

Adaptive Network-on-Chip with Wave-Front Train Serialization Scheme

Se-Joong Lee, Kwanho Kim, Hyejung Kim, Namjun Cho, and Hoi-Jun Yoo

Dept. of EE&CS, Korea Advanced Institute of Science and Technology (KAIST)
 373-1, Guseong-dong, Yuseong-gu
 Daejeon, 305-701, Republic of Korea
 E-mail : shocktop@eeinfo.kaist.ac.kr

Abstract

An adaptive network-on-chip (NoC) is implemented with self-calibration and dynamic bandwidth control. The chip calibrates skew between clock domains automatically for reliable mesochronous communication. Link bandwidth is controlled adaptively according to network traffic for energy-efficient packet transmission. A new on-chip serialization scheme, wave-front train (WAFT), is used in the NoC chip to realize high-performance serial link with minimum overhead. The chip is fabricated using 0.18 μm CMOS technology. The overall network and WAFT operations are successfully measured at 1.2Gb/s and 3Gb/s, respectively.

Introduction

Recently, the NoC has been actively studied to replace Bus architecture and to be used as a system-on-chip design platform. And, relatively few researches have been reported on the real implementation to date [1,2]. The first chip, which is fabricated with 0.38 μm technology [1], demonstrated feasibility of high-speed (800MHz) on-chip serialized networking with mesochronous communication. The second chip, which is fabricated with 0.18 μm technology [2], achieved 1.6GHz network operation, focusing on low-power consumption and system-level integration. The NoC chips, however, utilized conventional schemes for on-chip serialization and mesochronous communication. The conventional schemes suffer from overheads in the NoCs, because they do not reflect NoC-specific situation.

In this paper, we report an NoC with a new serialization scheme, wave-front train (WAFT), which provides high-performance serialization and requires less system overhead. Using the WAFT serializer the NoC provides adaptive link bandwidth control according to the network workload for energy-efficient packet transmission. For reliable mesochronous communication with minimum overhead, programmable delay synchronizer architecture is proposed, and the NoC utilizes the synchronizer for self-calibrating inter-clock domain skews.

In following sections, concept and detail operations of the WAFT will be presented, and the adaptive control features will be described in detail. Finally, experimental results and conclusion will be presented.

Wave-Front Train

A. Concept and basic operation

One of the fastest conventional serializer/deserializer (serdes) structures is shift-register type (Fig. 1) [3]. It consists of MUXs and serial D-FFs. The D-FFs load parallel data bits and shift them to one side using high-speed clock. The critical path delay is one MUX and one D-FF delay. The conventional




Fig. 1. Conventional shift-register type serdes

architecture has three problems; First, maximum clock frequency is limited by the delay time of the D-FF. Second, the high-speed clock for the serialization is system overhead. Third, data recovery at the destination requires additional synchronization process if the serial link is long enough to produce skew between the sender and receiver.

A new serdes architecture, wave-front train (WAFT), is proposed to overcome these limitations. The WAFT uses physical delay constant of delay elements (DEs) as a timing reference instead of clock, and utilizes signal propagation phenomenon instead of the shifting mechanism. And, sampling timing for data recovery at the receiver is embedded in the serialized signal.

A successful chip implementation using a technique utilizing delay time for high-speed logic operation is also found in race-logic-architecture, RALA [4], in which timing difference between two signal lines are utilized for Boolean operations.

Fig. 2-(a) shows a 4:1 WATF serializer circuitry. When the EN is low, the $D<3:0>$ is loaded to the $QS<3:0>$. The VDD input of the MUXP, which is called pilot signal, is also loaded to the QP. The pilot signal will be used to stop signal propagation at the deserializer. The GND input of the MUXO precharges the serial output (SOUT) to the ground while the serializer is disabled. If the EN is asserted, the $QS<3:0>$ and the pilot signal start to propagate to the output. Each signal forms a wave-front of the SOUT signal, and the timing distance between the wave-fronts is the DE and MUX delay. The series of wave-fronts propagates to the deserializer like a train.

When the SOUT arrives at the deserializer, it propagates through the deserializer until the pilot signal arrives at the end of the deserializer, i.e. STOP node. The unit delay (DE + MUX delay) of the deserializer is the same with that of serializer. Therefore, when the pilot signal arrives at the STOP node, $D<3:0>$ also arrives at $Q<3:0>$ on time as shown in Fig. 3. If the STOP signal is asserted, the MUXs feed back its output to its input, so that the output value is latched.

The MUX delay time varies depending on whether a signal transits low-to-high or high-to-low. To compensate the polarity-dependent delay difference, the DEs are implemented

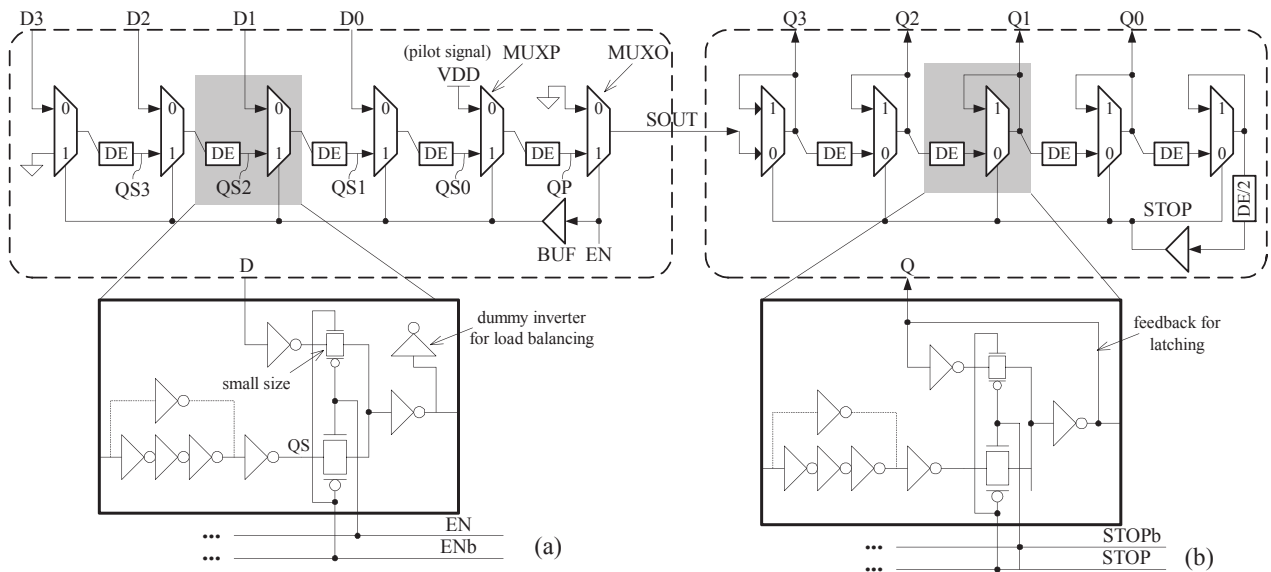


Fig. 2. Wave-Front Train (a) serializer and (b) deserializer




Fig. 3. Timing diagram of the WAFT deserializer.

with odd number inverters, i.e. 3, thus polarity of a propagating signal is inverted at every DE. When the 2 inverter delay is required for fast operation, three inverter chain is connected in parallel with one inverter, which is well-known technique in phase interpolation. To increase the speed, transmission gate of the data input port is small size.

The minimum pulse width of a data to propagate without pulse width degradation is 215psec in the 0.18 μ m CMOS technology. In a 8:1 serdes, this results 4.3Gb/s operation taking into account additional timing overheads like pilot signal and precharge, while the conventional one has maximum 2Gb/s throughput at the same condition. In power comparison, the WAFT serdes, which does not use power-consuming D-FFs, dissipates 47% less power for random input vector when the two serdes's are designed for the same 2Gb/s data rate. The WAFT does not require additional clock whereas the conventional scheme requires 2GHz clock for 2Gb/s operation. When multiple serial links are used in the conventional scheme, the skew between the multiple lines may cause sampling failure of some bits because the lines share single sampling clock. In the WAFT, however, every serial link has its own pilot signal, so that the skew does not cause such a problem.

B. PVT variation

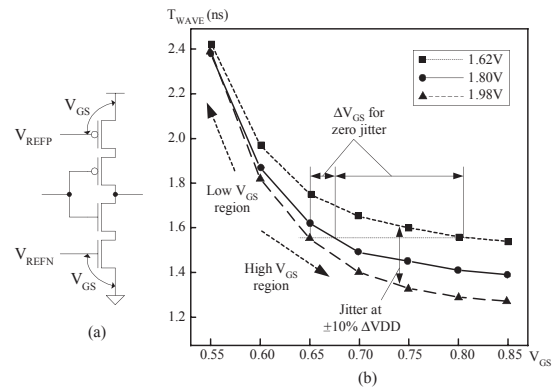


Fig. 4. (a) Current starved inverter and (b) T_{WAVE} vs V_{GS} for VDD variations.

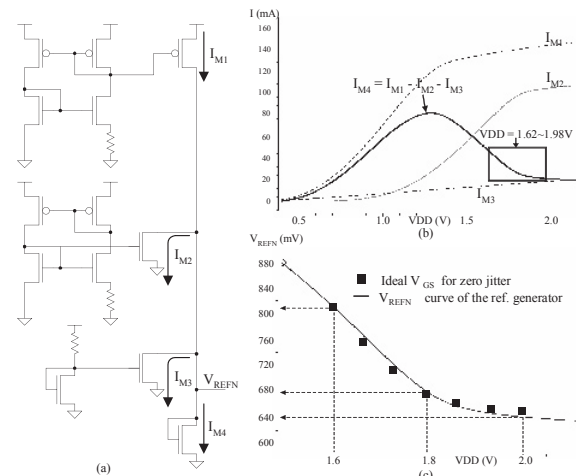


Fig. 5. (a) Reference voltage generator, (b) current, and (c) voltage profile

The WAFT serdes operates based on the fact that the unit delays of the serializer and deserializer are the same. Any difference will produce jitter at the receiver, which degrades performance of the WAFT serdes.

Among the PVT variations, the supply voltage variation affects the WAFT operation most seriously. For example,

$\pm 10\%$ variation of 1.8V VDD causes 30% jitter of the WAFT deserializer operation.

To resolve this problem, a DE can be designed with current starved inverters as shown in Fig. 4a. Fig. 4b shows T_{WAVE} (see Fig. 3) variation of a serializer according to the V_{GS} . When the V_{GS} is lower than 0.6v, the T_{WAVE} is almost independent from VDD variation. However, the T_{WAVE} is relatively high, or data rate of the serializer is low. As the V_{GS} increases, the T_{WAVE} is reduced but shows jitter on VDD variation: To keep the T_{WAVE} constant in the high V_{GS} region, the V_{GS} has to be changed adaptively according to the VDD variation.

For the adaptive V_{REFN} generation, the reference voltage generator shown in Fig. 5 is designed. In Fig. 5c ideal V_{GS} values for zero jitter and the actual V_{GS} generated by the reference voltage generator are plotted as dots and a solid line, respectively. The voltage profile is obtained using current subtraction technique in the reference voltage circuit: In Fig. 5a the I_{M4} is determined by subtracting the I_{M2} and the I_{M3} from the I_{M1} . The current profiles are shown in Fig. 5b. The I_{M4} curve has negative slope when the VDD is lower than 1.8v, and it is moderated as the VDD exceeds 1.8v. As a result, the reference voltage generator outputs proper reference voltage according to the VDD variation as shown in Fig. 5c. The V_{REFP} is generated using the same technique.

With the reference voltage generator, 30% jitter is reduced to 11% when the supply voltages have static $\pm 10\%$ differences from 1.8v. When the supply voltages dynamically varies with 200MHz frequency, 17% jitter is observed.

Adaptive Control Features

A. Adaptive link bandwidth control

Maximum bandwidth of a link wire (B_{MAX}) is given by $B_{MAX} = 0.35 / t_r$, where the t_r is rising time.

Because the t_r is dependent on supply voltage, the B_{MAX} is a function of supply voltage. When a network transfers packet signals through a link, the output bandwidth (B_{OUT}) must be lower than B_{MAX} . In other words, the B_{MAX} only have to be slightly higher than the B_{OUT} . Since the B_{OUT} varies according to traffic demands of processing units in the NoC, the B_{MAX} is set to the highest B_{OUT} in conventional NoCs. The NoC of this work, however, adaptively controls supply voltage of a link to lower the B_{MAX} to currently required B_{OUT} . Therefore, the power consumption in links is reduced effectively.

Fig. 6 depicts the overall operation of the implemented link bandwidth control scheme. It controls the B_{OUT} and B_{MAX} with two steps. The queuing buffer controller inspects the buffer status. When the buffer utilization is low, it asserts low-bandwidth-enable (LBE) signal, and outputs 1 packet per 2-cycles. When the LBE is enabled, the WAFT serializer selects low-VDD (LVDD) to reduce B_{OUT} . The line driving buffers also use LVDD, so that the B_{MAX} as well as power consumption decreases. When the queuing buffer controller changes the LBE signal, it disables the packet-enable (PE) signal for 1-cycle, which is the time necessary for the serializer and drivers to settle supply voltage.

Fig. 7 shows analysis results of B_{MAX} and B_{OUT} variation according to the supply voltage. The B_{MAX} is for a link wire of fan-out 3, and the B_{OUT} is for 4:1 WAFT serializer. The B_{OUT}




Fig. 6. Adaptive bandwidth control

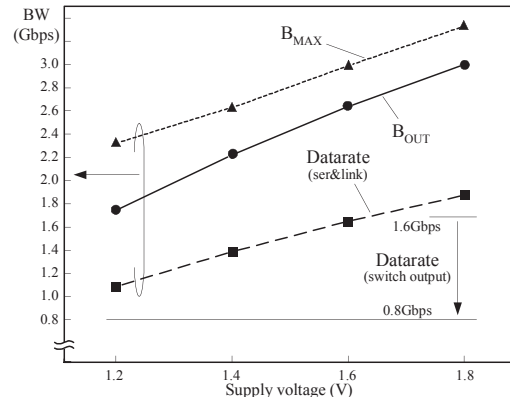


Fig. 7. Bandwidth and data-rate variation according to the supply voltage.

is determined by the delay time of DEs in the WAFT serializer. Because the delay time increases as the supply voltage drops, the B_{OUT} scales down together with the supply voltage. Therefore, the B_{OUT} is kept under the B_{MAX} automatically. The maximum data-rates that the serializer can support at 1.8V and 1.2V are 1.9Gb/s and 1.1Gb/s, respectively. And, those data-rates are enough to support switch output data-rates in normal mode (1.6Gb/s) and low-bandwidth mode (0.8Gb/s), respectively.

B. Self-calibrating phase difference

For mesochronous communication, the [1] utilizes FIFO synchronizers, whose overhead is considerable: The synchronizers consume more than 20% of overall network power consumption. In the [2], a single-stage pipeline synchronizer is used to minimize synchronizer overhead, and possible synchronization failure was resolved using intensive simulations. However, such a full-custom solution is not practical to be applied to a complex NoC system.

In this work, a programmable delay synchronizer is used. A variable delay (VD) is connected with a simple pipeline synchronizer, and the VD is controlled according to the network circumstance. An appropriate VD setting for each network clock frequency and configuration is programmed in a register file. The programming is done during initialization period using the phase detecting unit which finds the besting timing to sample the input signal.

Using the programmable delay synchronizer, the NoC performs phase difference calibration in initialization period, and performs reliable synchronization during operation time.

Experimental Results




Fig. 8. Microphotograph of the implemented NoC chip.

The NoC chip is fabricated using 0.18 μ m CMOS technology. Fig. 8 shows the microphotograph. The chip consists of four level-1 switches, one level-2 switch, 4 additional switch fabrics, network interface logics, traffic generators, and test patterns of WAFT circuits. The switches are interconnected through serial links with 4:1 WAFT serdes's. The chip size is 5mmx5mm including PADs, and total transistor count is 409k. The chip operates up to 400MHz which results 1.6Gb/s per a link wire.

Fig. 9 shows a waveform of 8:1 WAFT operation at 1.8V. The input parallel data is an output of 8b binary counter. According to the linearly increasing binary value, the corresponding serialized data are successfully measured. Eye-diagram of the serial output is shown in Fig. 10. The serialized 8b data follow the pilot signal. It is measured at 375MHz, which results 3Gb/s operation.

A switch output with adaptive link bandwidth control feature is captured as shown in Fig. 11. The waveform is measured at 300MHz. When the LBE signal is disabled, the switch outputs 1 packet data per 1 clock cycle, and the WAFT serializer outputs the serialized data using normal VDD. When the LBE is enabled, the switch outputs 1 packet data per 2 clock cycles. And, the bandwidths of the WAFT serializer and link are reduced using low-VDD, 1.2V. In the waveform, the signal levels of the high-bandwidth and low-bandwidth signals are shown as the same, because they are probed through a PAD driver.

Conclusion

An NoC chip with adaptive bandwidth control scheme and self skew-calibration is implemented using 0.18 μ m CMOS technology for application to SoC platform. The chip also employs a new serialization scheme, WAFT. The WAFT eliminates the use of high-speed serialization clock, thus minimize system overhead. It achieves high data-rate serialization with simple circuitry. It also supports adaptive bandwidth control by using dual supply voltages. The WAFT




Fig. 9. Measured waveform of a 8:1 WAFT serial output.




Fig. 10. Eye-diagram of the 8:1 serial output.




Fig. 11. Measured waveform of adaptive bandwidth control operation.

concept is proved to be feasible by its 3Gb/s operation at 1.8V, and the NoC also successfully operates at 1.2Gb/s.

References

- [1] Se-Joong Lee, et al., "An 800MHz Star-Connected On-Chip Network for Application to Systems on a Chip," Technical Digest of ISSCC, pp. 468-469, 2003.
- [2] Kangmin Lee, et al., "A 51mW 1.6GHz On-Chip Network for Low-Power Heterogeneous SoC Platform," Technical Digest of ISSCC, pp. 152-153, 2004.
- [3] Shinji Kimura, et al., "An On-Chip High Speed Serial Communication Method Based on Independent Ring Oscillators," Technical Digest of ISSCC, pp. 390-391, 2003.
- [4] Se-Joong Lee, et al., "Race Logic Architecture (RALA): A Novel Logic Concept Using the Race Scheme of Input Variables," JSSC, pp. 191-201, 2002.