EDA 3.0 Time to refactor Logic Synthesis

Leon Stok, VP EDA, IBM

[IWLS1993]

Tuning of Logic Synthesis Scenarios

Lukas P.P.P. van Ginneken Andreas Kuehlmann

IBM T.J. Watson Research Center
P.O. Box 218, Yorktown Heights, NY 10598

The manual tuning of a logic synthesis scenario, is a time consuming activity. Logic synthesis scenarios are dependent on many factors, such as the type of logic, design style, the set of transforms and the optimization goals. An automated method for the tuning of synthesis scenarios is presented. It is shown that using this method a manually designed standard scenario could be improved by an average of 8.9%. We also show that scenarios tuned to particular designs could improve the size by an additional average of 4.6%.

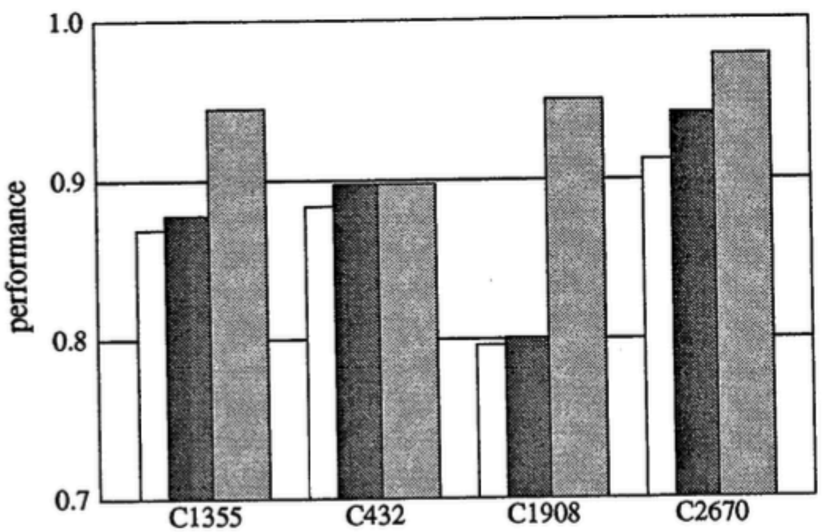

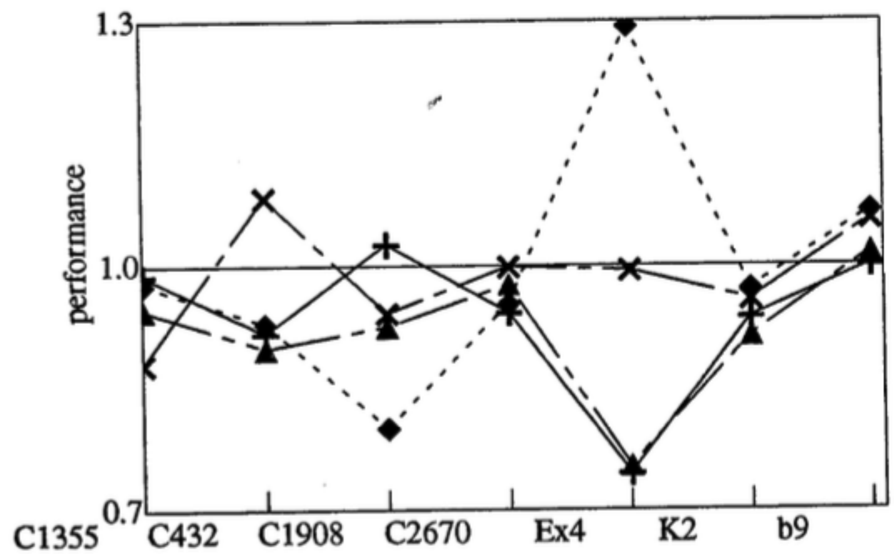

SynTuneSys

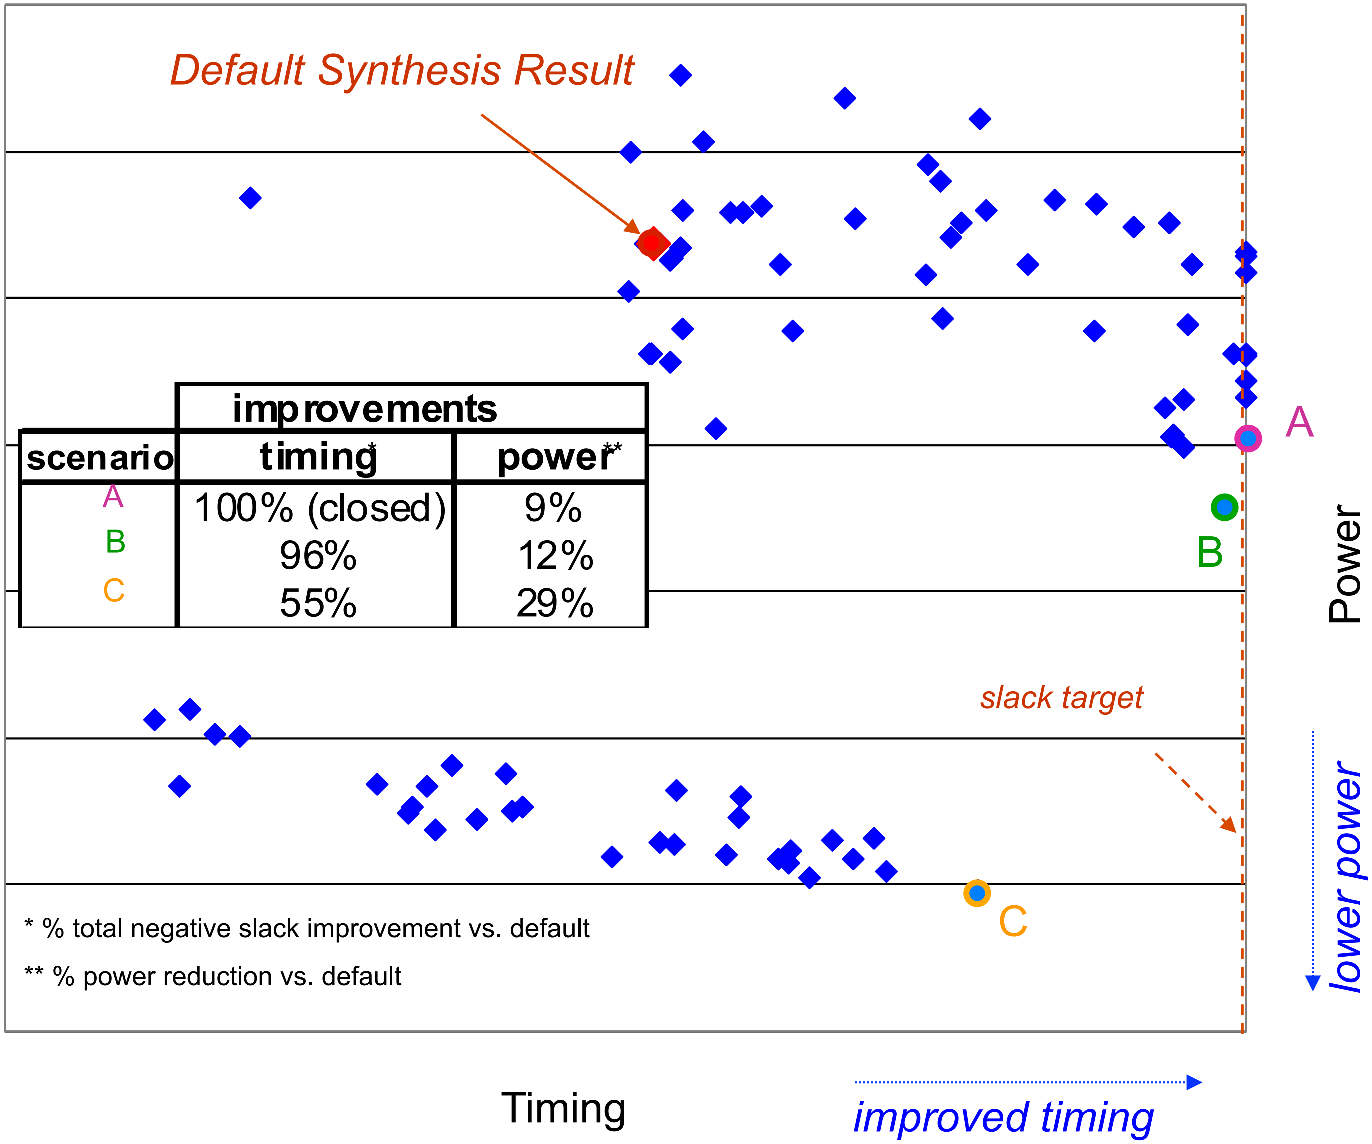

	improvements	
scenario	timing*	power**
A	100% (closed)	9%
B	96%	12%
C	55%	29%

* % total negative slack improvement vs. default

** % power reduction vs. default

Design Flow

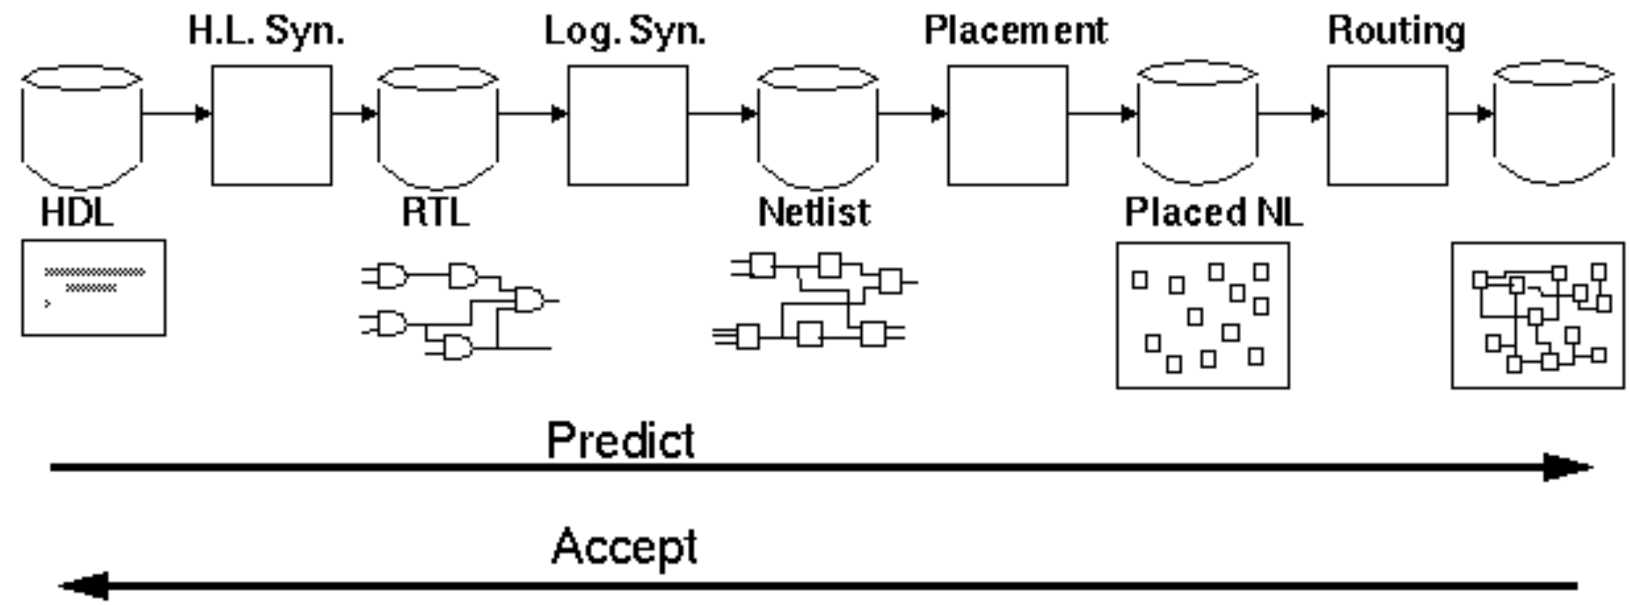

[Hathaway, Integrated, Incremental, and Modular EDA Tools,1996]

Prediction is still very difficult during synthesis

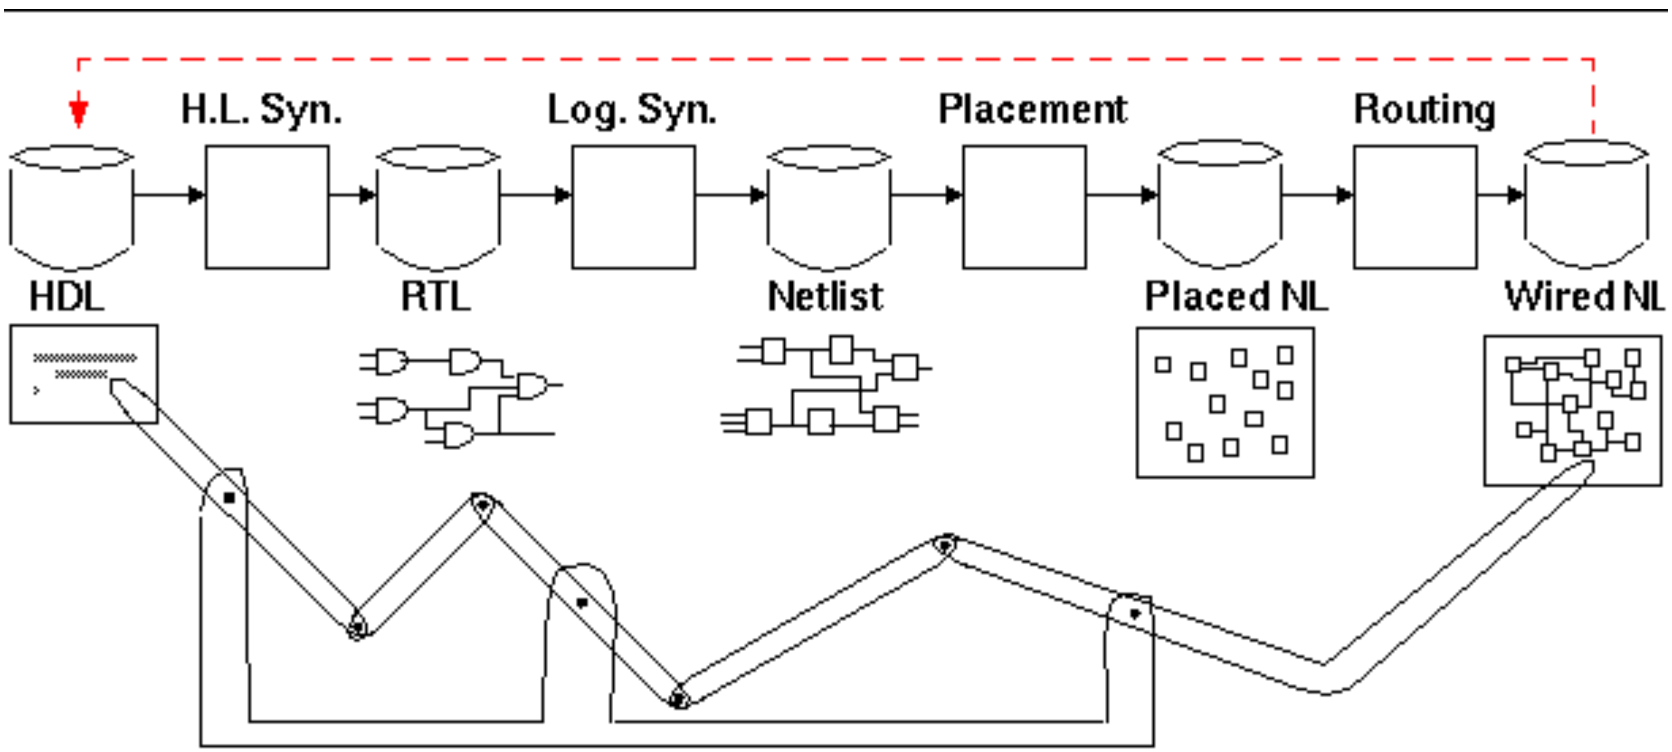

[Hathaway, Integrated, Incremental, and Modular EDA Tools,1996]

Question

- Why does it take about 2 seconds to look up the shortest path from my home to JFK on my phone, and does it take a couple of hours to find the critical path in a large chip in an advanced technology on a sizeable server?

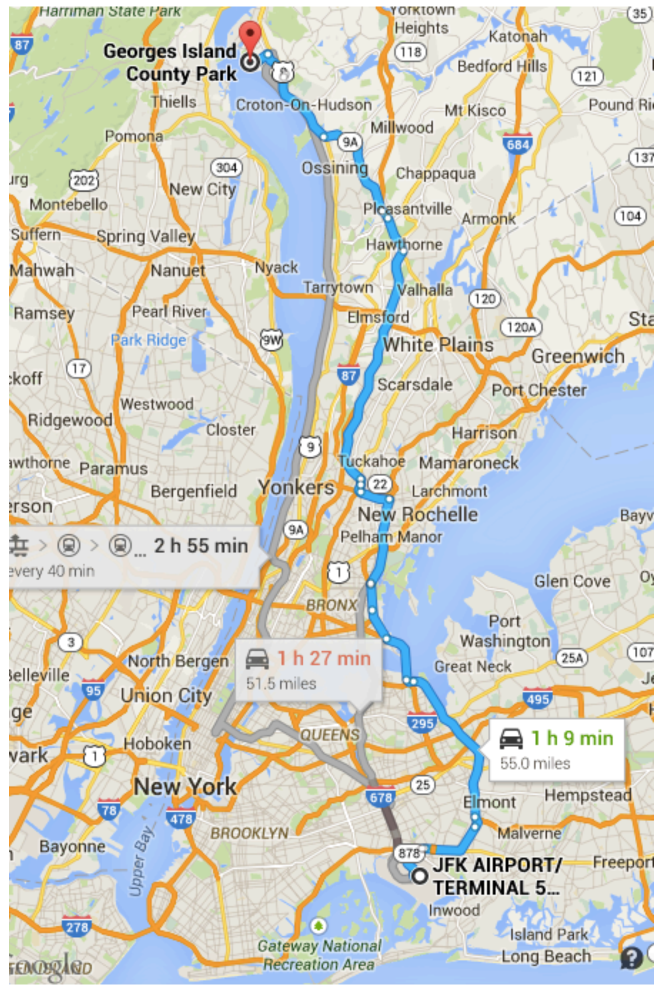

©LS 2015

11 fascinating facts about Google Maps

- **1. How much data has Google Maps accumulated**
  - Combining satellite, aerial and street level imagery, Google Maps has over 20 petabytes of data, which is equal to approximately 21 million gigabytes, or around 20,500 terabytes
- **2. How often are the images updated?**
  - Depending on data availability, aerial and satellite images are updated every two weeks. Street View images are updated as quickly as possible, though Google wasn't able to offer specific schedules, due to its dependence on factors such as weather, driving conditions, etc.
- **8. In the history of Google Maps, how many Street View images have been taken?**
  - The Street View team has taken tens of millions of images since the Street View project began in 2007, and they've driven more than 5 million unique miles of road.

[http://mashable.com/2012/08/22/google-maps-facts/]

Design Data Volume

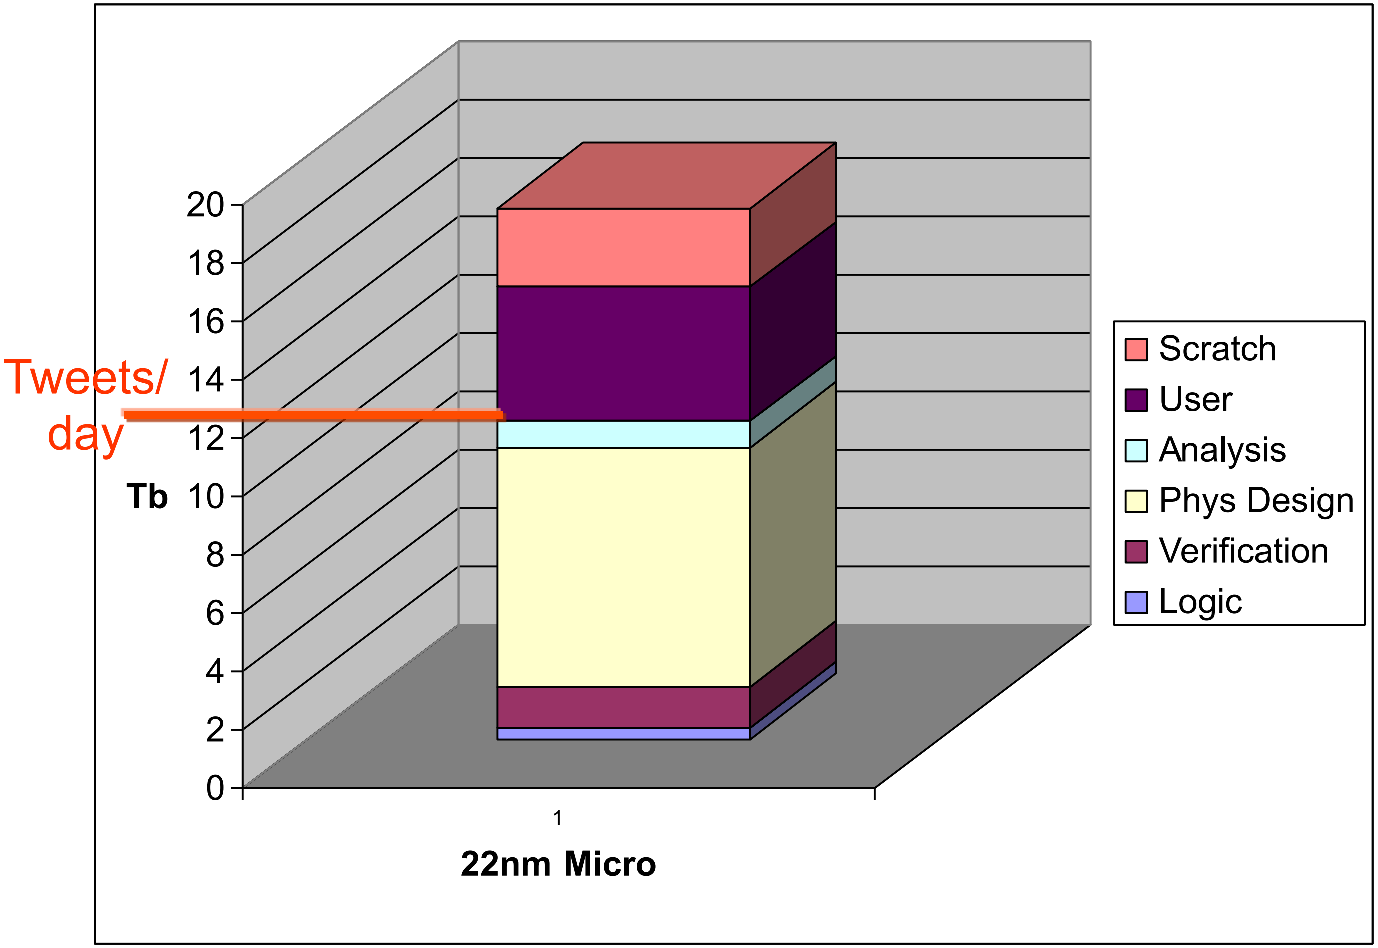

- 3Gb Oasis
- 1 Tb product engineering
- 5Tb test and diagnostics data
- 5km wire

EDA Evolution

- EDA1.0
  - Point tools on individual Workstations
- EDA 2.0
  - Integration of design tools on distributed Servers
- EDA 3.0
  - Integrated Design Flow on very large Clusters
- What will EDA 3.0 look like?

What will EDA3.0 look like?

- **Designer**, I want to
  - get to my DATA.. from anywhere any place
  - the DATA be there without me waiting for it
  - analyze the DATA with whatever tools I can lay my hands on
  - know how to improve my design [data]
  - know how to get from A to B through my design process
  - be like……..

- Brayton / Cong 2009 NSF workshop, EDA Past, Present, Future)
  - **Intuitive, simplified and standardized design environments.**
  - **Scalable design methodologies**
  - **Predictable design flows**

EDA Evolution

- EDA1.0
  - Point tools on individual Workstations
- EDA 2.0
  - Integration of design tools on distributed Servers
- EDA 3.0
  - Integrated Design Flow on very large Clusters
- What will EDA 3.0 look like?

- How do we get there?
  - Learn from analytics
  - Learn from Big Data: other data and graph parallel systems.
  - Capitalize on the Changing nature of IT so we will be able to run this effectively.
  - Change the way we interact with Design Data
  - Change our algorithms to take advantage of these changes

YARN

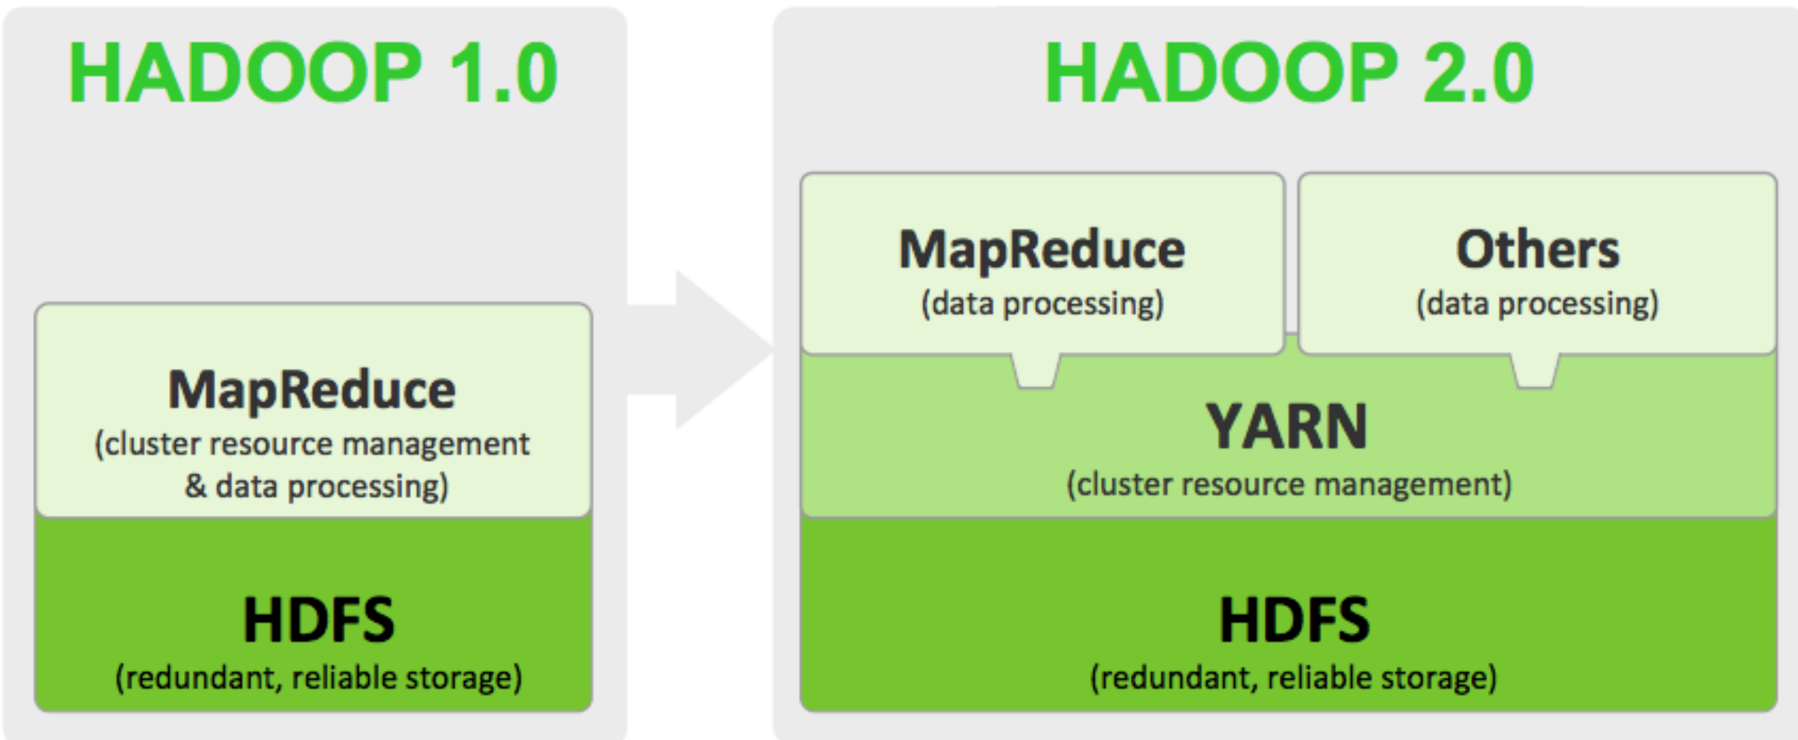

[tomsitpro.com]

- YARN: Yet Another Resource Negotiator
- Resource management and job tracking/scheduling got split up.

Scaling Apache Giraph to a trillion edges

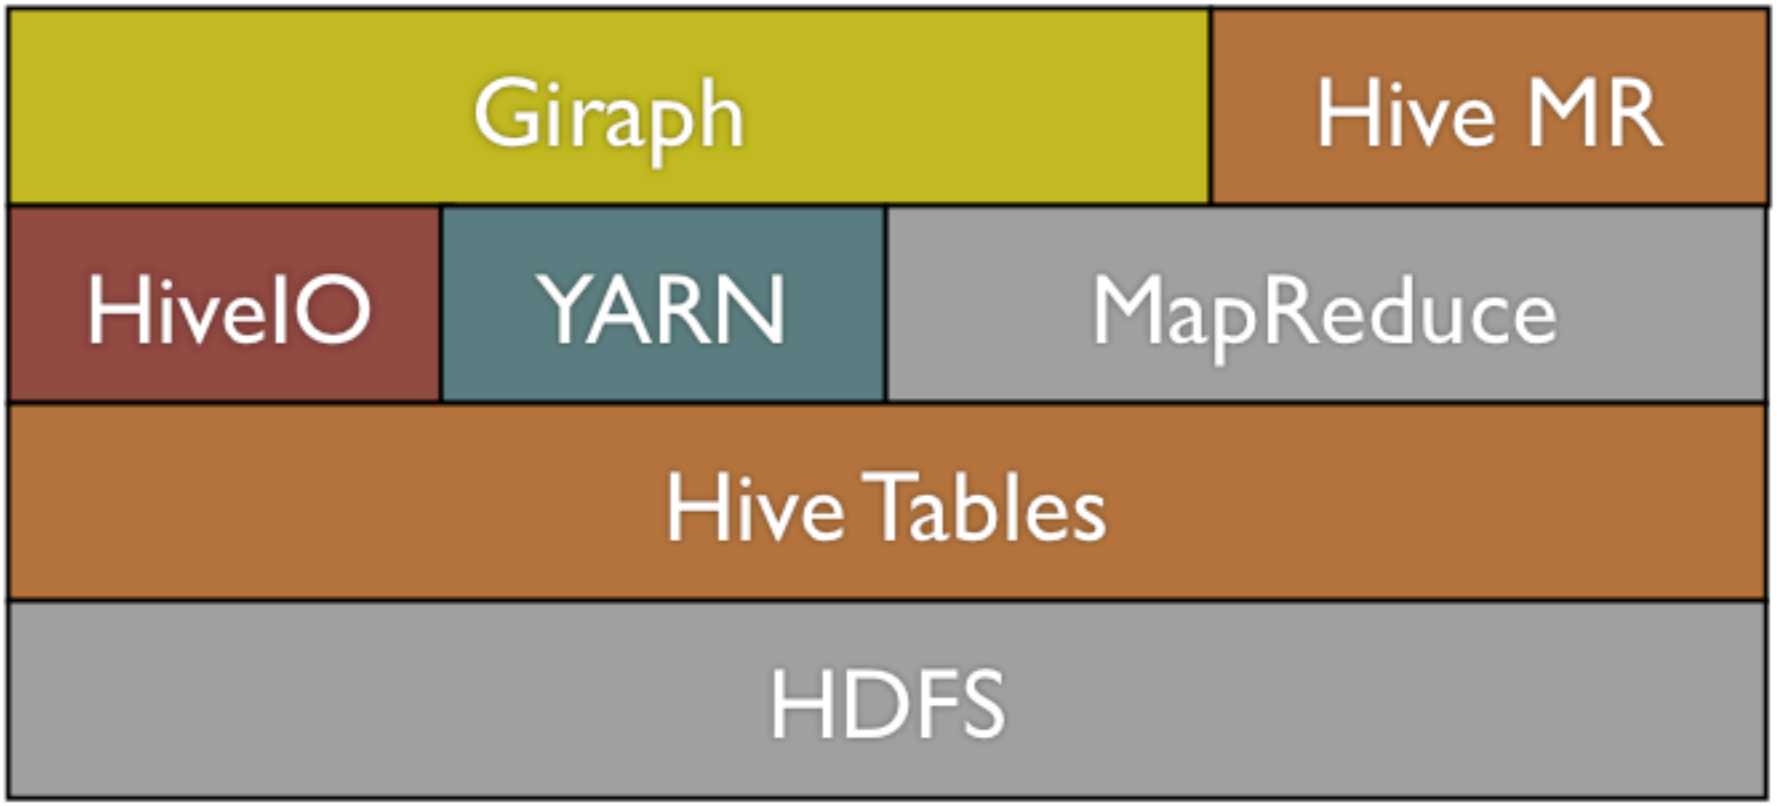

Four minutes
One Iteration of PageRank

https://www.facebook.com/notes/facebook-engineering/scaling-apache-giraph-to-a-trillion-edges/10151617006153920

Spark

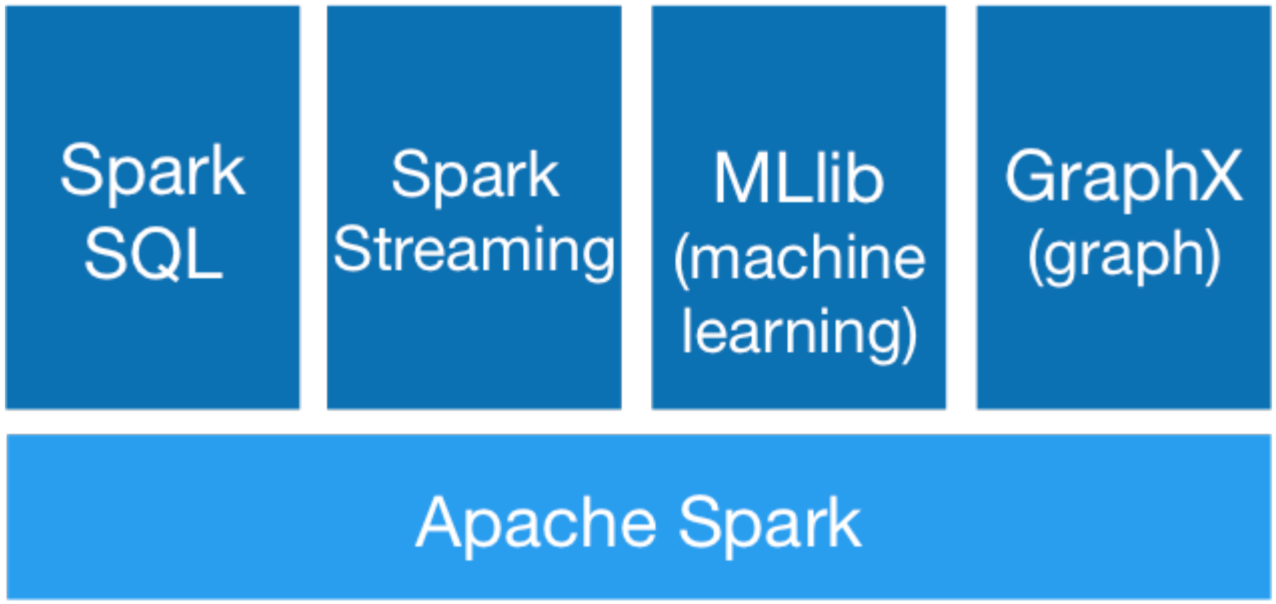

- RDDs: Resilient Distributed Datasets
- Transformations and Actions
- Lazy evaluation of transformations
  - Lineage graph

Spark Advantages

- Tight integration gives the ability
  - to build applications that seamlessly combine different processing models.
    - in Spark you can write one application that uses machine learning to classify data in real time as it is ingested from streaming sources.
    - Simultaneously, analysts can query the **resulting data**, also in real time, via SQL (e.g., to join the data with unstructured logfiles).
    - More sophisticated data engineers and data scientists can access the **same** data via the Python shell for ad hoc analysis.
    - Others might access the data in standalone batch applications.
- Second, the costs associated with running the stack are minimized, because instead of running 5–10 independent software systems, an organization needs to run only one. These costs include deployment, maintenance, testing, support, and others.
- Each time a new component is added to the Spark stack, every organization that uses Spark will immediately be able to try this new component.
- When Spark's core engine adds an optimization, SQL and machine learning libraries automatically speed up as well.
- This changes the cost of trying out a new type of data analysis from downloading, deploying, and learning a new software project to upgrading Spark.
- All the while, the IT team has to maintain only one system.

[Learning Spark, by Holden Karau, Andy Konwinski, Patrick Wendell, and Matei Zaharia]

Spark and Graphx

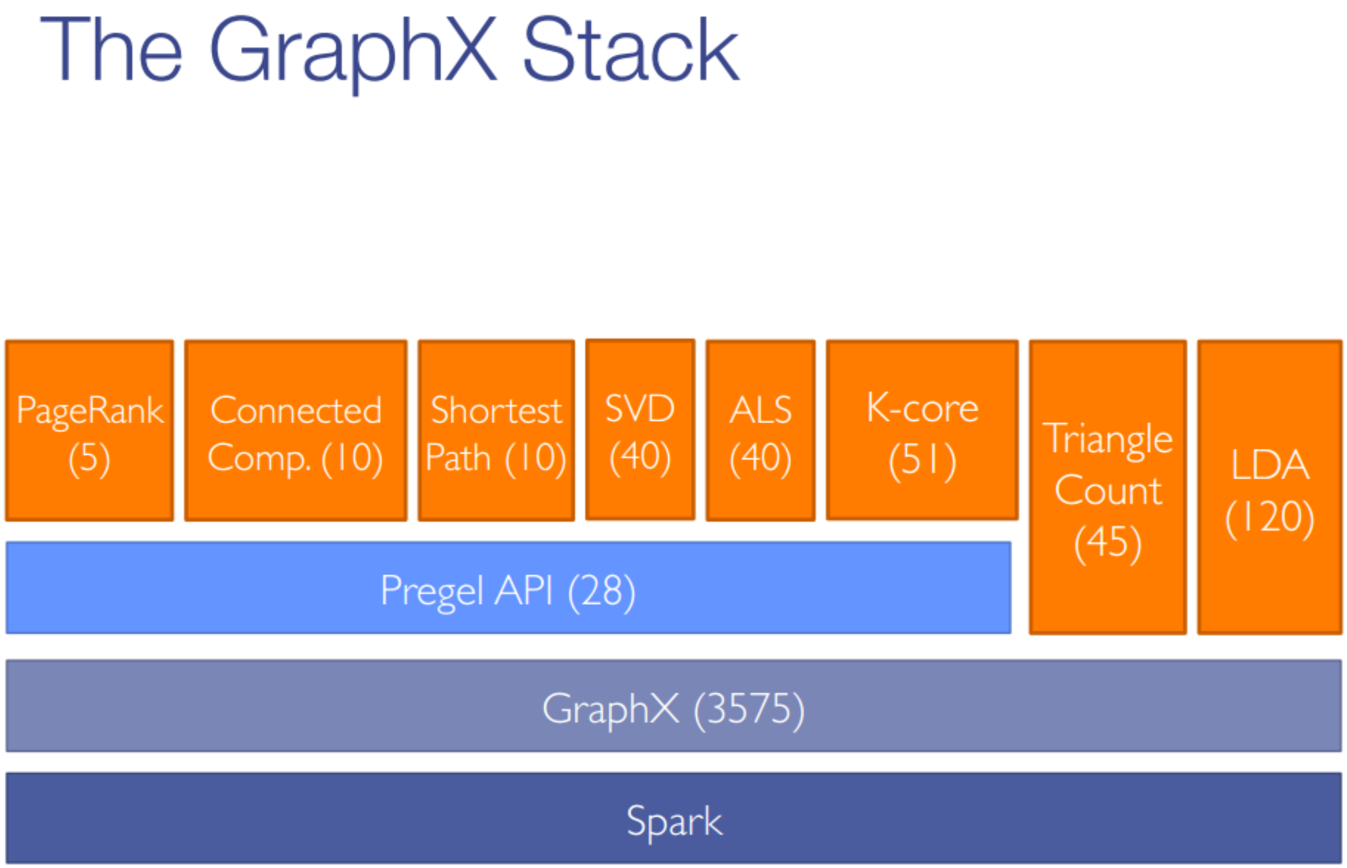

Quadratic Programing Solver for Non-negative Matrix Factorization with Spark
[http://spark-summit.org/2014/talk/quadratic-programing-solver-for-non-negative-matrix-factorization-with-spark]

DATABRICKS

Graph Databases

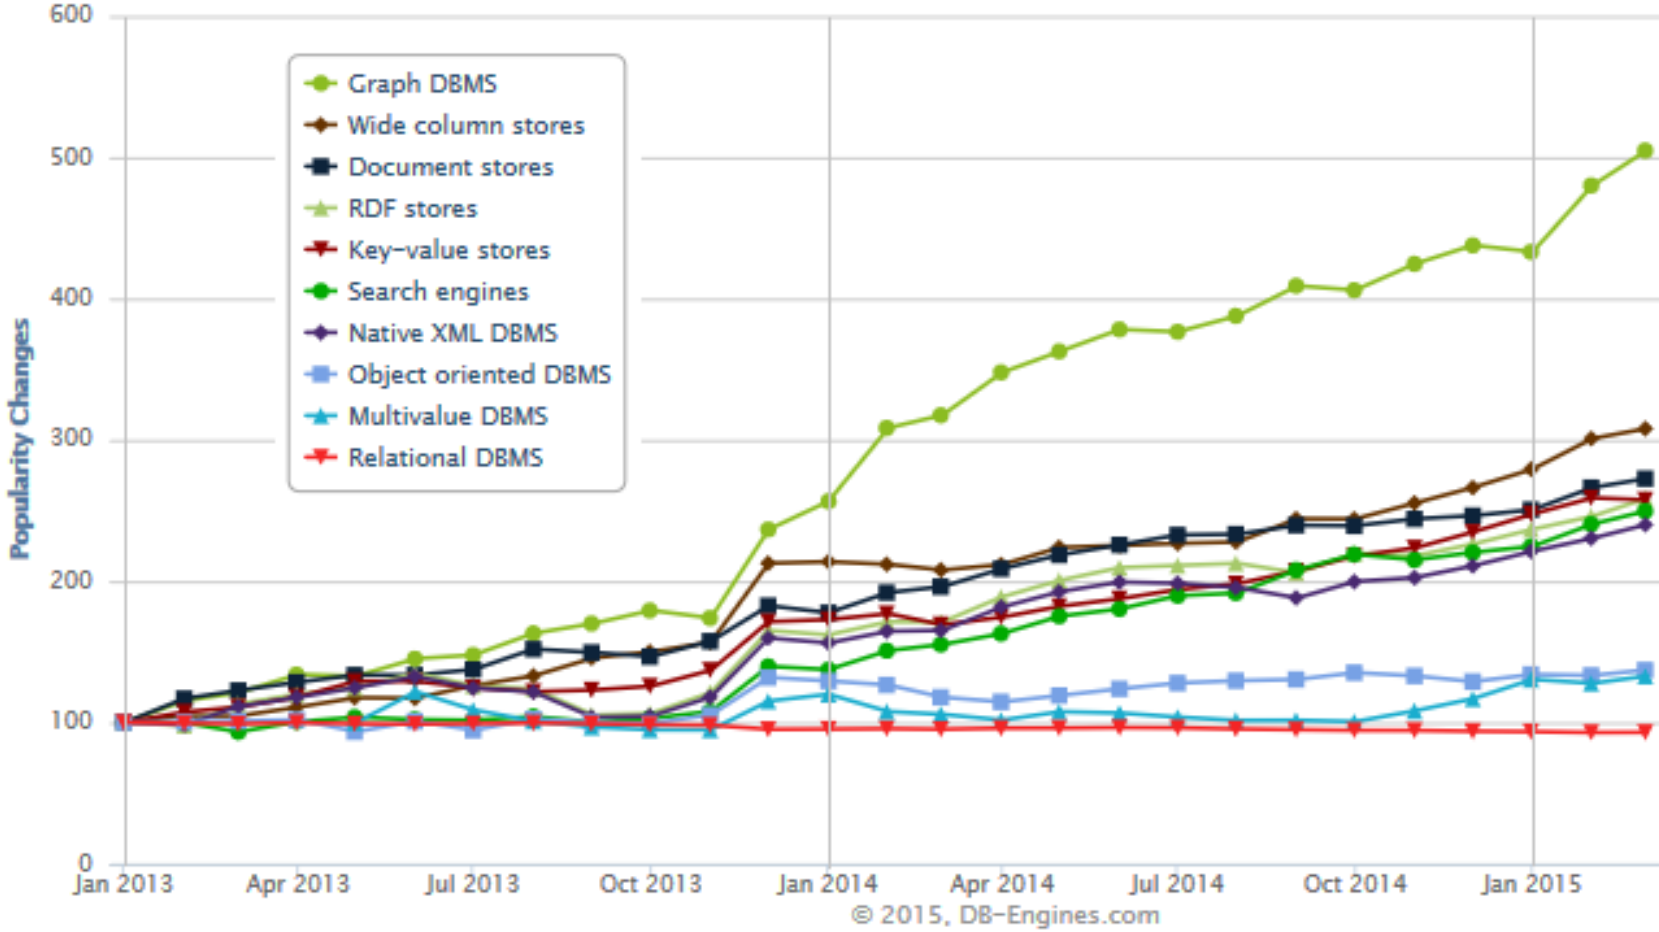

Wavefront Technology Mapping, Structural Bias

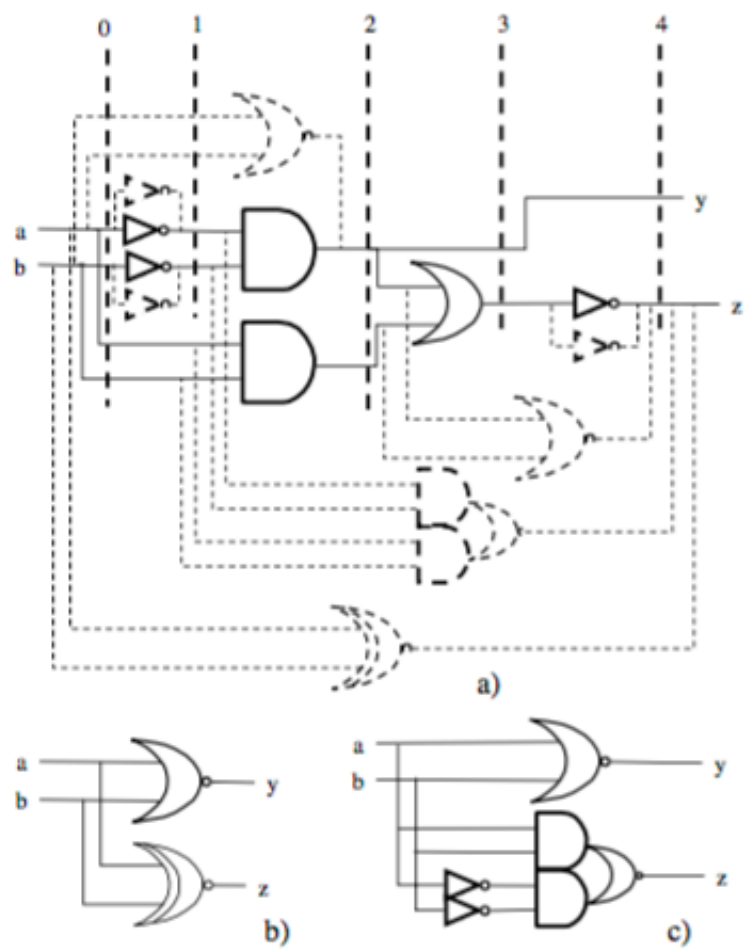

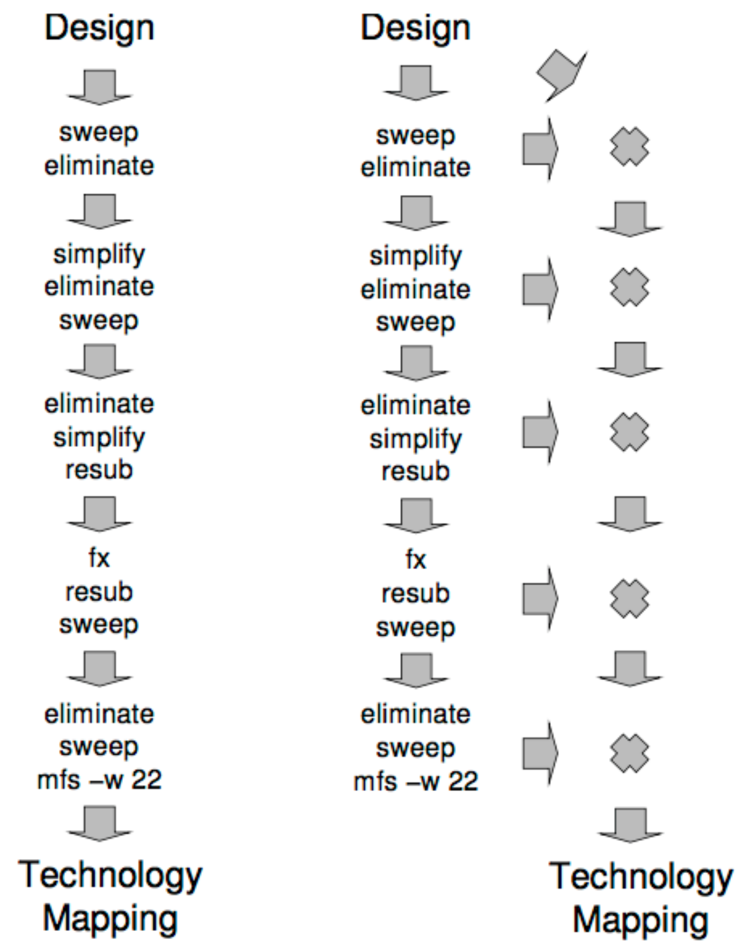

(a) Conventional Synthesis (b) Lossless Synthesis

[**Wavefront Technology Mapping ,** L. Stok, M. Iyer, A.J. Sullivan, 1998 IWLS]

Explore

Structural Bias in Technology Mapping , S. Chatterjee A. Mishchenko R. Brayton, X. Wang T. Kam, 2005 TCAD]

EDA 3.0 Time to refactor Logic Synthesis

- Let us harvest the enormous power of distributed analytics and optimization.
- Let us Create Fantastic new designs by harvesting the full power of Synthesis.
- Let us drive to Phenomenal design flow overall TAT improvements and efficiencies.