

We present a technique to maintain the pairwise cosine similarities of documents arriving as a stream of individual words. Words from different source documents interleave in the stream to either initiate new documents or grow existing ones, and may arrive out-of-order relative to their sequence in their respective source document. The technique conforms to a single-pass, constant memory data stream model and needs no prior information about the size of the incoming vocabulary. When applied to detect anomalous graphs arriving as a stream of edges derived from system call traces of executing processes, the technique was demonstrably performant (10,000 edges/second) and accurate (0.95 average precision) while consuming bounded memory (256MB) that is also user-configurable. Project website: <http://bit.ly/streamspot>