AALBORG UNIVERSITY

Streamspin: Mobile Services for The Masses

Christian S. Jensen

www.cs.aau.dk/~csj

daisy

Center for Data-intensive Systems

Overview

- Web 2.0
- The mobile Internet
- The Streamspin system
- Tracking of moving objects

Web 2.0

- Web 2.0 captures the sense that there is something qualitatively different about today's web.
- Leveraging the collective intelligence of communities
- New ways of interacting

- Sharing of user-generated content
- Text
  - Wiki’s, e.g., Wikipedia
  - Blogs
- Photos
  - E.g., Flickr, Plazes, 23
- Video
  - E.g., YouTube

Web 2.0

- Community concepts abound…
- Feedback and rating schemes
  - E.g., ratings of sellers and buyers at auctions, ratings of content
- Social tagging, tag clouds, folksonomies
- Wiki's
  - Collaborative authoring
- RSS feeds
- Active web sites, Ajax
- Fueled by Google-like business models
  - Google 2006 revenue: USD 10.6 billion; net income: USD 3.1 billion; 12k employees
  - Microsoft now has 8k people in Online Services

Flickr

- From the Flickr entry on Wikipedia

- “In addition to being a popular Web site for users to share personal photographs, the service is widely used by bloggers as a photo repository. Its popularity has been fueled by its innovative online community tools that allow photos to be tagged and browsed by folksonomic means.”

- Launched in February 2004. Acquired by Yahoo! in March 2005. Updated from beta to gamma in May 2006.

- “On December 29th, 2006 the upload limits on free accounts were increased to 100Mb a month (from 20Mb)”

YouTube

- From the YouTube entry on Wikipedia

- “The domain name "YouTube.com" was activated on February 15, 2005…”

- “According to a July 16, 2006 survey, 100 million clips are viewed daily on YouTube, with an additional 65,000 new videos uploaded per 24 hours.”

- “Currently staffed by 67 employees, the company was named *TIME* magazine's "Invention of the Year" for 2006. In October 2006, Google Inc. announced that it had reached a deal to acquire the company for US$1.65 billion in Google's stock.”

Other Video Sharing Sites

Angry Alien, AnimeEpisodes.Net, Blastro, Blennus, Blip.tv, Bofunk, Bolt, Break.com, Castpost, CollegeHumor, Current TV, Dachix, Dailymotion, Danerd, DailySixer.com, DevilDucky, Double Agent, eVideoShare, EVTV1, FindVideos, Free Video Blog, Google Video, Grinvi, Hiphopdeal, iFilm, Keiichi Anime Forever, Kontraband, Lulu TV, Metacafe, Midis.biz, Music.com, MusicVideoCodes.info, MySpace, MySpaceVideo Code, Newgrounds, NothingToxic, PcPlanets, Pixparty, PlsThx, Putle, Revver, Sharkle, SmitHappens, StreetFire, That Video Site, Totally Crap, Video-Codes4U, VideoCodesWorld, VideoCodeZone, vidiLife, Vimeo, vSocial, Yikers, and ZippyVideos

Wiki Systems

@Wiki, bitweaver, BrainKeeper, CanvasWiki, CentralDesktop, Clearspace, ClearWiki, Confluence, Corendal Wiki, Cospire, DekiWiki, DidiWiki, DokuWiki, EditMe, ErfurtWiki, FlexWiki, Friki, GeboGebo, Giki, IkeWiki, ikiwiki, Incentive, Instiki, JAMWiki, JaWiki, JSPWiki, KeheiWiki, KWikiKWiki, LunaWiki, MediaWiki, MicKI, Midgard Wiki, miniWiki, MoinMoin, MoniWiki, Netcipia, nexdo, Oddmuse, OpenWikiNG, PAUX, PBwiki, Perspective, PhpWiki, Pier, Pimki, PmWiki, PodWiki , ProjectForum, ProntoWiki, ProWiki, PukiWiki, Qwik, Riki, SamePage, ScrewTurn Wiki, SeedWiki, SnipSnap, Socialtext, StikiPad, SubWiki, telepark.wiki, TiddlyWiki, TikiWiki, TracWiki, TWiki, UniWakka, UseMod, VQWiki, WackoWiki, Wala Wiki, Wetpaint, Wiclear, Wikepage, Wiki-Toolkit, Wikia, WikiASP, WikiDoc, Wikidot, WikiSH, Wikispaces, WikkaWiki, WikyBlog, wxWikiServer, XWiki, yawiki, Zwiki

http://www.wikimatrix.org/

Web 2.0

- Another perspective: The technologies needed to be competitive in the economy of free services!

- “Free trumps quality all the time.” [Nicholas Carr]
- The devices may even become free!
- Some companies may even offer “free” mobile devices.
  - According to Google’s CEO Eric Schmidt, one day your phone could be free, subsidized by the growth of mobile advertising.

Overview

- Web 2.0
- The mobile Internet
- The Streamspin system
- Tracking of moving objects

The Internet Is Going Mobile

- We are at a unique point in history.
  - The Internet is just about ready to go mobile.
  - The mobile Internet has the potential for having more users than the conventional Internet.
- A mobile Internet infrastructure is emerging.
  - Mobile devices, e.g., mobile phones, PDAs, laptops, cameras, MP3 players, navigation systems, etc.
  - Communication networks, e.g., GPRS, EDGE, 3G, HSDPA, Wimax
  - Users with access
- Technologies are becoming practical/available that enable the accurate geo-positioning of all objects we care about.
  - The emerging network-assisted GPS reduces power consumption.
  - Galileo is underway.

Mobile Is Important

Mobile phones are cheaper than PCs, there are three times more of them, growing at twice the speed, and they increasingly have Internet access. What is more, the World Bank estimates that more than two-thirds of the world's population lives within range of a mobile phone network. Mobile is going to be the next big Internet phenomenon. It holds the key to greater access for everyone - with all the benefits that entails.

Eric Schmidt

CEO, Google, [Financial Times, May 2006]

Service Types

- Traffic and traffic-management related services
  - Emergency vehicle dispatching
  - Road pricing generalized: payment based on where, when, and how much one drives; taxes, insurance
  - Spatial pay per use, or metered services
- “Safety”-related services
  - Tracking of hazardous cargo
  - Warnings about accidents, slow-moving traffic ahead, icy or slippery road conditions
  - Monitoring of traffic offenders
  - Monitoring of tourists traveling in dangerous environments, reacting to emergencies

Service Types, cont.

- Games and ”-tainment” (edu-, info-, enter-)
  - Treasure hunting (geocaching)
  - Paintball (Botfighters)
  - Catch the monster (Raygun)
  - Escape the monster
  - Tell me about that!

Location-based Games: The Idea

- Move games from going on behind a small computer or phone screen into reality.
- Virtual objects, seen by the players on their displays, are given physical locations that are know to the system.
- Physical objects, the players, are being tracked by the system.

- The boundaries between reality and fantasy are being blurred.

- Mobile games brought in USD 74 million in the US in 2004. [IEEE Spectrum **43**(1), Jan 2006]
- 2009 projection: USD 430 million.

Raygun

- By GloVentures LLC, Redmond, WA
- Idea: Pac-Man, but with the player being the joystick!
- Ghosts are displayed as tiny colored dots on the phone's "geographical" display.
- Uses GPS for tracking the player.
- The digital world is embedded into the real world.

- GPS Version of Snake uses the same setup as Pac-Man

Raygun

- GPS games get players off their couches and into the real world

  ‘It's a drizzly Sunday in Marymoor Park, a leafy hangout for soccer kids and Ultimate Frisbee jocks, in Redmond, Wash. This afternoon, however, a new breed of outdoor enthusiasts has taken to the field.

  A half-dozen people are wandering the grounds while holding their cellphones at arm's length. They move in urgent and idiosyncratic trajectories, shifting directions on the fly without peeling their eyes from their phones. For the drivers on the highway nearby, it is an unusual sight: a group of oddballs apparently roaming in the rain for reception. But, in fact, the wanderers are not muttering, "Can you hear me now?" They're playing Raygun.

  The conceit is that you're hunting for ghosts. The phone displays a sort of supernatural radar screen that tracks surrounding ghosts as tiny colored dots. The object is to gobble up the dots before they get you. In a way, it's a little like Pac-Man, but with one key difference, as James Robarts of GloVentures LLC, in Redmond, Wash., the developers of the game, puts it: "The joystick is you.“’

  [IEEE Spectrum **43**(1), Jan 2006]

Mobile Is Different

- The conventional Internet
  - Computers with large screens and convenient qwerty keyboards
  - In controlled environments, at home or at work
- The mobile Internet
  - Small screens, inconvenient keyboards
  - The user is out and about – yields high variation in use situations
    - In a meeting or at a café
    - On the move, e.g., on foot, using collective transport, driving a car
  - Disruptive surroundings
  - Service use is often not the primary activity
    - Assist the user in accomplishing the primary activity
  - Push services
  - Delivery of the right service at the right time is important.

Context Awareness

- Context awareness will be important!

- Demographic user data
  - Age, gender, marital status, job, etc.
- Users may define profiles that may be (de-)activated
  - Interests and preferences
  - Subscriptions
- A user's social network
  - Friends, colleagues
- Geo-context
  - Current location (and speed)
  - Destination and route for users on the move

- Ranges from static to dynamic; ranges from user supplied to automatic

Overview

- Web 2.0
- The mobile Internet
- The Streamspin system
- Tracking of moving objects

Streamspin in a Nutshell

- Vision:

  To create data management technology that enables sites that are for mobile services what Flickr is for photos and YouTube is for videos.

- Challenges
  - Enable easy mobile service creation
  - Enable service sharing with support for community concepts
  - An open, extensible, and scalable service delivery infrastructure

- The streamspin project maintains an evolving platform that aims to serve as a testbed for exploring solutions to these challenges.

Service Creation

- Streamspin- or user-provided *templates* are available for service creation
  - Point-and-click service creation
  - Example templates: tour builder, e-mail, RSS push
- Streamspin-provided *web services* are available for creating custom services and templates
  - Content publishing
  - Service creation
  - Current location context for a user, using call back
  - Destination and route context, using call back (pending)
- Visual Studio *C# add-ins* for custom service creation
- Accessible from all web-service enabled languages

Example: Tour Services

- Users can create tours.
  - Tours are created using a point-and-click tour builder template.
  - Tours associate content (e.g., photos and text) with locations.
- Users can subscribe to tours.
  - They receive content when they get within a specified range of the locations associated with the content.
- Example tours
  - Walking and driving tours with directions
  - Tours that involve public transportation can tell their users when to get off busses in real time

Create Tour

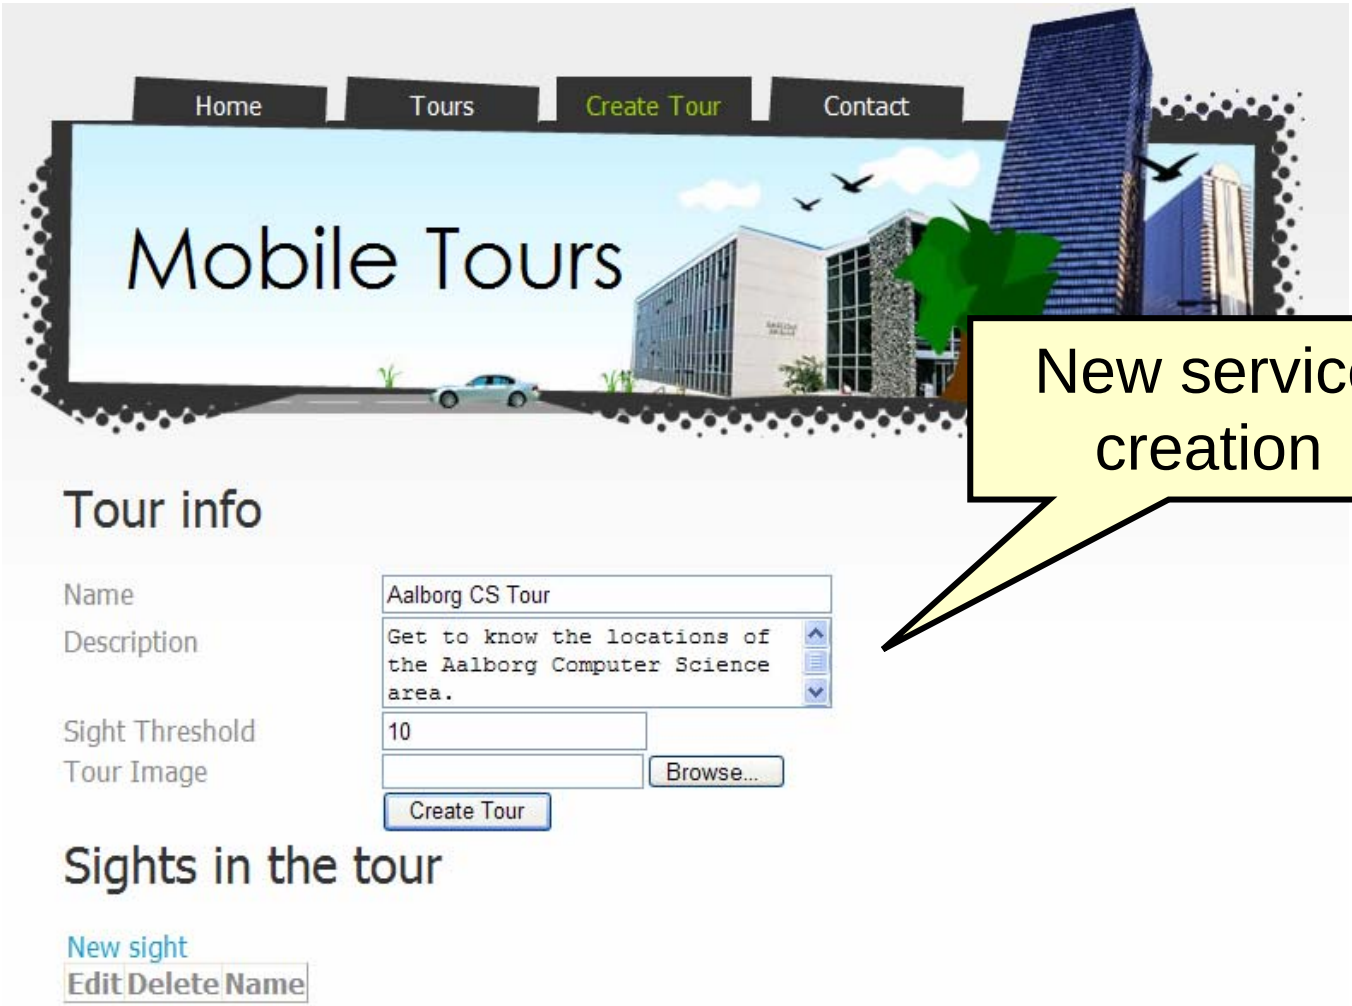
Home
Tours
Create Tour
Contact
Mobile Tours
New service creation
Tour info
Name
Aalborg CS Tour
Description
Get to know the locations of the Aalborg Computer Science area.
Sight Threshold
10
Tour Image
Browse...
Create Tour
Sights in the tour
New sight
Edit Delete Name

Tour Image


Home
Tours
Create Tour
Contact
Mobile Tours
Tour info
Name
Description
Sight Threshold
Tour Image
Sights in the t
New sight
Edit Delete Name
Choose file
Look in:
My Pictures
My Recent Documents
Desktop
My Documents
My Computer
My Network Places
Hillerbilly
hp
icon_discuss3
icon_discuss5
images
imgres
File name:
Files of type:
All Files (*.*)
Open
Cancel

Tour Sight

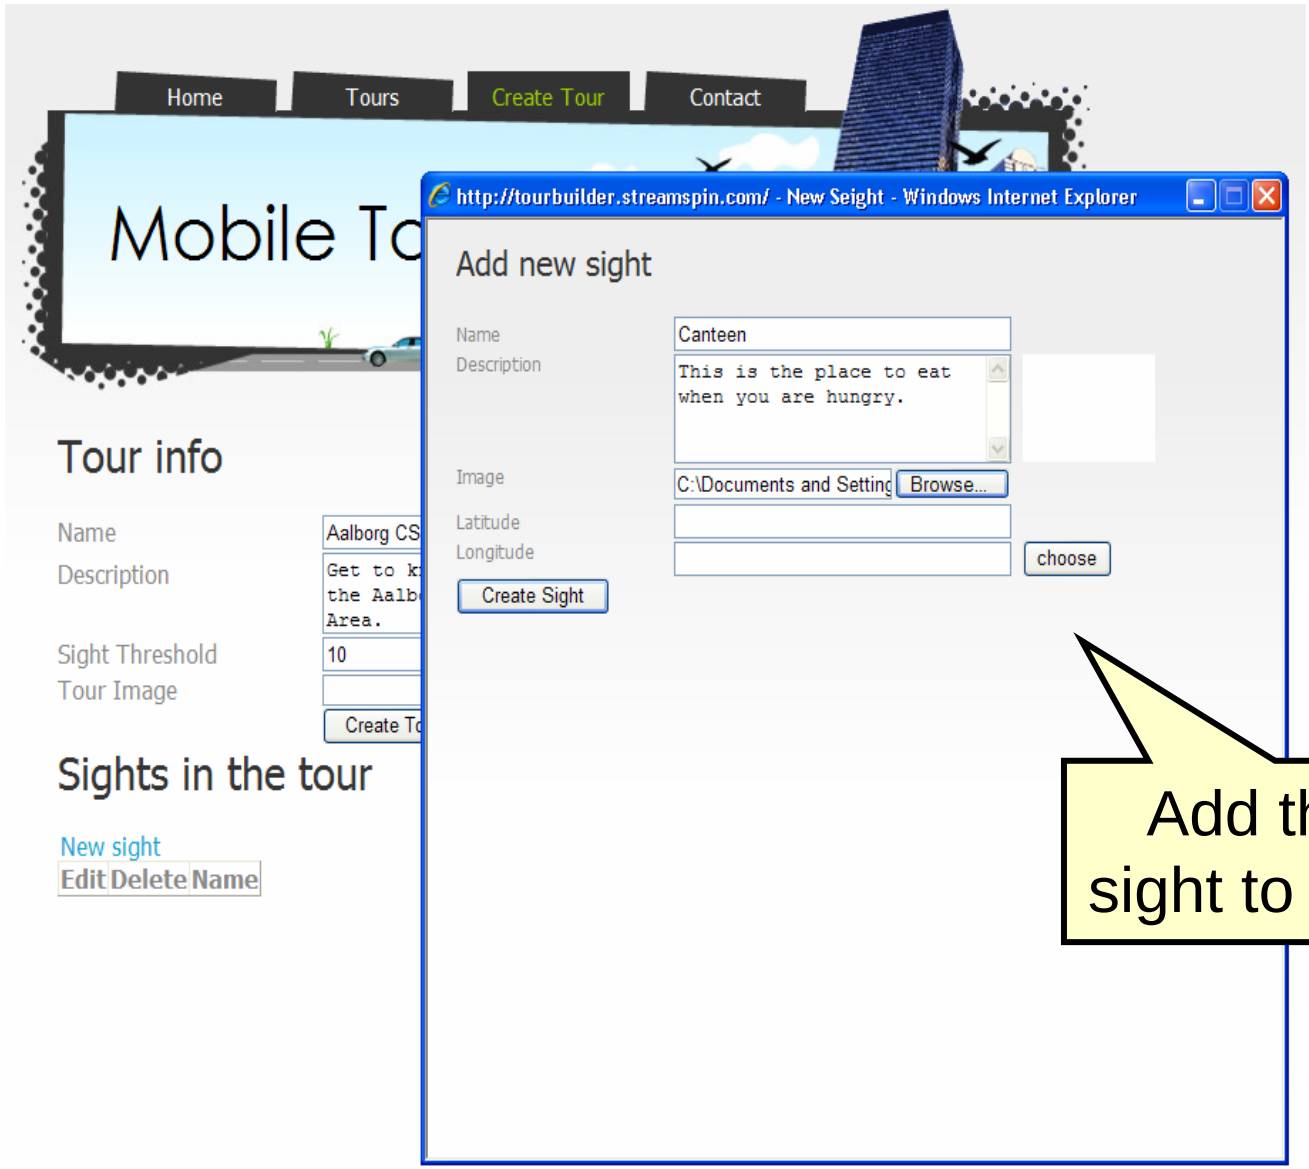
Home
Tours
Create Tour
Contact
Mobile To
Tour info
Name
Aalborg CS
Description
Get to k
the Aalb
Area.
Sight Threshold
10
Tour Image
Create To
Sights in the tour
New sight
Edit Delete Name
http://tourbuilder.streamspin.com/ - New Seight - Windows Internet Explorer
Add new sight
Name
Canteen
Description
This is the place to eat
when you are hungry.
Image
C:\Documents and Setting
Browse...
Latitude
Longitude
choose
Create Sight
Add the first sight to the tour.

Sight Location

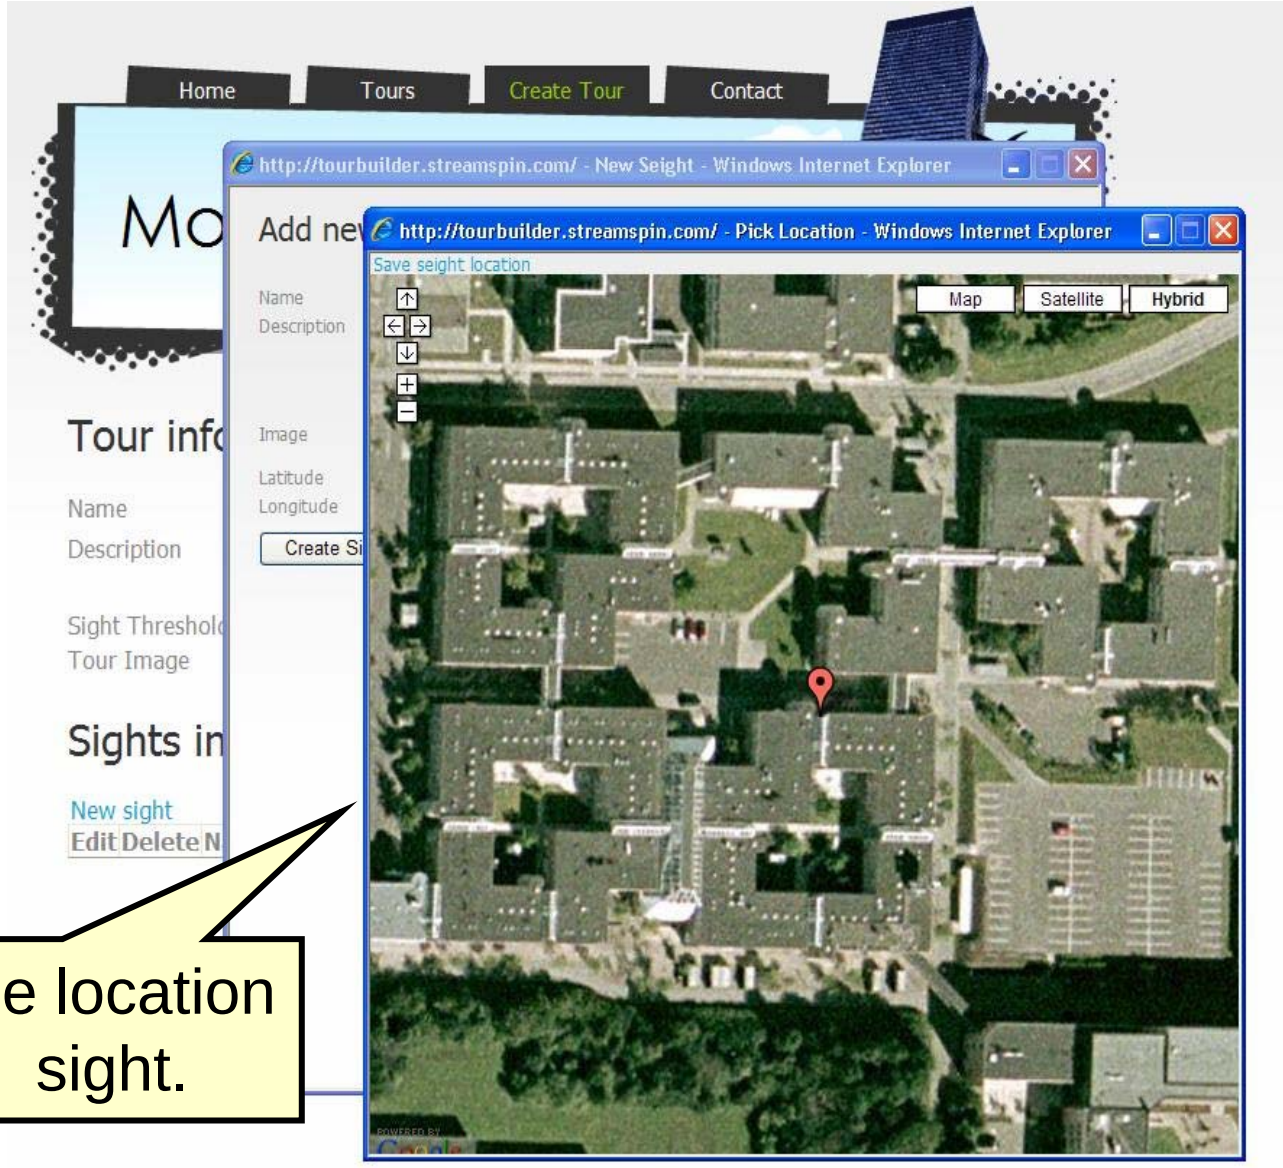

Select the location of the sight.

The Final Tour

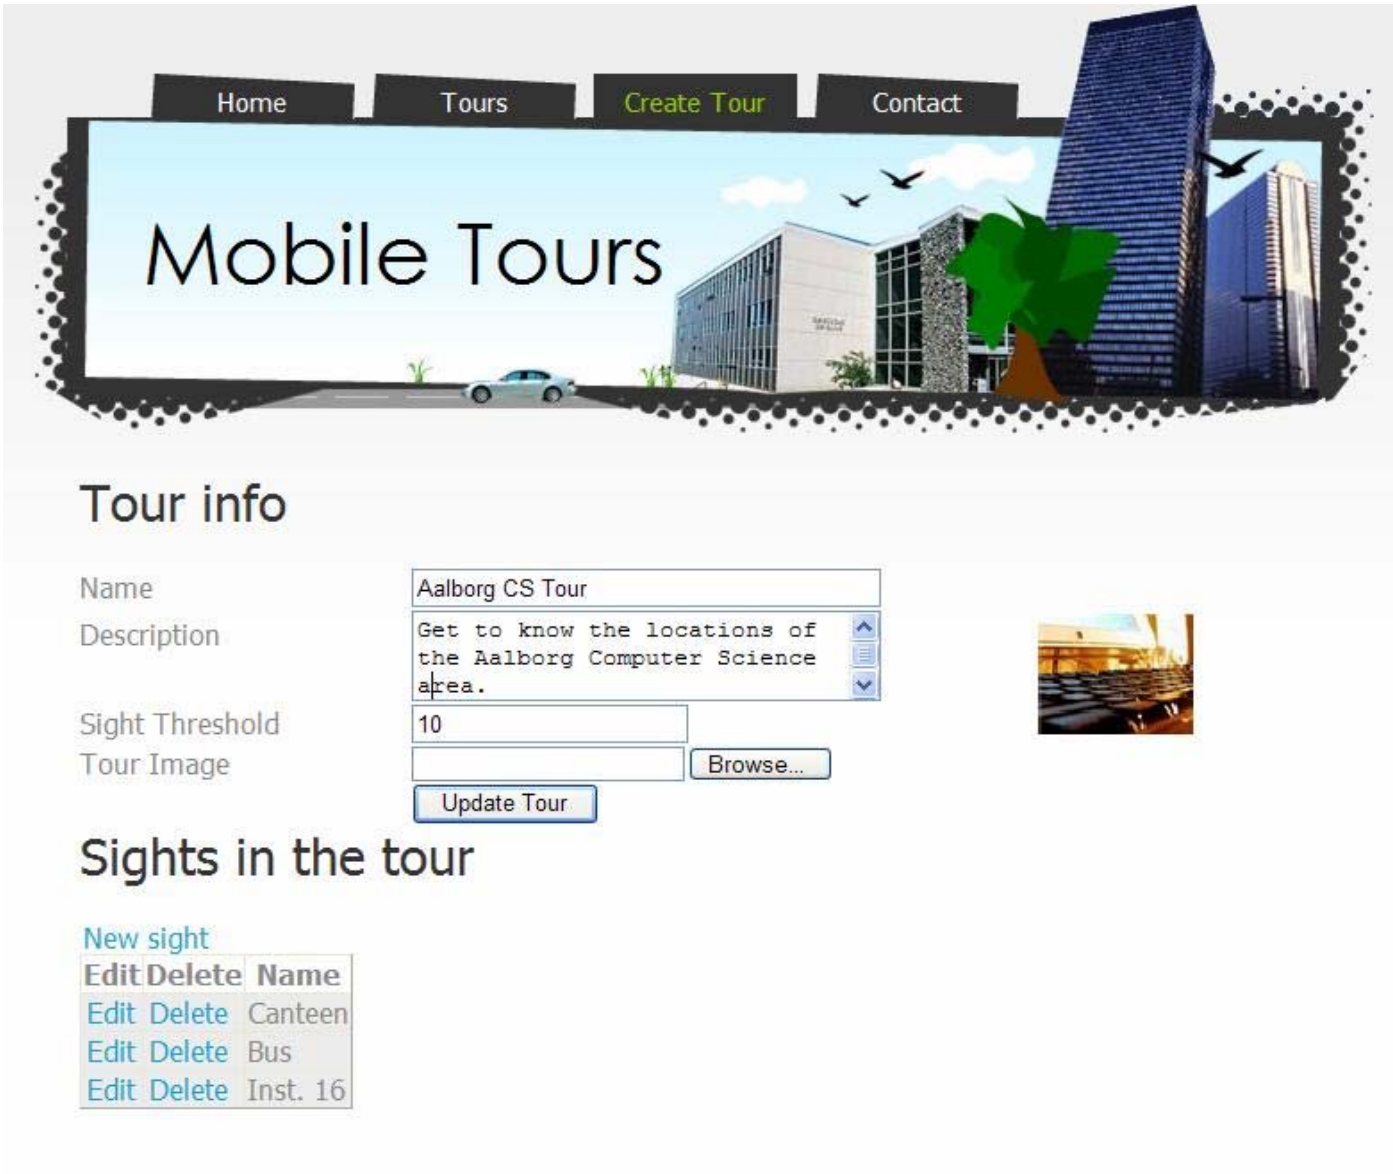

Example: Gasoline Service

- Services benefit from geo-context awareness.

- No current location awareness
  - The user is notified when a gas station near their normal location offers gas at least 10% below the “list” price.
- Current-location awareness
  - The user is notified when within 3 km of cheap gas.
  - Close-by gas
- Location- and route-awareness
  - The user is notified when cheap gas is close to the route ahead.
  - Cheaper gas and smaller detours.

Service Sharing

- Public content providers
  - Approval
  - Meta data: interest profile, location, location range, time to live, age range
- Publication of services in a service directory
- Interest hierarchy
  - For tagging of content by public providers
  - For specification of profiles by service users
- Content rating
  - Recipients of content can rate the content
- Content discussions
  - Recipients of content can comment on the content and see the comments provided by other users
- Friends

Service Delivery

- Filtering of public content
  - Based on the user's context and the meta data of the content
  - Publish/subscribe functionality
- Socket-based content push
  - Maintains socket-based connections to mobile clients.
  - No HTTP or web-service overhead (only TCP headers)
    - Text messaging at 2.5% of the normal cost, using current text messaging and GPRS pricing from DK
- Support for content that consists of a text header and a text (html) or URL body.
  - Text is used if the content is text.
  - Otherwise, a URL is provided that the client can then access.

Overview

- Web 2.0
- The mobile Internet
- The Streamspin system
- Tracking of moving objects

Problem Setting

Aim: To track moving objects with accuracy guarantees

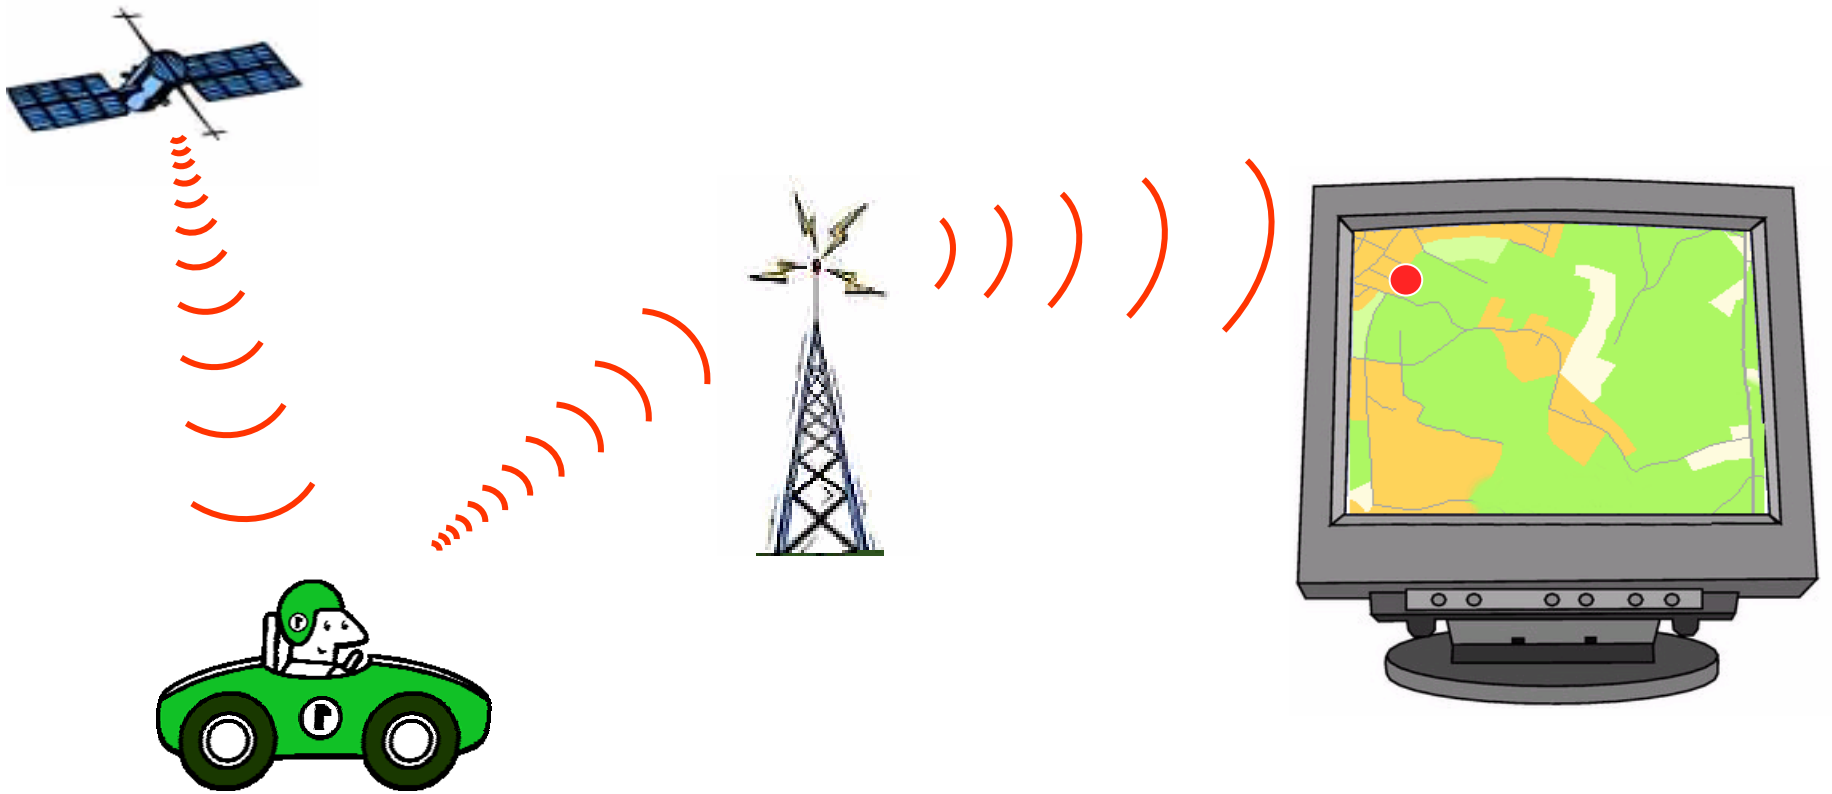

Objective: To reduce cost of communication between client and server and server-side update, client-side costs

Tracking Approach

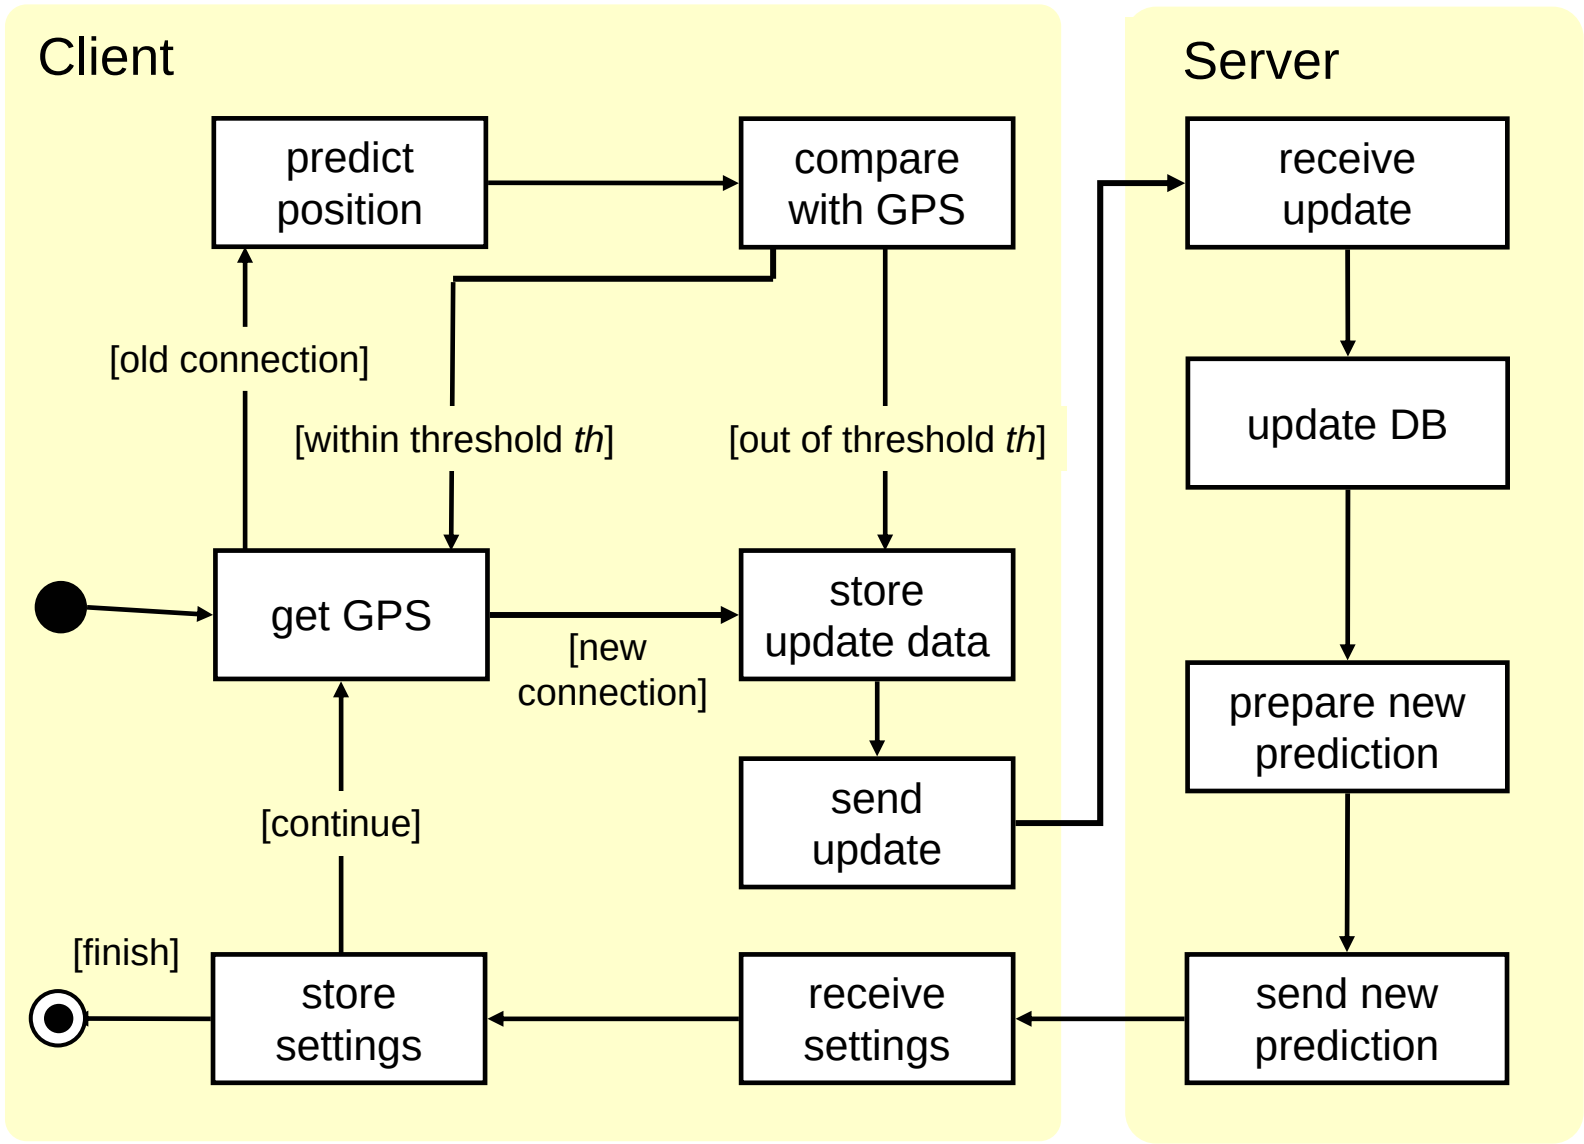

Point-Based Tracking

- The predicted position at time *t* is the most recently reported position, (*x*,*y*).
- An update occurs every time the agreed-upon threshold *th* is exceeded.
  - For “directed” movement, this is every time the vehicle has traveled *th* distance units.
  - Great for “undirected” movement, e.g., a person playing football while being tracked with a threshold of, say, 100m.
  - Great for non-moving objects.
  - Robust.
  - An important building block.

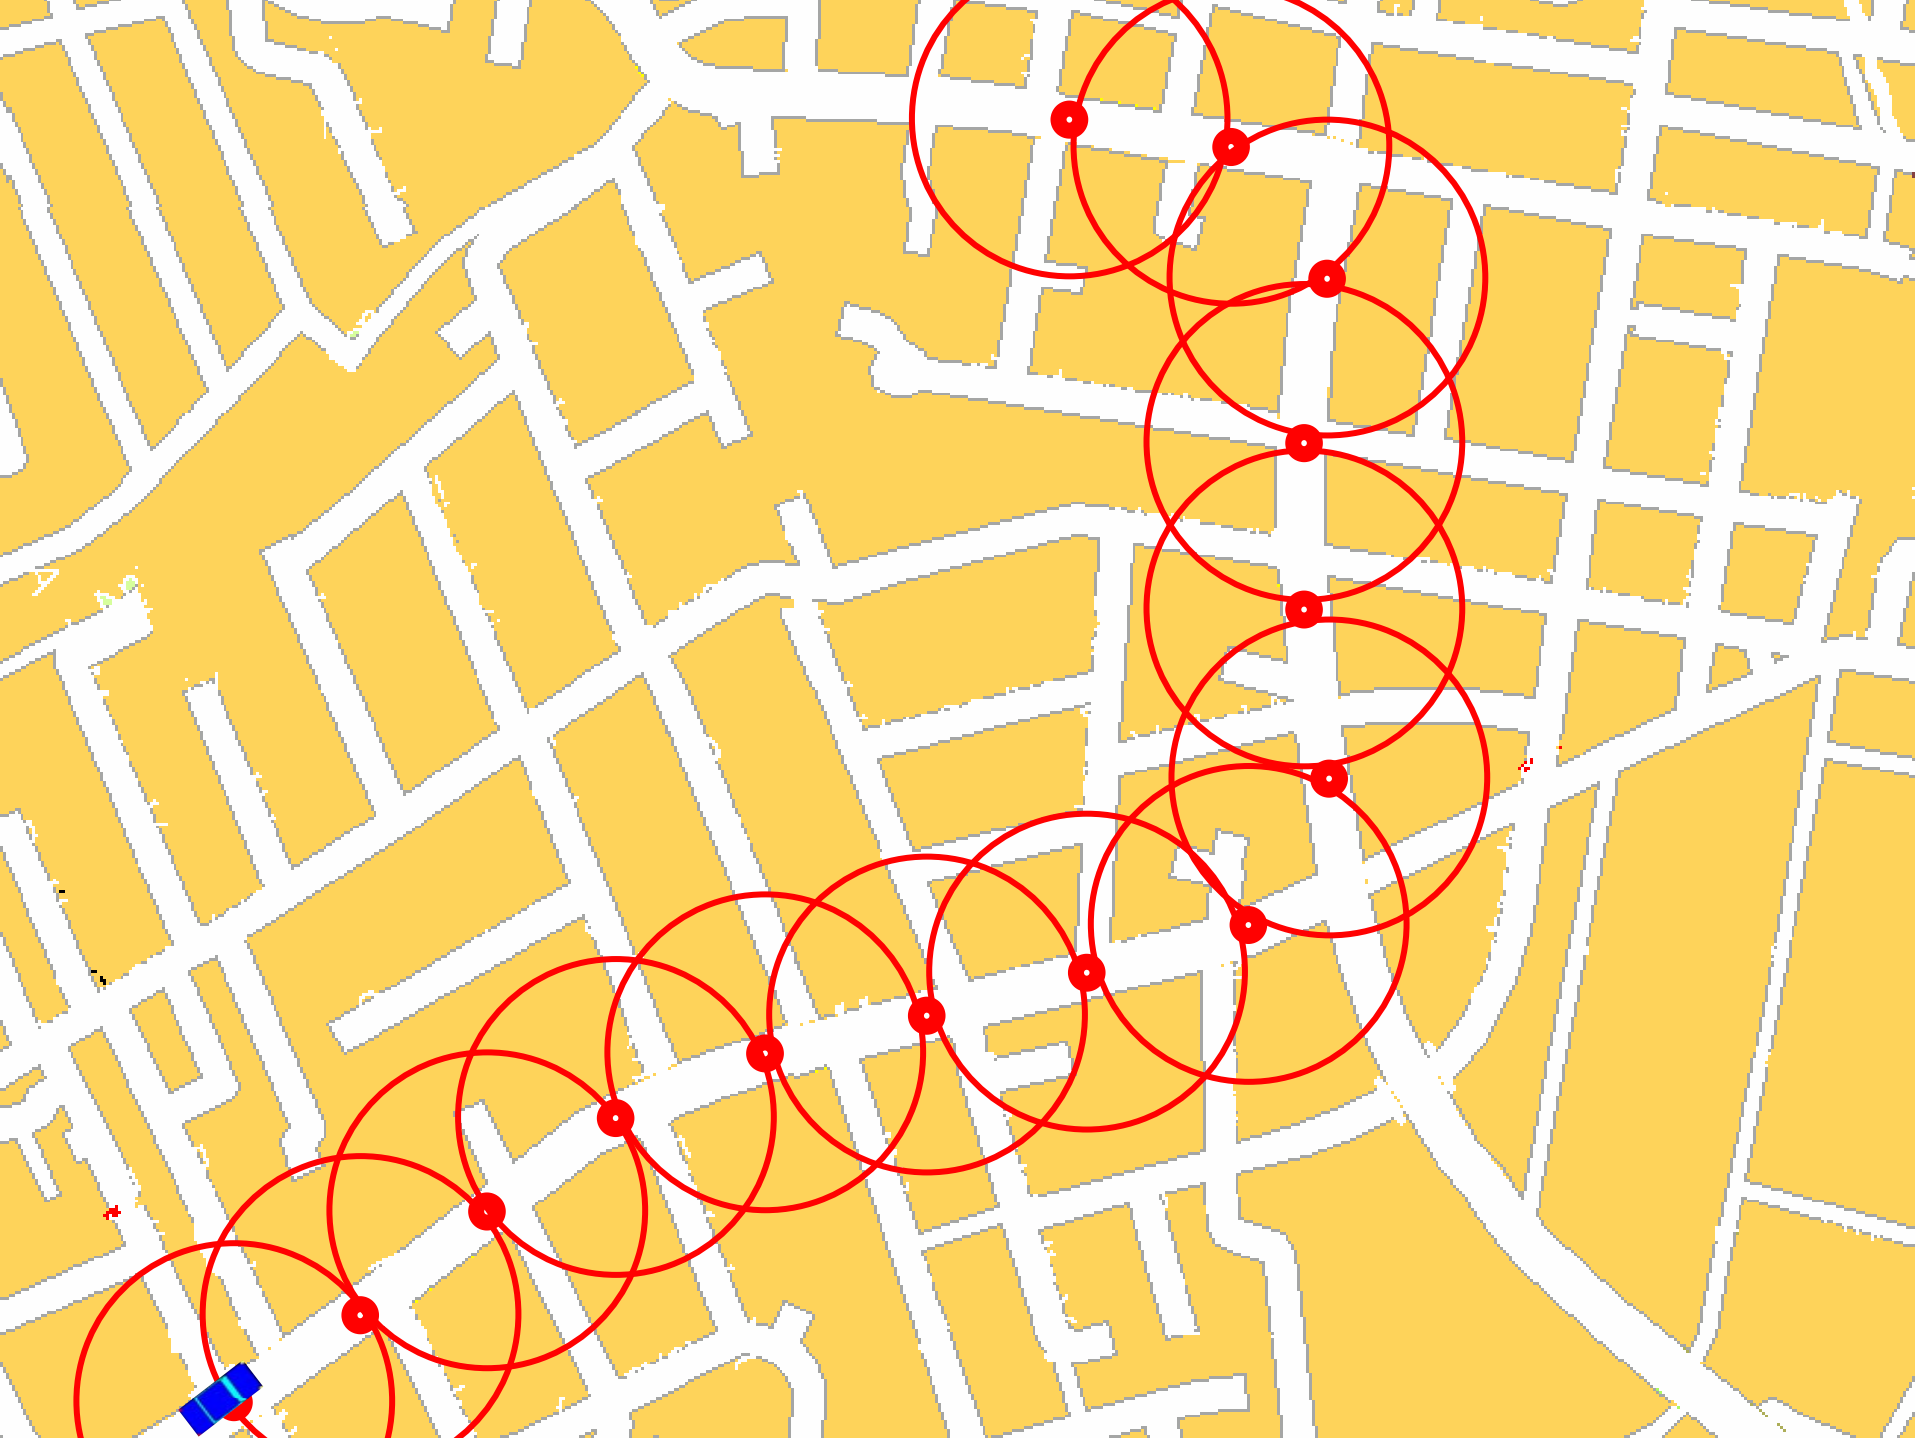

Vector-Based Tracking

- The predicted position at time t is

$$p(t) = p(t_i) + v(t_i)\,(t - t_i)$$

  where t_i is the most recent sample not after time t.
- Updates again occur every time threshold *th* is exceeded.
  - For “directed” movement, this is should be better than point-based tracking.
  - May be considered a generalization of point-based tracking.
  - Robust.
  - An important building block.

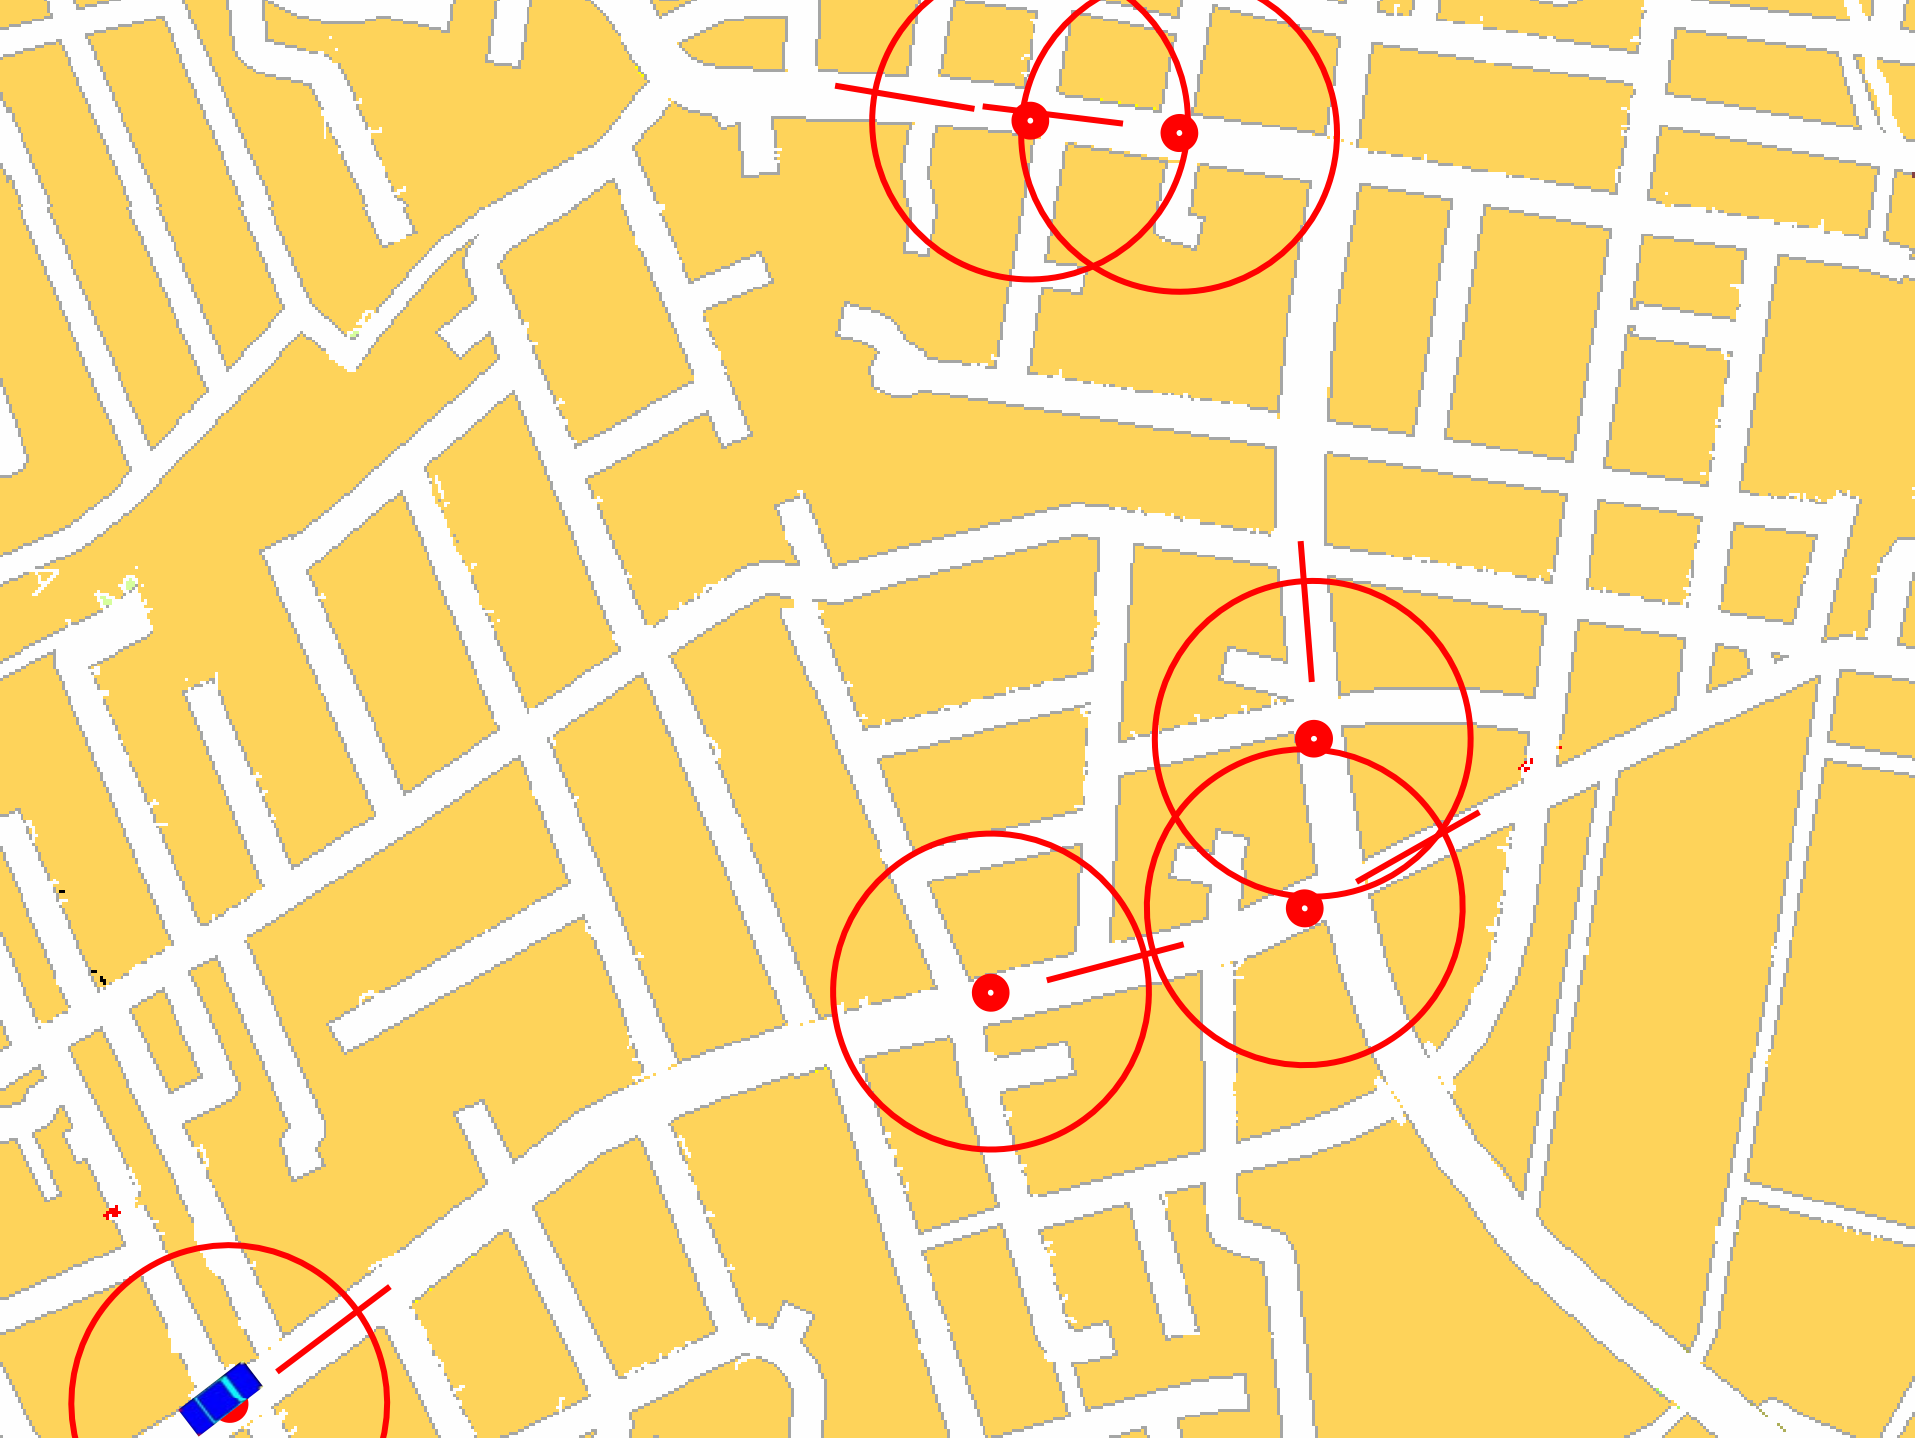

Segment-Based Tracking

- A road network is assumed that consists of segments, which are polylines.
- A sample reports the current segment, distance from the start of the segment, and a speed.
- Using the most recent sample, the predicted position at time *t* is given as a distance from the start of the current segment, assuming constant speed.
- Positions can be given in Euclidean coordinates.
- Object movement follows the shapes of the roads.

- Requires a map and successful map matching.
- Uses vector-based tracking when map matching fails.
- Uses point-based tracking when reaching segment ends.

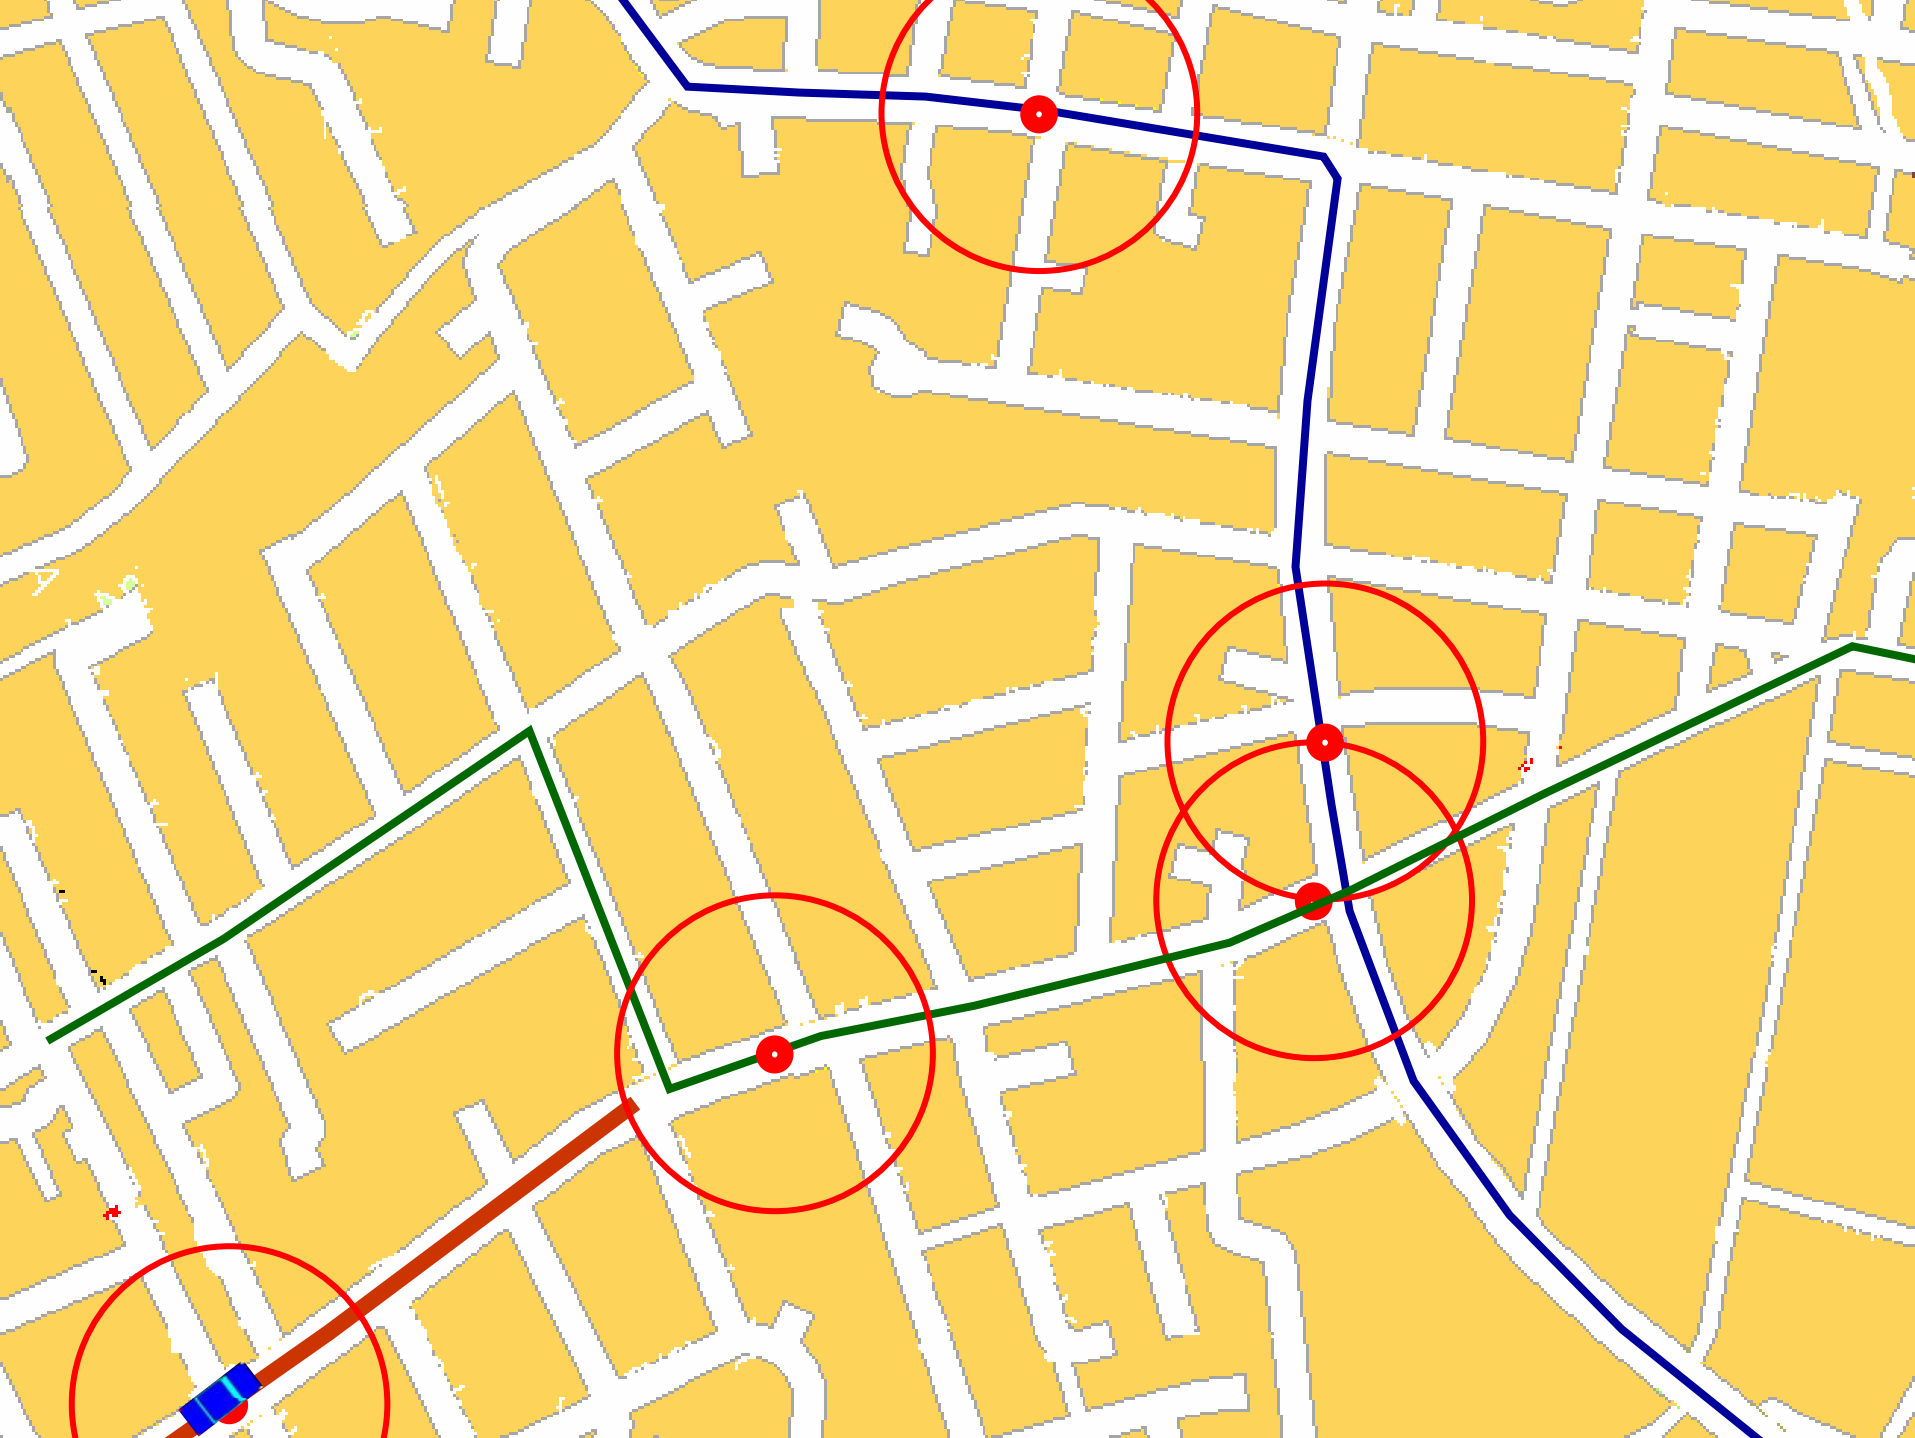

Two Implementations

- Centralized, Oracle-based implementation
  - Ideal for testing implementations of the algorithmic aspects of the techniques.
  - Well suited for simulation-based experiments with pre-recorded data.
- Real implementation
  - Involves a central server, mobile terminals, GPS receivers
  - More complex than the centralized implementation
  - Enables more detailed cost modeling, e.g., of data transmission cost and server and client side loads
  - Offers insight into the specifics, e.g., network delays
  - Offers the ultimate poof of concept

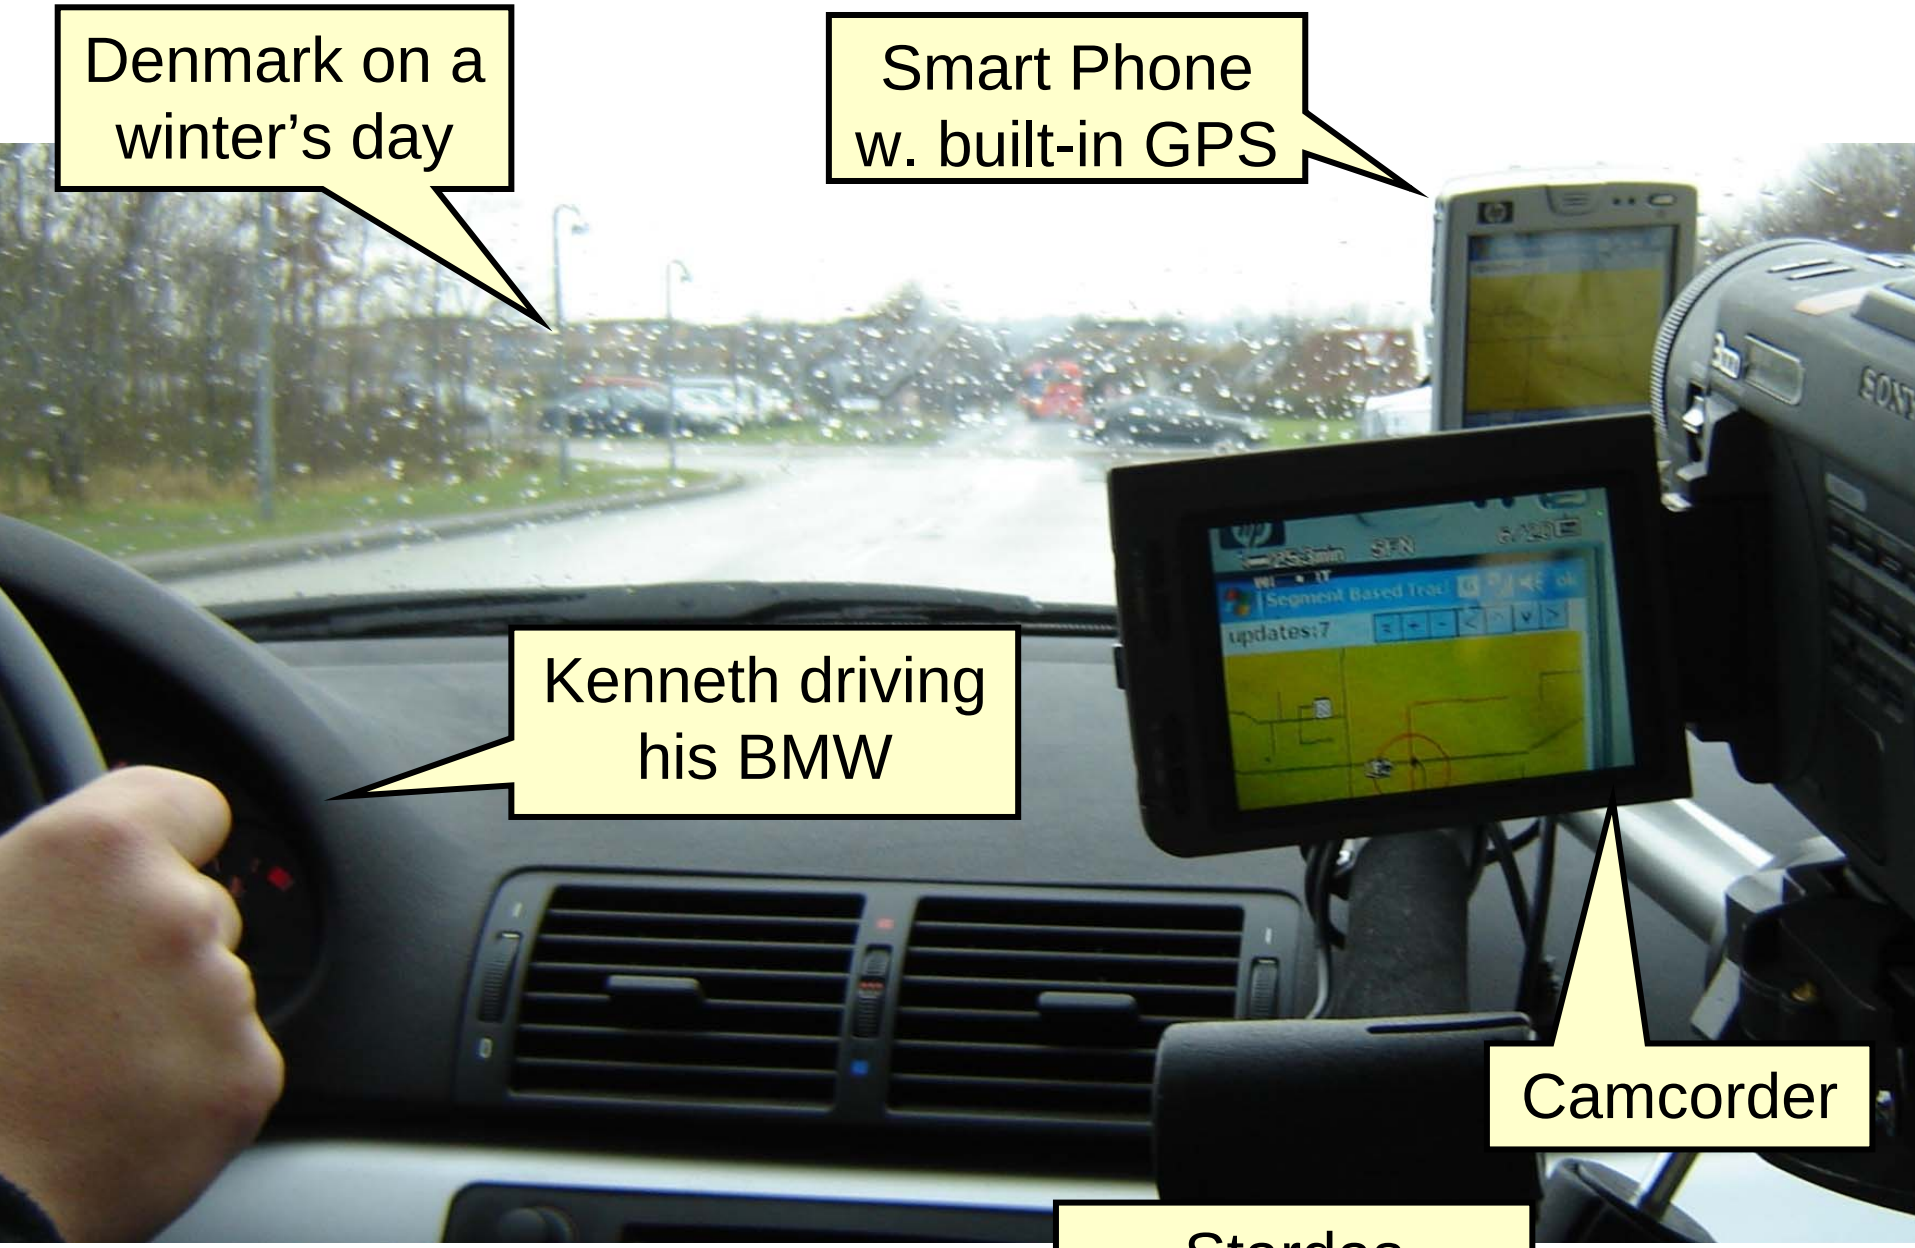
Denmark on a winter's day
Smart Phone w. built-in GPS
Kenneth driving his BMW
Camcorder
Stardas holding it all

skip

Data for Experiments

- GPS Data – the INFATI data is used for evaluation
  - GPS receivers and computers installed in cars
  - GPS coordinates are registered every second for ~6 weeks
  - The data used has ~100,000 records per car and ~458,000 in total
- Digital Road Network
  - Each segment corresponds to the road in-between two crossroads
  - The geometry of a segment is represented as a polyline

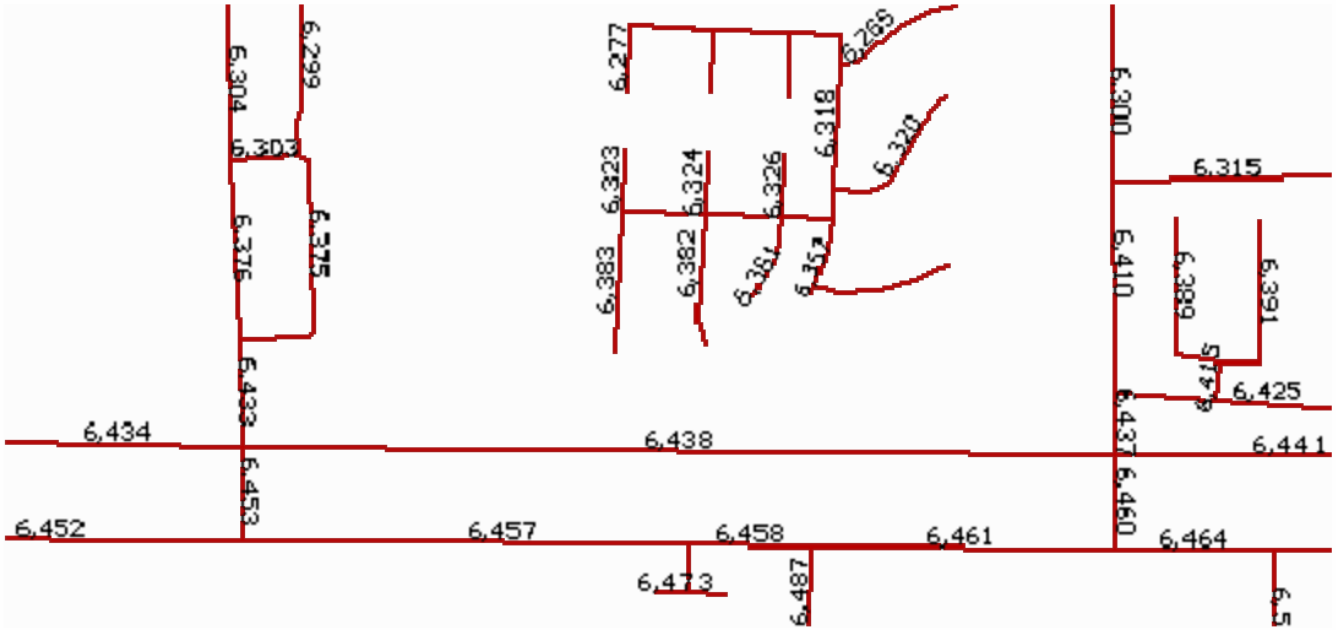

Map Matching and Inaccuracies

- Map matching is used in segment based tracking
- Non-trivial due to GPS and digital road network inaccuracies

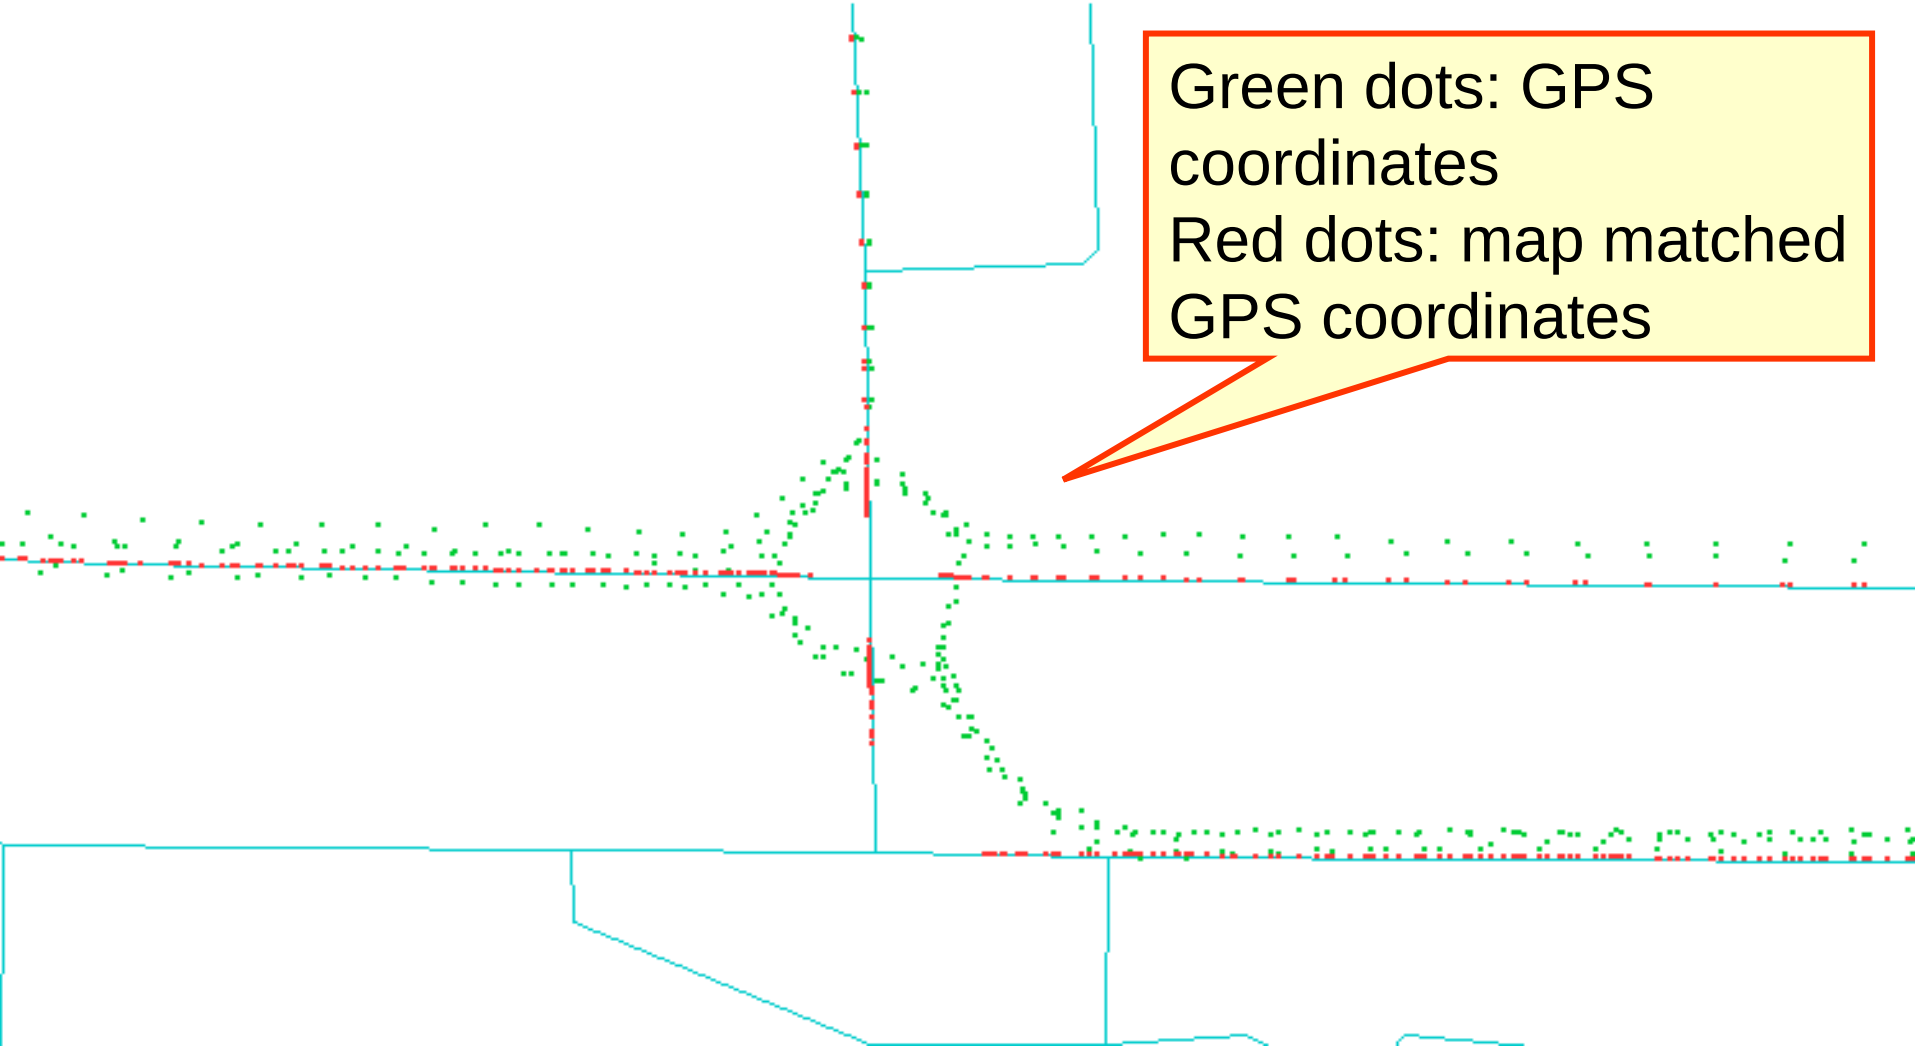

Comparison of Techniques

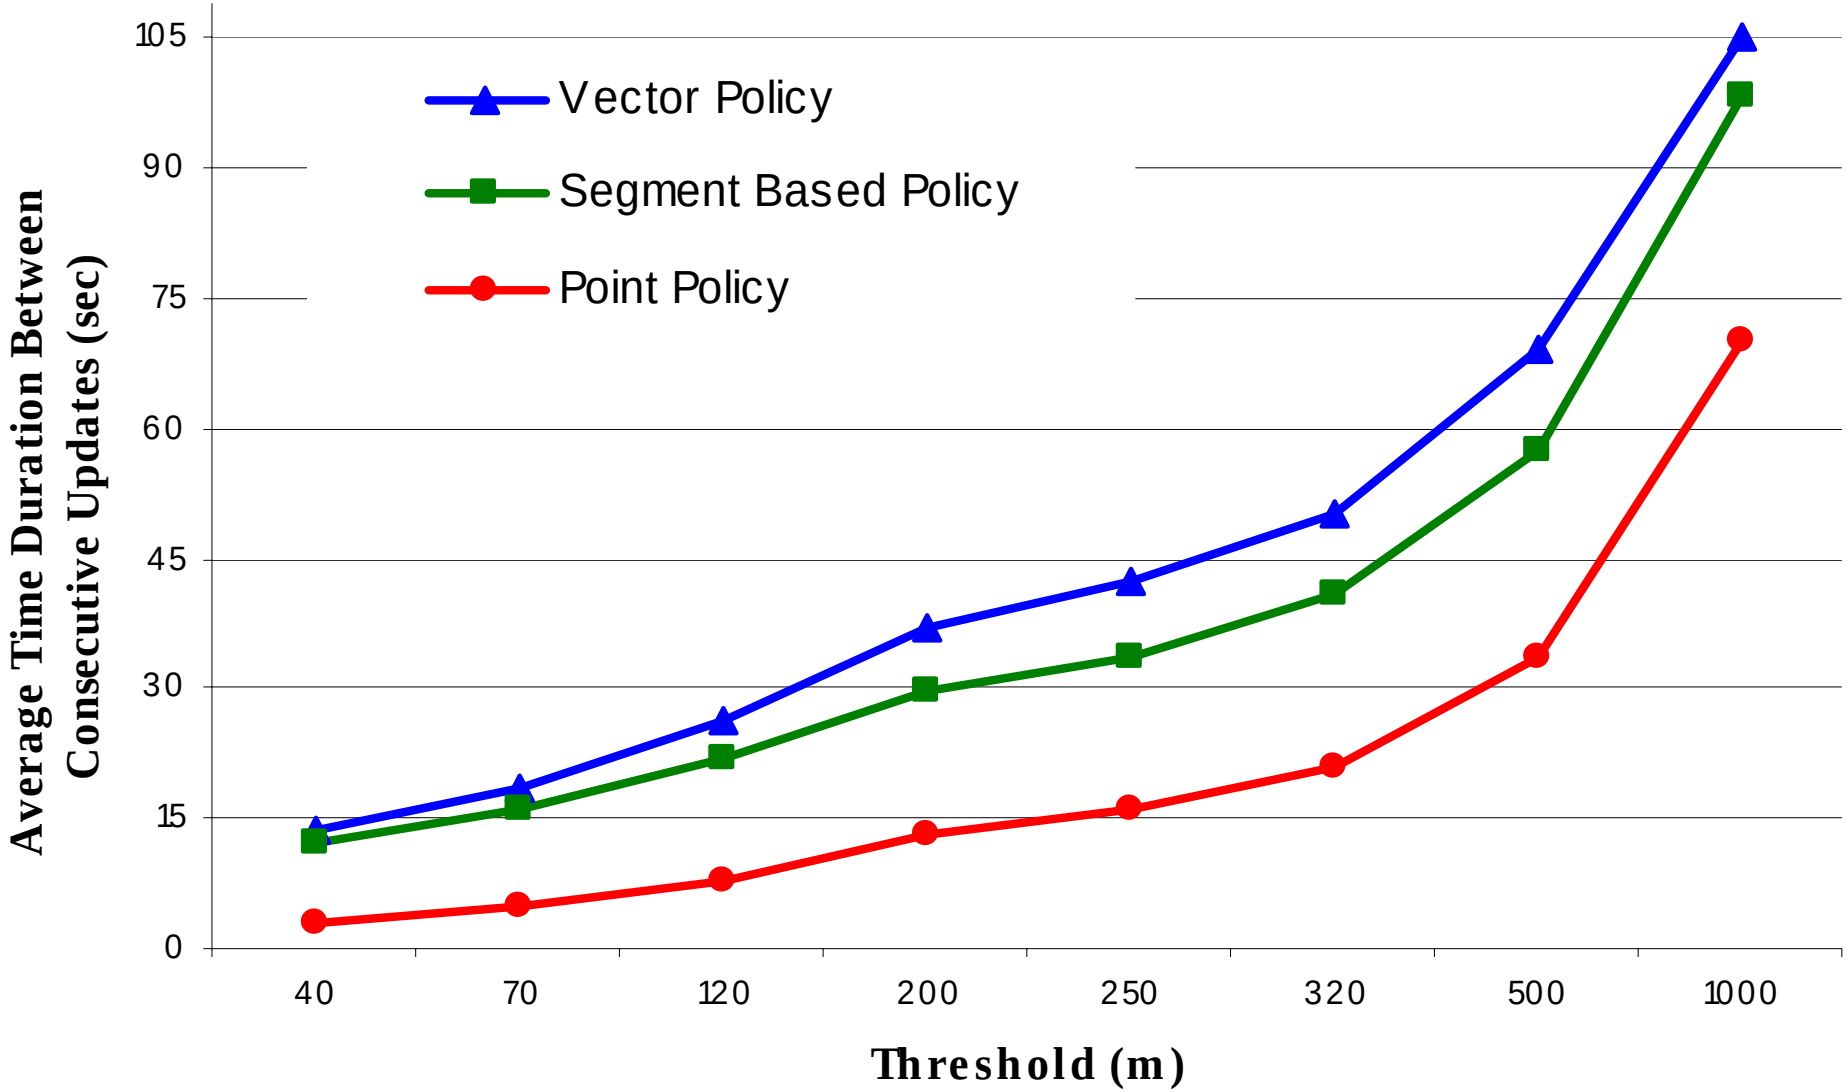

Network Re-Segmentation

- Goal: Create longer segments so that there are as few segment changes as possible.

- StreetID based modification
  - Connect segments with the same StreetID
  - Try create segments that are as long as possible.
  - Segments then tend to correspond to named streets.
- Tails based modification
  - Distinguishes between main streets and side streets, termed tails
  - Defines a tail level for each segment and gives preference to segments with high tail level
- Direction based modification
  - Assumes that most vehicles tend to go as direct as possible towards their destinations and thus tries to obtain straight segments

Use of Routes

- Users follow routes to reach their destinations.
- If we know the current route of a user, we can avoid segment changes altogether.
- As routes are (long) segments, segment-based tracking works.
- Routes may be obtained via a navigation system or a route acquisitioning and provisioning component (next!).

Results – Use of Routes

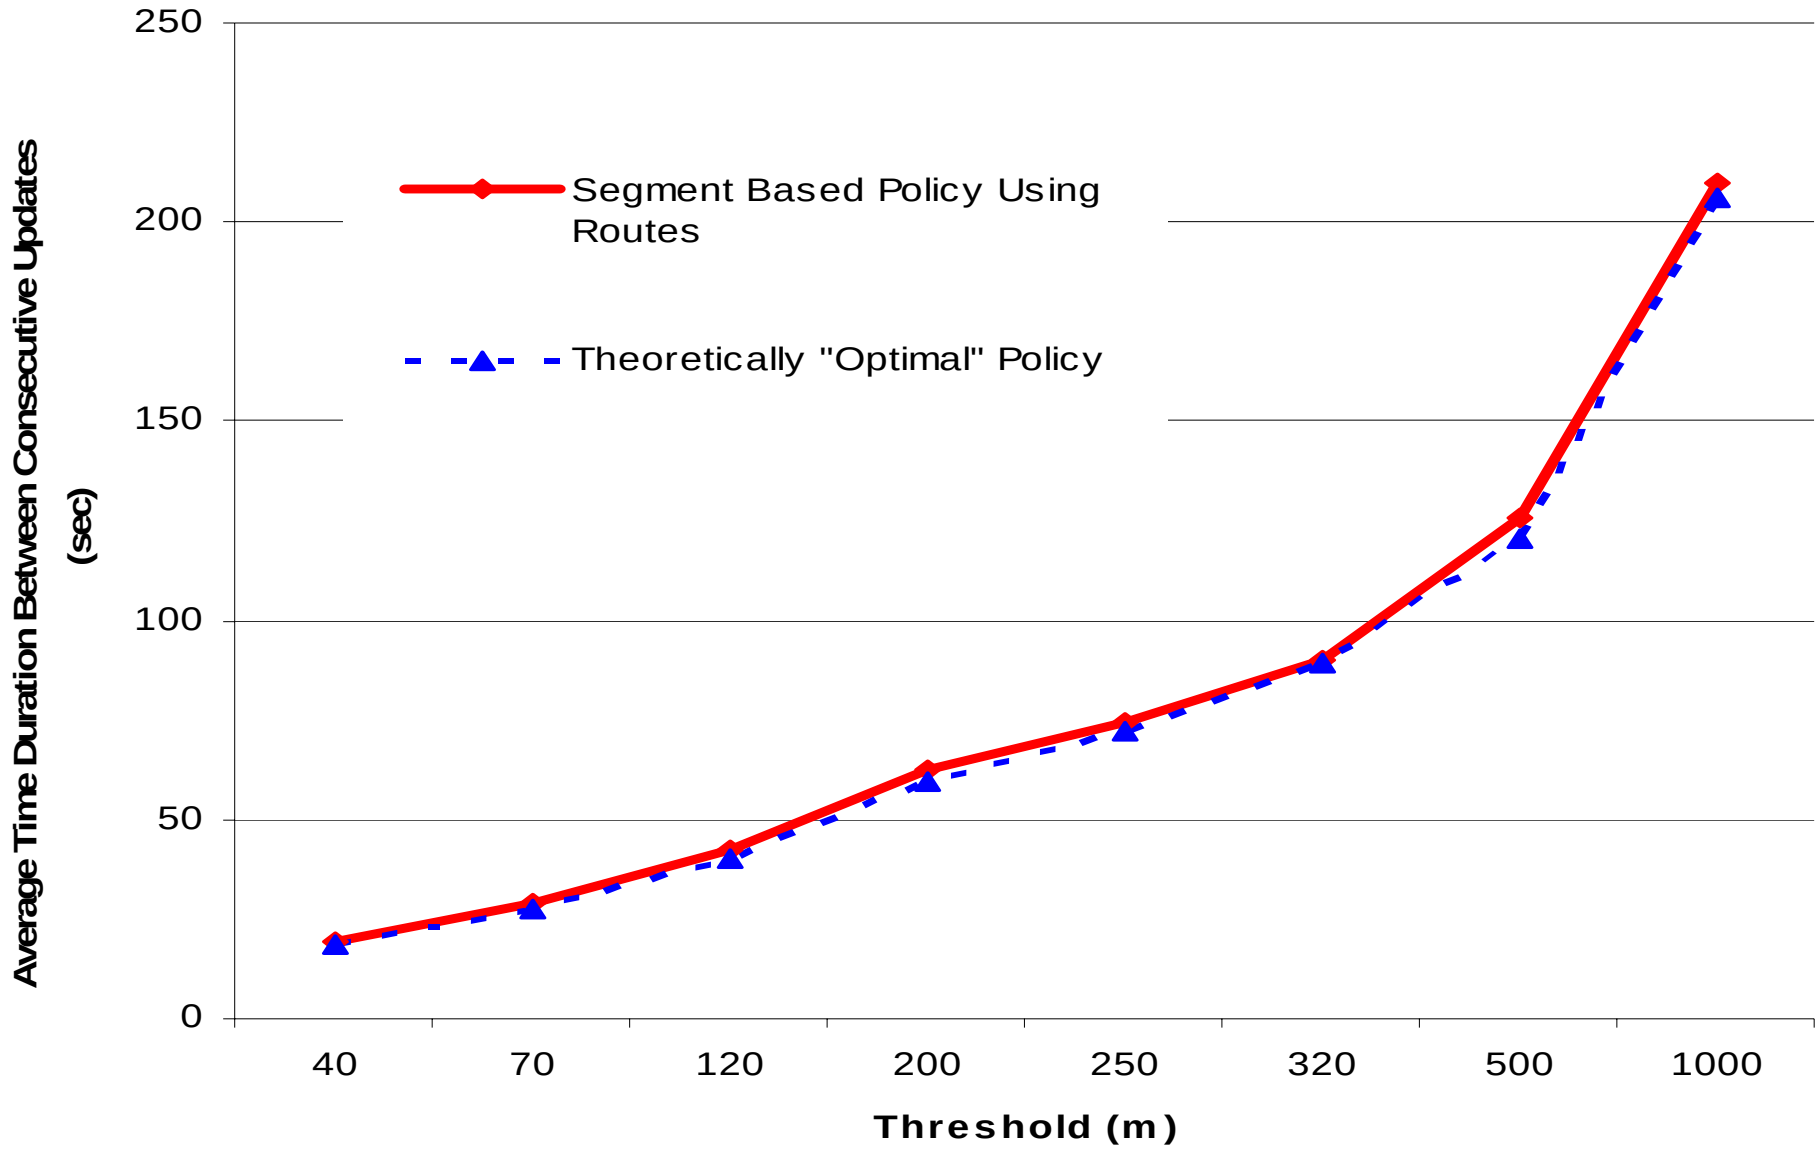

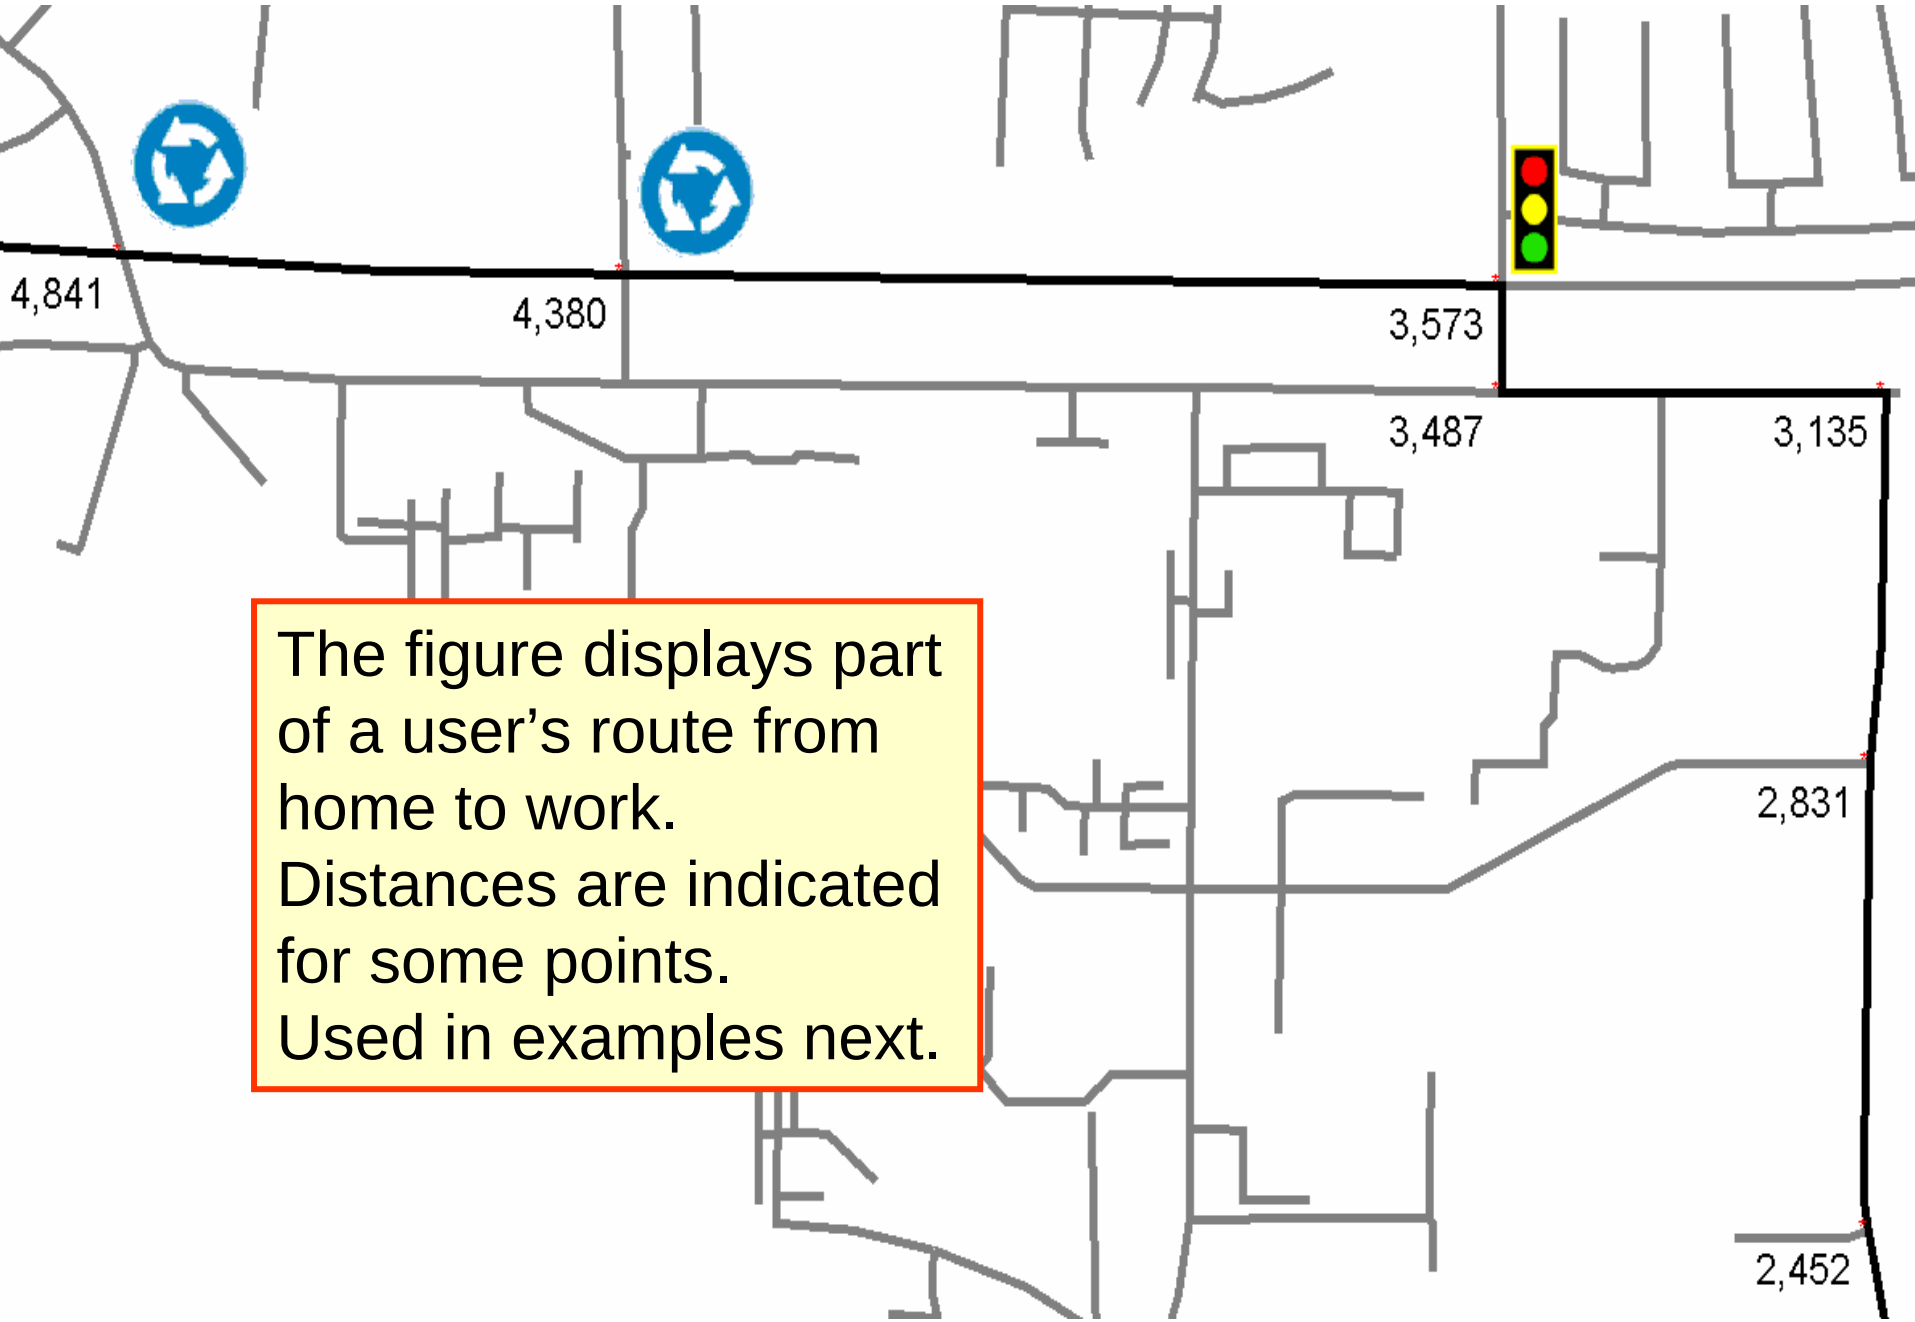

The figure displays part of a user's route from home to work.
Distances are indicated for some points.
Used in examples next.

Use of Acceleration Profiles

- Repeated route traversals exhibit a clear speed pattern.
- An acceleration profile is created for each route
  - Distance intervals with positive and negative acceleration are found using *average* speeds.
  - An average acceleration is calculated for each such interval.

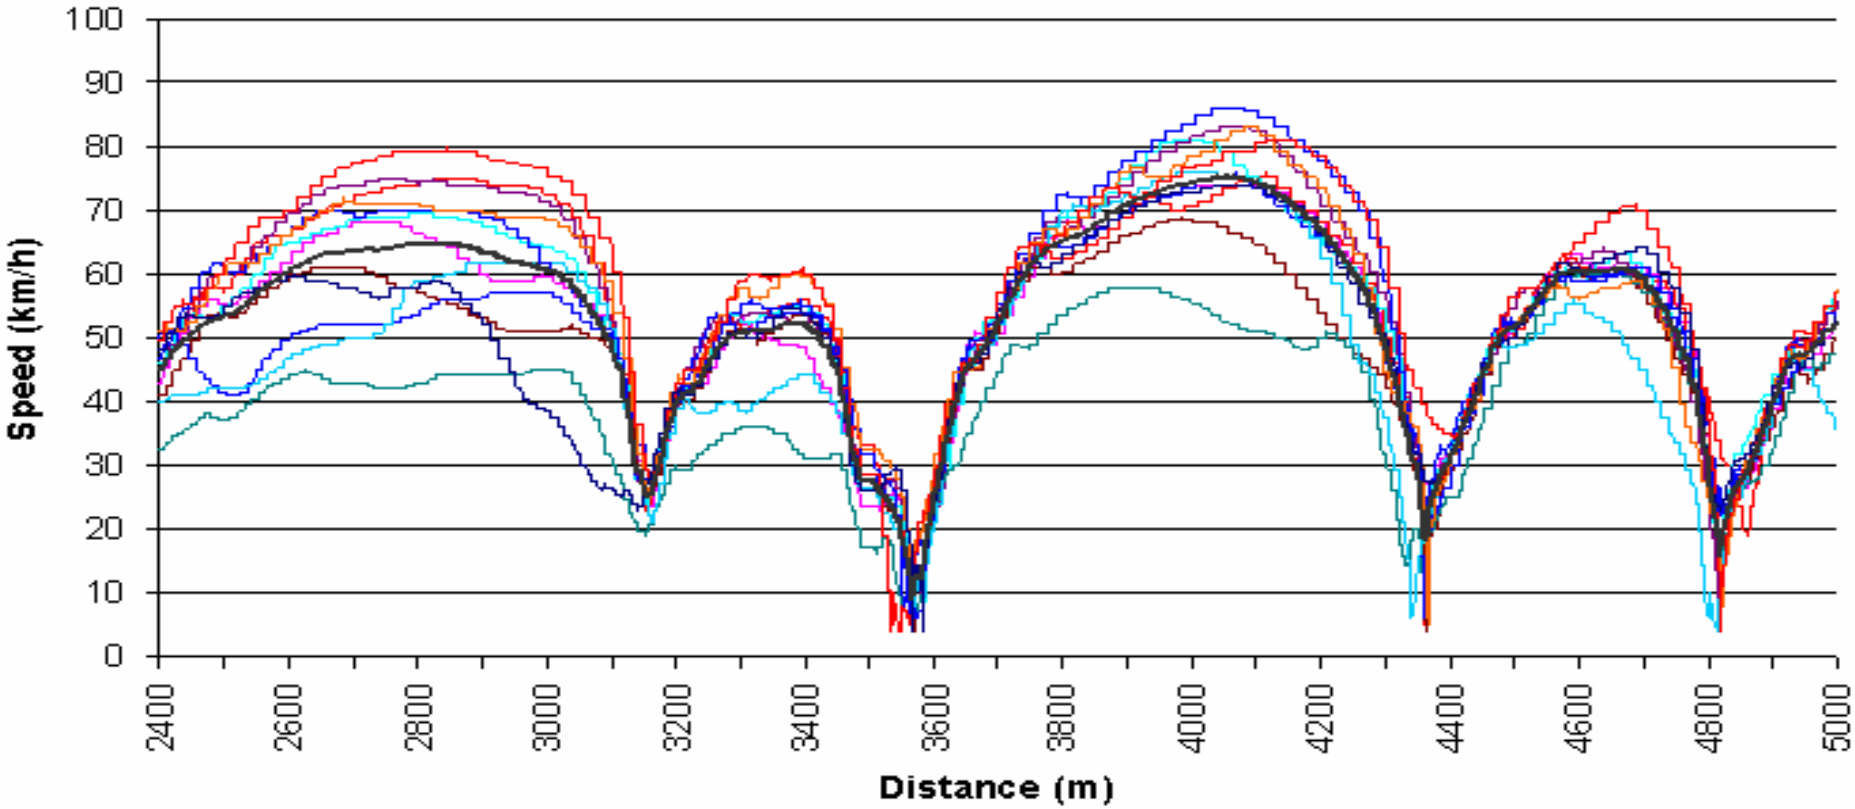

Average Speed

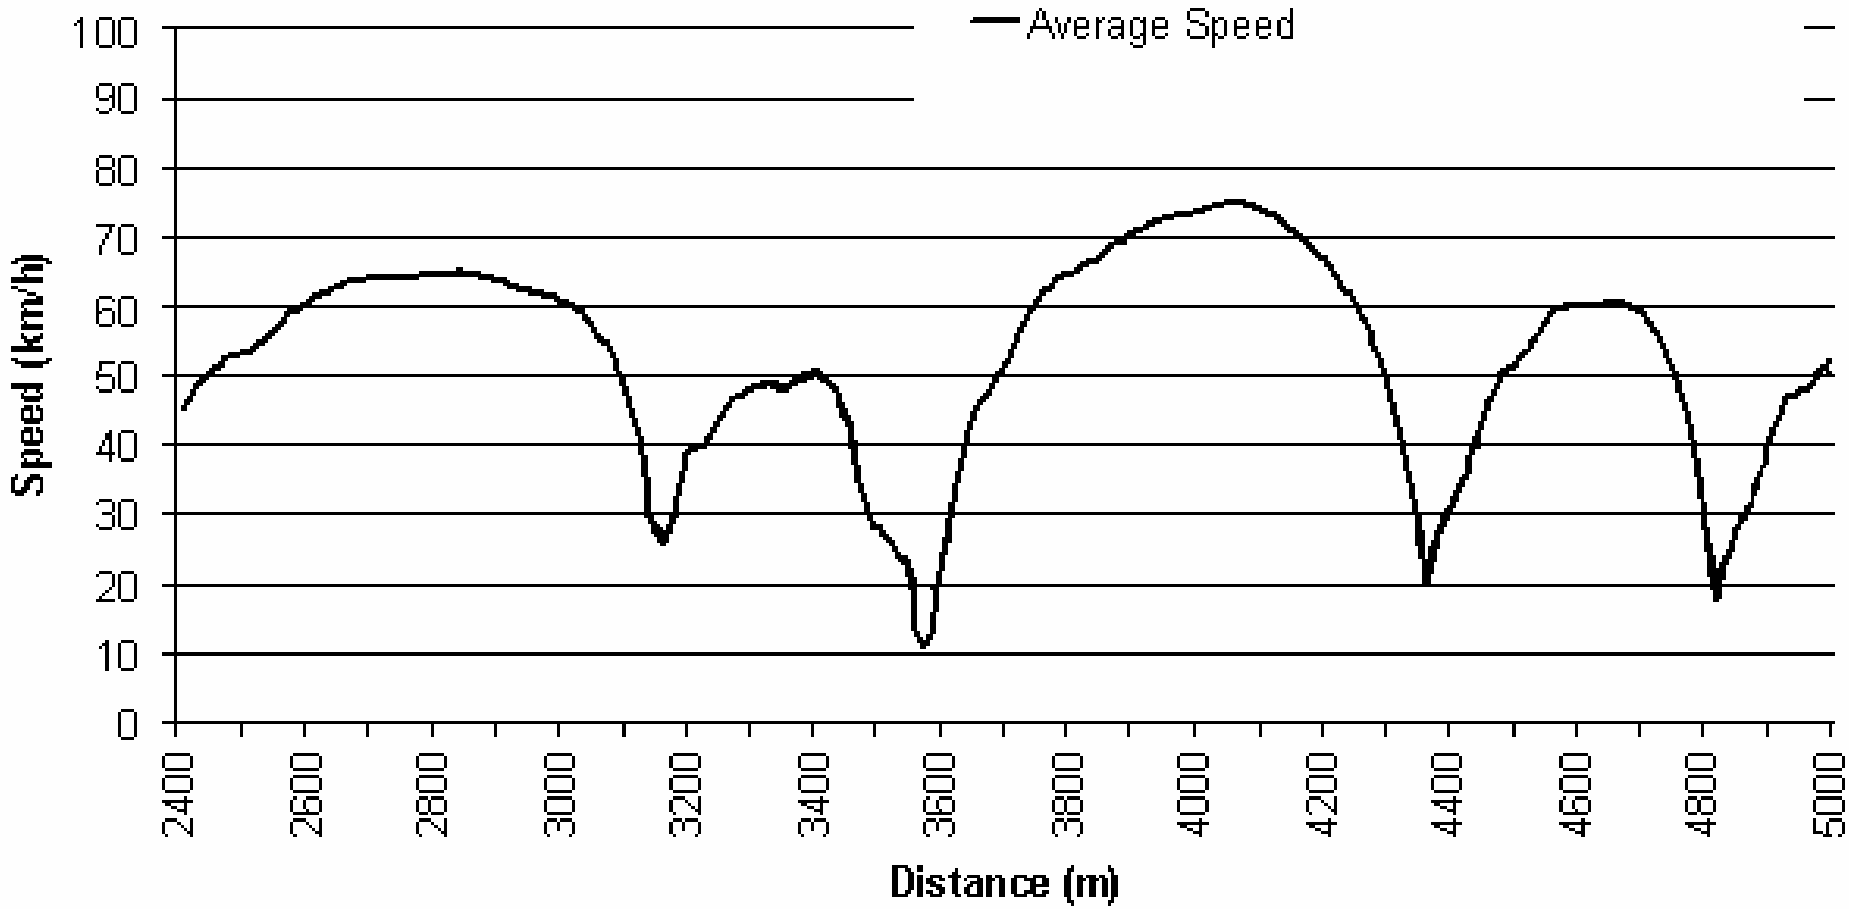

Use of Acceleration Profiles

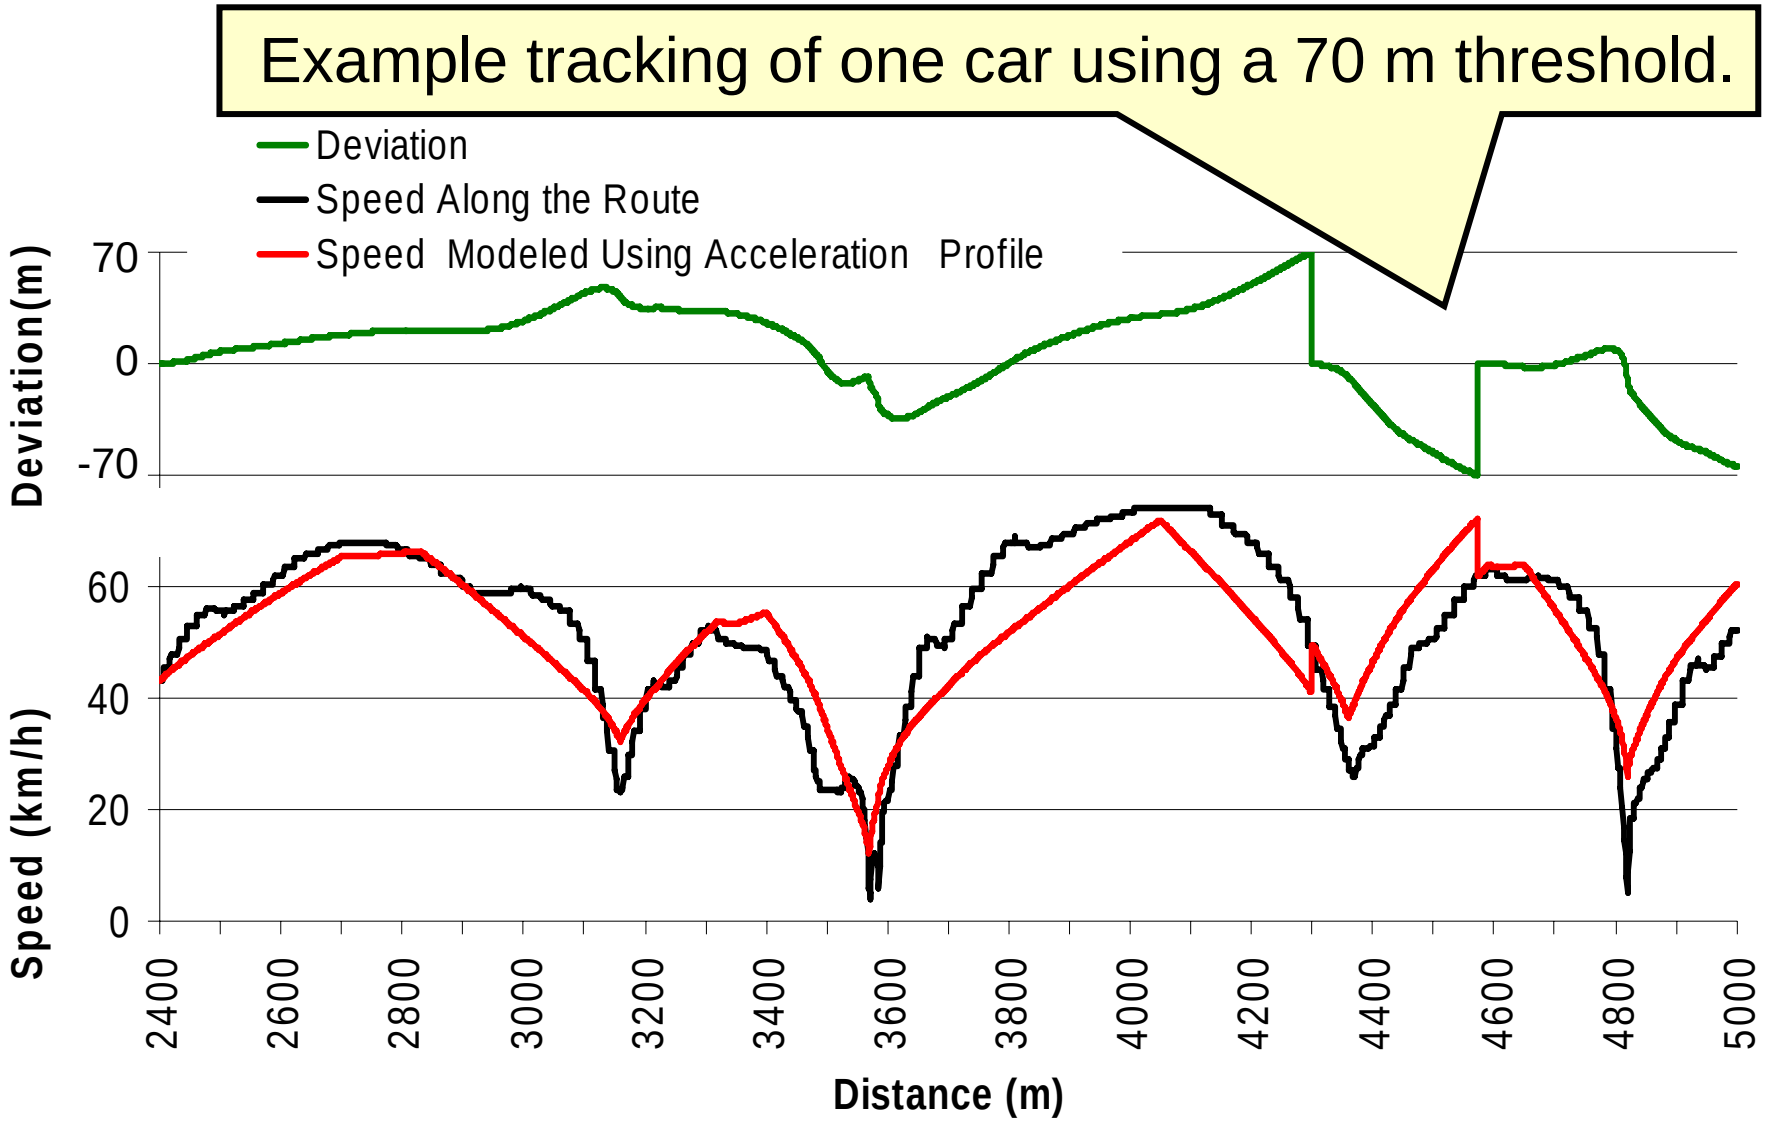

Results – Acceleration Profiles

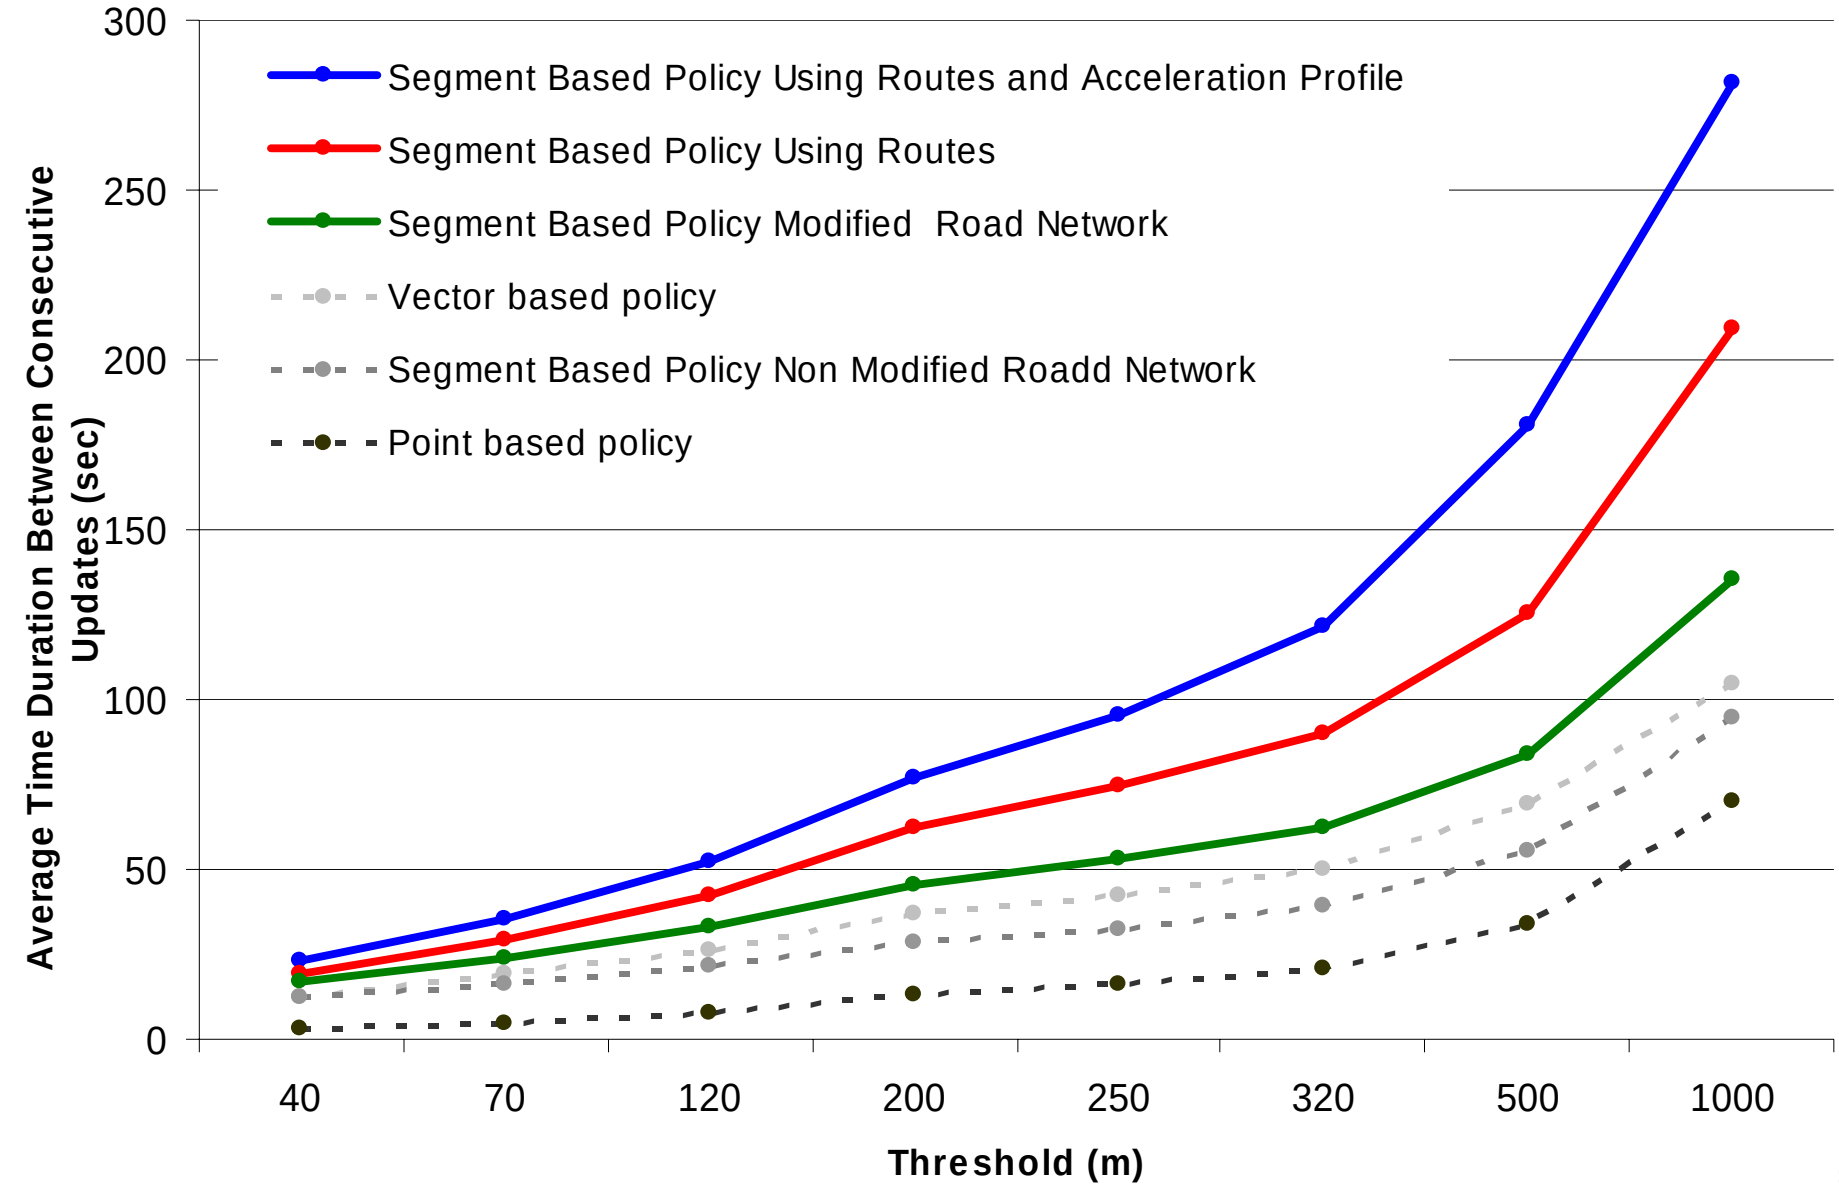

skip

Summary

- The Mobile Internet is emerging.
- Web 2.0 technologies will have a large impact.
- Streamspin applies Web 2.0 concepts to mobile services.
  - Easy creation and sharing of mobile services, scalable delivery of services.
- The foundations for one geo-context services were covered.
  - Efficient continuous tracking of moving objects with accuracy guarantees.
    - The use of real data was essential in guiding the design process.

Acknowledgments

- The Streamspin project
  - Aalborg University (*Kenneth H. Pedersen, Kristian Torp, Rico Wind*)
- The INFATI project
  - Aalborg University
- The M-Track and TAU-SeN projects, funded by ETRI, Korea.
  - Pusan National University (Ki-Joune Li, Si-Wan Kim, Kyoung-Sook Kim, Deuk-Chun Han)
  - Chungbuk National University (Keun Ho Ryu, Eung-Jae Lee)
  - Aalborg University (Dalia Tiesyte, Nerius Tradisauskas, Simonas Saltenis, Linas Bukauskas, *Alminas Civilis*, *Stardas Pakalnis*)
- The ContextIT project
  - Aalborg University (*Agne Brilingaite*)

Readings

- Papers co-authored by CSJ can be found here:
  - http://www.cs.aau.dk/~csj/Papers/
- The tracking video
  - On YouTube: http://www.youtube.com/watch?v=Ci9ZZ4FV77Y
  - Also available in higher quality on Daisy's website
- Websites
  - http://daisy.aau.dk
  - http://streamspin.com
  - http://www.cs.aau.dk/TRAX/
  - http://www.cs.aau.dk/DBTR/
  - http://www.cs.aau.dk/TimeCenter/

Readings

- Alminas Civilis, Christian S. Jensen, Stardas Pakalnis: Techniques for Efficient Road-Network-Based Tracking of Moving Objects. IEEE Trans. Knowl. Data Eng. 17(5): 698-712 (2005)
- Alminas Civilis, Christian S. Jensen, Jovita Nenortaite, Stardas Pakalnis: Efficient Tracking of Moving Objects with Precision Guarantees. MobiQuitous 2004: 164-173
- Christian S. Jensen, H. Lahrmann, Stardas Pakalnis, J. Runge: The Infati Data CoRR cs.DB/0410001: (2004)
- Agne Brilingaite, Christian S. Jensen, Nora Zokaite: Enabling routes as context in mobile services. GIS 2004: 127-136
- Agne Brilingaite, Christian S. Jensen: Enabling routes of road network constrained movements as mobile service context. GeoInformatica, 11(1): 55-102 (2007)
- Christian Hage, Christian S. Jensen, Torben Bach Pedersen, Laurynas Speicys, Igor Timko: Integrated Data Management for Mobile Services in the Real World. VLDB 2003: 1019-1030
- Christian S. Jensen: Database Aspects of Location-Based Services. Location-Based Services 2004: 115-148
- Christian S. Jensen: When the Internet Hits the Road. BTW 2007: 2-16
- Christian S. Jensen: Geo-Enabled, Mobile Services-A Tale of Routes, Detours, and Dead Ends. DASFAA 2006: 6-19