Leveraging 3D-IC for On-chip Timing Uncertainty Measurements

NC STATE UNIVERSITY

Randy Widialaksono, Wenxu Zhao, W. Rhett Davis, Paul Franzon
Department of Electrical and Computer Engineering
North Carolina State University, Raleigh, NC, USA

Introduction

Abstract

Modern high-performance designs require accurate on-chip timing uncertainty measurements for post-silicon validation of high speed interfaces and clock distribution networks. These measurements are facilitated by on-chip timing sensors, which incur area, routing, and power overhead.

With increasing design complexity and process variations, post-silicon validation and debug capabilities must keep up accordingly to meet competitive product time-to-market. However, enhancing post-silicon validation and debug capabilities cannot simply be met by proliferating on-chip structures, since the overhead would be prohibitively expensive.

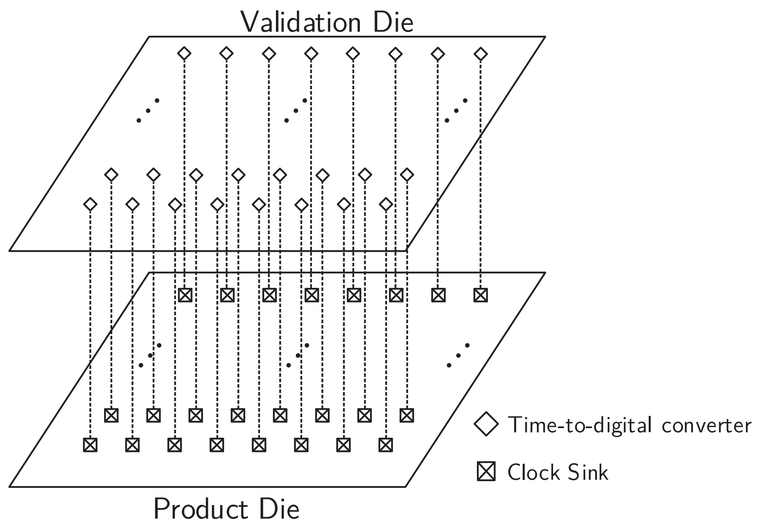

Solution

Move validation structures onto a separate die which would be stacked onto the product die.

The cost of die stacking would be justified by reducing the product die area or by accelerating silicon debug and validation for faster time-to-market.

Advantages

- Enhanced On-chip Timing Observability
- Reduced Area Overhead
- Noise Isolation

Circuit Implementation

Methodology

1. TDC
   - Sub-gate delay measurement resolution
   - Low Power
   - Small Area
2. Clock sink selection
3. Inter-die vias assignment
4. Reference signal for indirect skew measurement

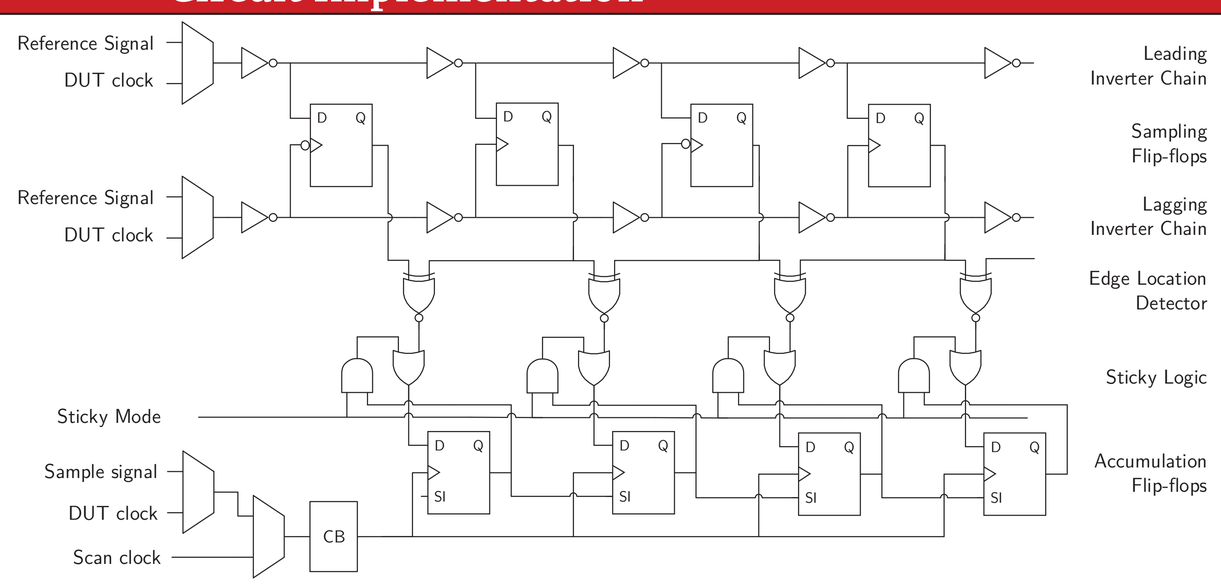

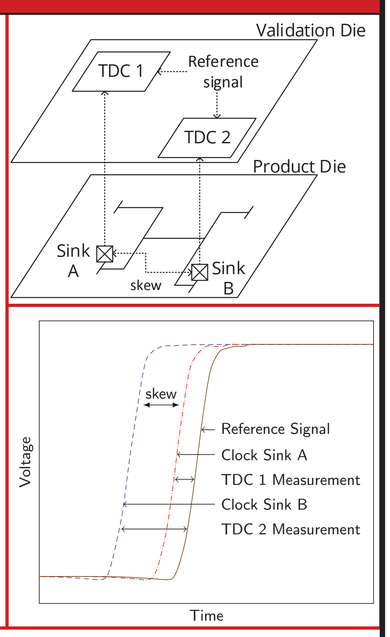

Clock Skew

Clock skew is measured by comparing different DUT clocks to a reference signal. The DUT clock and the reference signal are sent down the leading and lagging inverter chain respectively. The sampling flip-flop chain indicates how long the leading edge have propagated before the lagging edge arrives. Afterwards, the sampled signals are processed by the edge detection logic and latched into the accumulation flip-flops when triggered by the Sample signal. This Sample signal could be generated by delaying the reference signal for the edge detection logic propagation delay, or independently. The time between the reference edge and the DUT clock edge is indicated by multiplying the number of consecutive zeros with the Vernier delay line resolution. The skew between two clock sinks is finally calculated by comparing the outputs of its TDCs.

Clock Jitter

Jitter measurement is launched by setting the circuit to sticky mode and feeding the DUT clock into three inputs:the leading inverter chain, the lagging inverter chain and the clock buffer for the accumulation flip-flop chain. The changing location of edges indicates jitter. By asserting the sticky mode signal, the worst case jitter is recorded by the accumulation flip-flop chain. Jitter measurement is concluded by switching the circuit into scan-out mode.

Area & Power

Technology	Area	Power	Energy
130nm	16,931 µm2	2.11339 mW	1.057 pJ

Measurement results for an implementation of the TDC architecture which contains 150 inverters in each delay line, sampling and accumulation flip-flops, and processing logic.

Power consumption measured is for jitter measurements is assuming input clock frequency of 2 GHz.
Energy consumption measured is listed for a one-shot skew measurement.

Physical Implementation

Die Stack Configuration

Face-to-Face (F2F): F2F bondpoints have negligible delay parasitics, hence allowing on-chip timing measurements without imposing additional measurement offsets. F2F bondpoints commonly have 10-25 µm [6] pitch, hence providing fast, high throughput interface between the product and validation die.

For tier assignment, the validation die would be designated as the tier with its substrate thinned for I/O connections. This assignment is to avoid modifications to the design on the product die. The product die would then have its I/O connections routed through the F2F bondpoints, along with additional signals of interest for debug and validation.

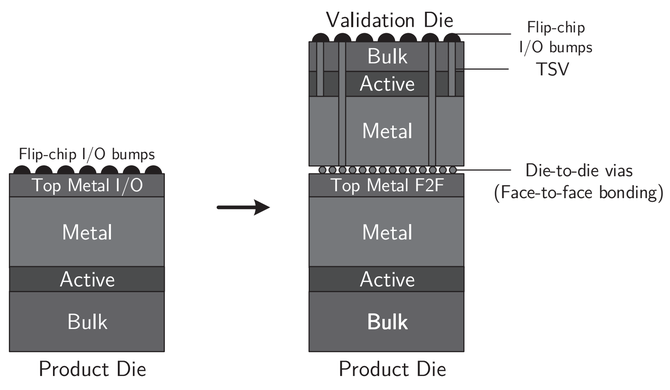

Design Flow

The physical design challenges include selecting which clock sinks to probe and assigning those selected sinks to available bondpoints/bumps. The design flow starts with a placed and routed design with its parasitic extracted. This is to ensure we obtain clock sink insertion delay analysis that are as accurate as possible. After the initial place and route step, the available bondpoints are those that are not reserved for routing the product's I/O signals in both die stack configurations shown on the right.

The bump assignment step considers the interconnect delay between the clock sink its assigned F2F bump. If the parasitic delay matches between two TDCs, hence the difference of both outputs would indicate actual skew. Otherwise, the parasitic delay must be de-embedded from the measurement readings. One way to avoid characterizing and de-embedding parasitic delay is to match the wire delay between a selected clock sink and its assigned bump.

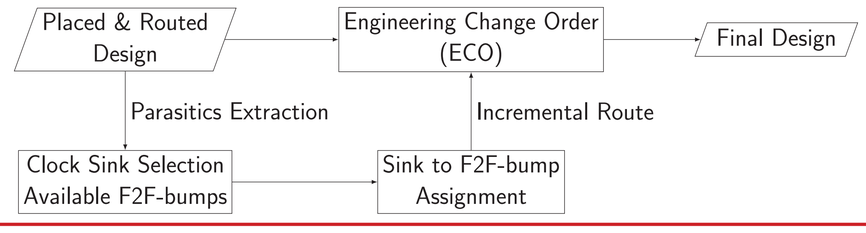

Input: sg_l: List of sink groups
Input: b_l: List of available bump locations
Input: d: Distance from sink to bump
Input: e: Distance tolerance
Input: n: Obtain n nearest bump per iteration
Ensure: Available bumps (b_l) > tapped sinks
Output: a: List of sink-bump assignments

function AssignAllGroups(sg_l, b_l, d, e, n)
  $kd \leftarrow kdtree(b\_l)$
  for all $sg \in sg\_l$ **do**
    append(a, AssignSinkGroup(sg, kd, d, e, n))
  end for
  return a
end function

function AssignSinkGroup(sg, kd, n, d, e)
  for all $s \in sg$ **do**
    $bl \leftarrow$ sort_ascend (n_nearest (s, kd, n))
    while tail(bl)< d **do**
      $n \leftarrow n * 2$
      $bl \leftarrow$ sort_ascend (n_nearest (s, kd, n))
    end while ▷ query n-nearest neighbor until list contains bump that reaches distance
    for all $b \in bl$ **do**
      if manhattan_distance$(s, b) = d \pm e$ **then**
        append$(m, (s, b))$ ▷ m: intermediate list
        remove(kd, b)
        break
      end if
    end for
  end for
  return m
end function

Matching wire delays could be achieved through custom routing or wire delay matching. Without custom routing, we approximate wire delay by calculating the Manhattan/rectilinear distance between a selected sink with a candidate bump, and then verify with parasitic extraction. In a scenario where the assignment fails due to available bump constraint, the designer could prioritize selected sinks into groups.

Bump Assignment

- k-d tree [7]: Fast nearest available bump search
- Specify sink groups for priority assignment.
- Specify distance and range.
- Algorithm finds nearest available bump to each sink
- Sink-to-bump assignments are incrementally routed with ECO.

References

[1] P. Dudek, S. Szczepanski, and J. Hatfield, "A high-resolution cmos time-to-digital converter utilizing a vernier delay line," Solid-State Circuits, IEEE Journal of, vol. 35, no. 2, pp. 240–247, Feb 2000.

[2] R. Franch, P. Restle, N. James, W. Huott, J. Friedrich, R. Dixon, S. Weitzel, K. Van Goor, and G. Salem, "On-chip timing uncertainty measurements on ibm microprocessors," in Test Conference, 2007. ITC 2007. IEEE International. IEEE, 2007, pp. 1–7.

[3] E. Rotenberg, B. Dwiel, E. Forbes, Z. Zhang, R. Widialaksono, R. Basu Roy Chowdhury, N. Tshibangu, S. Lipa, W. Davis, and P. Franzon, "Rationale for a 3d heterogeneous multi-core processor," in Computer Design (ICCD), 2013 IEEE 31st International Conference on, Oct 2013, pp. 154–168.

[4] N. Madan and R. Balasubramonian, "Leveraging 3d technology for improved reliability," in 40th Annual IEEE/ACM International Symposium on Microarchitecture, 2007. MICRO 2007, Dec. 2007, pp. 223–235.

[5] S. Mysore, B. Agrawal, N. Srivastava, S.-C. Lin, K. Banerjee, and T. Sherwood, "3d integration for introspection," Micro, IEEE, vol. 27, no. 1, pp. 77–83, Jan 2007.

[6] P. Enquist, G. Fountain, C. Petteway, A. Hollingsworth, and H. Grady, "Low cost of ownership scalable copper direct bond interconnect 3d ic technology for three dimensional integrated circuit applications," in 3D System Integration, 2009. 3DIC 2009. IEEE International Conference on. IEEE, 2009, pp. 1–6.

[7] J. L. Bentley, "Multidimensional binary search trees used for associative searching," Commun. ACM, vol. 18, no. 9, pp. 509–517, Sep 1975.