# LPX: *a low-power processor with a locally asynchronous (IPCMOS) execute pipe***

**Alper Buyuktosunoglu++, Michael Gschwind, Pradip Bose, Hans Jacobson, Philip Emma**
**IBM T. J. Watson Research Center**
**Yorktown Heights, NY**

*** contributors to LPX project (prior and current):*
P. Bose, D. Brooks, A. Buyuktosunoglu, P. Cook, K. Das, P. Emma, M. Gschwind, H. Jacobson,
P. Kudva, S. Schuster, J. Smith, V. Srinivasan, V. Zyuban: IBM. T. J. Watson Research Center.

D. Albonesi and S. Dwarkadas: Univ. of Rochester, NY.

K. Das: summer intern from U of Michigan
T. Karkhanis: summer intern from U of Wisconsin
J. Smith: was on sabbatical from U of Wisconsin

++Alper Buyuktosunoglu is also a Ph.D student at University of Rochester, NY

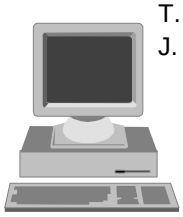

MICRO-35 Tutorial, November 18, 2002, Istanbul, Turkey.

# Current Generation High-End Processor Power Profile

(Quoted from prior published PACS-02 paper)

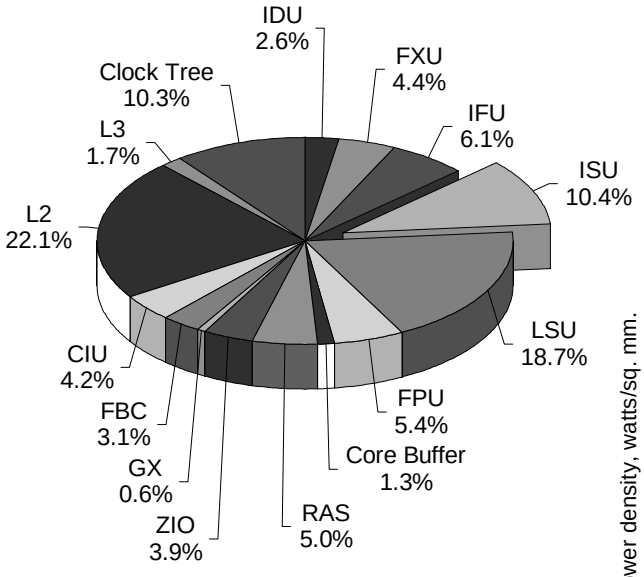


(Non-validated, pre-silicon, POWER4-like superscalar design)

**Instruction Seq. Unit Power Breakdown**

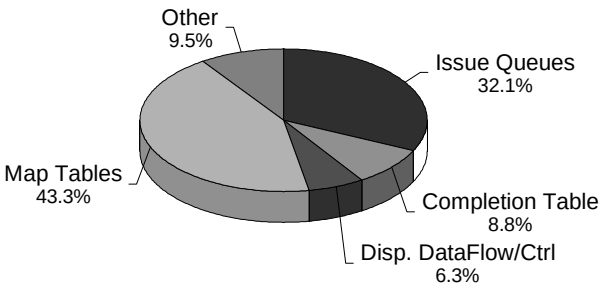


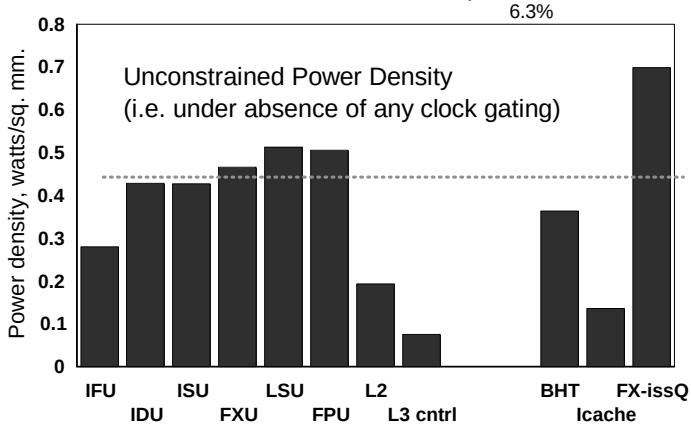

# Energy Waste in High-End Superscalar: Non-Gated

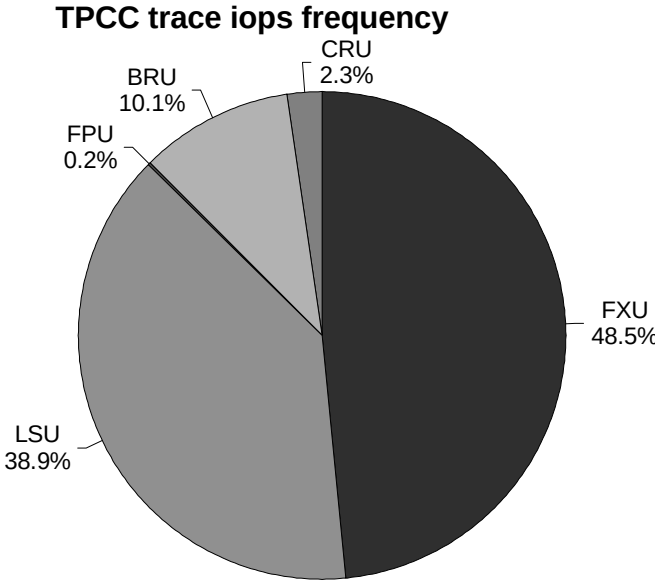


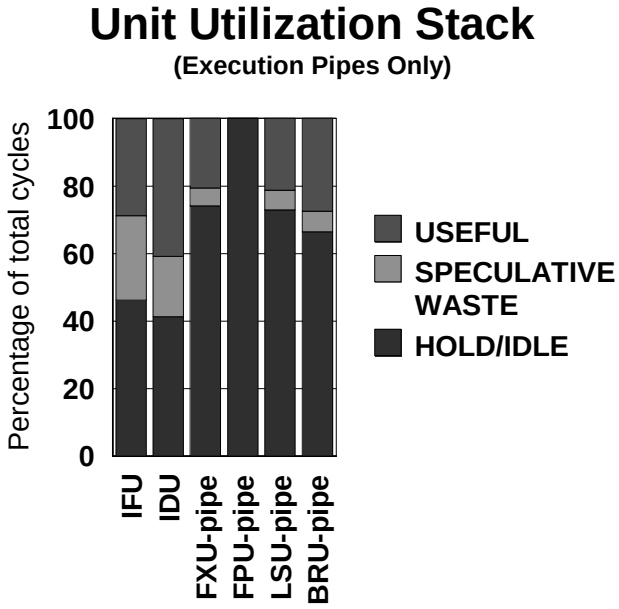


Coarse-grained clock gating opportunities

Pre-silicon, non-gated POWER4-like simulation mode results



# LPX: High-Level Block Diagram

**(Goal: demonstrate 5X power density reduction with *at most* 5 % IPC performance hit)**

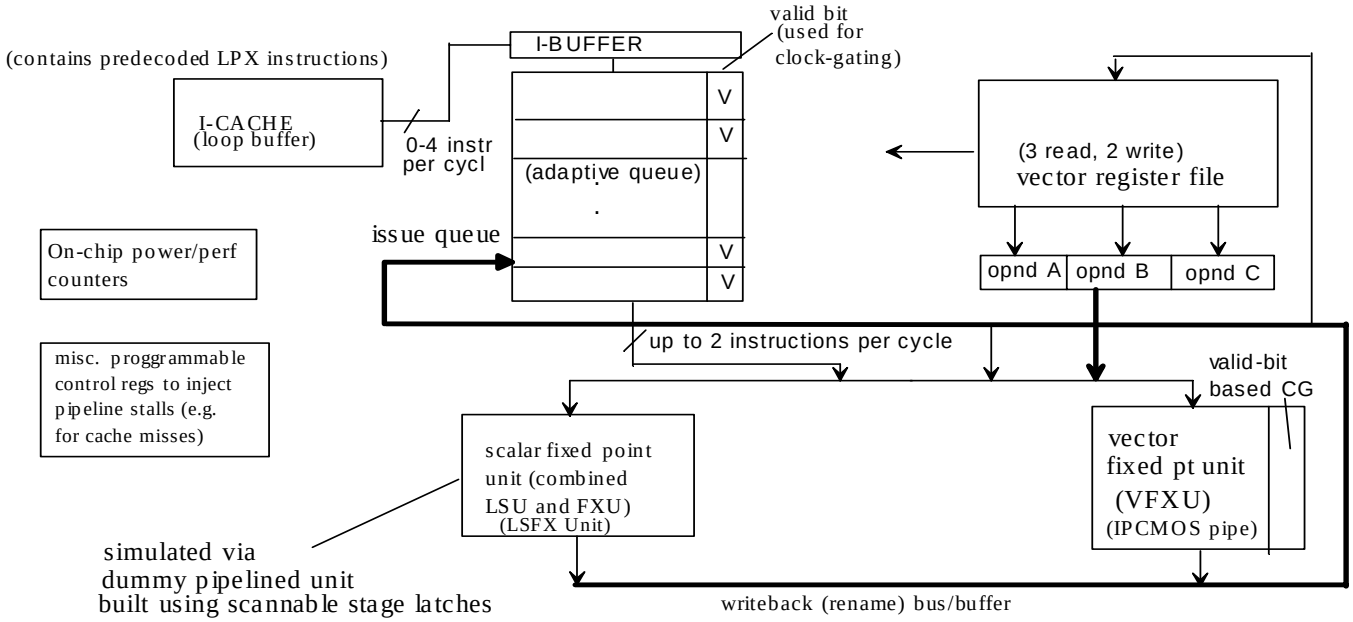

# Key Design Ideas in LPX

- Microarchitecture-level
  - adaptive resource sizing: dynamic, application-driven;
  - dynamic ifetch throttling: set of simple fetch-gating heuristics;
  - *new: combined fetch-gating + adaptive queues* (ongoing research)
  - microarchitectural support for clock gating: coarse-grain, fine-grain; "on-demand"
  - on-chip counter architecture for power/perf monitoring and feedback
- Logic/circuit-level
  - adaptive clocking: save power, maintain performance
    - locally generated, enabled by data (IPCMOS)
    - synchronous clock-gating: coarse and fine-grain (ISP)
    - variable or phased step-up/down: on-demand clock-gating (OD-CG)
    - Vdd-gating for banked issue queue (new addition - being incorporated)
  - power-efficient CAM/RAM structures - research only

***note: not all features above will be implemented; but are studied in high-level microarchitectural research work (simulation-based)***



# Methodology for LPX Definition and Tuning

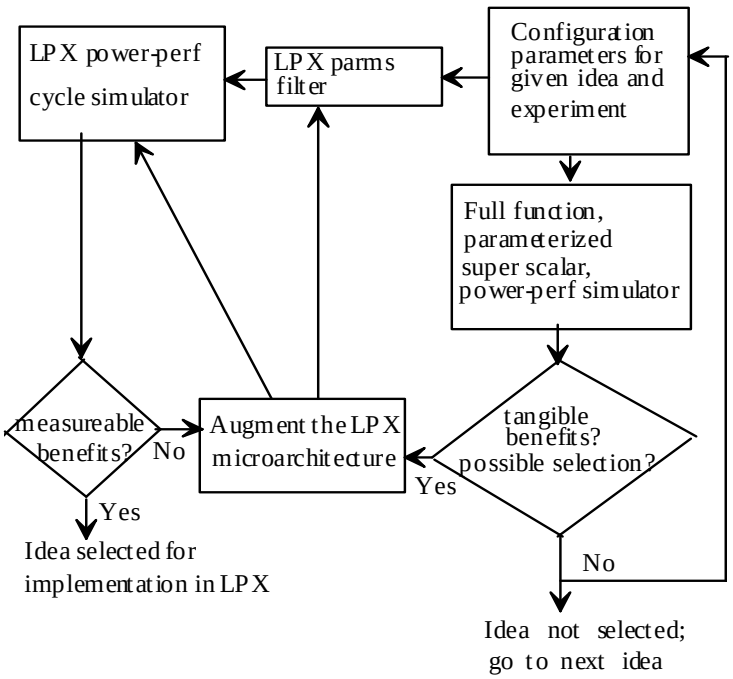


- Idea selected must have potential power-perf "worth" based on:
  *- simulation "in the large".*
- Basic idea is adapted to a much simpler hardware heuristic, appropriate for LPX:
  *- simulation and prototyping "in the small".*
- Extensive circuit simulation to validate power savings and overhead:
  *- done before microarch. def.*

# Interlocked Synchronous Pipelines (ISP)

PIPELINE STRUCTURE

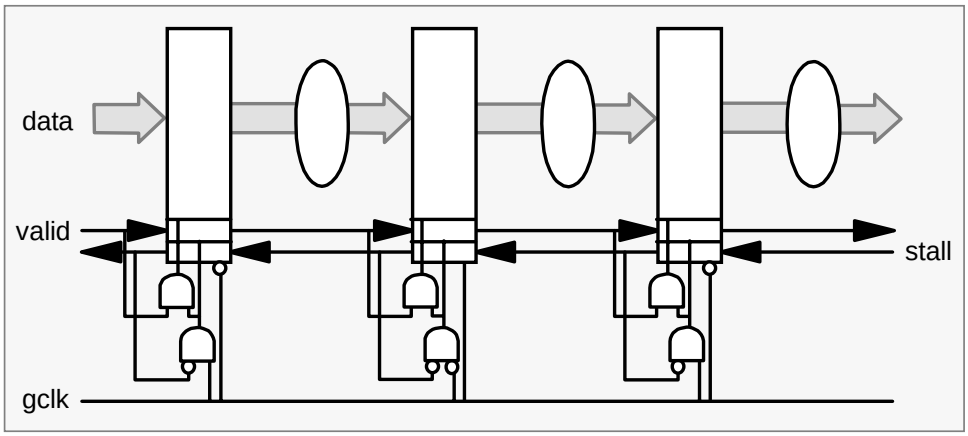


(Paper accepted for publication in ASYNC'2002:
H. Jacobson et al.)

- Clock gating is performed local to each stage on a cycle basis
- Clock is gated on both invalid data and stall conditions
- ISPs are testable using standard LSSD techniques



# ISP Results

## Clock Gating Results: FPU Datapath (RTL level)

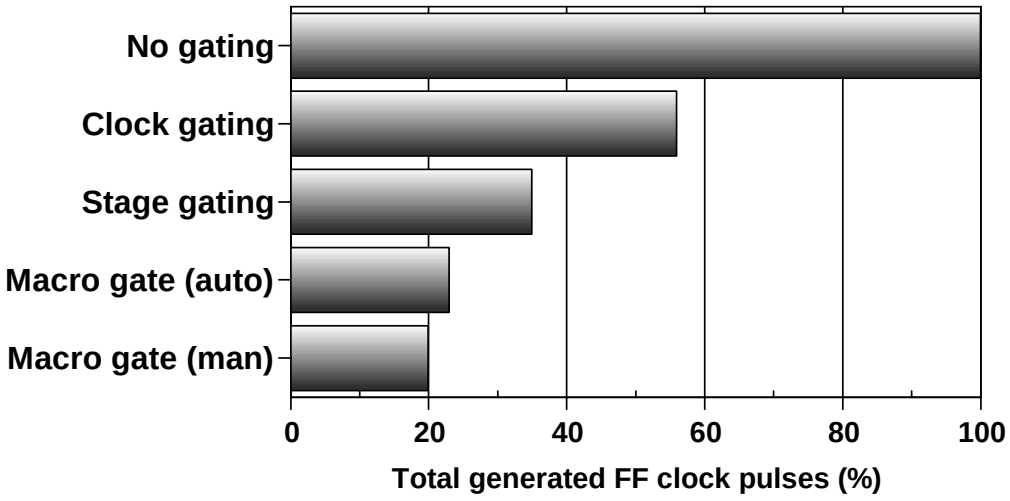


- ISP savings also being tested with LPX timer

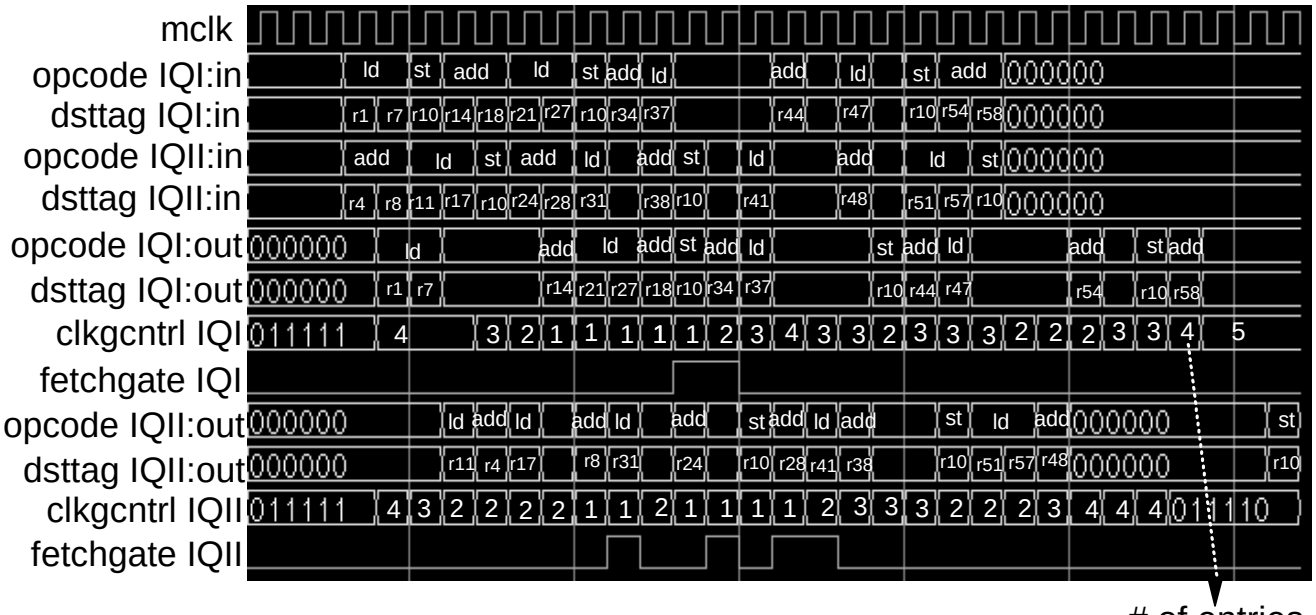


vector add loop:
ld r1 r2 r3 r r
add r4 r1 r6 nr r
ld r7 r2 r5 r r
add r8 r4 r7 nr nr
st r10 r8 r9 nr r

preset ready bits
mimic the rename
(r: ready, nr: not ready)

## *LPX VHDL results*





# Validation Methodology

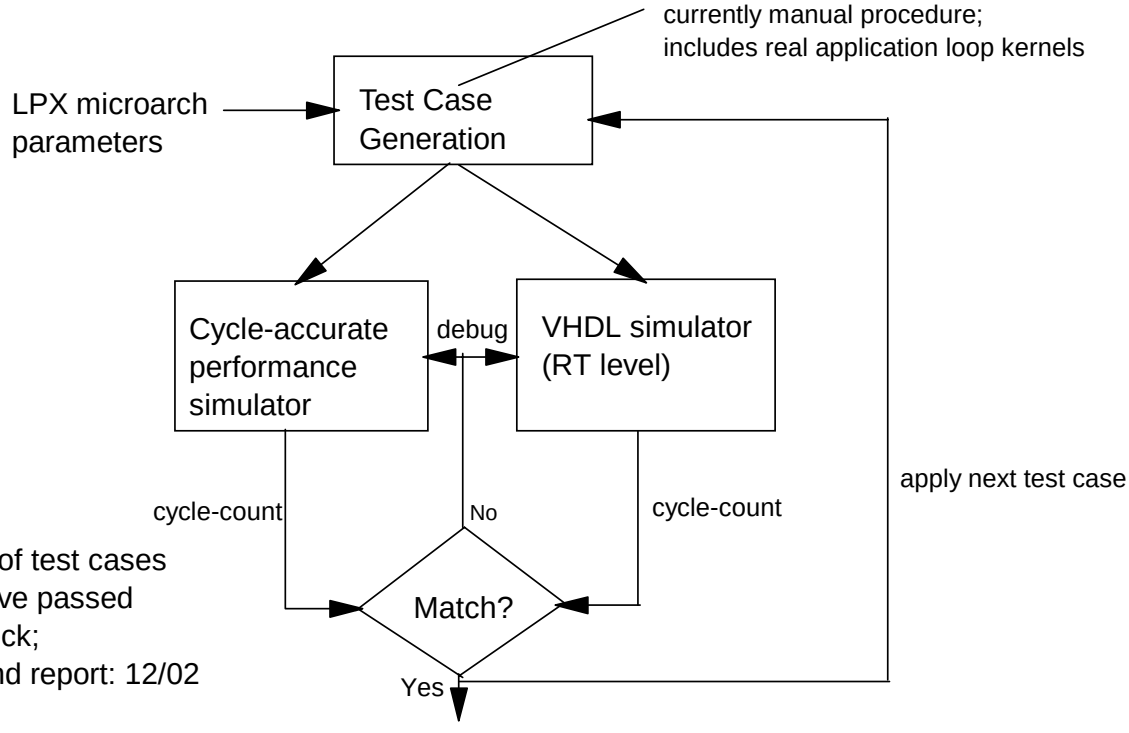


*Status:* 90 % of test cases developed have passed validation check;
completion and report: 12/02

# LPX Design Schedule (2001-03)

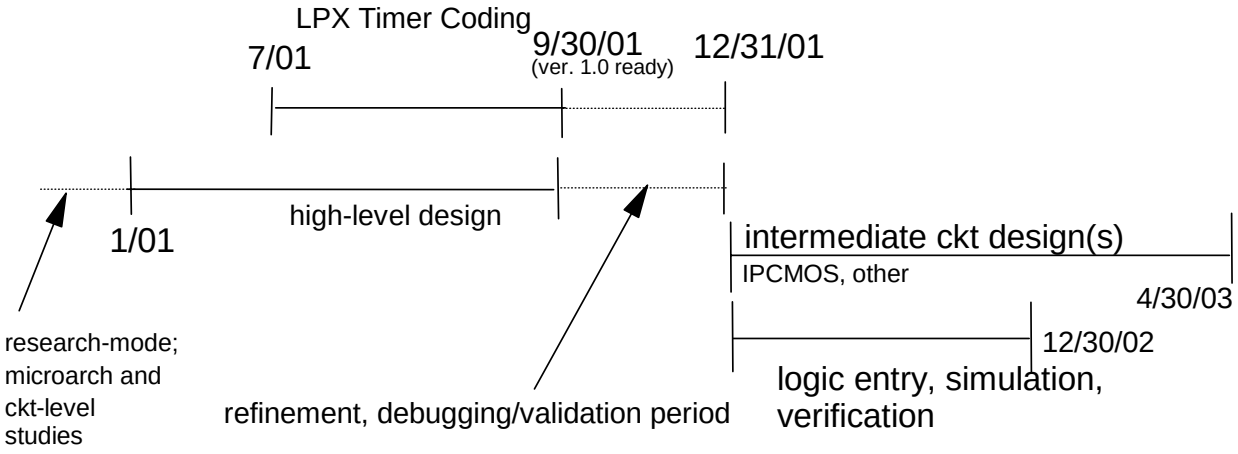


Physical design, implementation phase: through 2003



# Power versus Switching Factor

**(assuming 70% of power in clocks and latches as in POWER4)**

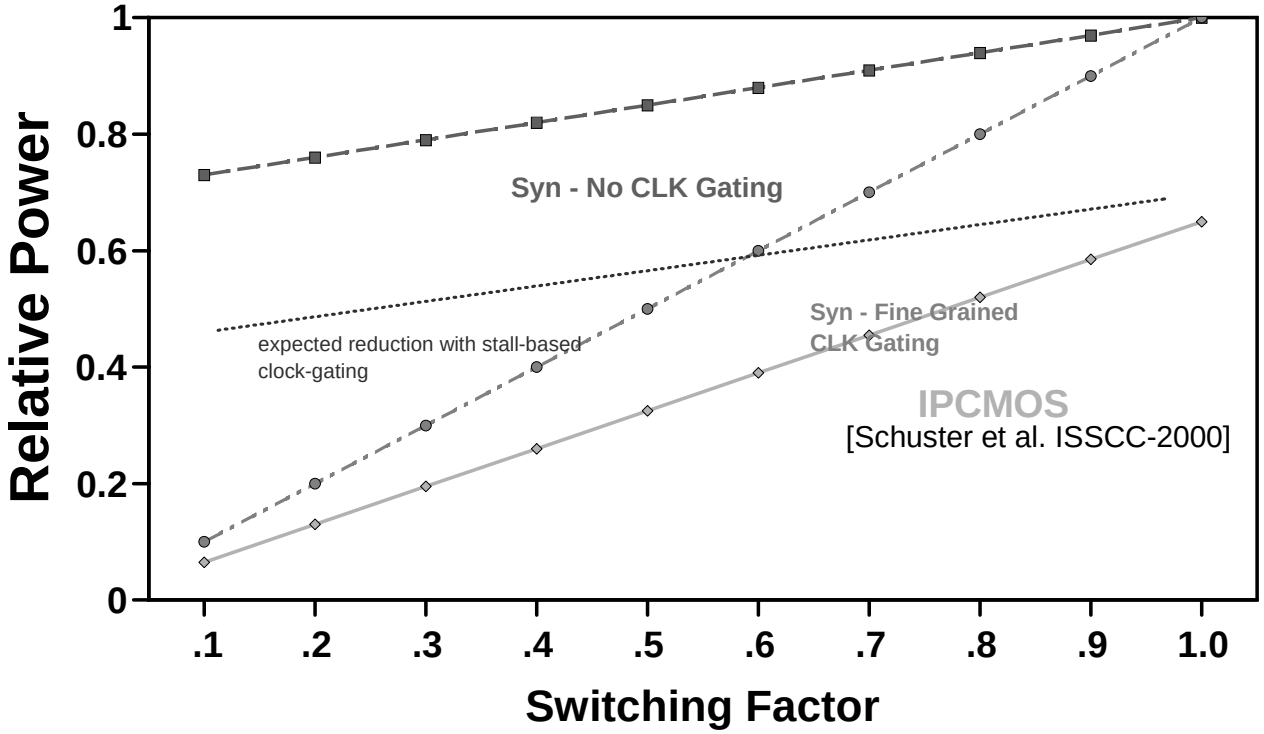

# Conclusion

## SUMMARY:

- Power density limits in high-end processors
  - ✓ power-aware microarchitecture + advanced clocking/ckt techniques
- LPX: a small research processor prototype to validate newer ideas
  - ✓ high-level microarchitecture definition and analysis reported
  - ✓ processor is scheduled for tapeout in 2003

## FUTURE/ONGOING WORK:

- Logic design/VHDL simulation and verification: by year-end 2002
- Power-efficient cache design research for LPX follow-on chip
- Power monitoring/measurement hardware architecture
- Measurement-based calibration of pre-silicon energy models



# Low Power Locally Asynchronous Interlocked Pipelined CMOS (IPCMOS) Clock Circuits Operating at 3.3-4.5GHz

# (0.18 micron; 1999)

Stan Schuster, Pete Cook

IBM Research Center
Yorktown Heights, NY

# IPCMOS Team Members

- Concepts
  - Stan Schuster, Pete Cook
- Test Site
  - Stan Schuster, Bill Reohr, Pete Cook, Dave Heidel, Mike Immediato, Keith Jenkins



# LPX: a low power issue-execute VMX processor

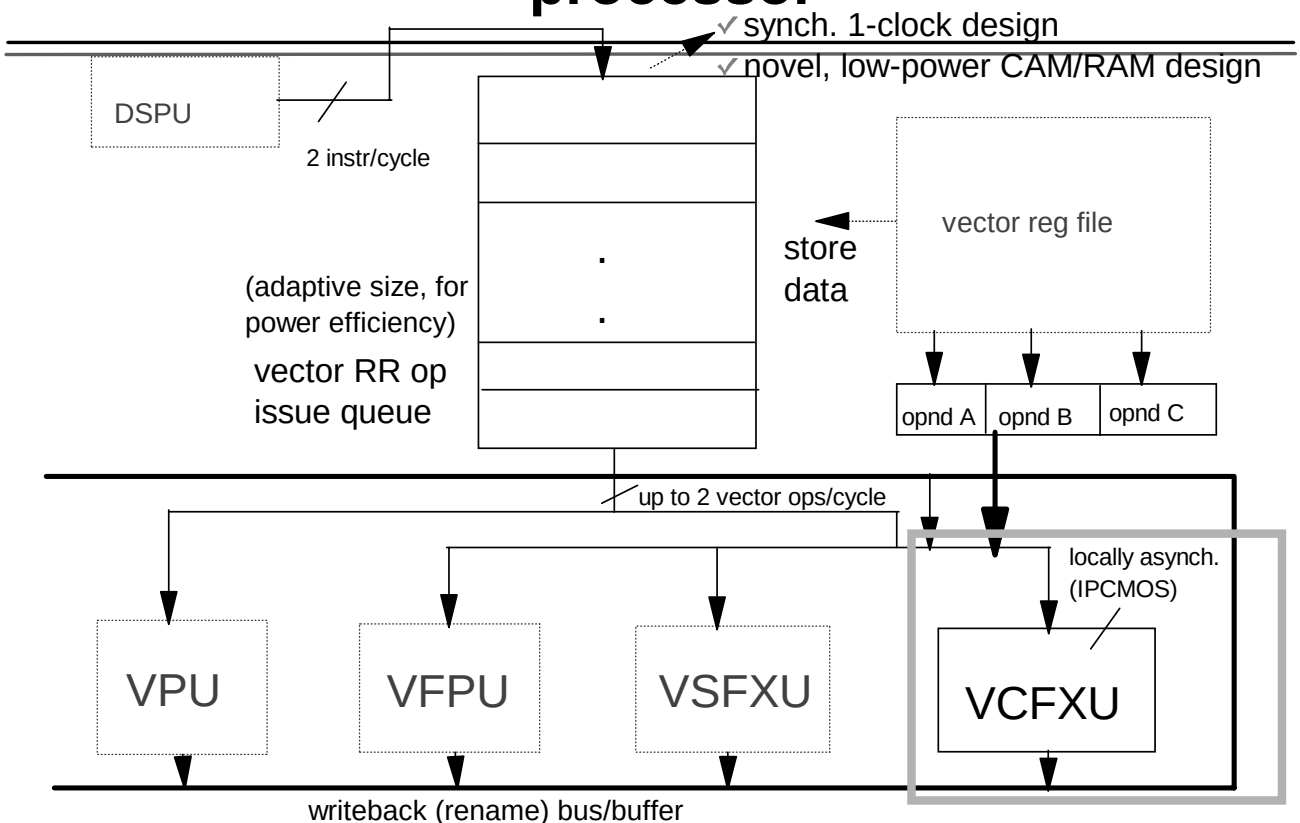

# Early Work

- Communications protocols
- Sutherland paper on "Micropipelines" in the Communications of the ACM 1989
- Muller C-element described by Mead and Conway 1980
- Williams et al. "-- 160ns 54-b CMOS Divider" JSSC 11/91
- Chappell et al. "-- SRAM with Fully Pipelined Architecture" JSSC 11/91
- Proceedings of the IEEE Special Issue on Asynchronous Circuits and Systems 2/99



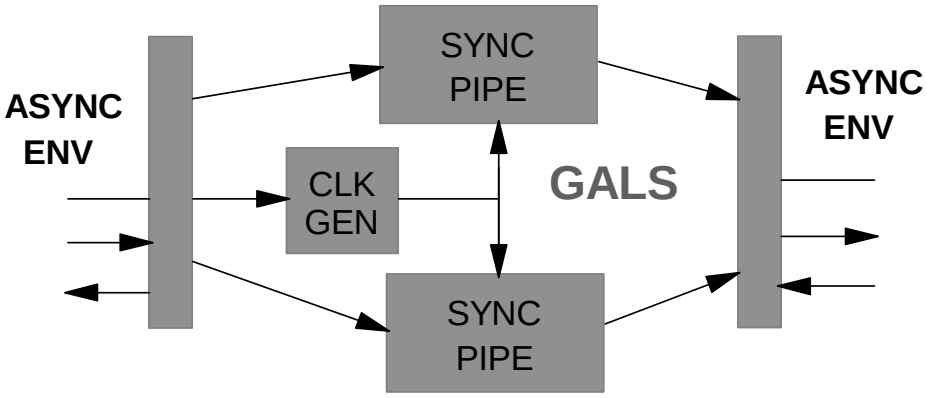


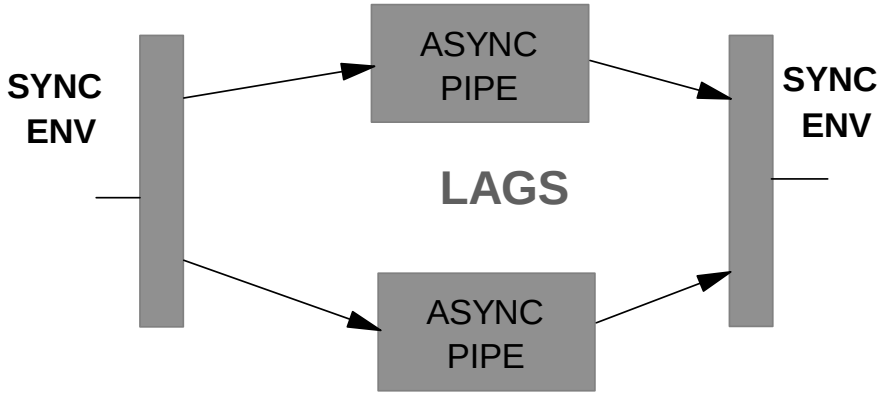

# Why IPCMOS?

- Lower Power
  - Clocks enabled only when operation to perform
  - Single stage transparent latch
- Higher speed
  - Single stage transparent latch
  - Interlocked local clocks
- Circuits deal with global timing issues
  - Delay variations from power supply noise
  - Delay variations from chip parametrics
- Asynchronous to synchronous interface possible
- Both ac and dc testing possible



S - Strobe R - Reset

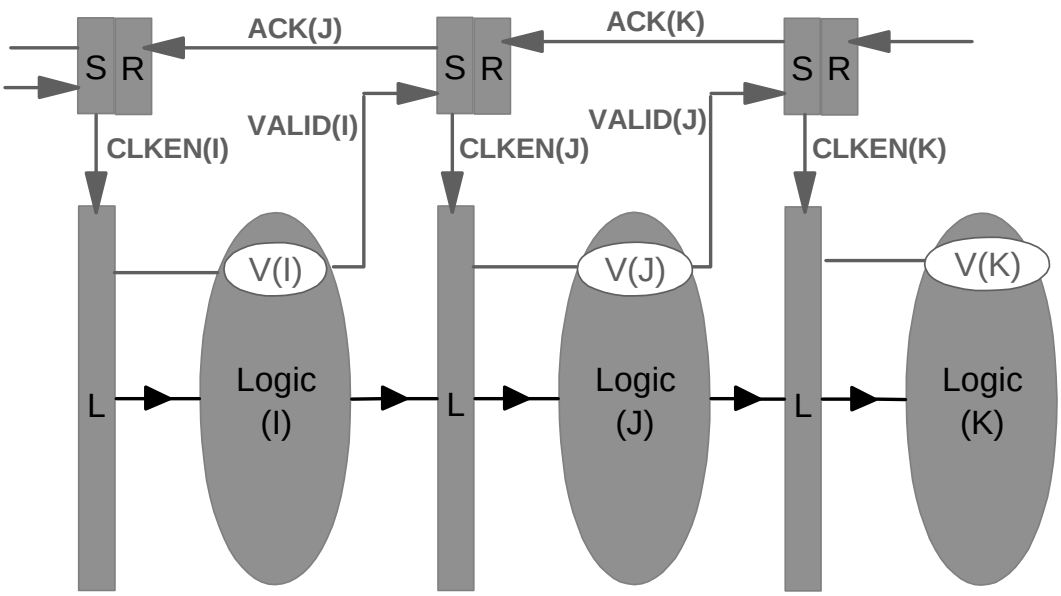


## *IPCMOS Interlocking*

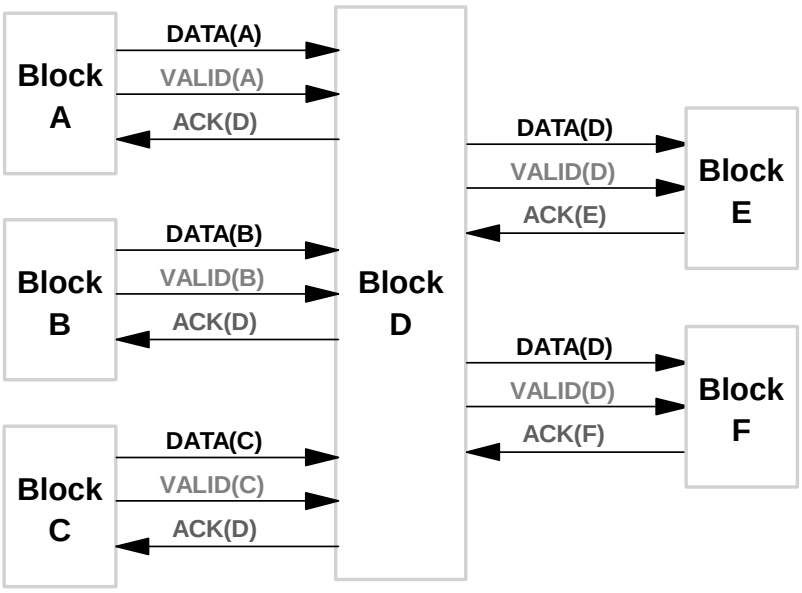


***Interlocking at the Block Level***



# IPCMOS Circuits

- Circuits are key to achieving high speed!
- Key IPCMOS Circuits
  - Strobe
  - Reset
  - Latch
  - Valid

# Strobe

- Unique "AND" circuit
  - Detects all inputs low
  - Detects all inputs have gone high
- Keeps track of cycles on inputs and outputs
- Switch between internal and external inputs minimizes interlocking overhead
- High speed operation for large number of inputs



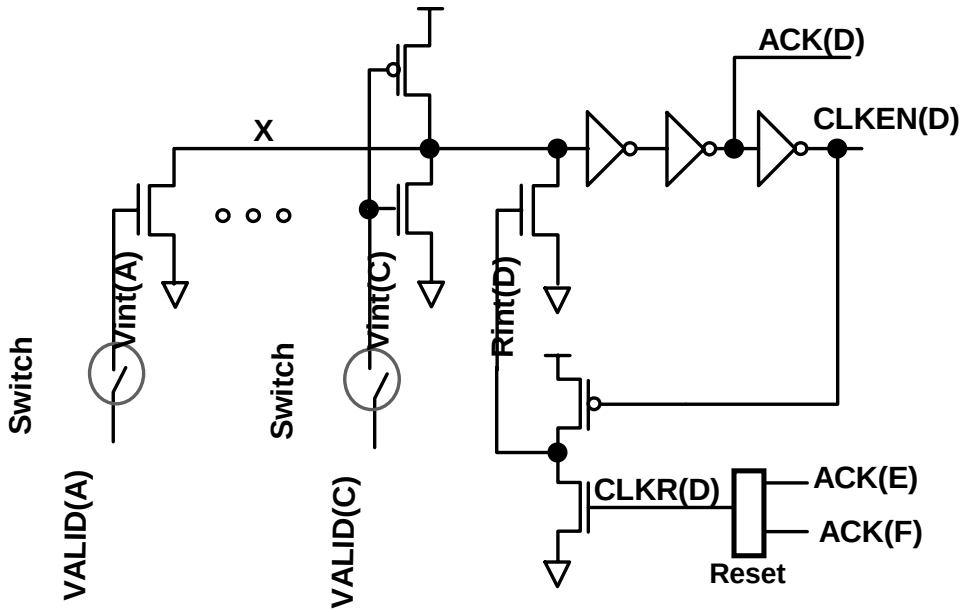


*Local Clock Strobe Circuit*

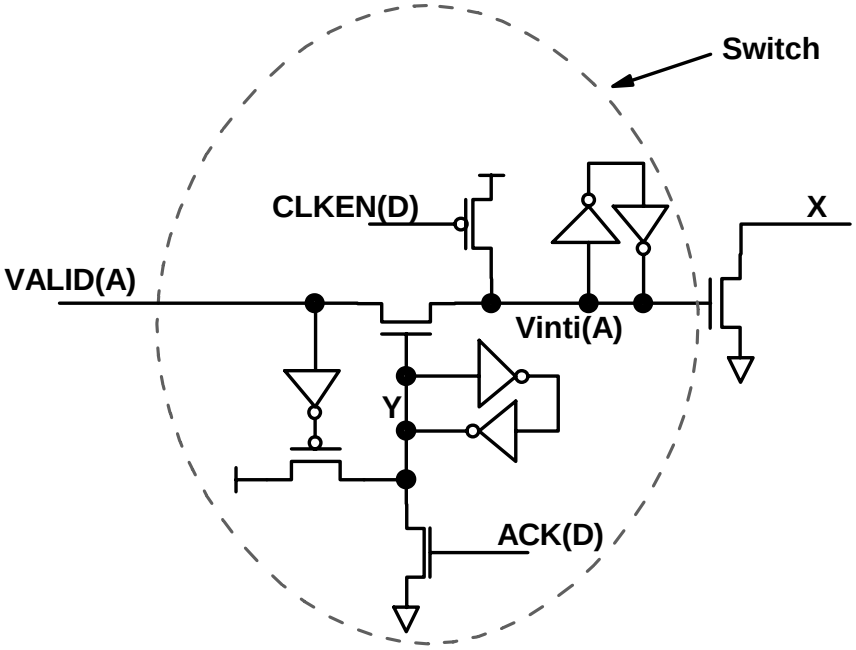


## *Strobe Circuit Switch*



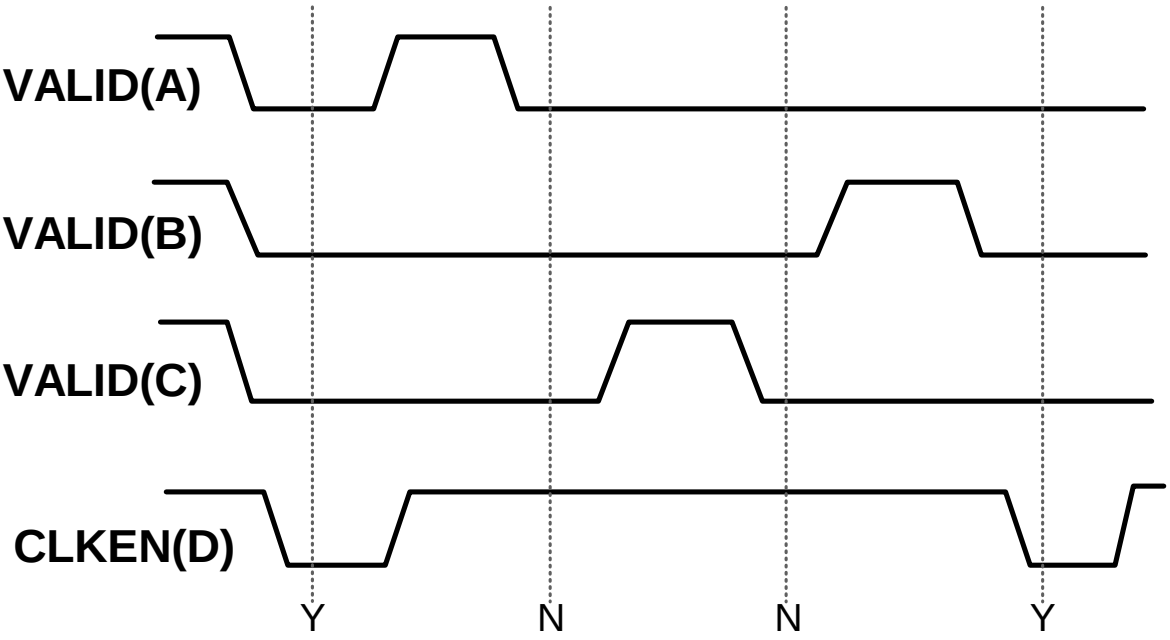


## *Strobe Circuit Unique "AND" Function*

# Strobe Delay versus # of Inputs

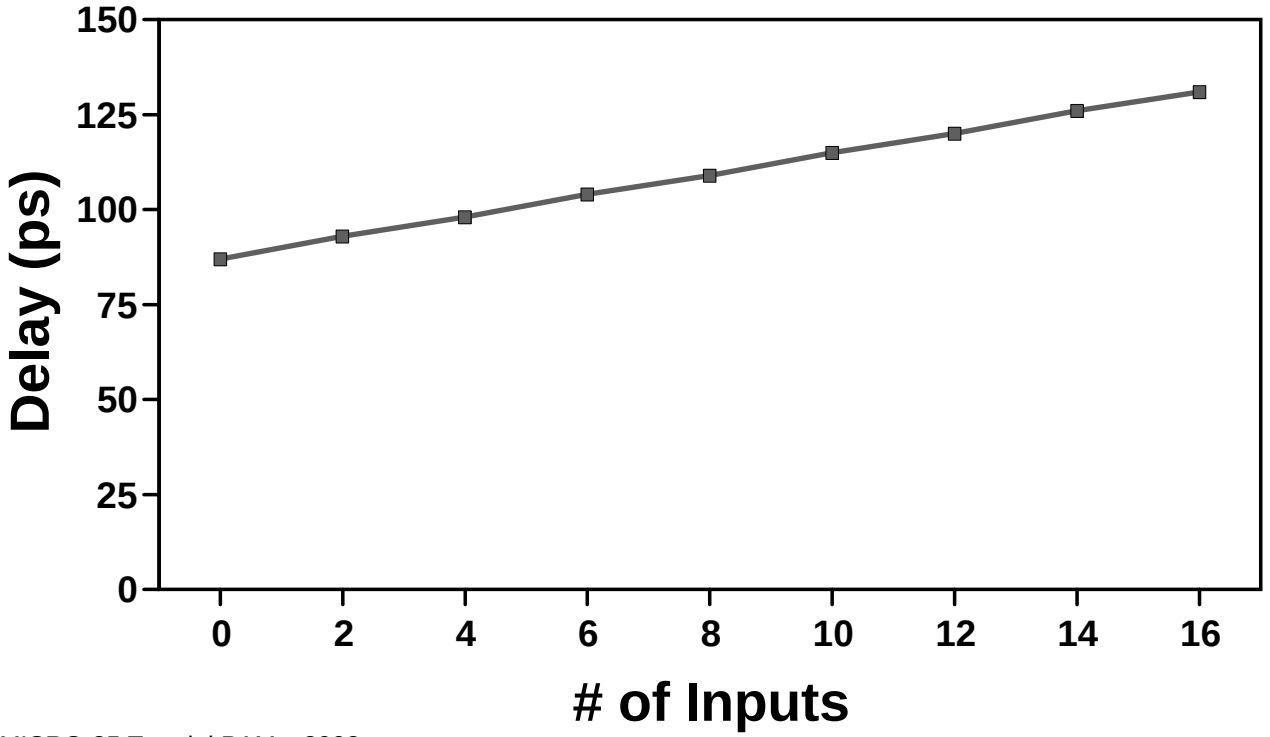




# Sensitivity Analysis with Device Variations

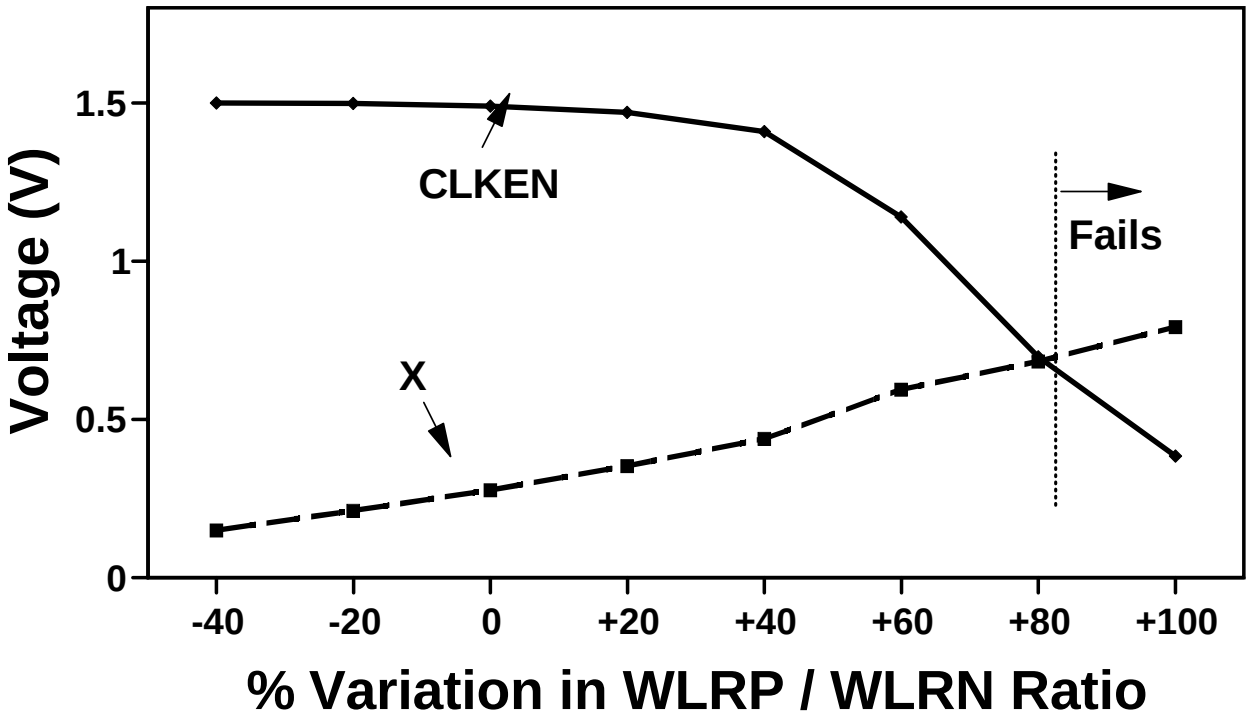

# Reset

- Combines the function of a static NOR and an input switch
- Performs unique cycle dependent AND function on the ACK signals
- High speed for high-fan is achieved
- Resets the Valid circuit each asynchronous cycle



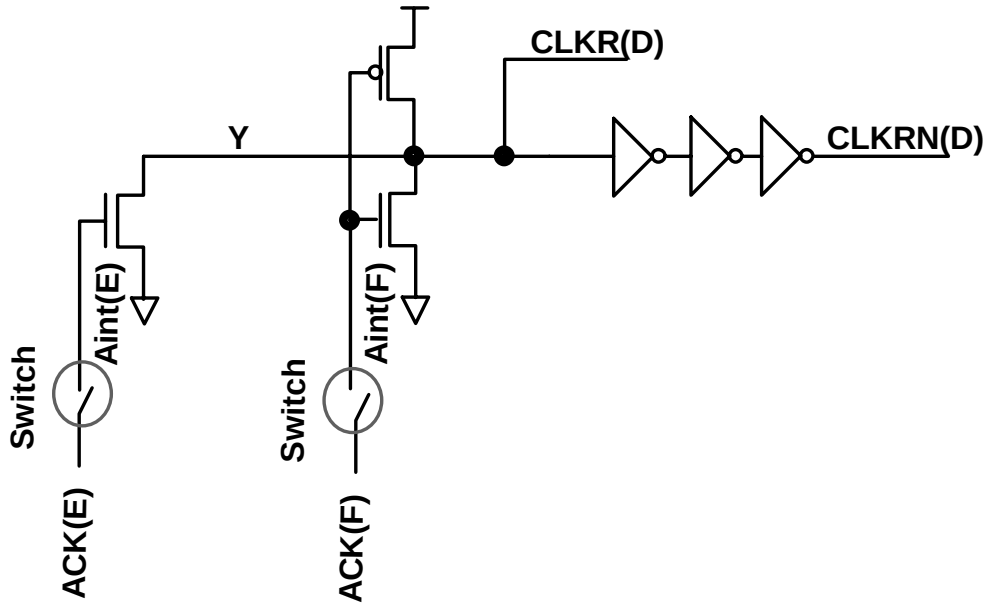


*Reset Circuit*

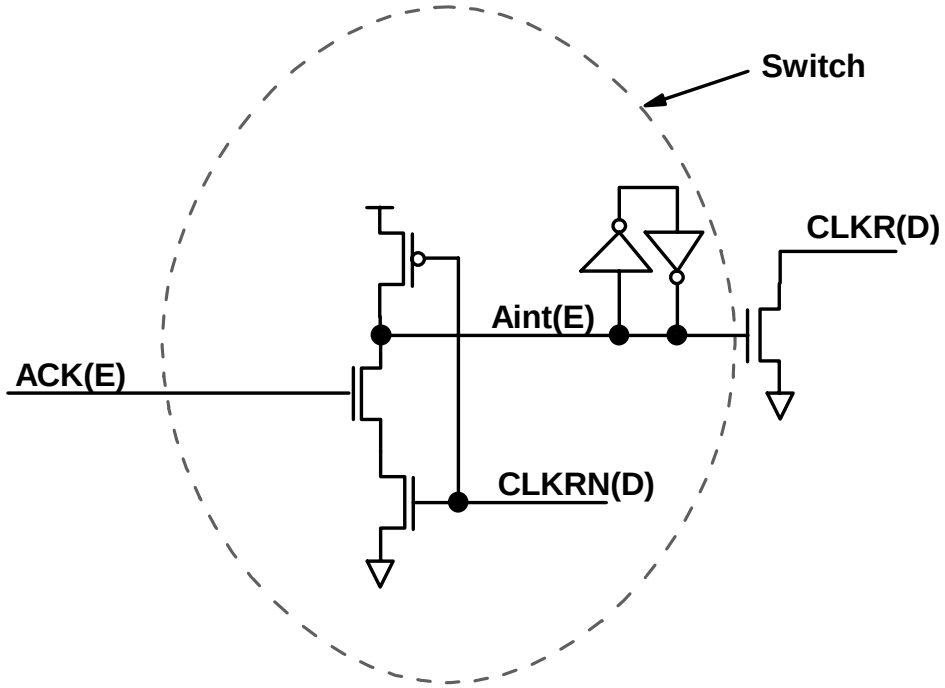


## *Reset Circuit Switch*



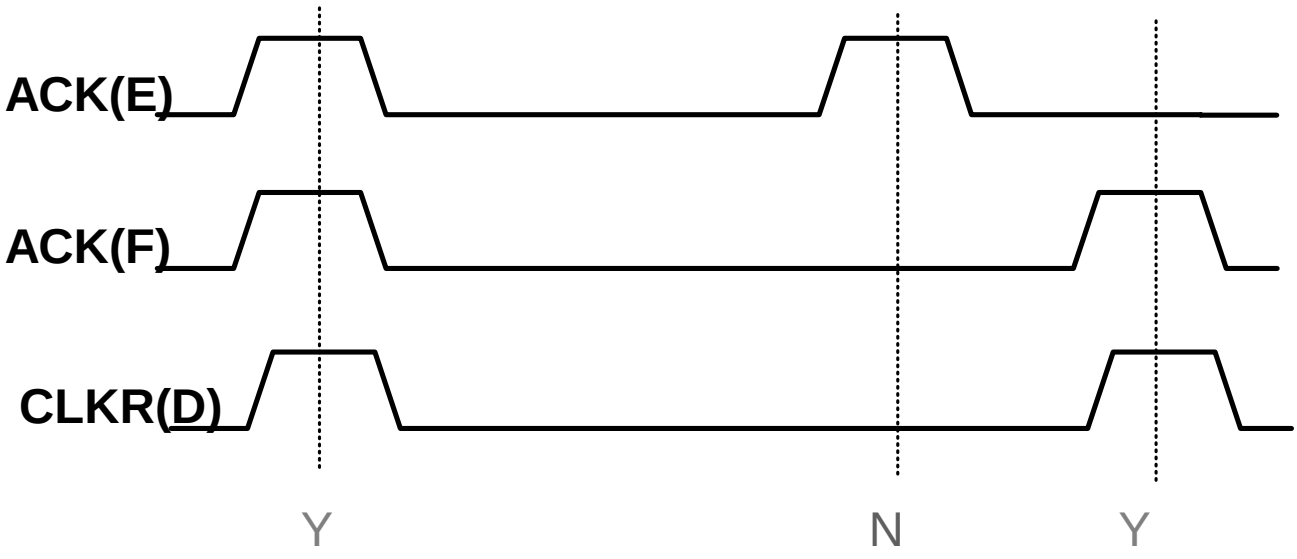


## *Unique Cycle Dependent "AND" Function*

# Latch

- Adds extremely small delay to logic path
- Single stage - not master / slave
  - Possible because adjacent latch stages are not enabled simultaneously
  - Latch enabled only when data valid
- Simple
- Simultaneous capture and launch
- Dynamic and static single rail and dual rail versions



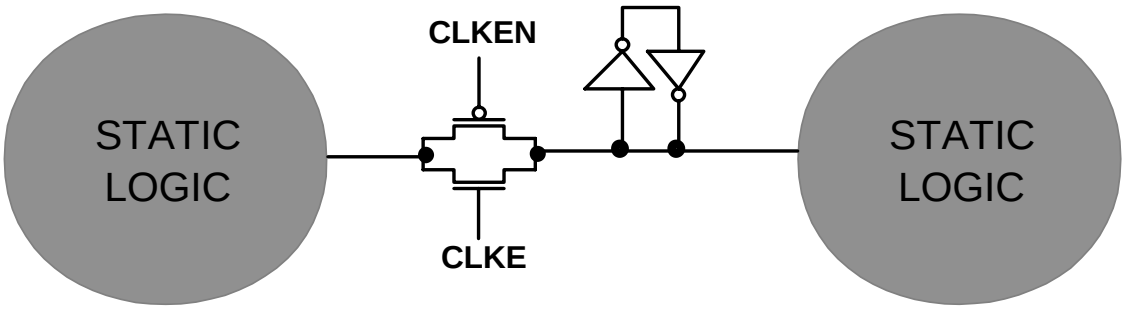


***Latch for Static Logic***

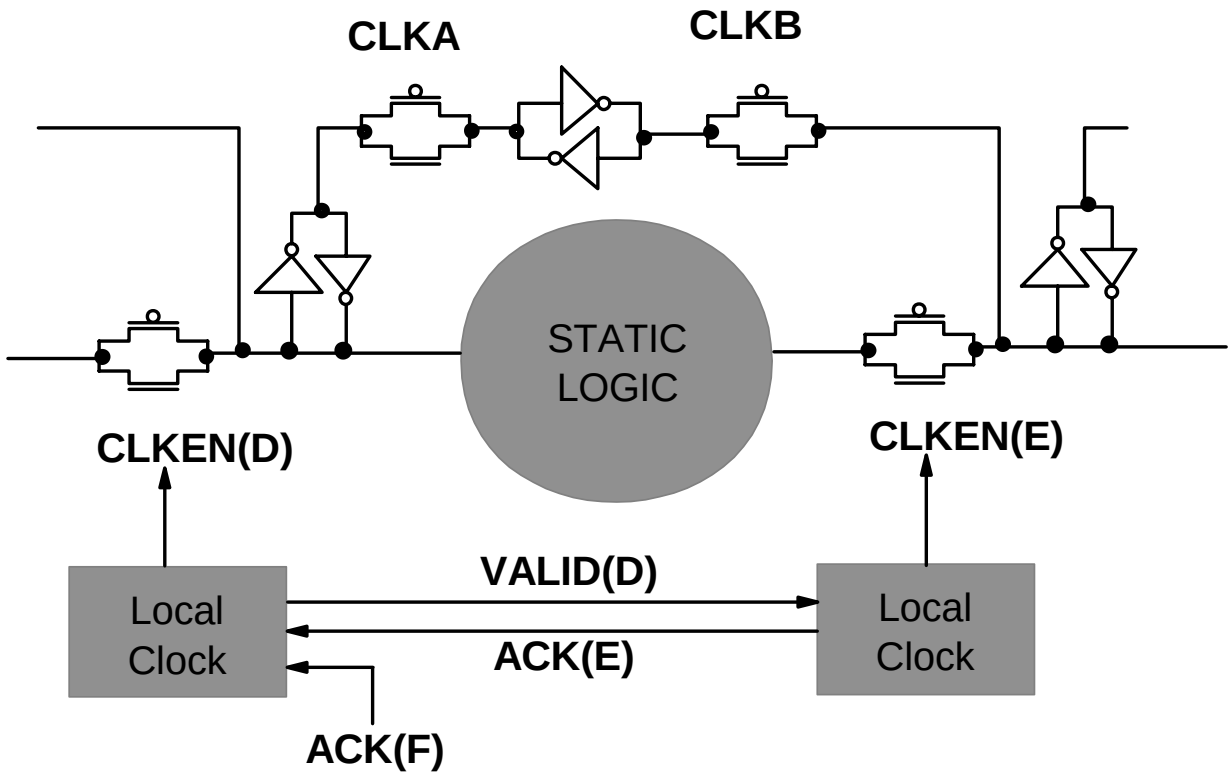


***Testing of IPCMOS with Static Logic***



# Data Valid

- Data valid initiated by CLKEN and reset by CLKR
- Timing emulates slowest data path minus strobe delay
- One valid needed per group of data outputs
- Critical timing signal
- Valid delay on test site made adjustable over wide range

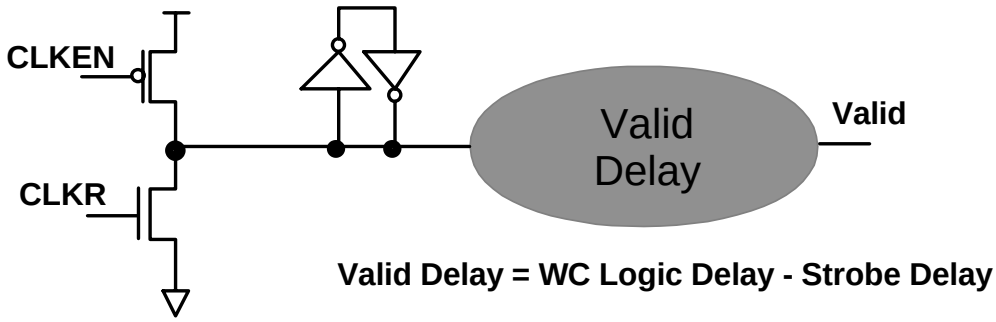


*Valid Circuit*



# Results

- Local clocks for a multiplier function implemented in 0.18 bulk CMOS technology
- High frequency operation
  - 4.5GHz under best case conditions
  - 3.3GHz under typical conditions
- Extremely robust operation, insensitive to
  - Power Supply Voltage
  - Temperature
  - Parametrics (Leff, Vt, ...)

# CMOS Technology Features

| | |
|---|---|
| Power supply | 1.5V |
| Gate oxide thickness | 2.8nm |
| n channel Leff | 0.096um |
| n channel threshold | 0.355V |
| p channel Leff | 0.12um |
| p channel threshold | 0.305V |
| Low threshold device delta | 0.07V |
| Minimum lithographic image | 0.18um |
| 4 - 6 levels of metal | Copper |

## Table I



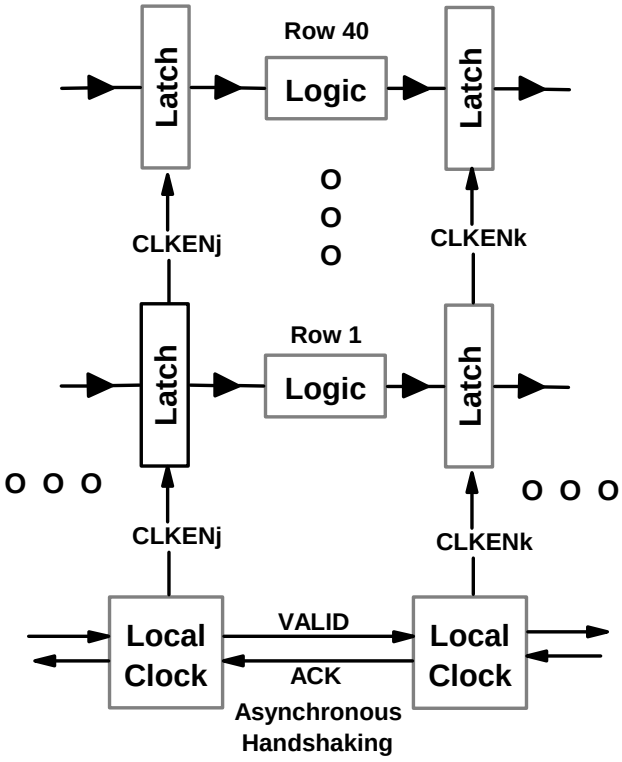


## Multiplier Core with Local Clocks

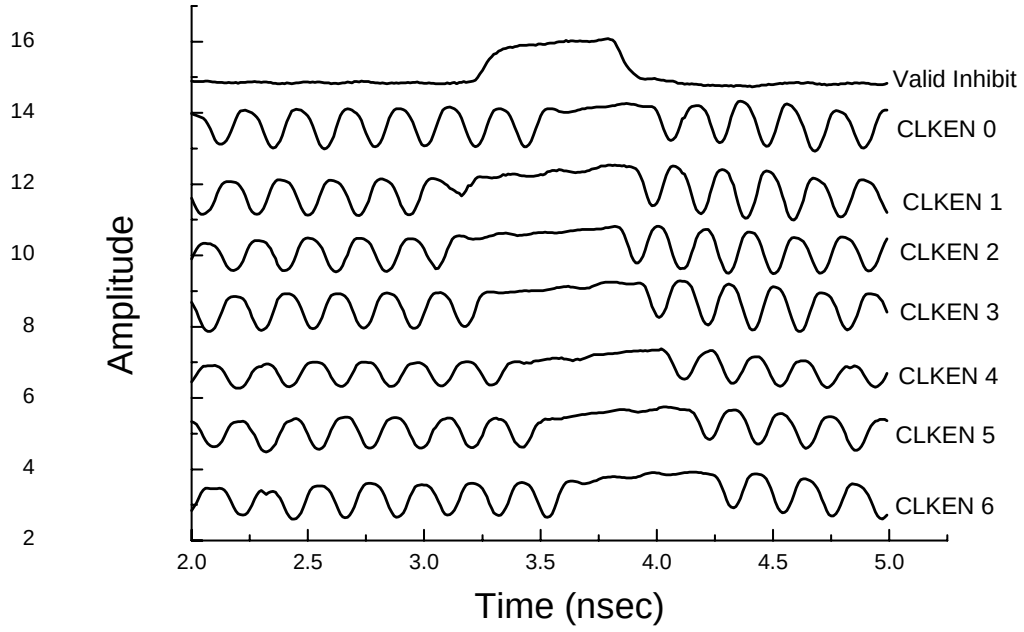


***Measured Local Clock Waveforms at 4.5GHz***



# Frequency Versus Channel Length

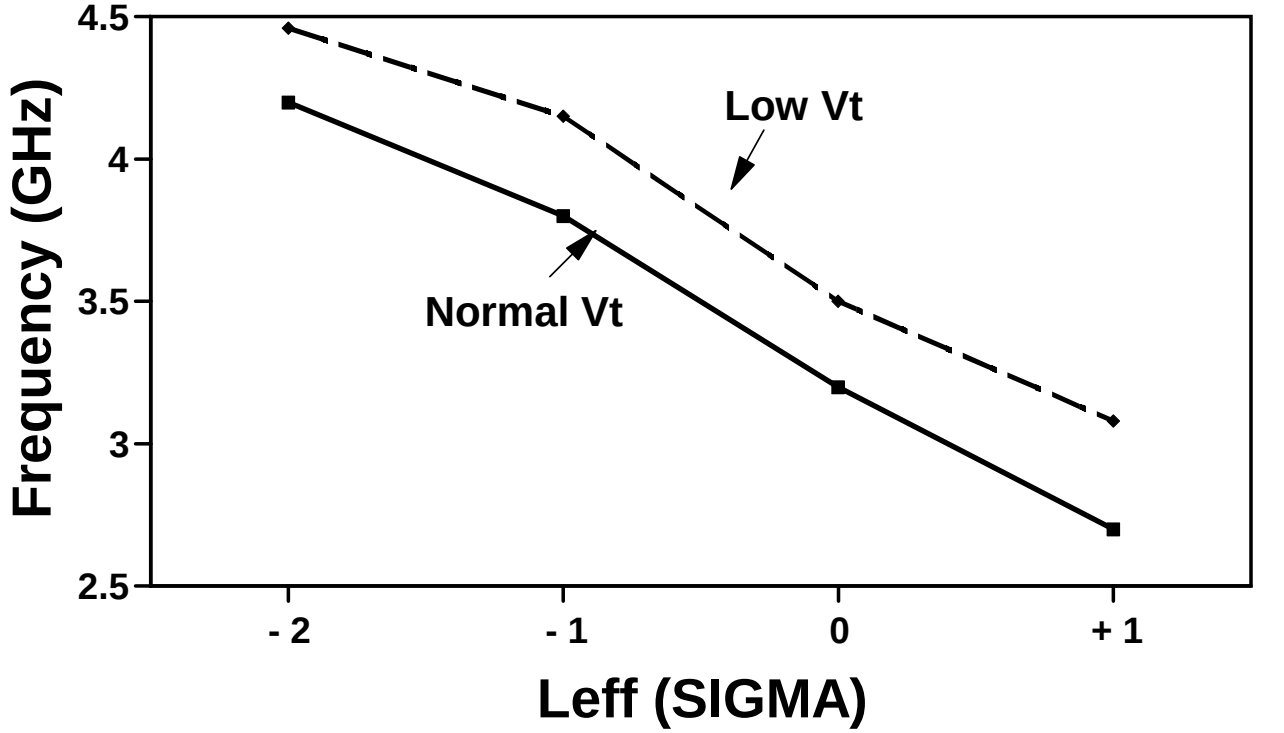

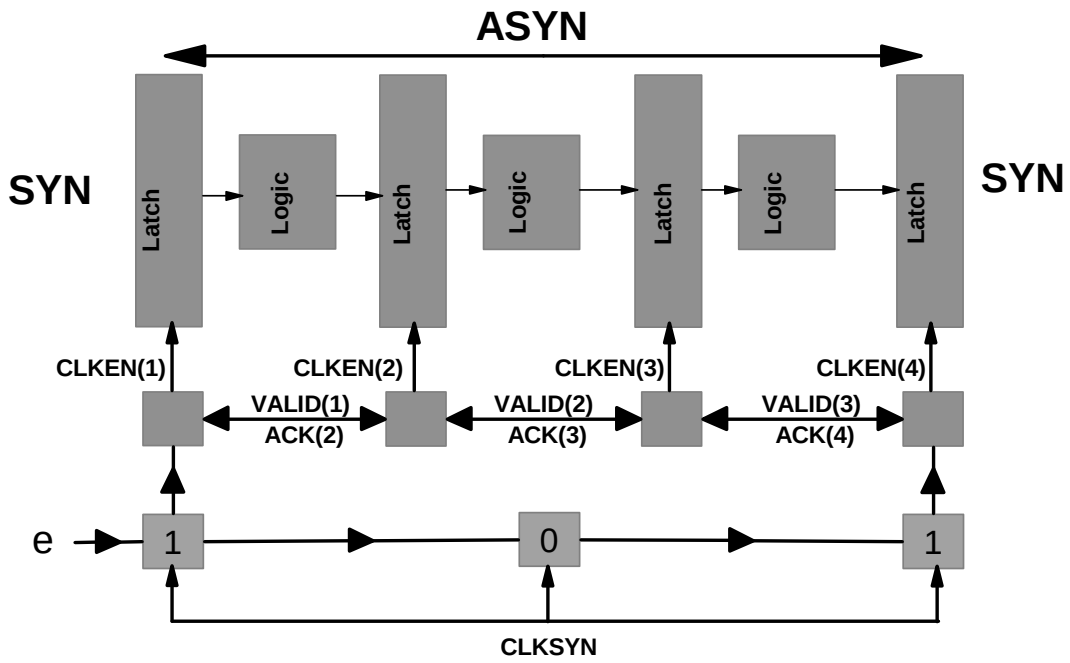


## *Syn to Asyn to Syn Interfaces*



## Power versus Switching Factor (assuming 70% of power in clocks and latches)

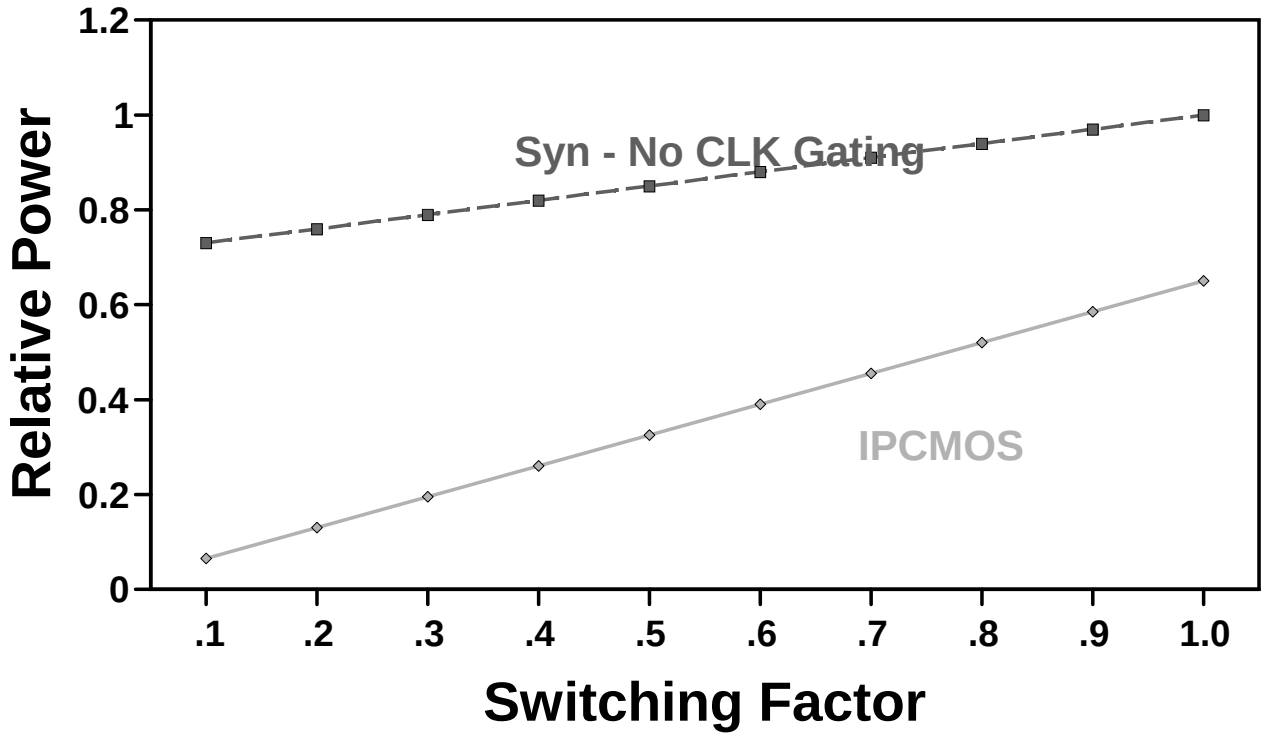

# Power and Noise

- IPCMOS offers significant power reduction
  - Clocks active only during an operation
  - Single stage register versus master/slave
- IPCMOS gives significant Ldi/dt noise reduction
  - Local clocks are staggered in time
  - Lower power results in lower current



# IPCMOS Testing

- Stuck fault and ac testing possible
- Timing for ac tests generated by macro
  - Do not need high speed external clocks
  - Simple scan chain
- Timing control in data valid signal gives additional ac test coverage

# Summary (IPCMOS)

- A syn to asyn to syn circuit technique suitable for multi-GHz operation has been developed
- Significant power reduction results from:
  - Enabling the local clocks only when data is valid
  - Reduced clock loading from the simplified latch structure
- Interlocked circuits deal with global timing issues
  - Delay variations from power supply noise
  - Delay variations from chip parametrics
- Helps alleviate the problem of stringent clock synchronization