amdb:
An Access Method Debugging and Analysis Tool

Marcel Kornacker, Mehul Shah,
Joe Hellerstein
UC Berkeley

Motivation

- **Access method (AM) design and tuning is a black art.**
  - Which AM do I use to index my non-traditional data type?
  - How well do existing AMs perform for my workload?
- **Generalized search trees (GiST) provide a framework for AM implementations**
- **amdb is a debugging and analysis tool for GiST-implemented AMs**

Overview of Generalized Search Trees

GiST AM Parameters

Bounding predicate (BP)=User defined key

Consistent( ) returns matches

PickSplit( ) splits page items into two groups

Penalty( ) returns insertion cost

Union( ) updates BPs

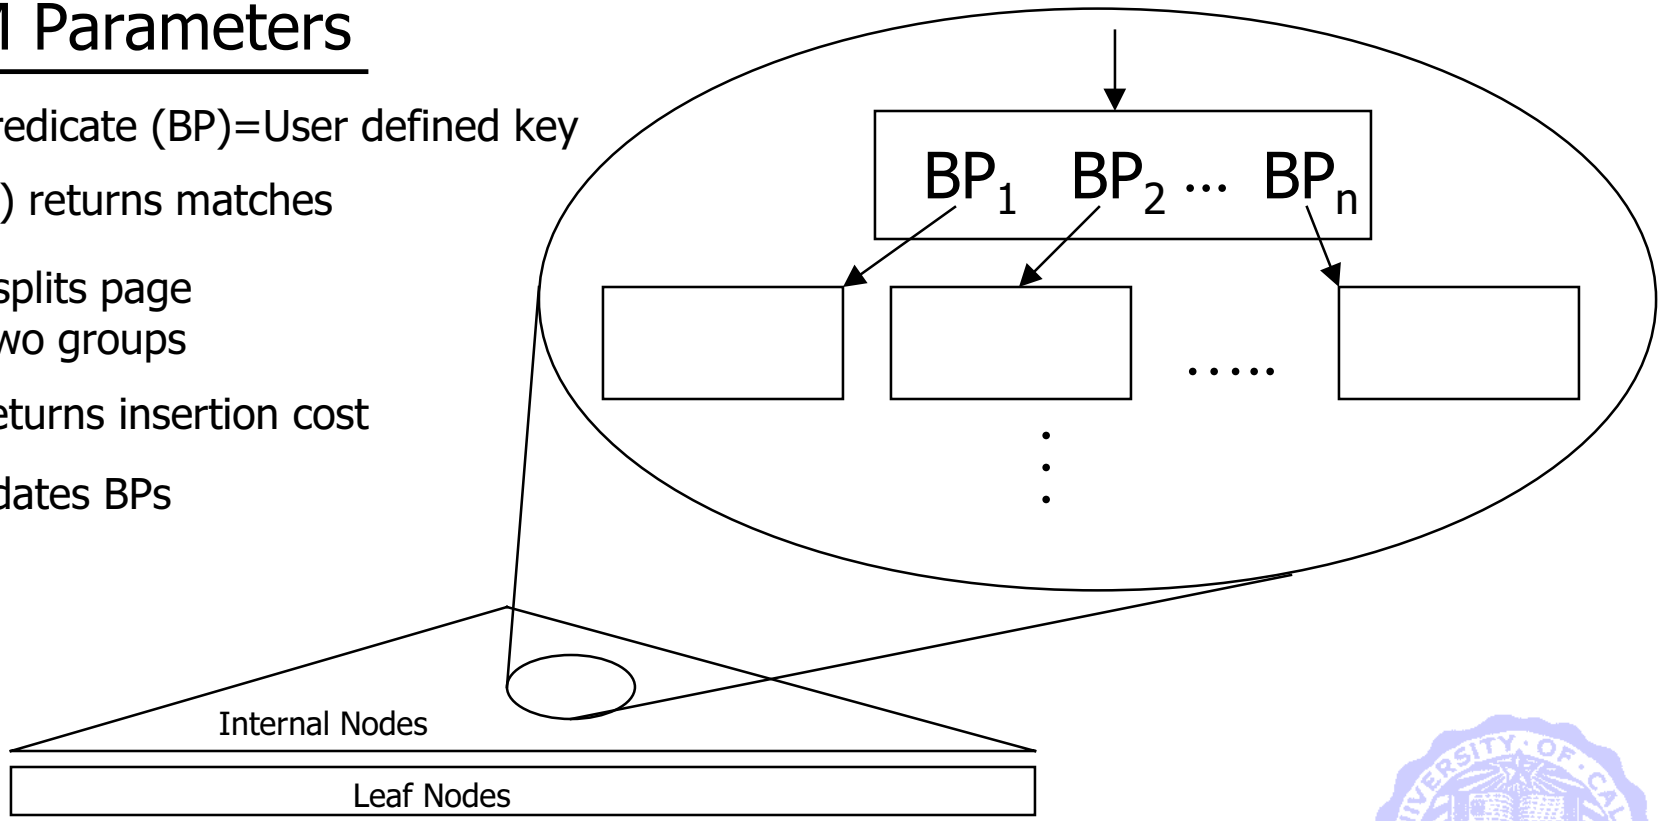

amdb Features

- **Tracing of insertions, deletions, and searches**
- **Debugging operations: breakpoints on node splits, updates, traversals, and other events**
- **Global and structural views of tree allow navigation and provide visual summary of statistics**
- **Graphical and textual view of node contents**
- **Analysis of workloads and tree structure**
- **Analysis of GiST AM parameters: BPs, PickSplit(), and Penalty()**

Debugging Operations

Stepping Controls

Tree View

Shows structural organization of index.

Highlights current traversal path during debugging steps.

Console window

Displays search results, debugging output, and other status info.

Breakpoint Table

Defines and enables breakpoints on events

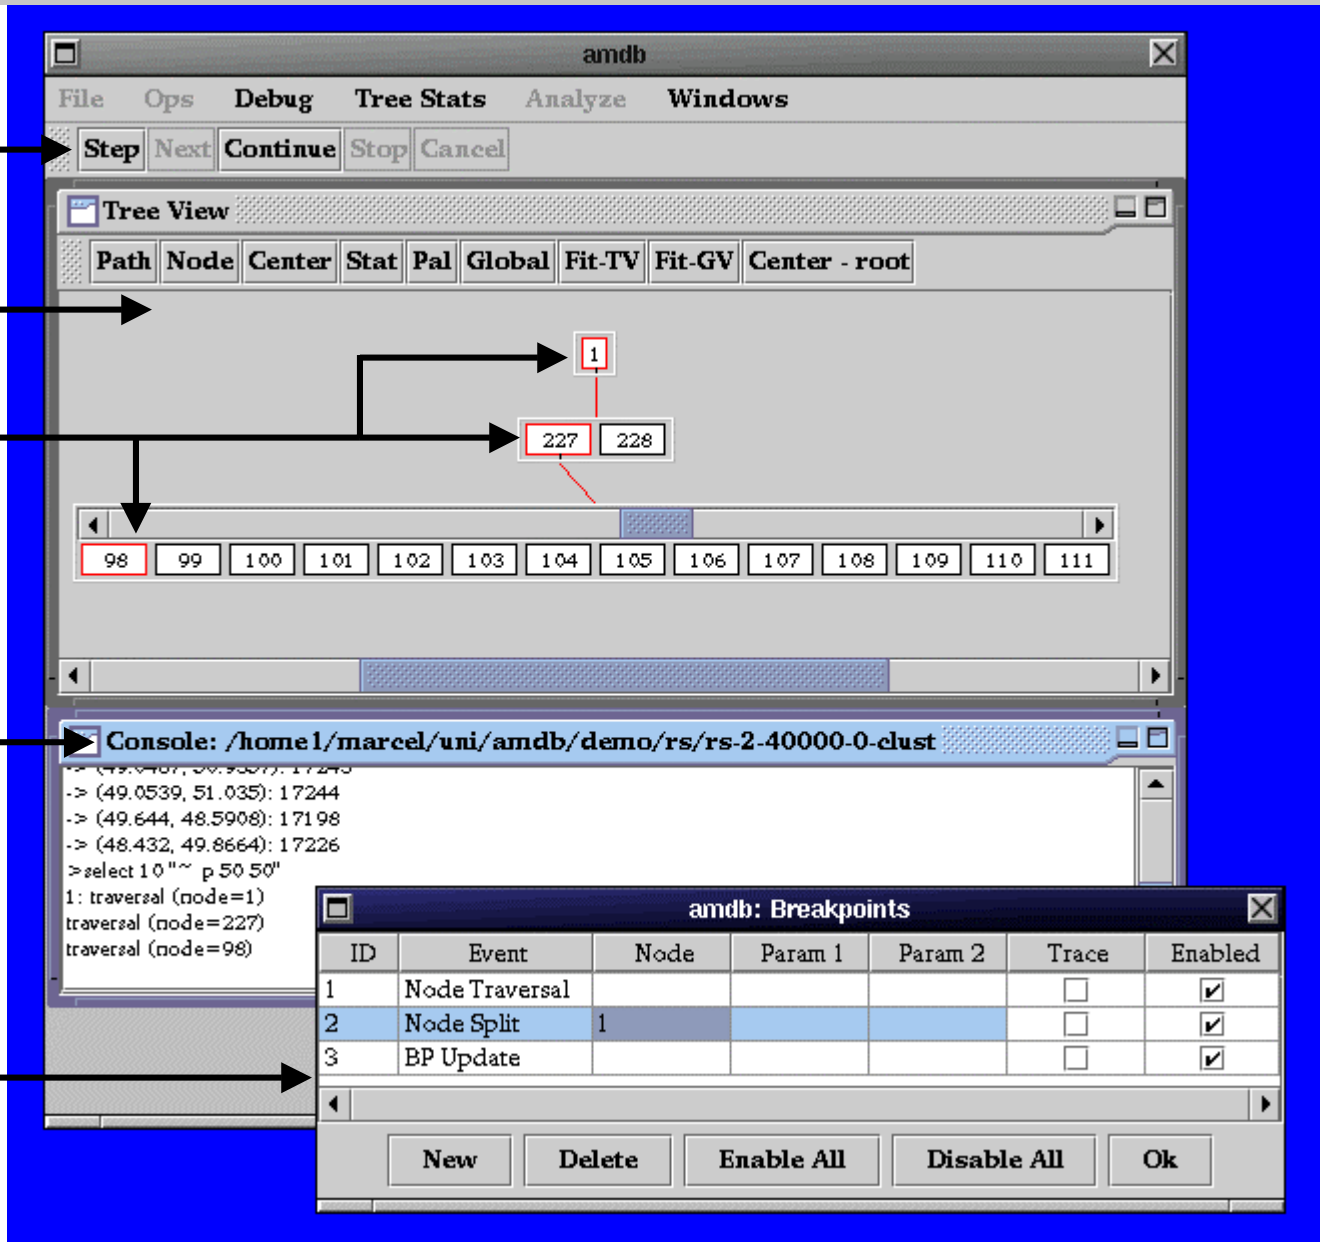

Node Visualization

Node View

Displays bounding predicates (BPs) and items within nodes.

Highlights BPs on current traversal path.

Split Visualization

Shows how BPs or data items are divided with PickSplit( )

Node Contents

Provides textual description of node

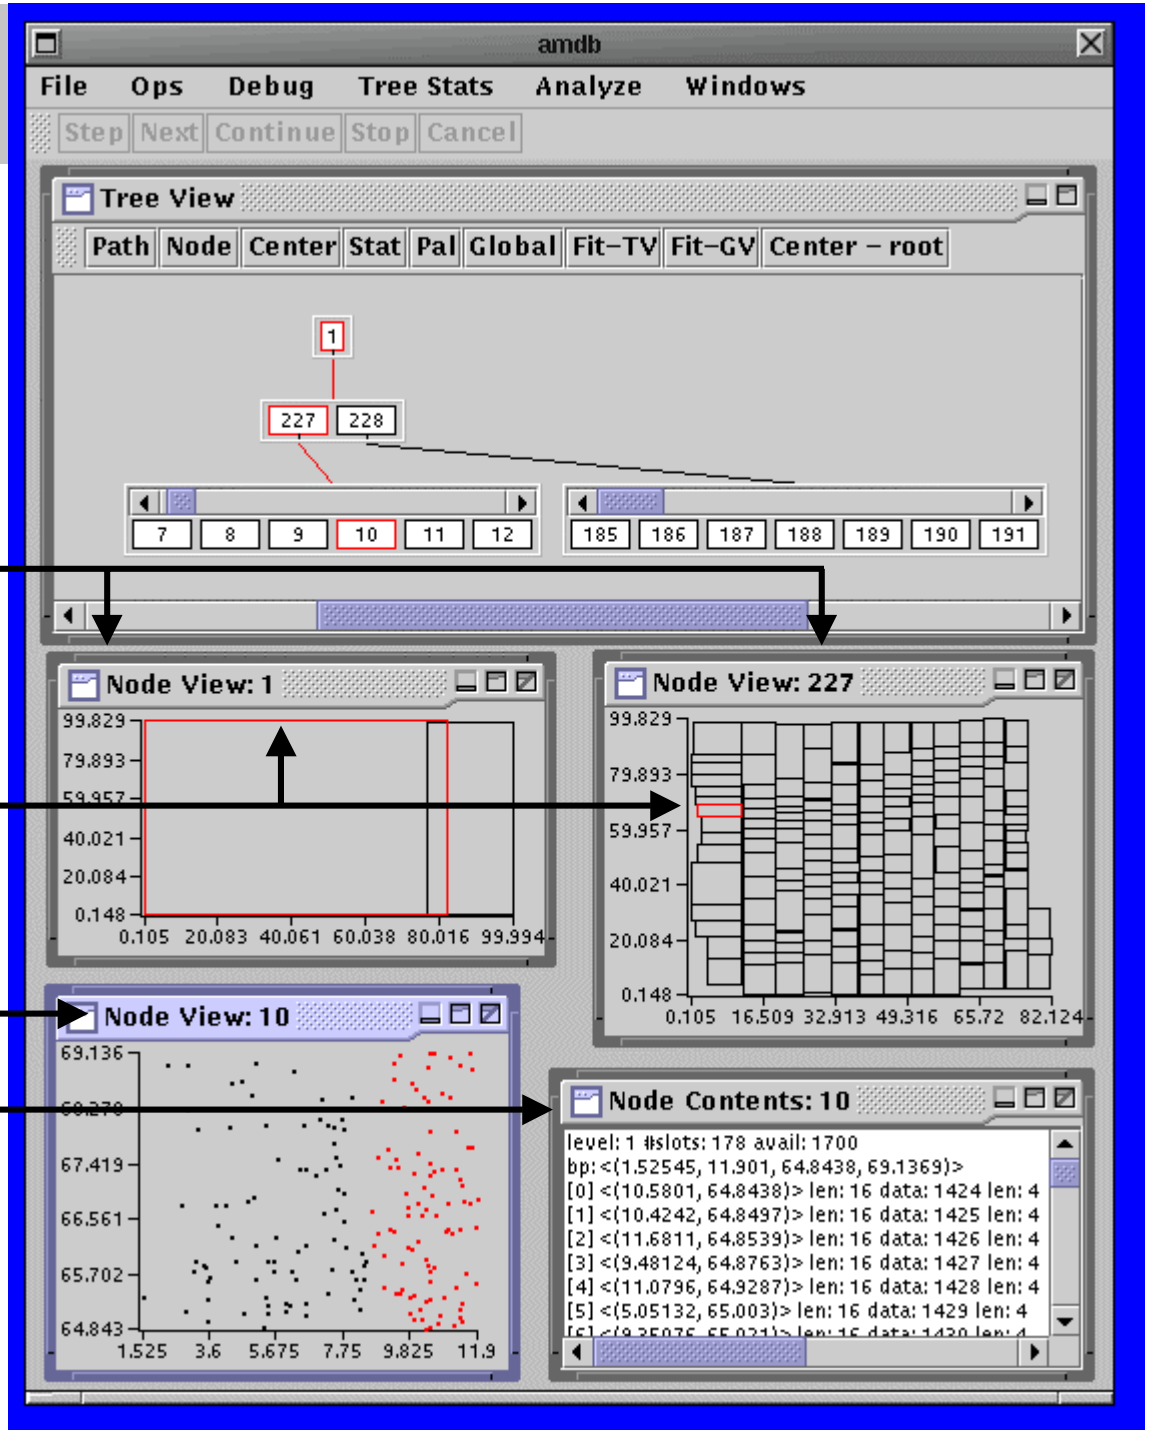

Analysis Framework

- **Analysis of index in context of user-specified workload:**
  - performance cannot be assessed independently of workload
  - metrics must reflect workload performance, not data semantics
- **Analysis procedure:**
  - Is data clusterable, given workload?
  - Assess tree and use metrics to pinpoint defects
  - Evaluate performance of PickSplit( ) and Penalty( ) methods

Performance Metrics

- **Factors affecting performance**
  - Clustering, Page Utilization, BP coverage and size
- **per-query metrics:**
  - based on required vs. observed I/Os
  - measure performance loss/overhead for each factor
- **per-node metrics:**
  - measure contributions to performance loss over entire workload
- **Penalty() and PickSplit() metrics:**
  - measure deterioration of workload performance

Leaf-Level Statistics

Global View

Provides summary of node statistics for entire tree

Tree View

Also displays node stats

Total or per query breakdown

I/O counts and corresponding overheads under various scenarios

Breakdown of losses against optimal clustering

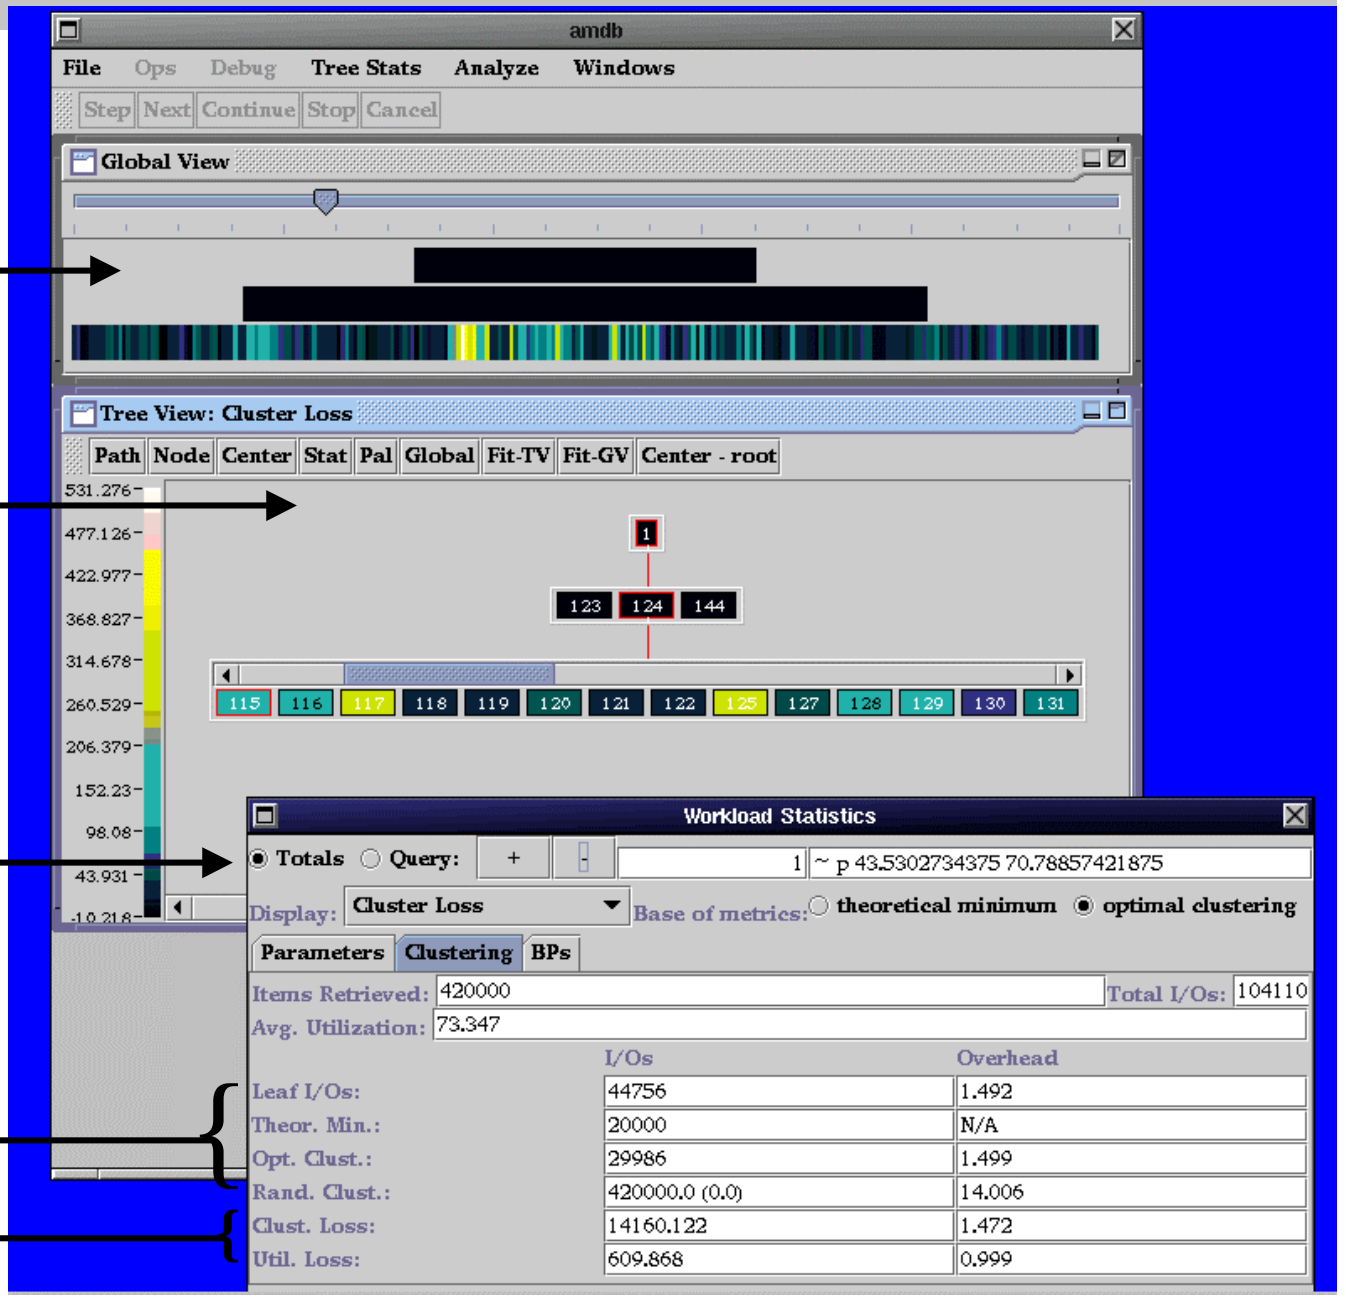

Bounding Predicate Statistics

Views highlight nodes traversed by query

Query breakdown in terms of "empty" and required I/O.

Excess Coverage
Overheads due to loose/wide BPs