

R-tree project

Albert-Jan Yzelman

August 15, 2007



Outline



- 1 Introduction
- 2 Context
- 3 R-trees
 - Demands
 - Minimal Bounding Rectangles (MBRs)
 - Putting it together: R-trees
 - Query demo
- 4 Questions?



What is the R-tree project

- Internal project at Alten NL, started March 2007
http://www.altennederland.nl/wiki/index.php/R-Tree_Project
- Applied for and received a WBSO grant for one year



What is the R-tree project

- Internal project at Alten NL, started March 2007
http://www.altennederland.nl/wiki/index.php/R-Tree_Project
- Applied for and received a WBSO grant for one year
- The first six months were mainly used for researching R-trees for use with oil reservoir simulation software at Shell



What is the R-tree project

- Internal project at Alten NL, started March 2007
http://www.altennederland.nl/wiki/index.php/R-Tree_Project
- Applied for and received a WBSO grant for one year
- The first six months were mainly used for researching R-trees for use with oil reservoir simulation software at Shell
- Goals:



What is the R-tree project

- Internal project at Alten NL, started March 2007
http://www.altennederland.nl/wiki/index.php/R-Tree_Project
- Applied for and received a WBSO grant for one year
- The first six months were mainly used for researching R-trees for use with oil reservoir simulation software at Shell
- Goals:
 - 1 Create a datastructure which will enhance Shell's simulation software



What is the R-tree project

- Internal project at Alten NL, started March 2007
http://www.altennederland.nl/wiki/index.php/R-Tree_Project
- Applied for and received a WBSO grant for one year
- The first six months were mainly used for researching R-trees for use with oil reservoir simulation software at Shell
- Goals:
 - 1 Create a datastructure which will enhance Shell's simulation software
 - 2 Obtain the theoretical knowhow and a generic, customisable R-tree software library required for using R-trees in many other areas
Alten is or will be active in



The project is managed by Eric Haesen.

Consultants from Alten who work or have worked on the R-tree project (chronologically):

- Arno Swart (supervisor)
- Guy Calluy (exploration, visual R-tree, amongst others)
- Ron van Kesteren (k -nearest neighbourhood (knn) query research & implementation, knn benchmarking, a.o.)
- Martijn Souman (optimisation)



Outline



- 1 Introduction
- 2 Context
- 3 R-trees
 - Demands
 - Minimal Bounding Rectangles (MBRs)
 - Putting it together: R-trees
 - Query demo
- 4 Questions?



Shell's oil reservoir simulation software: Dynamo

- Numerical modelling of oil reservoirs
- Simulates effects of, for example, well drilling at different locations



An oil reservoir discretisation (or *grid*) to work with is required for simulation:




Figure: Example oil reservoir grid

The simulation software and grid upscaling software included in Dynamo require efficient answers to at least the following types of queries:

- Point
- Line
- Box

The irregular input grids and the geometric nature of querying call for a specialised data structure.



Outline



1 Introduction

2 Context

3 R-trees

- Demands
- Minimal Bounding Rectangles (MBRs)
- Putting it together: R-trees
- Query demo

4 Questions?



Data structure demands:



Data structure demands:

- Can store d -dimensional *hypercubes*



Data structure demands:

- Can store d -dimensional *hypercubes*
- Performs point, line, and box queries as fast as possible...



Data structure demands:

- Can store d -dimensional *hypercubes*
- Performs point, line, and box queries as fast as possible...
- ...but also keeps memory usage in check!



Consider an object in 2D:



Figure: A two-dimensional object

Consider an object in 2D:




Figure: A two-dimensional object with its MBR

An MBR can also be defined on multiple objects:




Figure: Multiple objects

An MBR can also be defined on multiple objects:




Figure: Multiple objects contained in their MBR

An MBR can also be defined on multiple other MBRs:




Figure: Multiple MBRs

An MBR can also be defined on multiple other MBRs:





Figure: Multiple MBRs contained in another MBR

The tree data structure



We can let each node of a tree correspond to an MBR and obtain the so called R-trees.



We can let each node of a tree correspond to an MBR and obtain the so called R-trees.

- Let each internal node of the R-tree correspond to the MBR of all MBRs stored in the children of that internal node




We can let each node of a tree correspond to an MBR and obtain the so called R-trees.


- Let each internal node of the R-tree correspond to the MBR of all MBRs stored in the children of that internal node
- Let each leaf node of the R-tree correspond to the MBR of a single object stored in the R-tree




We can let each node of a tree correspond to an MBR and obtain the so called R-trees.




We can let each node of a tree correspond to an MBR and obtain the so called R-trees.



We can let each node of a tree correspond to an MBR and obtain the so called R-trees.



We can let each node of a tree correspond to an MBR and obtain the so called R-trees.




Actively researched R-tree variations:


- The original R-tree, introduced in 1984
- Top-down Greedy Split (TGS)
- Hilbert R-tree
- HilbertTGS




Demonstration of a line query.



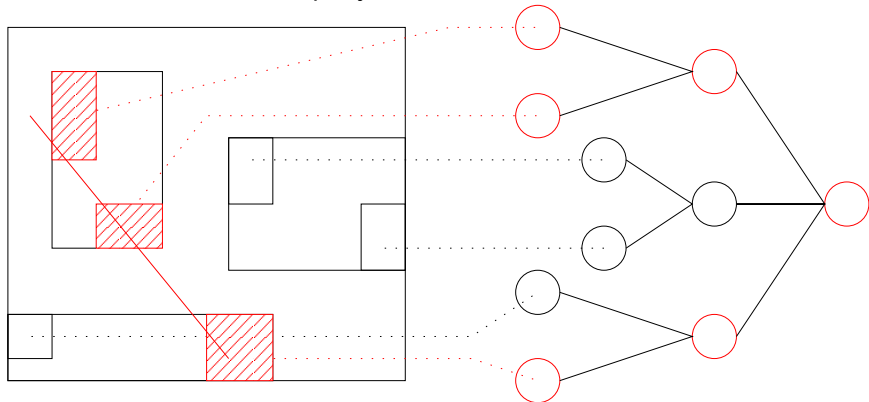
Demonstration of a line query.




Demonstration of a line query.



Demonstration of a line query.



Demonstration of a line query.



Possible other application areas:

- Databases
- Graphics/Gaming
- Chip Design
- Navigation Systems
- Image Processing
- ...?



Outline



- 1 Introduction
- 2 Context
- 3 R-trees
 - Demands
 - Minimal Bounding Rectangles (MBRs)
 - Putting it together: R-trees
 - Query demo
- 4 Questions?

