

# ReiserFS

Nodes of a BTree each covers many consecutive blocks

Files' contents are only stored in leaf nodes

Search keys (only D, F, O have to be used everywhere)

D = directory number - unique id for dir this file is in,

F = file number - unique id for this file

O = fragment offset

S = fragment size

L = total length of whole file

Many small files share a single leaf node

Big files are fragmented and spread over many nodes.

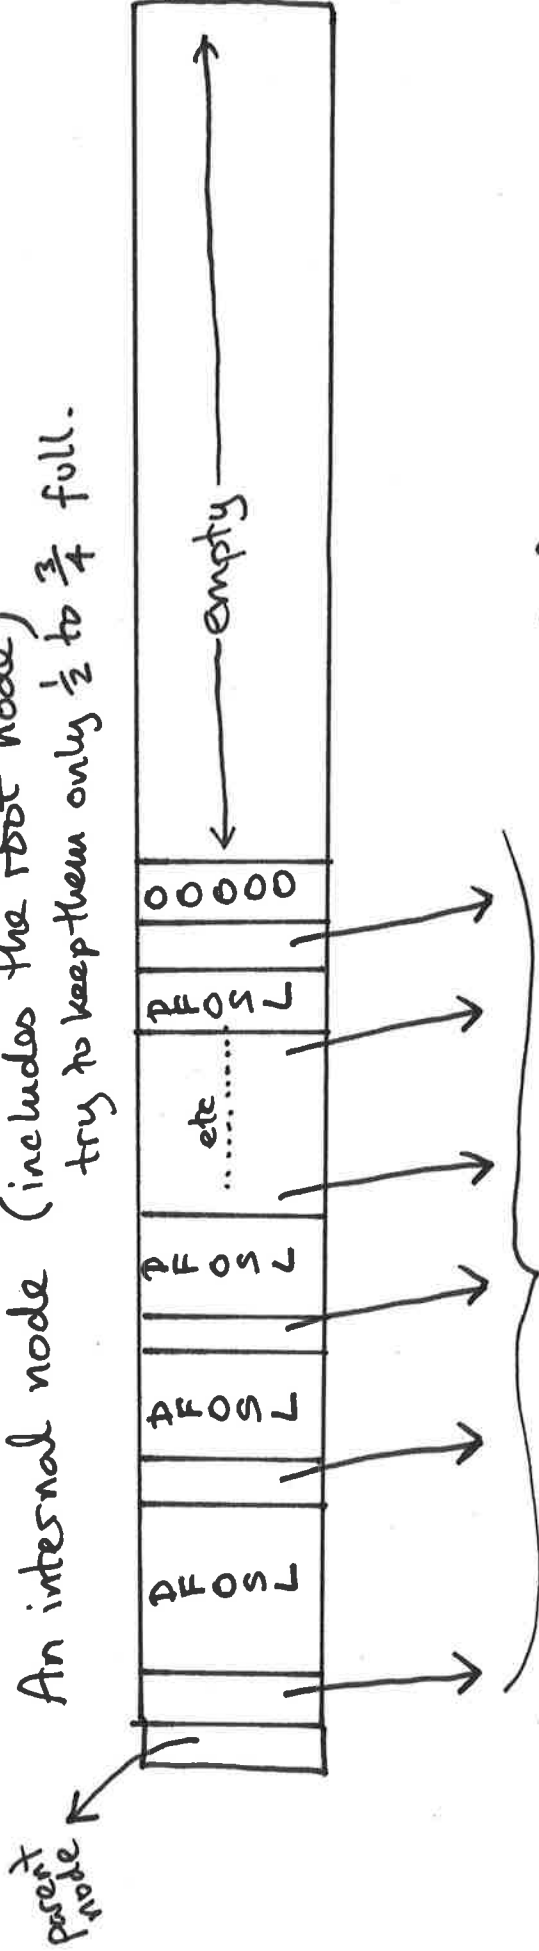
if node can contain up to 10,000 bytes of file content and we have a 27,300 byte file

Maybe first 3,100 bytes fit at the end of one node  
next 10,000 bytes (O=3,100) fill the next node  
next 10,000 bytes (O=13,100) fill the next node  
remaining 4,200 bytes (O=23,100) share the next node with the next file

Search keys are (D, F, O)

Tree is ordered primarily on increasing D,  
secondarily on increasing F,  
thirdly on increasing O.

An internal node (includes the root node)  
try to keep them only  $\frac{1}{2}$  to  $\frac{3}{4}$  full.



pointers to other nodes, internal or leaf.  
remember its all on disc, nodes are just big clusters,  
so these pointers are just cluster (block) numbers.

parent pointers are not really necessary.

