

How E.F. Rent influenced ten years of interconnect research from the prompter's box

Abstract

Although Rent's rule has been formulated in the 1960's and published worldwide in 1971, it took until the end of the 1990's before its value became clearly apparent and instigated an entire research field of a priori wire length predictions. Even when the rule itself became quite popular, its inventor E.F. Rent stayed in the background and it remained a mystery who the man actually was. So it was as if Mr. Rent stayed in the prompter's box while others took the stage and spread the word on his findings. In this article, I will show that Rent's rule can be viewed as a fundamental law of nature which reflects on electronic circuits. As there are many interpretations of the rule, this article will shed some light on the core of Rent's rule and the research that has been built on it.

1. Introduction

The scaling of computer technology has been driven by the well-known *Moore's law*, stating that the number of transistors on a chip doubles every technology generation.¹ In the previous century, Moore's law also meant an increase in the clock frequency every technology generation (Figure 1). However, since recently, this is no longer the case and clock frequencies almost stay the same with new technology generations. Few people are aware of it, but the reason for this has its basis in *Rent's rule*.

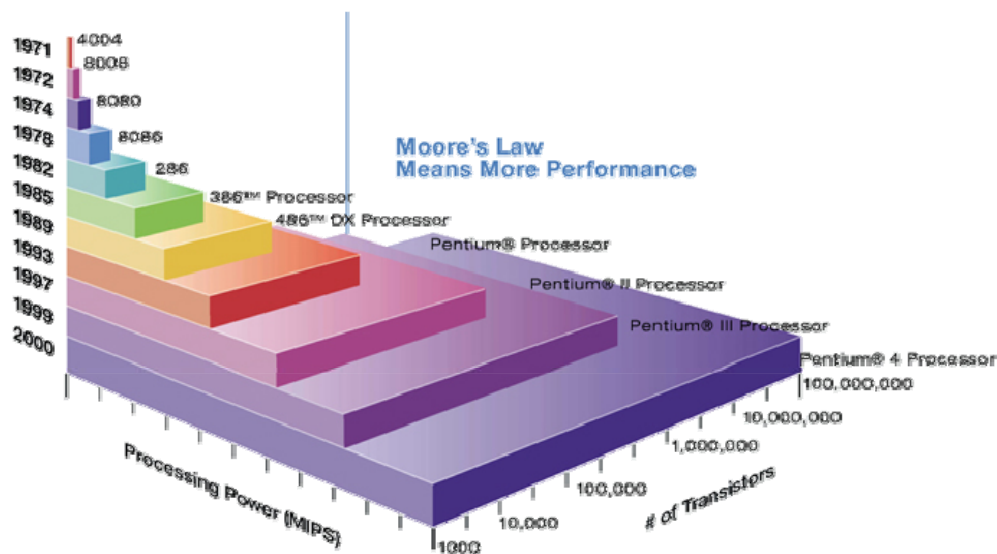


Figure 1: Moore's Law expresses the exponential progress of technology scaling as an increase in number of transistors and has had its effect on clock frequency as well.

Although there are many interpretations, Rent's rule basically states that, in a chip design, the number of wires emanating from a region containing B logic blocks (the basic computation elements on chips)

¹ In the 1970's, the number of transistors doubled every 18 months, later every 2 years, and the last decade every 3 years.

grows faster than the perimeter length increase when B grows (Figure 2). Note that this is explicitly a scaling argument and, in principle, assumes an infinitely large circuit. With some calculations, one can deduce from this that the average length of an interconnection on chip must increase with a growing number of gates [6]. Since the length of a wire has a large effect on the delay it induces on an electrical signal traveling the distance, the wire delay on chips has surpassed the transistor delay since the end of the 1990's. So while transistors keep on gaining in speed with every technology generation, the wires in between them induce a relatively larger delay and this has stopped the clock frequency upscaling.

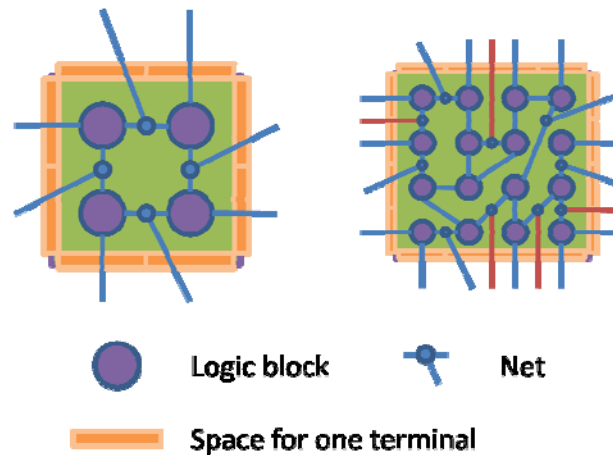


Figure 2: Rent's rule states that the number of terminals (wires emanating from a region) grows faster than the perimeter length increase when the number of blocks grows. If the space for one terminal on the boundary scales linearly with the block size, there is a shortage of space for all terminals (red lines lack space).

Rent's rule, first formulated in the 1960's, is not nearly as well-known as Moore's law but it is of a much more fundamental nature. Where Moore's law was a mere observation that has become a self-fulfilling prophecy (with major ASIC technology companies driving their roadmaps in accordance to it), Rent's rule has been largely neglected for a long time. Yet, there is no way to circumvent this fundamental rule and so it had a detrimental effect on clock speeds of new computer systems.

In this article, I will explain Rent's rule in more detail (in section 2), focus on the wire length estimations that result from it (in section 3) and present an overview of the myriad of research activities that sprouted from the initial research work on a priori wire length estimates (section 4). Section 5 wraps up with a short outlook on the future of Rent's rule.

2. Rent's rule

In the 1960's, IBM employee E.F. Rent wrote an internal memo where he described what later became known as Rent's rule. It wasn't until two other IBM employees, Landman and Russo, wrote a landmark paper in 1971, that Rent's rule was made public [5]. It is actually surprising and a bit mysterious that E.F. Rent never published his findings outside of IBM himself and it remains unknown (to me at least) why his name was not on the paper written by Landman and Russo as well (hence the title of this article).

2.1. The terminal-gate relationship of Rent's rule

In their paper, Landman and Russo discuss ways of optimally partitioning a circuit into modules in such a way that as few as possible interconnections between the modules are cut during the partitioning. The

rationale behind this way of partitioning is that connected gates in different modules will be placed further apart than the gates within each module and hence this partitioning strategy will keep the connections shorter. In such a partitioning strategy, Rent's rule relates the average number of terminals T of a part of any circuit (a module) and the average number of logic gates (basic logic blocks B) inside the module as $T = tB^p$. The parameter t is the average number of terminals per logic gate² and the exponent p is called the *Rent exponent*. Its value depends on the complexity of the interconnection topology (with higher values for more complex topologies) and on the quality of the placement (with higher values for less placement optimization). The maximal value of the Rent exponent p is 1 for a very complex topology or a random placement [1].

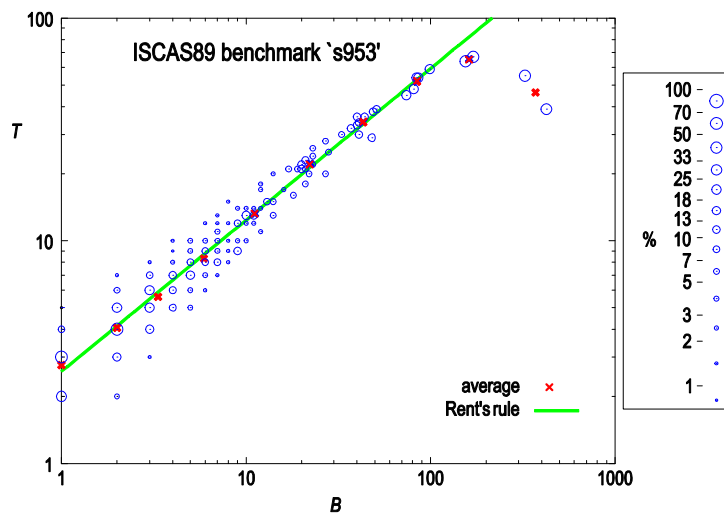


Figure 3: Rent's rule: number of terminals per module T versus number of gates per module B during the partitioning of a benchmark circuit (ISCAS89 benchmark 's953'). The size of the circles corresponds to the percentage of modules (on a total number of modules around an average number of gates, at equal distances in the log-log plot) that has B gates and T terminals.

Rent's rule was found by experimental analysis of many circuit partitions and proves to be valid for most designs. Figure 3 shows the result of a circuit partitioning in a log-log plot of number of terminals versus number of logic gates. The validity of Rent's rule follows from the fact that all points follow – on average – a straight line in the plot.

Note that there is a deviation from the straight line for high values of T and B which is known as Rent's region II and has been described in [1] and [5]. This deviation at the chip boundary is a direct result of the nature of Rent's rule itself. For circuits with an interconnection complexity p larger than 0.5, the number of pins (terminals at the outermost boundary) scales faster than the perimeter of the boundary (see Figure 2). In practice, the number of pins at the boundary of a chip is limited and hence the number of signals going out is intentionally lowered by techniques such as serialization of the information stream or encoding of the information in less bits. Therefore, the actual number of pins on a chip is significantly lower in real circuits than Rent's rule predicts. It is one of the many misconceptions about Rent's rule (and one that lead to false conclusions in many papers) that Rent's rule fixes the relation between the number of pins of a circuit and its number of internal blocks. Rather, Rent's rule is based on a scaling argument and really captures only the internal interconnection complexity. Another deviation at the low end has also been observed in [8] but this is much less frequent.

² If $B=1$, Rent's rule reduces to $T=t$.

2.2 Interpretations of Rent's rule

In the previous section, we in fact already presented two different interpretations of Rent's rule: the one presented in Figure 2 (about the scaling of the number of terminals according to the number of internal blocks if this number grows) and the interpretation based on an "optimal" partitioning (resulting in Figure 3). In fact, Rent's rule has been explained as a fundamental scaling law [10] and it has been shown in [1] that Rent's rule applies to any homogeneous design.

A less obvious interpretation, but the one that gives rise to all major uses of Rent's rule, is that it is a measure of interconnection complexity. The reasoning here is that a more complex structure of the wires between the logic blocks means that there are more wires that connect blocks that are less close (in terms of graph distances or, after placement, in terms of actual distances). Hence, it is harder to place such a circuit with short wire lengths. A circuit with simple interconnection complexity then is a circuit where all wires are between blocks that are close to each other (again in terms of graph distance). E.g., planar graphs are simpler than non-planar graphs because it is easier to place such circuits in 2D. It is clear that in this interpretation, a simpler circuit can be partitioned more easily with less connections to be cut than a more complex circuit. Therefore, the average number of terminals in a partitioned complex circuit will be higher than for a simple circuit and the Rent exponent will naturally be higher. Therefore, the Rent exponent is a measure of the interconnection complexity.

The complexity interpretation also gives rise to a relation between the Rent exponent p and a fractal dimension D that describes the geometric dimension that would be the "ideal" fit for implementing the circuit [10]. One can deduct (through partitioning) that for any D -dimensional mesh the Rent exponent is given by

$$p = \frac{D - 1}{D}$$

3. Interconnect length predictions

The main claim of fame for Rent's rule has come from Wilm Donath's 1979 paper on wire length estimation [3] and a followup paper in 1981 [4]. In these papers, Wilm Donath, another IBM employee, used Rent's rule to predict the average wire length and wire length distribution in computer chips, before the actual layout. The basic idea is simple and uses a partitioning scheme as the basis of the estimation. The three main models for the layout generation are (see Figure 4) (i) a circuit graph model with Rent's rule as the model of its interconnection complexity, (ii) a Manhattan grid architecture model where the circuit will be placed and routed, and (iii) a model for the placement and routing of the circuit on the architecture.

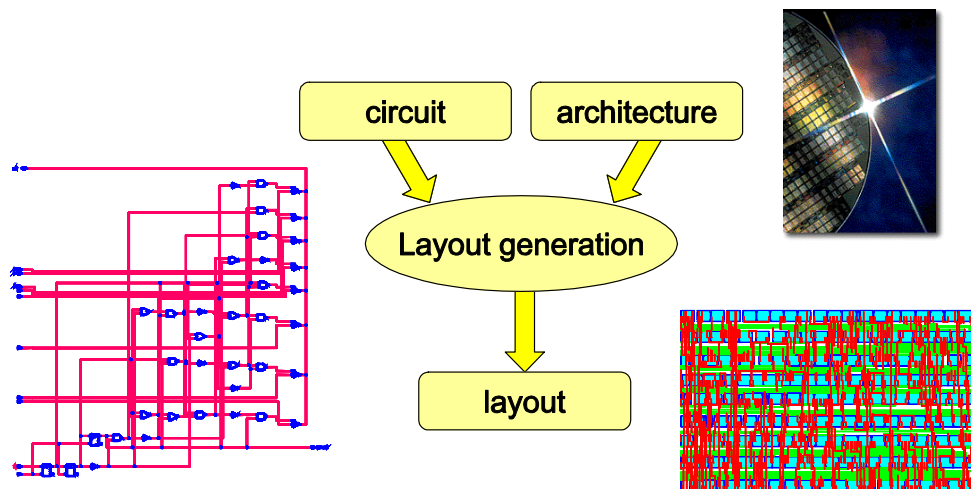


Figure 4: The three components of models for physical design: the circuit, the architecture and the layout generation. The combination of these models results in the (model for the) layout.

3.1. Donath's wire length prediction model

Donath's model is basically a hierarchical partitioning of both the circuit and the Manhattan grid architecture. The model starts with a partitioning of the circuit in 4 equal parts in such a way that the number of nets cut is minimized (Figure 5(a)). At the same time, the Manhattan grid is partitioned in 4 equal subgrids which are the four quadrants (Figure 5(b)). Then each of the subcircuits is mapped to a subgrid and for each of the subcircuit/subgrid pairs the partitioning steps are repeated until each subcircuit only contains a single gate and each subgrid contains a single cell location.

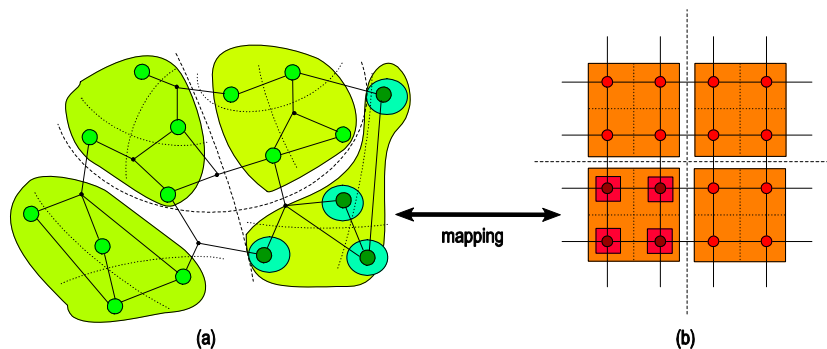


Figure 5: Donath's placement model: recursive partitioning of both circuit (a) and Manhattan grid (b) and mapping of circuit parts to grid parts.

The partitioning process as described above ensures that the number of longer interconnections between large subblocks is minimized in favour of shorter interconnections between smaller blocks (inside the larger blocks). This is exactly the same partitioning as was assumed in Rent's rule. Hence we can use Rent's rule and the corresponding Rent exponent as an estimate of the complexity of the interconnection structure.

Without delving into the mathematical details (see [3] and [10] for this), it is clear to see that, in each partitioning step, Rent's rule can be used to estimate the number of terminals from the number of gates inside the subcircuit. As each terminal represents a wire going out of the module under investigation and given the fact that two terminals are needed to represent one wire that is cut (under the restriction of two-terminal nets only) one can easily deduce the number of wires crossing the module boundaries at a

certain partitioning level. This number also contains the number of wires crossing the boundary from the previous partitioning in the hierarchy so one has to subtract that number to obtain the number of wires cut at each partitioning level.

The average length of a connection at a hierarchical level was assumed by Donath as the average of all possible connection lengths between each and every point from one subgrid to another one from the same partitioning level. The summation over all partitioning levels of these average lengths, weighted with the number of wires cut at that level, results in an estimate of the average wire length within the circuit. The detailed calculations can be found in [3] and [10].

3.2. Improved wire length prediction models

The placement and routing models used in Donath's prediction technique are very simple. The placement is modeled using the hierarchical partitioning model (which makes sense as partitioning is actually used in partitioning based placers to induce optimality of the wire lengths). The routing model is very simple as well as every connection is assumed to be routed as the shortest wire between its two endpoints. This is a very common assumption and provides at least a lower bound on actual wire lengths. However, the simple partitioning based placement model has its drawbacks. The main reason for this is the relatively large granularity of the partitioning steps. It is reasonable to assume wire lengths will be longer when gates are in different partitions. However, this is not necessarily the case. Two gates placed in different modules but near their border (at both sides) can have a much shorter length than two connected gates within the same module. Donath's simple model does not take this into effect as it assumes all possible interconnections within one partitioning level as likely as any other one (see Figure 6). This leads to an overestimation of the average wire length in Donath's model by a factor of 2 approximately, as has been noted by several authors in the 1980's and 1990's (an evaluation can be found in [9] and [10]).

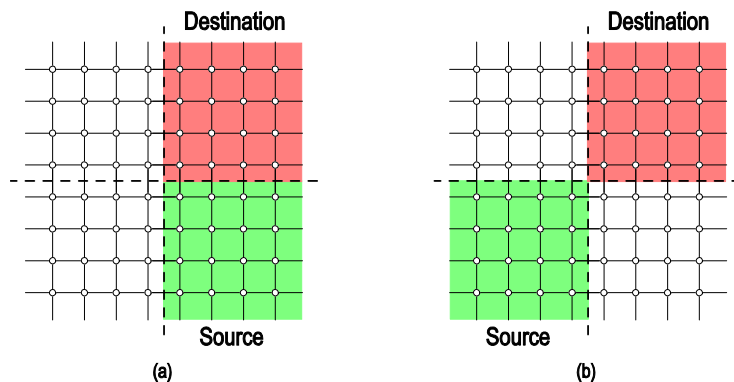


Figure 6: Donath's placement model: the average length on a hierarchical level is estimated by assuming that source and destination cells are uniformly distributed over the grid cells within the partition. We distinguish adjacent combinations (a) and diagonal combinations (b).

In my own Ph.D. research work [7] I found a way to remedy this discrepancy between the model and actual measurements by noting that an optimal placement (whether partitioning based or not) will prefer shorter wires over longer ones and will hence place gates in two modules of the same partitioning level preferably near the border of the modules. This has been represented in Figure 7 by a darker shade for more likely gate positions. We thus need a probability distribution for the placement of source and destination cells for all wires within a hierarchical level.

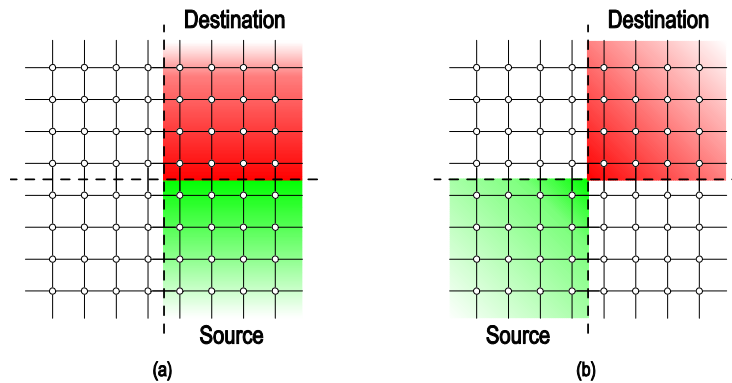


Figure 7: Placement of source and destination cells according to a probability distribution (darker regions have higher probability) for adjacent (a) and diagonal (b) combinations.

In [6] and [7], I deduced a probability distribution based on the overall wire length distribution found by Donath and statistical arguments that the local distributions should follow the same trends. It was very surprising to see that the exact same result was found around the same time by Jeff Davis at Georgia Tech [2] using a very different approach and another interpretation of Rent's rule. It even took until 1999 before we actually found out through careful analysis [1] that our results were essentially the same. It was undoubtedly this improved understanding of Rent's rule and its effects on wire length estimations that has re-ignited research on this fundamental rule of scaling providing us a figure of measure for the complexity of the interconnection structure of circuits.

4. The dawn of a new decade of system-level interconnection research

Although Rent's rule was known since the early 1970's, not a lot of work (research at universities or development at companies) was based on it in the next 30 years, not even after Donath presented the very interesting application of Rent's rule in a priori wire length predictions. Of course, the rule was mentioned in a few papers and there were some individual researchers who actually used it but we can hardly speak of a widespread proliferation of the rule at that time. This changed however with the introduction of the System Level Interconnect Prediction (SLIP) Workshop in 1999 [11]. This workshop started with a clear focus on Rentian interconnection models (our first keynote speaker in 1999 was Wilm Donath!) and gradually moved to a breeding ground for research on the boundary between physical interconnect modeling and interconnect technology, the impact of interconnects on CAD, and architectural interconnect issues.

Historical note

The idea of the SLIP workshop series came out of a form of frustration about the fact that my own research on Rent-based interconnection models always ended up at the strangest sessions at conferences because it was deemed interesting (hence it was accepted) but the only work of that kind. At the same time, there was a clear need for more interconnect-related research as the dominance of wire delays over gate delays began to show in real designs. When I first presented the idea of a new workshop to Prof. Andrew Kahng (who was then at UCLA), I meant it to be a workshop primarily on Rent's rule based research. Prof. Kahng rightfully thought the scope should be a bit broader than that and came up with the name SLIP, System level Interconnect Prediction workshop. But Rent's rule has played a major role in the workshop ever since.

The presentation of Rent-related work and a tutorial on Rent's rule at SLIP 99 seems to have ignited new research programs at several universities worldwide. It is impossible to list all uses of Rent's rule but in

the following I briefly state some of the topics where Rent's rule was used, with an indication of important contributors and the year they published their work at SLIP or the IEEE Transactions on VLSI (where several special issues on SLIP research have been published). Many more research papers addressing Rent's rule can be found on the world wide web.

- The interpretation and derivation/measurement of Rent's rule (Stroobandt, UGent, SLIP99/SLIP01/SLIP03; Davis, Georgia Tech, IEEE Trans. on Electron Devices 1998; Donath, IBM, SLIP99; Christie, Univ. Delaware, SLIP99)
- Improvements/validation of interconnection length models (Christie, Univ. Delaware, SLIP00/SLIP02; Najm, Toronto, SLIP00/SLIP03; Stroobandt/Dambre, UGent, SLIP99/SLIP00/SLIP01/SLIP02; Davis, Georgia Tech, SLIP00/SLIP03/SLIP06; Sarrafzadeh, Northwestern Univ., SLIP01; Otten, TU Delft, SLIP01; Cheng, UCSD, SLIP01/SLIP03; Bennebroek, Philips Research, SLIP03; Bhatia, UT Dallas, SLIP03, Zarkesh-Ha, LSI Logic, SLIP04; Chrzanowska-Jeske, Portland State University, SLIP04; Lanzerotti, IBM, SLIP05/SLIP07; Amakawa, Tokyo Institute of Technology, SLIP07; Behjat, Univ. Calgary, SLIP09)
- Generation of synthetic benchmark circuits (Stroobandt, UGent, TCAD vol.19, no.9 2000)
- Wire length models for 3D systems (Rahman, MIT, SLIP99/SLIP01; Saraswat, Stanford, SLIP00; Davis, Georgia Tech, SLIP00; Chandrakasan, MIT, SLIP05; Christie, Univ. Delaware, SLIP05)
- Timing estimations (Christie, Univ. Delaware, SLIP02; Amakawa, Tokyo Institute of Technology, SLIP05; Brown, Altera Toronto, SLIP06; Luk, Imperial College London, SLIP08)
- Routing/routability/congestion optimization (Chong, UC Berkeley, SLIP99; Stroobandt/Kahng, UCLA, SLIP00; Scheffer, Cadence, SLIP00; He, Univ. Wisconsin, SLIP01; Kahng/Stroobandt, UCLA, TCAD vol.20, no. 5, 2001; Sapatnekar, Univ. of Minnesota, SLIP02; Teig, Simplex Solutions, SLIP02; Becer/Blaauw, Motorola/Univ. of Michigan/Univ. of Illinois, SLIP02; Christie, Univ. Delaware, SLIP02; Kravets/Kudva, IBM, SLIP03; Karypis Univ. Minnesota, SLIP03, Kahng, UCSD, SLIP03; Groeneveld, TU Eindhoven, SLIP05; Sarrafzadeh, UCLA, SLIP07)
- Placement optimization (Cong, TVLSI vol. 9, no 6, 2001; Christie, TVLSI vol. 9, no. 6, 2001 ; Marek-Sadowska/Xilinx, UCSB, SLIP01/SLIP02/SLIP03)
- Floorplanning (Sarrafzadeh, Northwestern Univ., SLIP99)
- Manufacturability and yield (Christie, Univ. Delaware, SLIP01; Zarkesh-Ha, LSI Logic, SLIP03; Zarkesh-Ha, Univ. New Mexico, SLIP07)
- Rent-based system/architectural analysis and technology extrapolations (Sylvester, UC Berkeley, SLIP99; Kahng, UCLA, DAC 2000; Rose, Rensselaer Polytechnic Institute, SLIP01; Hutton, Altera, SLIP01/SLIP03; DeHon, CalTech, SLIP01; Maex, IMEC, SLIP02/SLIP04; Cheng, UCSD, SLIP03; Bergamaschi, IBM, SLIP04; Kumar, Cornell University, SLIP04; Greene, Actel, SLIP06)
- On-chip power distribution/optimization (Friedman, Univ. of Rochester, SLIP02; Nassif, SLIP02; Kolodny, Intel, SLIP04; Saraswat, Stanford, SLIP04; Kahng/Sylvester, UCSD/Univ. Michigan, SLIP04; Schmit, CMU, SLIP04)
- Networking and NoCs (Muddu, Sanera Systems, SLIP02; Verbauwhede, UCLA, SLIP02; Tenhunen, KTH Sweden, SLIP03; Davis, Georgia Tech, SLIP04; Burlison, Univ. of Massachusetts, SLIP04; Kolodny, Technion-Israel Institute of Technology, SLIP07; De Micheli, Ecole Polytechnique Federale de Lausanne, SLIP07; Smit, Univ. Twente, SLIP08; Heirman/Dambre/Stroobandt, UGent, SLIP08; Reda, Brown University, SLIP09)
- Optical systems (O'Connor, Ecole Centrale de Lyon, SLIP04; Heirman/Dambre/Stroobandt, UGent, SLIP05/SLIP06)

Apart from the papers presented at SLIP, there have been a number of special issues of journals dedicated to System Level Interconnect Prediction: IEEE Transactions on VLSI Systems published one in

2002 (vol.10, no. 2), in 2003 (vol. 11, no. 1), and 2007 (vol. 15, no. 8). Integration, the VLSI Journal, had a special SLIP issue in 2007 (vol. 40, issue 4)

As can be seen from the above list, several companies also became involved in discussing and using Rent's rule within their company: notably IBM but also Cadence, Altera, Xilinx, Simplex Solutions, Sanera Systems, Intel, Actel ...

Historical note

Before SLIP, I had read a few papers from prof. Phillip Christie (then at the University of Delaware) but I had never met him. He was, at the time, the one person in the world that had written a serious of papers related to Rent's rule. He came to SLIP 99 and had a presentation there but what I remember most is our hours-long discussion on Rent's rule in the local Irish pub that Saturday evening. Our interpretations of Rent's rule were very different yet so alike. The next morning Phillip asked me for 15 minutes of the workshop program time to explain to the audience the unifying interpretation that he came up with during the night after our discussions. The truly workshop-like atmosphere allowed for such an intervention and laid the basis for the highly cited paper [1] we wrote together afterwards.

5. An outlook on future interconnect research issues

As stated in the previous sections, Rent's rule has found its way mainly to a priori interconnect length estimation and related extrapolations. As wire lengths are ever more dominating circuit delays as well as power and area usage, the importance of interconnects will surely stay. One can question if Rent's rule will still be sufficient as a basis of the predictions as the accuracy of Rentian predictions may not be high enough. If one needs accurate estimates, one needs to revert to simulation and actual (albeit fast) synthesis methods. Using actual synthesis as an estimation may not be as problematic as it used to be and hence will probably gain more importance over using Rent's rule. However, as hardware design is moving up to higher hierarchy levels (ESL or Electronic System Level design), the early steps of architecture exploration before any synthesis has been done will again require very fast and simple estimates to weed out the inferior solutions and keep only the promising ones. In this domain, a simple estimate based on Rent's rule, even if not very accurate, provides the only plausible way to obtain estimates that are fast enough. It is difficult to predict the future but I believe Rent's rule has a bright future as design is moving to higher abstraction levels. But even if we risk to forget about its power, the rule fundamentally governs our designs and we will be forced to listen to the prompter and take it into account.

References

- [1] P. Christie and D. Stroobandt, "The Interpretation and Application of Rent's Rule", IEEE Trans. on VLSI Systems, Special Issue on System-Level Interconnect Prediction, vol. 8, no. 6, December 2000, pp. 639–648. IEEE Computer Society Press.
- [2] J.A. Davis, V.K. De, and J.D. Meindl, "A Stochastic Wire-Length Distribution for Gigascale Integration (GSI) – Part I: Derivation and Validation", IEEE Trans. on Electron Devices, vol. 45, no. 3, pp. 580 – 589, 1998.
- [3] W.E. Donath, "Placement and Average Interconnection Lengths of Computer Logic", IEEE Trans. Circuits & Syst., CAS-26, pp. 272 – 277, 1979.

- [4] W.E. Donath, "Wire Length Distribution for Placements of Computer Logic", IBM J. of Research and Development, vol. 25, pp. 152 – 155, 1981.
- [5] B.S. Landman and R.L. Russo, "On a Pin Versus Block Relationship For Partitions of Logic Graphs", IEEE Trans. on Comput., vol. C-20, 1971, pp. 1469-1479.
- [6] D. Stroobandt, H. Van Marck, and J. Van Campenhout, "An Accurate Interconnection Length Estimation for Computer Logic", Proc. 6th Great Lakes Symp. on VLSI, pp. 50 – 55. IEEE Computer Society Press, 1996.
- [7] D. Stroobandt, "Analytical Methods for a priori Wire Length Estimates in Computer Systems." Ph.D. thesis, Faculty of Engineering, Ghent University, 1998.
- [8] D. Stroobandt, "On an efficient method for estimating the interconnection complexity of designs and on the existence of region III in Rent's rule", Proc. 9th Great Lakes Symposium on VLSI, pp. 330 – 331. IEEE Computer Society Press, March 1999.
- [9] D. Stroobandt and J. Van Campenhout, "Accurate Interconnection Length Estimations for Predictions Early in the Design Cycle", VLSI Design, Special Issue on Physical Design in Deep Submicron, vol. 10, no. 1, pp. 1 – 20, 1999.
- [10] D. Stroobandt, "A Priori Wire Length Estimates for Digital Design", Kluwer Academic Publishers. ISBN 0-7923-7360-X. 2001. pp. 298
- [11] Presentations from previous SLIP workshops can be downloaded from <http://www.SLIPonline.org>

About the author



Dirk Stroobandt obtained the Ph.D. degree in electrotechnical engineering in 1998 from Ghent University. Since 2002, he is Professor at Ghent University, affiliated with the Department of Electronics and Information Systems (ELIS). He currently leads the Hardware and Embedded Systems (HES) research group of about 10 people with interests in semi-automatic hardware design methodologies and tools, run-time reconfigurable hardware and applications and reconfigurable multiprocessor networks.

Dirk Stroobandt is the inaugural winner of the ACM/SIGDA Outstanding Doctoral Thesis Award in Design Automation (1999) and also received the 'Scientific prize Alcatel Bell' for his work on "Structural and behavioural aspects of short optical interconnects in electronic systems" in 2002.

He was visiting researcher at the lab of Prof. Fadi J. Kurdahi at the University of California at Irvine (1997) and post-doctoral researcher at the group of Prof. Andrew B. Kahng at the University of California at Los Angeles (UCLA) (1999-2000).

Dirk Stroobandt initiated and co-organized the International Workshop on System-Level Interconnect Prediction (SLIP) since 1999. He is guest editor of two special issues of the IEEE Transactions on VLSI Systems on System-Level Interconnect Prediction and a special issue on SLIP for Integration, the VLSI Journal.