EE273 Lecture 11
Clock Distribution

February 21, 2001

William J. Dally
Computer Systems Laboratory
Stanford University
billd@csl.stanford.edu

Logistics

- Project feedback (strongly recommended)
  - Thursday 9:00AM to 11:00AM – 15min appointments
  - Make an appointment to meet with the course staff to get feedback on your project ideas and direction
- Problem set 5 due today
- Problem set 6 (the last one)
  - Problems 10-1 and 10-2
- Reading
  - Sections 10.1 and 10.2
  - Complete before class on Monday 2/26

How are people doing on their projects?

A Quick Overview

- The clock-distribution problem
  - want lots (10^5 - 10^8) of clock loads to see the clock rise at precisely the same picosecond
  - noise and variations in delay of distribution network make this difficult
- Off-chip Clock Distribution
  - wires are good LC transmission lines
  - typically use a tree of clock buffers
- On-chip Clock Distribution
  - wires are slow RC transmission lines
  - buffer delay is modulated by power-supply noise
  - still use a clock tree, but concerns are different

The Clock Distribution Problem

- Some systems have large synchronous clock domains
  - 10s - 100s of chips
  - 10^3 - 10^5 clock loads per chip
  - need to distribute the clock to within 10% of t_{ck}
    - 200ps for a 500MHz clock
- Two step process
  - get the clock to each chip with low skew
  - distribute the clock on each chip with low skew
- Other issues
  - may not want a single point of failure
    - can't have just one clock oscillator or master buffer
  - need to adjust for manufacturing variations with a minimum of labor

Off-Chip Clock Distribution

- Relatively easy because wire delays are
  - stable
  - easily characterized
- Most common strategy is a clock tree
  - single oscillator
  - buffer tree with a fanout of N
  - low-skew fanout buffers used
    - e.g., Motorola 100LVE111
  - tuning needed to get very low skews
    - self-tuning series terminated drivers
- For fault tolerance can use an array of oscillators
  - phase lock multiple clock generators to one another
- Traveling waves and standing waves can also be exploited
  - round-trip clock distribution
  - salphasic clock distribution

A Typical Clock-Tree Driver
The 100LVE111

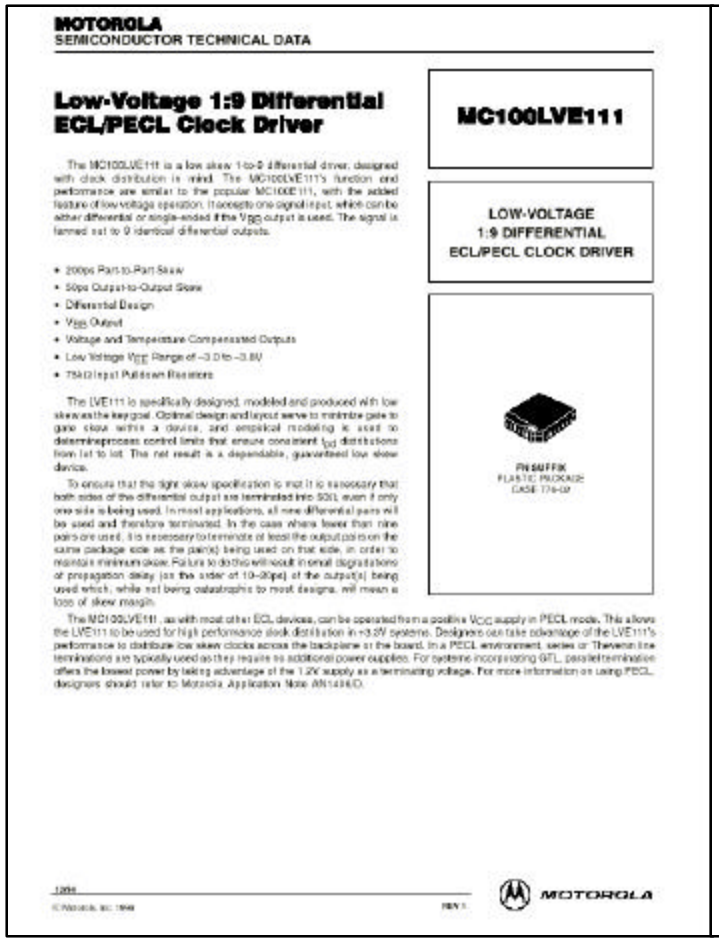

MOTOROLA
SEMICONDUCTOR TECHNICAL DATA

Low-Voltage 1:9 Differential ECL/PECL Clock Driver

MC100LVE111

LOW-VOLTAGE
1:9 DIFFERENTIAL
ECL/PECL CLOCK DRIVER

The MC100LVE111 is a low skew 1-to-9 differential driver, designed with clock distribution in mind. The MC100LVE111's function and performance are similar to the popular MC100E111, with the added feature of low voltage operation. It accepts one signal input, which can be either differential or single-ended if the V_{BB} output is used. The signal is fanned out to 9 identical differential outputs.

- 200ps Part-to-Part Skew
- 50ps Output-to-Output Skew
- Differential Design
- V_{BB} Output
- Voltage and Temperature Compensated Outputs
- Low Voltage V_{EE} Range of −3.0 to −3.8V
- 75kΩ Input Pulldown Resistors

The LVE111 is specifically designed, modeled and produced with low skew as the key goal. Optimal design and layout serve to minimize gate to gate skew within a device, and empirical modeling is used to determine process control limits that ensure consistent t_{pd} distributions from lot to lot. The net result is a dependable, guaranteed low skew device.

To ensure that the tight skew specification is met it is necessary that both sides of the differential output are terminated into 50Ω, even if only one side is being used. In most applications, all nine differential pairs will be used and therefore terminated. In the case where fewer than nine pairs are used, it is necessary to terminate at least the output pairs on the same package side as the pair(s) being used on that side, in order to maintain minimum skew. Failure to do this will result in small degradations of propagation delay (on the order of 10–20ps) of the output(s) being used which, while not being catastrophic to most designs, will mean a loss of skew margin.

The MC100LVE111, as with most other ECL devices, can be operated from a positive V_{CC} supply in PECL mode. This allows the LVE111 to be used for high performance clock distribution in +3.3V systems. Designers can take advantage of the LVE111's performance to distribute low skew clocks across the backplane or the board. In a PECL environment, series or Thevenin line terminations are typically used as they require no additional power supplies. For systems incorporating GTL, parallel termination offers the lowest power by taking advantage of the 1.2V supply as a terminating voltage. For more information on using PECL, designers should refer to Motorola Application Note AN1406/D.

FN SUFFIX
PLASTIC PACKAGE
CASE 776-02

1/99
© Motorola, Inc. 1999
REV 1

MOTOROLA

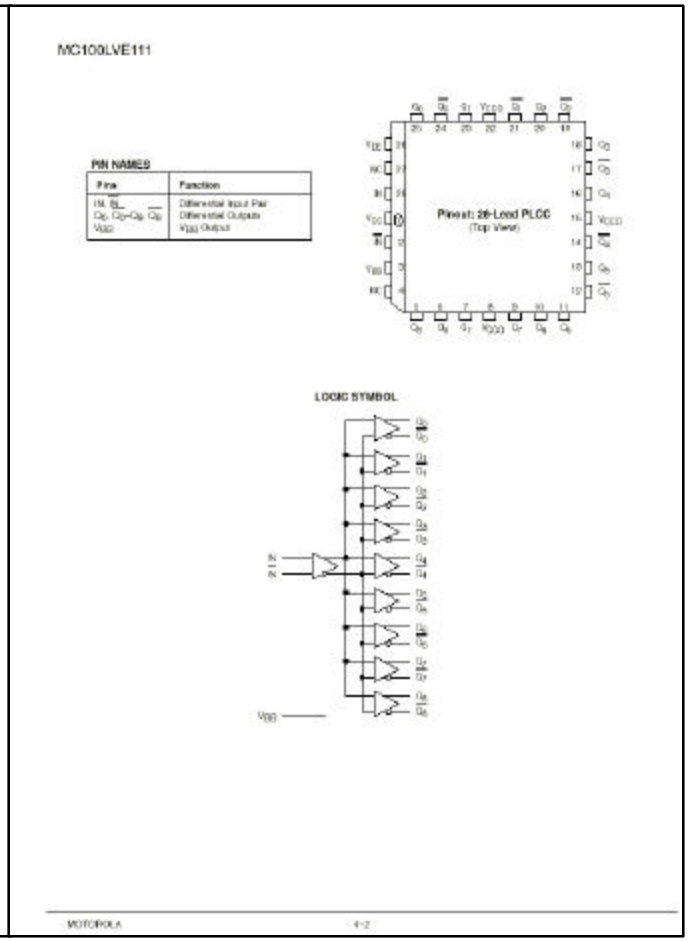

MC100LVE111

PIN NAMES

Pins	Function
IN, $\overline{IN}$	Differential Input Pair
Q_0, $\overline{Q_0}$–Q_8, $\overline{Q_8}$	Differential Outputs
V_{BB}	V_{BB} Output

Pinout: 28-Lead PLCC
(Top View)

LOGIC SYMBOL

MOTOROLA 4-2

Skew:
50ps within a part
200ps across parts

An Example Clock Tree

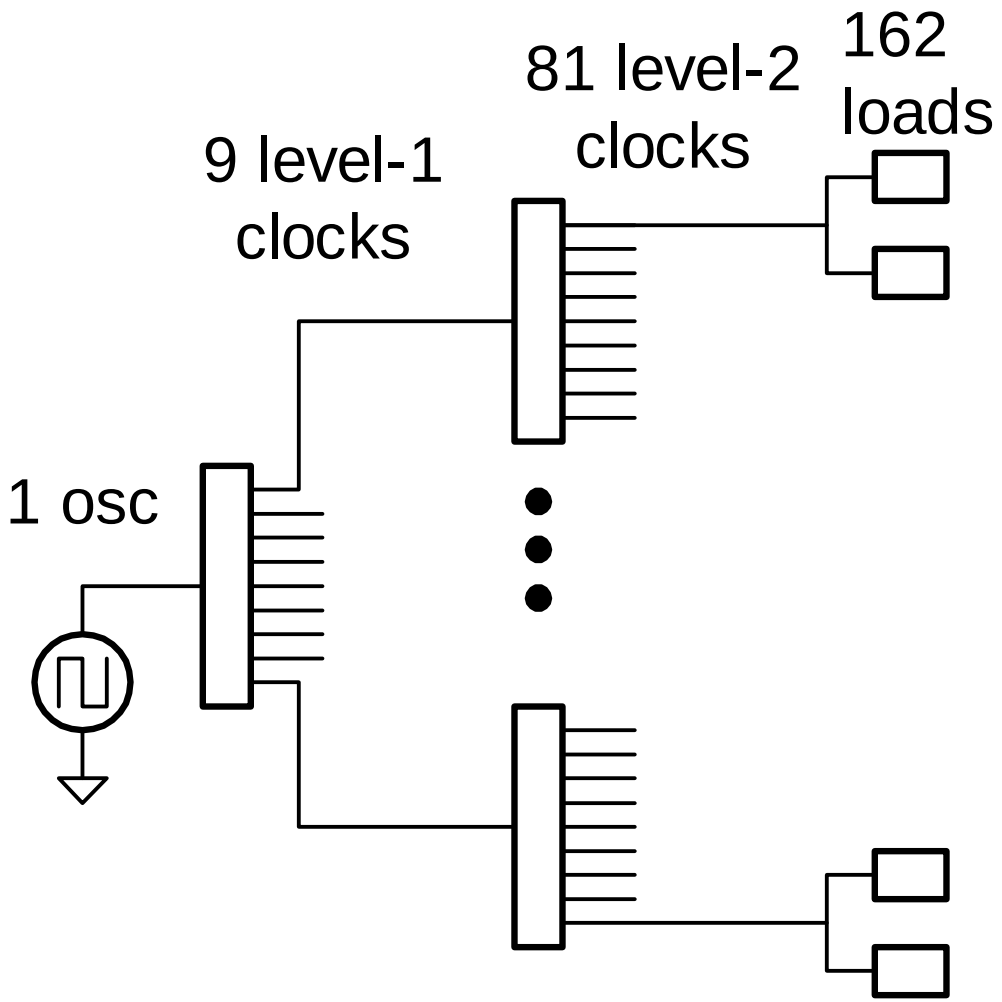

All lines are differential and terminated into 50Ω

Clock Trimming

- If open-loop skew is not good enough, clocks can be *trimmed*
  - adjust the *length* of each fanout to equalize delay
  - trombones
  - RC padding
  - modulation of buffer delay
- This can be done manually or automatically

Coupled Oscillators

- Avoid single point of failure
- Average phase of several clocks

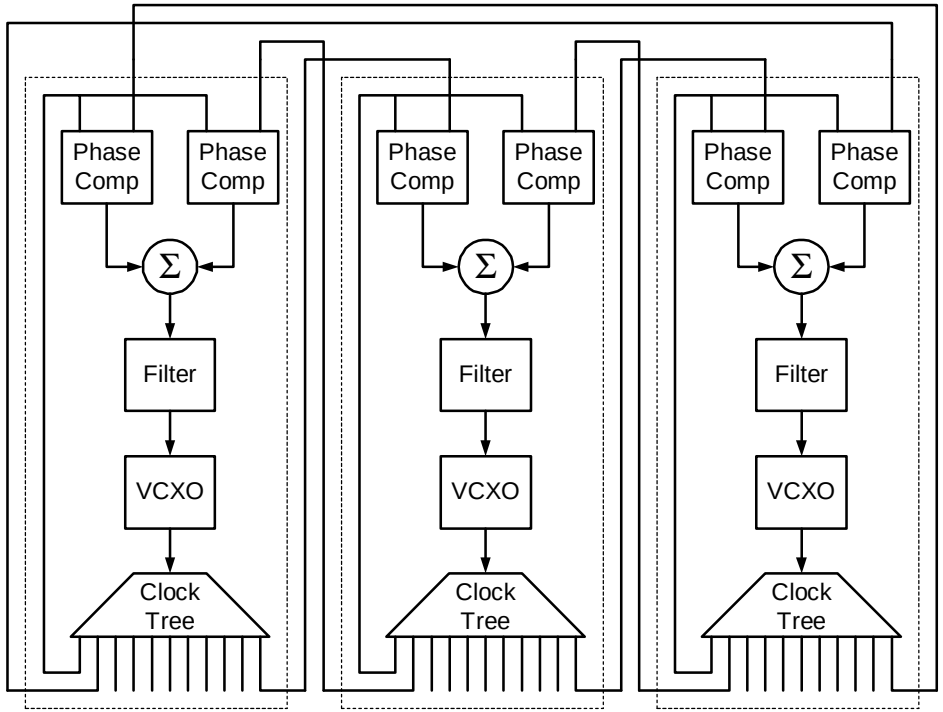

Salphasic Distribution

- Drive an unterminated (or shorted) line with a sine-wave
- Standing wave has the same phase everywhere
  - just different amplitude
- Amplify and limit this wave to generate a clock with negligible skew!

$$V_f(x,t) = \sin(wt + q - x/v)$$

$$V_r(x,t) = \sin(wt + q + x/v - 2l/v)$$

$$V = V_f + V_r = 2\sin(wt)\cos\left(\frac{x-l}{v}\right)$$

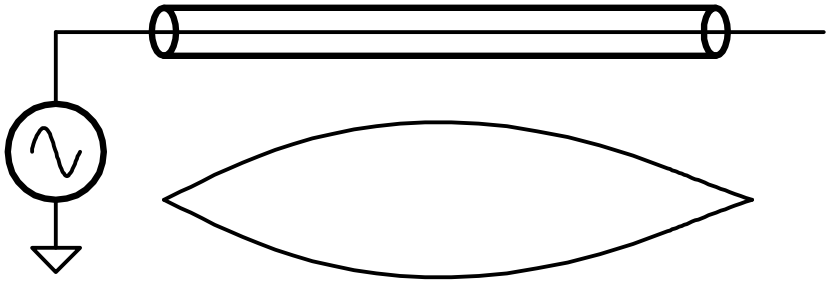

Round-Trip Distribution

- Send clock down an array of modules (e.g., a bus) and back
- Each module sees forward and reverse traveling clock
- The average time of arrival is the same for all modules
- Modules can interpolate the forward and reverse traveling clocks *or*
- Data can be sent forward with the forward clock and backward with the reverse clock

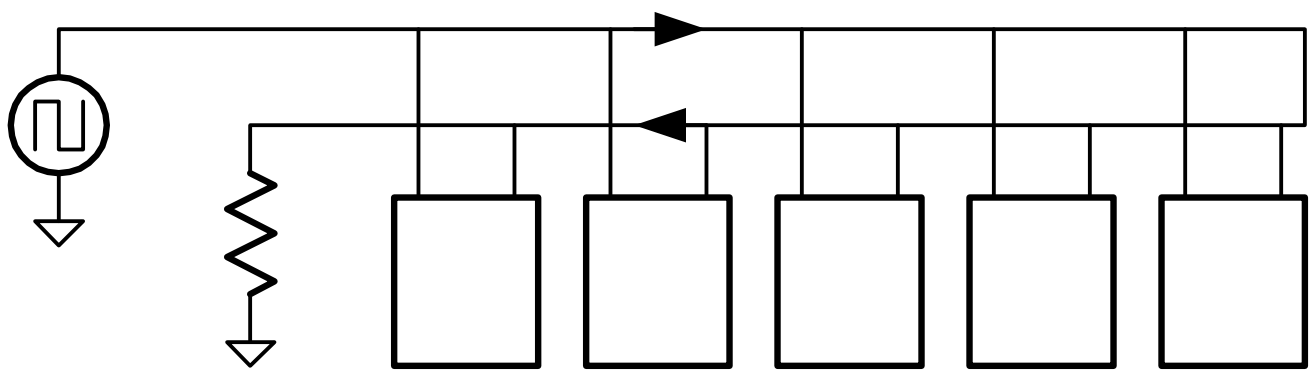

On-Chip Clock Distribution

- A large ASIC may have 10^4 to 10^5 clock loads
- Each is about 20fF
  - total load is 200pF to 2nF!
  - this load switches every cycle
- Need to
  - amplify a small clock signal to drive a 2nF load
  - distribute the clock over resistive wires to the amplifiers and loads
    - repeaters needed on long runs
  - avoid jitter
- Jitter comes from 3 main sources
  - differential supply variation modulates the delay of the clock buffers
  - single supply variation modulates the thresholds of the clock buffers
  - crosstalk to adjacent lines changes the delay of the clock fanout lines
- Each of these sources can be dealt with

An On-Chip Clock Tree

- Distribute clock from center of chip
- Fanout first in X, then in Y
  - H-tree layout also possible
- Line lengths balanced at each stage of fanout
- Repeaters needed in long lines
- Wider than minimum width lines usually used
- Adjacent ground shields to eliminated crosstalk
- Differential distribution to avoid threshold shift

- Buffer sizing adjusted to match delay to varying spatial load

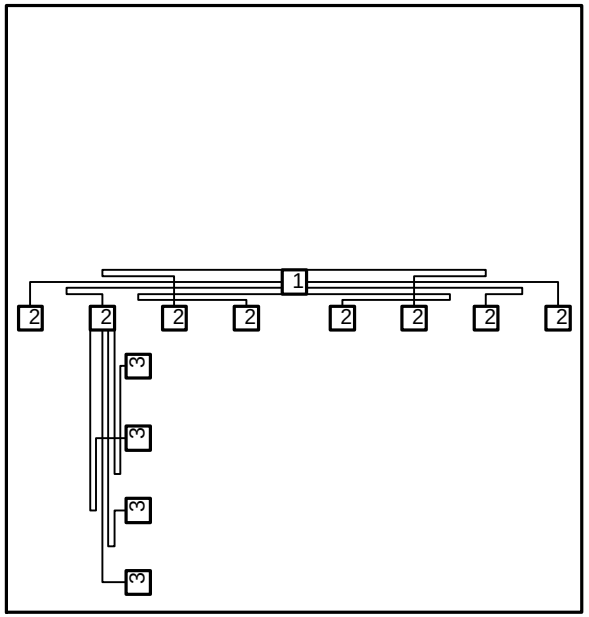

Reducing Jitter

- Address each of the three sources
- Differential Supply Noise
  - use a separate ‘clock’ power supply
  - adequately *bypass* this supply on chip
  - use insensitive buffers
- Single Supply Noise
  - use differential signaling
- Crosstalk
  - surround clock lines by parallel *shields*

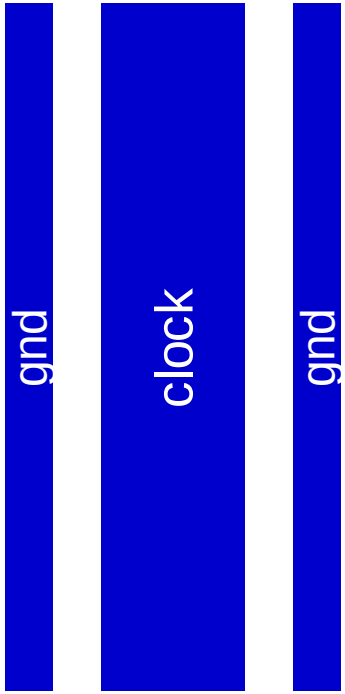

Clock Grids

- Tie points of clock distribution network together in a grid
- Averages phases of nearby regions
- Performance depends on RC time constant of grid
- Need to avoid supply overlap current

On-Chip Clock Trim

- Can locally adjust the phase of clock drivers to compensate for
  - Load variation
  - Process variation in distribution network
- Adjust via
  - Phase locking local oscillators
  - Trimming variable delay

Next Time

- Synchronization