

3.6-GHz 0.2-mW/ch/GHz 65-nm Cross-Correlator for Synthetic Aperture Radiometry

E. Ryman^{†‡}, A. Emrich[†], S. Andersson[†], J. Riesbeck[†], L. Svensson[‡], P. Larsson-Edefors[‡]

[†]Omnisys Instruments AB, August Barks gata 6B, SE-421 32 Västra Frölunda, Sweden

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

Abstract—A high-speed low-power cross-correlator ASIC has been implemented in a 65-nm CMOS process for the purpose of synthetic aperture radiometry from geostationary orbiting earth observation satellites. The chip performs cross-correlation on all individual signal pairs from 64 digital 1-bit inputs, which amounts to 2016 individual cross-correlation products. The experimental evaluation, using a specially developed PCB, demonstrates that the 3-mm² chip has a top performance of 3.6 GHz at a 1.2 V supply, at which it dissipates 790 mW.

I. INTRODUCTION

For future improvements in weather forecasting and climatology, earth observation in the microwave band can give important information on temperature and humidity distribution. Performing microwave sounding from a geostationary earth orbit (GEO) would give the additional advantage of continuous coverage of a large part of the earth surface, making it possible to study dynamic weather phenomena. One difficulty with observation in the microwave band is the large antenna aperture needed to achieve the required spatial resolution, especially if performed from GEO. Image acquirement with a parabolic dish antenna is performed by scanning the surface of the earth; here the half power beam width (HPBW) sets the limit of the resolution that can be achieved. To reach a spatial resolution of 30 km, an antenna aperture of approximately 8 m would be required.

Synthetic aperture by interferometry has been used in ground-based radio astronomy observatories for decades. An array of smaller antennas is used to emulate one large antenna. The signals from each antenna pair are cross-correlated; this gives samples of the visibility function, which is the Fourier transform of the intensity image. The cross-correlation function in Eq. 1 describes the similarity between two signals depending on the time skew between them. There are several advantages to using this technique in space. First, the aperture required can be achieved by placing a number of receiving elements on folding booms, making the construction more compact during launch and lighter than a large parabolic dish. Additionally, taking images of the earth using a single dish scanning introduces time skew within the image. Instead, using synthetic aperture interferometry, the entire image can be acquired simultaneously, avoiding the time skew.

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

The excessive amount of calculations required to perform synthetic aperture makes the technique unsuitable for satellite deployment. In 2009, the Soil Moisture and Ocean Salinity (SMOS) satellite was launched carrying the first polar-orbiting 2-D interferometric radiometer [1]. While this was an achievement, taking the same concept into GEO will set even higher requirements for signal processing performance.

Two missions using synthetic aperture on satellites in GEO are considered; Geostationary Synthetic Thinned Aperture Radiometer (GeoSTAR) [2] and Geostationary Atmospheric Sounder (GAS) [3]. GeoSTAR is a NASA project, aimed at improved forecasting and understanding of extreme weather such as hurricanes, and has been in development since 2003. GAS is another weather and climate satellite mission, funded by ESA and developed by RUAG Aerospace and Omnisys Instruments.

Applying synthetic aperture radiometry using cross-correlation in space requires great reductions in power dissipation and hardware size compared to the large ground-based correlators. Using process technologies in the sub-65-nm range makes it possible to fit a cross-correlator, with performance enough for synthetic aperture imaging from GEO, into a single low-power ASIC.

II. ARCHITECTURE

The presented correlator is a demonstrator that has 64 1-bit inputs; this translates to 2016 unique signal pairs, each of which is to be cross-correlated. One of the main objectives of this demonstrator is to verify functionality of a cross-correlator layout that is sufficiently large to make synchronization of all 2-input correlators inputs a challenge.

All 2-input correlators are identical, which makes most of the chip very regular. The regularity of the chip made it possible to estimate both power and performance figures in advance [4]. A 2-input correlator was RC-extracted and subjected to Monte Carlo simulations to acquire timing information. The mean and variation of the clock delay was inserted into a MATLAB model of the entire 2016-correlator core, where the delay for each 2-input correlator was randomly generated. The maximum skew that would appear on the chip marked the performance limit in the simulation. This way 10,000 chips could be simulated in a matter of minutes, indicating top speeds of up to 4 GHz with 99% yield. Further simulations in MATLAB indicated a 20% increase of the maximum skew for

every doubling of the number of inputs. This means that for half the maximum operating frequency of the current chip, a correlator with as many as 2048 inputs can be implemented. Since every doubling means roughly four times the number of correlator elements, the die area however becomes an issue.

The power dissipation of the 2-input correlator cell, including input buffers and routing, was obtained from circuit simulation. The cell dissipation was multiplied by the number of cells to estimate the total chip power. At 4 GHz and a 1.0 V supply, with random input vectors, the total power was estimated to 1.05 W which translates to 0.13 mW/ch/GHz.

The correlator performs 0-lag correlation; hence no delay will be implemented between signals. The integrators are 30 bits long but only the 24 most significant bits are read out; the remaining 6 bits are considered to be mostly noise. At 2 GHz this gives an integration time of up to 0.5 seconds and a data compression ratio of up to $7.9 \cdot 10^{-7}$. Readout is performed through an SPI interface, thus, all data is read out through a single pin. All IOs are single ended, because the design is pad limited.

The chip is based on a commercial 65-nm process technology with a 1.0-1.2 V nominal voltage range. The layout uses both standard- and high- V_T devices; the latter to reduce standby power. Neither the architecture nor the process is radiation hardened, but with technology scaling comes an intrinsic hardening [5]. While the design will fail in the case of a hard error, a single-event upset (SEU) will not cause any major problem. In the event of an SEU, the bit flip is treated as noise during the integration or it will cause an unnatural spike, which is easily detected and removed in the resulting data.

III. IMPLEMENTATION

A synchronization methodology was devised to keep inputs arriving synchronously at all 2-input correlators. There is a total of eight correlator clock inputs: All of these are routed, each together with eight data signals, from the input pads to one edge of the correlator. The data is at regular intervals synchronized with the clock using flip-flops. At the edge, which all paths are routed to, the clocks are split up: Each data signal is then routed together with its own clock. The data-clock pair is split two ways, one going straight, and one going sideways. Fig. 1 shows an 8-input example of the routing pattern used. At each 2-input correlator, the clocks of the incoming signals are synchronized. The idea is that this will




Fig. 1. 8-input routing example.

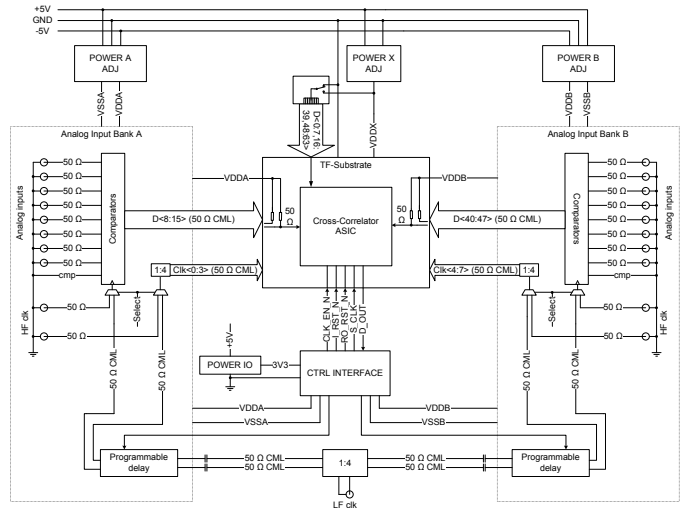


Fig. 3. Test PCB block diagram.

mitigate the impact of device mismatch. This clocking scheme has another advantage; the chip will be asynchronous in one dimension, thus, reducing the system peak current.

The 2-input correlator and synchronization is depicted in Fig. 2. The clock synchronization is performed by C-elements that wait for the latest arriving clock. The data is synchronized to this clock using dynamic flip-flops. The multiplication is performed by a dynamic XOR-gate that propagates pulses. A series of six high-speed semi-static pulse-triggered flip-flops constitutes a prescaler. The rest of the integrator consists of conventional static flip-flops connected to the readout bus. The 23 most significant flip-flops and all readout logic are implemented using high- V_T transistors, while all input routing, synchronization, multiplication and the first seven flip-flops in the integrator are implemented as standard- V_T transistors.

IV. TEST SETUP

A PCB has been designed for measurements of the chip (Fig. 3). The PCB has 16 analog input ports connected to comparators, which act as samplers, divided into two input banks. The remaining 48 inputs can be set to a common logical 1 or 0. The input banks use separate power and clocking domains. Two clocking options are available; for frequencies up to ~ 2 GHz an input with four programmable delay channels to the respective input banks sampler and correlator clocks can be used. For higher clock frequencies, four SMA connectors connect directly to the sampler and correlator clock paths. The advantage of using the programmable delays is the possibility to do frequency sweeps and keep all clocks in correct phase without the need for manual tuning.

The supply voltage to the correlator, and the offset between correlator and input banks can be independently adjusted. A 0.05Ω shunt resistor is used for measuring the chip current.

The 3-mm² correlator chip has been mounted on a custom thin-film substrate to allow for high-speed operation and good thermal dissipation (Fig. 4). The substrate contains decoupling capacitances and etched termination resistors. The substrate is glued and wire-bonded to the PCB.

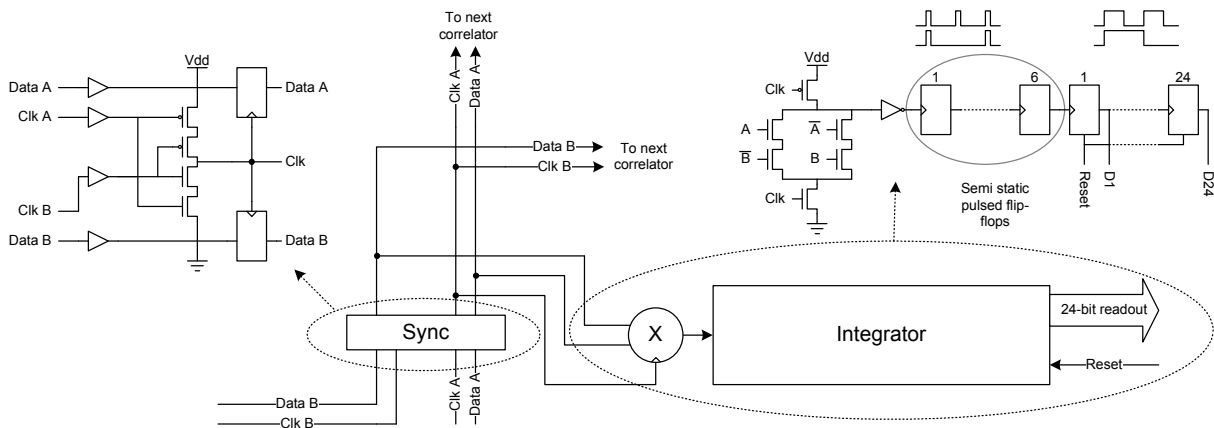


Fig. 2. 2-input correlator schematic.

An Arduino Uno board with an ATmega328 microcontroller connects the control interface of the PCB to a computer via USB for programming of the four clock delay lines, controlling correlation and performing readout. A LabVIEW interface acts as a GUI, controlling both the PCB, a frequency synthesizer for clocking, and a multimeter for supply current measurement.

V. TEST RESULTS

The performance was measured with the 16 analog inputs set at opposite voltage to the remaining inputs. This gave a predictable resulting pattern and allowed counting the number of errors. The performance measurement is summarized in Fig. 5. The PCB did not allow for automated tests at lower supply voltages than 0.8 V. However, the correlator has been verified to perform correct correlations for clock rates of at least 1.25 GHz at 0.65 V and 0.75 GHz at 0.6 V.

The lower limit of the operating frequency is set by the semi-static prescalers; for low frequencies self oscillation will occur. The lower limit is around 300 MHz for full functionality, and it depends only weakly on supply voltage. The top frequency is 2.7 GHz for the nominal voltage of 1.0 V, and it increases to 3.6 GHz for 1.2 V. Around these frequencies the clock distribution starts to fail and this will propagate in the chip according to the routing scheme. Occasional errors do also occur in the range between 2.6 and 3.1 GHz and are subject of further investigation.

To test how well the synchronization scheme handles skew between incoming correlator clocks, a sweep of the time delay between clocks passing through the two delay circuits is performed. A functional breakdown is expected when the phase shift between clocks is close to 180 degrees. According to earlier simulations [4], the synchronization block should ideally be able to handle up to 32% skew between clock signals, when operating at 4 GHz. This would translate to the chip not functioning for 9% of a clock period at 1 GHz. In the measurement, the chip is not functioning for 18% of the period, as shown in Fig. 6. Note, however, that this figure also includes skew caused by the test setup and the input routing, which was not accounted for in the simulations.

For power measurements all 16 analog inputs were connected to a noise source to emulate a case with high switching activity. Fig. 7 shows how the supply current scales with frequency and supply voltage. The curves are nearly linearly increasing with frequency, however, some drop-off occurs since the shunt resistor voltage drop increases with higher currents. Note that the breakdown of the functionality at high frequencies is visible as a non-linear behavior.

By measuring the power dissipation at 2 GHz, both with and without noise applied to the 16 analog inputs, the chip power for uncorrelated noise on all 64 inputs—the realistic use case—can be calculated. The total power dissipation in this scenario, with 1 V supply, amounts to 530 mW which translates to 0.13 mW/ch/GHz. At 1.2 V supply, the corresponding number is 0.20 mW/ch/GHz. The standby power, at 1 V, was measured to 40 mW for the whole chip, with active clock and noise on the 16 analog inputs.

Table I summarizes some specifications, operating conditions and measurement results of the chip.

VI. DISCUSSION

The measured performance did not meet the top speed indicated by simulation at 1.0 V, and there are several possible reasons for this. First, the circuit simulations that used extracted timing information did not accurately account for crosstalk. Second, signal degradation throughout the clock chain was not accounted for; the circuit simulations indicated that 4 GHz was at the limit of what the clock buffers could handle. Furthermore, the test setup is not an ideal environment; while previous simulations accounted for some skew between inputs, the test setup may very well have a larger skew. The power estimation done before tape-out did very closely match the measured value, which marks a significant improvement over current cross-correlators.

A future correlator with an order of magnitude more correlations would still be feasible in terms of both area and power, especially considering that the chip presented in this paper is pad limited. A future migration to, for example, a 40-nm process technology will improve on both area and power.

TABLE I
DEVICE SUMMARY

Process	65 nm
Supply/IO voltage	1.0-1.2 V / 1.0-1.2 V
Size	3 mm ² / 3M transistors
Correlation inputs	64 data, 8 clocks
Number of correlations	2016
Integrator depth	30 bit (24 bit readout)
Operating frequency range 1.0/1.2 V	0.3-2.6/0.3-3.6 GHz
Power dissipation 1.0/1.2 V	0.13/0.20 mW/ch/GHz
Standby power 1.0 V	40 mW

VII. ACKNOWLEDGMENT

This work was supported in part by the Swedish Research Council (VR) under Contract 621-2007-4727.

REFERENCES

- [1] J. Font *et al.*, "SMOS: The Challenging Sea Surface Salinity Measurement from Space," *Proc. of the IEEE*, vol. 98, no. 5, pp. 649–665, May 2010.
- [2] B. Lambrigtsen, A. Tanner, T. Gaier, P. Kangaslahti, and S. Brown, "Developing a GeoSTAR Science Mission," in *IEEE Int. Geoscience and Remote Sensing Symp.*, July 2007, pp. 5232–36.
- [3] A. Carlström, J. Christensen, J. Embretsén, A. Emrich, and P. D. Maagt, "A Geostationary Atmospheric Sounder for Now-Casting and Short-Range Weather Forecasting," in *IEEE Antennas and Propagation Society Int. Symp.*, June 2009.
- [4] E. Ryman, P. Larsson-Edefors, L. Svensson, A. Emrich, and S. B. Andersson, "A Single-chip 64 Input Low Power High Speed Cross-Correlator for Space Application," in *Microwave Technology and Techniques Workshop*, May 2010.
- [5] P. Roche *et al.*, "A Commercial 65 nm CMOS Technology for Space Applications: Heavy Ion, Proton and Gamma Test Results and Modeling," *IEEE Trans. on Nuclear Science*, vol. 57, no. 4, pp. 2079–88, Aug. 2010.




Fig. 4. Cross-correlator ASIC mounted on thin-film substrate.

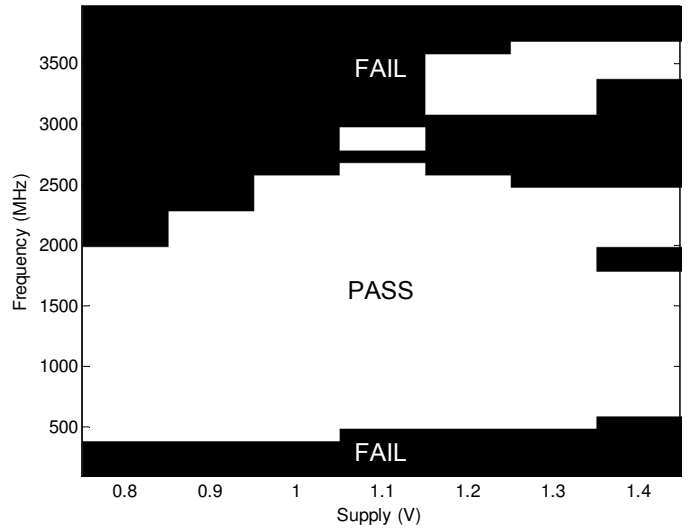


Fig. 5. Shmoo plot for performance vs supply voltage.




Fig. 6. Clock skew sweep.




Fig. 7. Supply current vs frequency and supply voltage.