



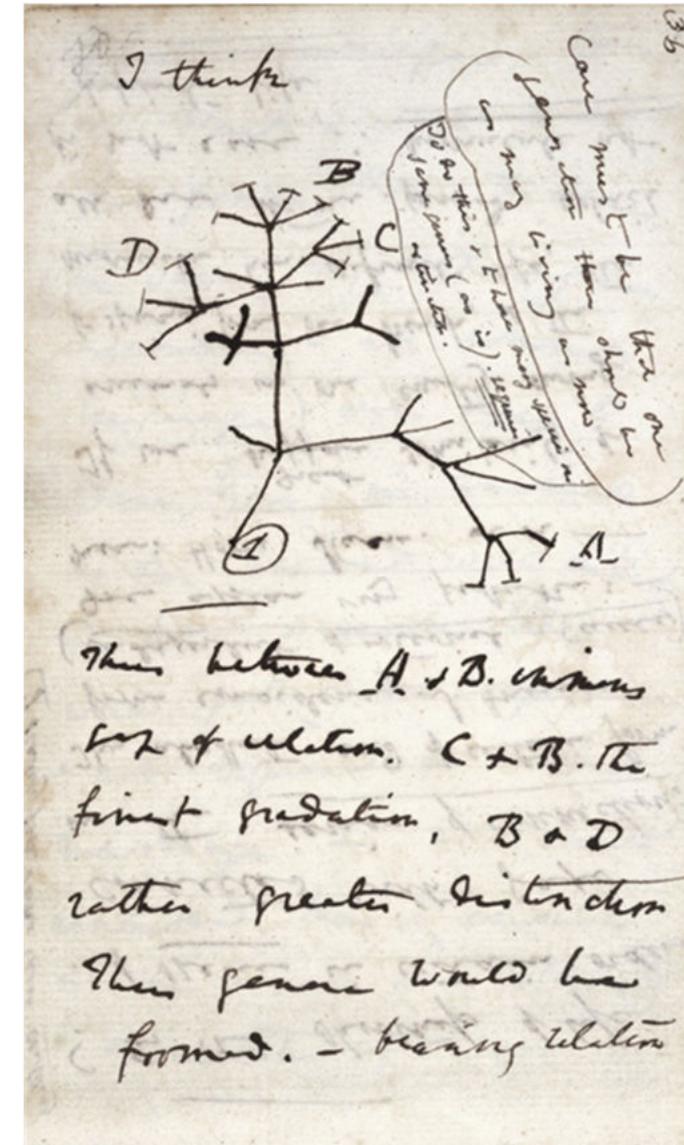
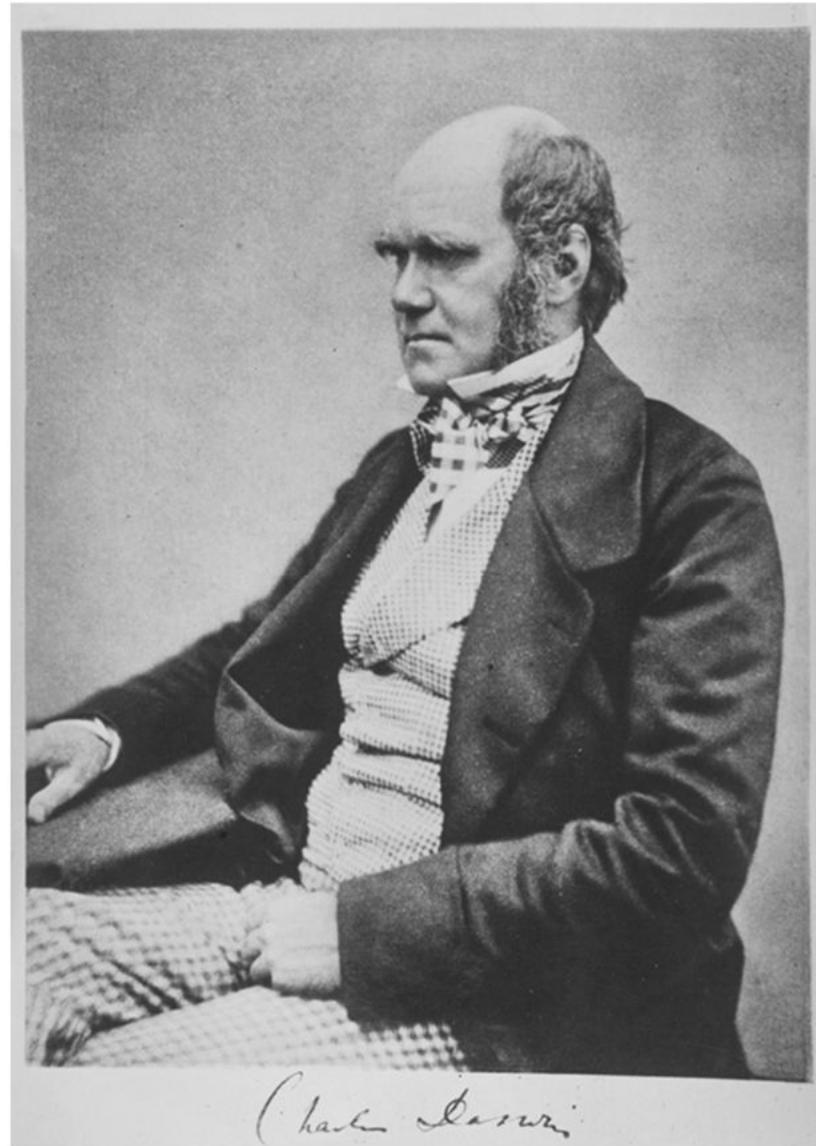
The Evolution Of Logic Synthesis

Dots and Dashes, ... Zeroes and Ones

Dr. Antun Domic
EVP & GM, Design Group
Synopsis, Inc.

Charles R. Darwin

“The Great Tree of Life”



Smartphones & Logic Synthesis

Have More in Common than You May Think !

Smartphones



Logic Synthesis

Claude E. Shannon
"A Symbolic Analysis of Relay and Switching Circuits"

Logic Compiler, ca. 1986
Optimal Solutions, Inc. a.k.a. Synopsys, Inc.

Power Compiler, 1997
Convergence of Logic & [Dynamic] Power Synthesis

Physical Compiler, 1999
Convergence of Synthesis & Implementation

Design Compiler, 2001-2015
The Evolution of Synthesis !

	Area	Timing	Power		Runtime
			Dynamic	Static	
2001	100	100	100	100	100
2005	78	91	82	78	20
2010	65	74	66	49	1.6 ⁽¹⁾
2015	53	63	66	38	0.6 ⁽¹⁾

(1) Design Compiler Multicore

© 2015 Synopsys, Inc. 23 Source: Synopsys Research 2015

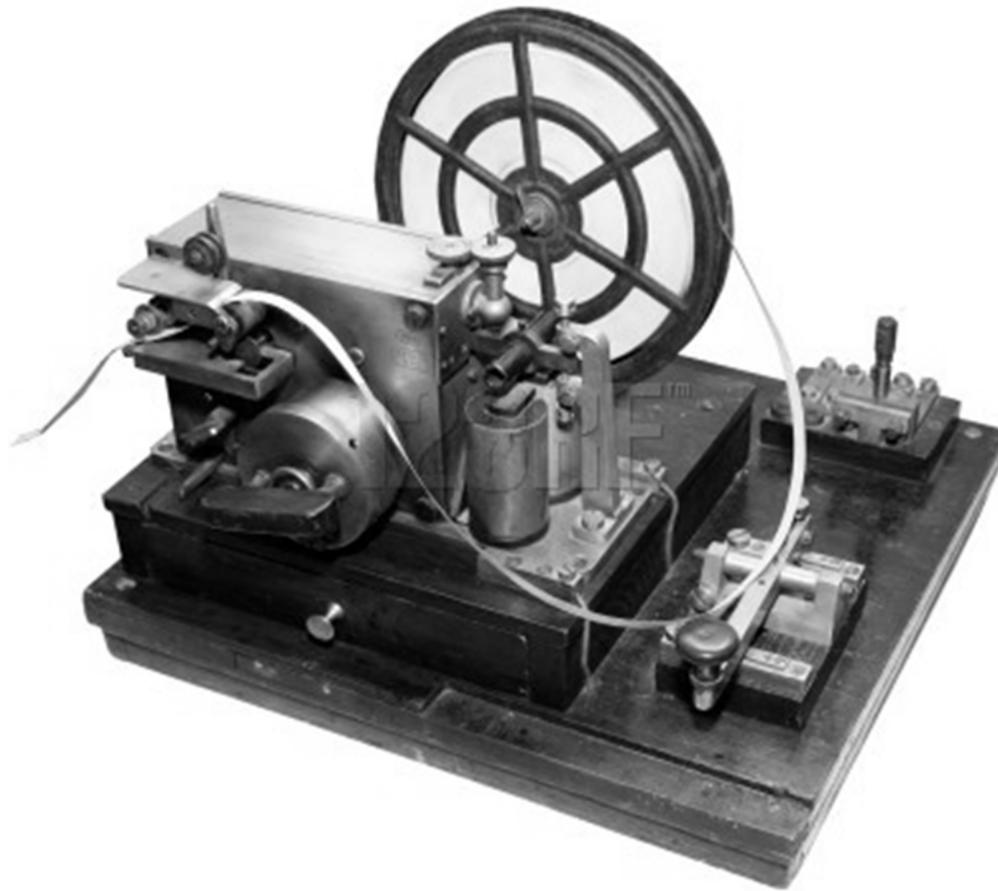
SYNOPSYS

Once Upon A Time...

Samuel F.B. Morse's Telegraph and George Boole's Algebra

Dots and Dashes...

...Zeros and Ones

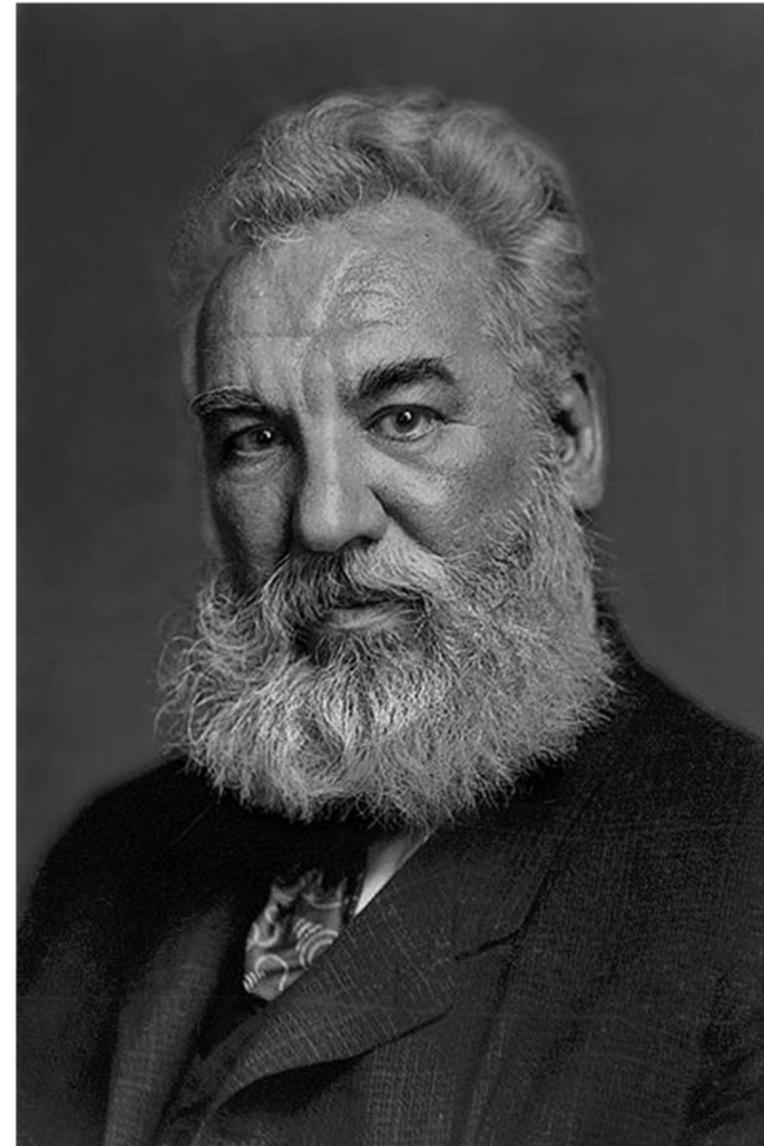
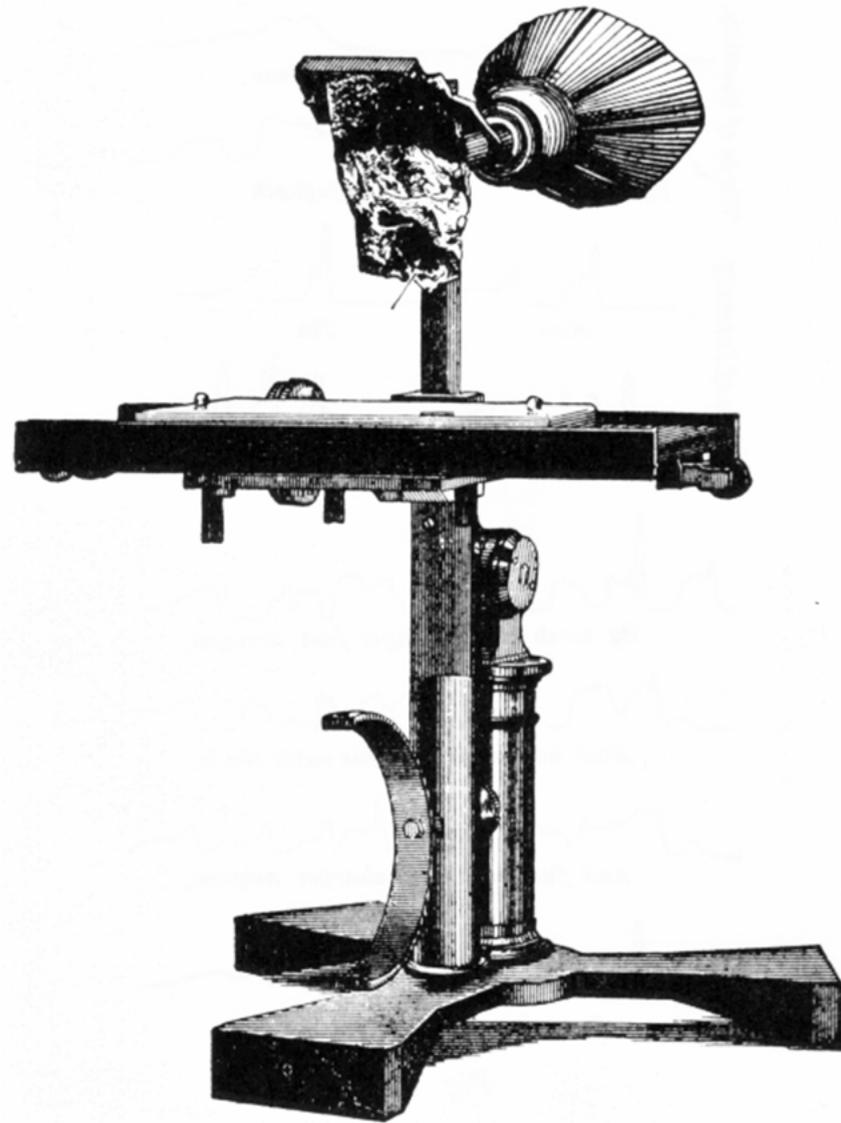


“An Investigation of the Laws of Thought, on Which are Founded the Mathematical Theories of Logic and Probabilities”, 1854

$$\begin{aligned} f(A, B, C, D) = & \\ = & (\overline{A}BC\overline{D}) + (\overline{A}BCD) + (\overline{A}\overline{B}C\overline{D}) + (\overline{A}\overline{B}CD) + \\ & + (\overline{A}B\overline{C}D) + (\overline{A}BC\overline{D}) + (\overline{A}B\overline{C}D) + (\overline{A}BC\overline{D}) \end{aligned}$$

Once Upon A Time...

A. Graham Bell's Telephone – the "Speaking Telegraph"



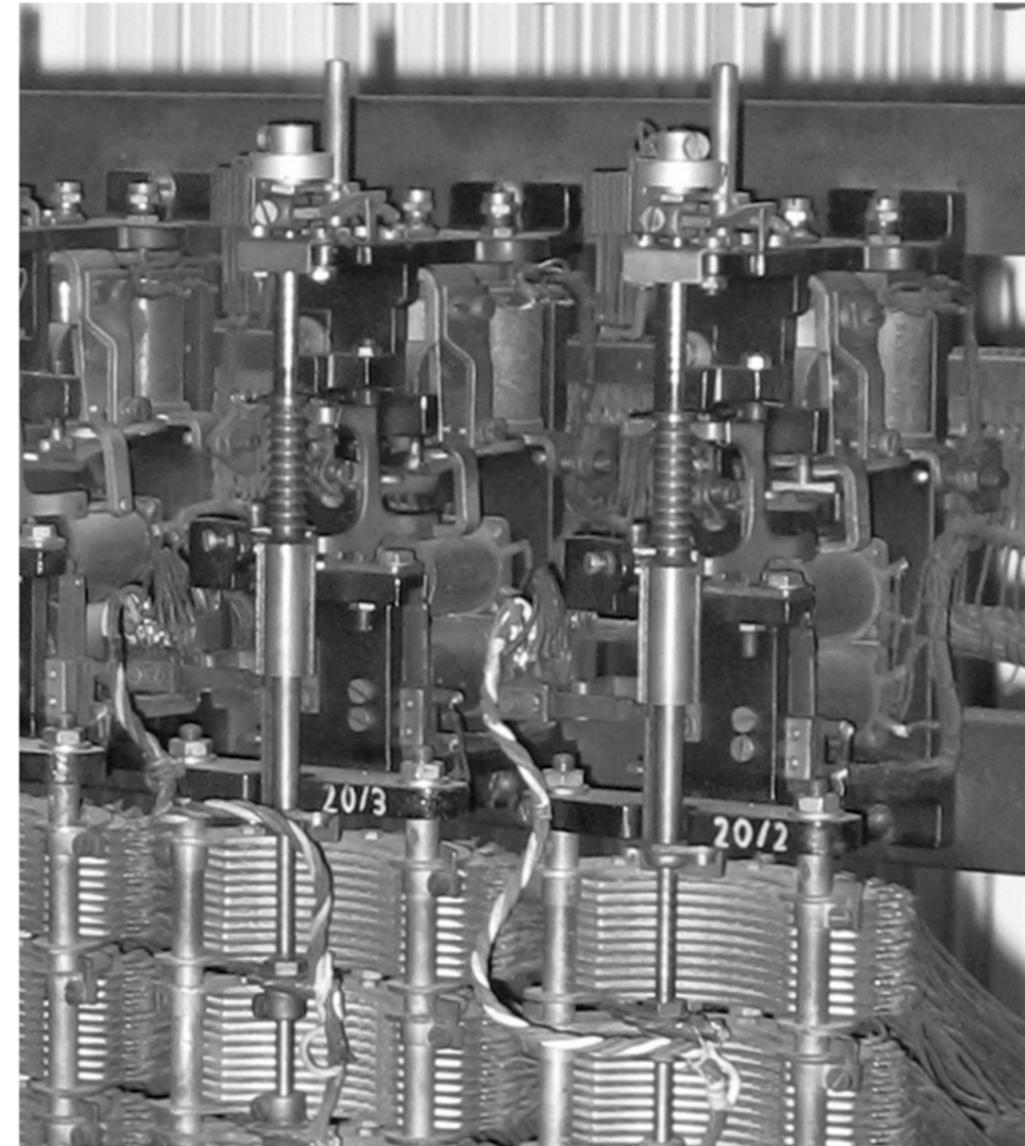
Once Upon A Time...

Bell Labs Radio-Telephone – the “Mobile Phone”



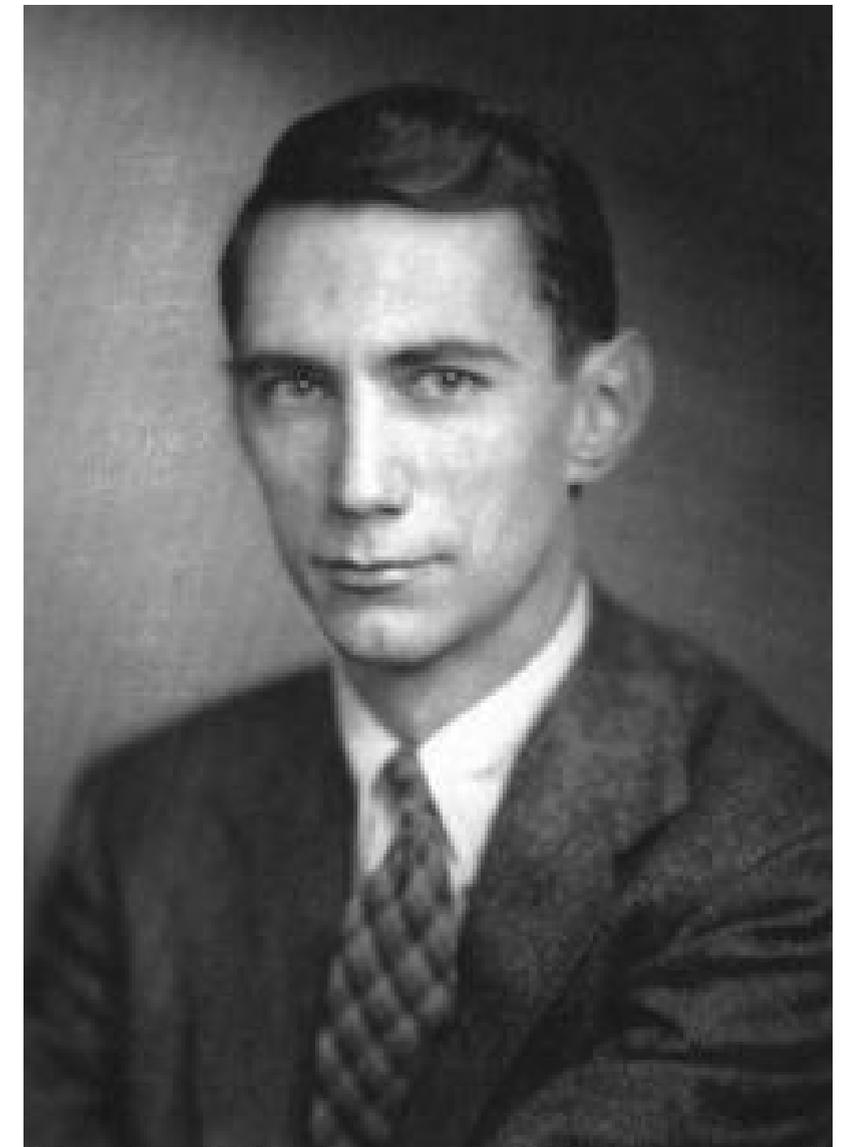
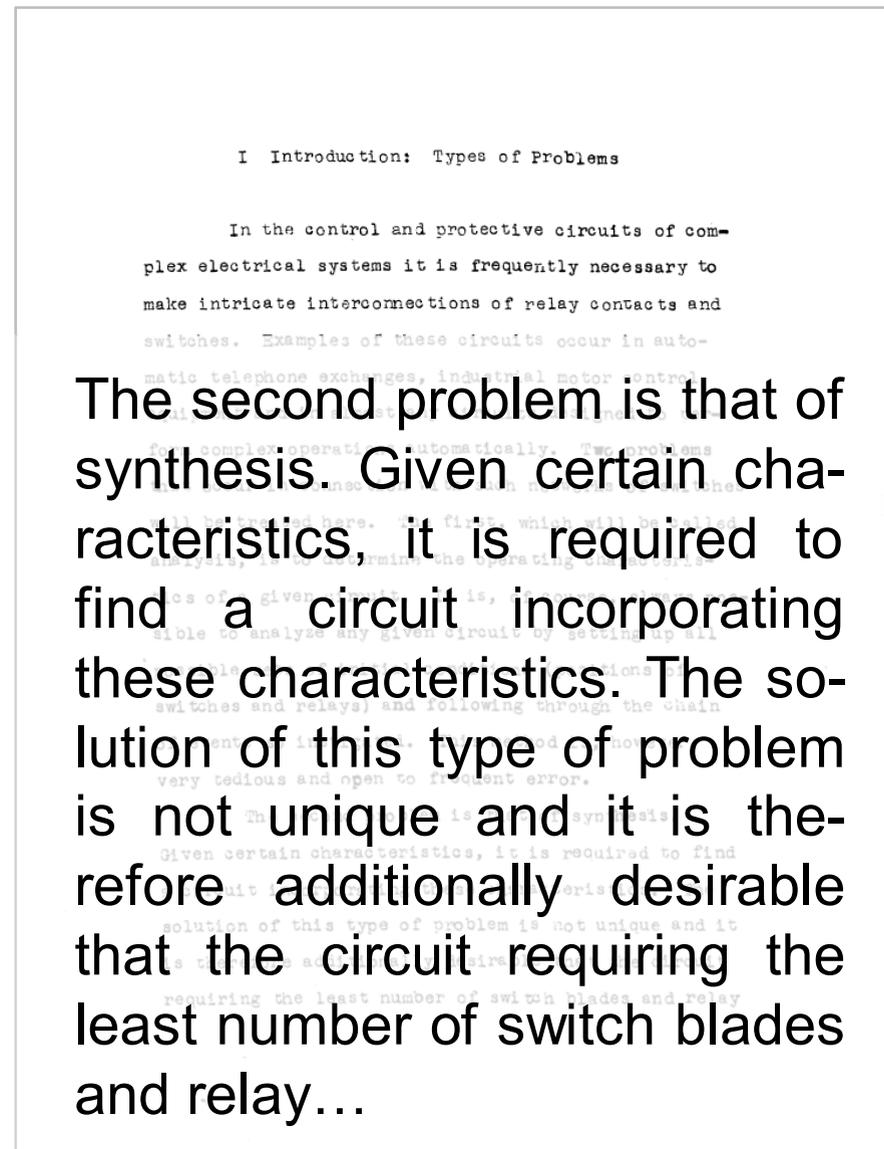
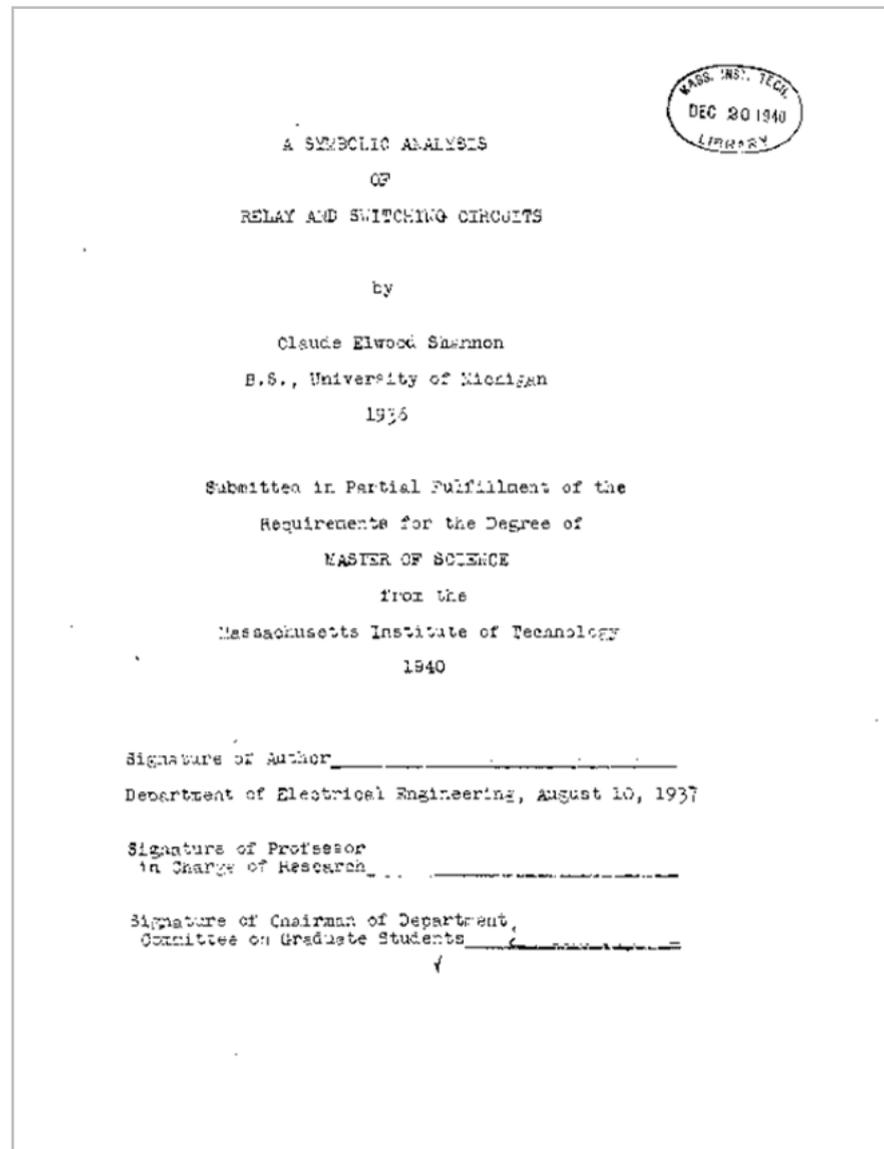
Once Upon A Time...

*Almon B. Strowger 's Rotary Dial Telephone, and...
Electromechanical Telephone Exchange !*



Claude E. Shannon

*“A Symbolic Analysis of Relay and Switching Circuits”
Master’s Thesis, MIT, 1937*



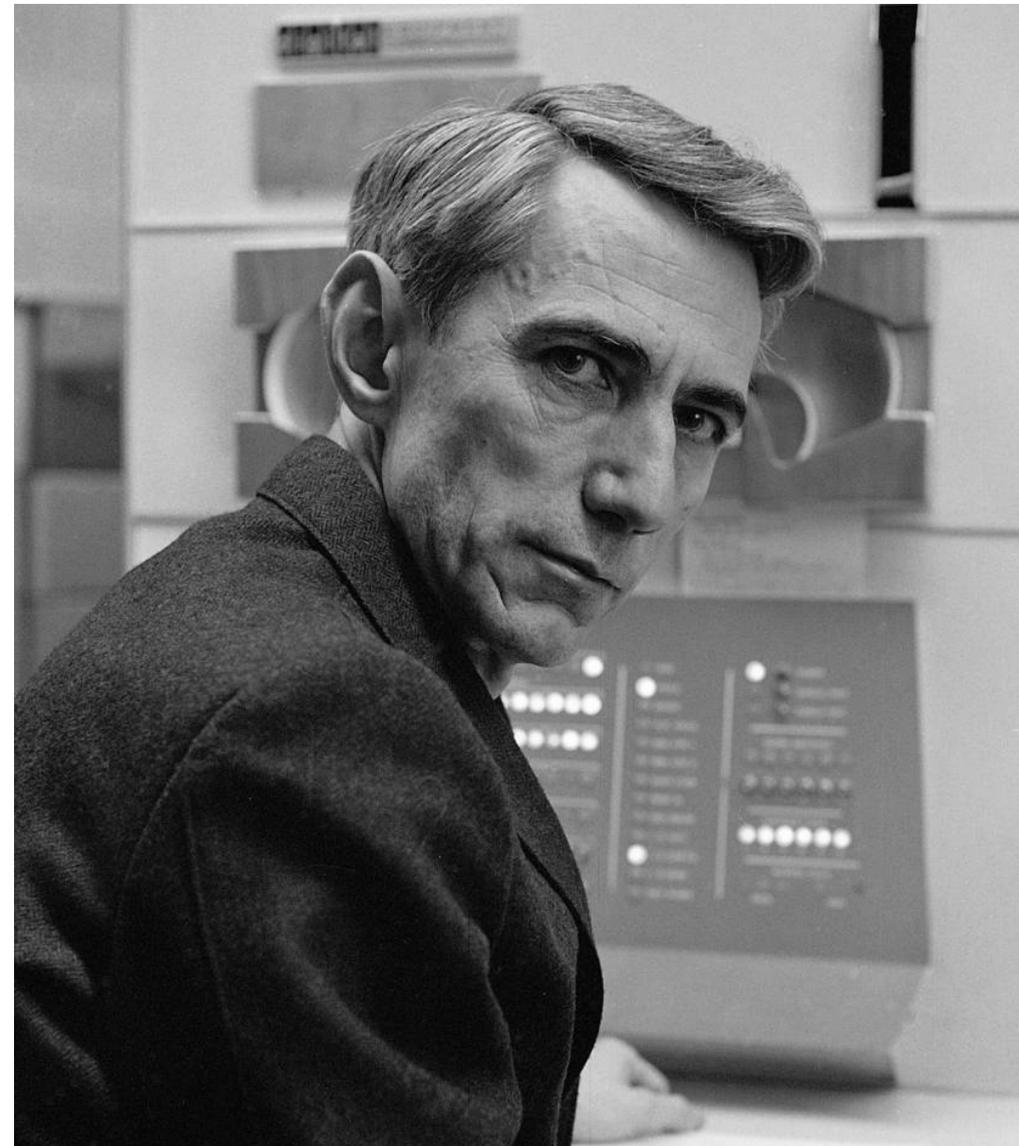
The Sampling Theorem

a.k.a. Nyquist-Shannon-Kotelnikov Theorem

The Foundation of Digital Communications

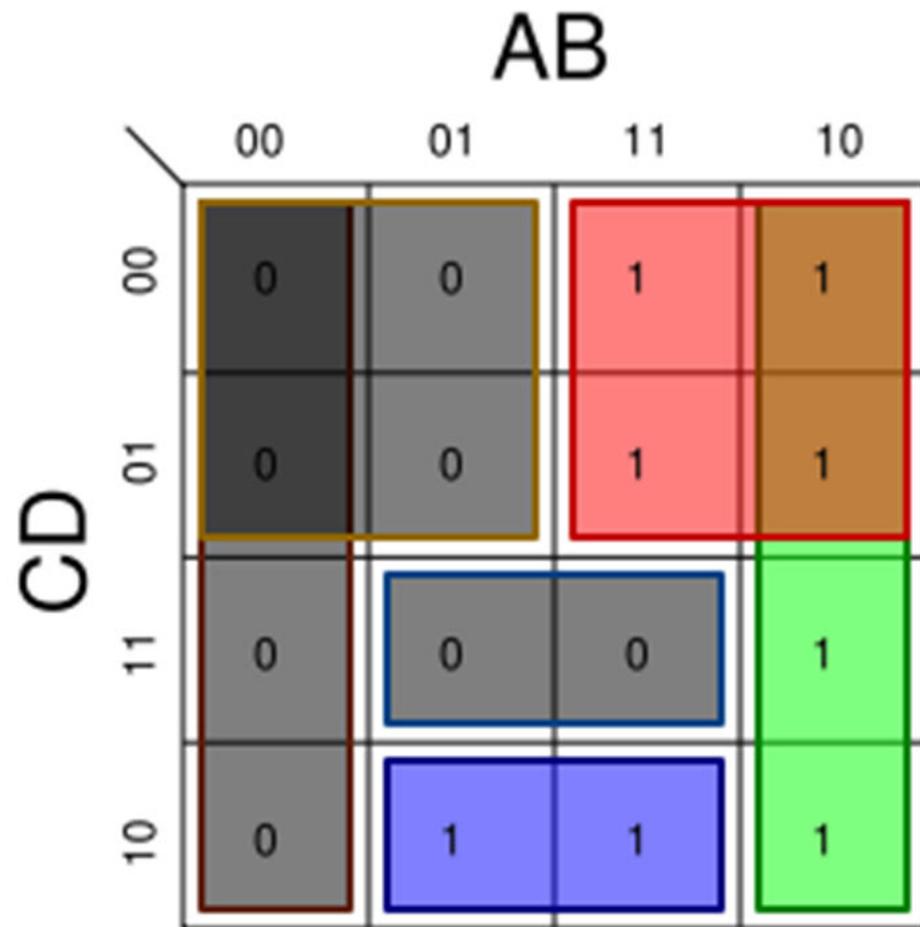
If a function $x(t)$ contains no frequencies higher than B Hertz, it is completely determined by giving its ordinates at a series of points spaced $1/(2B)$ seconds apart.

$$f\left(\frac{n}{2W}\right) = \frac{1}{2\pi} \int_{-2\pi W}^{2\pi W} F(\omega) e^{i\omega \frac{n}{2W}} d\omega.$$



Maurice Karnaugh

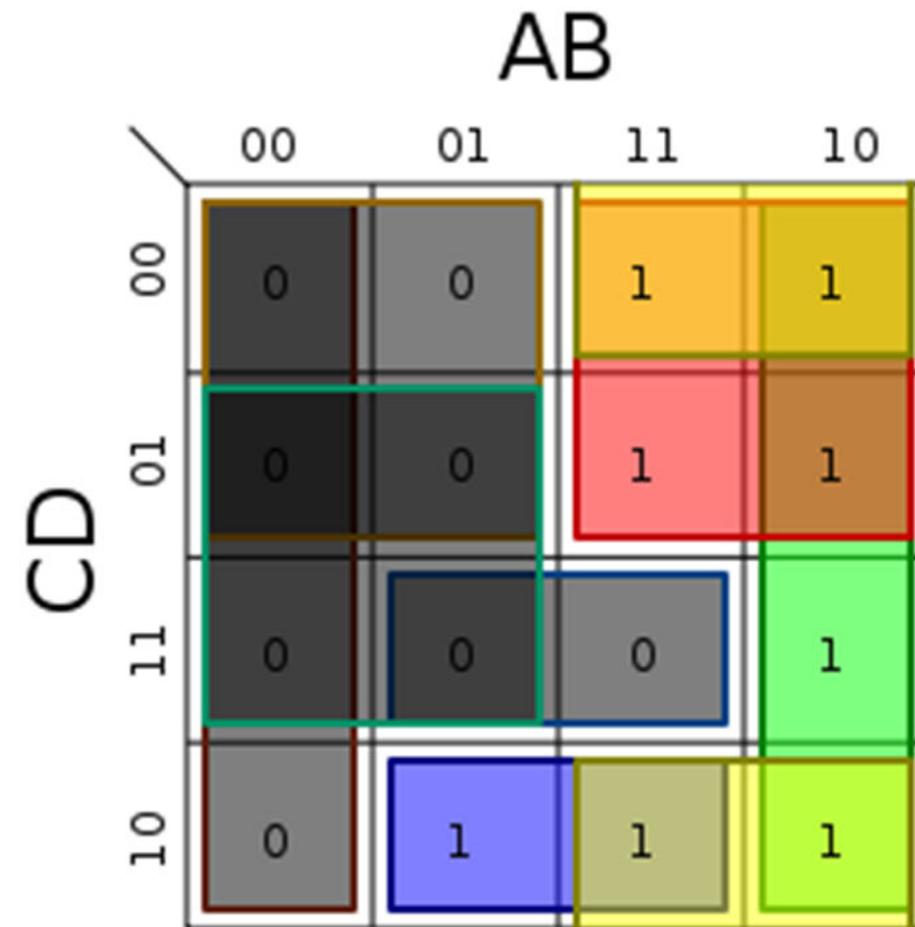
“The Map Method for Synthesis of Combinational Logic Circuits”, Bell Labs, 1953



$$f(A,B,C,D) = \sum m(6,8,9,10,11,12,13,14)$$

$$F = AC' + AB' + BCD'$$

$$F = (A+B)(A+C)(B'+C'+D')$$



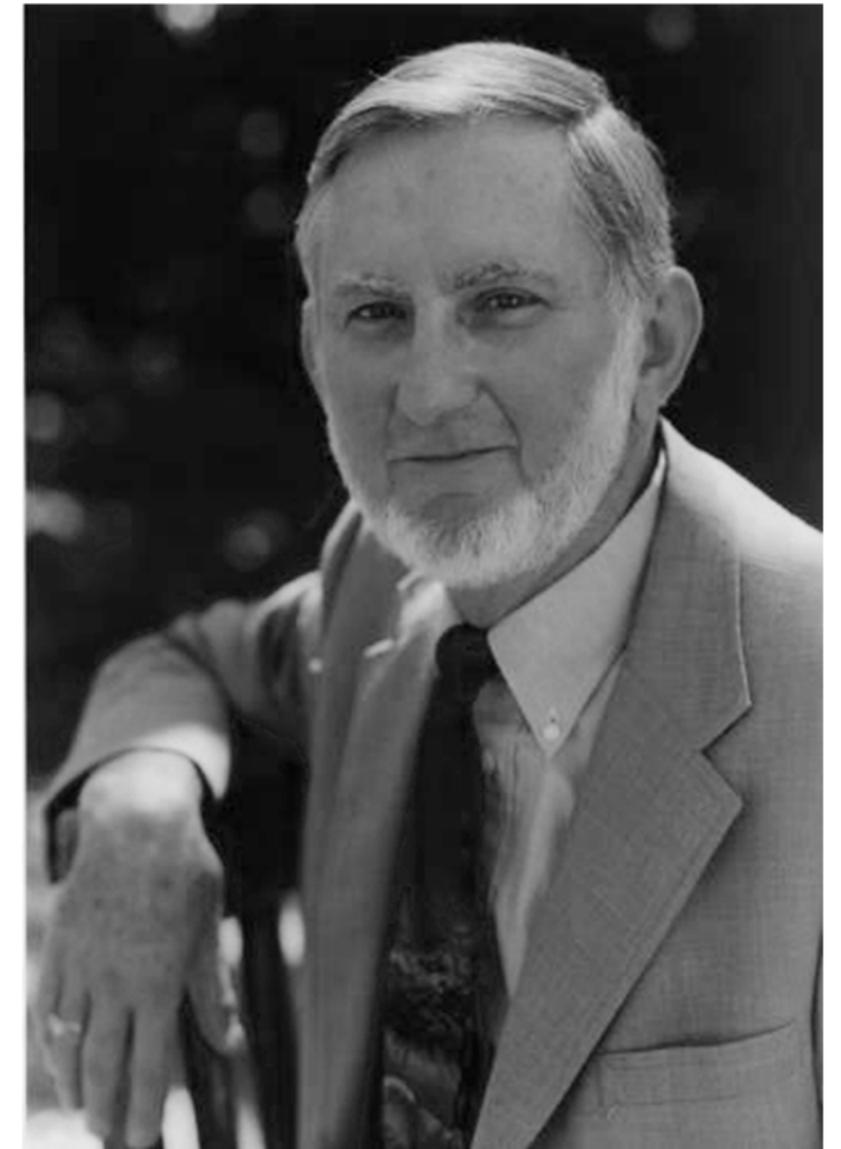
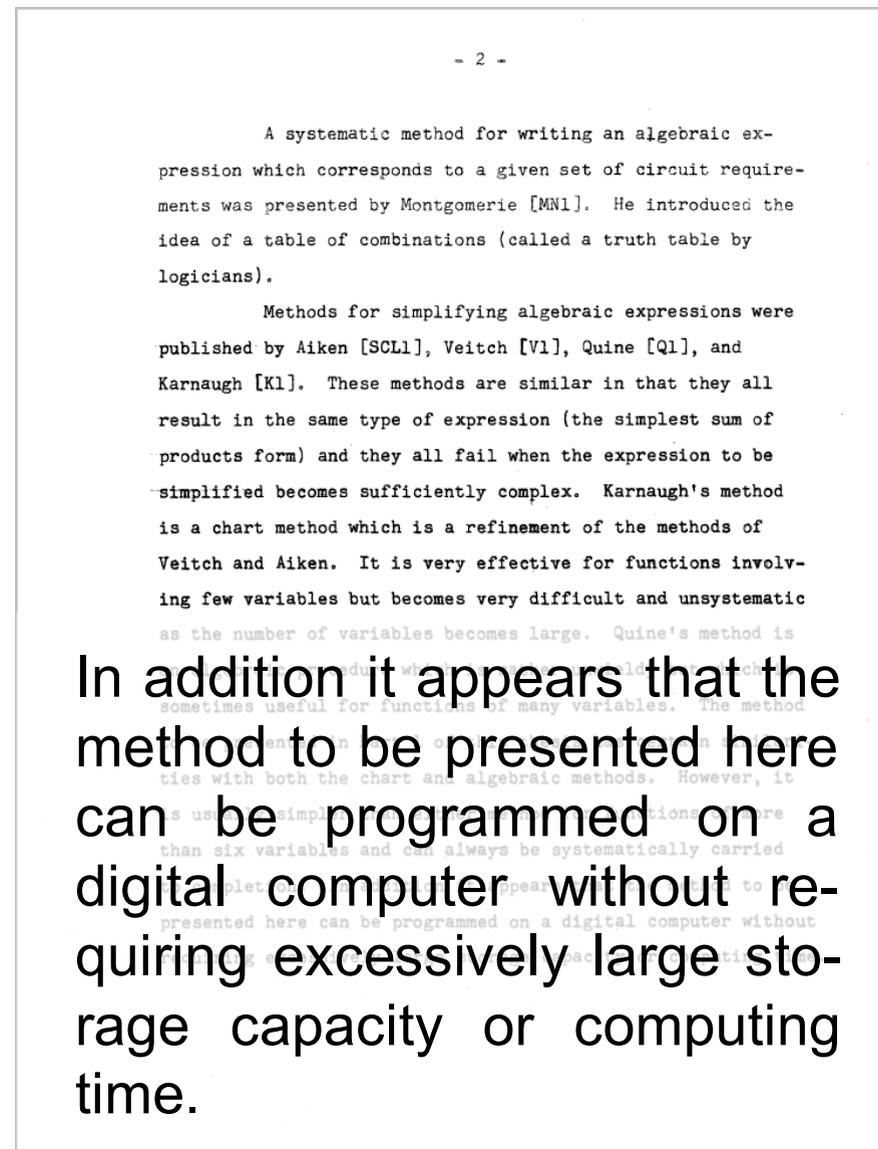
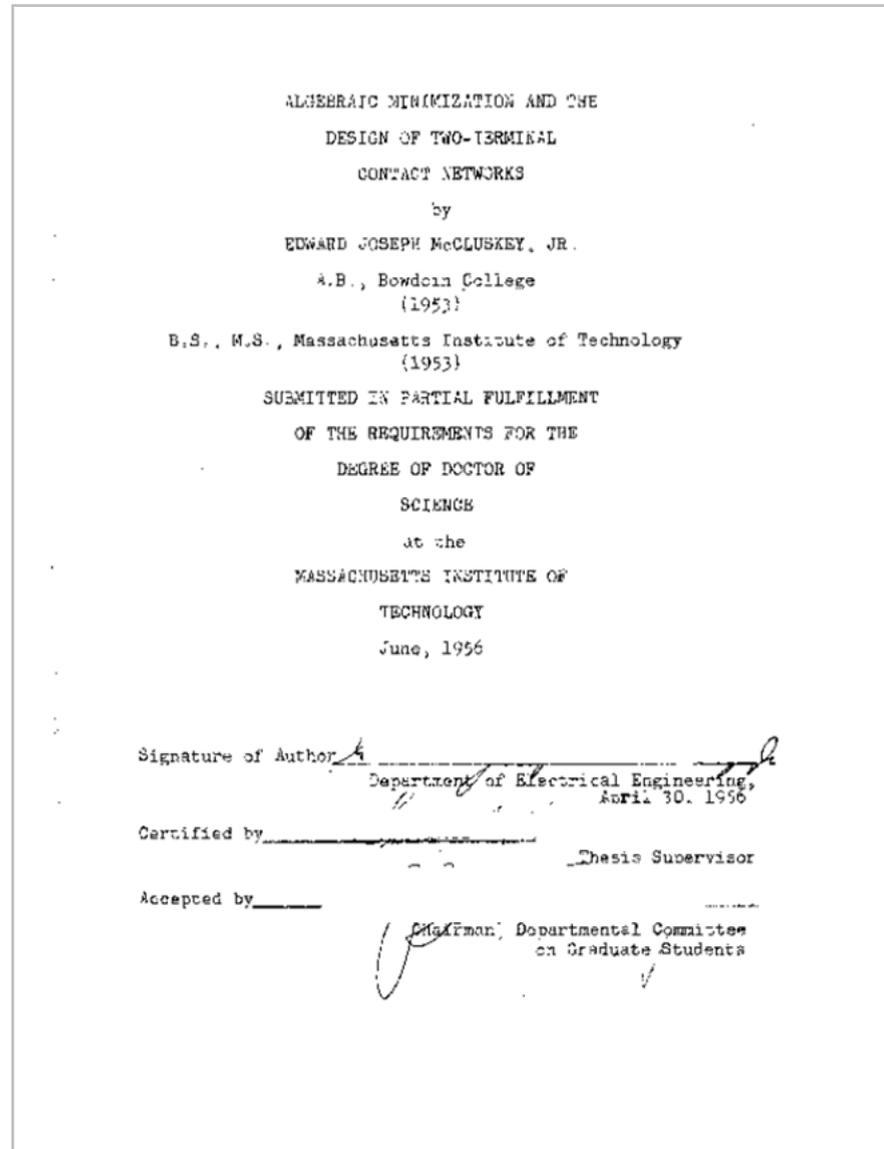
$$f(A,B,C,D) = \sum m(6,8,9,10,11,12,13,14)$$

$$F = +++ + AC'AB'BCD'AD'$$

$$F = (A+B)(A+C)(B'+C'+D')(A+D')$$

Edward J. McCluskey

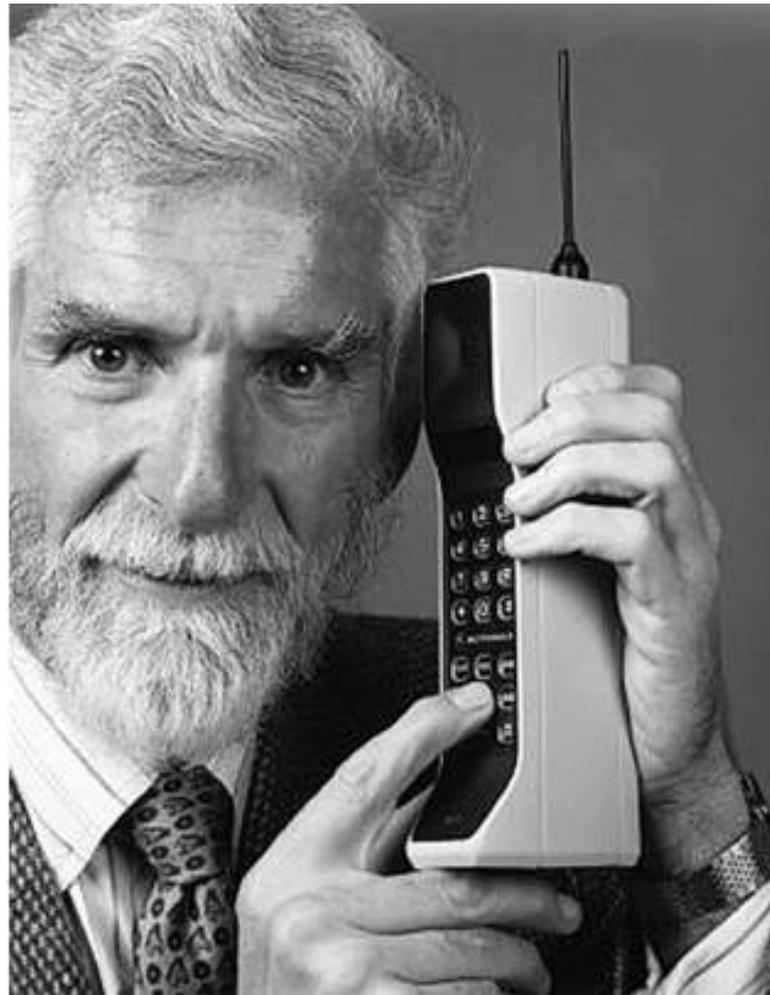
*“Algebraic Minimization and the Design of Two-Terminal Contact Networks”
Ph.D. Thesis, MIT, 1956*



...Multi-Level Minimization,...

IBM, UCB, and University of Colorado at Boulder

Motorola DynaTAC



John Darringer, William H. Joyner, and Louise H. Trevilyan, e.g. “Logic Synthesis Through Local Transformations”, IBM, 1981

Robert K. Brayton, Gary D. Hachtel, A. Richard Newton, and Alberto L. Sangiovanni-Vincentelli, e.g. “Logic Minimization Algorithms for VLSI Synthesis”, UCB, 1984

...And, Eventually : Logic Synthesis !

General Electric, and AT&T

Nokia 100

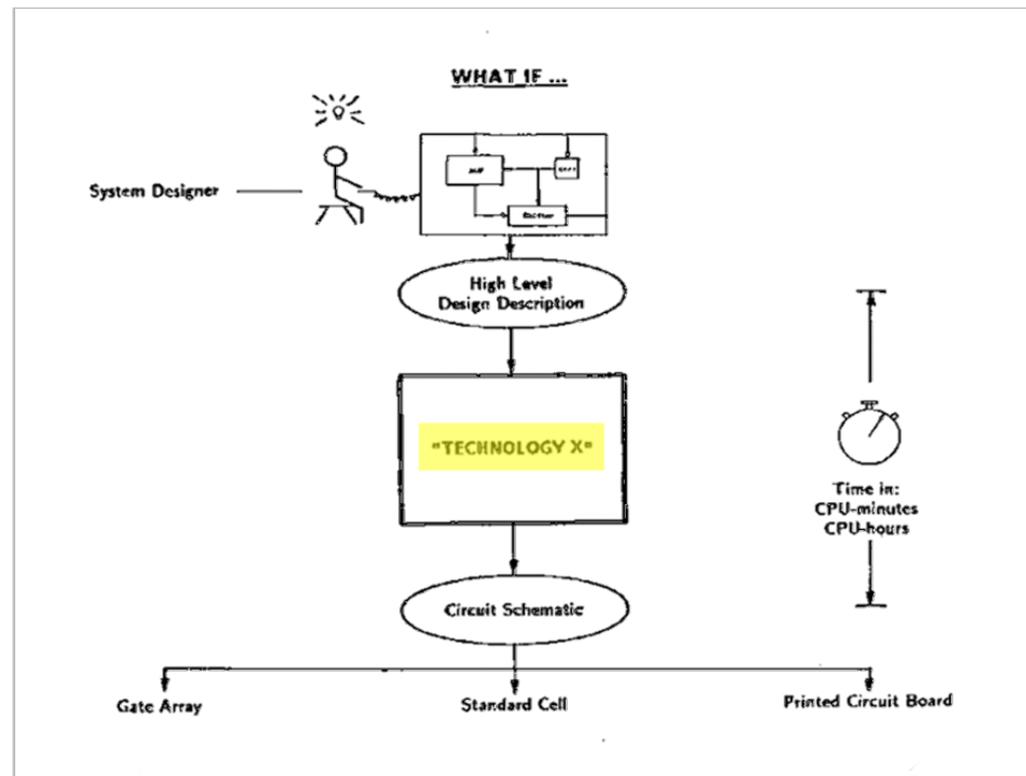


David Gregory, Karen Bartlett, Aart J. de Geus, Gary D. Hachtel, e.g. “SOCRATES: a System for Automatically Synthesizing and Optimizing Combinational Logic”, GE Calma, 1986

Kurt Keutzer, e.g. “DAGON: Technology Binding and Local Optimization by DAG Matching”, AT&T, 1988

Logic Compiler, ca. 1986

The Beginning of a New Era ! Just the Beginning...

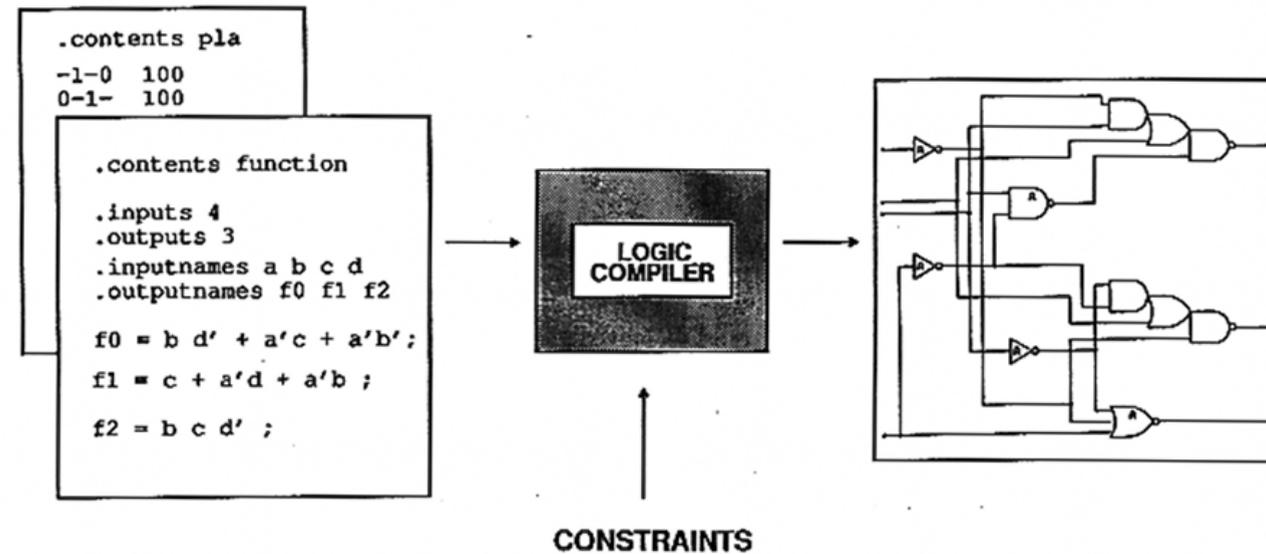


Logic Compiler, ca. 1986

Optimal Solutions, Inc. a.k.a. Synopsys, Inc.

What can LOGIC COMPILER do?

Convert equations into working circuits:



Optimal Solutions, Inc

“MIS: A Multiple-Level Logic Optimization System”, 1987

Robert K. Brayton, Richard L. Rudell, Alberto L. Sangiovanni Vincentelli, Albert R. Wang



...Binary Decision Diagrams (BDD)...

Randal E. Bryant \Rightarrow Olivier Coudert, and Jean-Christophe Madre
Logic Synthesis (and Formal Verification)

IEEE TRANSACTIONS ON COMPUTERS, VOL. C-39, NO. 2, AUGUST 1990 677

Graph-Based Algorithms for Boolean Function Manipulation

RANDAL E. BRYANT, MEMBER, IEEE

Abstract—In this paper we present a new data structure for representing Boolean functions and an associated set of manipulation algorithms. Functions are represented by directed, acyclic graphs in a manner similar to the representations introduced by Lee [1] and Akers [2], but with further restrictions on the ordering of decision variables in the graph. Although a function requires, in the worst case, a graph of size exponential in the number of arguments, many of the functions encountered in typical applications have a more reasonable representation. Our algorithms have time complexity proportional to the sizes of the graphs being operated on, and hence are quite efficient as long as the graphs do not grow too large. We present experimental results found applying these algorithms to problems in logic design verification that demonstrate the practicality of our approach.

Index Terms—Boolean functions, binary decision diagrams, logic design verification, symbolic manipulation.

1. INTRODUCTION

BOOLEAN Algebra forms a cornerstone of computer science and digital system design. Many problems in digital logic design and testing, artificial intelligence, and combinatorics can be expressed as a sequence of operations on Boolean functions. Such applications would benefit from efficient algorithms for representing and manipulating Boolean functions symbolically. Unfortunately, many of the tasks one would like to perform with Boolean functions, such as testing whether there exists any assignment of input variables such that a given Boolean expression evaluates to 1 (satisfiability), or two Boolean expressions denote the same function (equivalence) require solutions to NP-complete or co-NP-complete problems [3]. Consequently, all known approaches to performing these operations require, in the worst case, an amount of computer time that grows exponentially with the size of the problem. This makes it difficult to compare the relative efficiencies of different approaches to representing and manipulating Boolean functions. In the worst case, all known approaches perform as poorly as the naive approach of representing functions by their truth tables and defining all of the desired operations in terms of their effect on truth table entries. In practice, by utilizing more clever representations and manipulation algorithms, we can often avoid these exponential computations.

Our representation has several advantages over previous approaches to Boolean function manipulation. First, most commonly encountered functions have a reasonable representation. For example, all symmetric functions (including even and odd parity) are represented by graphs where the number of vertices grows at most as the square of the number of arguments. Second, the performance of a program based on our algorithms when processing a sequence of operations degrades slowly, if at all. That is, the time complexity of any single operation is bounded by the product of the graph sizes for the functions being operated on. For example, complementing a function requires time proportional to the size of the function graph, while combining two functions with a binary operation (of which intersection, subtraction, and testing for

Implicit and Incremental Computation of Primes and Essential Primes of Boolean functions

O. Coudert and J. C. Madre
Bull Corporate Research Center
Rue Jean Jaurès
78340 Les Clayes-sous-bois, FRANCE

Abstract
Recently introduced implicit set manipulation techniques have made it possible to formally verify finite state machines with state graphs too large to be built. This paper shows that these techniques can also be used with success to compute and manipulate implicitly extremely large sets of prime and of essential prime implicants of incompletely specified Boolean functions. These sets are denoted by *meta-products* that are represented with binary decision diagrams. This paper describes two procedures. One is based on the standard BDD operators, and the other, more efficient, takes advantage of the structural properties of BDDs and of meta-products to handle a larger class of functions than the former one.

Introduction

We have recently introduced a technique [4] for verifying finite state machines that can deal with machines with state graphs too large to be built. The key concepts that make this verification possible are to denote subsets of $\{0,1\}^n$ with their characteristic functions, and represent these Boolean functions with binary decision diagrams (BDDs) that are a very compact graph representation of such functions [3]. Since there is no relation between the size of a set and the size of the BDD that denotes it, the computation cost of this technique is completely independent from the number of states of the machines [4].

In this paper we show that these concepts can be used with success to implement implicit computation procedures of the sets of prime and of essential prime implicants of incompletely specified Boolean functions. These procedures have costs that are independent of the sizes of these sets, and thus they overcome the limitations of methods based on explicit prime manipulations [2, 8, 11, 12].

An *incompletely specified* Boolean function f_c , also called a function with a *core set*, is defined by a couple (C, f) where f is a function from $\{0,1\}^n$ into $\{0,1\}$, and the core set C is a subset of $\{0,1\}^n$. This function is defined by $f_c(x) = f(x)$ if $x \in C$, and $f_c(x) = *$ if $x \notin C$, where the symbol $*$ can be either 0 or 1 [2].

The product p of P_c is a cube of the function $f_c = (C, f)$ if and only if the set $S_p \cap C$ is not empty, and for any element of this set, the function f evaluates to 1. The cube p of f_c is a *prime implicant* or *prime* of f_c if and only if p is a maximal element of the set of cubes of f_c with respect to the partial order " \supseteq ". Finally the

A New Viewpoint on Two-Level Logic Minimization

Olivier Coudert Jean Christophe Madre Henri Fraisse
Bull Corporate Research Center
Rue Jean Jaurès
78340 Les Clayes-sous-bois, FRANCE

Abstract
This paper presents a new implicit cyclic core computation procedure that covers the limitations of the standard Boolean manipulations [15, page 26]. This procedure is independent of the number of minterms of the function f and of its number of prime implicants, which allows us to treat functions for which these numbers are so large that it has never been possible to perform a 2-level minimization.

2. Definitions and Notations

An implicit cyclic core computation procedure has been developed by Bryant, Coudert, and Madre [15]. The idea here is to build, at each step of the cyclic core computation, the BDDs of the dominance relations, and then to extract from these BDDs a reduced covering matrix. Though this procedure is able to treat 2 of the 20

This paper presents a new implicit cyclic core computation procedure that covers the limitations of the standard Boolean manipulations [15, page 26]. This procedure is independent of the number of minterms of the function f and of its number of prime implicants, which allows us to treat functions for which these numbers are so large that it has never been possible to perform a 2-level minimization.

This paper presents a new implicit cyclic core computation procedure that covers the limitations of the standard Boolean manipulations [15, page 26]. This procedure is independent of the number of minterms of the function f and of its number of prime implicants, which allows us to treat functions for which these numbers are so large that it has never been possible to perform a 2-level minimization.



A Revolutionary... Evolution

Convergence !

Jackson Pollock, Convergence, 1952, Albright-Knox Art Gallery, Buffalo, NY

Convergence Of Logic And Test Synthesis

Srinivas Devadas, Hi-Keung Tony Ma,

A. Richard Newton, Alberto L. Sangiovanni-Vincentelli



FOR IMMEDIATE RELEASE

MOMENTUM BUILDS FOR SYNOPSIS TEST SYNTHESIS

— Test Compiler orders surpass \$2 million since announcing the product —

“Our goal was to solve the problem of developing highly testable designs with minimum penalties in chip speed and size,” said Aart de Geus, co-founder and senior vice-president of marketing at Synopsys.

Although design for test (DFT) techniques have been available in the last decade, several obstacles have slowed down the adoption of DFT by engineers. Previous techniques made designs larger by learning complex design rules. Test Compiler, however, relaxes the design rules, automatically creates the test structures, and enables engineers to optimize the design for speed and size with the test structures in place. As a result, designers are able to hand over a highly testable design to the manufacturing test engineer and help reduce the risk of undetected manufacturing defects.

Convergence Of Logic And Power Synthesis

L. Benini, P. Siegel, G. De Micheli, "Automatic Synthesis of Gated Clocks for Power Reduction in Sequential Circuits", Stanford University, 1994

Automatic Synthesis of Gated Clocks for Power Reduction in Sequential Circuits

L. Benini P. Siegel G. De Micheli

Center for Integrated Systems
Stanford University, Stanford CA 94305

Abstract

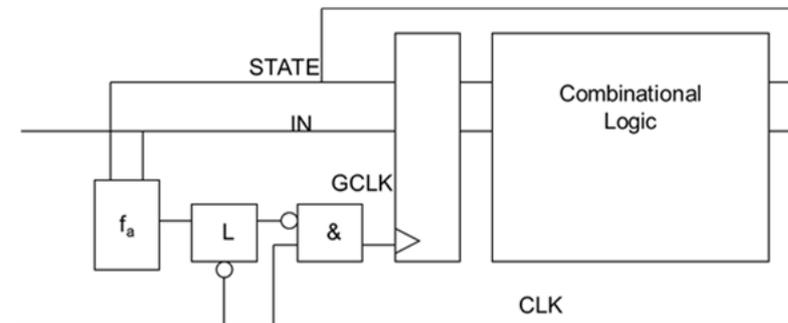
With the proliferation of portable devices and increasing levels of chip integration, reducing power consumption is becoming of paramount importance. We describe a technique to automatically synthesize gated clocks for finite-state machines (FSMs) to reduce power in the final implementation. This technique recognizes self-loops in the FSM (either from the state diagram or from a synchronous network) and uses the function described by the self-loops to gate the clock. The clock activation function is then used as *don't-care* information to minimize the logic in the FSM for additional power savings. We applied these techniques to standard MCNC benchmarks and found an average reduction in power dissipation of 25%, at the cost of a 5% increase in area.

1 Introduction

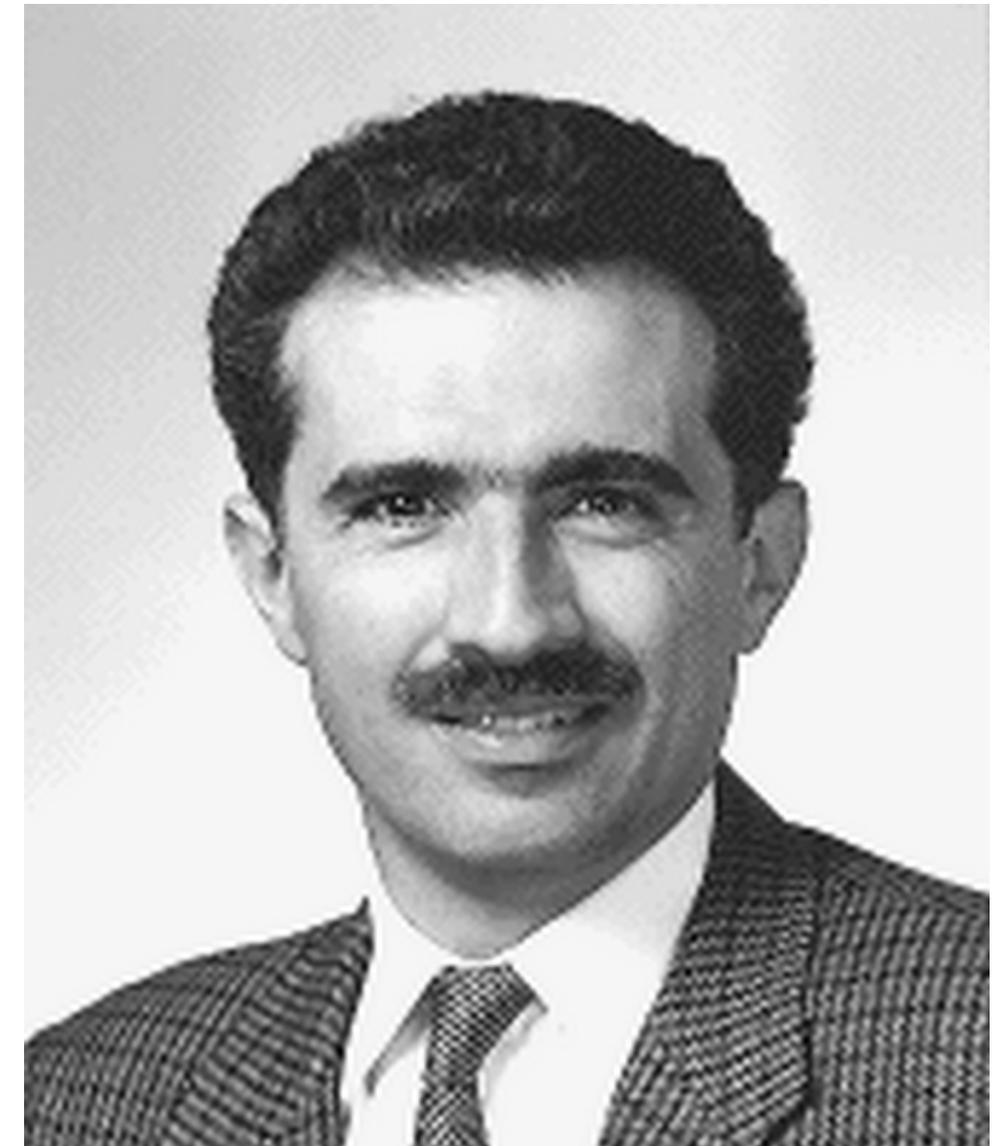
As portable devices proliferate and device sizes continue to shrink, allowing more devices to fit on a chip, power consumption has taken on increased importance. Much recent work has focused on accurate estimation of power consumption and on its concomitant reduction at all levels of abstraction, from high level synthesis down to physical layout [1, 2, 3, 4, 5, 6, 7].

Most power reduction techniques have emphasized reducing the level of activity in some portion of the circuit. We extend this research by concentrating on reducing the activity level of the clock by selectively stopping the clock. Because many sequential machines are implementations of reactive systems which wait for a certain event to occur before changing state, much power is wasted during this waiting period [8]. Latches are still clocked and the next state function is computed, consuming unnecessary power since nothing can change until the requisite event arrives. By stopping the clock during this period, we can realize substantial savings in many finite state machines (FSMs).

1

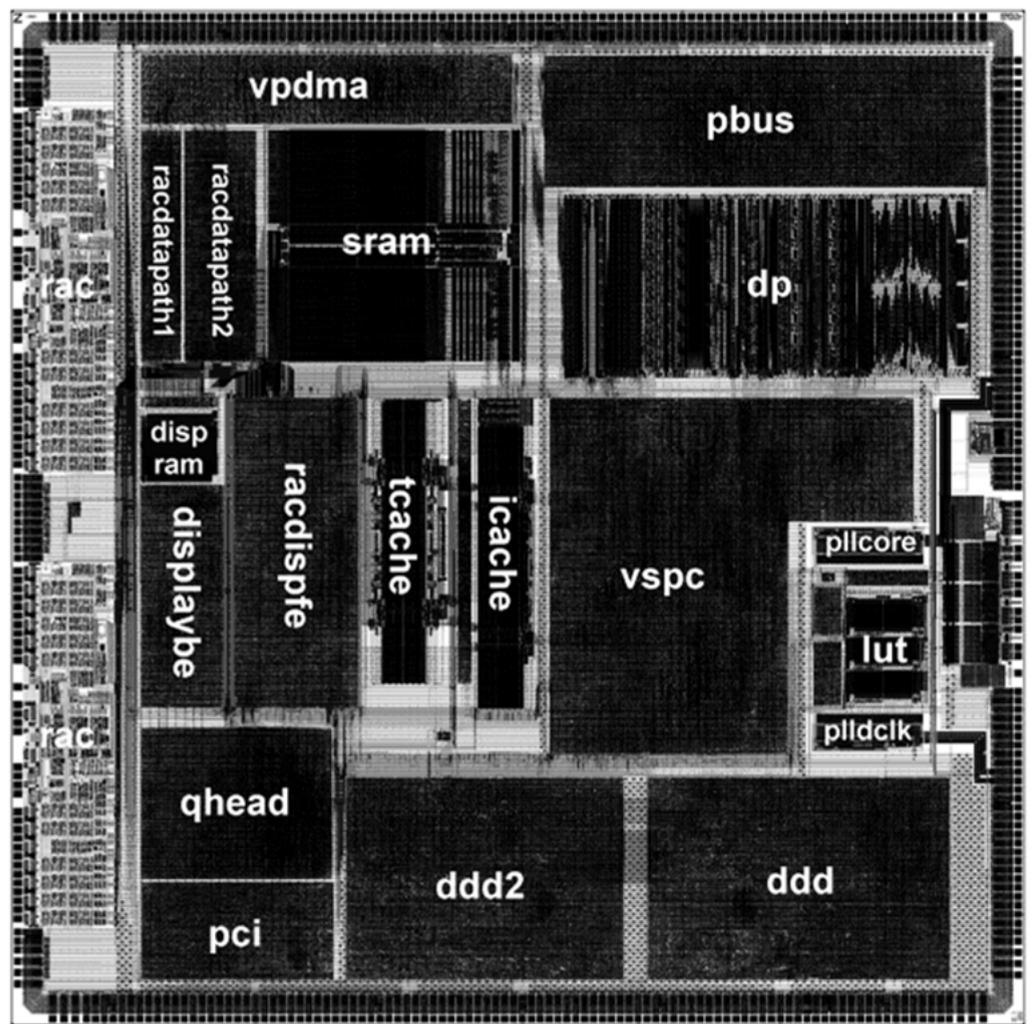


Circuit	Original		Gated		Power Reduction (%)
	Size	Power	Size	Power	
ex	128	51 μ W	118	27 μ W	48%
bbsse	966	212 μ W	1002	190 μ W	11%
bbara	348	328 μ W	390	127 μ W	61%
bbtas	178	62 μ W	188	50 μ W	19%
sse	912	175 μ W	946	150 μ W	15%
s386	856	179 μ W	882	140 μ W	21%
cse	1320	125 μ W	1406	70 μ W	44%
dk14	758	212 μ W	852	211 μ W	1%
s27	146	60 μ W	178	54 μ W	10%
mc	182	73 μ W	225	61 μ W	17%
sand1	2220	265 μ W	2108	180 μ W	32%

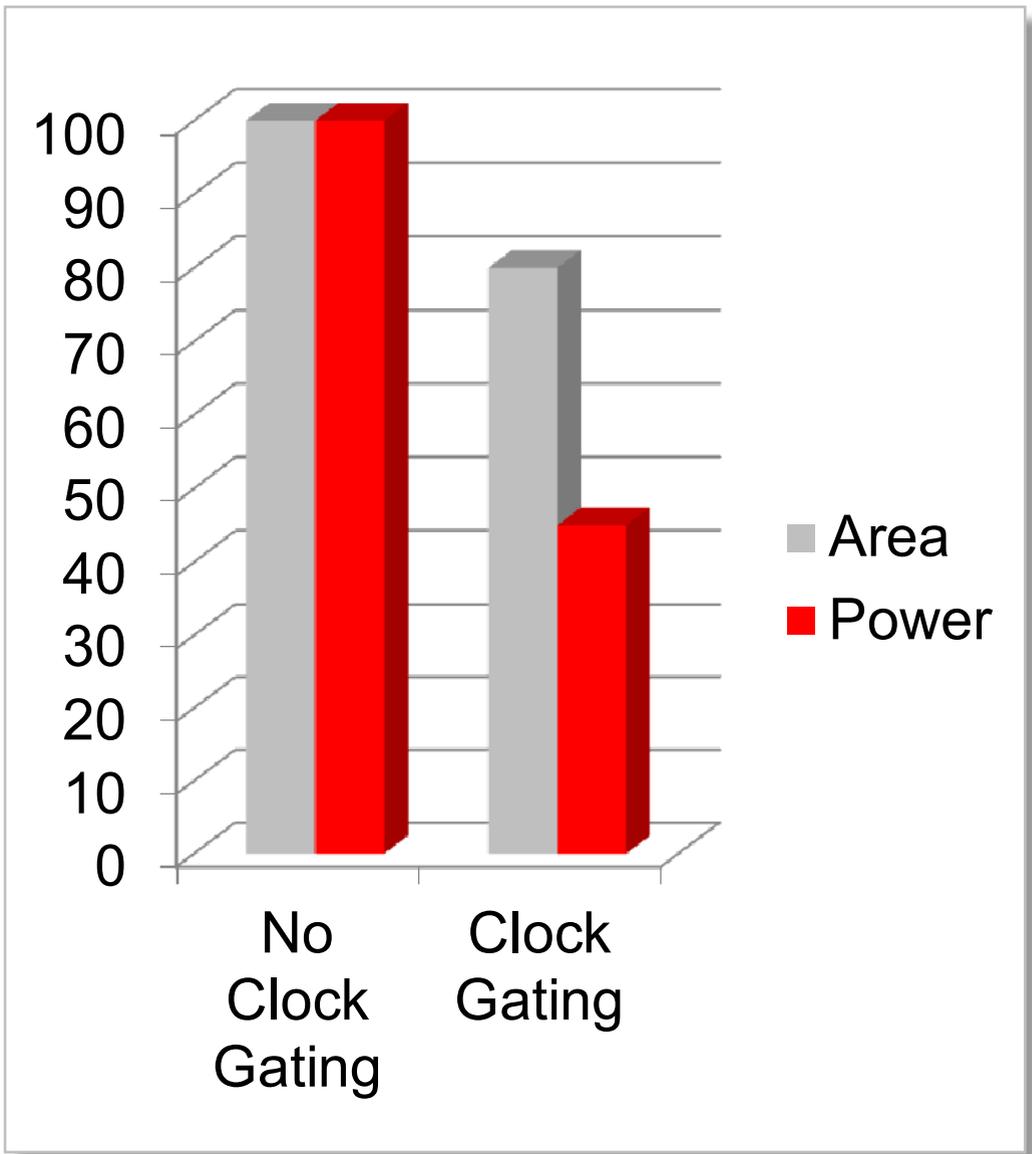


Convergence Of Logic And Power Synthesis

Power Compiler, 1997



Wireless Application, 250nm



Convergence of Synthesis And Implementation

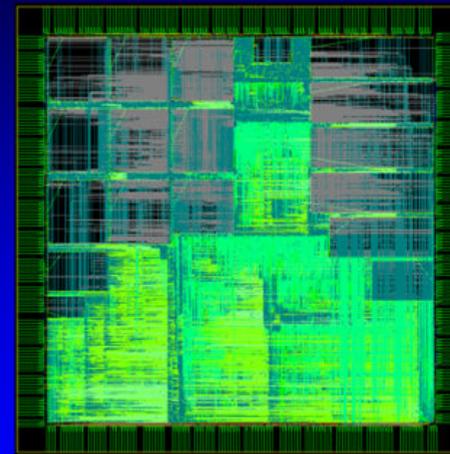
Physical Compiler, 1999



Physical Compiler

Evaluation Vehicle #1

- Evaluation Objective
 - Stress robustness, capacity and limits of the tool
- Design Characteristics
 - 250 Nanometers
 - Transmission application
 - 220K placeable instances, ~800K gates, 18 RAMs
 - 70% utilization, irregular, non-rectilinear floor-plan
 - 70+ Clocks, fastest @ 65 MHz
 - Traditional flow converged after 4 weeks
- Evaluation Results
 - Compiled in 27 hours
 - Met timing in one pass
 - Design routable, no congestion

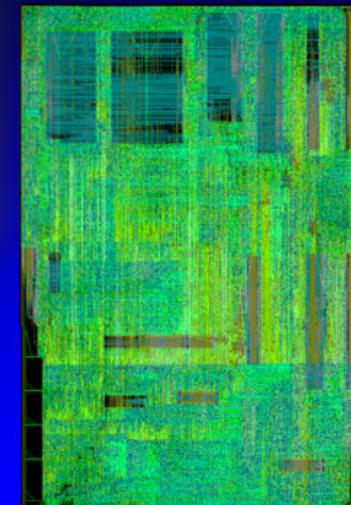


gates, 22 RAMs

- 35 combinatorial levels between 2 sequential stages on the critical path
- 2 Clocks, fastest @ 80 MHz
- Traditional flow converged after 2 weeks

□ Evaluation Results

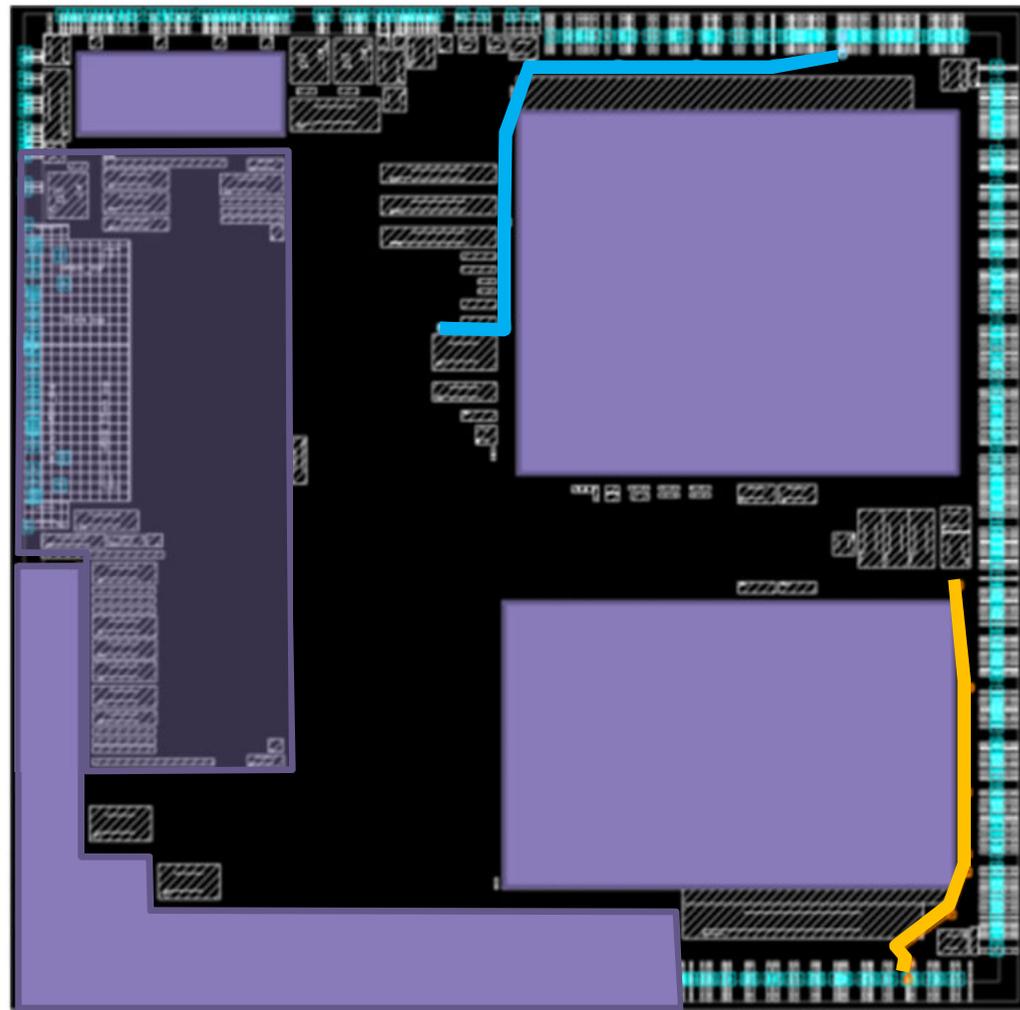
- Compiled in 10 hours
- 300ps worst negative slack post-routing after one pass



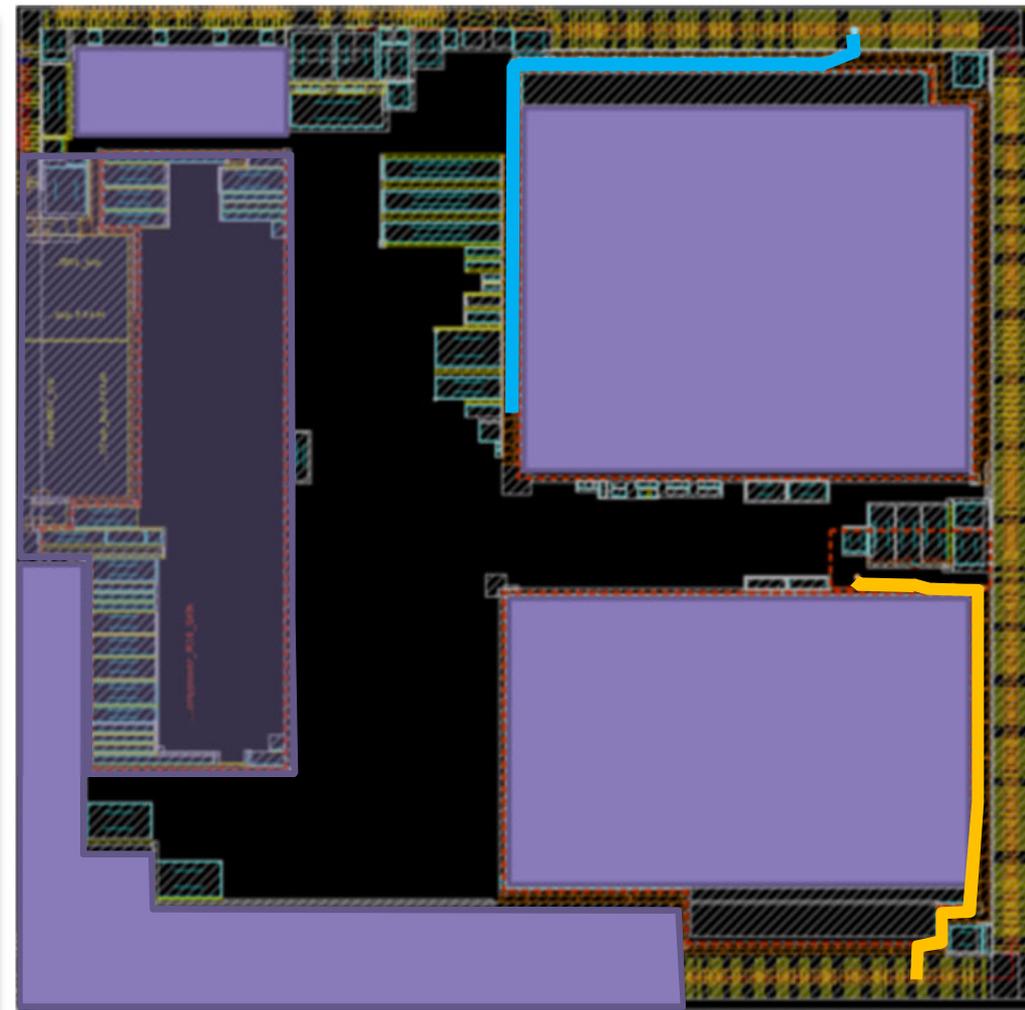
Convergence Continues, 2005

“Some” Placement in Synthesis: Same Critical Paths, Correlation $\pm 5\%$

Design Compiler

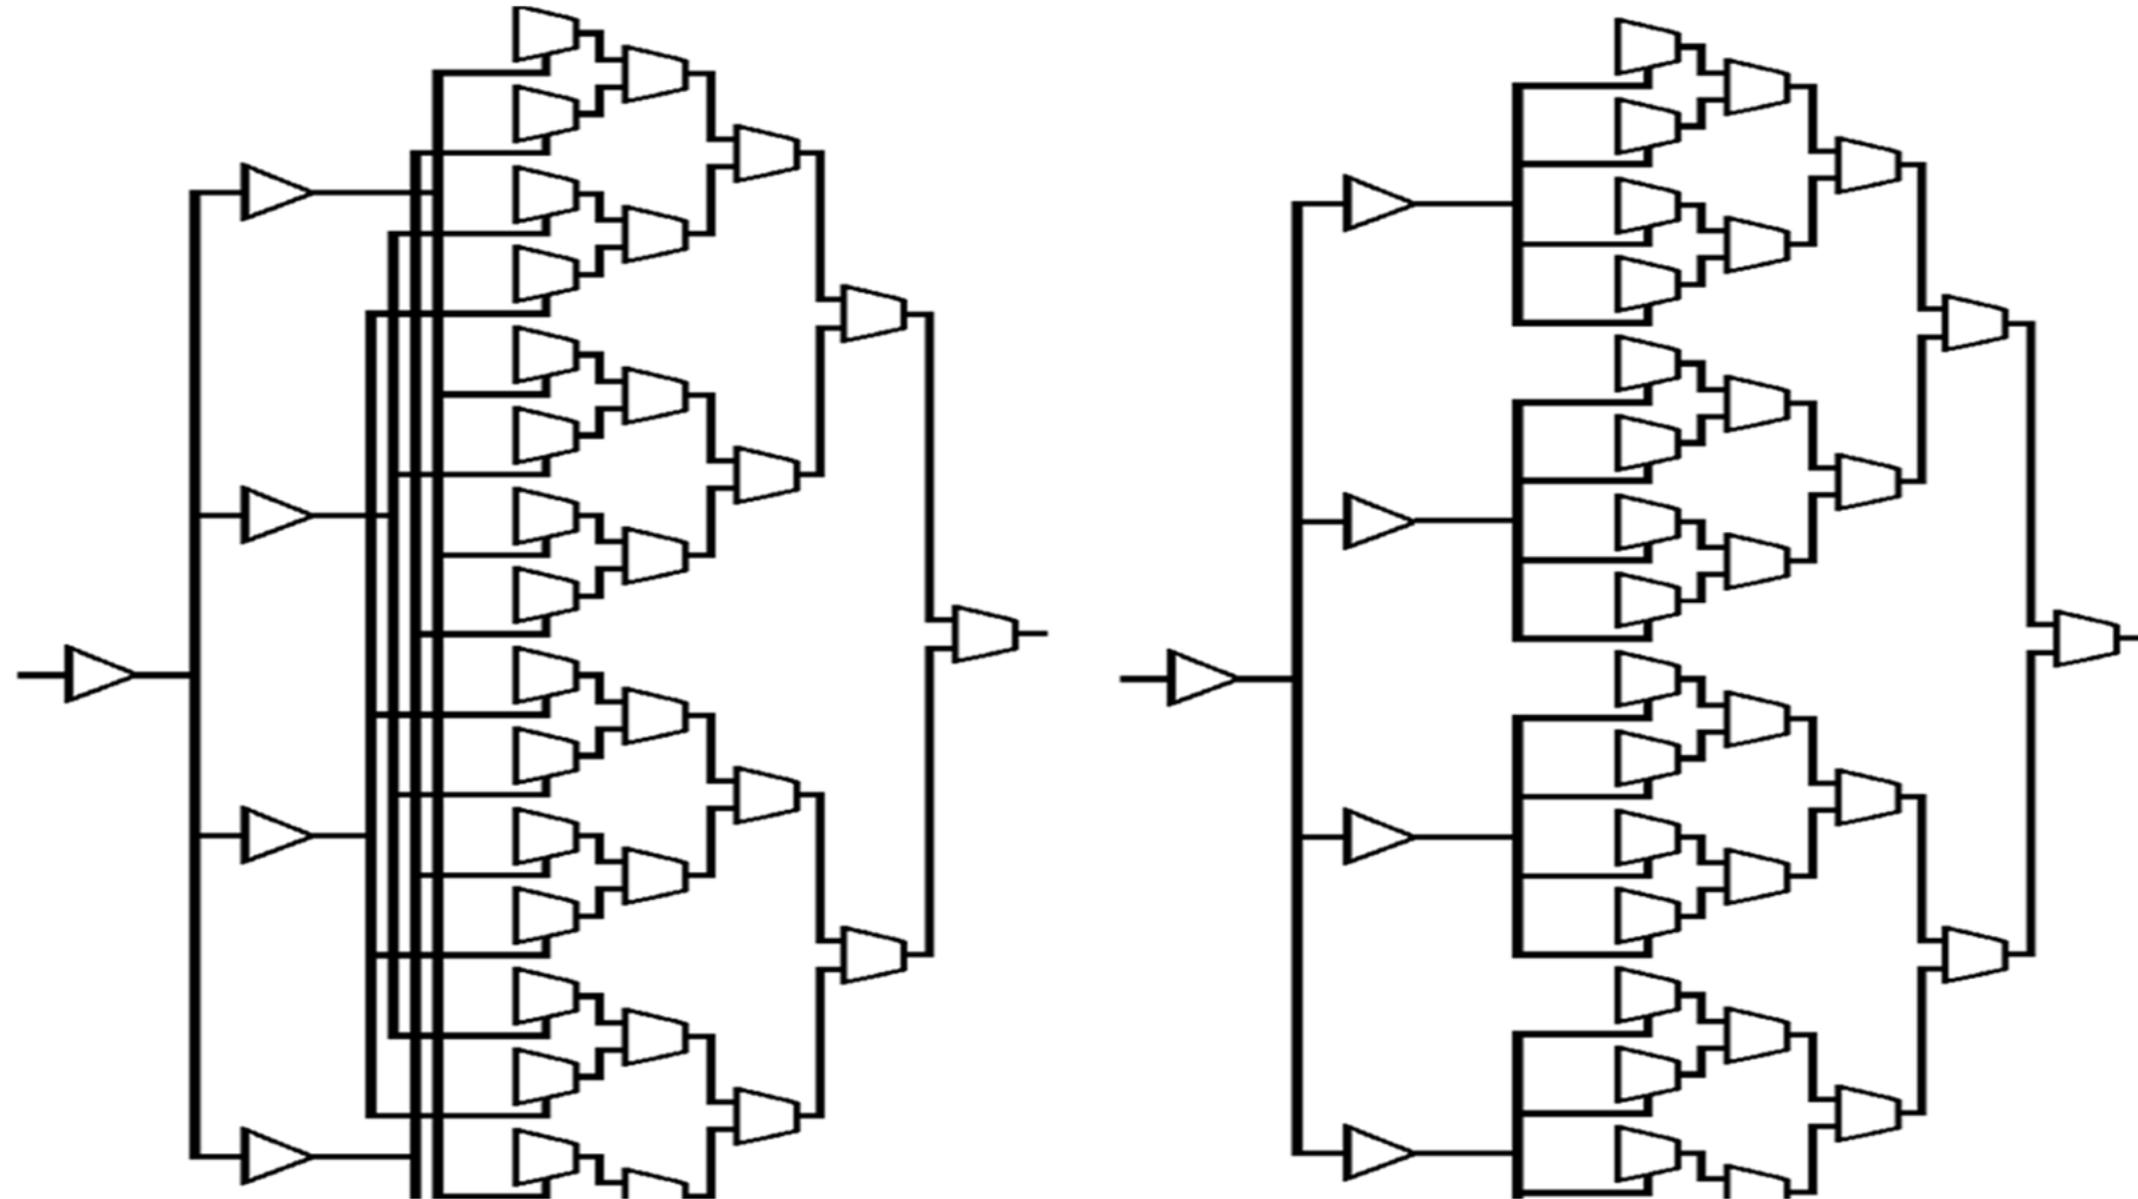


IC Compiler



Convergence Continues, 2007

“Some” Global Routing In Synthesis: Improving Planarity in a Graph

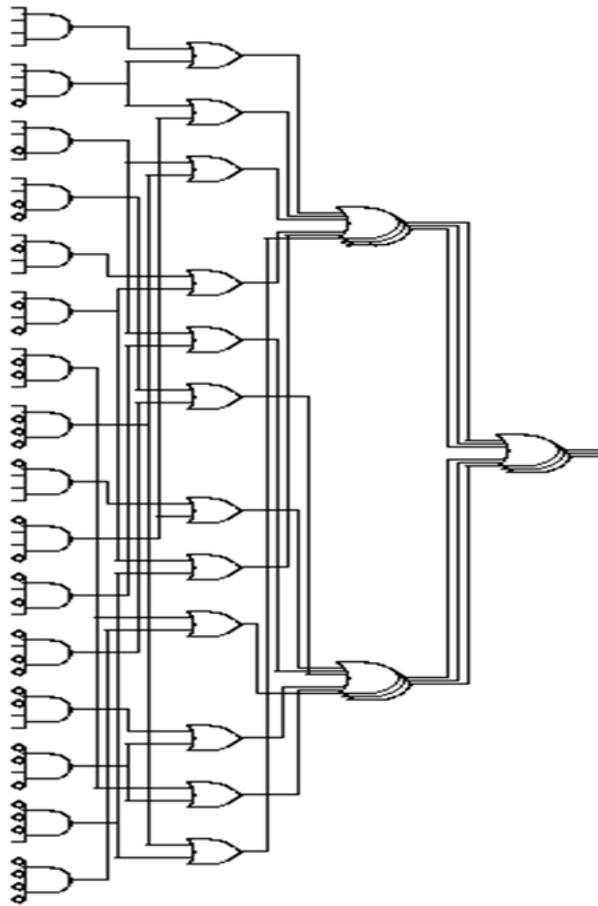


Convergence Continues, 2014

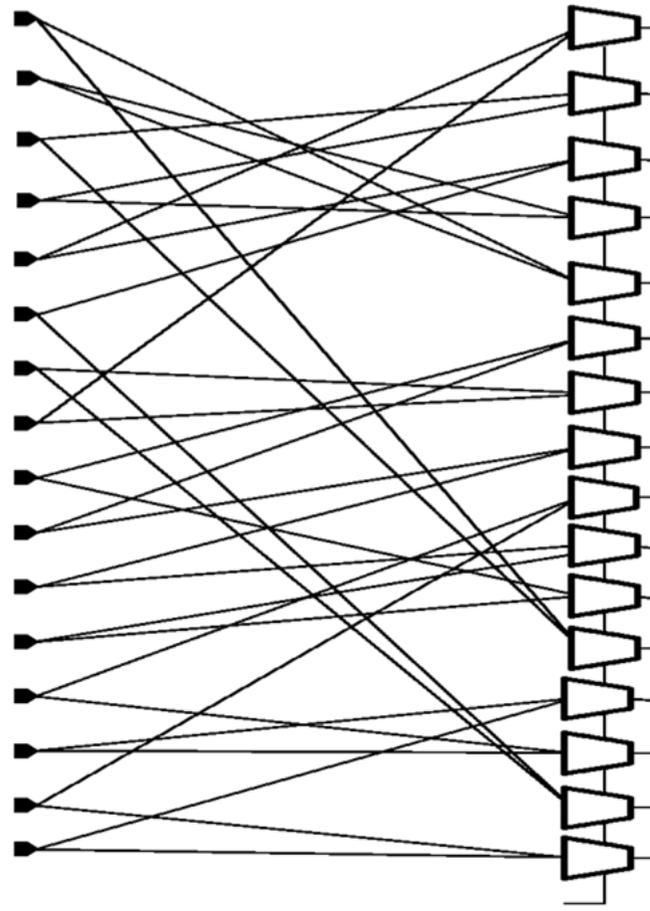
Congestion Analysis Pre-Synthesis Technology

Detects RTL Structures Causing Congestion down in the Implementation Flow

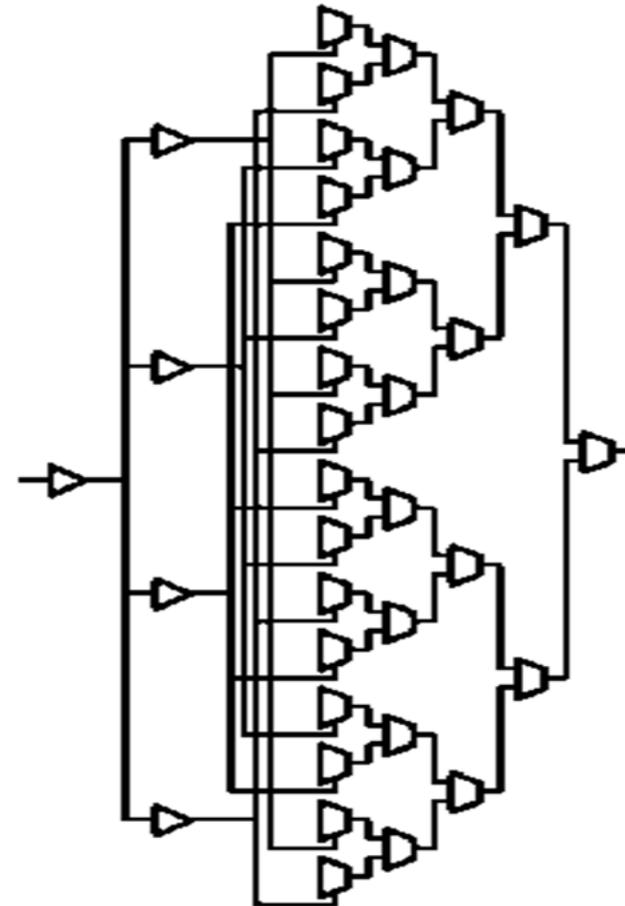
Large ROMs



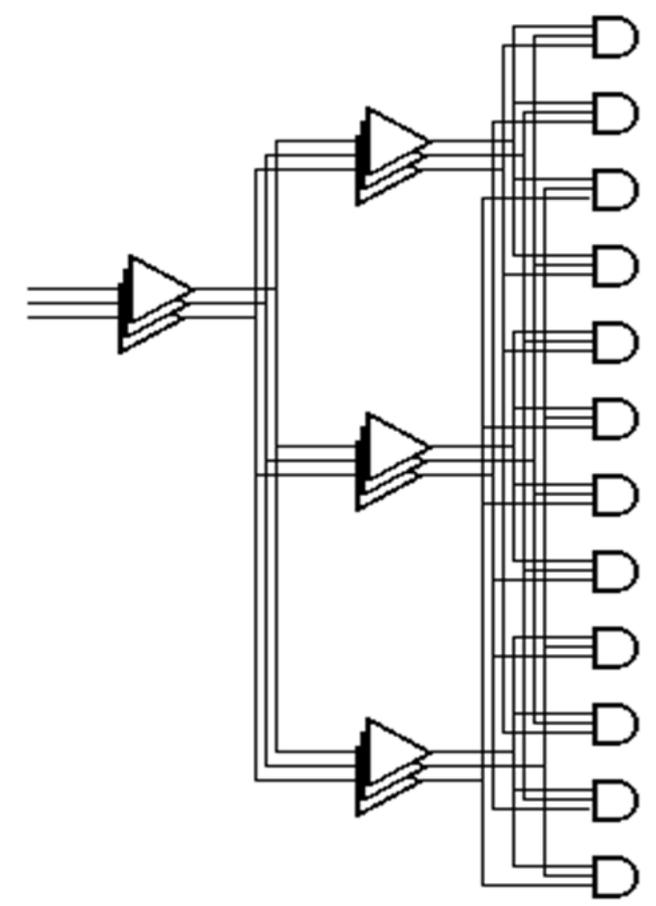
Large Data Switches



Large MUXs



Large Selectors

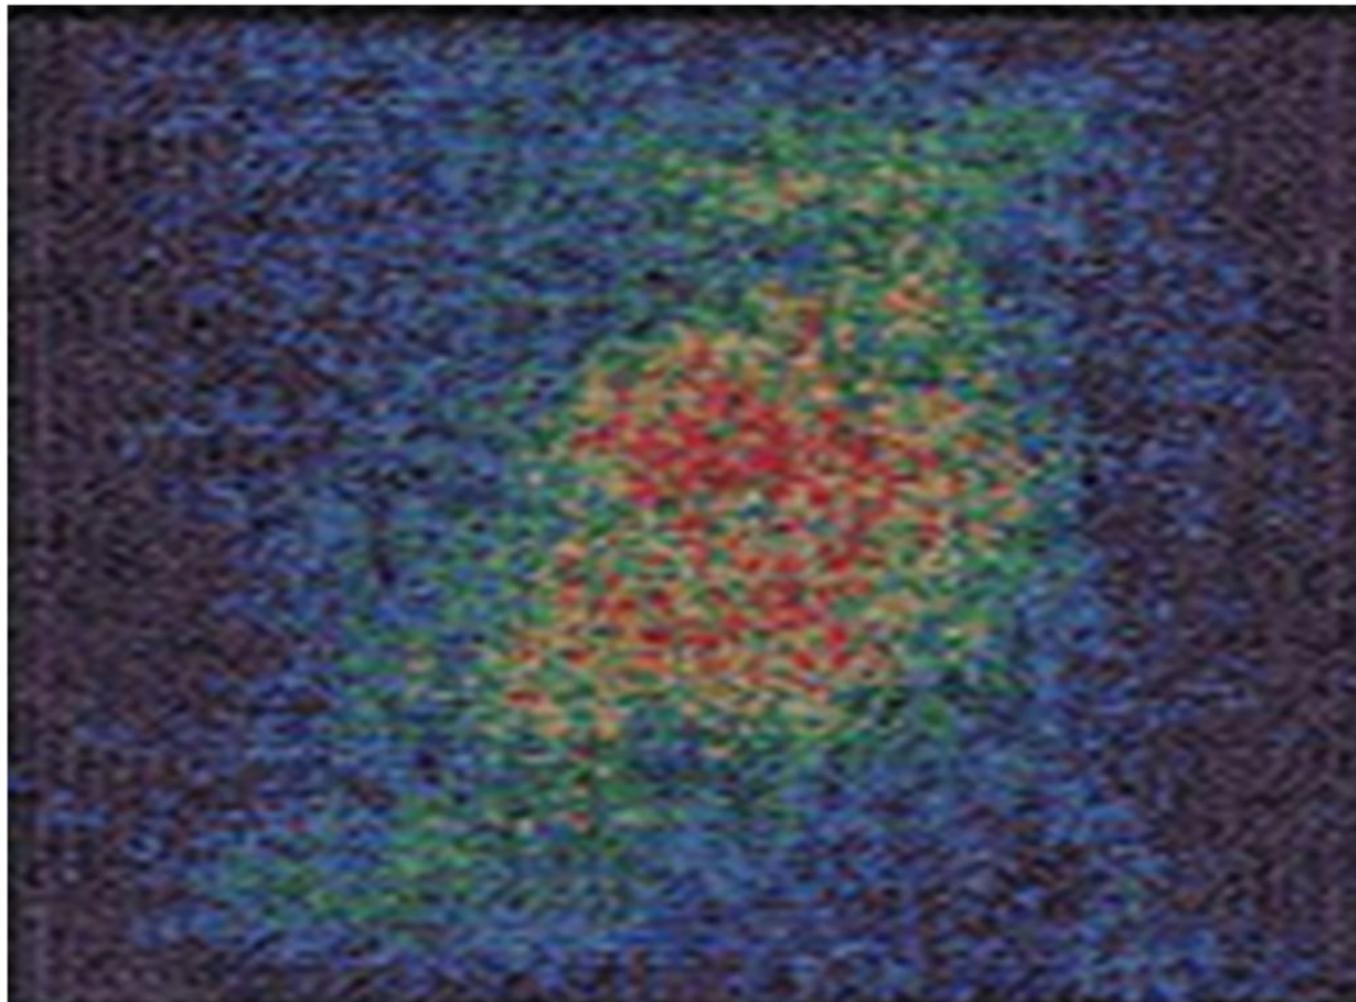


Congestion Analysis Pre-Synthesis Technology

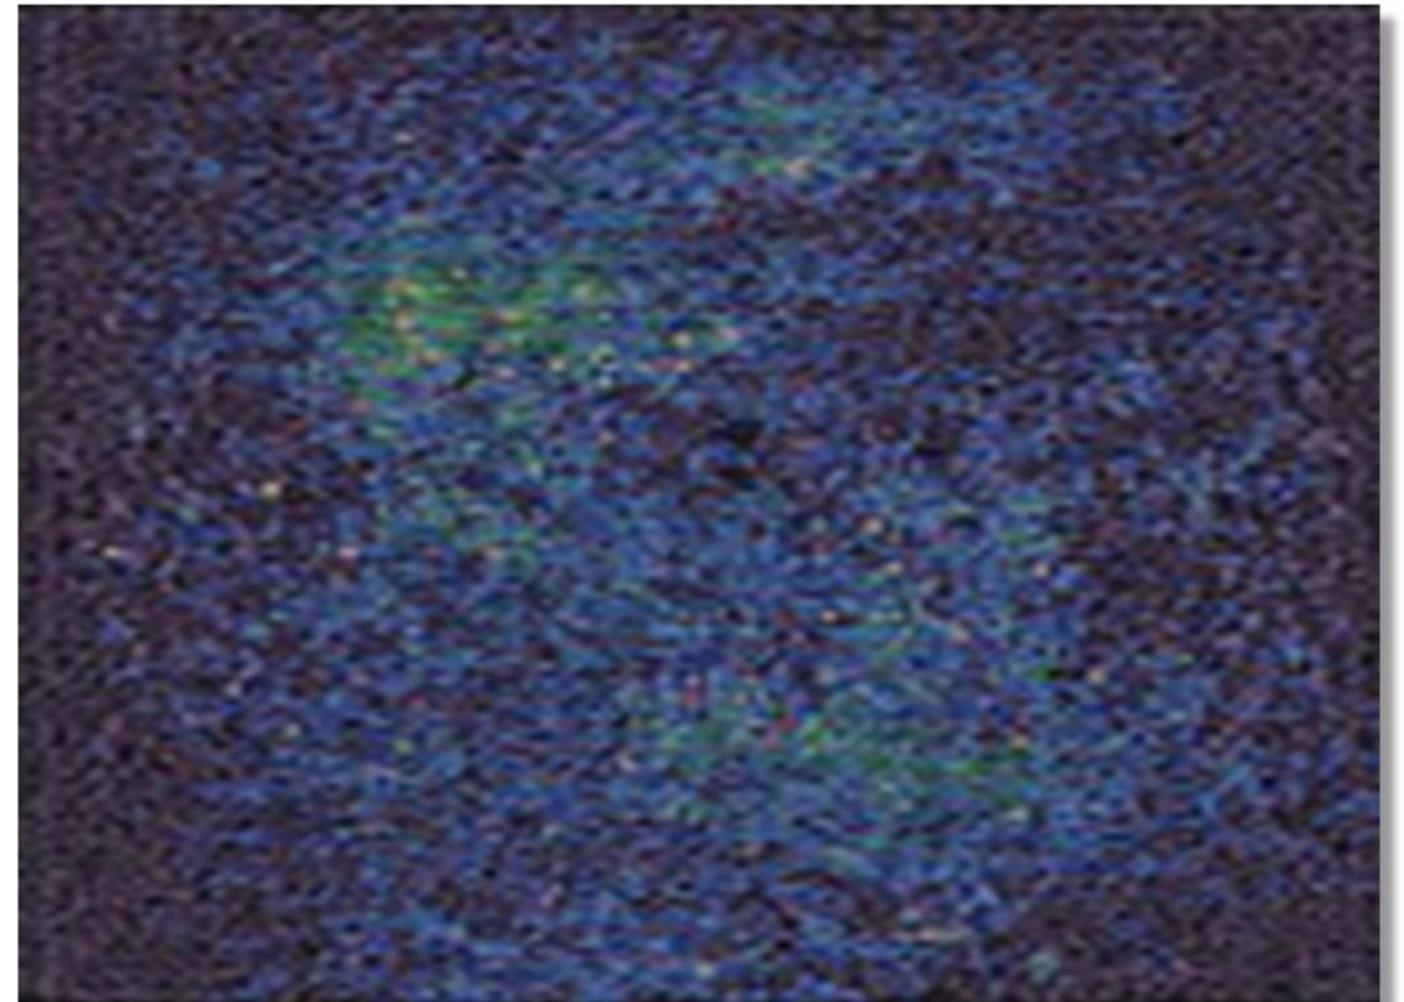
E.g. Large MUX Structures

Same-Inputs Sharing & Connectivity-Based Decomposition

Before



After



Design Compiler, 2001-2015

The Evolution of Synthesis !

	Area	Timing	Power		Runtime
			Dynamic	Static	
2001	100	100	100	100	100
2005	78	91	82	78	20
2010	65	74	66	49	1.6 ⁽¹⁾
2015	52	63	66	38	0.6 ⁽¹⁾

(1) Design Compiler Multicore

Where Do We Stand ?

Well, Telephones Have Made a Great Deal of Progress !



Smartphones

A Revolutionary... Evolution : Convergence !

50 Things We Won't Do/Use Anymore (Mostly Because of/Thanks to Smartphones)

Yesterday

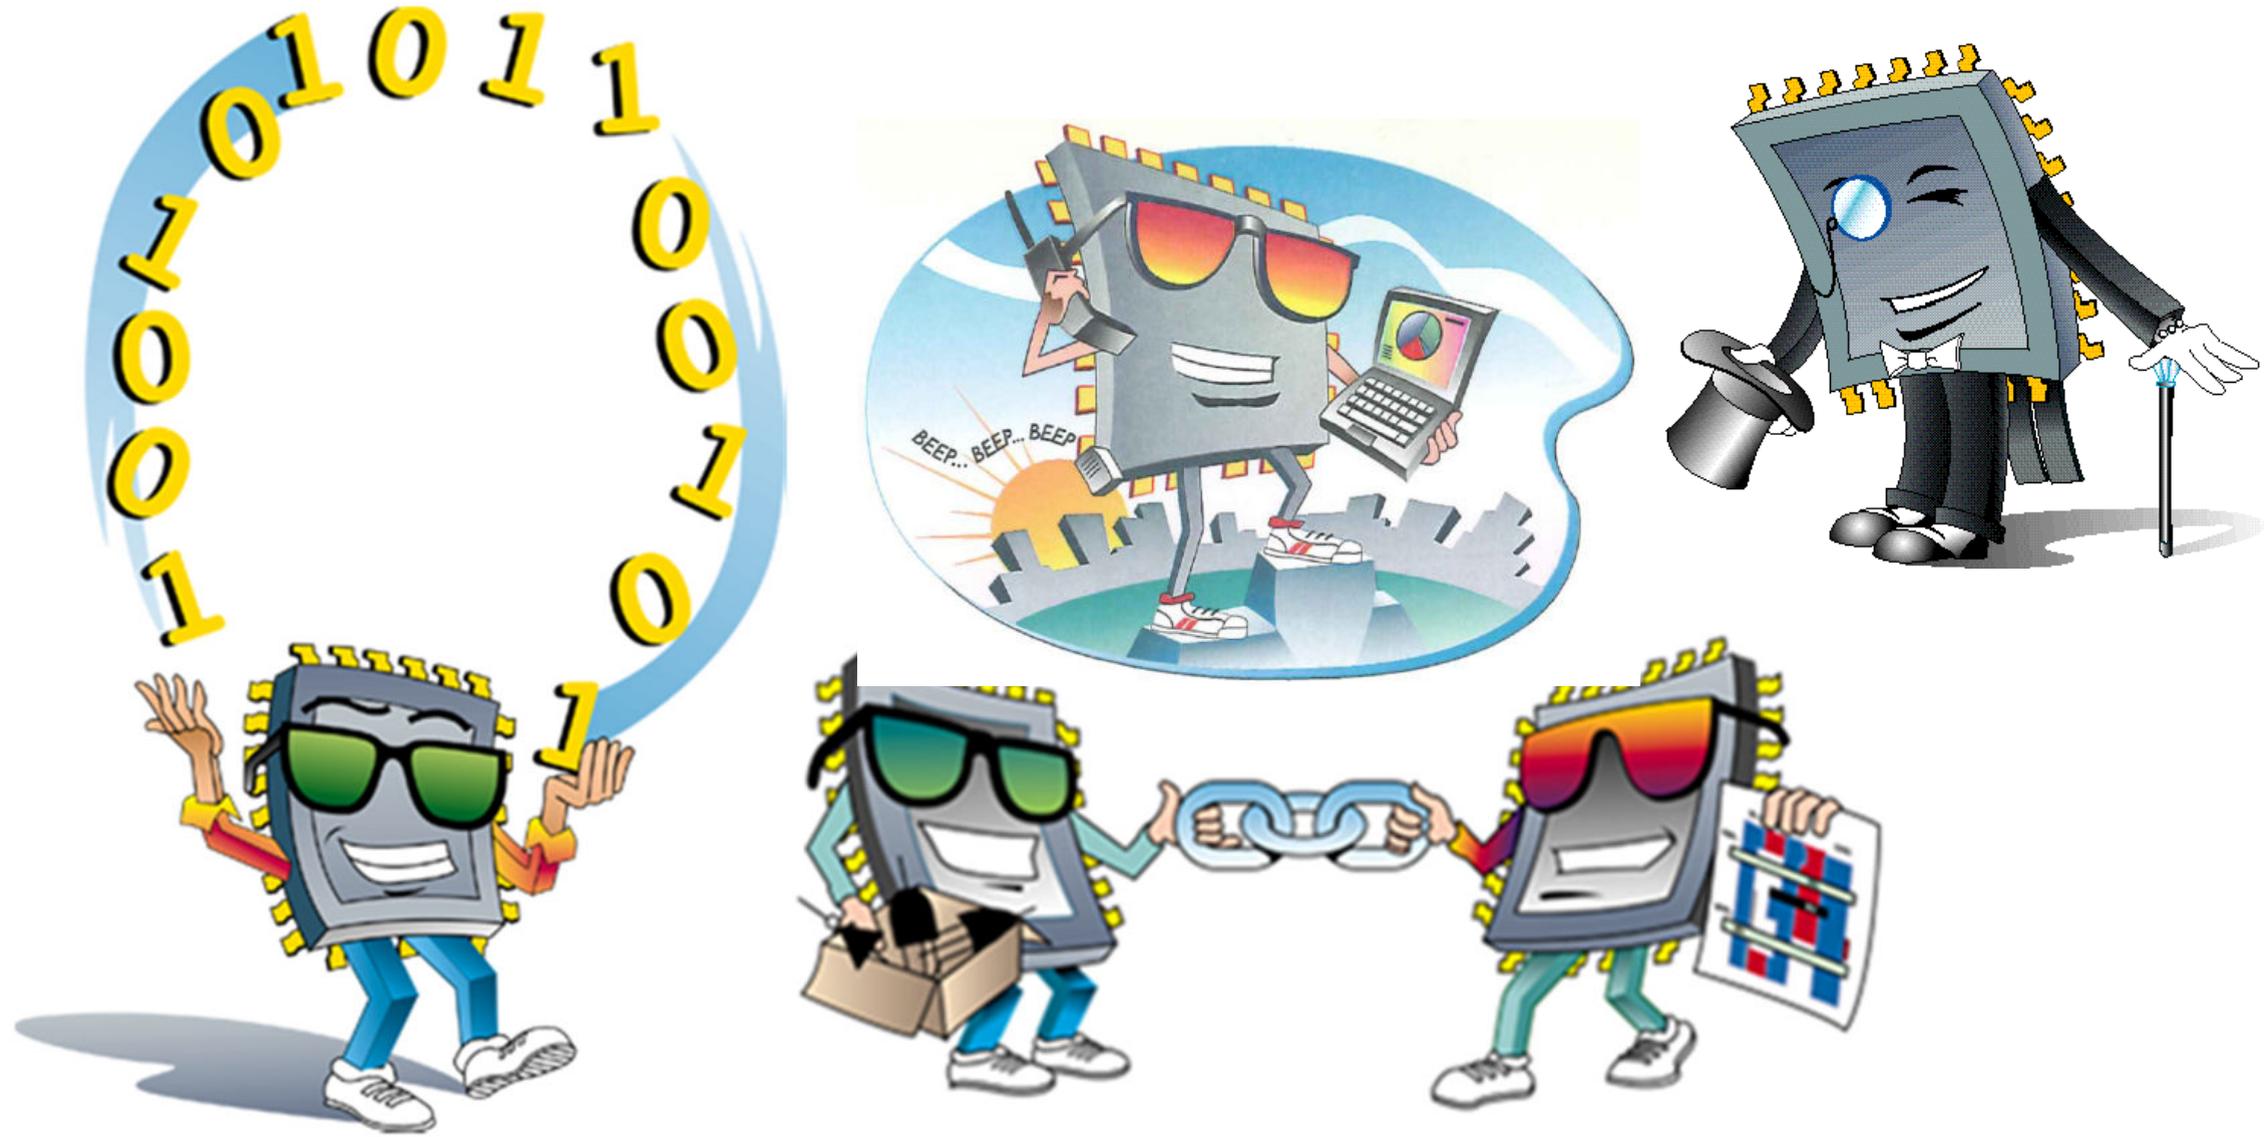


Today



Where Do We Stand ?

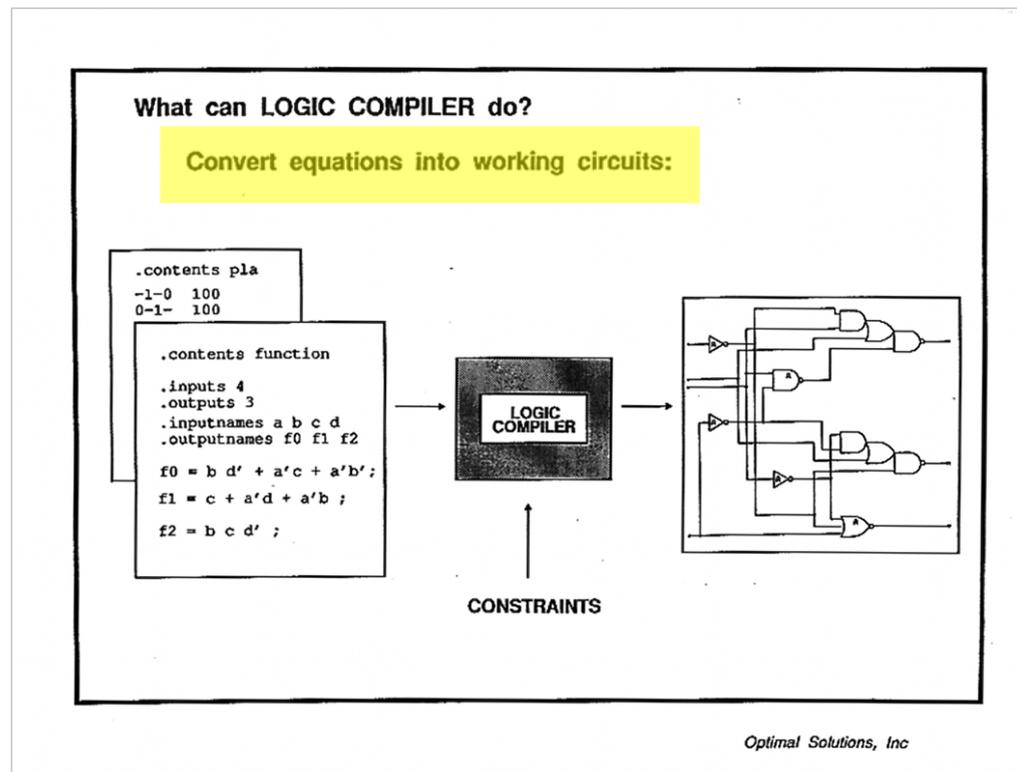
Even Synthesis Went Far Beyond Logic, Isn't It ?



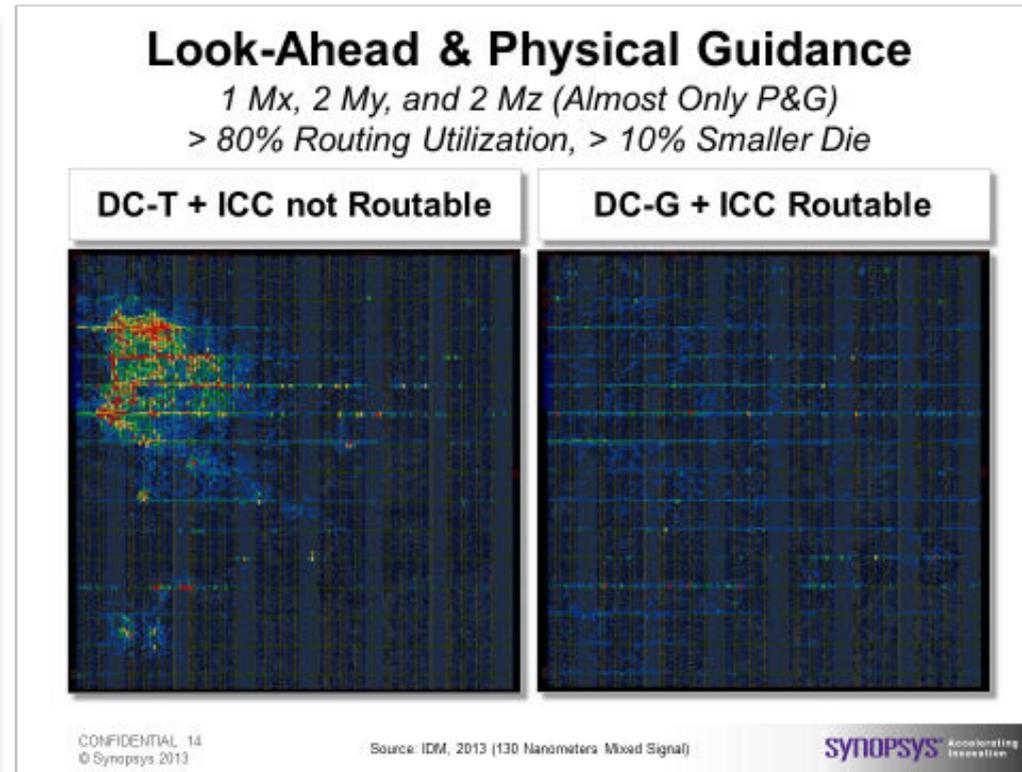
Logic Synthesis

A Revolutionary... Evolution Too : Convergence !

Logic Compiler



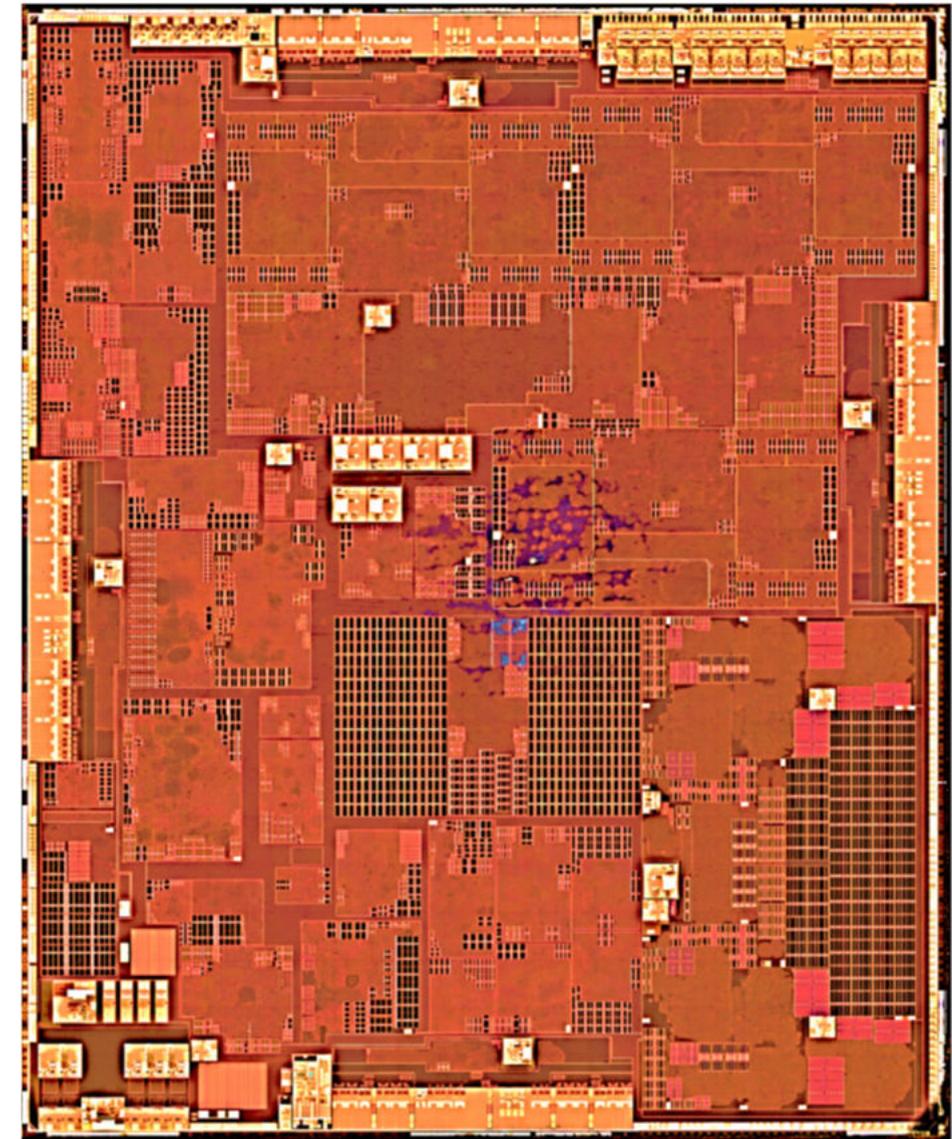
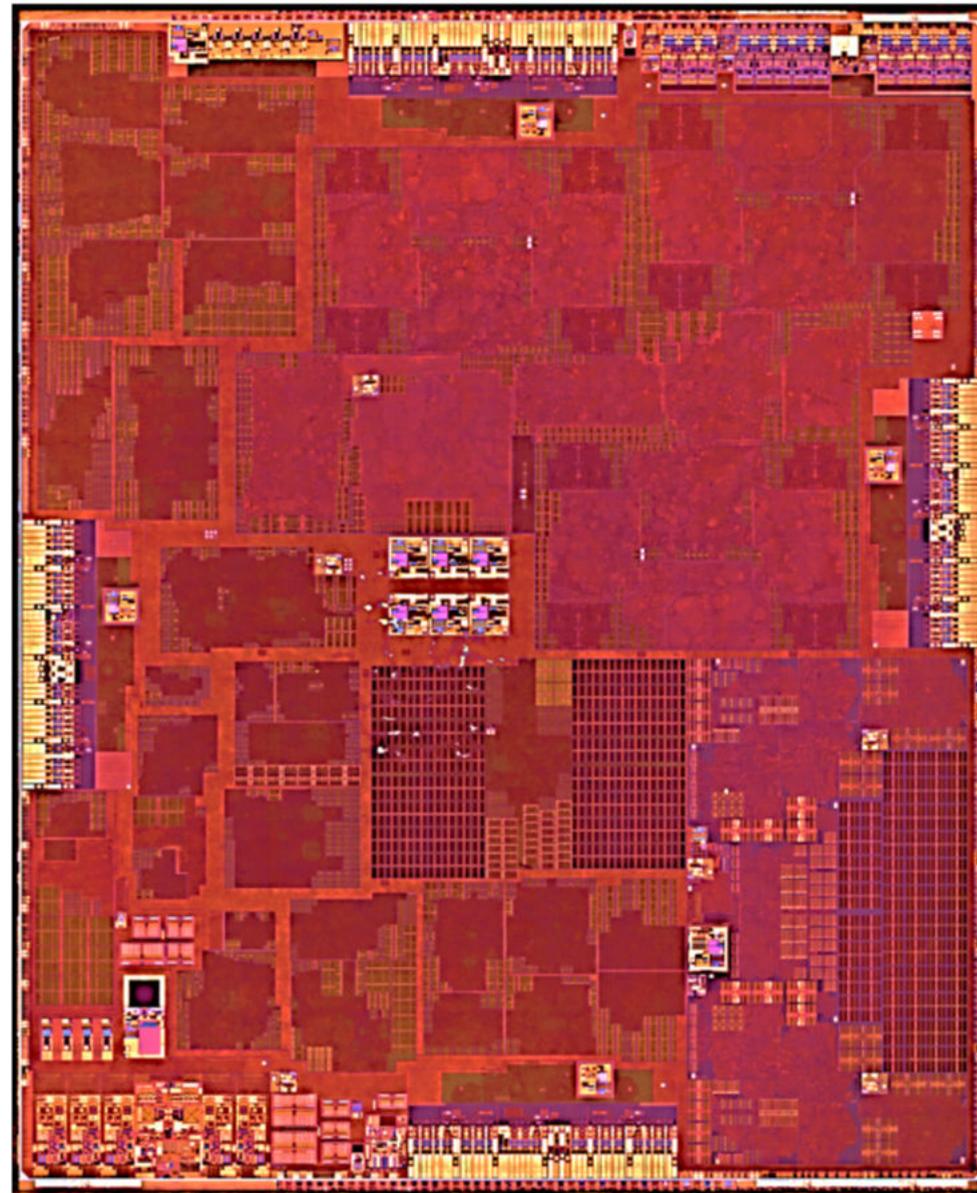
Design Compiler



From Equations to Gates, to... "Placed" and Routable Gates

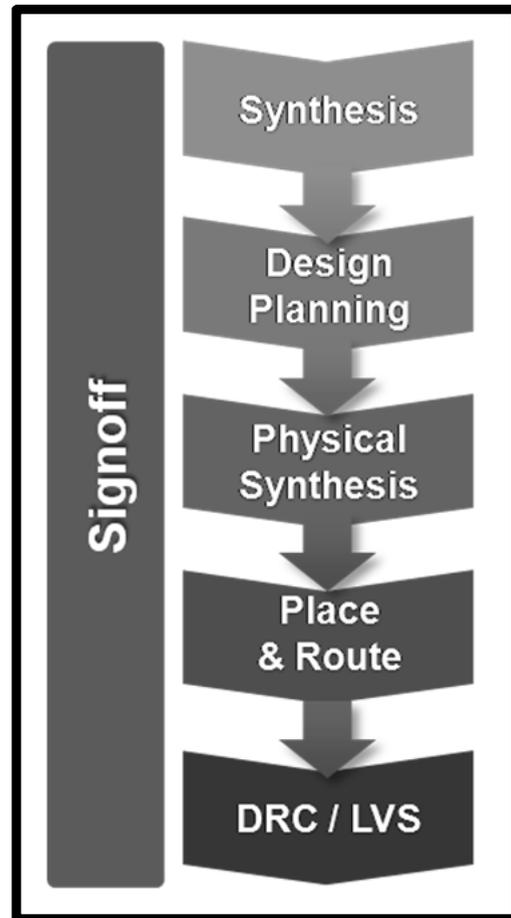
From Telephone To Logic Synthesis...

...And from Logic Synthesis to... Smartphones

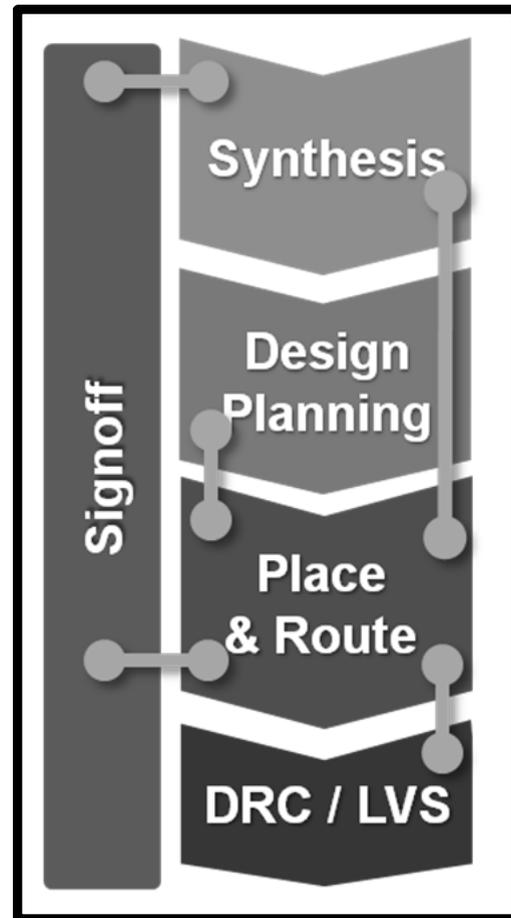


The Evolution Of Logic Synthesis

Convergence !



2003
90 Nanometers
“Interoperability”



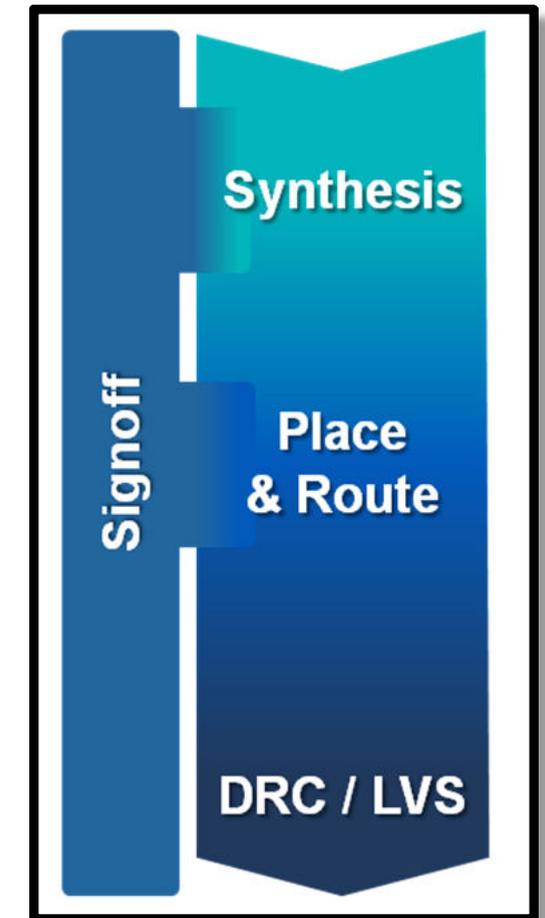
2005
65 Nanometers
“Correlation”



2007
45/40 Nanometers
“Look Ahead”



2009
32/28 Nanometers
“In-Design”



2011
22/20 Nanometers
“Convergence”

The Evolution Of Logic Synthesis

Convergence... But Not Done Yet : Still Differences !



inoff

$$G_{ij} = \bigcup_{k=0}^{p_j-1} (r^k)_j^k c_j^k = 0$$
$$(c \# p)_j = \begin{cases} c_j \cap \bar{p}_j & \text{if } i = j \\ c_j & \text{if } i \neq j \end{cases}$$
$$f(x, y) = \frac{k}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} D(x', y') \frac{\vec{r} - \vec{r}'}{|\vec{r} - \vec{r}'|^2} dx' dy'$$

DRC / LVS



Looking Into The Next Decade

There Is a Great Deal Of New Technology Ahead !

Looking Into The Next Decade

*1T Transistors per Die by the End of This Decade, “1” nm by the End of the Next
Infinitely Large, Infinitely Small, Infinitely Many*

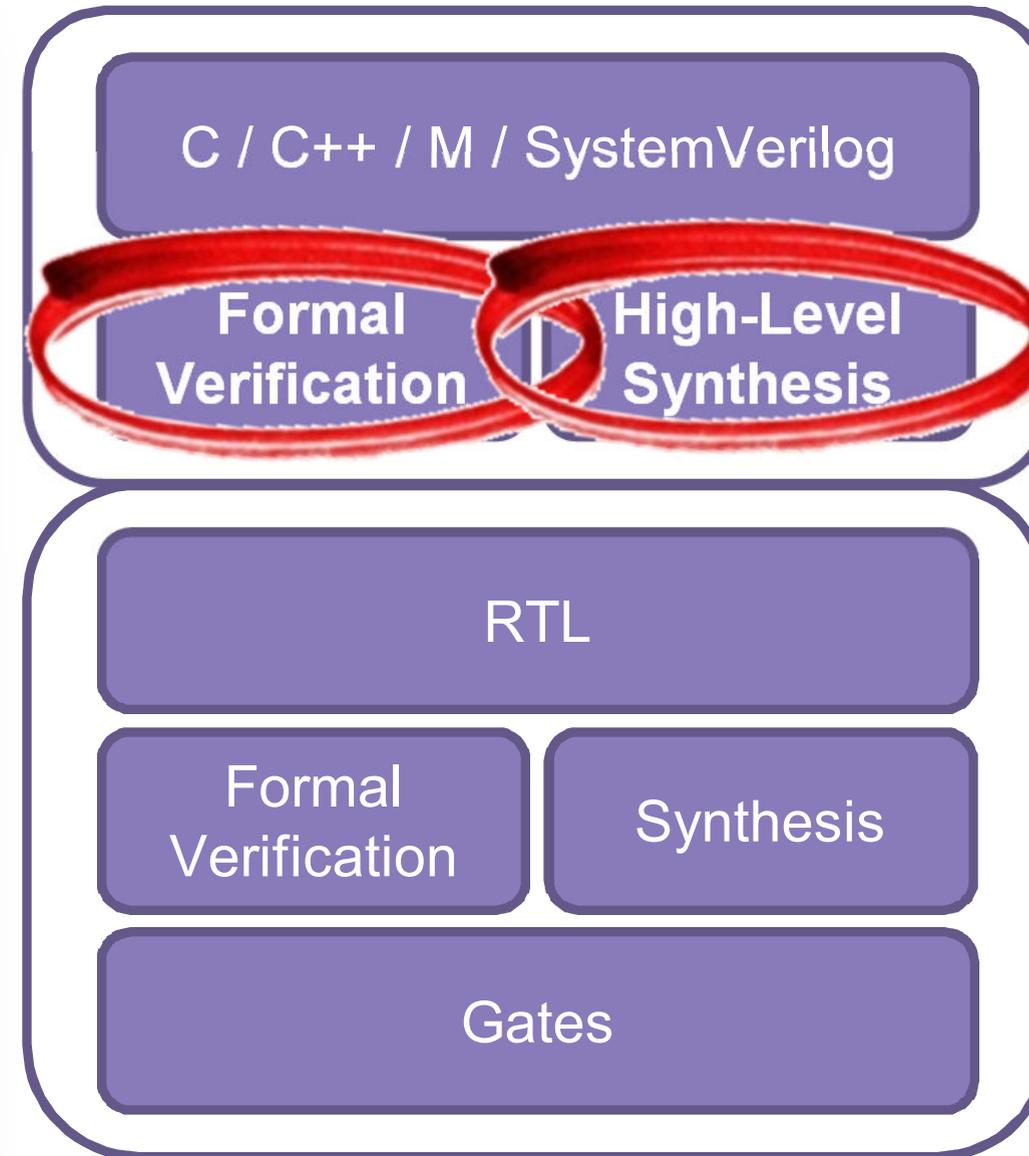
	2013	2015	2017	2019	2021	2023	2025	202x
“Technology Node” (nm)	“14”	“10”	“7”	“5”	“3.5”	“2.5”	“1.8”	“1.3”
DRAM ½ Pitch (nm)	28	24	20	17	14	12	10	7.7
MPU/ASIC ½ Pitch (nm)	40	32	25	20	16	13	10	7
FLASH ½ Pitch (nm)	18	15	13	11	9	8	8	8
MPU Printed Gate Length (nm)	28	22	18	14	11	9	7	5
MPU Physical Gate Length (nm)	20	17	14	12	10	8	7	5
Theoretical Integration Capacity (BT)	64	128	256	512	1024	2048	4096	8192

(Assuming 450mm Wafers in Production in 2018, and EUV in Production after 10nm)

	2013	2015	2017	2019	2021	2023	2025	202x
Size of Internet (IP Addresses)		25B		~ 50B			100B	

Looking Into The Next Decade

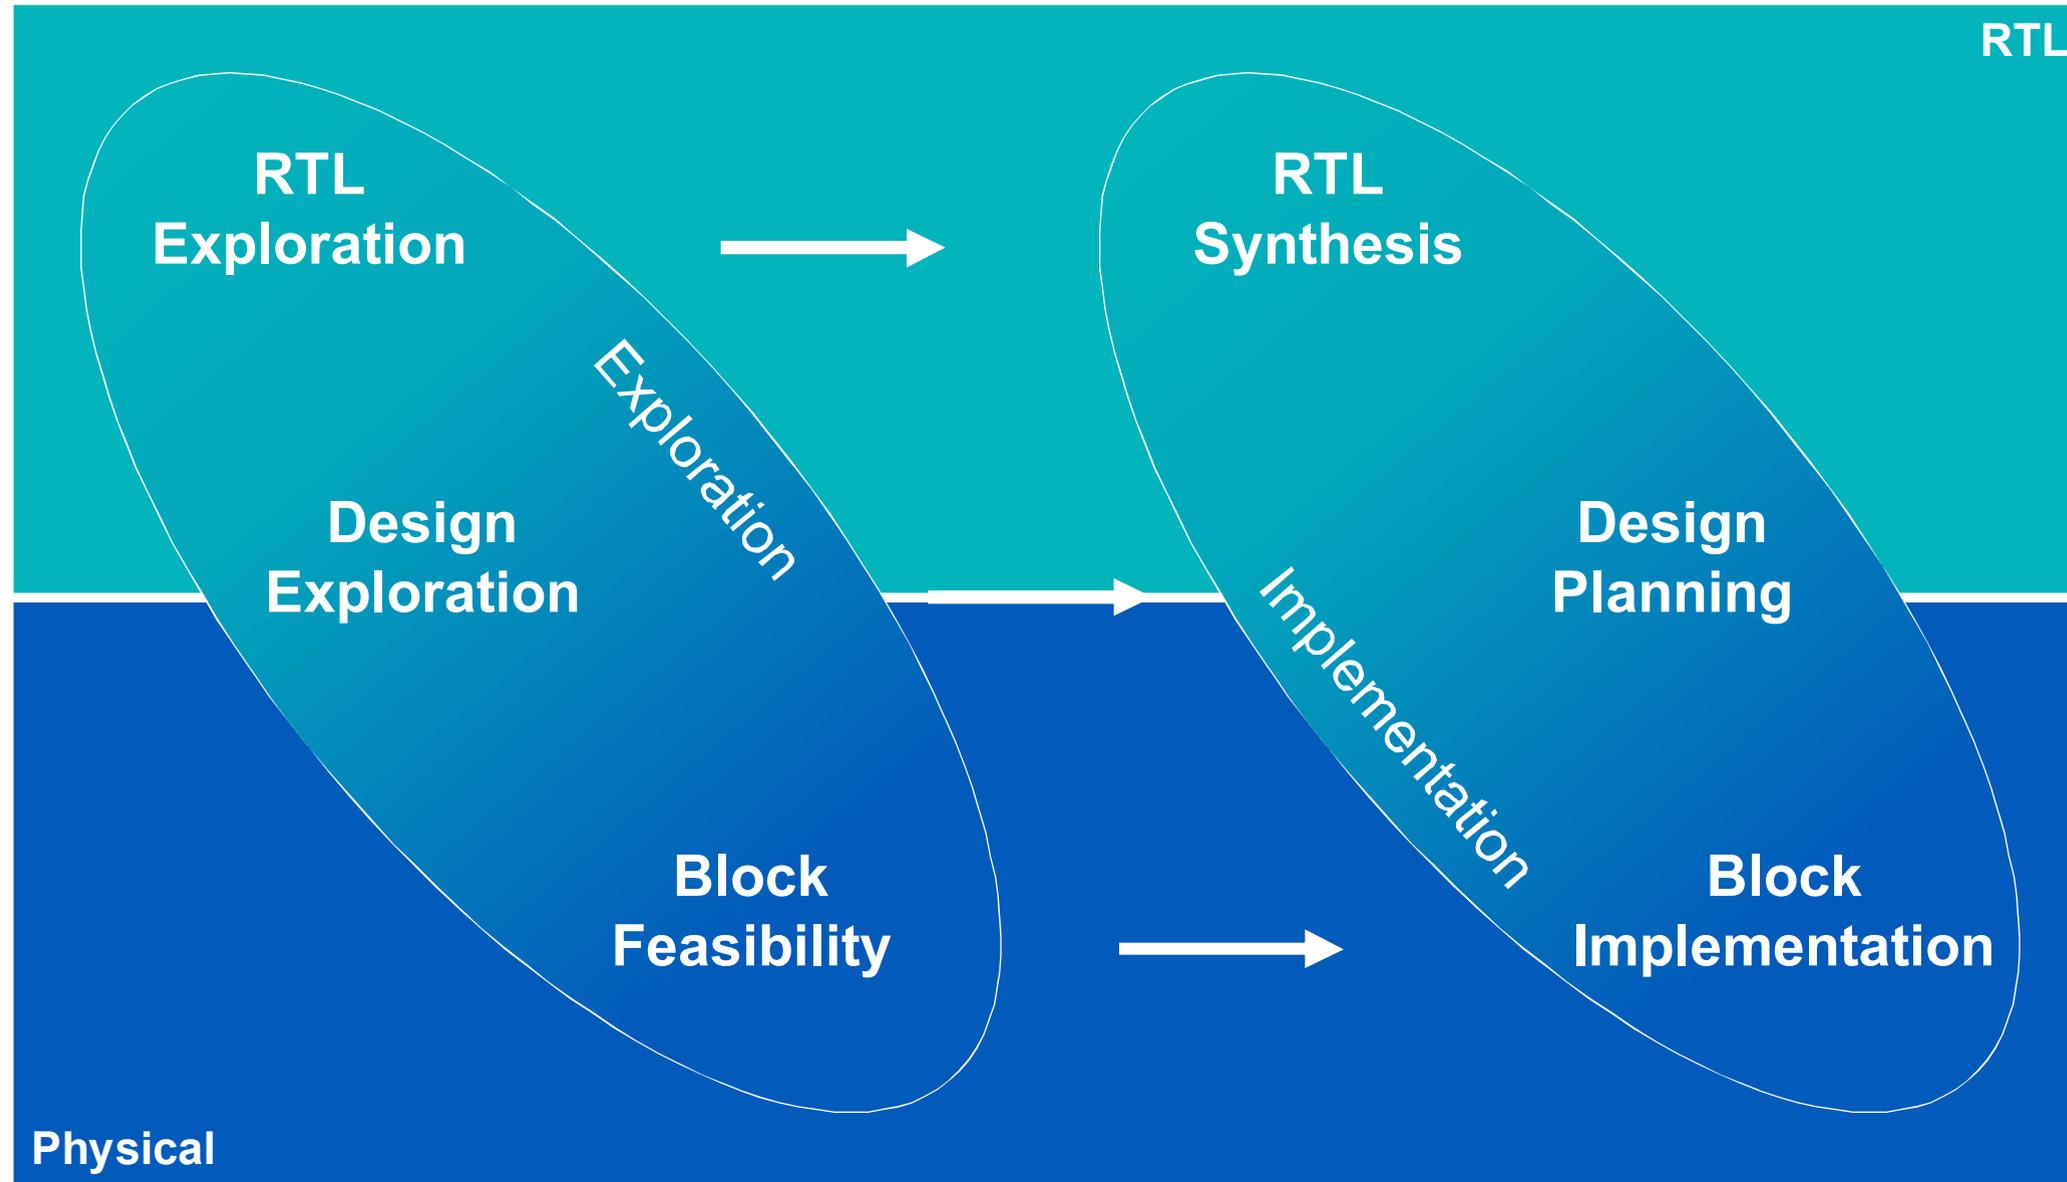
Making the Transition to High-Level Design... Again !



Looking Into The Next Decade

Evolving Towards Two, Closely Connected “Sub-Systems”

Explore & Analyze, then Implement



Looking Into The Next Decade

Evolving Towards Two, Closely Connected “Sub-Systems”

Explore & Analyze, then Implement

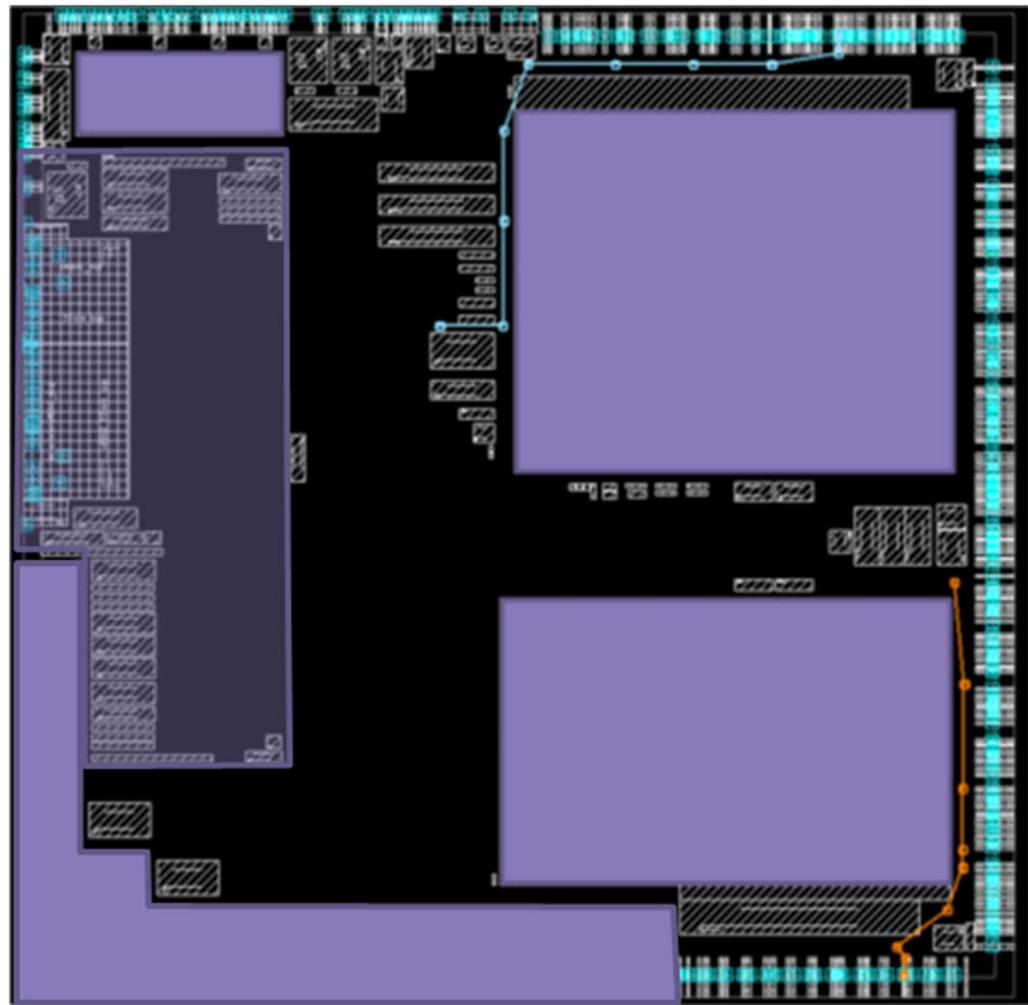
- Until now :
 - Mostly “preserve” the gates
 - “Change” placement, and/or routing to close timing, power, ...
 - Very, very time consuming, and
 - Leads to an infinite number of iterations
- In the future :
 - Once the objective is “within reach” ...
 - “Hold” placement and routing
 - Systematically “change” the gates
 - Same footprint, different timing, power, temperature inversion point, etc.
 - The richness of the library is fundamental

Looking Into The Next Decade

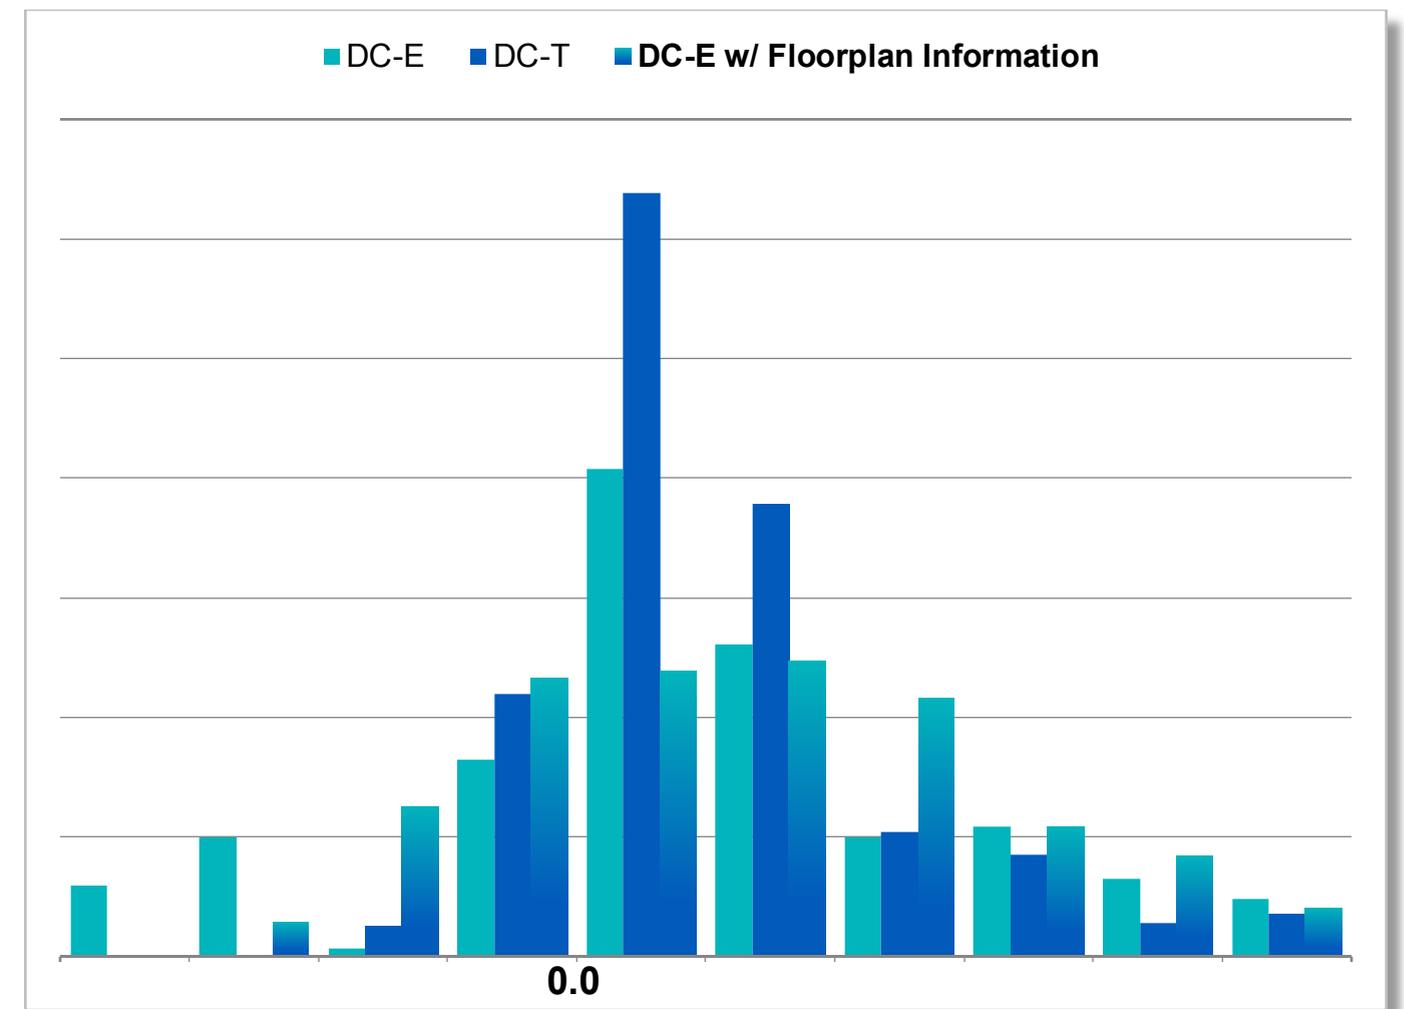
RTL Exploration Is ~ 6X Faster Than Full Synthesis

Slack Distribution Comparison, Correlation $\pm 8\%$

Floorplan Information



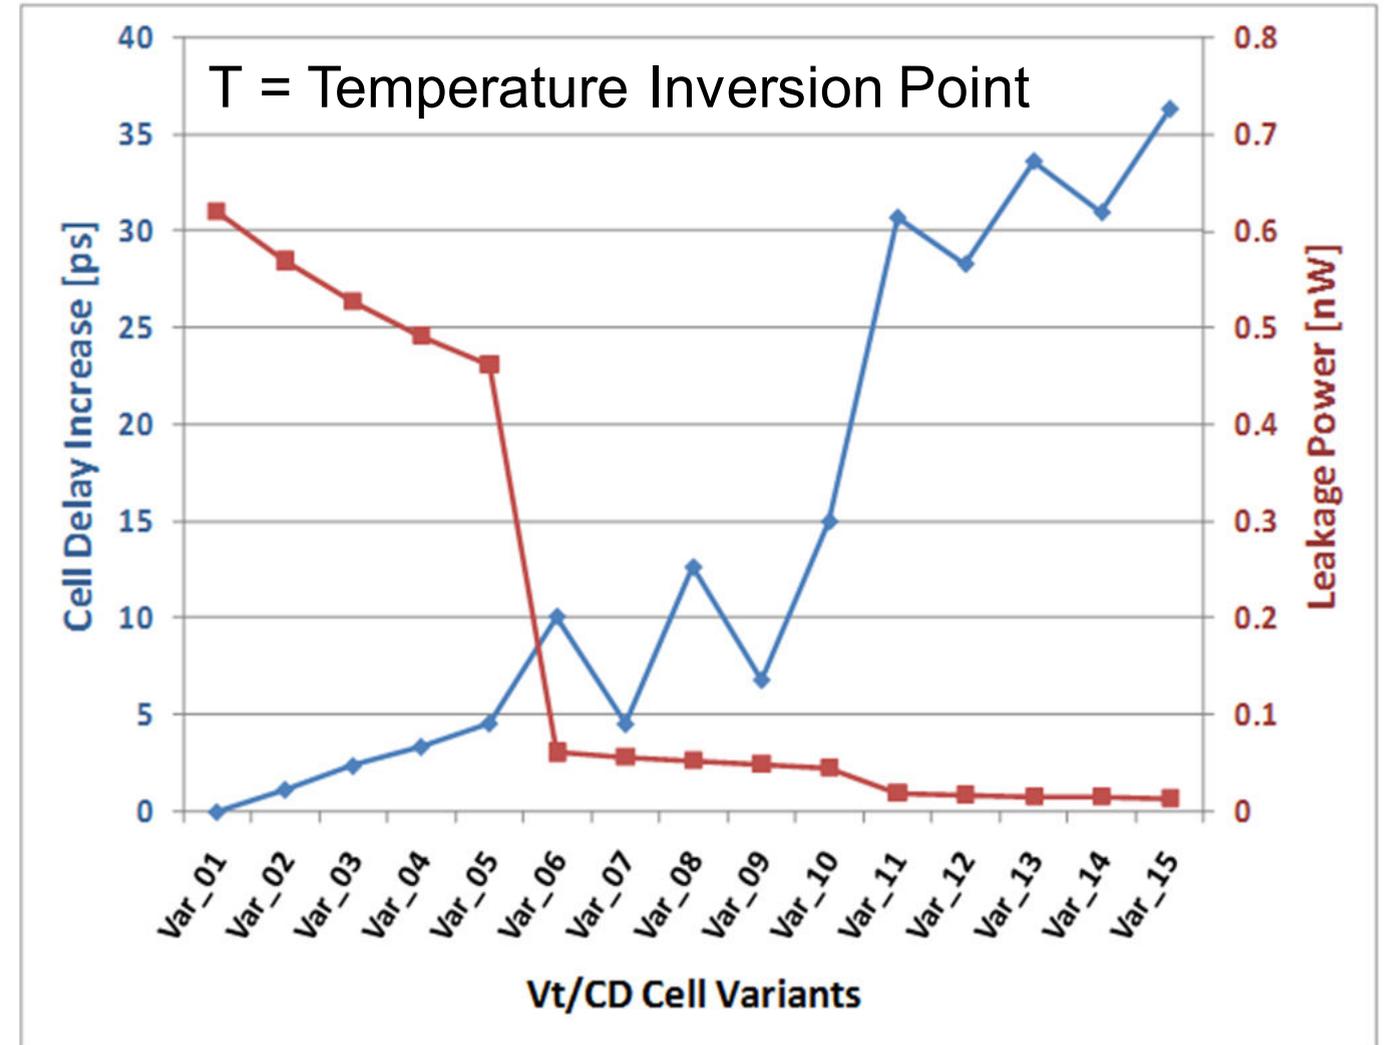
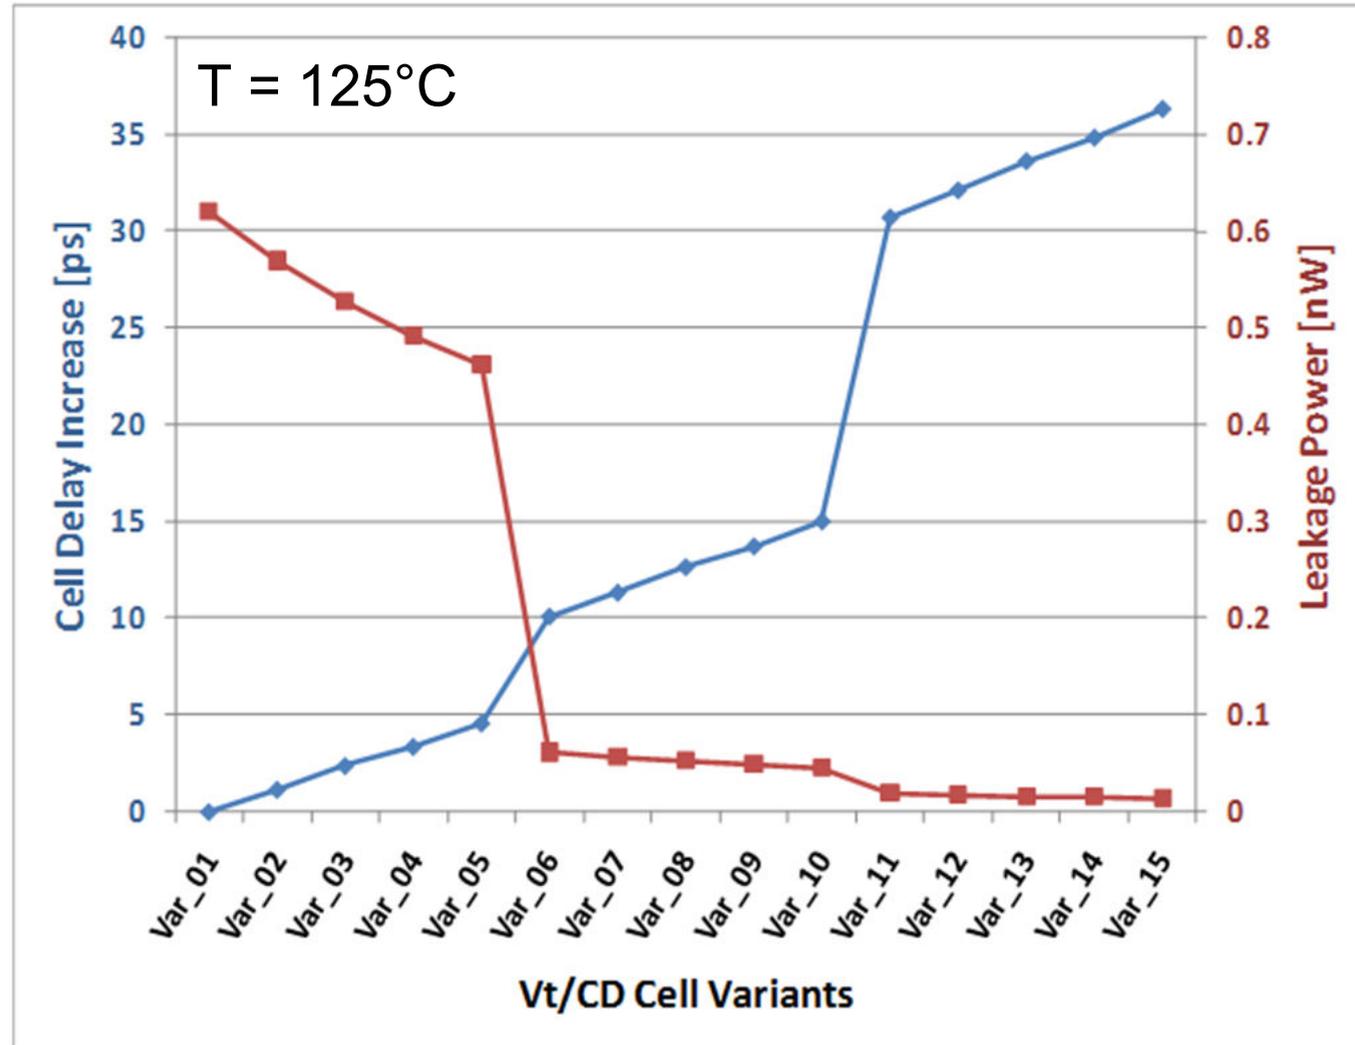
Slack Histogram



Looking Into The Next Decade

Today Libraries Contain Thousand of Elements

Many Variants with the Same Footprint but Different “Performance”

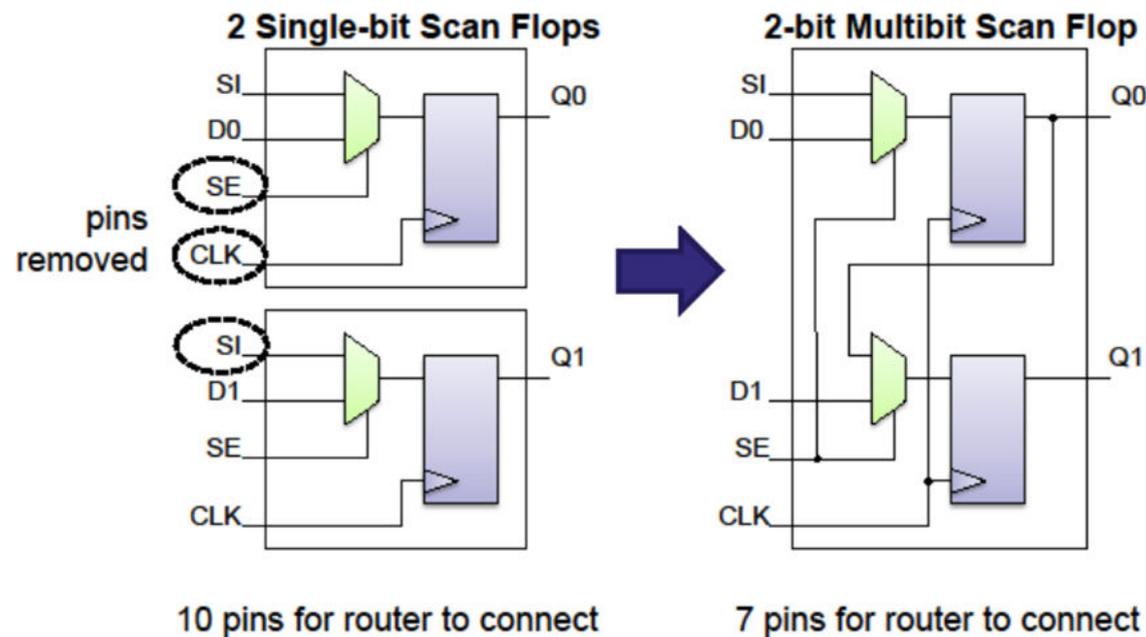


Looking Into The Next Decade

Extending the Use of Multi-Bit Structures Will Save Area and Power, and Will Alleviate Congestion, Simplifying Routing

Multi-bit Flip-Flop (MBFF)

Fewer pins for the router to connect

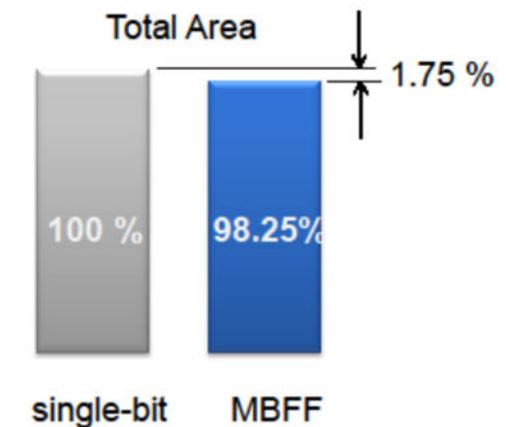


Non-MBFF vs. MBFF - Area Results

1.75% Area Savings

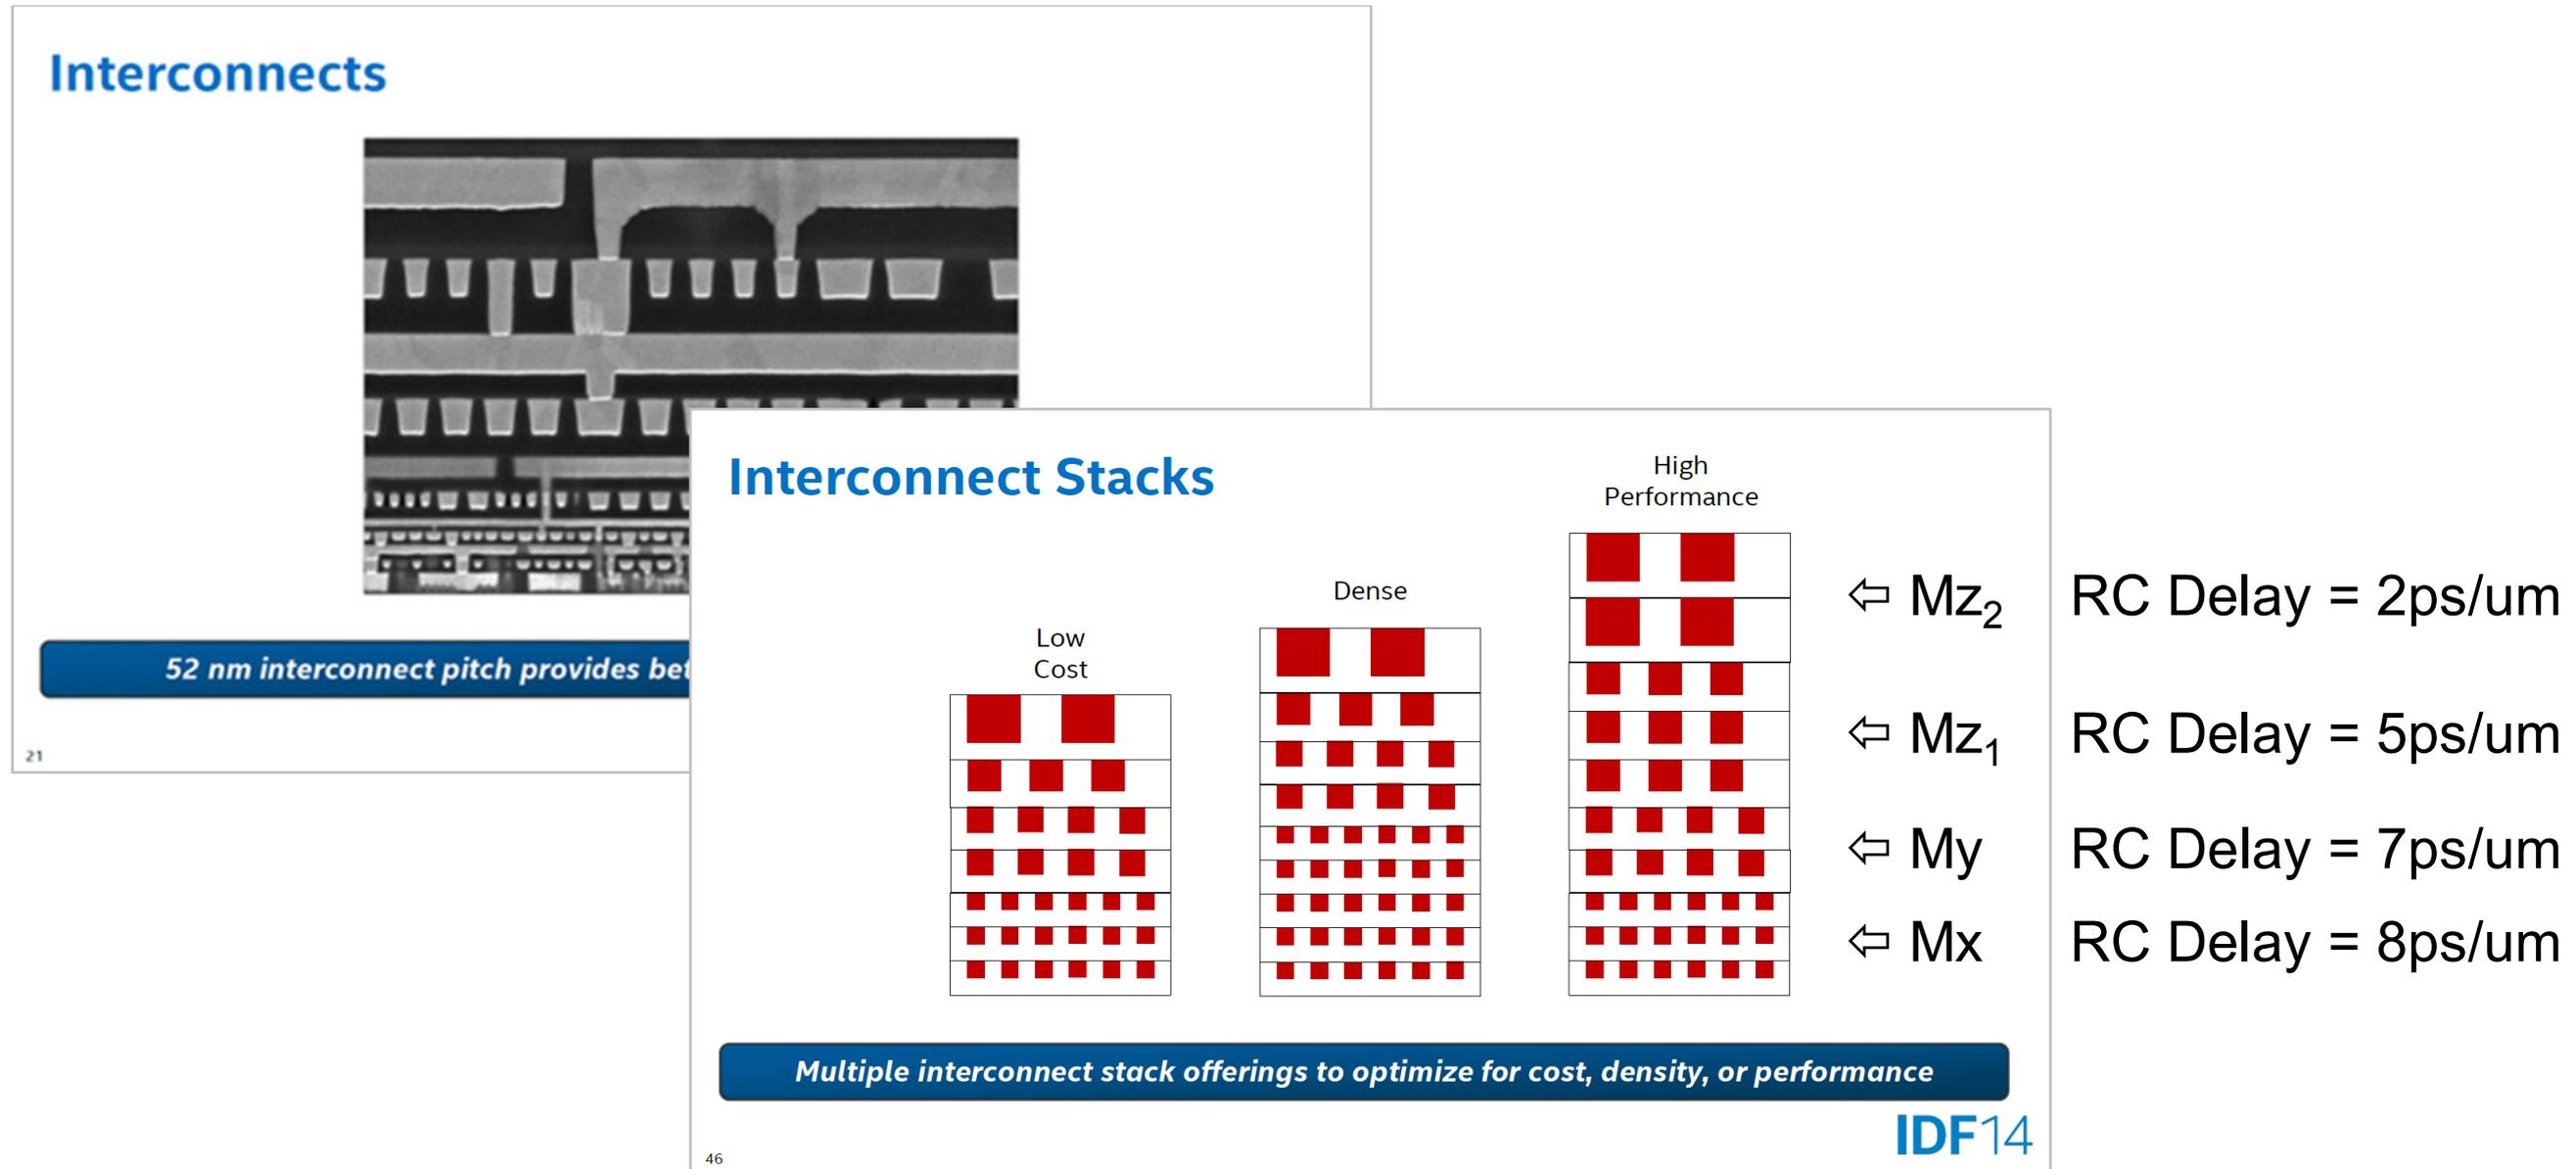


Savings	DC-G	ICC (post-CTS)
Register Area	- 7.3%	- 7.10%
Overall Std-cell Area	- 2.6%	- 1.75%



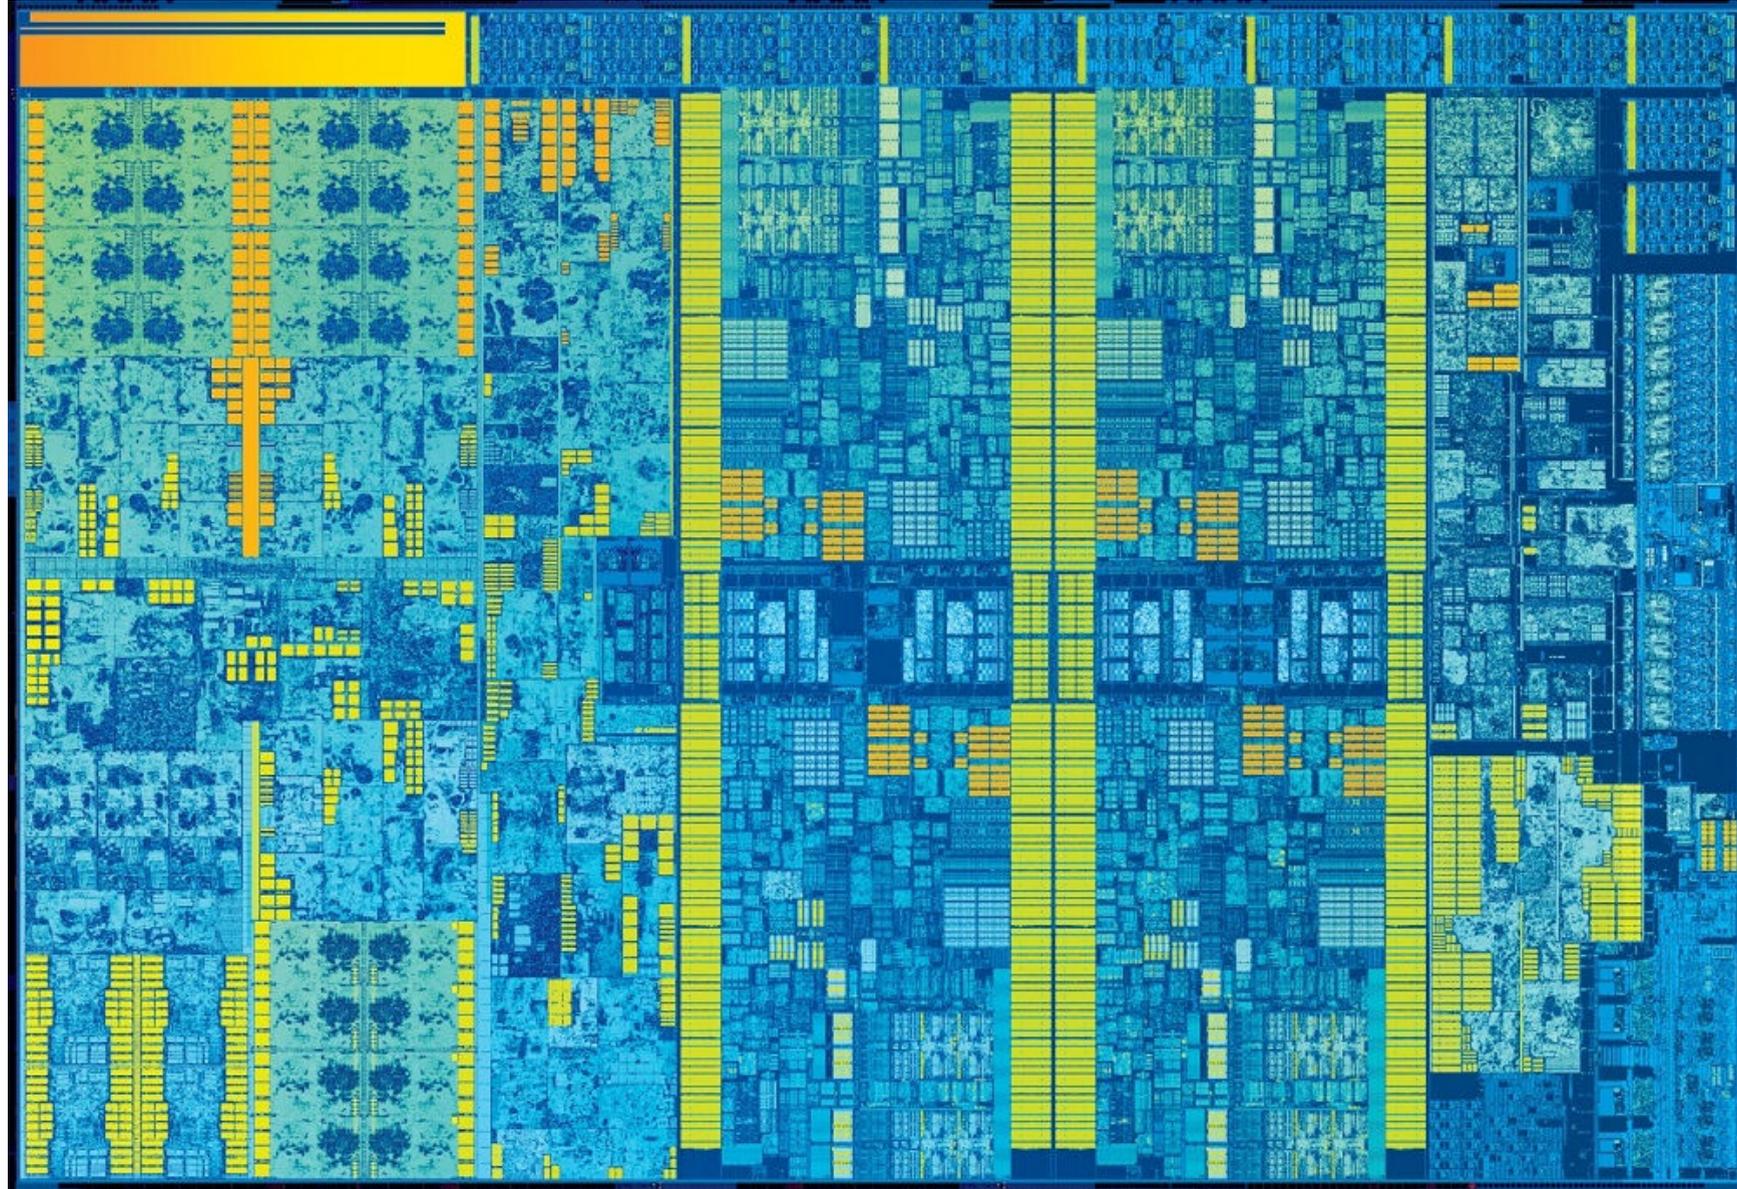
Convergence Of Synthesis And Implementation

RC Variation Among the Many Metal Layers Makes Estimates Extremely Difficult Forcing to Bring Global Routing, and Probably Detailed, into the Synthesis Picture



Looking Into The Next Decade

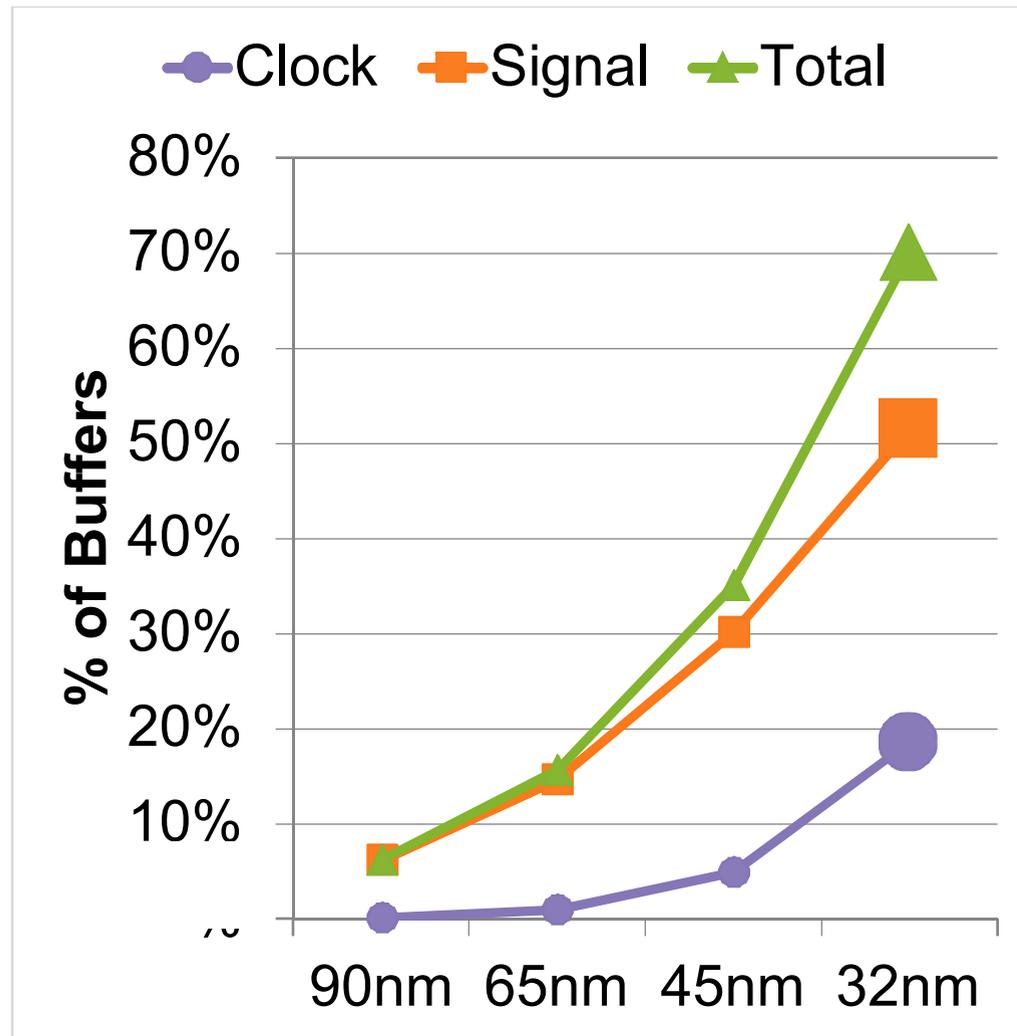
*Today MPU Have 4-8 Cores, 16 Cores Are Just Around the Corner
Logic Synthesis Algorithms Must Be Suitable for Fine Parallelization*



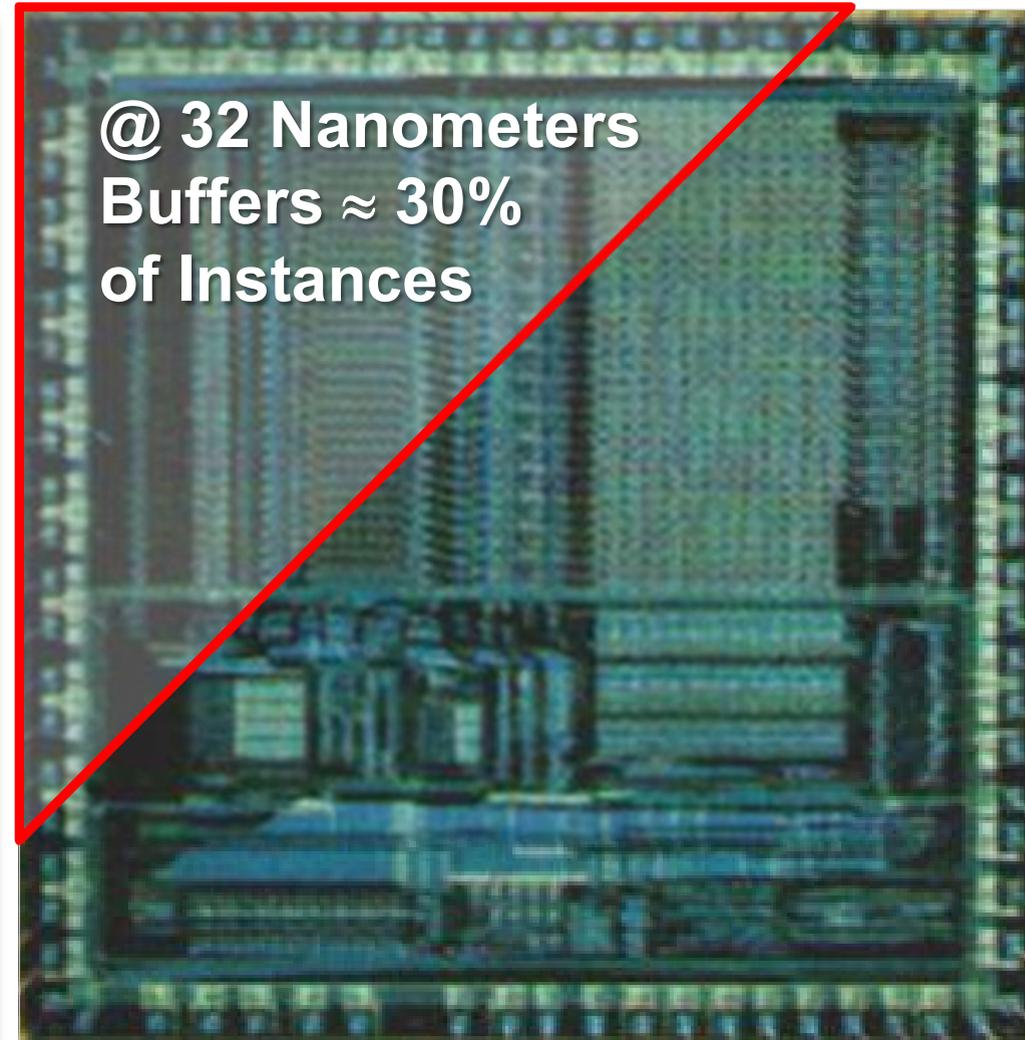
Looking Into The Next Decade

Buffers Insertion Optimization

A 2003 Estimate...

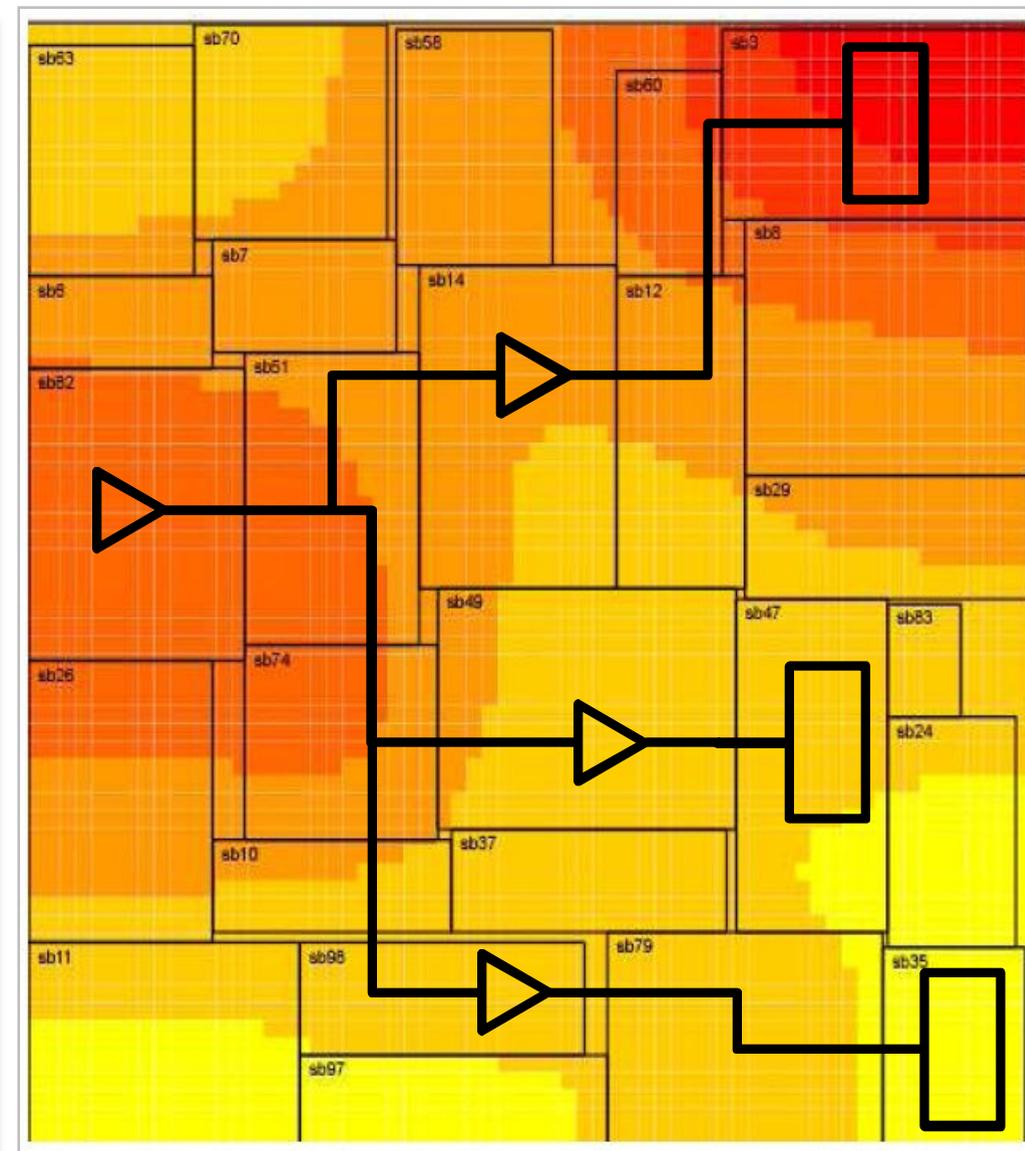
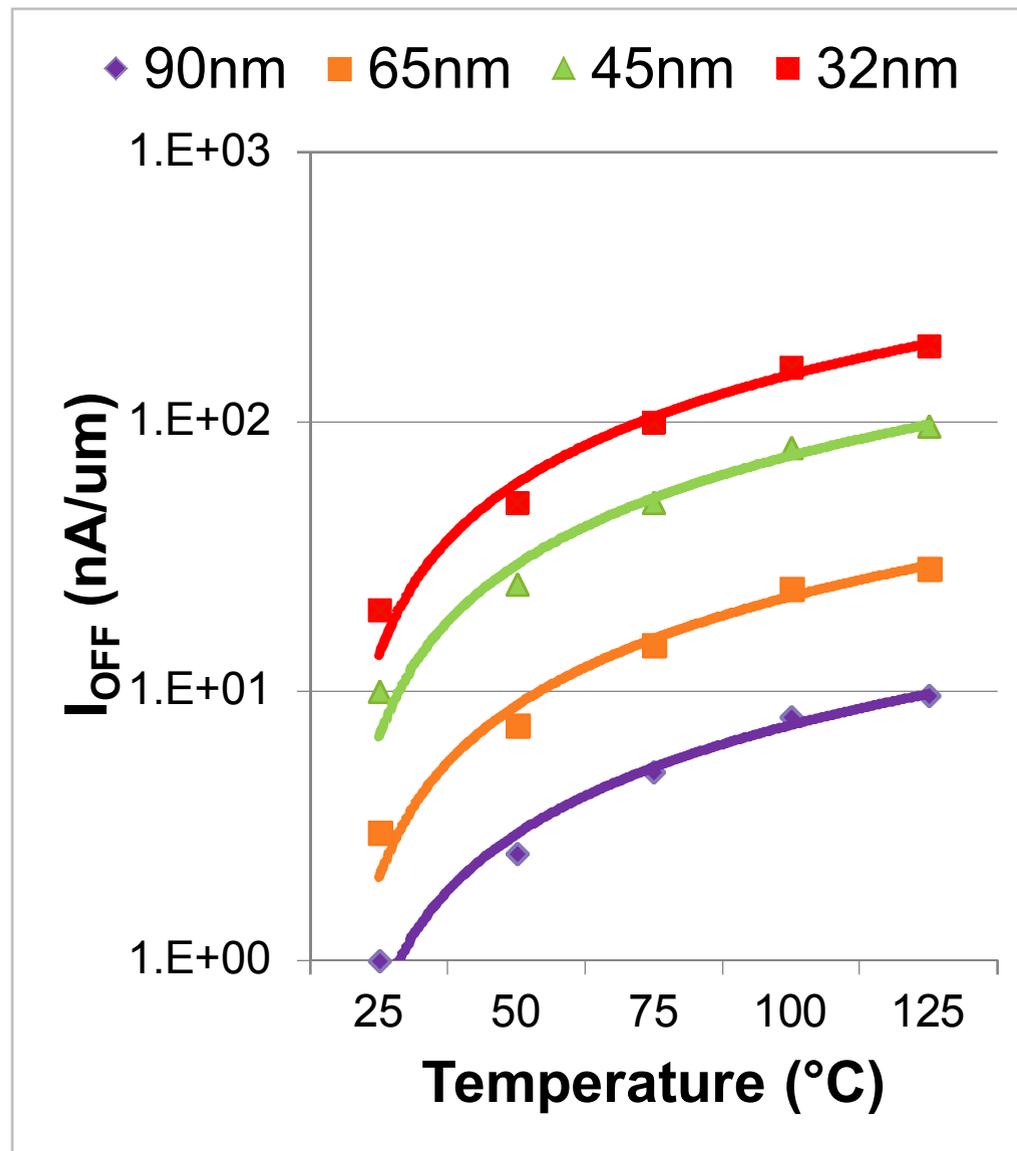


A 2010 Data Point...



Looking Into The Next Decade

Temperature-Aware Synthesis



Looking Into The Next Decade

Andreas Kuehlmann, Robert K. Brayton, Alan Mishchenko,
Luca Amarù, Pierre-Emmanuel Gaillardon, Giovanni De Micheli

IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS, VOL. 31, NO. 12, DECEMBER 2002

1377

Robust Boolean Reasoning for Equivalence Checking and Functional Property Verification

Andreas Kuehlmann, Senior Member, IEEE, Viresh Paruthi, Florian Krohm, and Malay K. Ganai, Member, IEEE

Abstract—Many tasks in computer-aided design (CAD), such as equivalence checking, property checking, logic synthesis, and false paths analysis, require efficient Boolean reasoning for problems derived from circuits. Traditionally, canonical representations, e.g., binary decision diagrams (BDDs), or structural satisfiability (SAT) methods, are used to solve different problem instances. Each of these techniques offer specific strengths that make them efficient for particular problem structures. However, neither structural techniques based on SAT, nor functional methods using BDDs offer an overall robust reasoning mechanism that works reliably for a broad set of applications. The authors present a combination of techniques for Boolean reasoning based on BDDs, structural transformations, an SAT procedure, and random simulation actively working on a shared graph representation of the problem. The described intertwined integration of the four techniques results in a powerful summation of their orthogonal strengths. The presented reasoning technique was mainly developed for formal equivalence checking and property verification but can equally be used in other CAD applications. The authors' experiments demonstrate the effectiveness of the approach for a broad set of applications.

Index Terms—BDD, Boolean reasoning, equivalence checking, formal verification, property checking, SAT.

I. INTRODUCTION

MANY tasks in computer-aided design (CAD) such as equivalence or property checking, logic synthesis, timing analysis, and automatic test-pattern generation, require Boolean reasoning on problems derived from circuit structures. There are two main approaches used alternatively for such applications. First, by converting the problem into a functionally canonical form such as binary decision diagrams (BDDs), the solution can be obtained from the resulting diagram. Second, structural satisfiability (SAT) procedures perform a systematic search for a consistent assignment on the circuit representation. The search either encounters a solution or, if all cases have been enumerated, concludes that no solution exists. Both approaches generally suffer from exponential worst case complexity. However, they have distinct strengths and weaknesses which make them applicable to different classes of practical problems.

Manuscript received December 28, 2001; revised April 18, 2002. This paper was recommended by Associate Editor J. H. Kukula.
A. Kuehlmann is with the Cadence Berkeley Labs, Berkeley, CA 94704 USA (e-mail: kua@cadence.com).
V. Paruthi is with the IBM Enterprise Systems Group, Austin, TX 78758 USA (e-mail: vparuthi@us.ibm.com).
F. Krohm is with the IBM Microelectronic Division, Hopewell Junction, NY 12533 USA (e-mail: florian@usmail.fsk.ibm.com).
M. K. Ganai is with the NEC C&C Research Labs, Princeton, NJ 08540 USA (e-mail: malay@nec-lab.com).
Digital Object Identifier 10.1109/TCAD.2002.804386

A monolithic integration of SAT and BDD-based techniques could combine their individual strengths and result in a powerful solution for a wider range of applications. Additionally, by including random simulation its efficiency can be further improved for problems with many satisfying solutions.

A large fraction of practical problems derived from the above-mentioned applications have a high degree of structural redundancy. There are three main sources for this redundancy: first, the primary netlist produced from a register transfer level (RTL) specification contains redundancies generated by language parsing and processing. For example, in industrial designs, between 30% and 50% of generated netlist gates are redundant [1]. A second source of structural redundancy is inherent to the actual problem formulation. For example, a miter structure [2], built for equivalence checking, is globally redundant. It also contains many local redundancies in terms of identical substructures used in both designs to be compared. A third source of structural redundancy originates from repeated invocations of Boolean reasoning on similar problems derived from overlapping parts of the design. For example, the individual paths checked during false paths analysis are composed of shared subpaths which get repeatedly included in subsequent checks. Similarly, a combinational equivalence check of large designs is decomposed into a series of individual checks of output and next-state functions which often share a large part of their structure. An approach that detects and reuses structural and local functional redundancies during problem construction could significantly reduce the overhead of repeated processing of identical structures. Further, a tight integration with the actual reasoning process can increase its performance by providing a mechanism to efficiently handle local decisions.

In this paper, we present an incremental Boolean reasoning approach that integrates structural circuit transformation, BDD sweeping [3], a circuit-based SAT procedure, and random simulation in one framework. All four techniques work on a shared AND/INVERTER graph [3] representation of the problem. BDD sweeping and SAT search are applied in an intertwined manner both controlled by resource limits that are increased during each iteration [4]. BDD sweeping incrementally simplifies the graph structure, which effectively reduces the search space of the SAT solver until the problem can be solved. The set of circuit transformations get invoked when the sweeping causes a structural change, potentially solving the problem or further simplifying the graph for the SAT search. Furthermore, random simulation can efficiently handle problems with dense solution spaces.

This paper is structured as follows. Section II summarizes previous work in the area and contrasts it to our contributions. Section III presents the AND/INVERTER graph representation,

ABC: An Academic Industrial-Strength Verification Tool

Robert Brayton Alan Mishchenko

EECS Department, University of California, Berkeley, CA 94720, USA
(brayton, alanm)@eecs.berkeley.edu

Abstract. ABC is a public-domain system for logic synthesis and formal verification of binary logic circuits appearing in synchronous hardware designs. ABC combines scalable logic transformations based on And-Inverter Graphs (AIGs), with a variety of innovative algorithms. A focus on the synergy of sequential synthesis and sequential verification leads to improvements in both domains. This paper introduces ABC, motivates its development, and illustrates its use in formal verification.

Keywords: Model checking, equivalence checking, logic synthesis, simulation, integrated sequential verification flow.

1 Introduction

Progress in both academic research and industrial products critically depends on the availability of cutting-edge open-source tools in the given domain of EDA. Such tools can be used for benchmarking, comparison, and education. They provide a shared platform for experiments and can help simplify the development of new algorithms. Equally important for progress is access to real industrial-sized benchmarks.

For many years, the common base for research in logic synthesis has been SIS, a synthesis system developed by our research group at UC Berkeley in 1987-1991. Both SIS [35] and its predecessor MIS [8], pioneered multi-level combinational logic synthesis and became trend-setting prototypes for a large number of synthesis tools developed by industry.

In the domain of formal verification, a similar public system has been VIS [9], started at UC Berkeley around 1995 and continued at the University of Colorado, Boulder, and University of Texas, Austin. In particular, VIS features the latest algorithms for implicit state enumeration [15] with BDDs [11] using the CUDD package [36].

While SIS reached a plateau in its development in the middle 90's, logic representation and manipulation methods continued to be improved. In the early 2000s, And-Inverter Graphs (AIGs) emerged as a new efficient representation for problems arising in formal verification [23], largely due to the published work of A. Kuehlmann and his colleagues at IBM.

In that same period, our research group worked on a multi-valued logic synthesis system, MVSIS [13]. Aiming to find better ways to manipulate multi-valued relations, we experimented with new logic representations, such as AIGs, and found that, in addition to their use in formal verification, they can replace more-traditional representations in logic synthesis. As a result of our experiments with MVSIS, we developed a methodology for tackling problems, which are traditionally solved with SOPs [35] and BDDs [37], using a combination of random-guided simulation of AIGs and Boolean satisfiability (SAT) [25].

Majority-Inverter Graph: A Novel Data-Structure and Algorithms for Efficient Logic Optimization

Luca Amarù, Pierre-Emmanuel Gaillardon, Giovanni De Micheli
Integrated Systems Laboratory (LSI), EPFL, Switzerland.

Abstract—In this paper, we present Majority-Inverter Graph (MIG), a novel logic representation structure for efficient optimization of Boolean functions. An MIG is a directed acyclic graph consisting of three-input majority nodes and regular/complemented edges. We show that MIGs include any AND/OR/Inverter Graphs (AIGs), containing also the well-known AIGs. In order to support the natural manipulation of MIGs, we introduce a new Boolean algebra, based exclusively on majority and inverter operations, with a complete axiomatic system. Theoretical results show that it is possible to explore the entire MIG representation space by using only five primitive transformation rules. Such feature opens up a great opportunity for logic optimization and synthesis. We showcase the MIG potential by proposing a delay-oriented optimization technique. Experimental results over MCNC benchmarks show that MIG optimization reduces the number of logic levels by 18%, on average, with respect to AIG optimization performed by ABC academic tool. Employed in a traditional optimization-mapping circuit synthesis flow, MIG optimization enables an average reduction of (22%, 14%, 11%) in the estimated [delay, area, power] metrics, before physical design, as compared to academic/commercial synthesis flows.

Categories and Subject Descriptors
B.6.3 [Design Aids]: Automatic Synthesis, Optimization

General Terms
Algorithms, Design, Performance, Theory.

Keywords

Majority Logic, Boolean Algebra, DAG, Logic Synthesis.

I. INTRODUCTION

The performance of today's digital integrated circuits largely depends on the capabilities of logic synthesis tools. In this context, efficient representation and optimization of Boolean functions are key features. Some data structures and algorithms have been proposed for these tasks [1]–[8]. Most of them consider, as basis operations, inversion (INV), conjunction (AND), disjunction (OR) [2]–[5] and if-then-else (MUX) [6], [7]. Other Boolean operations are derived by composition. Even though existing design automation tools, based on original optimization techniques [1]–[8], produce good results and handle large circuits, the possibility to push further the efficacy of logic synthesis continues to be of paramount interest to the Electronic Design Automation (EDA) community. With this aim in mind, we approach the logic optimization problem from a new angle.

In this paper, we propose a novel methodology to represent and optimize logic, by using only majority (MAJ) and inversion (INV) as basis operations. We present the Majority-Inverter Graph (MIG), a logic representation structure consisting of three-input majority nodes and regular/complemented edges. MIGs include any AND/OR/Inverter Graphs (AIGs), therefore containing also AIGs [8]. To provide native manipulation of MIGs, we introduce a novel Boolean algebra, based exclusively on majority and inverter operations. A set of five primitive transformations forms a complete axiomatic system. Using

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

DAC '04, June 1-5 2004, San Francisco, CA, USA.
Copyright 2004 ACM 978-1-4503-2730-5/04/06\$15.00.
http://dx.doi.org/10.1145/259069.2591158

a sequence of such primitive axioms, it is possible to explore the entire MIG representation space. This remarkable property opens up great opportunities in logic optimization and synthesis. We showcase the potential of MIGs by proposing a delay-oriented optimization technique. Experimental results, over the MCNC benchmark suite, show that MIG optimization decreases the number of logic levels by 18%, on average, with respect to AIG optimization run by ABC academic tool. Applied in a standard optimization-mapping circuit synthesis flow, MIG optimization enables a reduction in the estimated [delay, area, power] metrics of (22%, 14%, 11%), on average before physical design, as compared to academic/commercial synthesis flows.

The study of majority-inverter logic synthesis is also motivated by the design of circuits in emerging technologies. In the quest for increasing computational performance per unit area [9], majority/minority gates are natively implemented in different nanotechnologies [10]–[12] and also extend the functionality of traditional NAND/NOR gates. In this scenario, MIGs and their algebra represent the natural methodology to synthesize majority logic circuits in emerging technologies. In this paper, we focus on standard CMOS, to first showcase the interest of MIGs in an ordinary design flow.

The remainder of this paper is organized as follows. Section II provides a background on logic representation and optimization. Section III presents MIGs and their new associated Boolean algebra. Section IV describes the optimization of MIGs using primitive transformation rules. Section V validates, through experimental results, MIG-based optimization and also presents and compares synthesis results to state-of-art academic/commercial tools. Section VI concludes the paper.

II. BACKGROUND AND MOTIVATION

This section presents relevant background on logic representations and optimization for logic synthesis. Notations and definitions for Boolean algebra and logic networks are also introduced.

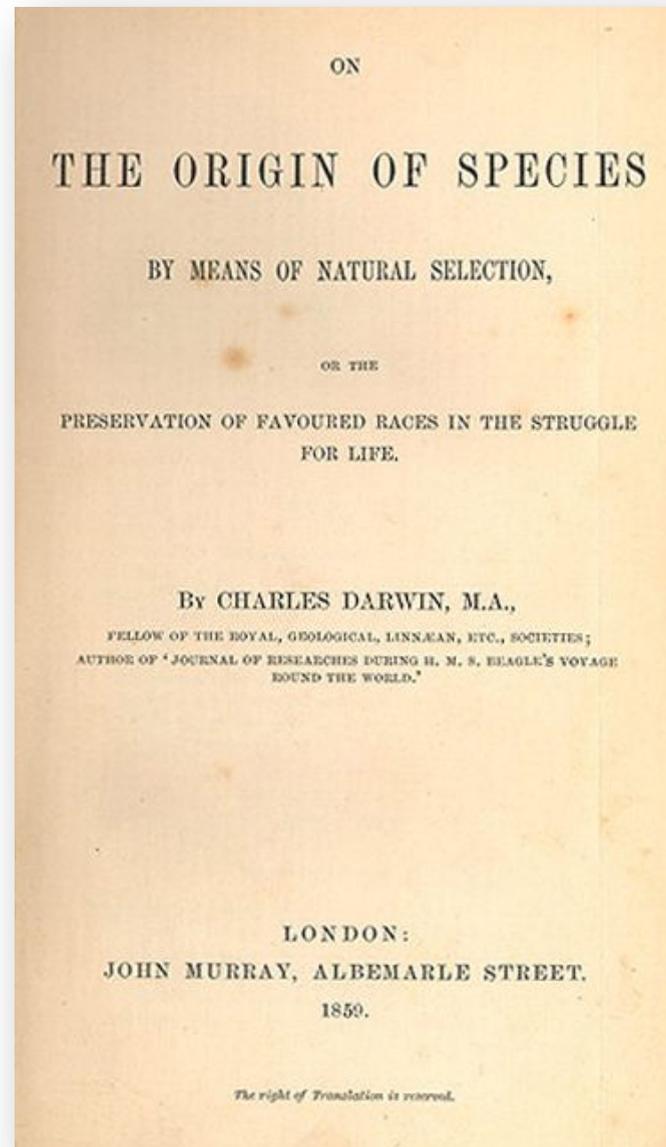
A. Logic Representation and Optimization

Virtually, all digital integrated circuits are synthesized thanks to efficient logic representation forms and associated optimization algorithms [1]. Early data structures and related optimization algorithms [2] are based on two-level representation of Boolean functions in Sum Of Product (SOP) form, which is a disjunction (OR) of conjunctions (AND) where variables can be complemented (INV). Another pioneering data structure is the Binary Decision Diagram (BDD) [6]: a canonical representation form based on nested if-then-else (MUX) formulas. Later on, multi-level logic networks [3], [4] emerged, employing AND, OR, INV, MUX operations as basis functions, with more scalable optimization and synthesis tools [4], [7]. To deal with the continuous increase in logic designs complexity, a step further is enabled by [5], where multi-level logic networks are made homogeneous, i.e., consisting of only AND nodes interconnected by regular/complement (INV) edges. The tool ABC [8], which is based on the AND-Inverter Graphs (AIGs), is considered the state-of-art academic software for (large) optimization and synthesis.

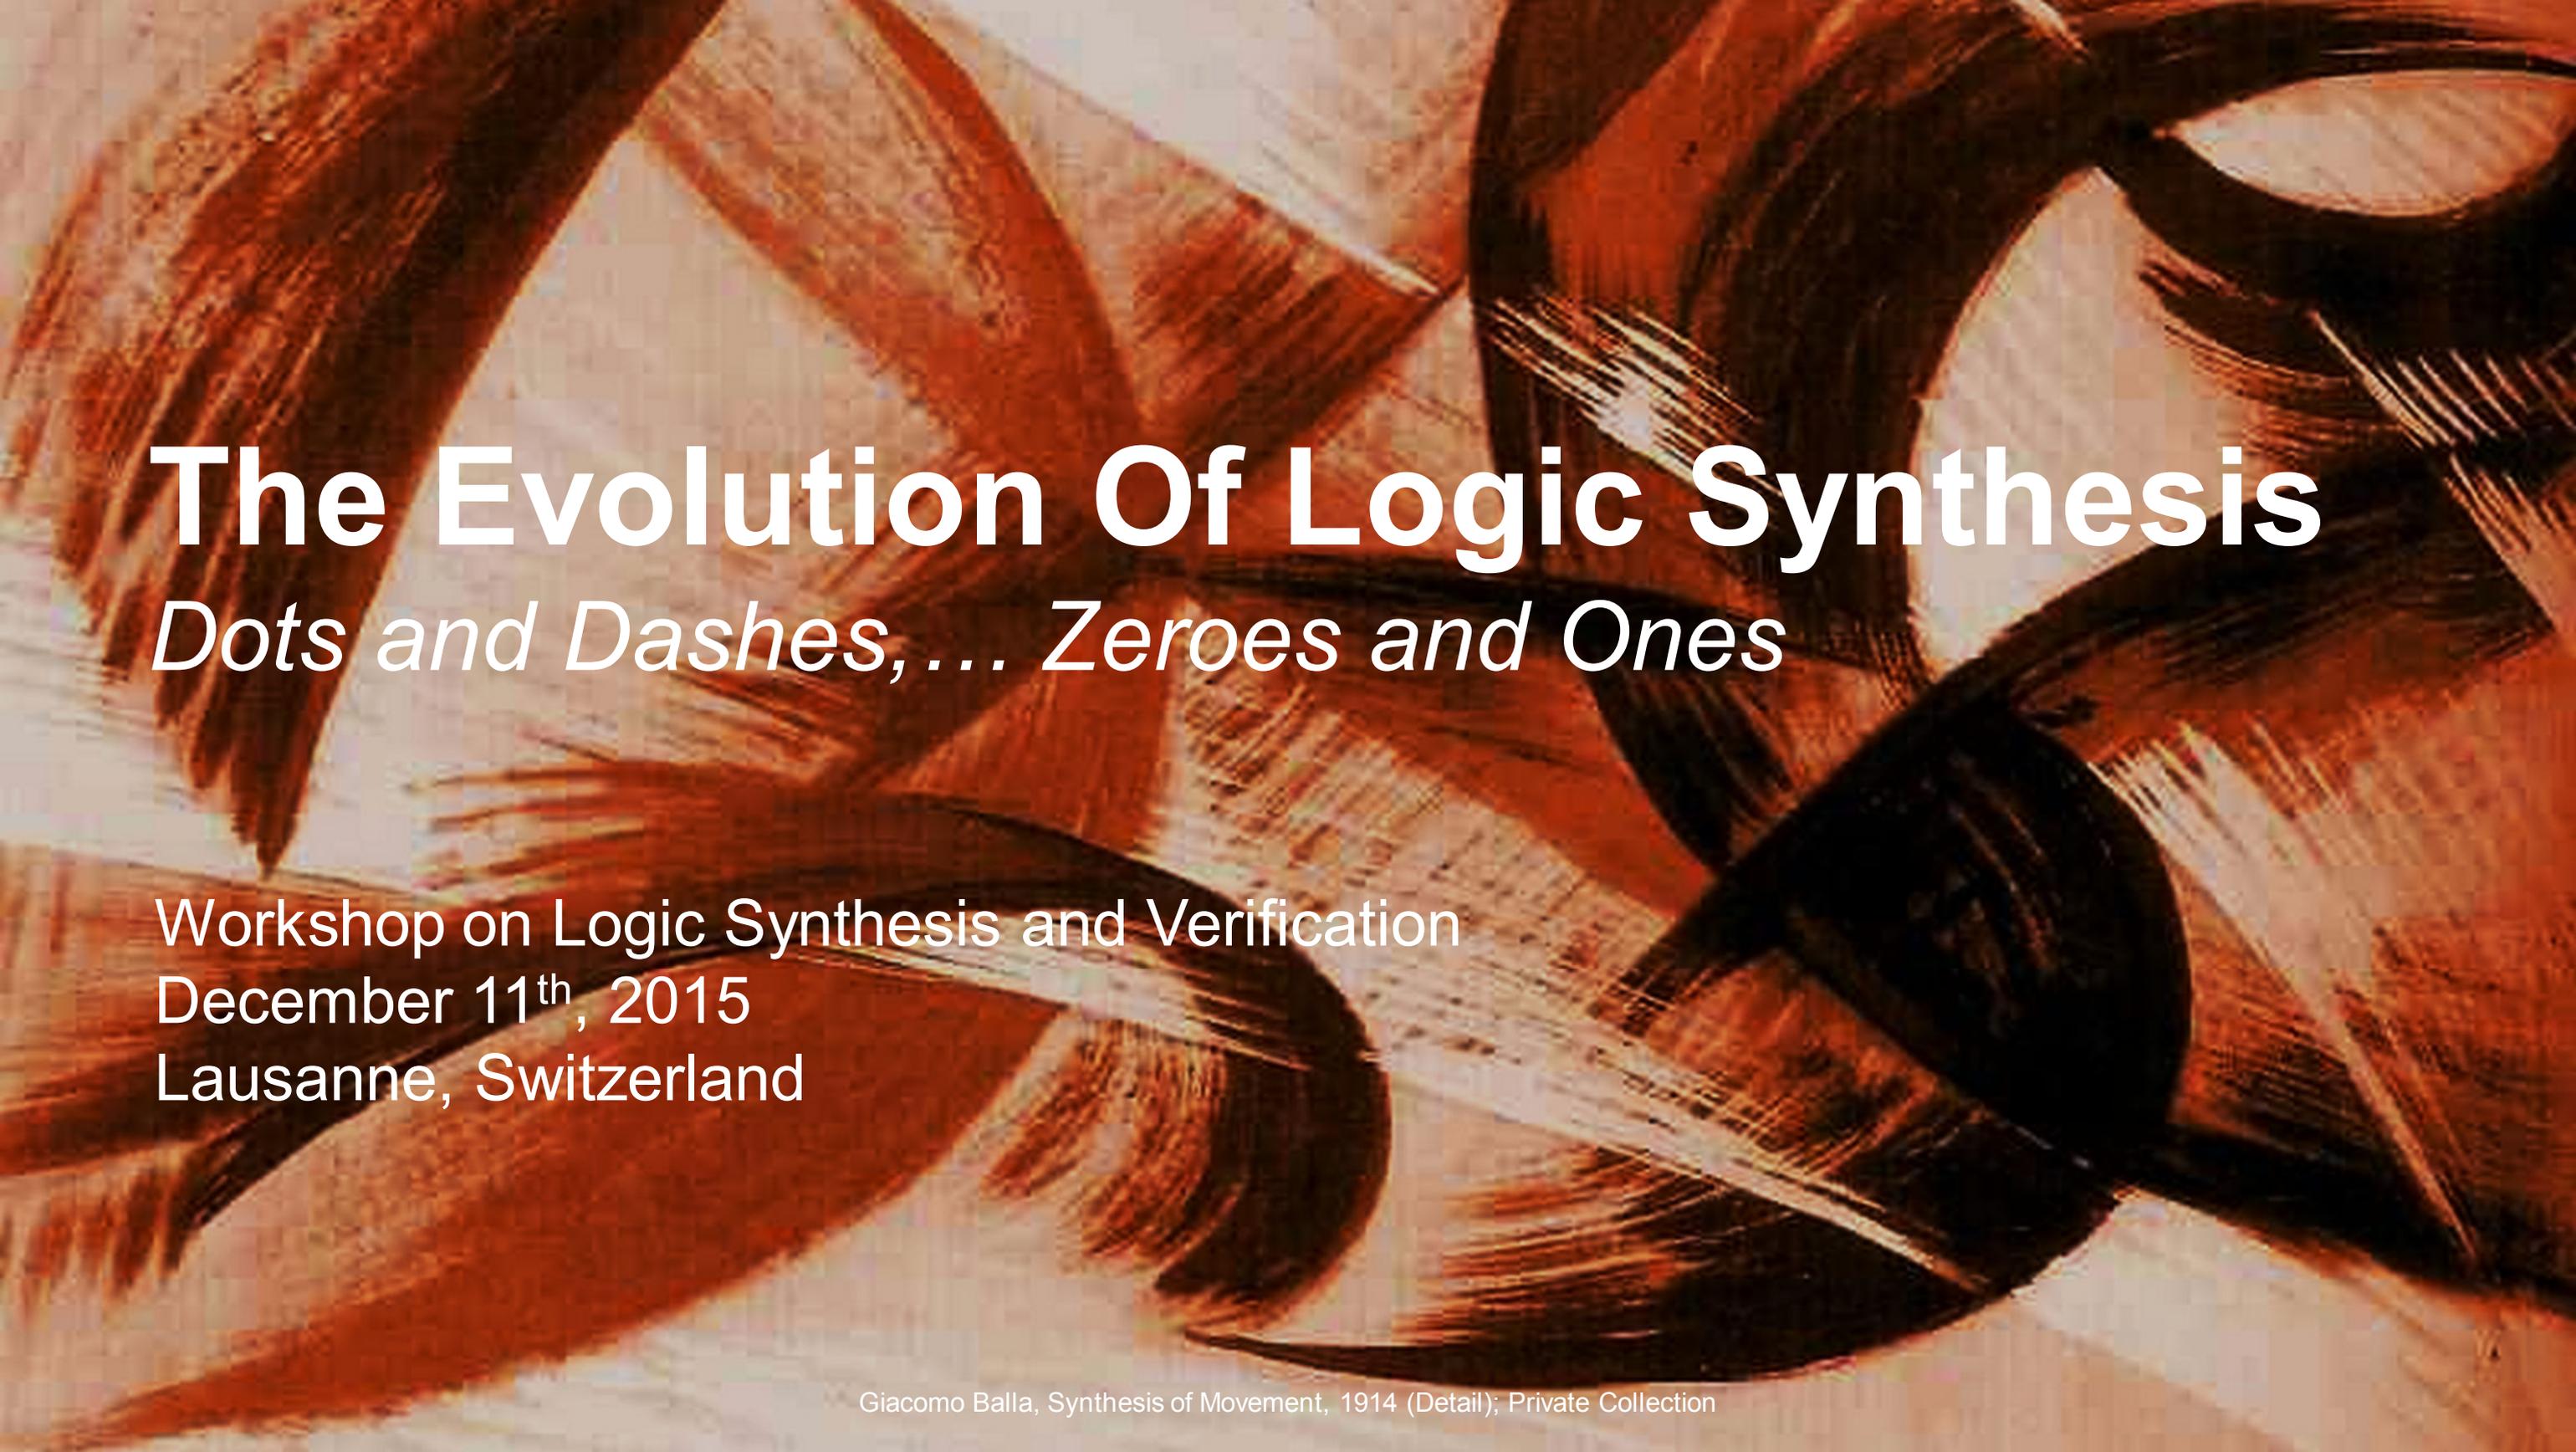
We propose, in this paper, a new logic optimization paradigm that aims at extending the capabilities of modern synthesis tools.

Back To Charles R. Darwin

Things Don't Happen Overnight !



“It may be said that natural selection is daily and hourly scrutinizing every variation, even the slightest; rejecting that which is bad, preserving and adding up all that is good ... silently and insensibly working [...] at the improvement of each [...]. We see nothing of these slow changes in progress, until the hand of time has marked the long lapses of ages, and then [...] we only see that the forms of life are now different from what they formerly were.”



The Evolution Of Logic Synthesis

Dots and Dashes, ... Zeroes and Ones

Workshop on Logic Synthesis and Verification

December 11th, 2015

Lausanne, Switzerland