High-Performance 64-input Cross-Correlator

Erik Ryman†‡   Per Larsson-Edefors‡   Anders Emrich†   Lars "J" Svensson ‡

†Omnisys Instruments AB, Gruvgatan 8, SE-421 30 Västra Frölunda, Sweden

‡Dept of Computer Science and Engineering, Chalmers University of Technology,
SE-412 96 Gothenburg, Sweden

Abstract— **Cross-correlation is a signal processing method that requires immense amounts of calculations. Advancements in integrated circuit process technologies enable much smaller, less power-dissipating cross-correlators. A single-chip 64-input cross-correlator is designed and simulated, demonstrating the feasibility of creating large single-chip cross-correlators. A methodology for simulating large regular layouts has been developed to verify that performance variation, caused by process variation, can be kept under control. Improvements in power efficiency over previous cross-correlators are also shown.**

Fig. 1.   The cross-correlation scheme.

I. Introduction

Cross-correlation is a signal processing method that is likely to be important in future THz interferometry for astronomical observations. Previous and current ground-based cross-correlators, such as the ones used in VLA [1] or ALMA [2], are manufactured on multiple PCBs, each using multiple ASICs. Since these cross-correlators are optimized for maximum computational performance, they are far from being portable and power-efficient enough for space application.

Ever-shrinking process technologies enable new ways to construct cross-correlators. The possibility has now emerged to fit a complete, large cross-correlator in a single-chip solution. This will enable powerful satellite-borne cross-correlators, where size and power dissipation are critical factors. Geo MS (also known as GAS) [3], currently under development, is such an instrument. The Geo MS is a satellite for weather and climate observation in the mm and sub-mm wavelengths. The satellite will have a rotating Y-shaped array of receivers, and the signals of these are to be cross-correlated. Another satellite project, also under development, in which the results from our research could be applicable is GeoSTAR [4], also an atmospheric sounder. A cross-correlator aimed at usage in the GeoSTAR satellite is being developed at the University of Michigan [5].

The cross-correlation of two signals describes the similarity between them. The discrete cross-correlation of two signals $f[n]$ and $g[n]$ is:

$$(f * g)[m] = \frac{1}{N} \sum_{n=0}^{N-1} f[n]g[n+m]$$

The signals are multiplied and integrated at different time lags, see Fig. 1. The result is a function of time lag.

One of the greater challenges of implementing a cross-correlator with many input signals, is the need for the signal pairs at each correlation point to be synchronous. To be able to control the skew between signals, it has to be accurately
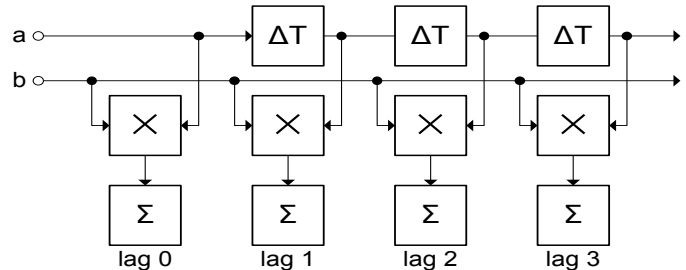
simulated. For a chip of this size a circuit simulation of the entire cross-correlator would be impractical. However, the very regular structure of the cross-correlator enables us to create simplified MATLAB models. While Sec. II describes the cross-correlator design, the simulation approach and subsequent results are presented in Sec. III and IV.

II. The 64-input Cross-Correlator

A. Background and Requirements

The design presented in this article is a first step in studying the possibility of making larger (such as 100-input) one-chip, ultra-high performance cross-correlators for space applications. This scaled-down cross-correlator has 64 digital 1-bit input signals. Each of the input signals has to be cross-correlated with each of the other input signals. The cross-correlator has no lags, meaning that each signal will only be multiplied once with each of the other signals. The cross-correlator is designed in a STMicroelectronics 65-nm technology using a mix of general-purpose, standard threshold voltage and low-power, high threshold voltage transistors. This process has a 1 V nominal supply voltage.

A larger cross-correlator with higher input precision would need significantly more IO pins. Assume a 100-input design with 1.5 bit precision. Such a design would require 200 IOs for data inputs alone; if differential signaling is to be used, 400 pins are required. On top of this we need to accommodate clocks, power supply and ground, readout and control pins. This is a difficult challenge, but not an impossibility: As a comparison up to 1,200 IOs are used in recent FPGAs [6].

B. Correlator Structure

All incoming signals are synchronized at the lower edge of the cross-correlator, whose structure is shown in Fig. 2. The signals are then weaved in a pattern, to make all signals

connect at some point. At each point the two signals connect to a 2-input cross-correlator. The entire cross correlator consists of 2016 such correlation sub-blocks. The pattern is scalable, meaning the path length of clocks and signals between correlation sub-blocks is not dependent on the number of input signals. The structure is also nearly rectangular, which makes it area efficient when fitted onto a chip. At each row of correlation sub-blocks, the data and clock signals will have to be nearly synchronous to enable passing of data to neighboring correlation sub-blocks. Too much skew between the different columns would create incorrect cross-correlation results. There is, however, no requirement for different rows to be synchronous to each other. This makes the cross-correlator horizontally synchronous and vertically asynchronous.

The readout is done through a scan chain passing back and forth through the rows of the cross-correlator.
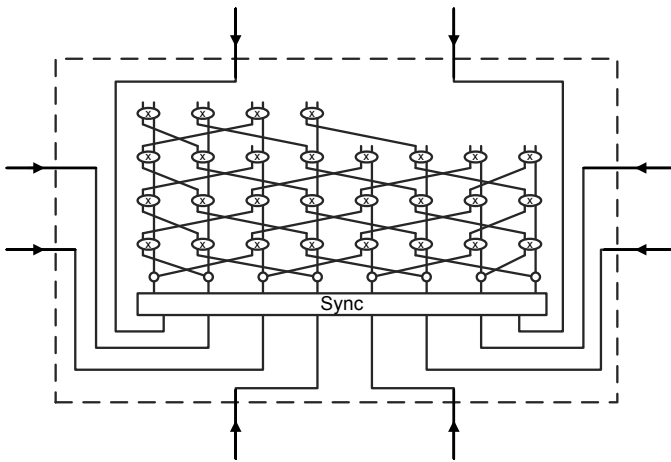

Fig. 2. 8-input example of data and clock routing.

Fig. 3. Cross-correlator core layout.

The total size of the cross-correlator layout (Fig. 3), excluding pad frame, is 1.4 mm^2.

C. The Correlation Sub-Block

Early simulations showed that a significant clock skew would appear between signals in the same row. To reduce this clock skew the two incoming clock signals in each of the correlation sub-blocks (Fig. 4) are compared and the slowest is passed through. This is done using a C element.

The data signals are synchronized to the clock signal using fast dynamic D-flip-flops [7]. Since all signals are 1-bit, the multiplication in each correlation sub-block is performed by a simple XOR circuit. The integration (Fig. 5) is made in two steps. The first few bits of the integration chain, the prescaler, will be noise dominated. This means there is no need to read the data from the prescaler. Also, the prescaler will need to be significantly faster than the later flip-flops in the chain. In this cross-correlator design the prescaler consists of six flip-flops. The rest of the chain consists of 24 flip-flops. These will only have to work at speeds 64 times lower than that of the prescaler; significant amounts of power can therefore be saved by using small, high-V_T transistors.
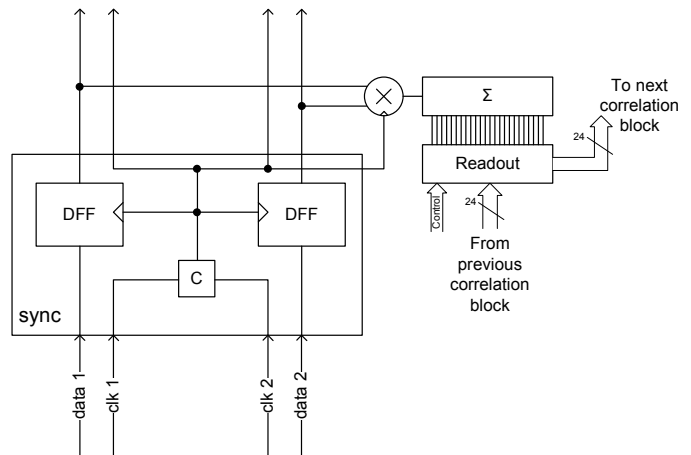

Fig. 4. One correlation sub-block.

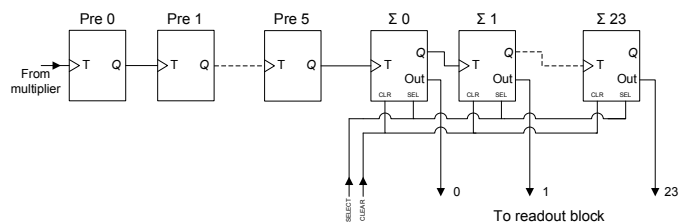

Fig. 5. The integrator.

III. SIMULATIONS

As previously discussed, the speed performance of the design is limited by the skew between the different clock paths. While the design is the same for all paths, process variation has to be accounted for.

A simplified model of the cross-correlator is constructed in MATLAB. For each correlation sub-block, only a time delay

is needed. This greatly simplifies the simulation, since a full Monte-Carlo simulation would require random generation of all process parameters, followed by a huge circuit simulation of the signal propagation.

To generate the time delay only two values are needed: The mean and the variation of the delay for one correlation sub-block. To find these delay parameters, Monte-Carlo simulations are performed on extracted netlists of one correlation sub-block layout. To include wire variations, the delay difference between RCMAX and RCMIN corner extractions are also simulated and included in the variability of the MATLAB simulation.

Fig.6 depicts a simulation of the entire cross-correlator based on the extracted variation and mean. 10,000 variations of the cross-correlator are simulated, at 4 GHz. The maximal cross-correlator clock skew, due to process variation stays below $0.3T_{clk}$ for the majority of the 10,000 chip simulations.
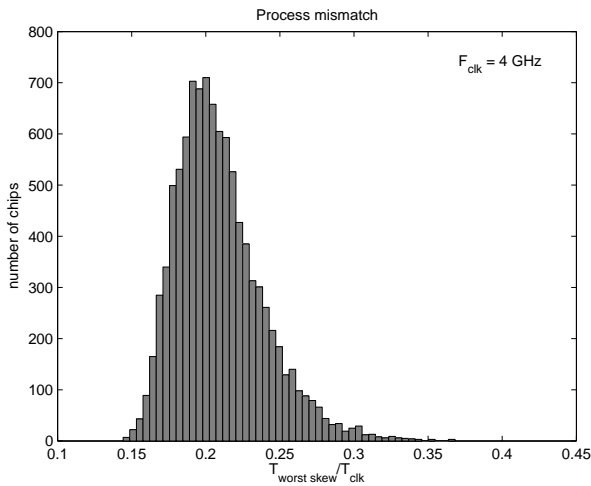

Fig. 6.   Simulation of clock signal skew due to process variation.

IV. DISCUSSION

The C-element used for combining clock signals should be able to handle skews of up to 30% of the clock period. Simulations of the C-element show that it can handle skews of up to 32% of the clock period. While this is not a big margin, high yield is not as important here as in mass-produced ASICs. Throwing away all chips with skews above 25% of the clock period would still leave a sufficient amount of chips.

To get an estimation of the power dissipation of the cross-correlator one of the correlation sub-blocks is simulated at typical conditions. At a clock speed of 4 GHz with random input vectors and a 1 V supply, the average current of one

correlator sub-block is 0.336 mA. Since the cross-correlator consists of 2016 correlation sub-blocks (or channels), the total supply current should be close to 0.68 A. More important to compare is the power dissipation per channel and GHz: 0.084 mW/ch/GHz.

Another, FPGA-based cross-correlator with similar properties is currently being developed at Omnisys Instruments. This design uses 1-bit inputs, no lags and performs 2016 cross-correlations, at a top rate of 330 MHz. The power dissipation of this FPGA-based cross-correlator is estimated to lie between 5 and 10 W. At 5 W this would give 7.52 mW/ch/GHz, which is almost 90 times more than the ASIC solution. The cross-correlator being developed for the previously mentioned GeoSTAR satellite also has 1-bit precision inputs. In [5] a state-of-the-art cross-correlator is estimated to have a total power dissipation of less than 20 W when performing 20 trillion 1-bit multiplications per second. This would equal a power dissipation of less than 1 mW/ch/GHz.

V. CONCLUSION

This cross-correlator design project hints at the possibility of fabricating larger single-chip cross-correlators with more lags and input signals. The simulations in this article show that power dissipation per calculation could be greatly improved over other cross-correlators. Going into smaller process technologies could even further improve the power efficiency. The power delivery and cooling of cross-correlators should no longer be considered to be the major design challenges for satellite-borne interferometers.

REFERENCES

[1] P. J. Napier, A. R. Thompson, and R. D. Ekers, "The Very Large Array: Design and Performance of a Modern Synthesis Radio Telescope," in *Proceedings of the IEEE*, November 1983, vol. 71, pp. 1295–1320.
[2] R. P. Escoffier, G. Comoretto, J. C. Webber, A. Baudry, C. M. Broadwell, J. H. Greenberg, R. R. Treacy, P. Cais, B. Quertier, A. W. Gunst, and A. Bos, "The Atacama Large Millimeter Array Correlator," in *XXVIIIth URSI General Assembly*, 2005.
[3] J. Christensen, A. Carlstrom, H. Ekstrom, P. de Maagt, A. Colliander, A. Emrich, and J. Embretsen, "GAS: the Geostationary Atmospheric Sounder," in *Proceedings of Geoscience and Remote Sensing Symposium*, July 2007, pp. 223–226.
[4] B. Lambrigfsen, W. Wilson, A. Tanner, and T. Gaier, "GeoSTAR - A Synthetic Aperture Approach for a Geostationary Microwave Sounder," in *Proceedings of IEEE Aerospace Conference*. Jet Propulsion Laboratory, March 2004, vol. 2, pp. 1008–1014.
[5] A. B. Tanner, S. T. Brown, S. J. Dinardo, T. M. Gaier, P. P. Kangaslahti, B. H. Lambrigtsen, and W. J. Wilson, "Initial results of the GeoSTAR Prototype (Geosynchronous Synthetic Thinned Array Radiometer)," in *Proceedings of IEEE Aerospace Conference*. Jet Propulsion Laboratory California Institute of Technology, 2006.
[6] Xilinx, *DS100 Virtex-5 Family Overview*, February 2009.
[7] S.-H. Yang, Y. You, and K.-R. Cho, "A New Dynamic D-Flip-Flop Aiming at Glitch and Charge Sharing Free," *IEICE TRANS. ELECTRON.*, vol. E86-C, no. 3, pp. 496–505, March 2003.