Conventional solution and proof

First consider the usual external-graph model in which two N-entry rank vectors fit in memory but the E links do not. Keep the previous vector p fixed during a pass. Initialize $s = 0$ and dangling mass $d = 0$. Read each adjacency list once in its stored order. For a nonempty list of node i, add $p_i/\deg^+(i)$ to s_j for each outgoing neighbour j; for an empty list add p_i to d. Under uniform dangling completion set

$$p_{\text{new}} = (1-\alpha)(s + d\mathbf{1}/N) + \alpha\mathbf{1}/N.$$

Do not update p in place while reading: that would change the iteration. By the definition of the adjacency matrix, the accumulated $s + d\mathbf{1}/N$ is exactly $W^\top p$, so each pass realizes the PageRank fixed-point map. Stop when $\|p_{\text{new}} - p\|_1/\alpha \le \varepsilon$; this bounds the error of p, and also bounds the new error by contraction. The a priori bound $2(1-\alpha)^t$ supplies an alternative pass limit.

If even rank vectors do not fit, store rank records ordered by node identifier and merge them with the sequential adjacency stream. Emit $(j, p_i/\deg^+(i))$ contribution records, externally sort by destination, sum each destination group, add teleport/dangling mass and write the next rank file. Sorting and merging, or an equivalent external aggregation, are necessary to combine contributions to nonconsecutive destinations. This variant respects the inability to load the graph or the rank vector into memory.

Structural or analytic derivation

The graph traversal is a sparse matrix-vector product and the teleport term has rank one. The contraction proof separates mathematical accuracy from the storage representation and applies to either streaming variant in exact arithmetic.

Comparative cost and preferred method