LPX: *a low-power processor with a locally asynchronous execute pipe***

PRESENTED BY:

Alper Buyuktosunoglu and Pradip Bose,

IBM T. J. Watson Research Center

*** contributors*

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.

A. Buyuktosunoglu: summer intern from U of Rochester
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

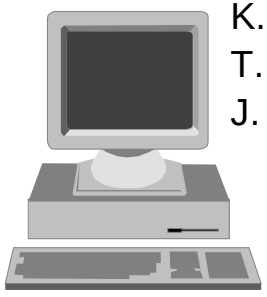

ICS-02 Tutorial T4, June 23, 2002, NY City.

LPX: High-Level Block Diagram

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

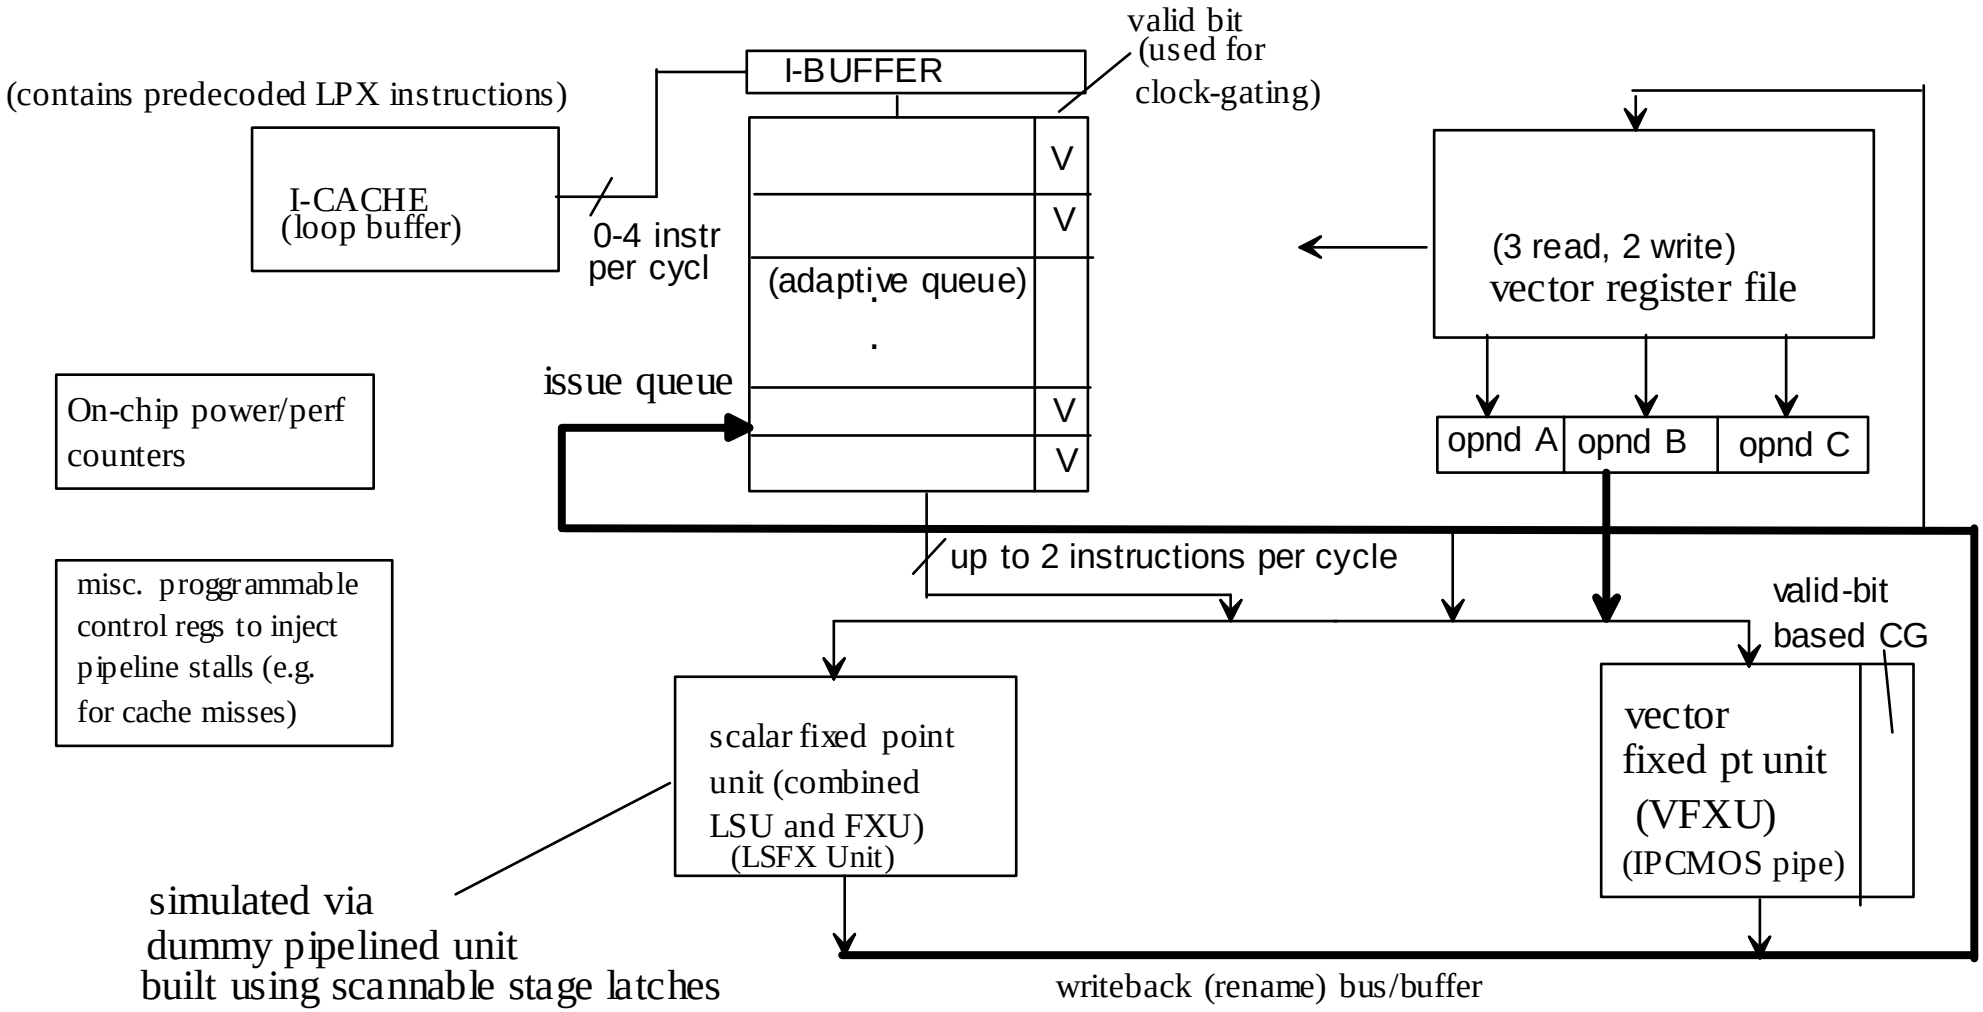

Key Design Ideas in LPX

- Microarchitecture-level
  - adaptive resource sizing: dynamic, application-driven
    - unified issue queue; register file
  - dynamic throttling of fetch/dispatch/issue/compl width
  - microarchitectural support for clock gating: coarse-grain, fine-grain..
  - on-chip counter architecture for power/performance monitoring and feedback for adaptive control
- Logic/circuit-level
  - adaptive clocking: save power, maintain performance
    - locally generated, enabled by data (IPCMOS)
    - synchronous clock-gating: coarse and fine-grain
  - power-efficient CAM/RAM structures

note: not all features above will be implemented; but will be studied in high-level design phase (simulation-based)

Methodology for LPX Definition and Tuning

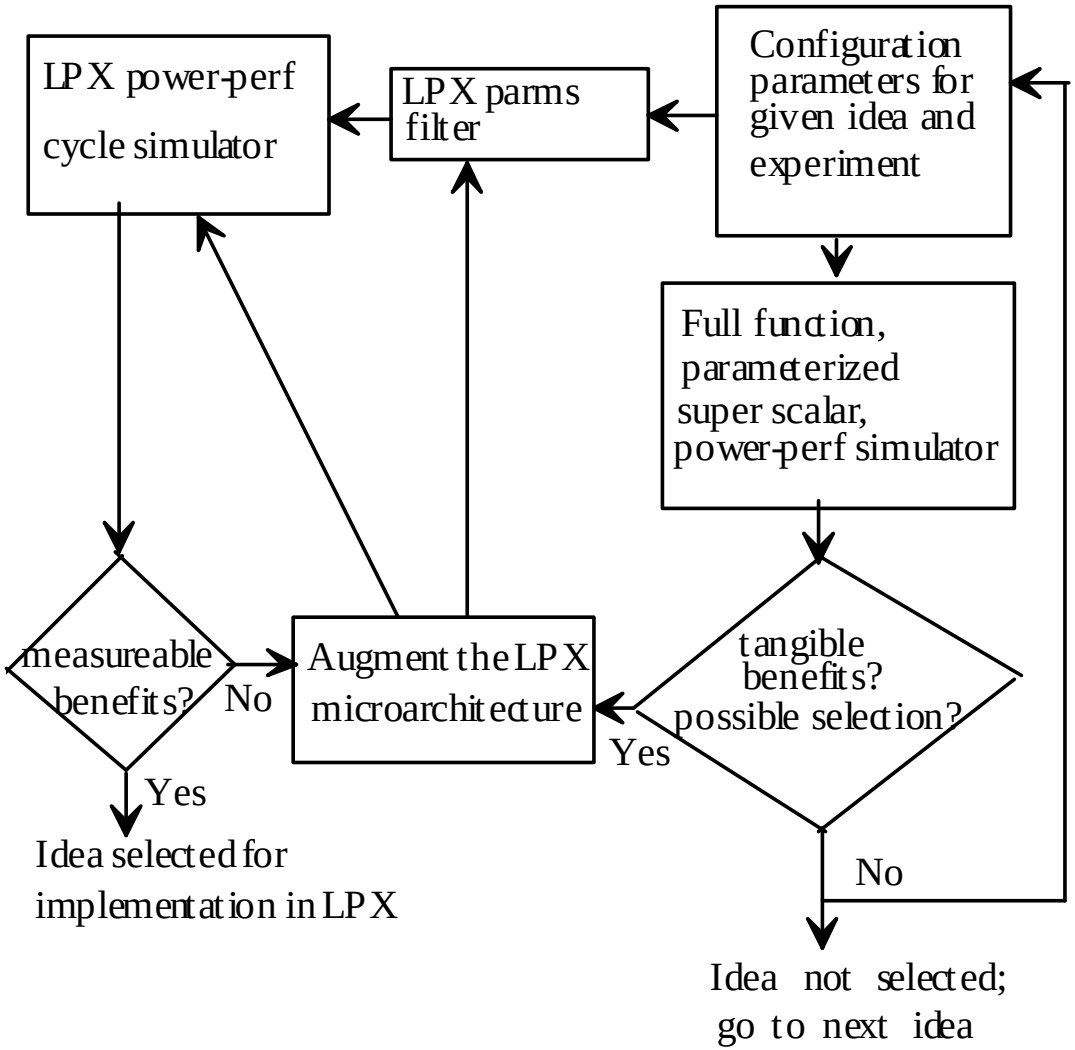

- Idea selected must have potential power-perf "worth" based on:
  - simulation "in the large".
- Basic idea is adapted to a much simpler hardware heuristic, appropriate for LPX:
  - simulation and prototyping "in the small".
- Extensive circuit simulation to validate power savings and overhead:
  -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

Issue Queue Design Choices

(Simulation in the large for adaptive issueQ)

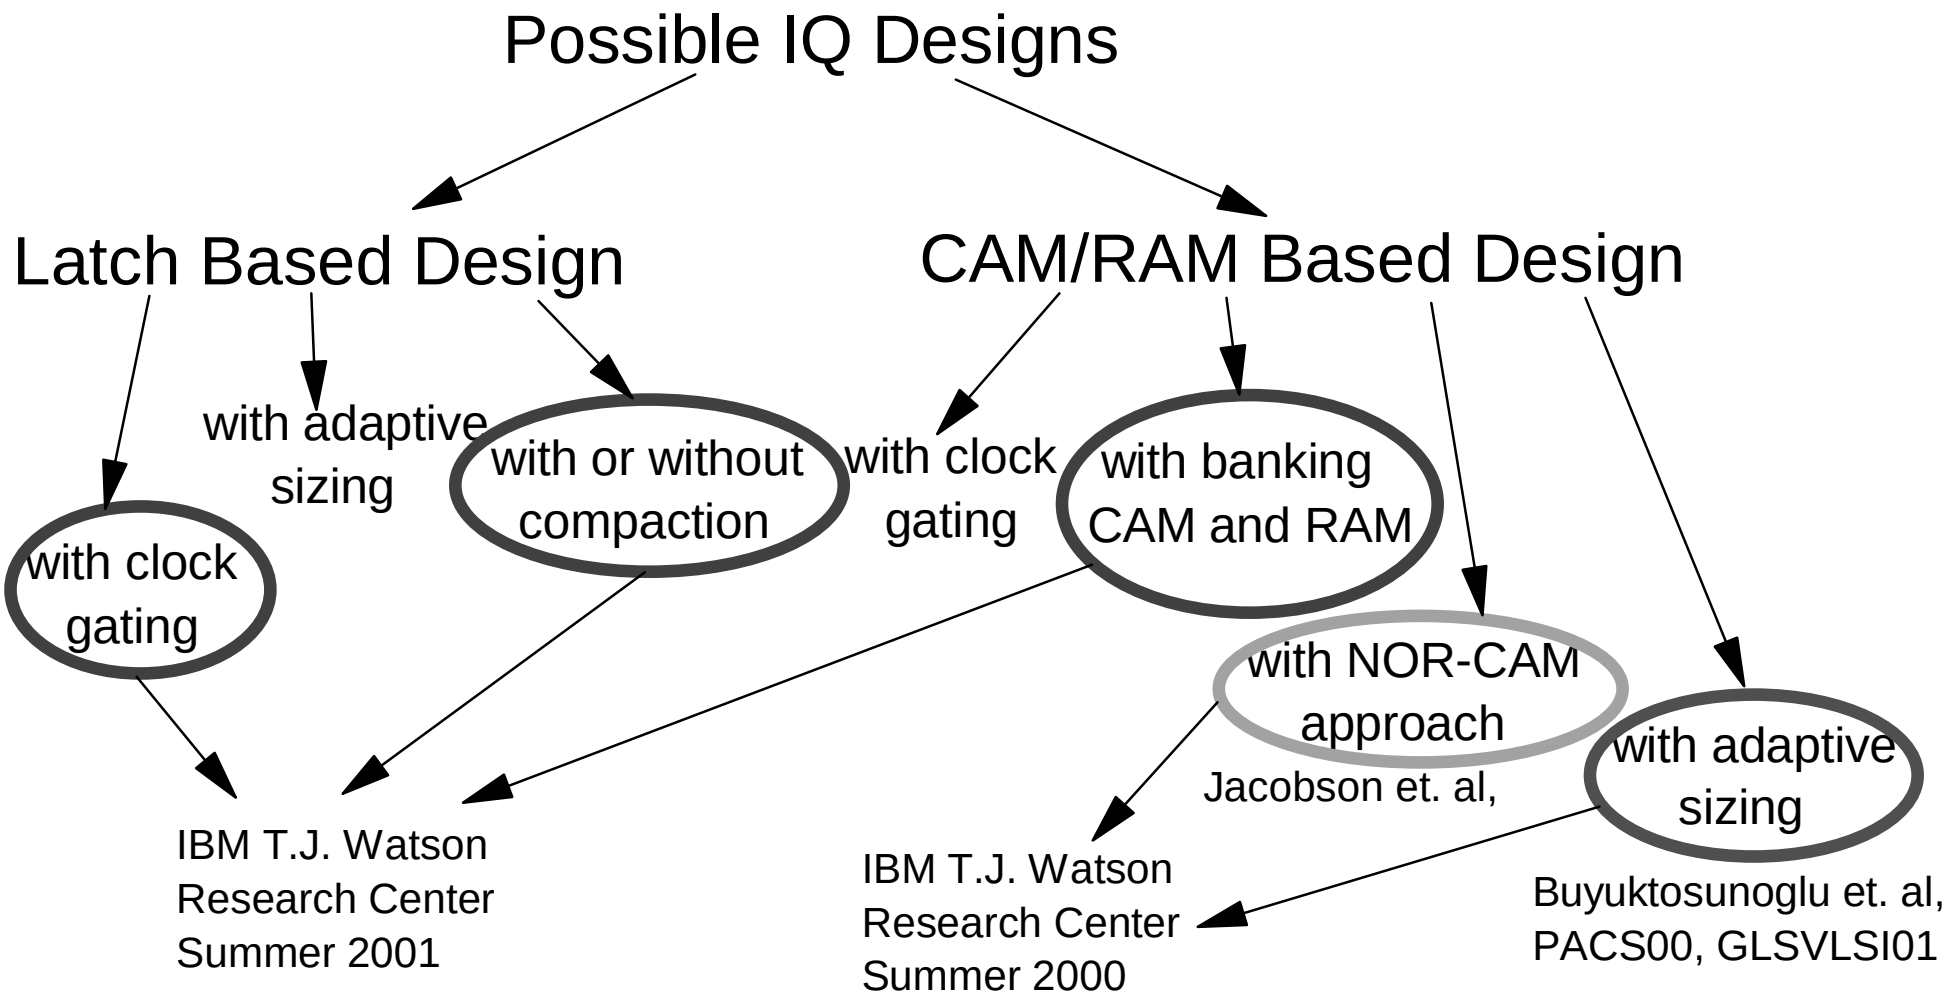

Power-efficient issue queue research: Buyuktosunoglu et al.

as in POWER4 issue queue (baseline)

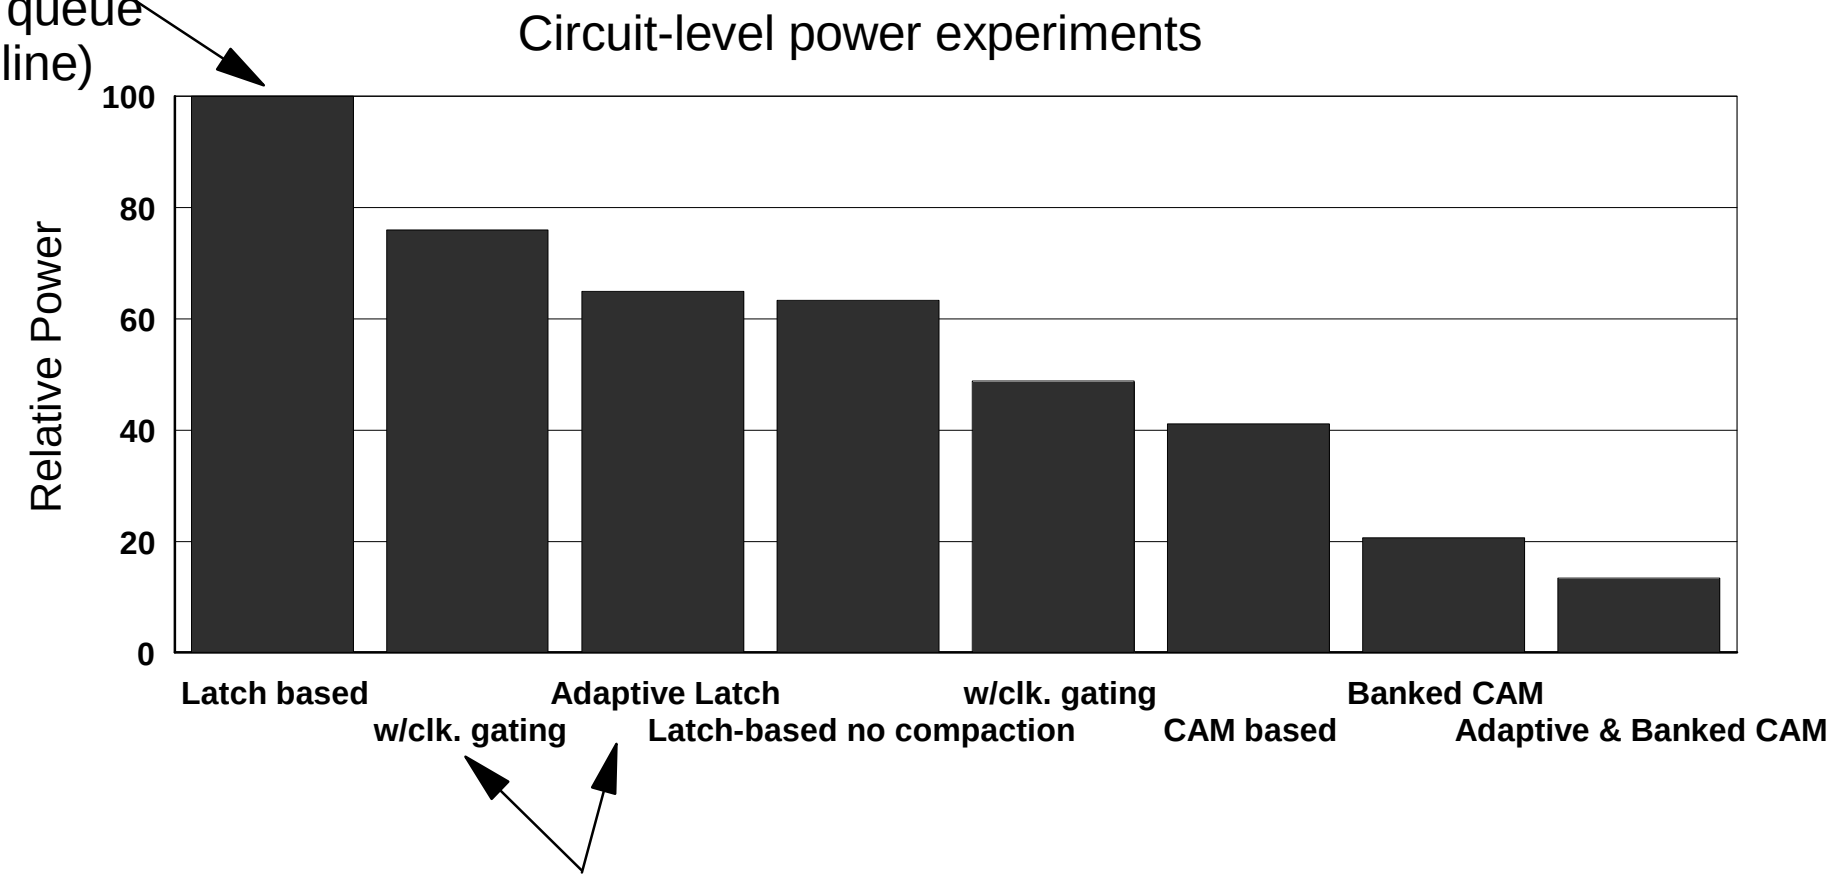

Adaptive Issue Queue in LPX

vect_add loop trace on LPX

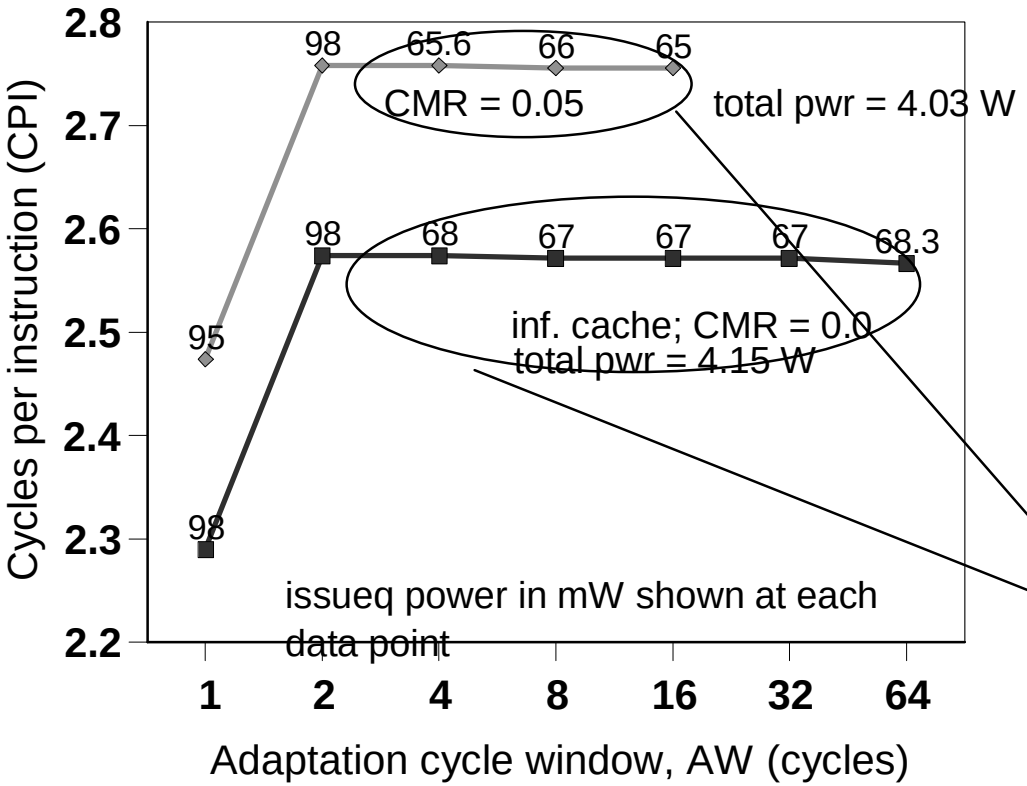

Adaptation Heuristic

```
if (current-cycle-window-issuecount
    <  0.5 * last_cycle_window_issuecount)
then
     increase_size (* if possible*);

else decrease_size (* if possible *);
```

issueq power reduction ~ 51 %
CPI degradation ~ 12.4 %

baseline: OO, non-adaptive; CPI (inf. cache) = 2.29; CPI (finite cache) = 2.45
total power = 4.46 W; IssQpwr = 137 mW

Power-Efficient Ifetch

- Pipeline Gating via Speculation Control
  - S. Manne, D. Grunwald et al. 1998 (ISCA)

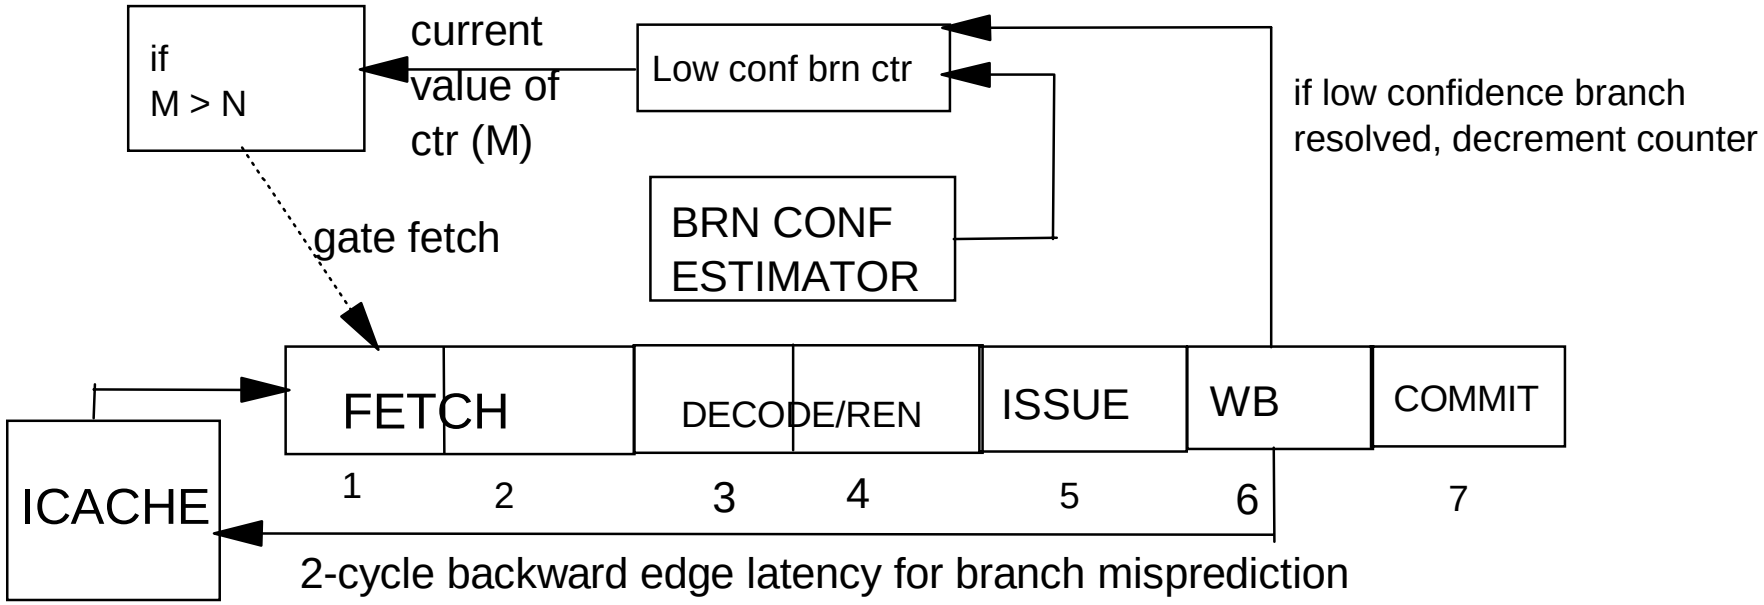

- Up to 38 % reduction in speculative waste, with ~1 % perf. loss
  - **but.... high complexity/overhead.**
- For LPX, we experimented with simpler ifetch throttling control
  - ✓ e.g., stall-based: IBUFFER stall; ifetch throttled GW cycles after stall signal clears.

Conditional Ifetch in LPX

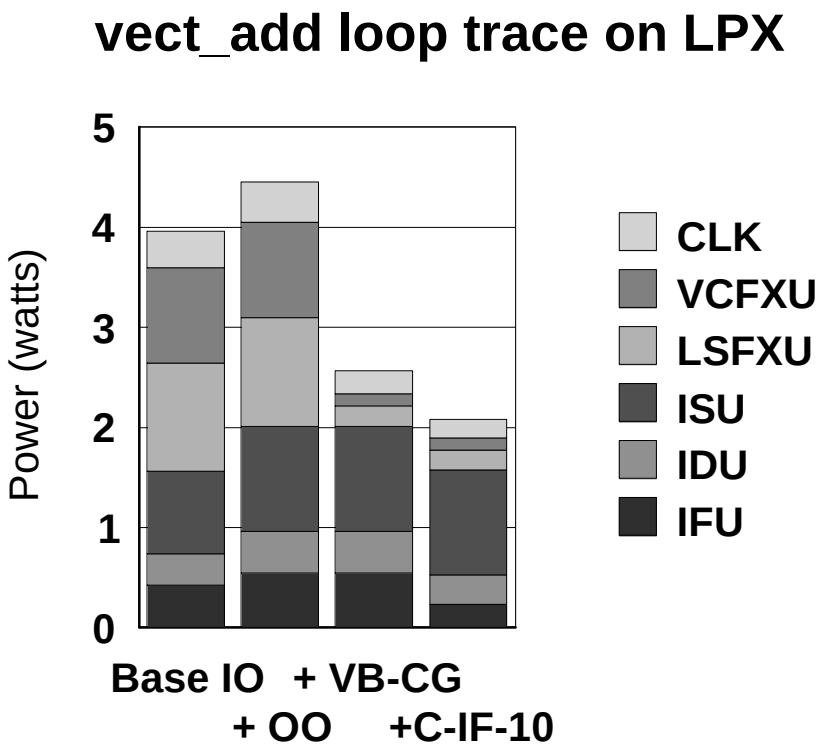

CPI of baseline in-order (IO) = 3.00
CPI of all the other out-of-order (OO) = 2.29

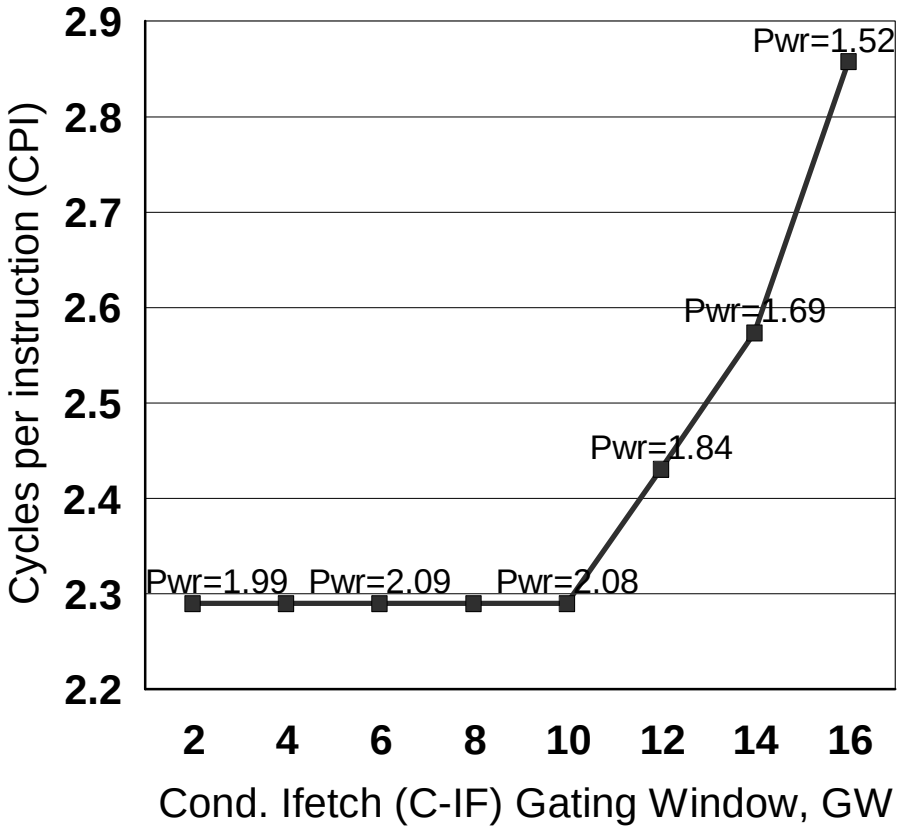

Other simple hardware throttle heuristics: being reported separately, *T. Karkhanis et al., 2002*

Interlocked Synchronous Pipelines (ISP)

PIPELINE STRUCTURE

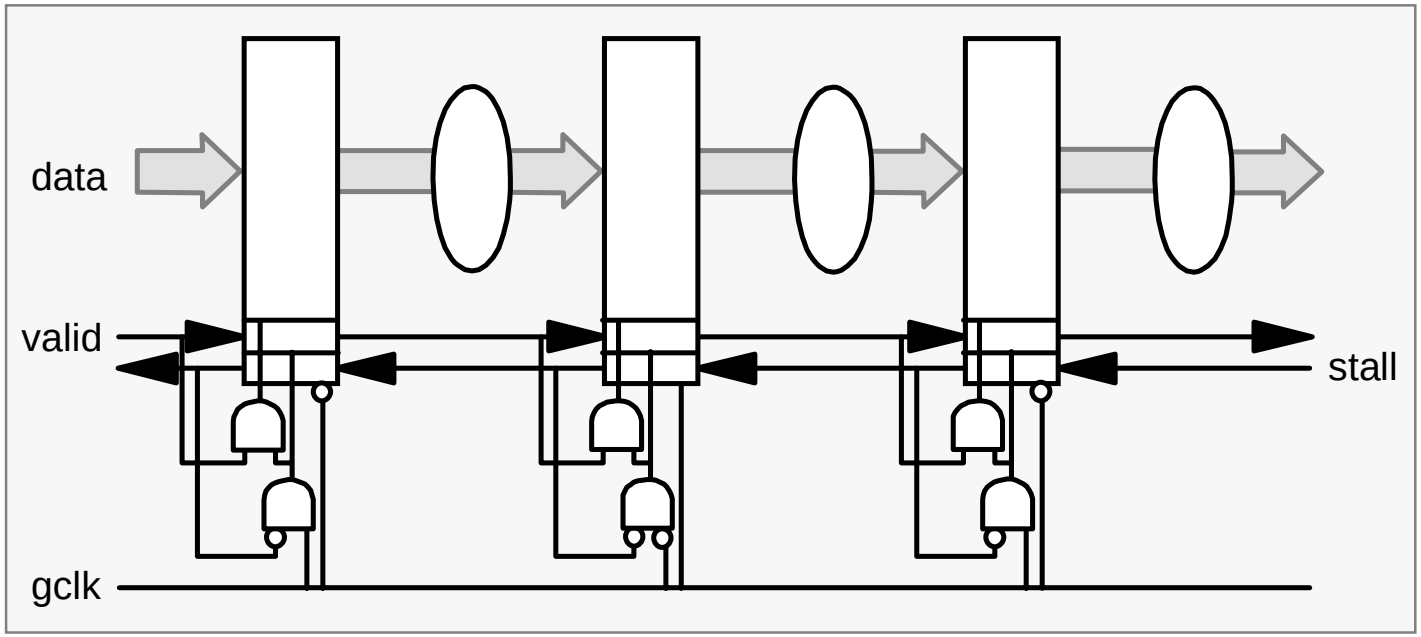

(Paper accepted for publication in ASYNC'2002:
H. Jacobson et al.)

- 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

LPX Design Schedule (2001-02)

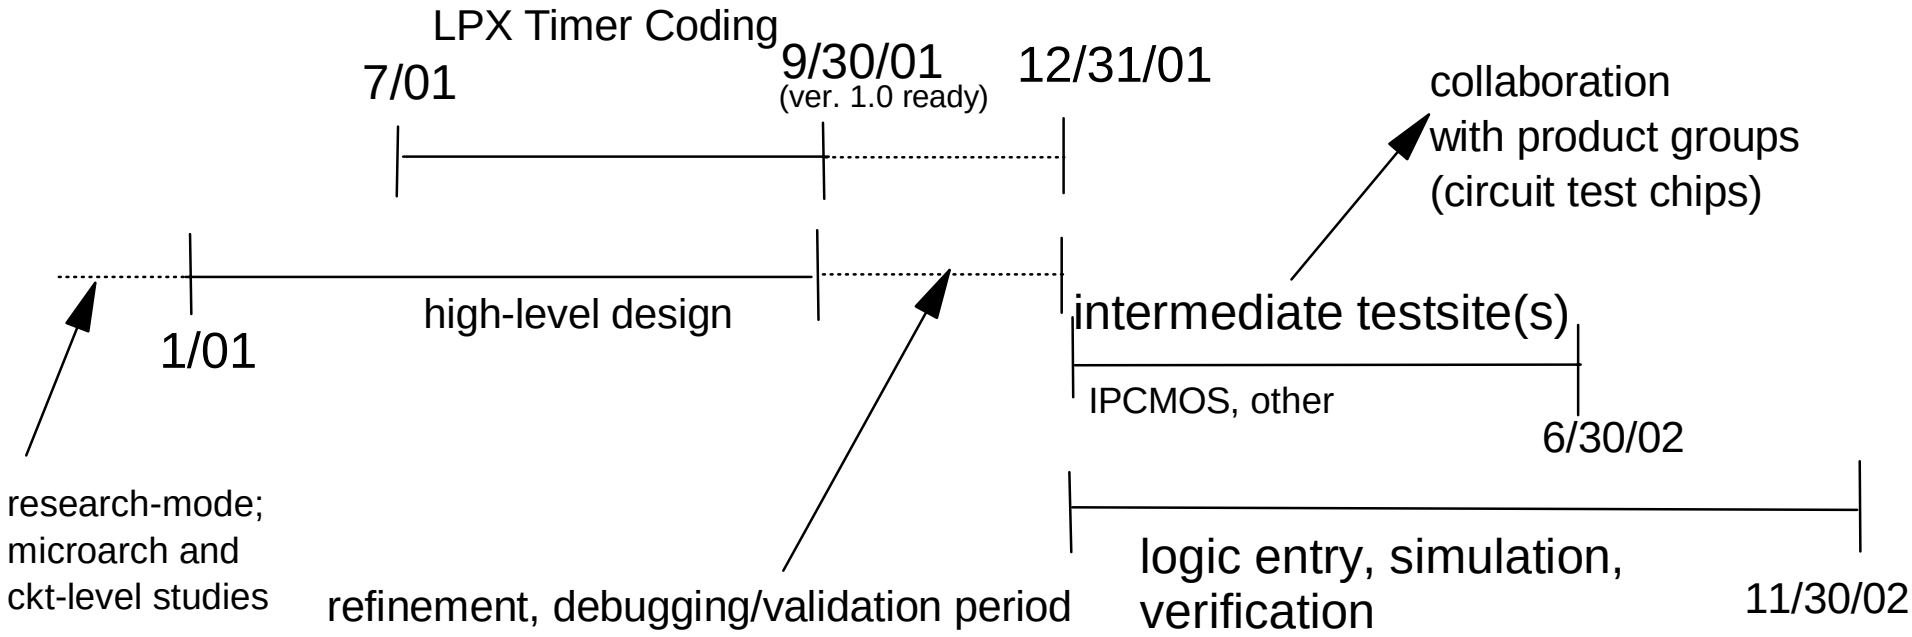

Power versus Switching Factor

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

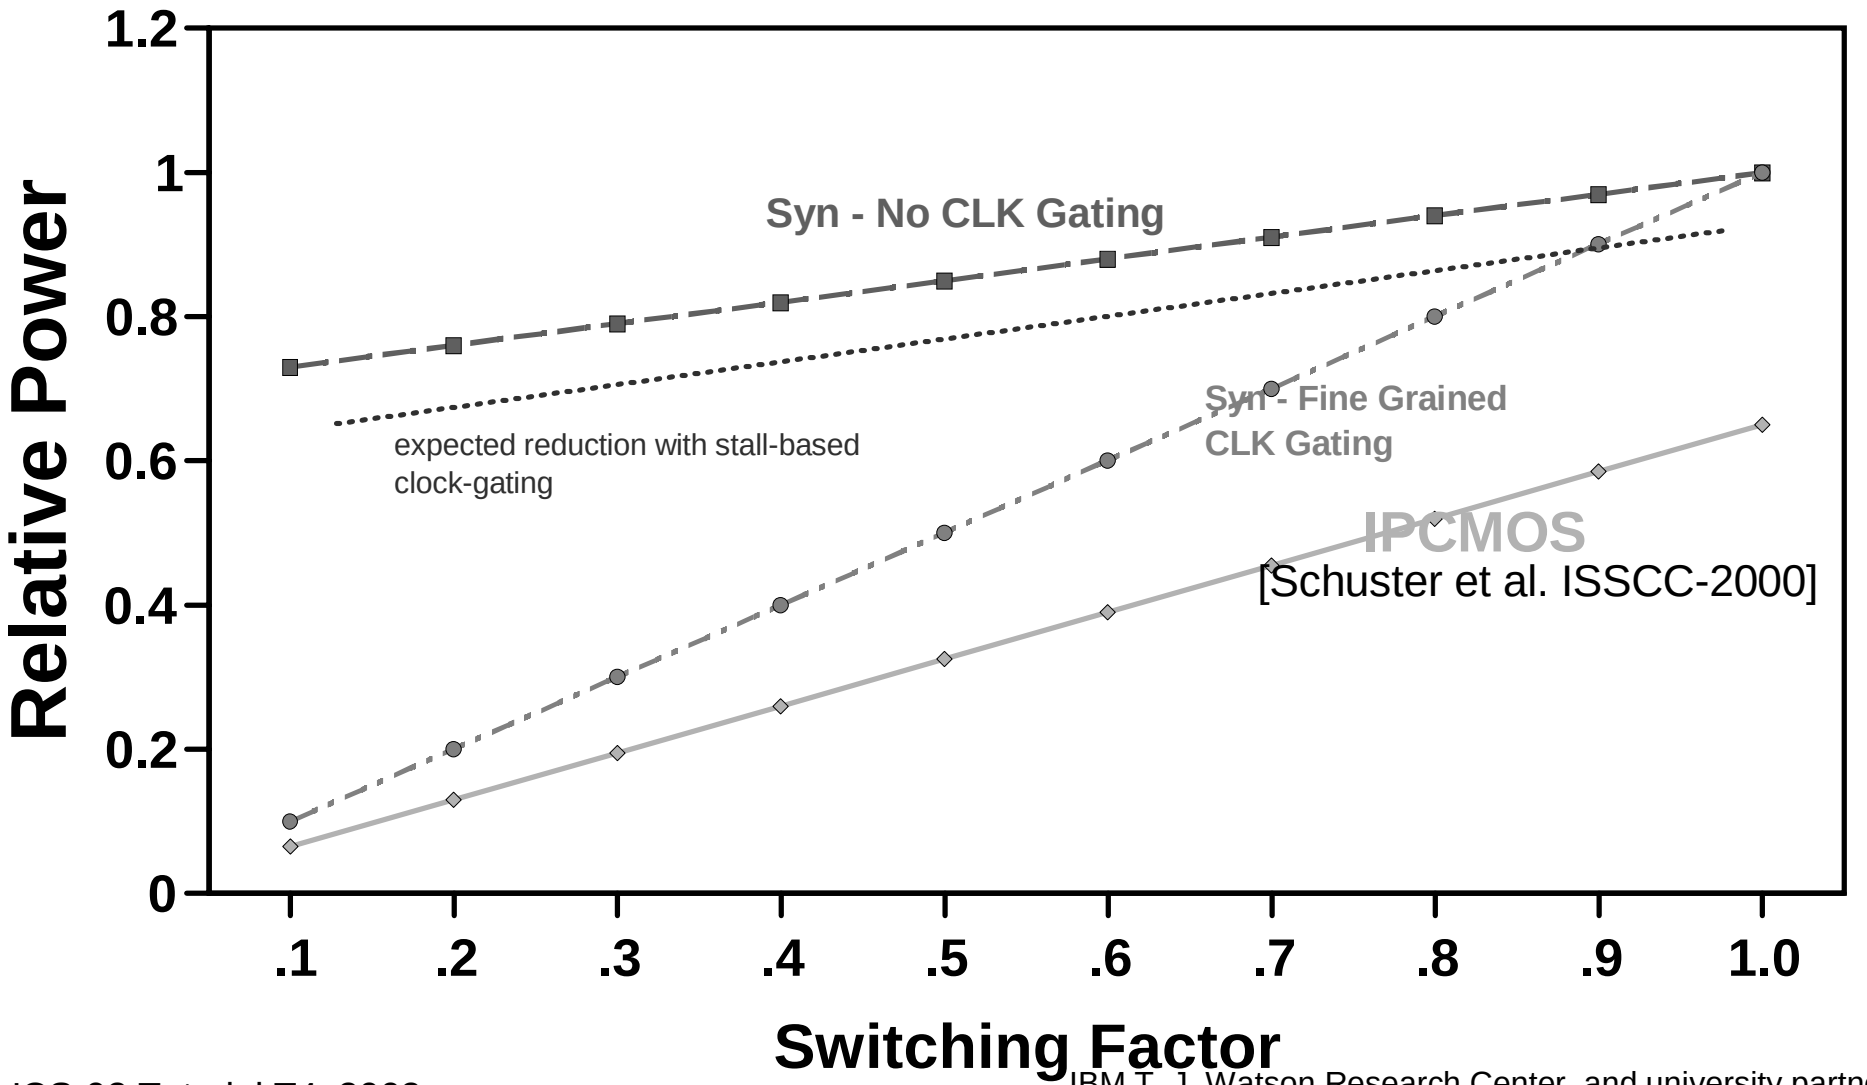

Conclusion

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 early 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

Stan Schuster, Pete Cook

IBM Research Center
Yorktown Heights, NY

(presented earlier at ISCA WCED-02 conference, June 2002)

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

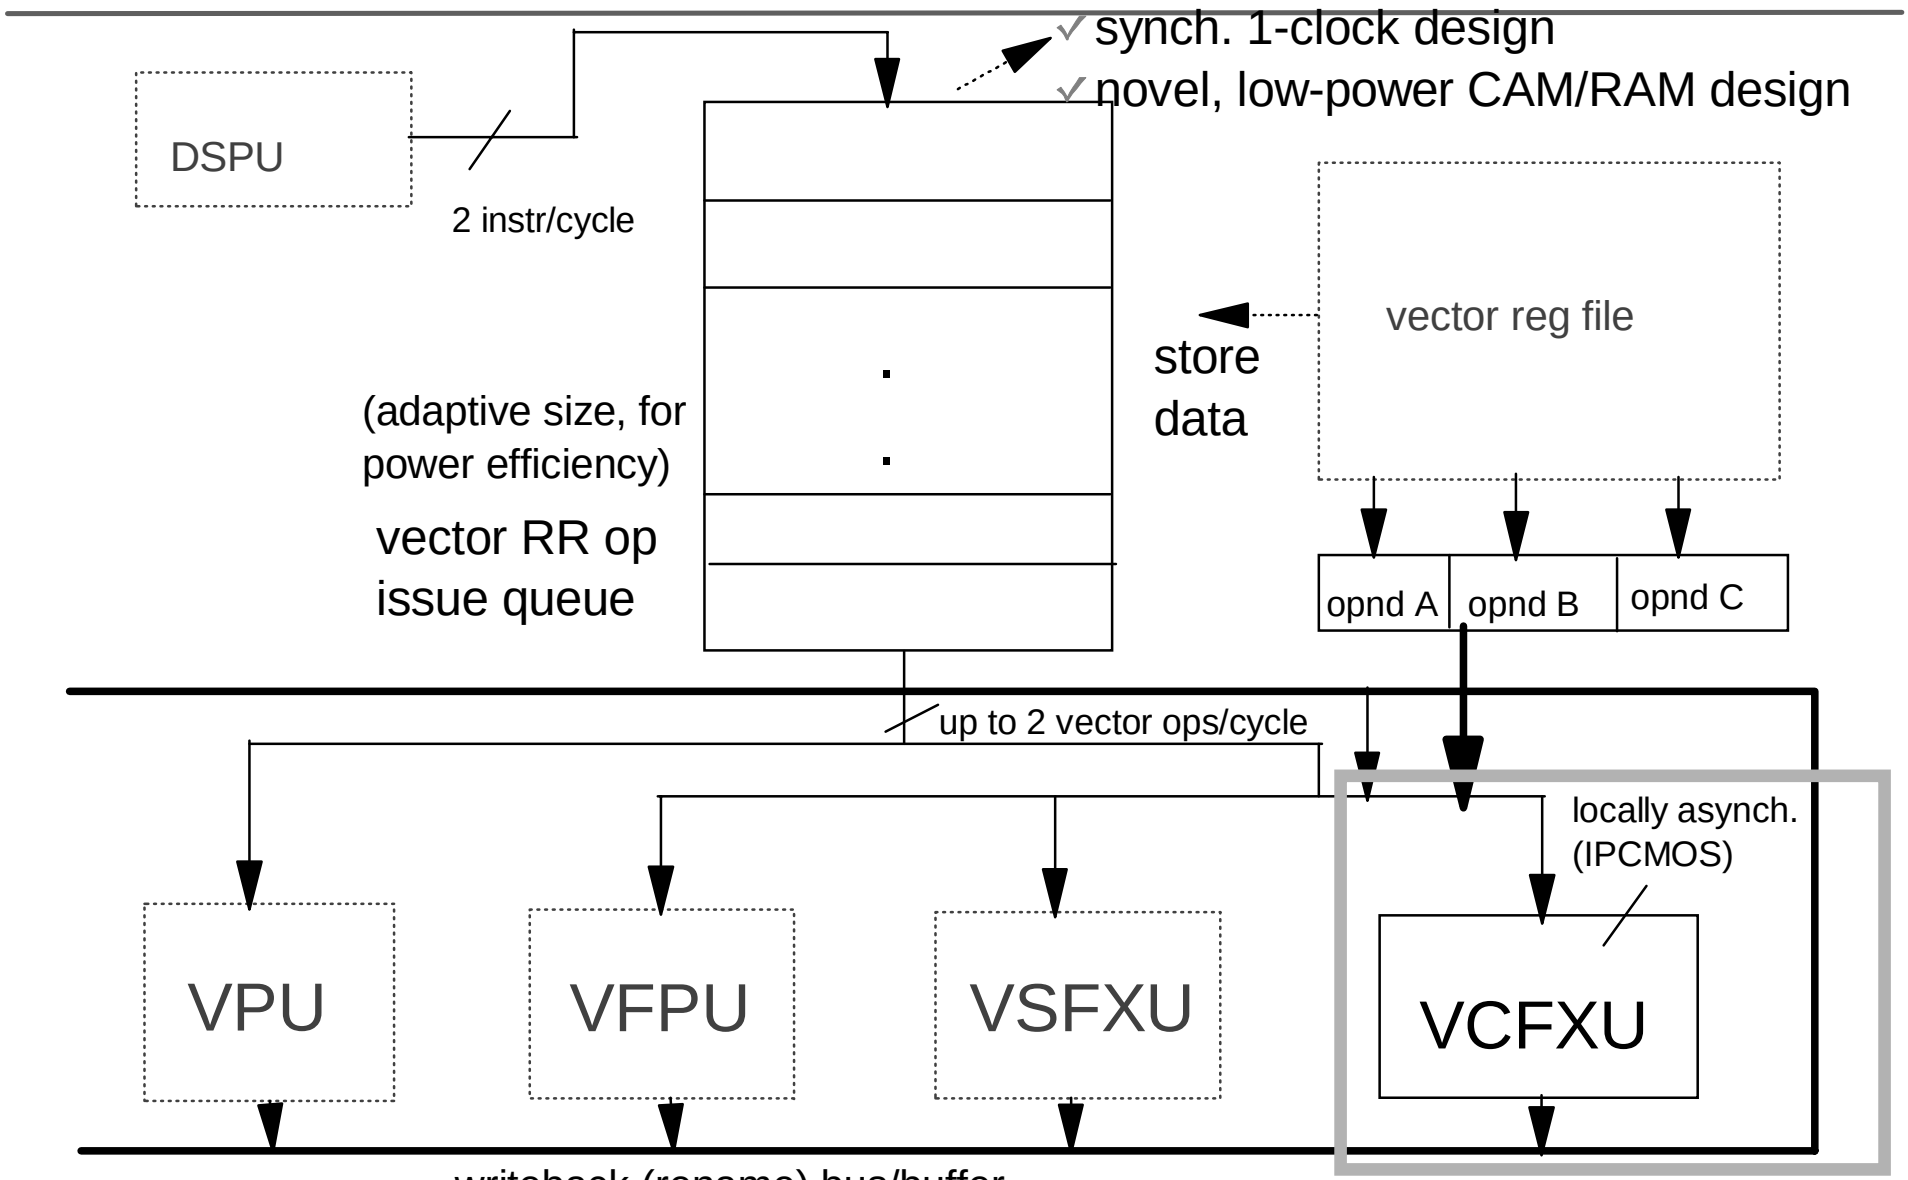

IBM T. J. Watson Research Center and university partners
IPCMOS section of Tutorial, Schuster and Cook, IBM

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

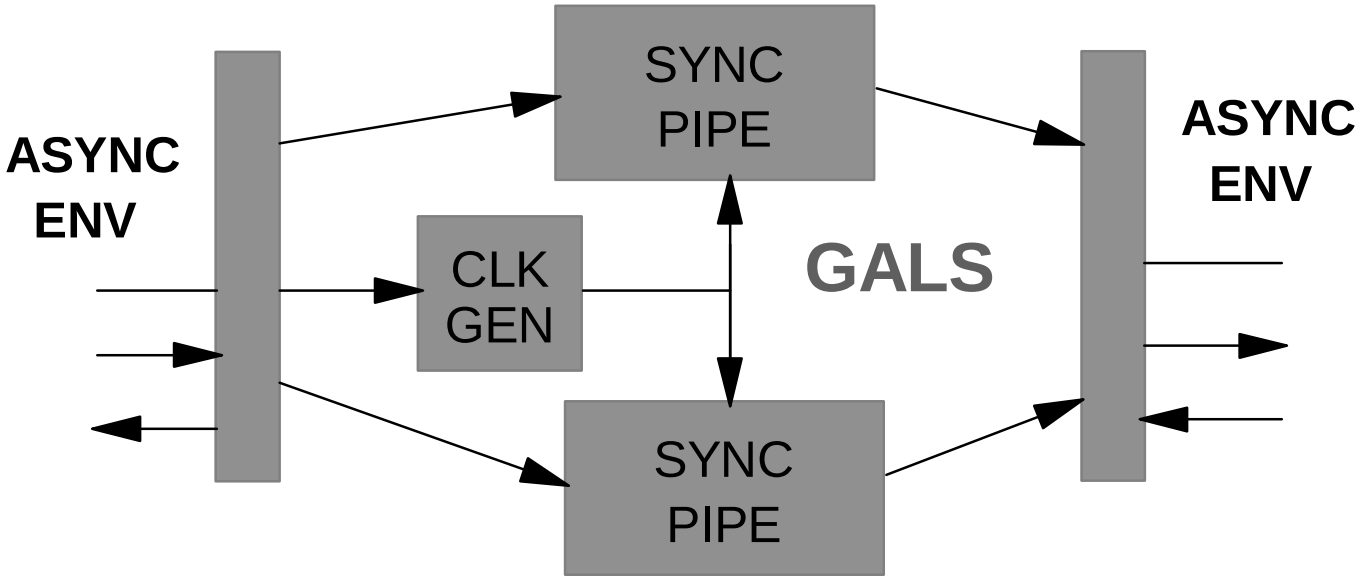
ASYNC
ENV
SYNC
PIPE
CLK
GEN
GALS
SYNC
PIPE
ASYNC
ENV

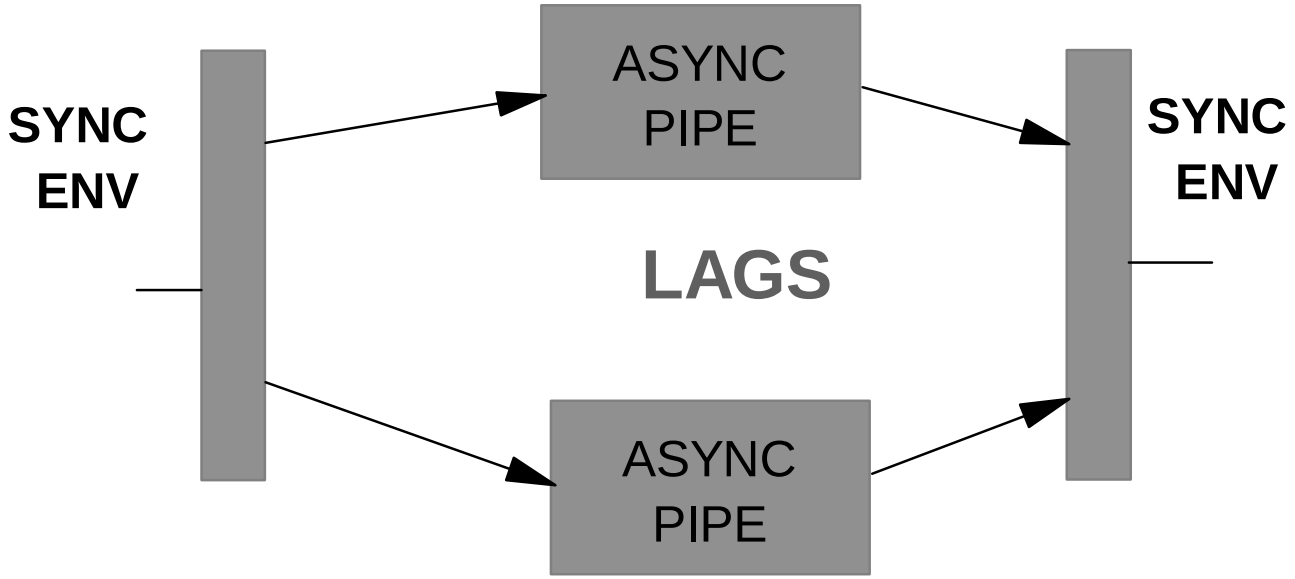
SYNC
ENV
ASYNC
PIPE
LAGS
ASYNC
PIPE
SYNC
ENV

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

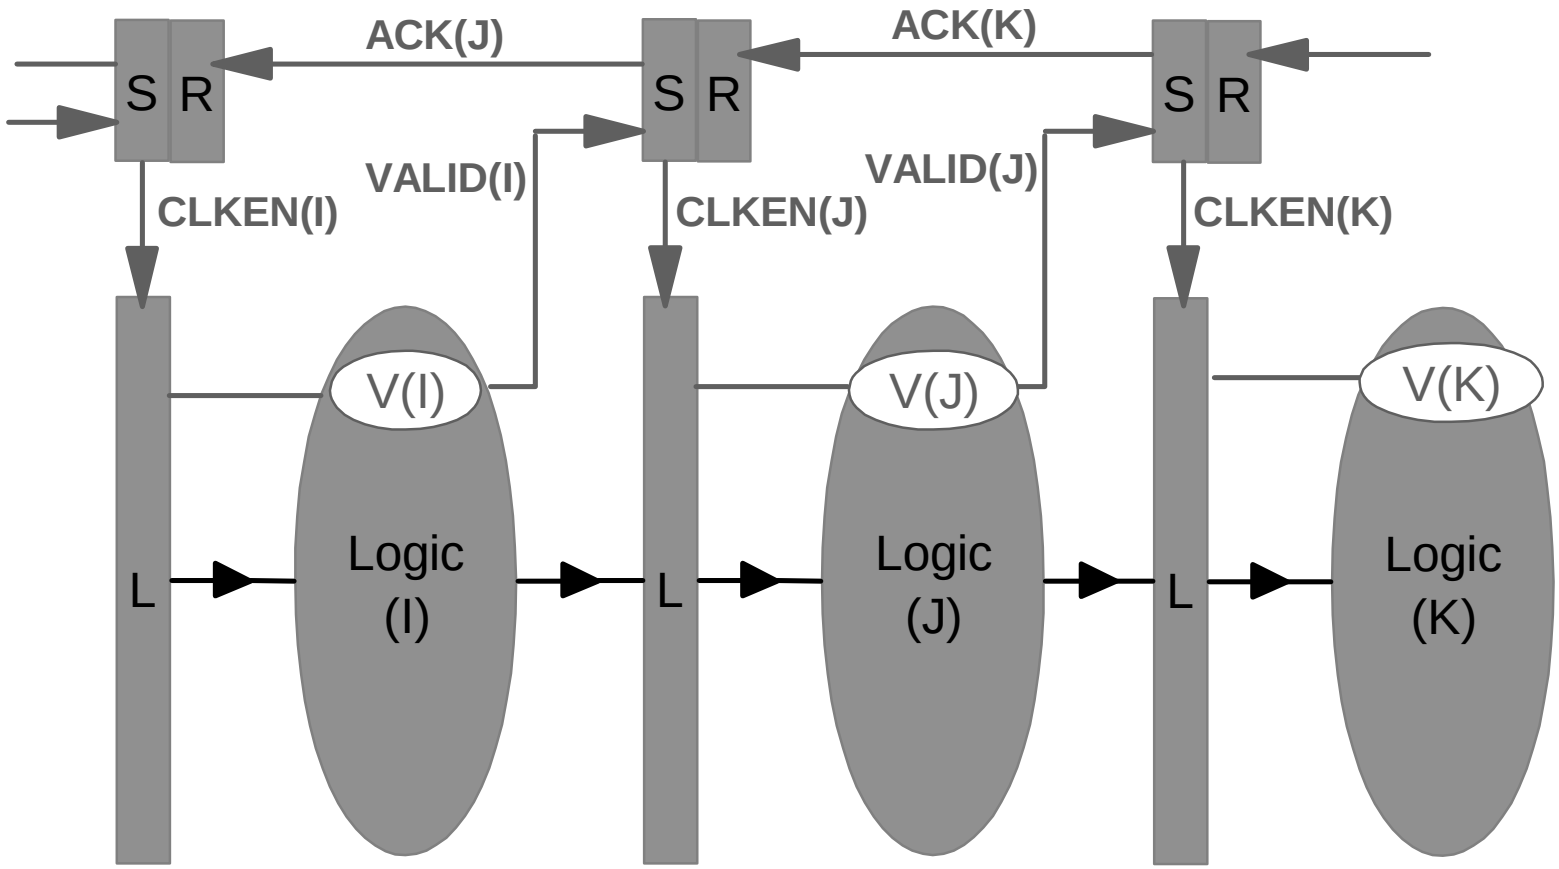

IPCMOS Interlocking

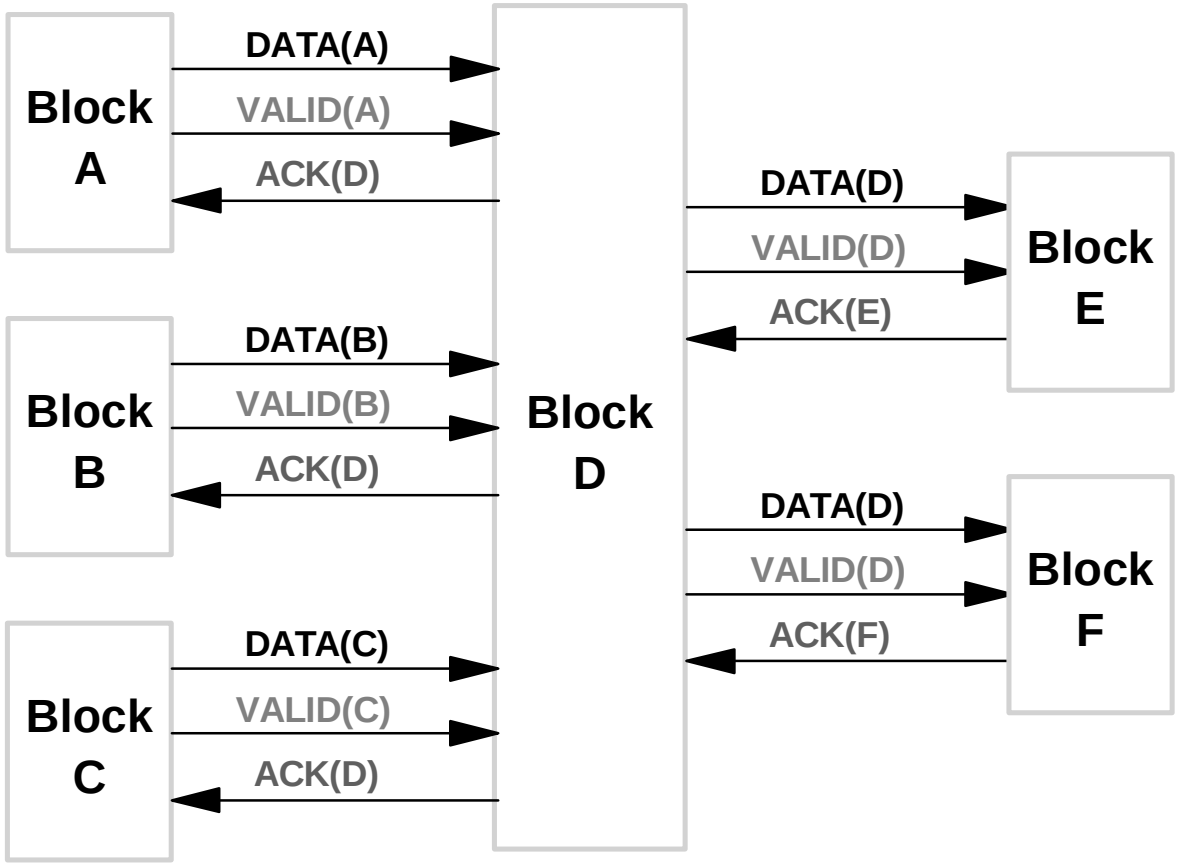

Interlocking at the Block Level

IPCMOS Circuits

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

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

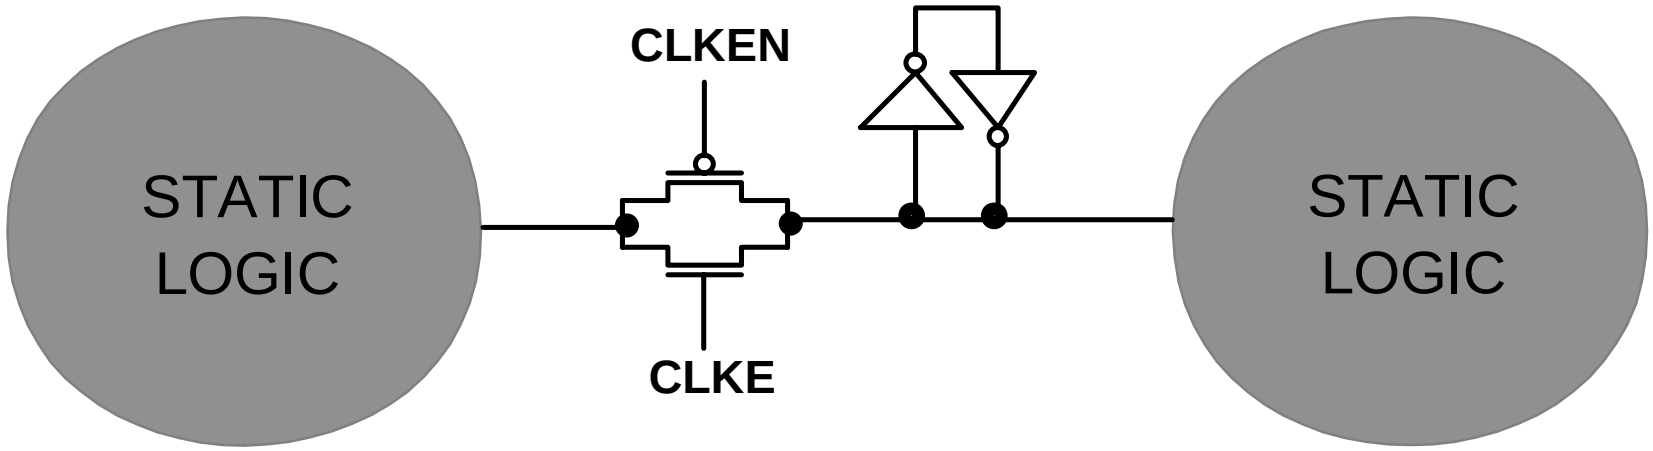

Latch for Static Logic

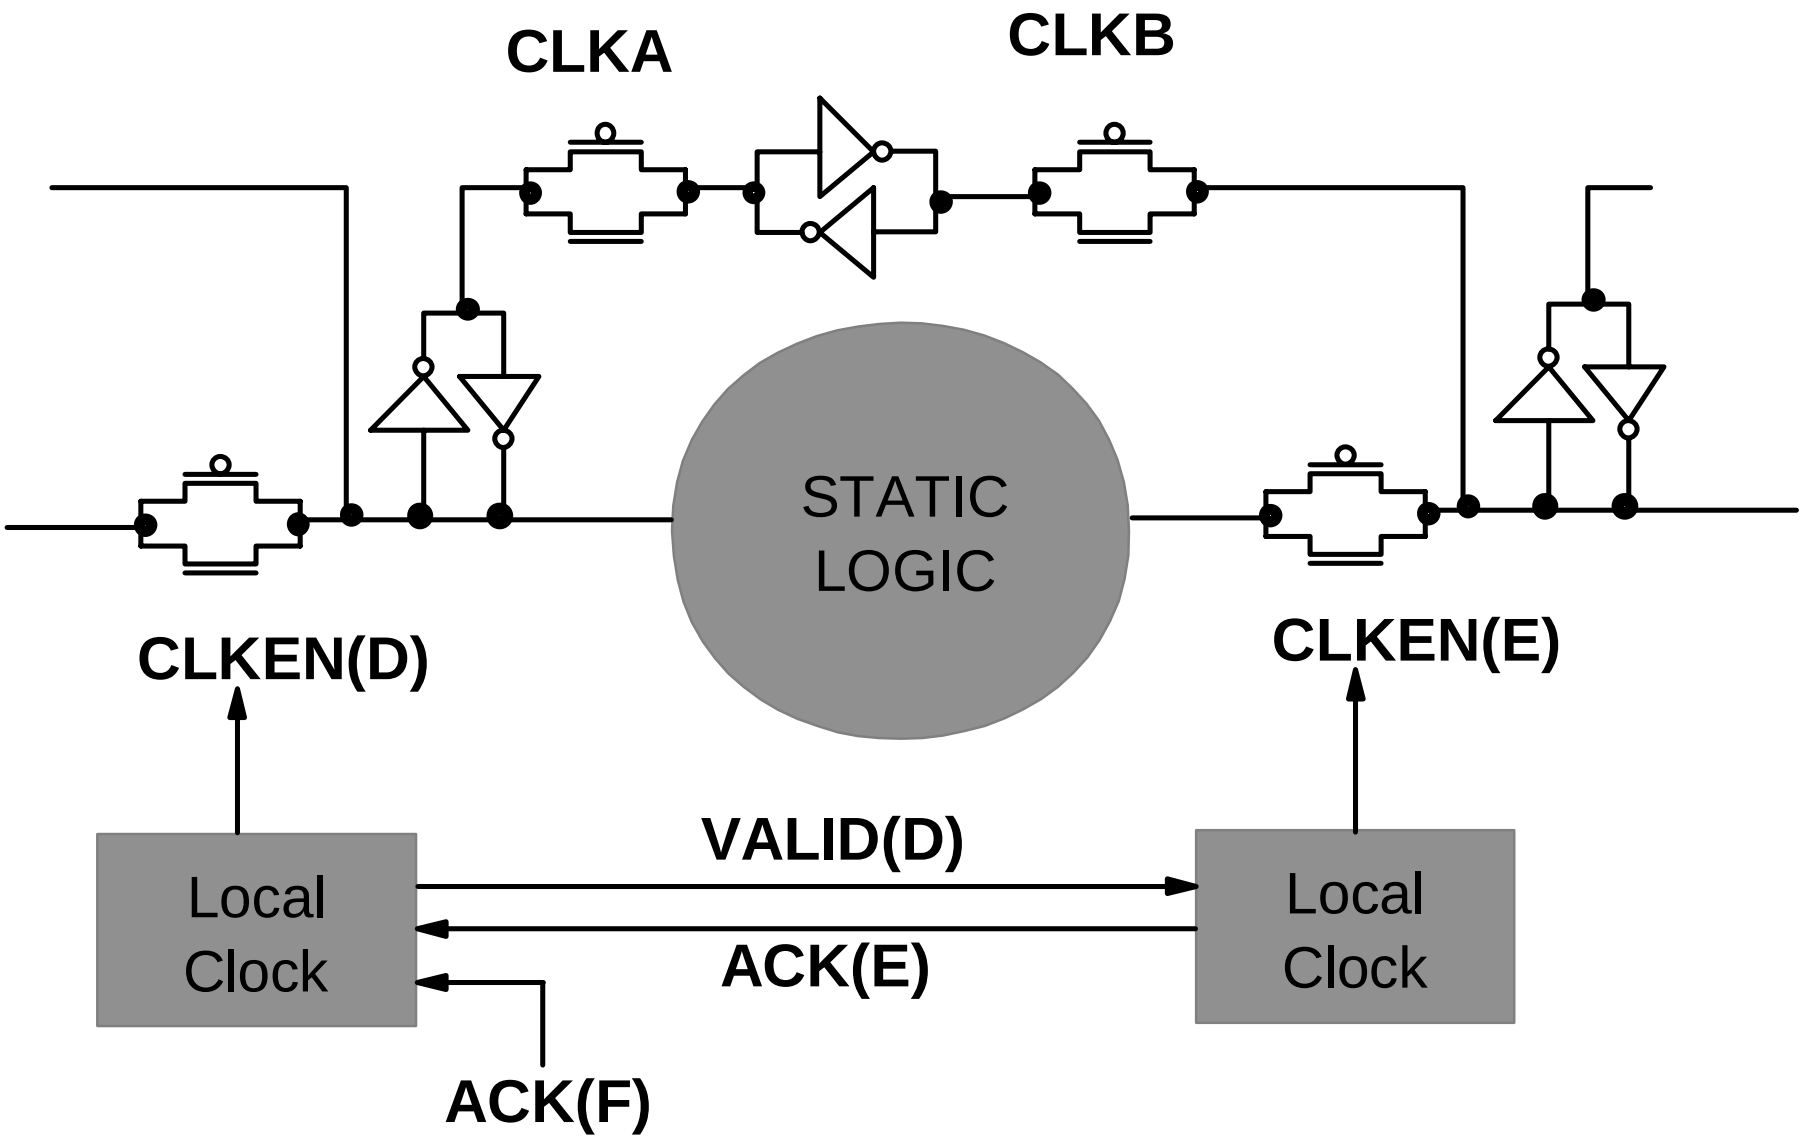

Testing of IPCMOS with Static Logic

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

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

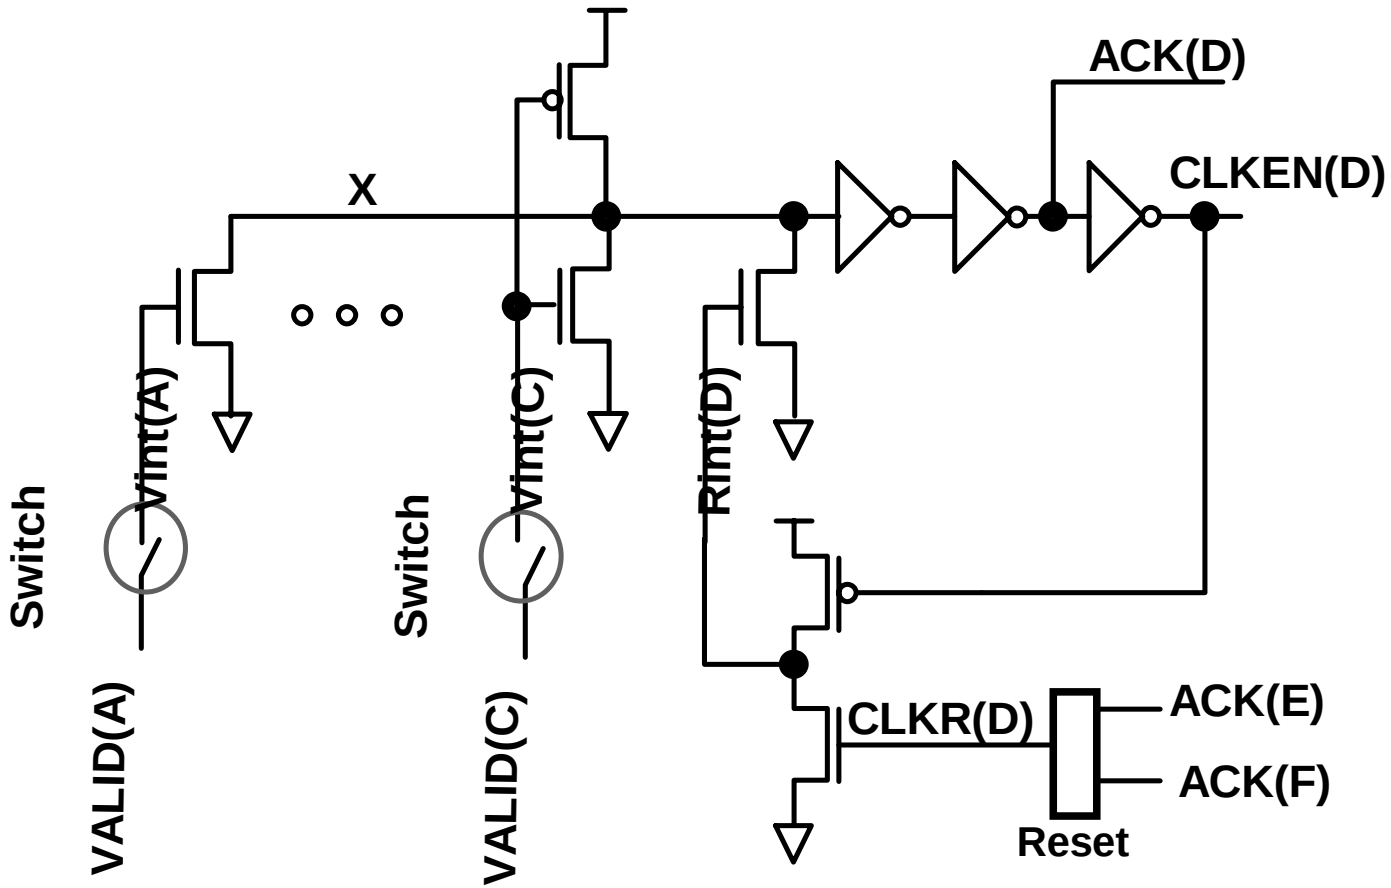

Local Clock Strobe Circuit

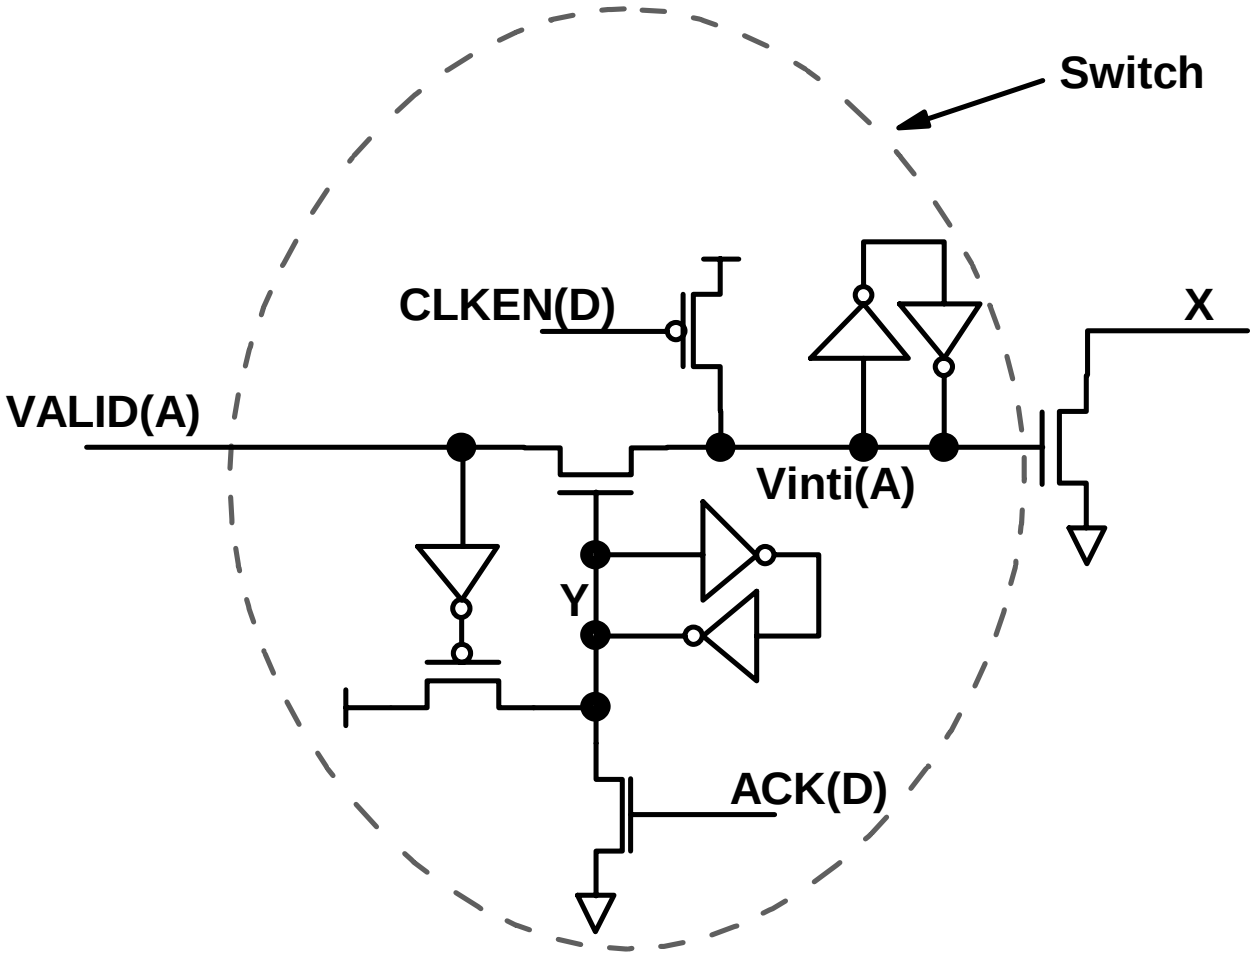

Strobe Circuit Switch

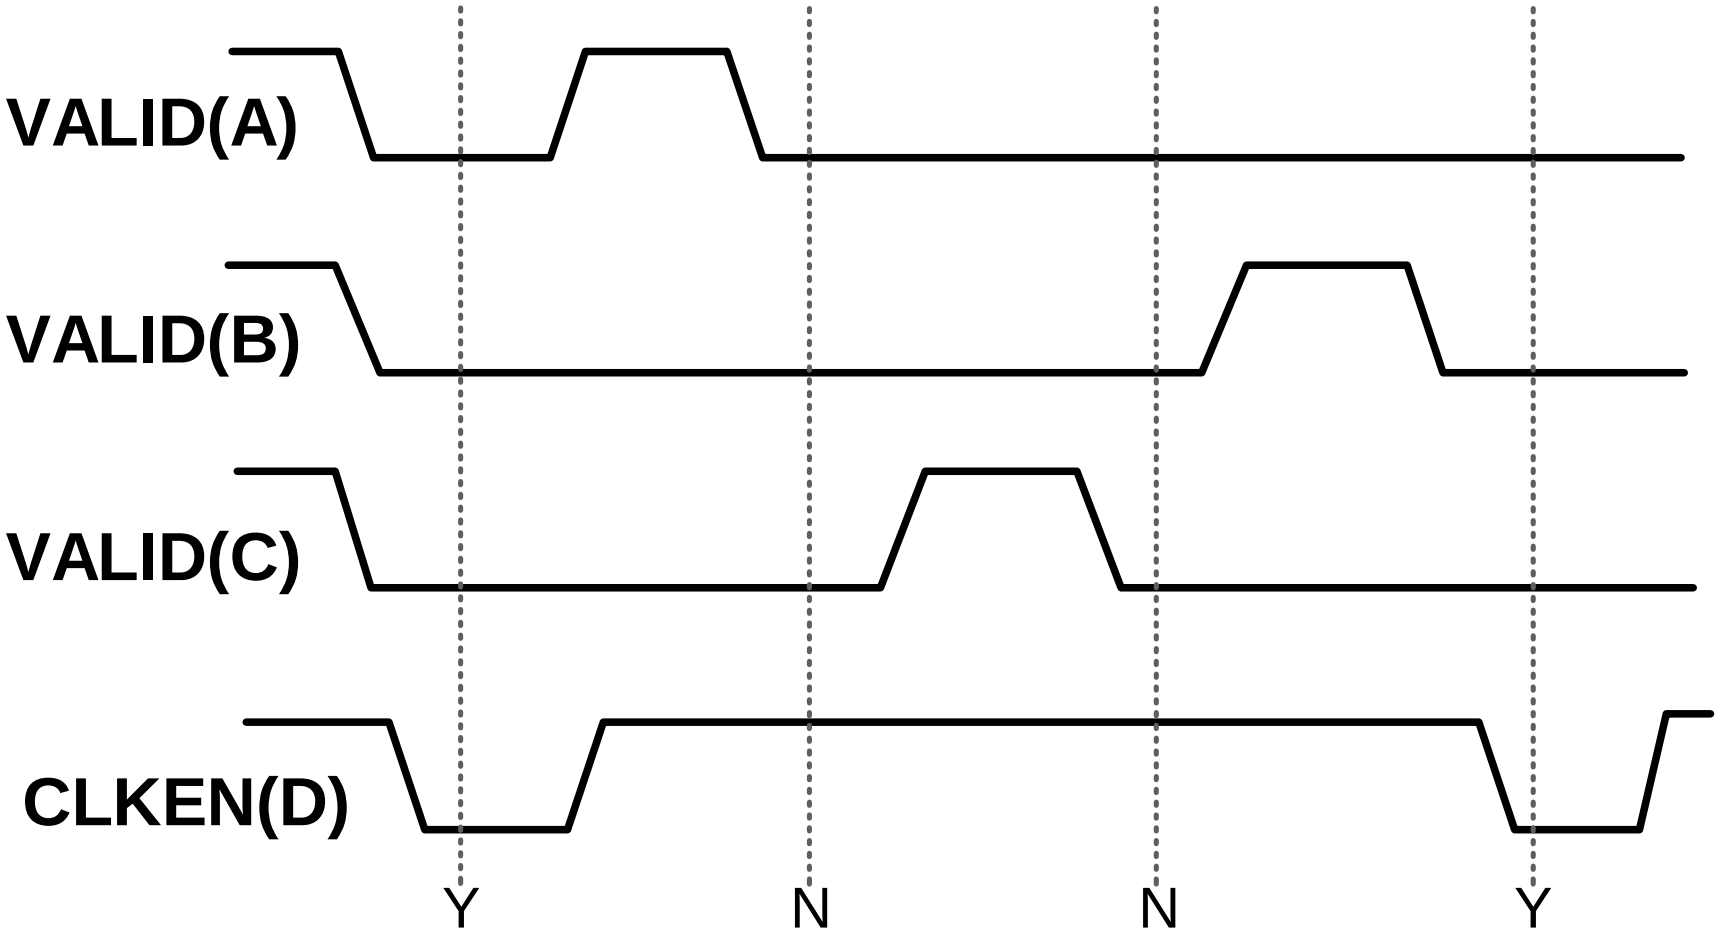

Strobe Circuit Unique "AND" Function

Strobe Delay versus # of Inputs

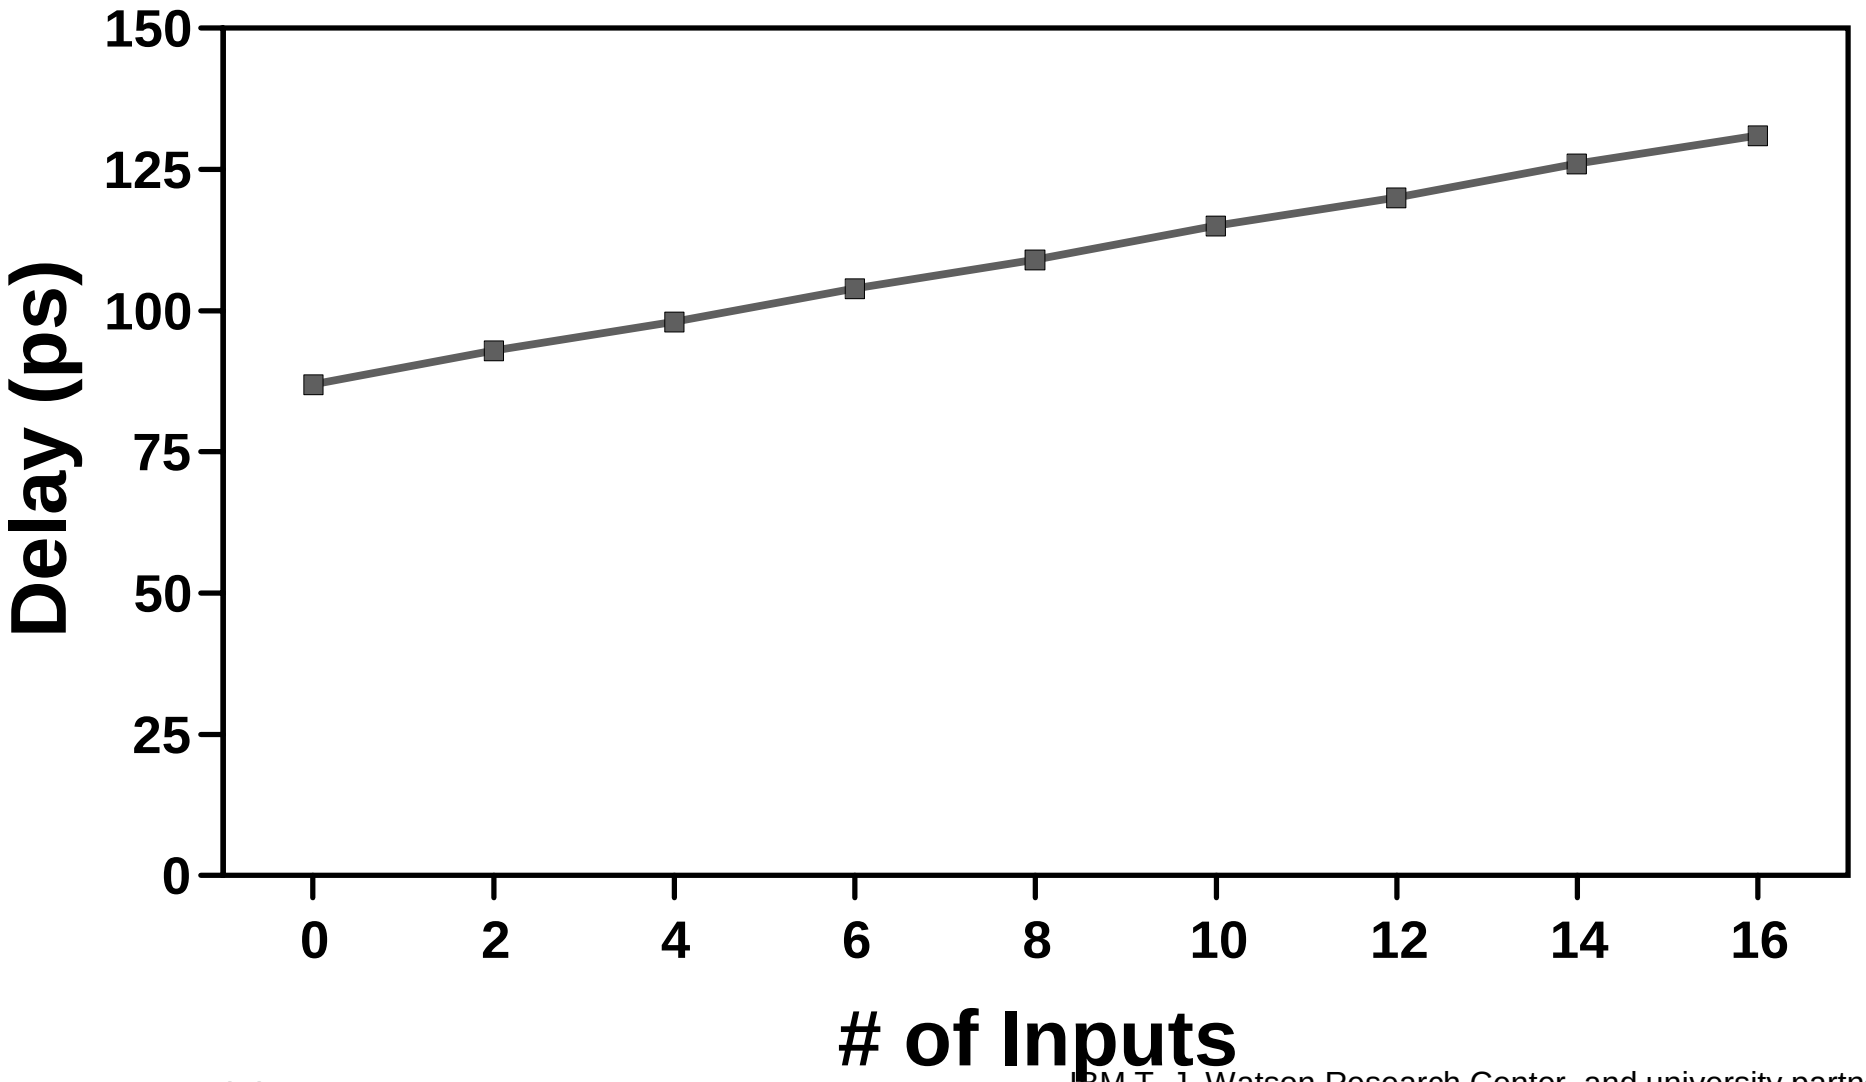

Data Valid

- 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

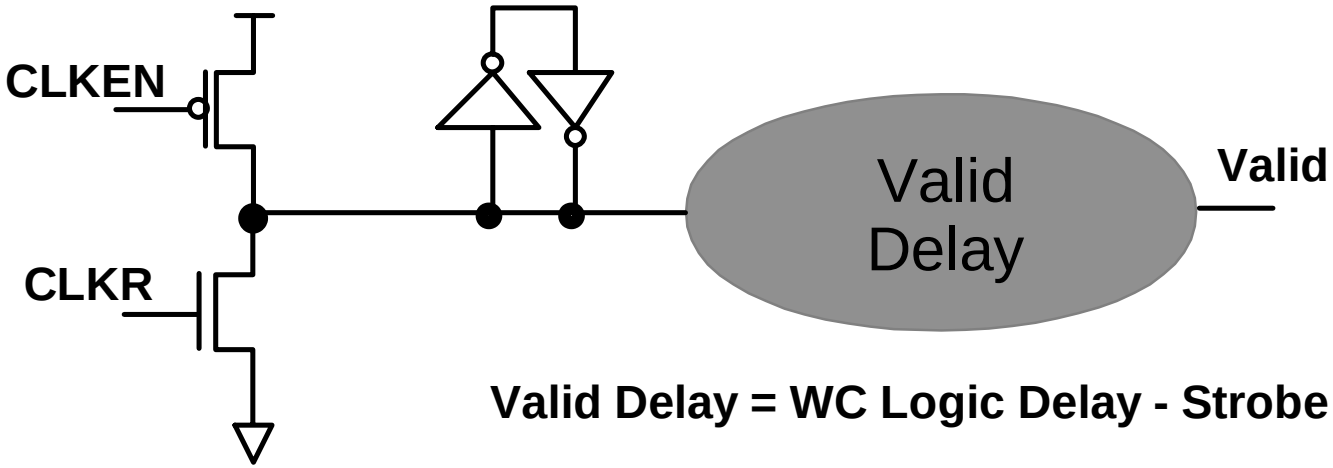

Valid Delay = WC Logic Delay - Strobe Delay

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

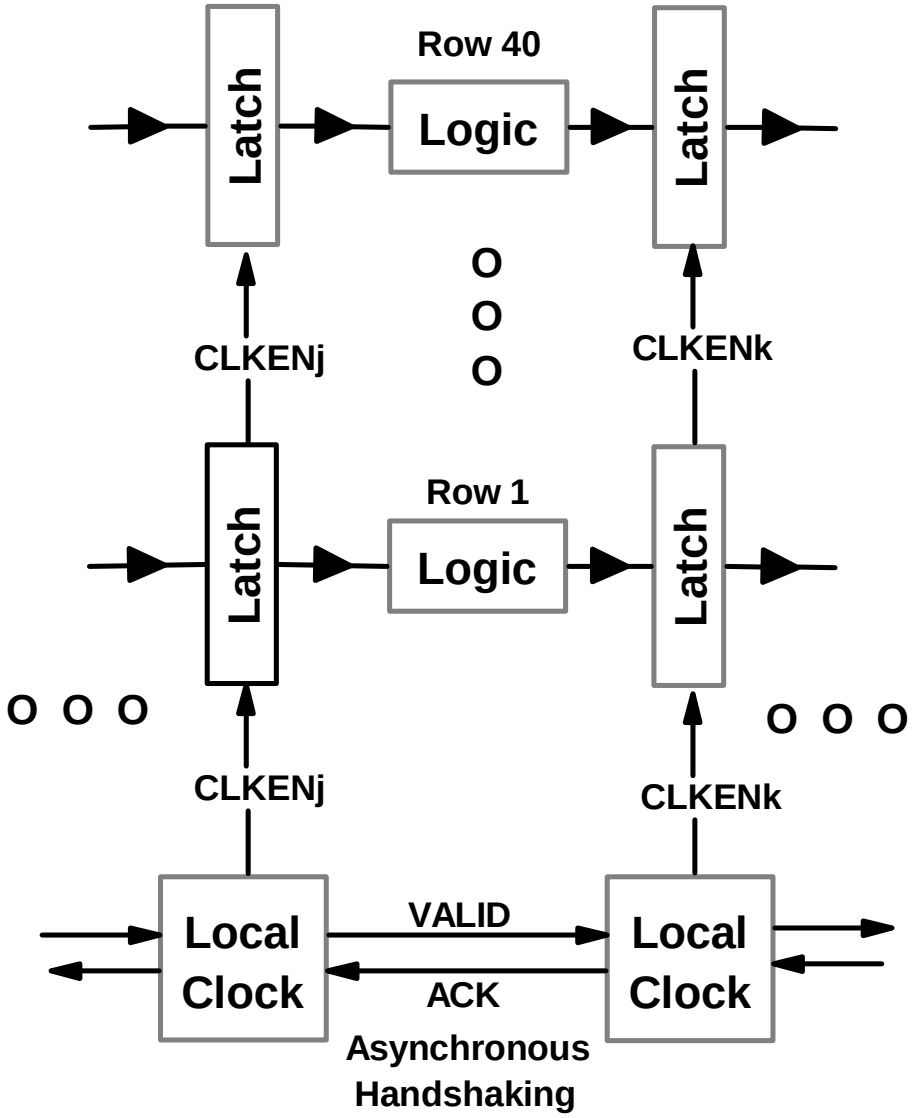

Multiplier Core with Local Clocks

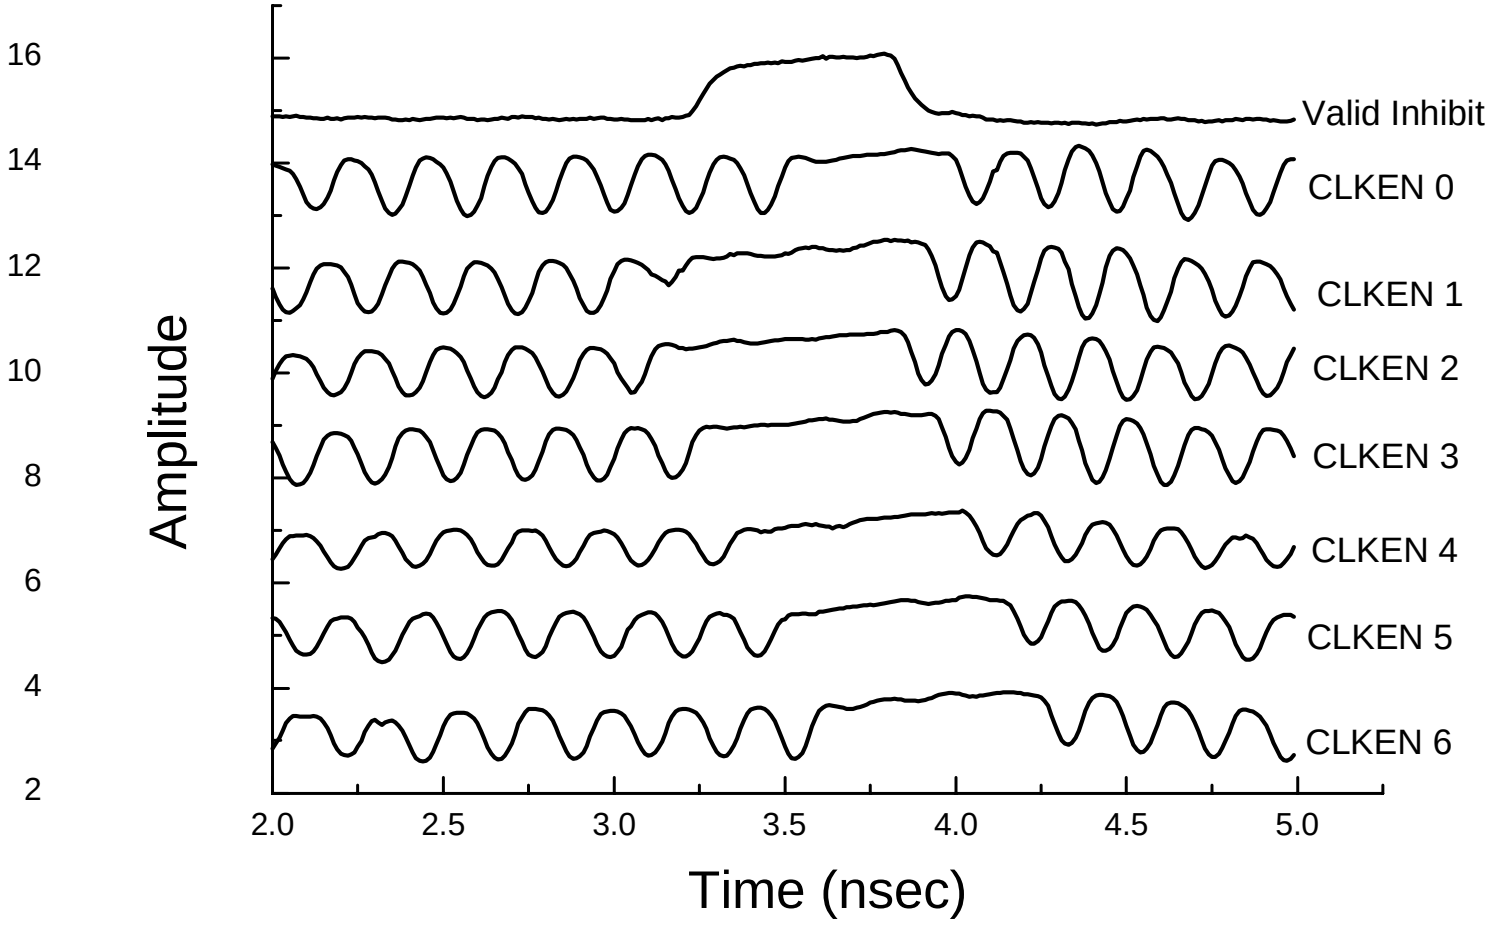

Measured Local Clock Waveforms at 4.5GHz

Frequency Versus Channel Length

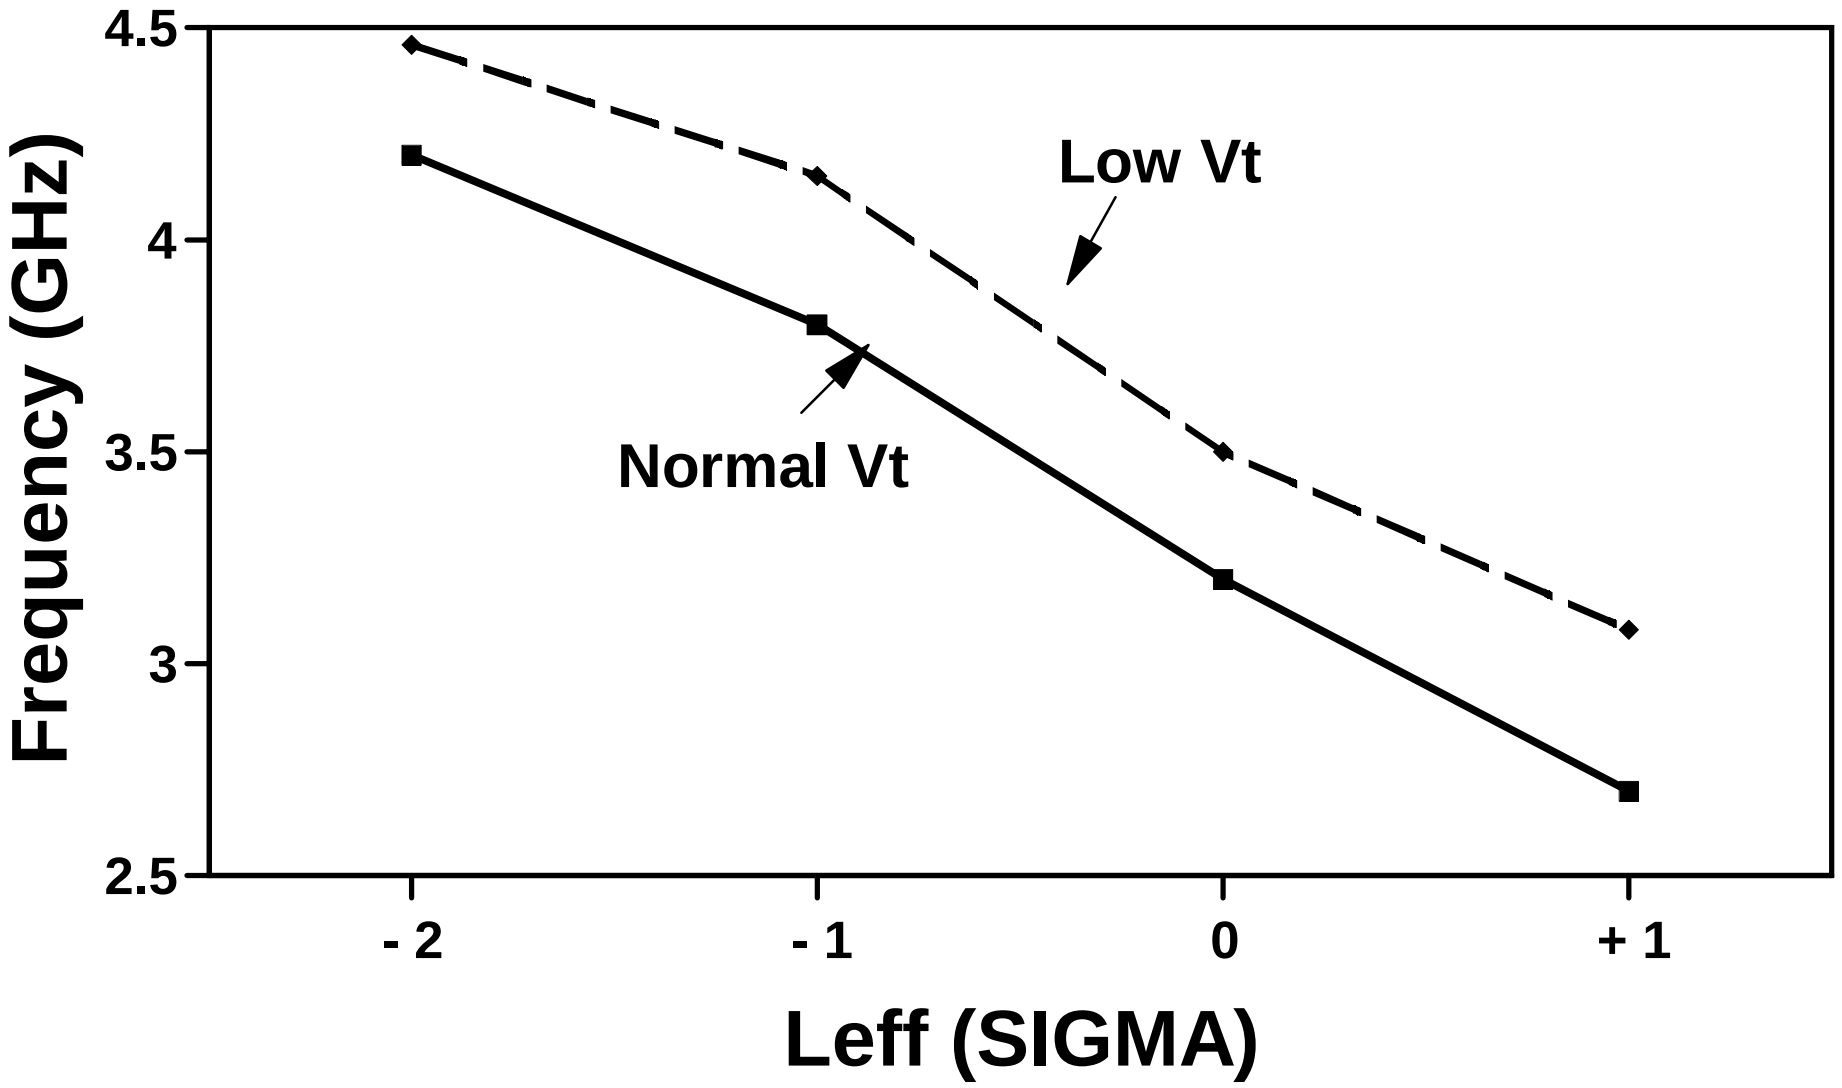

Syn to Asyn to Syn Interfaces

Fine Grained Clock Gating

Power and Noise

IBM T. J. Watson Research Center and university partners
IPCMOS section of tutorial, Schuster and Cook, IBM

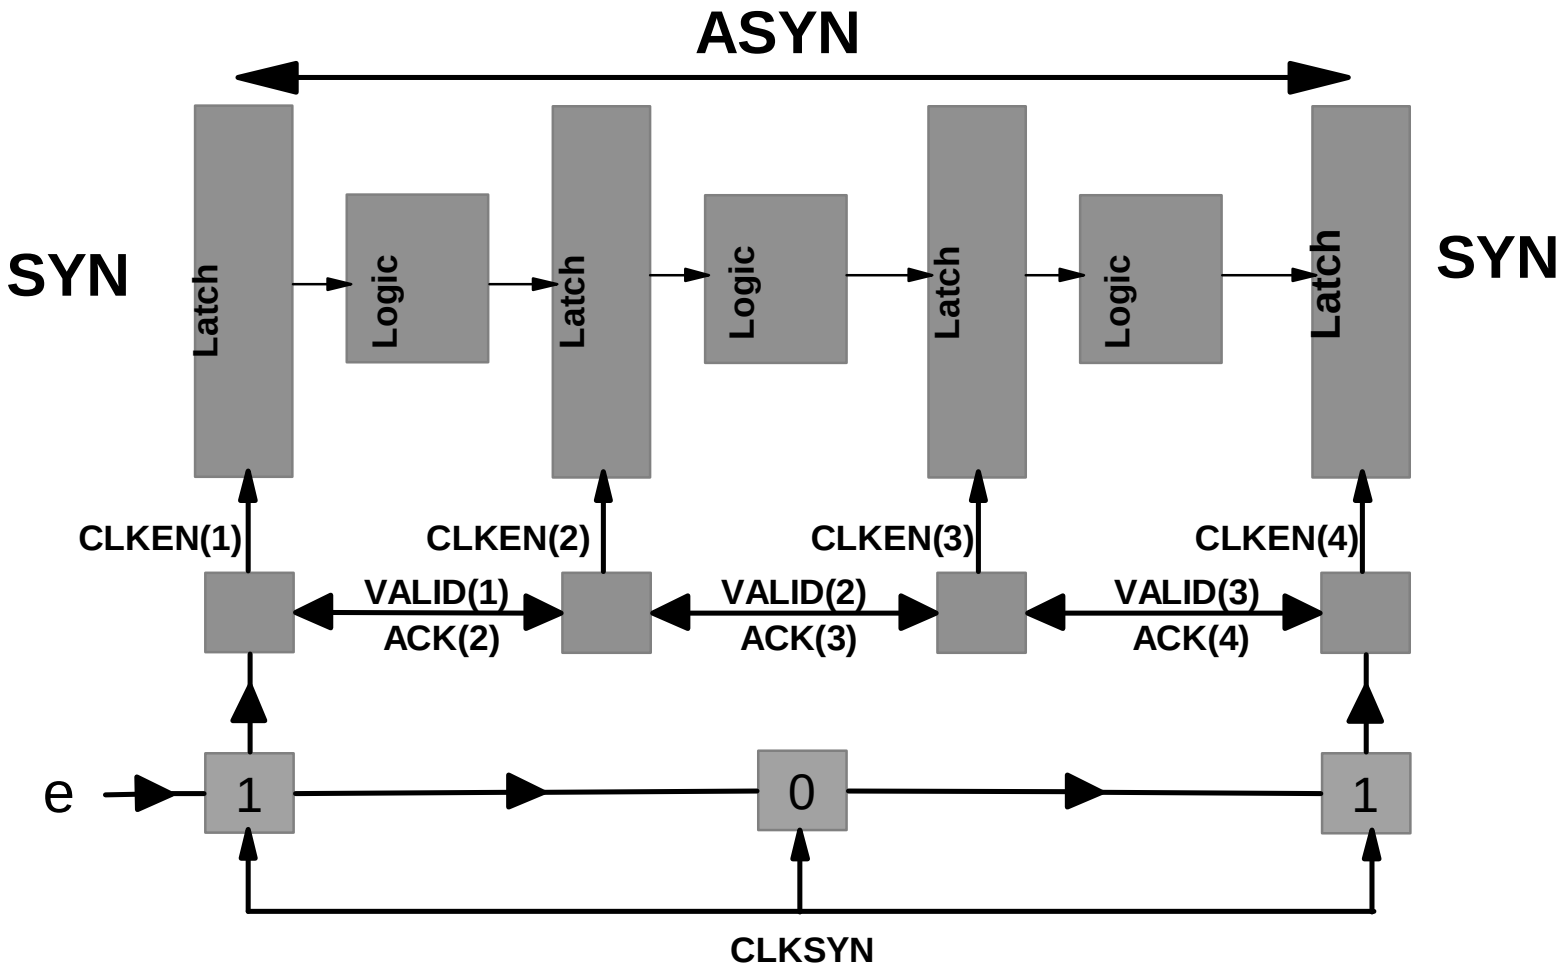

Syn to Asyn to Syn Interfaces

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

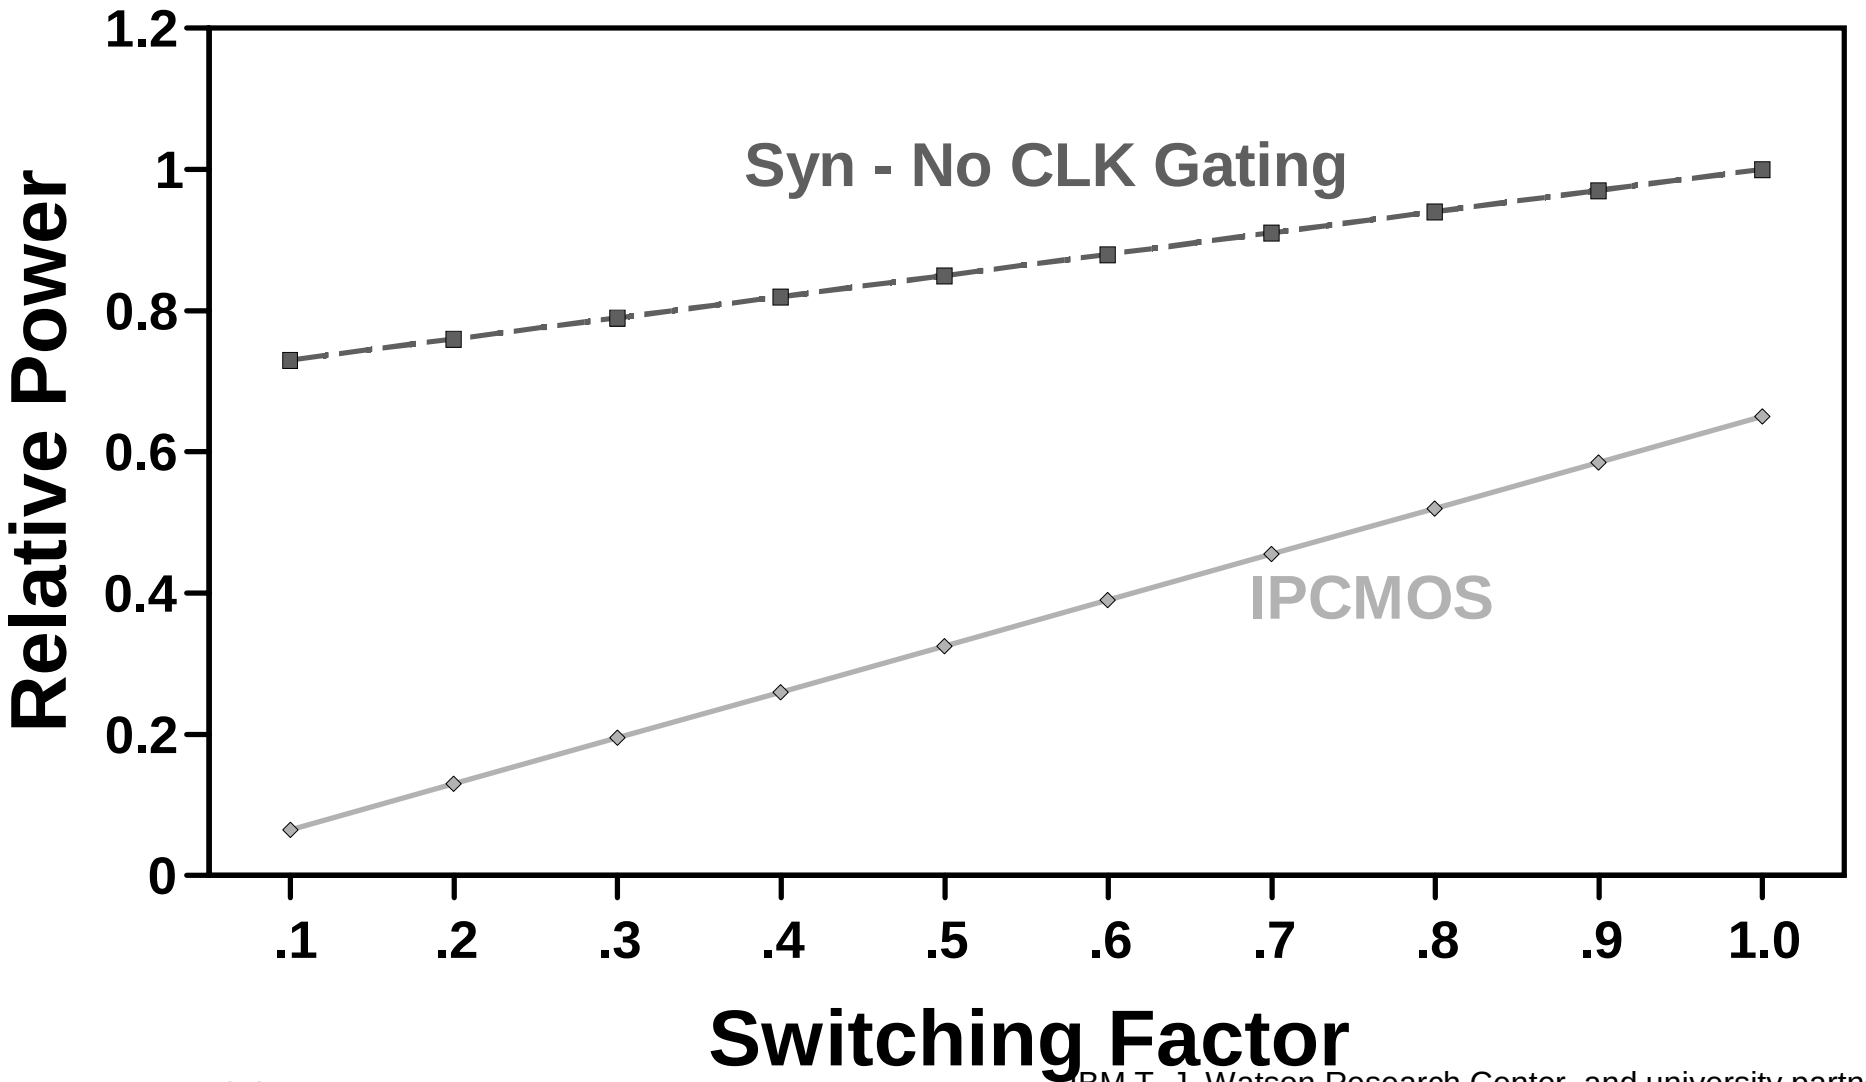

IBM T. J. Watson Research Center and university partners
IPCMOS section of tutorial, Schuster and Cook, IBM

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

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