOGY


accurate search down (ASD) navigation




accurate search down (ASD) navigation




accurate search down (ASD) navigation



accurate search down (ASD) navigation




accurate search down (ASD) navigation



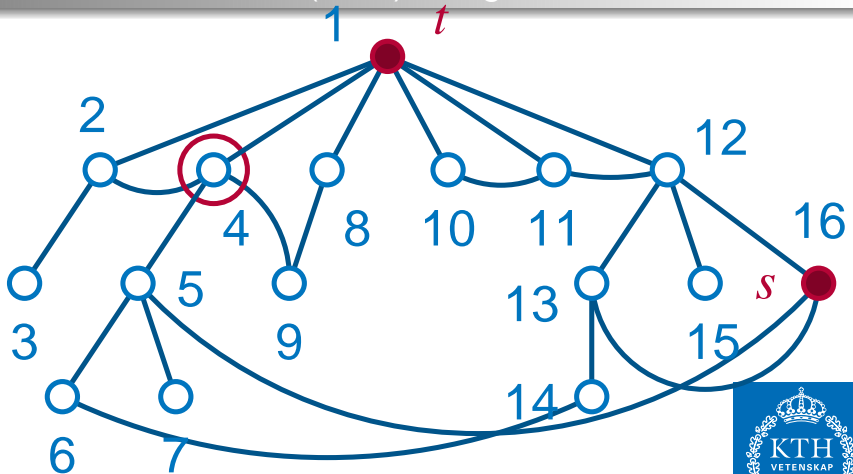
ROYAL INSTITUTE
OF TECHNOLOGY

accurate search down (ASD) navigation




ROYAL INSTITUTE
OF TECHNOLOGY

accurate search down (ASD) navigation



ROYAL INSTITUTE
OF TECHNOLOGY

accurate search down (ASD) navigation



accurate search down (ASD) indexing

- choose one of the vertices with least eccentricity as root
- index the rest with a depth first search
- increment index at the first encounter of a vertex



accurate search down (ASD) indexing

- choose one of the vertices with least eccentricity as root
- index the rest with a depth first search
- increment index at the first encounter of a vertex



accurate search down (ASD) indexing

- choose one of the vertices with least eccentricity as root
- index the rest with a depth first search
- increment index at the first encounter of a vertex



accurate search down (ASD) indexing

- choose one of the vertices with least eccentricity as root
- index the rest with a depth first search
- increment index at the first encounter of a vertex



accurate search down (ASD) navigation

- first, move towards smaller indices until the root, 1, is reached
- then, move to the neighbor with largest index $\geq t$
- if t is in the neighborhood during the upward search, go there



accurate search down (ASD) navigation

- first, move towards smaller indices until the root, 1, is reached
- then, move to the neighbor with largest index $\geq t$
- if t is in the neighborhood during the upward search, go there



accurate search down (ASD) navigation

- first, move towards smaller indices until the root, 1, is reached
- then, move to the neighbor with largest index $\geq t$
- if t is in the neighborhood during the upward search, go there




accurate search down (ASD) navigation

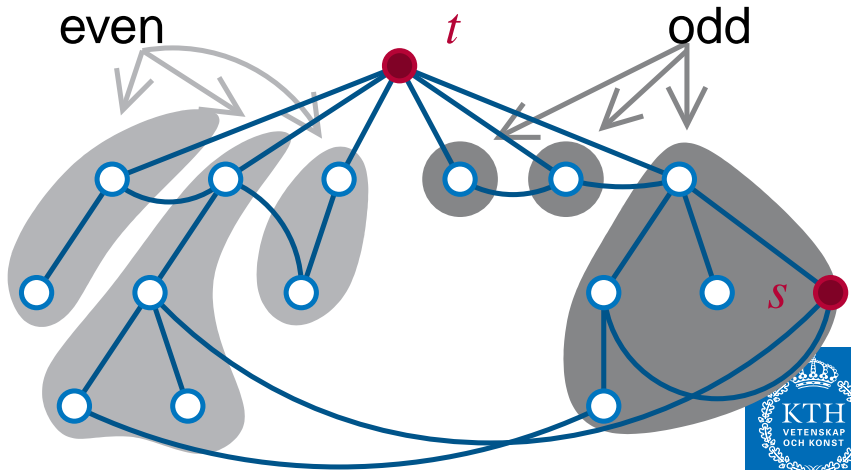
- first, move towards smaller indices until the root, 1, is reached
- then, move to the neighbor with largest index $\geq t$
- if t is in the neighborhood during the upward search, go there



accurate search up (ACU) indexing & navigation



accurate search up (ACU) indexing & navigation



ROYAL INSTITUTE
OF TECHNOLOGY

random navigation

- random depth first search
- if t is in the neighborhood, go there



random navigation

- random depth first search
- if t is in the neighborhood, go there



random navigation

- random depth first search
- if t is in the neighborhood, go there



degree biased navigation

- depth first search in order of degree
- if t is in the neighborhood, go there



degree biased navigation

- depth first search in order of degree
- if t is in the neighborhood, go there



degree biased navigation

- depth first search in order of degree
- if t is in the neighborhood, go there



Network models

- **Barabási–Albert model** Power-law degree distribution, vanishing clustering (fraction of triangles)
- **Holme–Kim model** Power-law degree distribution, tunable clustering
- **Erdős–Rényi** Poisson degree distribution, vanishing clustering
- **Square grid**



Network models

- **Barabási–Albert model** Power-law degree distribution, vanishing clustering (fraction of triangles)
- **Holme–Kim model** Power-law degree distribution, tunable clustering
- **Erdős–Rényi** Poisson degree distribution, vanishing clustering
- **Square grid**



Network models

- **Barabási–Albert model** Power-law degree distribution, vanishing clustering (fraction of triangles)
- **Holme–Kim model** Power-law degree distribution, tunable clustering
- **Erdős–Rényi** Poisson degree distribution, vanishing clustering
- **Square grid**



Network models

- **Barabási–Albert model** Power-law degree distribution, vanishing clustering (fraction of triangles)
- **Holme–Kim model** Power-law degree distribution, tunable clustering
- **Erdős–Rényi** Poisson degree distribution, vanishing clustering
- **Square grid**

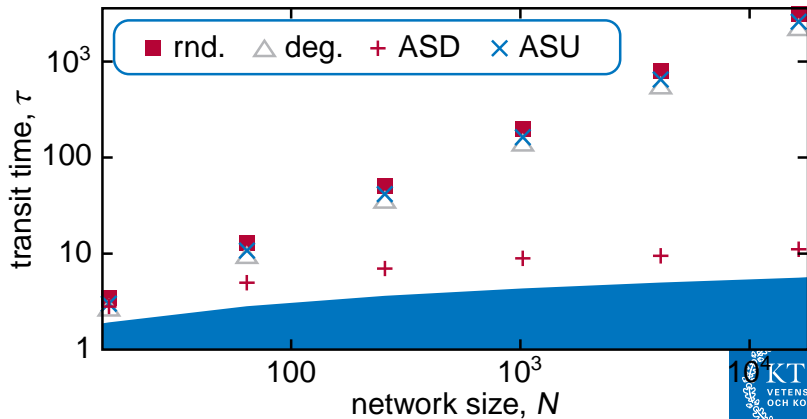


Network models


- **Barabási–Albert model** Power-law degree distribution, vanishing clustering (fraction of triangles)
- **Holme–Kim model** Power-law degree distribution, tunable clustering
- **Erdős–Rényi** Poisson degree distribution, vanishing clustering
- **Square grid**



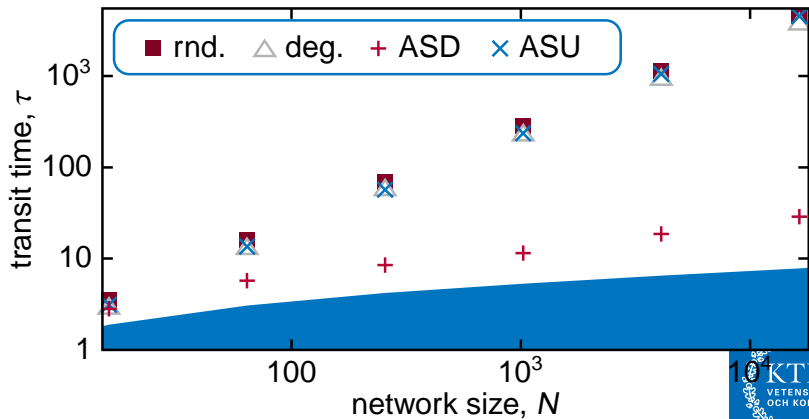
Barabási–Albert model




Holme–Kim model



Erdős–Rényi model



square grid



Possible application: P2P protocols.

Things needed to be addressed:

- how to update the indices
- the root is vulnerable
- . . . and a bottleneck



Possible application: P2P protocols.

Things needed to be addressed:

- how to update the indices
- the root is vulnerable
- . . . and a bottleneck



P2P

Possible application: P2P protocols.

Things needed to be addressed:

- how to update the indices
- the root is vulnerable
- . . . and a bottleneck



ROYAL INSTITUTE
OF TECHNOLOGY

Possible application: P2P protocols.

Things needed to be addressed:

- how to update the indices
- the root is vulnerable
- . . . and a bottleneck



summary

- we have presented two indexing & navigation schemes (+ two reference schemes) for the indexed-graph-navigation problem on networks
- these schemes were tested on four types of model network
- one of the schemes, ASD, performed consistently better than the others
- the other search-tree based scheme does not perform significantly better than the reference schemes



summary

- we have presented two indexing & navigation schemes (+ two reference schemes) for the indexed-graph-navigation problem on networks
- these schemes were tested on four types of model network
- one of the schemes, ASD, performed consistently better than the others
- the other search-tree based scheme does not perform significantly better than the reference schemes



summary

- we have presented two indexing & navigation schemes (+ two reference schemes) for the indexed-graph-navigation problem on networks
- these schemes were tested on four types of model network
- one of the schemes, ASD, performed consistently better than the others
- the other search-tree based scheme does not perform significantly better than the reference schemes



summary

- we have presented two indexing & navigation schemes (+ two reference schemes) for the indexed-graph-navigation problem on networks
- these schemes were tested on four types of model network
- one of the schemes, ASD, performed consistently better than the others
- the other search-tree based scheme does not perform significantly better than the reference schemes



summary

- we have presented two indexing & navigation schemes (+ two reference schemes) for the indexed-graph-navigation problem on networks
- these schemes were tested on four types of model network
- one of the schemes, ASD, performed consistently better than the others
- the other search-tree based scheme does not perform significantly better than the reference schemes

