

A Heuristic for Optimizing Stochastic Activity Networks with Applications to Statistical Digital Circuit Sizing

S.-J. Kim, S. Boyd, S. Yun, D. Patil, and M. Horowitz

Optimization and Engineering, 8(4):397-430, December 2007.

- [heur_san_opt.pdf](#)

A deterministic activity network (DAN) is a collection of activities, each with some duration, along with a set of precedence constraints, which specify that activities begin only when certain others have finished. One critical performance measure for an activity network is its makespan, which is the time required to complete all activities. In a stochastic activity network (SAN), the durations of the activities and the makespan are random variables. The analysis of SANs is quite involved, but can be carried out numerically by Monte Carlo analysis. This paper concerns the optimization of a SAN, *i.e.*, the choice of some design variables that affect the probability distributions of the activity durations. We concentrate on the problem of minimizing a quantile (*e.g.*, 90%) of the makespan, subject to constraints on the variables. This problem has many applications, ranging from project management to digital integrated circuit (IC) sizing (the latter being our motivation). While there are effective methods for optimizing DANs, the SAN optimization problem is much more difficult; the few existing methods cannot handle large scale problems. In this paper we introduce a heuristic method for approximately optimizing a SAN, by forming a related DAN optimization problem which includes extra margins in each of the activity durations to account for the variation. Since the method is based on optimizing a DAN, it readily handles large scale problems. To assess the quality of the resulting suboptimal designs, we describe two widely applicable lower bounds on achievable performance in optimal SAN design. We demonstrate the method on a simplified statistical digital circuit sizing problem, in which the device widths affect both the mean and variance of the gate delays. Numerical experiments show that the resulting design is often substantially better than one in which the variation in delay is ignored, and is often quite close to the global optimum (as verified by the lower bounds).