Do power efficient asynchronous circuits have to be slow?

Simon Moore
simon.moore@cl.cam.ac.uk
University of Cambridge, Computer Laboratory

9 June, 2003

Abstract

Asynchronous circuits have been promoted as a low power and higher performance alternative to synchronous designs. I conclude that asynchronous circuits are unsuitable for high performance systems design, but may have application for ultra low power designs which are very slow.

1 Introduction

This extended abstract argues that the asynchronous circuit design community should only work on slow circuits. It is designed to be contentious in order to raise debate at the UK Async. Forum. I do look forward to the debate and will provide a small reward for each proposition disproved.

2 Background

2.1 Power Consumption

Power is consumed switching the load provided by gates and wires, which is referred to as *dynamic power*. Increasingly leakage, or *static power*, is becoming significant.

Asynchronous circuits do have the advantage that computations need only be performed as required, leaving circuits in a standby state when there is no activity. Synchronous circuits, on the other hand, have the penalty of distributing a clock. Clock pausing is used to reduce power and this can be applied over a range of circuit areas. None-the-less, a reasonable proportion of the clock tree still oscillates, thereby consuming power. Clock distribution networks on high end processors can consume 40% of the power budget. For smaller systems which have lower clock frequency requirements, clock distribution is much easier and consumes less power.

Whilst asynchronous circuits eliminate the clock, power is consumed in request/acknowledge signals and completion detection circuits.

Proposition 1: *given an architectural design, the dynamic power consumed is proportional to the switching activity and circuit area.*

2.2 Control Overhead

For high performance circuits, the control overhead on the critical path must be minimised. For a clocked system, the dynamic control components are typically computed in parallel with data manipulation. Clock jitter imposes a control overhead since data must be stalled in latches between the minimum and maximum predicted jitter (typically 3% to 5% for PLLs). For asynchronous circuits it is possible to make the control overhead zero [1], though this is often not the case particularly for synthesised designs (e.g. from Tangram, Balsa, etc.).

Proposition 2: *high performance systems require a low control overhead.*

2.3 Parallelism

Parallelism is a key architectural technique to improve performance. Asynchronous circuits offer more concurrency than clocked circuits, though merging data requires synchronisation to be imposed. Synchronisation is almost for free in clocked designs (less so for very large designs) but a control overhead penalty sometimes has to be paid in asynchronous designs. It is not clear whether asynchronous circuits can exploit more parallelism than synchronous ones.

Proposition 3: *high performance designs have lots of parallel activity.*

3 Speed Independent Design

Speed independent design relies on a data encoding scheme (e.g. typically 1-of-2 or 1-of-4) to propagate a request token along with the data. Computation on data encoded in this way requires 50% to 100% (sometimes more) circuit area compared to a conventional binary encoded data.

For random data, the number of transitions per bit of data is typically 1/2 per cycle where as 1-of-2 encoded data is always 2 transitions per cycle and 1-of-4 is always 1 transition per bit per cycle.

Proposition 4: *given that SI implementations of high performance systems are larger and require more transitions to perform the same computation, they consume more power than their synchronous counterpart.*

4 Bundled Data Design

Bundled data (BD) design (including GasP) uses binary encoded data. As a consequence the number of data transitions and circuit area are comparable with synchronous design and, therefore, consume similar power. Rather than clock power, BD designs consume power distributing and delaying (timing) request/acknowledge signals. None-the-less, these signals are easier to distribute (no global skew issues).

Proposition 5: *For high performance BD designs with lots of parallel activity, the power consumed is comparable with their clocked equivalent but is no better.*

5 Conclusion

High performance asynchronous circuits consume more (or at best similar) power to their clocked counterpart. However, for ultra low power sequential designs where little if any parallelism is exploited, perhaps asynchronous design will win.

References

[1] T. E. Williams and M. A. Horowitz, "A zero-overhead self-timed 160ns 54b CMOS divider," *IEEE Journal of Solid-State Circuits*, vol. 26, pp. 1651–1661, Nov. 1991.