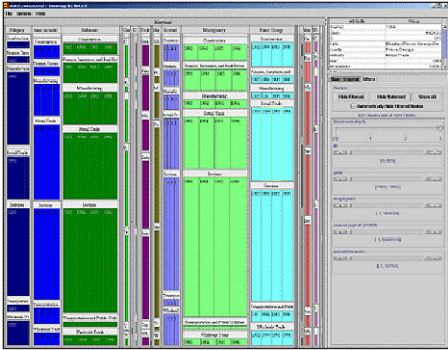


RESEARCH BRIEF

TREEMAPS: SPACE-CONSTRAINED VISUALIZATIONS OF HIERARCHIES

The potential



HCIL's Treemap 4.1 has rich controls for the layout, color, labeling and filtering of data items.

Many professions need to be able to analyze large, complex, hierarchical collections of quantitative data. Financial analysts, medical researchers, inventory managers, petroleum engineers, sports reporters, and criminal investigators, for example,

typically work with large data sets and need to be able to spot trends and exceptions in the data.

The ability to see tens of thousands of pieces of data in a fixed space and discover patterns, clusters and outliers is very powerful. Increasingly, the average person also can benefit from this kind of analysis, for example to organize large groups of digital photos or keep track of the stock market.

The challenge

The challenge is to develop a visualization interface that can dynamically create and manipulate hierarchies in large and complex data sets.

The research

In 1990, in response to the common problem of a filled hard disk, Dr. Ben Shneiderman decided to produce a compact visualization of directory tree structures. He explored ways to show a tree in a space-constrained layout, rejecting strategies that left blank spaces or those that dealt with only fixed levels or fixed branching factors. He developed a recursive algorithm that split the screen into rectangles in alternating horizontal and vertical directions as levels are traversed.

The Treemap visual interface for two-dimensional

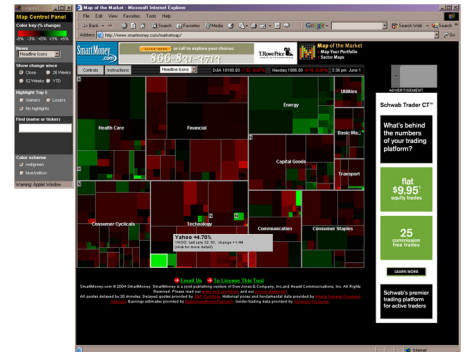
displays is the result. Treemap lets users view thousands of files at once as proportionately sized rectangles, which can be grouped in many different configurations.

Treemaps are a space-filling method of visualizing large hierarchical collections of quantitative data. A treemap works by dividing the display area into a nested sequence of rectangles whose areas correspond to an attribute of the data set, effectively combining aspects of a Venn diagram and a pie chart.

The algorithm used to create the nested rectangles that make up the map is a key ingredient. The slice-and-dice algorithm of the original space uses parallel lines to divide the rectangle representing an item into smaller rectangles representing its children. At each level of hierarchy the orientation of the lines—horizontal or vertical—is switched.

Treemap has unmatched utility for certain tasks. It provides many features to allow users to customize the display to their particular needs. There is a rapid overview of the relative size of nodes. Dynamic query filters allow users to filter out unwanted items by dragging sliders or selecting values with buttons. Hierarchies can be created using categorical or numerical attributes. The data can be aggregated by controlling the level of detail using a depth slider. Users also can specify what data attribute should be mapped to the size or color of the rectangles.

Recently, specialized techniques have been developed to visualize up to 1 million items on a treemap without aggregation.



SmartMoney.com's 'Map of the Market' displays the activity of more than 500 stocks in near-real time. The brighter the green, the more the price has gone up. The brighter the red, the more it has gone down. Black shows little to no change. The number of shares is indicated by the size of the rectangle. Passing the cursor over a stock brings up more information.

