CCIT Report #642 October 2007

Zero latency synchronizers using four and two phase protocols

Rostislav (Reuven) Dobkin and Ran Ginosar
VLSI Systems Research Center, Technion—Israel Institute of Technology, Haifa 32000, Israel
rostikd@tx.technion.ac.il
October 2007

Abstract

Synchronizers typically incur long latency of multiple clock cycles, resulting in low throughput. This work presents a number of novel fast synchronizers, based on four and two-phase protocols: a four and two-phase two-flop synchronizer which reduces the data cycle from 6-12 down to 2-4 clock cycles, and a four- and two-phase LDL synchronizer which strives for maximum throughput and "zero-latency," namely data transfers that incur no extra penalty due to synchronization. These synchronizers are useful for data transfers over long interconnects. Simulations of best- and worst-case scenarios are presented which demonstrate the improved performance of the novel synchronizers. The results are compared to two-clock FIFO and to conventional two-flop synchronizers.

1. Introduction

Advanced silicon technologies enable the incorporation of an increasing number of modules in a single chip, constituting Systems on Chip (SoCs) devices. A SoC typically consists of multiple-clock domains, resulting from requirements for different external frequencies, the integration of modules that were designed to operate on different frequencies, and clock gating and partitioning of large and fast clock trees. Moreover, in order to reduce power consumption, frequency and voltage may also be changed dynamically in DVFS systems [1]-[3], leading to dynamically changing clock relations during chip operation.

A SoC constructed of multiple modules, each working with uncorrelated local clock, is termed a Globally Asynchronous, Locally Synchronous (GALS) system [4][5]. Data synchronization and communication across clock domains in GALS architectures is a major challenge. Non-scaling interconnect complicates even more the synchronization problem. Increasing global wire delays incur high overhead, directly affecting communication performance. In addition, long wires have high variability in delay, both due to process variations and noise [6][7].

Asynchronous solutions for global communication across clock domains are preferred over synchronous ones since they eliminate the need for re-synchronization when crossing clock domains, do not require complex clock distributions and are more flexible under changing voltage and temperature conditions [8]-[12]. Thanks to these advantages, ITRS [13] predicts that by the year 2020, 40% of SoC global signaling will be performed asynchronously.

Dynamically changing clock frequencies and wire delay variations call for robust synchronizers that provide high data rate and low latency. The mutual relationships of pairs of clock domains are classified in Table 1 according to the frequency and phase differences of the two domains. Mesochronous domains share the same frequency and have a constant phase difference between them, which can be compensated by relatively simple synchronizers [24][42], e.g. by a small FIFO. Adaptive phase compensation can be employed to connect multi-synchronous domains, in which the phase drifts slowly over time [40][41], as well as plesiochronous domains [43], where a very small frequency difference can be viewed as a phase drift. When two different-frequency clocks are used in the periodic