	Event	**Entity**
credit cards	credit card transaction	credit card account
frequent shopper cards	purchases	shopper
telephone calls	call detail record	telephone number
air plane tickets	passenger name record	passenger
online browsing	viewing a web page	Internet user

Table 5.10: This tables contains some common examples of event-entity streams.

good statistical characterization of the cardholder can be generated.

Notice that this statistical characterization doesn't require using the name, address, or any other information that could identify the cardholder. All that is required is a mechanism for grouping together credit card transactions that belong to the same person. There are several ways of doing this that do not require using personal information.

Graphs. You can think of a graph as a diagram that consists of nodes, viewed as small circles, and edges, viewed as lines that connect the nodes. The web is a familiar example of data that can be viewed as a graph. Think of each web page as a node and connect two nodes with an edge in case there is a hyperlink from the first page to the second. With powerful enough web crawlers, you could compute a graph like this, consisting, say of 8 billion nodes and 20 billion edges.

Independent of the content of a page, the graph defined above can provide some interesting information about the importance of a page. For example, the more links to a page, the more authoritative you might expect the page to be. Google was the first commercial search engine to