Toward the accurate prediction of placement wire length distributions in VLSI circuits

Joni Dambre, Student Member, IEEE, Dirk Stroobandt, Member, IEEE and Jan Van Campenhout, Member, IEEE
Ghent University, ELIS Dept., Ghent, Belgium, e-mail: {jdambre, dstr, jvc}@elis.UGent.be

Abstract— Since its introduction, Donath's technique for predicting placement wire length distributions has become one of the most popular techniques for a priori wire length estimation. However, in its original form, it was heavily constrained by the underlying circuit and architecture models. In this paper, we show how a careful relaxation of those constraints results in very high correlations between predicted and experimentally measured average wire lengths as well as in a much improved accuracy in predicting wire length distributions.

Because the availability of the Rent characteristic is crucial for the quality of our model, we investigate how the prediction quality degrades when only an estimated characteristic is available. Such an estimation can be required to save computation time or when the complete netlist is not yet available (partial use of typical values). It turns out that a fitted β-model, based on only a few partitioning levels, can still result in a relatively high prediction quality. In particular with respect to the wire length distribution, the results are considerably better than when Rent's rule is used.

Index Terms— Interconnect prediction, design space exploration, Donath's wire length estimation technique, placement

I. Introduction

As technology pushes forward, feature sizes continue to decrease, but chip sizes tend to stay the same or grow even larger. As a consequence, the length of semiglobal and global wires increases, relative to the feature size and the relative cost of interconnect in the cost of VLSI chips is becoming ever more important [1]. This causes large problems for designers, because most interconnect properties relate to their length and only become known during physical design. Hence, they are very difficult to optimize before at least some part of physical design (e.g., the floorplanning stage) has been executed.

The study of the link between the interconnect topology of a (logic-level) circuit netlist and the cost of its physical implementation has been the main subject of system level interconnect prediction (SLIP) since the 1970's (introduction of Rent's rule, [2]). The earliest contributions to SLIP research provided models for the average wire length and the distribution of wire lengths in a circuit implementation [3], [4], [5]. Such predictions can be combined with models for length related interconnect costs such as delay, power dissipation, area or routability (e.g., [6], [7]). Since most of these models are non-linear functions of the wire length, estimating their average requires the wire length distribution [8], [9].

To be useful for speeding up the design process, interconnect prediction techniques should be applicable as early in the design cycle as possible. Hence, it must be possible to abstract them for use in earlier design stages, e.g., by quickly estimating characteristic parameters of the interconnect topology or by replacing them by typical values. Since the prediction quality unavoidably decreases toward higher abstraction levels, it is important that the logic-level predictors are as accurate and reliable as possible. In the context of design space exploration, i.e., the comparison of different design alternatives, the predicted cost metrics are reliable if they correlate well with measured costs after physical design, across implementations in different layout substrates (e.g., different materials, buffering strategies, wiring schemes, cell aspect ratios, ...) as well as across netlist alternatives.

The reality of state-of-the-art interconnect prediction is still far removed from this ideal picture. Rent's rule has become the corner stone of most current prediction techniques. However, as we will show in this work, predictions based on Rent's rule are not sufficiently accurate to obtain the desired correlation across different netlist topologies. This is one of the main reasons why, although some of the standard prediction techniques date from several decades ago, the majority of applications still lies in the area of technology evaluation. Well known examples are the International Technology Roadmap for Semiconductors [10], the Berkely Advanced Chip Performance Calculator (BACPAC, [11]), the Marco GSRC Technology Extrapolation (GTX) System [12] and evaluations of the benefits of using (optical) 3-D interconnect, or 3-D integration (e.g., [13], [14], [15]). All of these try to assess the impact of technological evolutions or variations in the geometry of the layout substrate for classes of circuits (e.g., general purpose processors).

Besides the bad correlation with individual circuit results, many other problems remain to be solved in the area of interconnect prediction. Some of the most pressing are the modelling of wires that connect more than two gates (multi-terminal nets, as opposed to two-terminal nets) and an accurate modelling of routing and routing congestion.[1]

In this paper, we address what we consider to be the first and most important step toward accurate interconnect prediction for design space exploration: the prediction of wire length distributions and average wire lengths for specific netlist topologies (with two-terminal nets only) and specific physical implementation substrates.

To obtain this goal, we turn to the existing work. Most current prediction techniques focus on predicting placement wire length distributions from a logic-level netlist. They are based on either Donath's [3] or Davis' [4] prediction

[1] In most existing work on interconnect prediction the impact of routing and routing congestion is not modelled.