

An Efficient Method for Large-scale Slack Allocation

S. Joshi and S. Boyd

Engineering Optimization, 41(12)1163-1176, Dec. 2009.

- [slack_alloc.pdf](#)

We consider a timing or project graph, with given delays on the edges and given arrival times at the source and sink nodes. We are to find the arrival times at the other nodes; these determine the timing slacks, which must be nonnegative, on the edges. The set of possible timing slacks is a polyhedron; to choose one we maximize a separable concave utility function, such as the sum of the logarithms of the slacks. This slack allocation problem, for which we give a simple statistical interpretation, is convex, and can be solved by a variety of methods. Gradient and coordinate ascent methods are simple and scale to large problems, but can converge slowly, depending on the topology and problem data. The Newton method, in contrast, reliably computes an accurate solution, but typically cannot scale beyond problems with a few thousand nodes. In this paper we describe a custom truncated Newton method that efficiently computes an accurate solution, and scales to large graphs (say, with a million or more nodes). Our method typically requires just a few hundred iterations, with each iteration requiring a few passes over the graph; in particular, our method has approximately linear complexity in the size of the problem. The same approach can be used to solve slack allocation problems with constraints, using an interior-point method that relies on our custom truncated Newton approach.

Page generated 2018-11-24 09:00:16 PST, by jemdoc.