# Algorithm Engineering for Large Graphs 

## Fast Route Planning

Veit Batz, Robert Geisberger, Dennis Luxen, Peter Sanders, Christian Vetter

Universität Karlsruhe (TH)

Zurich, October 22, 2008

## Route Planning

## Goals:

$\square$ exact shortest (i.e. fastest) paths in large road networksfast queries (point-to-point, many-to-many)fast preprocessinglow space consumptionfast update operations

## Applications:


$\square$ route planning systems in the internet, car navigation systems,
$\square$ ride sharing, traffic simulation, logistics optimisation

Sanders et al.: Route Planning

## Overview

$\square$ Exact Contraction Hierarchies - a very simple approach
$\square$ Transit Node Routing - getting really fast
$\square$ Mobile Contraction Hierarchies
$\square$ Many-to-many Routing
$\square$ Ride Sharing
$\square$ Dynamic Scenario
$\square$ Time-dependent Contraction Hierarchies
$\square$ Future Work

## Contraction Hierarchies (CH)



## Main Idea

Contraction Hierarchies (CH)

- contract only one node at a time $\Rightarrow$ local and cache-efficient operation
in more detail:
- order nodes by "importance", $V=\{1,2, \ldots, n\}$
- contract nodes in this order, node $v$ is contracted by foreach pair $(u, v)$ and $(v, w)$ of edges do if $\langle u, v, w\rangle$ is a unique shortest path then add shortcut $(u, w)$ with weight $w(\langle u, v, w\rangle)$
- query relaxes only edges to more "important" nodes $\Rightarrow$ valid due to shortcuts

Sanders et al.: Route Planning

## Example: Construction



Sanders et al.: Route Planning

## Example: Construction



Sanders et al.: Route Planning

## Example: Construction



Sanders et al.: Route Planning

## Example: Construction



Sanders et al.: Route Planning


Sanders et al.: Route Planning
Example: Construction


## Construction

to identify necessary shortcuts

- local searches from all nodes $u$ with incoming edge ( $u, v$ )
- ignore node $v$ at search
- add shortcut ( $u, w$ ) iff found distance $d(u, w)>w(u, v)+w(v, w)$



## Construction

to identify necessary shortcuts

- local searches from all nodes $u$ with incoming edge ( $u, v$ )
- ignore node $v$ at search
- add shortcut ( $u, w$ ) iff found distance $d(u, w)>w(u, v)+w(v, w)$



## Node Order

use priority queue of nodes, node $v$ is weighted with a linear combination of:

- edge difference \#shortcuts - \#edges incident to $v$
- uniformity e.g. \#deleted neighbors
integrated construction and ordering:

1. remove node $v$ on top of the priority queue

$2-3=-1$
2. contract node $v$
3. update weights of remaining nodes

## Query

- modified bidirectional Dijkstra algorithm
- upward graph $\quad G_{\uparrow}:=\left(V, E_{\uparrow}\right)$ with $E_{\uparrow}:=\{(u, v) \in E: u<v\}$ downward graph $G_{\downarrow}:=\left(V, E_{\downarrow}\right)$ with $E_{\downarrow}:=\{(u, v) \in E: u>v\}$
- forward search in $G_{\uparrow}$ and backward search in $G_{\downarrow}$



## Query

- modified bidirectional Dijkstra algorithm
- upward graph $\quad G_{\uparrow}:=\left(V, E_{\uparrow}\right)$ with $E_{\uparrow}:=\{(u, v) \in E: u<v\}$ downward graph $G_{\downarrow}:=\left(V, E_{\downarrow}\right)$ with $E_{\downarrow}:=\{(u, v) \in E: u>v\}$
- forward search in $G_{\uparrow}$ and backward search in $G_{\downarrow}$



## Query

- modified bidirectional Dijkstra algorithm
- upward graph $\quad G_{\uparrow}:=\left(V, E_{\uparrow}\right)$ with $E_{\uparrow}:=\{(u, v) \in E: u<v\}$ downward graph $G_{\downarrow}:=\left(V, E_{\downarrow}\right)$ with $E_{\downarrow}:=\{(u, v) \in E: u>v\}$
- forward search in $G_{\uparrow}$ and backward search in $G_{\downarrow}$



## Query

- modified bidirectional Dijkstra algorithm
- upward graph $\quad G_{\uparrow}:=\left(V, E_{\uparrow}\right)$ with $E_{\uparrow}:=\{(u, v) \in E: u<v\}$ downward graph $G_{\downarrow}:=\left(V, E_{\downarrow}\right)$ with $E_{\downarrow}:=\{(u, v) \in E: u>v\}$
- forward search in $G_{\uparrow}$ and backward search in $G_{\downarrow}$



## Outputting Paths

- for a shortcut $(u, w)$ of a path $\langle u, v, w\rangle$, store middle node $v$ with the edge
- expand path by recursively replacing a shortcut with its originating edges



## Stall-on-Demand

- $v$ can be "stalled" by $u$
(if $d(u)+w(u, v)<d(v)$ )
- stalling can propagate to adjacent nodes
- search is not continued from stalled nodes

- does not invalidate correctness (only suboptimal paths are stalled)


## Experiments

environment

- AMD Opteron Processor 270 at 2.0 GHz
- 8 GB main memory
- GNU C++ compiler 4.2.1
test instance
- road network of Western Europe (PTV)
- 18029721 nodes
- 42199587 directed edges


Sanders et al.: Route Planning

## Performance



HNR: 594 s / $802 \mu \mathrm{~s}$

Sanders et al.: Route Planning

## Worst Case Costs



## Contraction Hierarchies

$\square$ foundation for our other methodsconceptually very simple
$\square$ handles dynamic scenarios

## Static scenario:


$\square 7.5$ min preprocessing
$\square 0.21 \mathrm{~ms}$ to determine the path length
$\square 0.56 \mathrm{~ms}$ to determine a complete path description
$\square$ little space consumption (23 bytes/node)

## Transit-Node Routing

[DIMACS Challenge 06, ALENEX 07, Science 07]
joint work with H. Bast, S. Funke, D. Matijevic
$\square$ very fast queries (down to $1.7 \mu s, 3000000$ times faster than DIJKSTRA)

$\square$ winner of the 9th DIMACS Implementation Challenge
$\square$ more preprocessing time (2:37 h) and space (263 bytes/node) needed


SciAm50 Award


## Mobile Contraction Hierarchies

preprocess data on a personal computerhighly compressed blocked graph representation 8 bytes/node$\square$ compact route reconstruction data structure

+ 8 bytes/node
experiments on a Nokia N800 at 400 MHz

$\square$ cold query with empty block cache 56 ms
$\square$ compute complete path
73 ms
$\square$ recomputation, e.g. if driver took the wrong exit
14 ms
joint work with S. Knopp, F. Schulz, D. Wagner [ALENEX 07]

$\square$
efficient many-to-many variant of hierarchical bidirectional algorithms
$\square 10000 \times 10000$ table in 10 s

\[

\]



## Ride Sharing

## Current approaches:

match only ride offers with identical start/destination (perfect fit)sometimes radial search around start/destination
## Our approach:

$\square$ driver picks passenger up and gives him a ride to his destination
$\square$ find the driver with the minimal detour (reasonable fit)

## Efficient algorithm:

$\square$ adaption of the many-to-many algorithm

## Highway-Node Routing

$\square$ generalization of contraction hierarchies
$\square$ allow multiple nodes in the same 'importance'-level i.e., select node sets $S_{1} \supseteq S_{2} \supseteq S_{3} \ldots$
$\square$ construct multi-level overlay graph
$\square$ perform multi-level query
$\square$ designed for dynamic scenarios


## Overlay Graph

[Holzer, Schulz, Wagner, Weihe, Zaroliagis 2000-2007]
$\square$ graph $G=(V, E)$ is given
$\square$ select node subset $S \subseteq V$


## Overlay Graph

[Holzer, Schulz, Wagner, Weihe, Zaroliagis 2000-2007]
$\square$ graph $G=(V, E)$ is given
$\square$ select node subset $S \subseteq V$

$\square$ overlay graph $G^{\prime}:=\left(S, E^{\prime}\right)$ where
$E^{\prime}:=\{(s, t) \in S \times S \mid$ no inner node of the shortest $s-t$-path belongs to $S\}$

## Dynamic Scenarios

$\square$ change entire cost function (e.g., use different speed profile)

$\square$ change a few edge weights (e.g., due to a traffic jam)


## Constancy of Structure

## Assumption:

$\square$ structure of road network does not change
(no new roads, road removal = set weight to $\infty$ )
$\rightsquigarrow$ not a significant restriction
$\square$ classification of nodes by 'importance' might be slightly perturbed, but not completely changed
(e.g., a sports car and a truck both prefer motorways)
$\rightsquigarrow$ performance of our approach relies on that (not the correctness)

## Dynamic Highway-Node Routing

change entire cost function

$\square$ keep the node sets $S_{1} \supseteq S_{2} \supseteq S_{3} \ldots$
$\square$ recompute the overlay graphs

| speed profile | default | fast car | slow car | slow truck | distance |
| :--- | ---: | ---: | ---: | ---: | ---: |
| constr. [min] | $1: 40$ | $1: 41$ | $1: 39$ | $1: 36$ | $3: 56$ |
| query [ms] | 1.17 | 1.20 | 1.28 | 1.50 | 35.62 |
| \#settled nodes | 1414 | 1444 | 1507 | 1667 | 7057 |

## Dynamic Highway-Node Routing

change a few edge weights

$\square$ server scenario: if something changes,

- update the preprocessed data structures
- answer many subsequent queries very fast

mobile scenario: if something changes,
- it does not pay to update the data structures
- perform single 'prudent' query that takes changed situation into account



## Dynamic Highway-Node Routing

## change a few edge weights, server scenario


$\square$ keep the node sets $S_{1} \supseteq S_{2} \supseteq S_{3} \ldots$
$\square$ recompute only possibly affected parts of the overlay graphs

- the computation of the level- $\ell$ overlay graph consists of $\left|S_{\ell}\right|$ local searches to determine the respective covering nodes
- if the initial local search from $v \in S_{\ell}$ has not touched a now modified edge $(u, x)$, that local search need not be repeated
- we manage sets $A_{u}^{\ell}=\left\{v \in S_{\ell} \mid v\right.$ 's level- $\ell$ preprocessing might be affected when an edge $(u, x)$ changes $\}$


## Dynamic Highway-Node Routing

change a few edge weights, server scenario



## Dynamic Highway-Node Routing

change a few edge weights, mobile scenario


1. keep the node sets $S_{1} \supseteq S_{2} \supseteq S_{3} \ldots$
2. keep the overlay graphs
3. $C:=$ all changed edges
4. use the sets $A_{u}^{\ell}$ (considering edges in $C$ ) to determine for each node $v$ a reliable level $r(v)$
5. during a query, at node $v$
$\square$ do not use edges that have been created in some level $>r(v)$
$\square$ instead, downgrade the search to level $r(v)$ (forward search only)

## Dynamic Highway-Node Routing

change a few edge weights, mobile scenario

reliable levels: $r(x)=0, \quad r\left(s_{2}\right)=r\left(t_{2}\right)=1$


## Dynamic Highway-Node Routing

change a few edge weights, mobile scenario

iterative variant (provided that only edge weight increases allowed)

1. keep everything (as before)
2. $C:=\emptyset$
3. use the sets $A_{u}^{\ell}$ (considering edges in $C$ ) to determine for each node $v$ a reliable level $r(v)$ (as before)
4. 'prudent' query (as before)
5. if shortest path $P$ does not contain a changed edge, we are done
6. otherwise: add changed edges on $P$ to $C$, repeat from 3 .

## Dynamic Highway-Node Routing

change a few edge weights, mobile scenario


|  | affected queries | single pass query time [ms] | iterative |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| \|change set| (motorway edges) |  |  | query time <br> [ms] | \#iter avg | tions <br> max |
| 1 | 0.4 \% | 2.3 | 1.5 | 1.0 | 2 |
| 10 | 5.8 \% | 8.5 | 1.7 | 1.1 | 3 |
| 100 | 40.0\% | 47.1 | 3.6 | 1.4 | 5 |
| 1000 | 83.7\% | 246.3 | 25.3 | 2.7 | 9 |

Sanders et al.: Route Planning
static routing in road networks is easy
$\leadsto$ applications that require massive amount or routing
$\rightsquigarrow$ instantaneous mobile routing
$\leadsto$ techniques for advanced models
$\rightsquigarrow$ updating a few edge weights is OK

## Current / Future Work

$\square$ Time-dependent edge weights challenge: backward search impossible (?)
$\square$ Multiple objective functions and restrictions (bridge height,...)
$\square$ Multicriteria optimization (cost, time,...)
$\square$ Integrate individual and public transportation
$\square$ Other objectives for time-dependent travel
$\square$ Routing driven traffic simulation

