Spatial Indexing and Cars

Dr John Fawcett

Indexing

- Index = additional data structure
- Contains pointers to objects that are often stored in databases
- Index entries are arranged/grouped/sorted to speed up insertion/deletion/searching
- *Spatial* indexes speed up operations involving geometric spaces in 1, 2, or 3 dimensions

Nature of Spatial Datasets

- Type of data
  - point data (e.g. spot-height samples)
  - continuous spatial properties (e.g. contours)
  - discrete spatial data/regions (e.g. RCL databases)
- Use of data
  - location as place/context for activity
  - spatial relationships/relative positioning
- Examples: see Prof Robert Haining, “Spatial Data Analysis Theory and Practice” Cambridge University Press; 2003.

Geographic Information Systems

- GIS databases store data and construct indexes based on spatial attributes.
- Market leaders:
  - AutoCAD—architecture and scientific modelling/drawing/rendering
  - Oracle *Spatial*—models co-ordinate data and uncertainty/error. 2D: Cartesian and Lat/Long. Fuzzy comparison operators.

Spatial Indexing

- Typical situation:
  - Static data (e.g. road map) but mobile focus
  - Searches are “range queries”
- Region-Trees (R-Trees)
  - Like B^+-Trees
  - Ref: A. Guttman, “R-Trees: a dynamic index structure for spatial searching”, Proceedings of SIGMOD, Boston MA, June 1984, pp 45—57

R-Trees [1]

- Shapes of arbitrary objects too complex for efficient processing
  - Use Bounding Boxes or “Minimum Bounding Rectangles” (MBRs)
- Leaf node of R-Tree:
  - Contains tuples: (MBR, pObject)
- Internal node of R-Tree:
  - Contains tuples: (MBR, pChildNode)

Example R-Tree Leaf Node

Tuples:

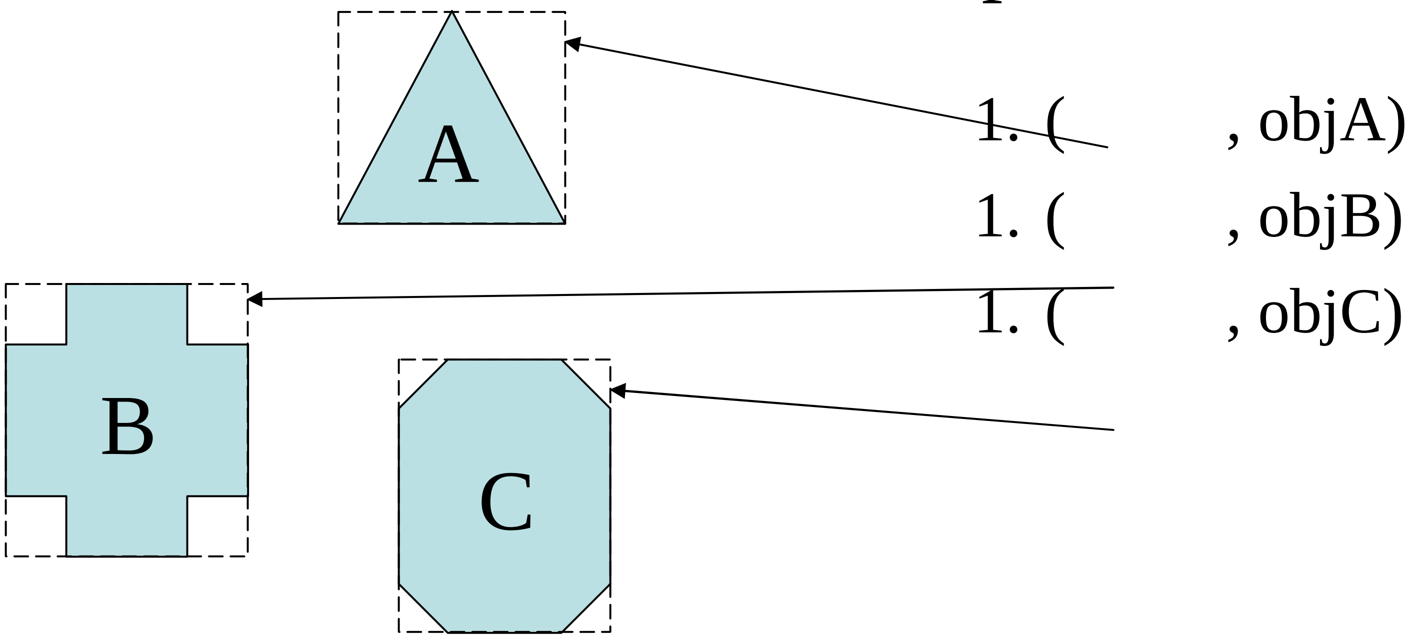

Example R-Tree

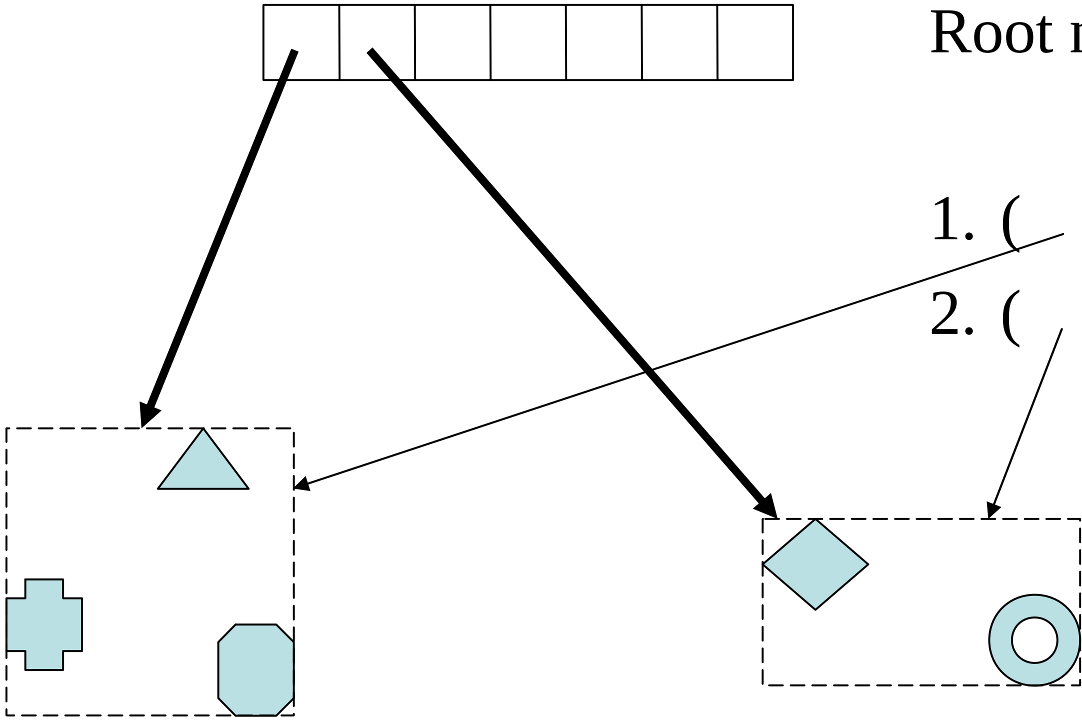

Root node tuples:

1. (, child-1)
2. (, child-2)

(m,M) R-Trees [2]

- Leaf and internal nodes contain between m and M entries where 2<=m<=ceil(M/2)
  - Similar storage flexibility to B-Trees
  - There’s a special case…
- Tree with only one node (the root) may house fewer than m items.
- Terminology: as items are added and removed, a node might *overflow* or *underflow* and require *splitting* or *merging*

R-Trees [3]

- Tree is height-balanced
- Bounding boxes of internal nodes are:
  - hierarchically nested
  - permitted to overlap
  - found to offer max searching efficiency when:
    - coverage is minimised
    - overlap is minimised

Another R-Tree

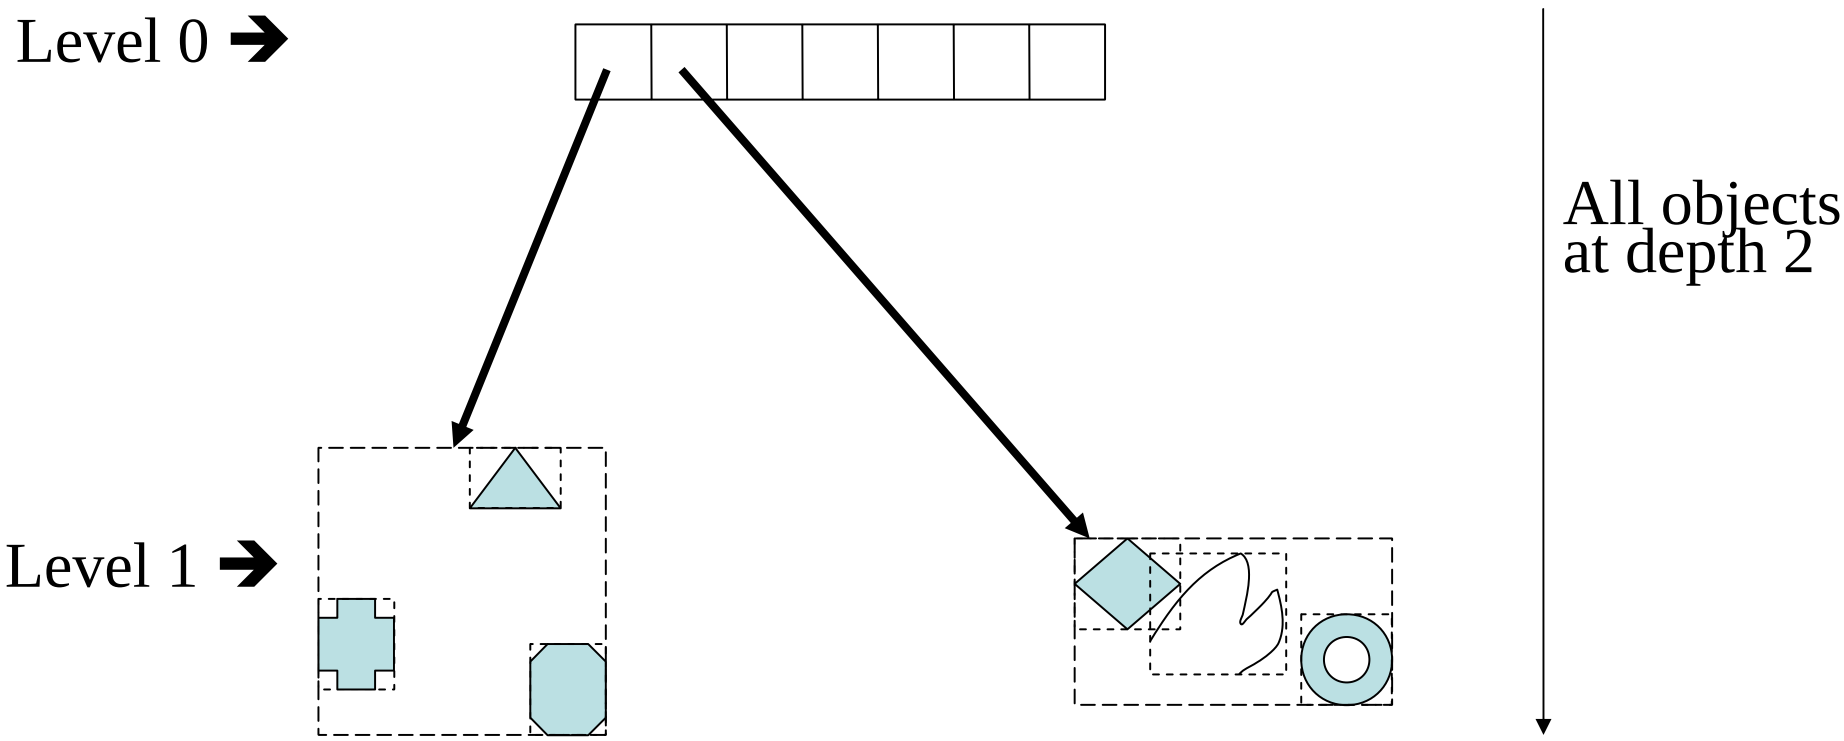

Constructing an R-Tree

- Build it dynamically...
  - Insert objects in order of arrival into system
  - Need a splitting strategy

- ...or build it statically
  - Remember objects as they appear
  - Then construct an optimally-divided tree
  - “Packing”

Algorithms on R-Trees

- Insertion, splitting strategies

  [demonstration given in lectures]
- Searching

  [demonstration given in lectures]

- Exercises: work out how to…
  - Delete an item from R-Tree
  - Merge two R-Trees

Optimising R-Trees: m

- Consider large m…
  – At least m of M entries used—few unused positions in data structure
  – Splitting a node can be expensive
- Consider small m…
  – (m,M) gives wider population range—fewer over/under flows expected
- Which best suits static/changing datasets?

Variants

- R^+-Trees (1987): hierarchically nested containers, overlap not permitted, objects indexed in each container spanning their position
- R*-Trees (1990): hierarchically nested containers, overlap permitted, objects indexed in each container spanning their position
- TV-Tree (1994), X-Tree (1997), …

Special-Purpose Variants

- SMR-Trees: range queries are infrequently demanded
- QSF-Trees: dataset *very* big; not even the index fits in main memory. Arranged to minimise page faults when descending levels of the tree.
- K-Trees, B-Trees, kdB-Trees, … enough.
- What if my dataset is dynamic?
  - Predictable changes can be pre-optimised!
  - See Kollios, Gunnopulos and Tsotras [96]

Fun Stuff

- Install in your car: GPS, a PC, sensors, GSM/GPRS/3G phone, Wireless LAN, Bluetooth, USB hub in dashboard, CanBUS adaptor, EMC link, ByteFlight host, …
- Naïve to think of the PC as just a means of collecting sensor data

- Instead, consider the car to be a piece of the Internet…

Navstar GPS in 1 slide

- All-weather, round-the-clock, timing and ranging to an unlimited number of simultaneous users with anti-jamming
- Block 1 SVs formed Demo System. Block 2 achieved FOC. 24 SVs in space segment+spares. 3 orbital planes.
- SV=3-Axis stabilized, nadir pointing using reaction wheels. Dual solar arrays 400W+ NiCd batteries. S-Band (SGLS) communications for control and telemetry. UHF cross-link between spacecraft. Hydrazine propulsion system. *Translation to English=>flying atomic clock.*
- Fix complex—corrections for relativistic velocity, photon pressure, local G anomalies, ionosphere diffraction...

Portable?


Esc
Del
Q W E R T Y U I O P
Enter
Tab
A S D F G H J K L
Shift
Z X C V B N M
Shift
Ctrl
Fn
Menu


TONE
MODE
SB
RDS
TP
AM
FM
BMW
SC
MPH

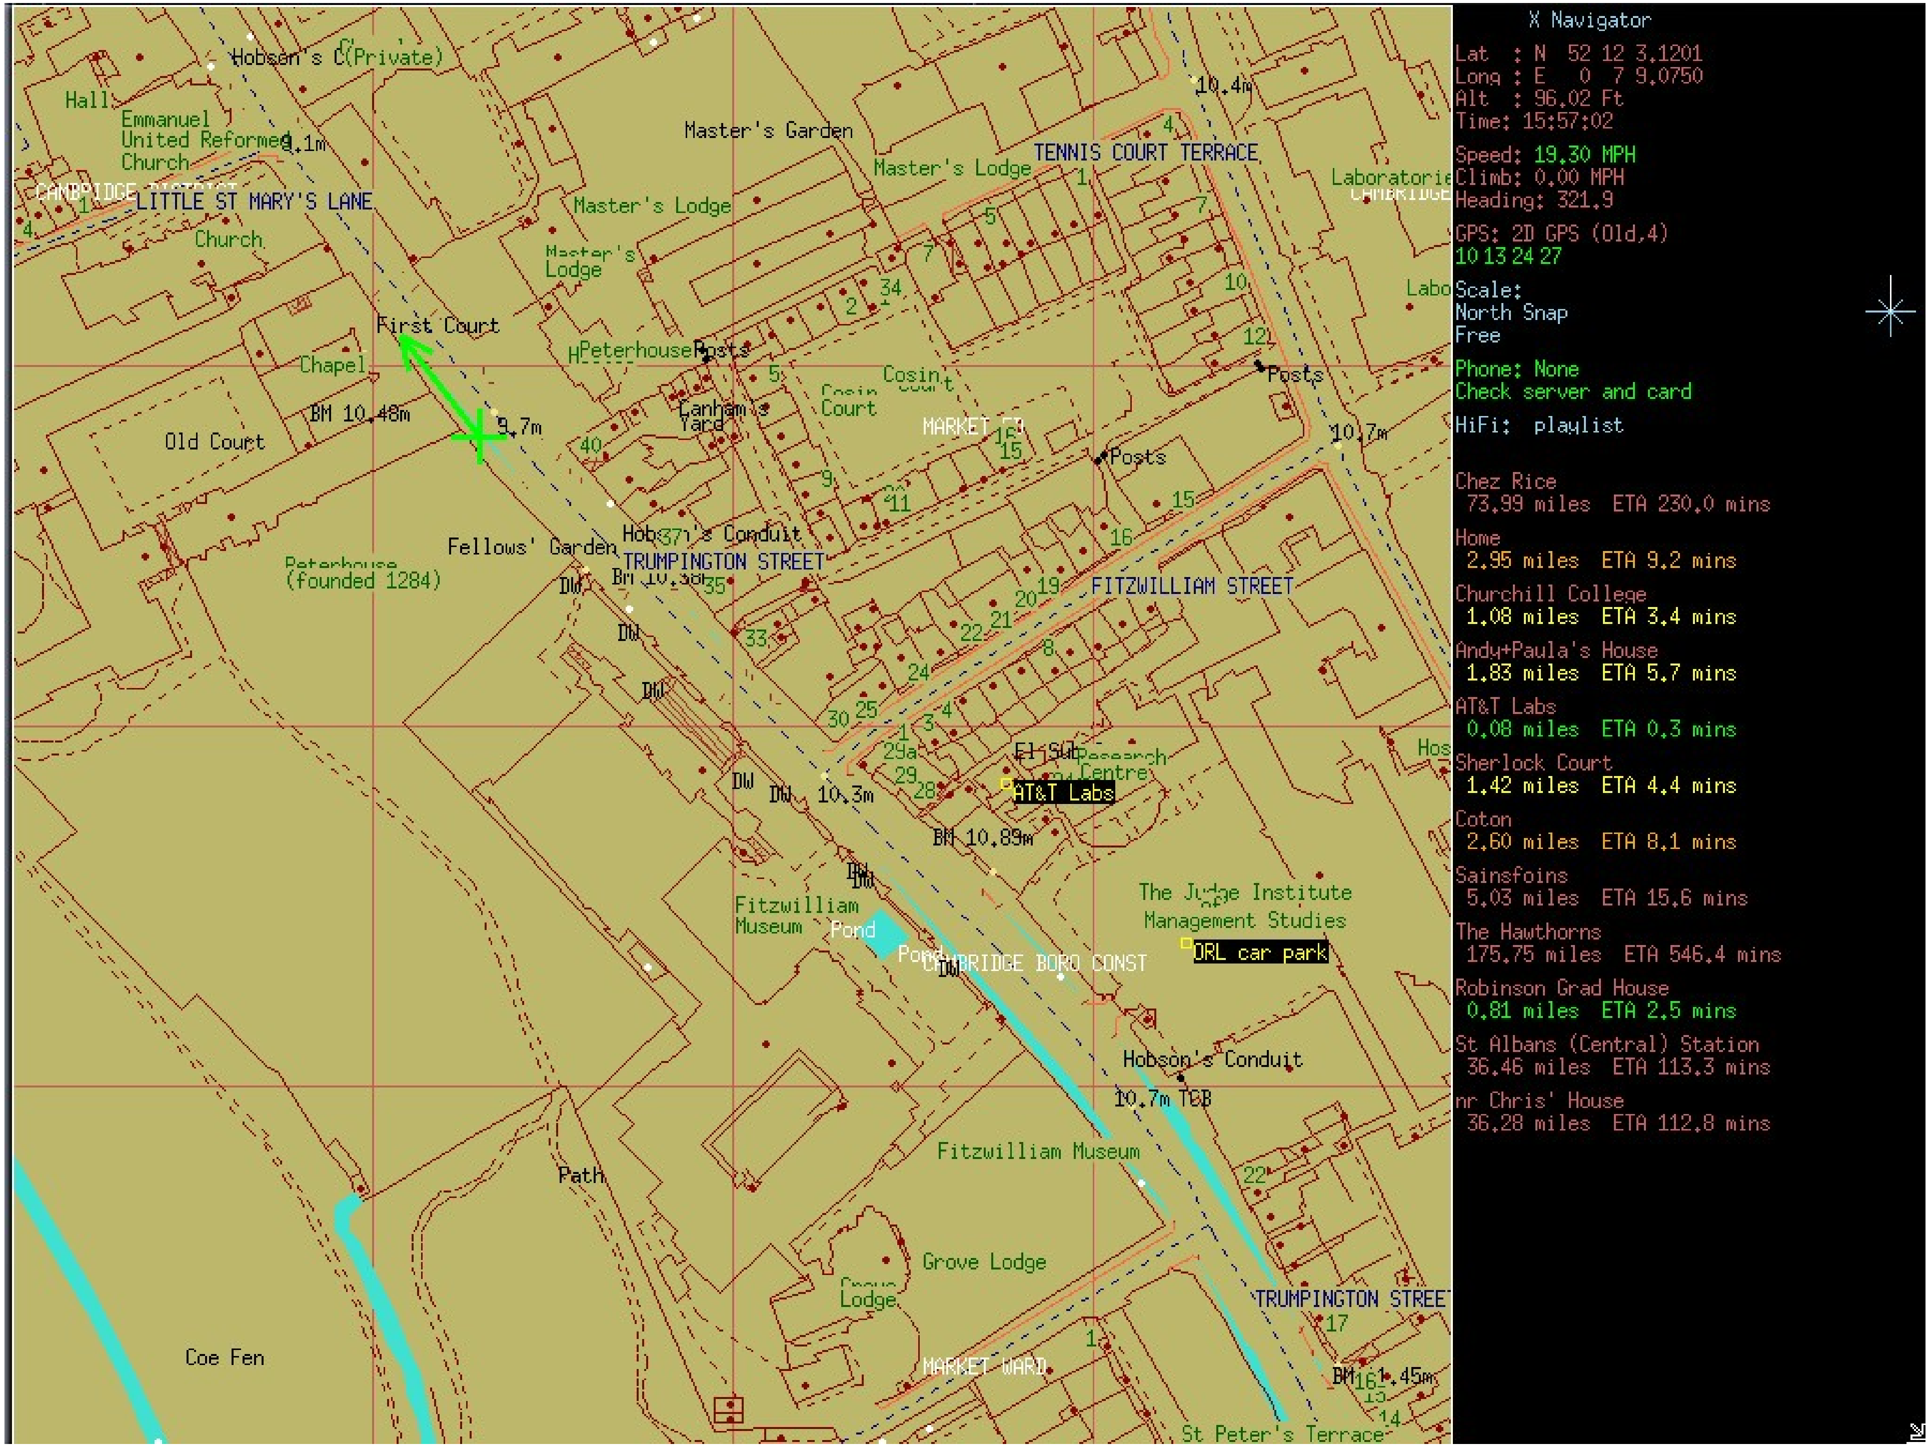
X Navigator
Lat : N 52 12 3.1201
Long : E 0 7 9.0750
Alt : 96.02 Ft
Time: 15:57:02
Speed: 19.30 MPH
Climb: 0.00 MPH
Heading: 321.9
GPS: 2D GPS (Old.4)
10 13 24 27
Scale:
North Snap
Free
Phone: None
Check server and card
HiFi: playlist
Chez Rice
73.99 miles ETA 230.0 mins
Home
2.95 miles ETA 9.2 mins
Churchill College
1.08 miles ETA 3.4 mins
Andy+Paula's House
1.83 miles ETA 5.7 mins
AT&T Labs
0.08 miles ETA 0.3 mins
Sherlock Court
1.42 miles ETA 4.4 mins
Coton
2.60 miles ETA 8.1 mins
Sainsfoins
5.03 miles ETA 15.6 mins
The Hawthorns
175.75 miles ETA 546.4 mins
Robinson Grad House
0.81 miles ETA 2.5 mins
St Albans (Central) Station
36.46 miles ETA 113.3 mins
nr Chris' House
36.28 miles ETA 112.8 mins
Hobson's C(Private)
Hall
Emmanuel
United Reformed
Church
LITTLE ST MARY'S LANE
Church
Master's Garden
Master's Lodge
TENNIS COURT TERRACE
First Court
Chapel
Peterhouse
Posts
Old Court
BM 10.48m
9.7m
Canham's
Yard
Cosin Court
MARKET
Fellows' Garden
TRUMPINGTON STREET
FITZWILLIAM STREET
(founded 1284)
10.3m
BM 10.89m
AT&T Labs
The Judge Institute
Management Studies
ORL car park
Fitzwilliam
Museum
Pond
CAMBRIDGE BORO CONST
Hobson's Conduit
10.7m TCB
Fitzwilliam Museum
Path
Grove Lodge
Lodge
TRUMPINGTON STREET
Coe Fen
MARKET WARD
St Peter's Terrace

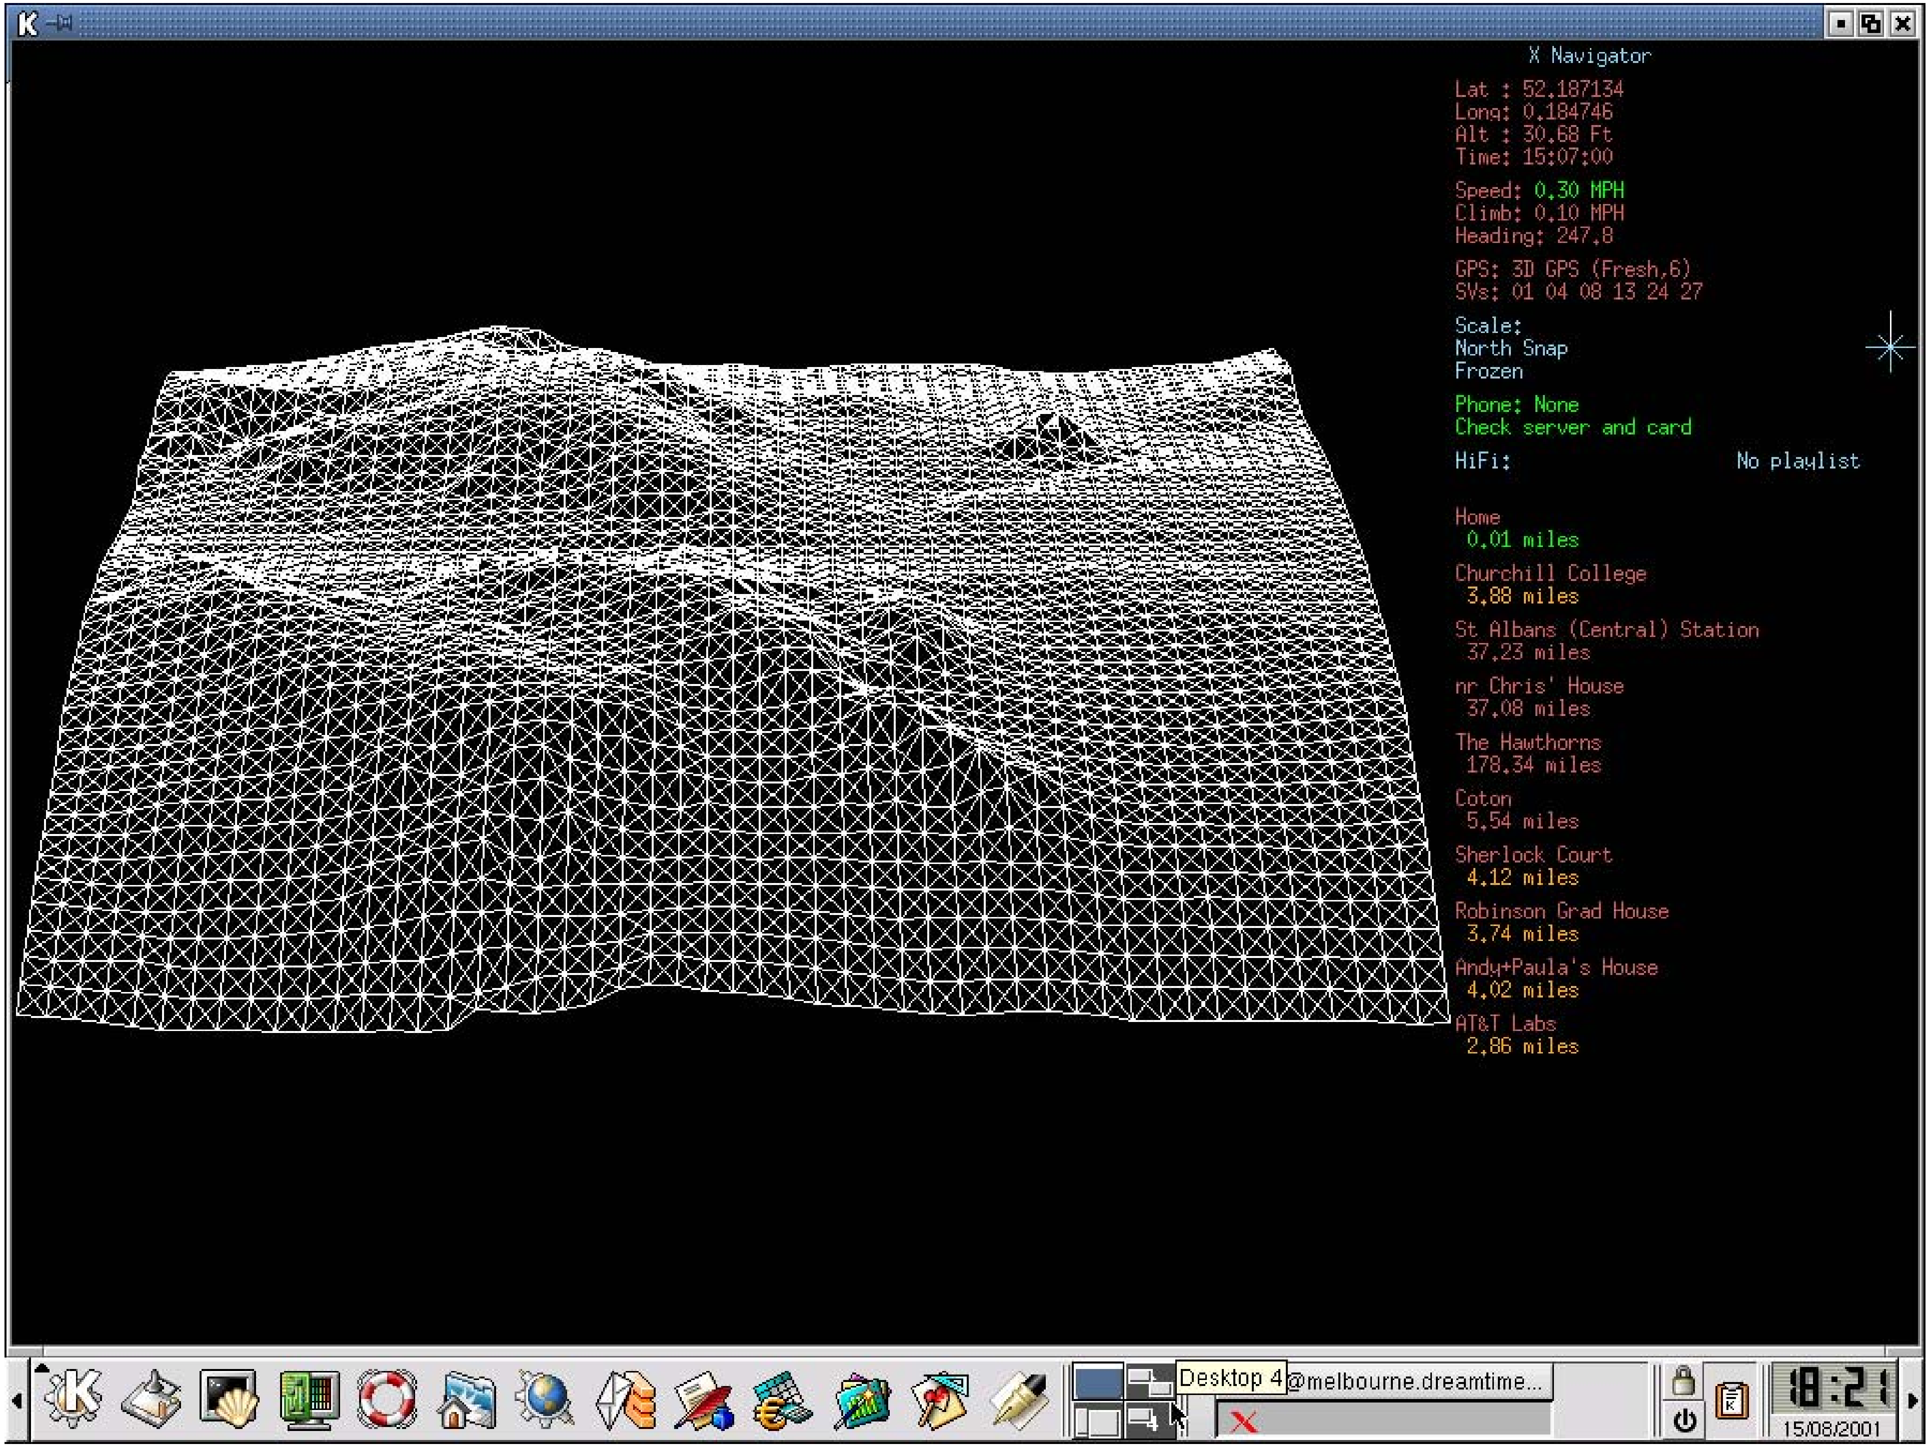
X Navigator
Lat : 52.187134
Long: 0.184746
Alt : 30.68 Ft
Time: 15:07:00
Speed: 0.30 MPH
Climb: 0.10 MPH
Heading: 247.8
GPS: 3D GPS (Fresh,6)
SVs: 01 04 08 13 24 27
Scale:
North Snap
Frozen
Phone: None
Check server and card
HiFi:
No playlist
Home
0.01 miles
Churchill College
3.88 miles
St Albans (Central) Station
37.23 miles
nr Chris' House
37.08 miles
The Hawthorns
178.34 miles
Coton
5.54 miles
Sherlock Court
4.12 miles
Robinson Grad House
3.74 miles
Andy+Paula's House
4.02 miles
AT&T Labs
2.86 miles
Desktop 4
@melbourne.dreamtime...
18:21
15/08/2001