# Pathfinder: A Negotiation-Based Performance-Driven Router for FPGAs

Larry McMurchie, Carl Ebeling

**Year of publication:** *1995*
**Area:** *CAD*

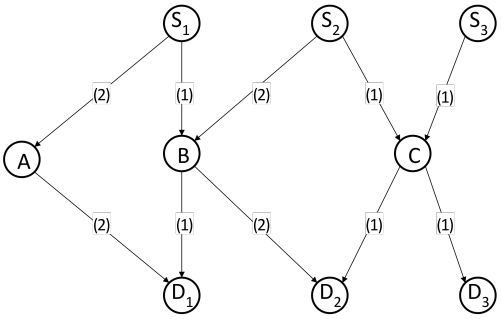


**Figure 2.** Second order congestion

I personally consider this to be the single most important paper for FPGAs at any technical conference (certainly at the FPGA Symposiums) in the past twenty years. This assertion is based on the accumulated impact of this paper on the FPGA industry and the academia alike. This paper changed the FPGA routing from a major headache with wildly fluctuating results to a reasonably well controlled optimization problem. Today, all FPGA vendors have routers in production that are based on Negotiated Congestion or based on some generalization of the idea. It is also the cornerstone of VPR, the most commonly used tool for Academic Research.

Some papers inspire the audience immediately. Others, like this one, go underappreciated for a while before their significance is fully understood. Most CAD papers propose some new idea, try it on some benchmarks and report 5% to 10% better results than what's available at the time. There are a very large number of such CAD papers. Most of them are good, but they make an incremental impact, which is frequently temporary in nature, until some other CAD paper reports slightly better results once again. Back in 1995, most FPGA researchers expected this to be another 10% step forward. (I certainly did.) Very few understood that this was not an incremental step: this was a game changing fundamental idea that will withstand the challenges of decades and will not be surpassed by any other router, except by its own extensions and generalizations. In the years followed, slowly but surely, the academia and the industry both understood the magnitude of the milestone achieved by the concepts advanced in this paper.

The paper starts by explaining the basic negotiation idea and how the first order congestion can be dealt with. Then the second order congestion is analyzed (see Figure 2 above) and the need for the "history cost" is introduced. Then the concept is generalized to account for routing delays. The paper ends with a pseudo-code of the algorithm and some experimental results (11% better than what was known from a commercial tool.) This is a very good paper, but it is not a model of clarity. This is the kind of paper that one needs to read many times, each reading revealing some nuance that was not apparent in previous readings.

Today, we are able to use the negotiated congestion widely and very successfully. Despite that, why it works so well eludes us at a theoretical level. (For example we rigorously understand why and how annealing works and converges. We do not understand negotiated congestion at an equivalent level.) We do not, yet, have full theory. As such, this paper will continue stimulating further research on this topic, experimental and especially theoretical.

Sinan Kaptanoglu

**DOI:** http://doi.acm.org/10.1145/201310.201328