

RD Engineers Insider Views on Best of Test Award

By [Christophe Zaretti](#), Product Manager
[Agilent Technologies](#)
SA – Switzerland

The Agilent 12-bit, 8 channels AXIe digitizer was elected as the 2013 TM Best in Test winner of the category signal analyzer. Meet with the different RD engineers that developed it.

Best in test summary: RD engineers insider views

We have recently let our RD engineers speak about how they were seeing the different innovations that have been developed during the conception of the [12-bit AXIe digitizer](#) which has been elected as the 2013 TM Best in Test in the product category “signal analyzer”. Here is a summary of the different posts.

Meet the proud project manager



Following the win of the T&M Award “Best in Test” in the category signal analyzer of our latest product, the 12-bit 8 channels AXIe digitizer, we will take you during the next 2 weeks on a tour of the technology breakthroughs that needed to be developed in order to release this new digitizer. It should also help you understand how the Agilent high-speed ADC technology team responds to the highest demands in terms of performances, reliability and time to market. It was obvious to us to start with the interview of our AXIe R&D project manager, Pierre-François.

Pierre-François

Pierre-François, tell us a little bit about your technical background and history.

Pierre-François (PF): I have a strong technical background in electronics design ranging from high-reliability embedded designs for Defense and Aerospace applications to state-of-the-art FPGA based architectures. As a hardware design engineer I have been involved in many area including digital processing, high speed serial data transmission, high speed digital design within FPGA and obviously high speed digitizer design in several form factors.

Pierre-François, as the M9703A project manager how did you feel when you learnt that your product won the “BEST in Test” award in the category signal analyzer?

PF: I'm very proud of all the work that has been delivered by the cross-functional team: It was the most complex digitizer design we ever carried out but it was really exciting to develop a new product in an emerging AXIe form factor standard.

Did the fact that this digitizer is based on AXIe, the new modular industry's open factor, influence your approach and design of the product architecture compared to other products?

PF: Yes, it did definitively. AXIe is basically a measurement-oriented standard layer based on the well-established ATCA standard. It offers both a wide PCB area that let us implement 4 front-end analog mezzanines on a single carrier board and enough power efficiently drawn from a single power rail. The AXIe infrastructure provides also convenient triggering schemes and reference clock distribution across multiple boards and the AXIe Local bus has been used to provide a unique data processing scalability and synchronization over multiple adjacent boards. We hear that the M9703A implements a number of technology breakthroughs. Can you tell us a little bit about what they are? Which ones are you particularly proud of?

PF: This is correct: As part of this new platform we have integrated a number of breakthrough technologies such as an innovative front-end architecture, an enhanced clocking scheme based on an in-house developed IC that delivers clean sampling clock distribution with unmatched added-data jitter performance, a very precise trigger positioning also based on a custom IP chip and a multi-channel multi-board scalable architecture providing phase-matched channels with de-skewing capabilities. I'm very proud of our multi-channel multi-board scalable architecture that provides not only synchronous sampling over multiple channels located on several boards but also synchronous processing over multiple boards as required for example on multiple DDC channels applications. The M9703 provides a true advantage for multi-channels applications and having the ability to scale-up seamlessly multiple boards to create a 40 synchronous processing channels system is really exiting.

For you, what are the key performance and/or features a potential user should look at when considering a high-speed digitizer? How is the M9703A situated in this context?

PF: The analog data conversion performances are of course of primary concerns, but the user should also pay attention to the some details such as the offset range capability, the proposed bandwidth and especially the true performances over the full bandwidth (not only limited as low Fin frequency). The user shall also look at the proposed clocking schemes and their related phase noise performances because the quality of the clock distribution is directly impacting the data conversion performances at high input frequencies .In addition for multi-channel applications the channel phase-matching or channel deskewing capability is of course of significant concern. The M9703A addresses the points above by especially providing consistent data conversion performances over its large input bandwidth and supporting an external clocking scheme that only adds 25fs of added jitter to the user clocking source.

We guess you are now working on new projects. We understand that they are certainly still quite confidential but could you provide some indications to our readers?

PF: With the M9703A, we've built the ground foundation for our AXIe digitizer family. We are now looking at providing additional firmware features such as enabling the Customer to program

his own firmware within our digitizer platform and to extend the product portfolio with other data conversion characteristics...

If you were projected in 10 years time, how do you think the high-speed digitizer technology will have evolved and on what types of projects do you imagine working on?

PF: Over the last decade the technology rapidly evolved from 8-bit multi GS/S ADC to 12-bit multi GS/s ones and from the first simple FPGA equipped with a couple of MGT to very large with dozen of MGT. Thus I imagine working on multi GS/s 16-bit digitizer with high-end FPGA equipped with multiple 100Gb/s transceivers...

Thanks for your time and congratulations again for the Award.

Meet the FE designer engineer



FE (Front End) engineer master, Neil

In this second chapter of “Best in test” series we will focus on the analog front-end (FE) which, as mentioned by Pierre-François the project manager in his interview, is one of the technology breakthrough of the signal analyzer winner, the [M9703A](#).

For the R&D engineering team, the analog front-end (FE) is one of the most critical parts of the digitizer design. Every stage of the design process will influence the quality of the signal conditioning, and hence the signal integrity and data conversion quality. This is why you find in Agilent high speed digitizers very carefully designed analog FEs, from a well-defined and innovative architecture, to a meticulously implemented layout, followed by a rigorous qualification. But let us introduce you the analog FE engineer master, Neil, to speak about it.

Neil, tell us a bit about your experience. How long have you been concentrating on analog FE design?

Neil: I'm embarrassed to tell you the number of years – I'll just say it's longer than most Special Reserve Scottish Single Malt Whiskies are left to mature! I started in FE design for scopes, then digitizer FEs for (at that time) Acqiris, who is now Agilent.

What is the purpose of the analog front-end and what type of functionalities does it offer?

Neil: The FE's job is to scale and offset the Input Signal so that it is correctly positioned for optimal ADC performance. This will normally include single ended to differential conversion, and any other range of possible options such as choice of 50 Ω /1 M Ω input impedance, AC/DC coupling and switchable bandwidth filters. A trigger pickoff is often added for "Internal" trigger operation.

Why do not all of the FE's offer all of these functionalities?

All engineering is a compromise – and FE design is one of the most challenging. Some functionalities are simply not compatible – for instance a 1 M Ω input impedance FE which also has a very high bandwidth with low noise and low distortion (Or if some genius out there knows how to do it – please let me in on the secret !). Also – as in all analogue design - less complexity in a FE is always better than more complexity from the point of view of preserving signal fidelity. The expertise lies in finding the right compromise.

Another constraint for us is that we have very tight physical space and power restrictions for our FEs – which is not generally a problem for oscilloscope manufacturers. This can also lead to a reduced number of options.

What are the other objectives and challenges behind the design of an analog front-end?

One challenge that is overlooked (but always expected by the customer) is having some form of Input Overload protection. It is of course possible to blow up anything if you try hard enough, but we must try and protect up to some reasonable level. For high bandwidth designs this input protection can create performance issues.

The required single ended to differential signal conversion is just in itself not an easy thing to do well. All previous designs used either a Balun (transformer) coupled input – which would then be automatically AC coupled. This gives optimal analogue performance – but a restricted range of usage scenarios. The second option is the use of an active single ended to differential amplifier. This gives true DC to HF functionality, but with an inevitable signal degradation of noise and distortion, especially at high frequency.

What is in your opinion the key thing that distinguishes a good design from a bad one?

As the Channel analogue performance is a combination of the FE and ADC behaviour, a good design is one that does NOT degrade the ADC performance too much, while offering a maximum of functionality. This has become much more difficult as ADC technology is progressing rapidly, and now 12 bit converters exist with very good analogue performance, but which also have high input bandwidths and sample rates. The module specification should show the channel performance over the whole useable frequency range, and these always come down to ENOB (effective number of bits), SNR (signal-to-noise ratio), SINAD (signal-to-noise and

distortion ratio), SFDR (spurious-free dynamic range), DNL/INL (differential/integral non-linearity), etc.

You have apparently succeeded to design a versatile FE on the M9703A which still guarantees extremely high performance over the full input frequency range (typically 58 dB SNR, more than 9 ENOB, 63 dBc SFDR). How did you achieve this?


The M9703A features an innovative DC-coupled input architecture for the single ended to differential conversion process. This combines the best analogue performance features of the AC type at high frequency, with the functionality of the DC type at low frequencies. This allows the very best noise and distortion performance over the whole Input Frequency range. It also means we can offer a very large range of Input Offset capability which drastically increases the general purpose functionality of the module. The available offset range is +/- 2 x the Full Scale input range, for both 1V and 2V Full Scale ranges.

As for Input Overload Protection – the M9703A has been evaluated by customers who frequently have large signal overvoltages and they have found the module to resist without being damaged – whereas other digitizers or oscilloscopes did not.

What is the future in analog FE design? Do you still see open space for technology improvements?

One essential area of IC technology that has not kept pace with ADC improvements is signal amplifier technology. Some IC manufacturers are addressing this problem right now. Over time, the FE design process will become more plug and play, but putting these components together will always require design expertise, which is what sets us apart from the competition !

Zoom in the clock-chip



Analog & IC Design Manager, Daniel

Today and in the next post, we will zoom in the core of Agilent digitizer technology mentioned by Pierre-François as another breakthrough: the development of in-house performance dedicated

IC's . Let us introduce you to our Analog & IC Design Manager, Daniel, who will talk about the clock-chip that was implemented on the M9703A.

Daniel, before we start, can you tell us a bit more of your background? I spent about half of my life doing Analog IC Design. I started as young engineer in a scope company designing FE, clock and Trigger circuits and moved to a large IC Manufacturer company doing Analog IC designs in the area of telecom, automotive and consumer electronic. I then got back in the measurement world dedicated my work to the design of digitizers.

Why has Agilent decided to design a new clock circuit? Wouldn't have been easier to buy and implement available IC's?

Daniel: Agilent development's team #1 objective is to guarantee that our high-speed digitizers reach high-level of performances. Therefore, when we cannot use or find available IC's that respond to this demand, we design in-house dedicated full custom Analog IC. The clock chip is good example: Achieve very good Signal-to-Noise Ratio (SNR) performance has always been one of the top requirement priorities. However, with the development of higher and higher resolution analog to digital converters (ADC's) and with the need of measuring higher and higher frequency signals, we realized that there were no circuits available on the market that could achieve such performance. Some available clock chip or ADC have built-in delay control, but while used this feature usually deteriorate the clock noise floor or jitter and then the SNR. Hence, we decided to create a complete new IC in order to respond to this need of performance for our high-speed digitizers.

A chart showing the SNR vs frequency & ADCs SNR at Nyquist showing the influence of the clock-chip and the Time Base on the SNR is available.

Is the design of a clock circuit also the result of a design compromise like the analog FE is [as explained by Neil in our last post](#)? What are the main challenges and objectives behind it?

Daniel: In the design of clock-chip, the compromise is to compose with the different numerous analog circuits between maximum frequency of operation and noise in particular while implementing delay control function. Our target was to design a clock-chip with multiple clock output and delay control to drive multiple high resolutions ADC with the lowest noise floor or added-jitter as possible. In doing so, the M9703A, which is the 1st digitizer on which we implemented this new IC, achieves the unprecedented noise floor leading to a phase noise of -154 dBc/Hz at 2GHz and 100MHz offset even when using the whole delay control range.

We guess you are working closely with the others R&D groups. Is your development work continuously influenced by the others team's requirements or the needs and definition of your design are clearly defined before you start working on it?

Daniel: Yes of course, we are working closely together and seating in the same space than the hardware and firmware design teams. This helps define the requirements and implement new concepts. Other R&D team members, but as well as sales and marketing people, are fully involved in the definition of our new ICs.

What are the main innovations that were developed with this new circuit? What can you tell us about them?

Daniel: One of the main innovations of our clock-chip is the new architecture design to control the phase and delay of the multiple clock outputs with no added jitter. This allows the best performances while interleaving ADCs and aligning channel sampling time.

What are the main added-values brought by the new clock circuit on the M9703A performances?

Daniel: The clock chip is a key block used in our Time Base mezzanine to provide ultra-low jitter clock to our four ADC mezzanines. The low added jitter of the clock-chip in the order of 25fs leads to a Time Base with an overall jitter of about 100fs. The on-chip delay adjustment allows us to align the sampling time between the ADC mezzanine and the multiple channel of M9703A.

What would be the graal of IC design on high-speed digitizers? Do you see room for others breakthrough innovations in the coming years in these field or rather continuous improvements based on your experience?

Daniel: The Monolithic ADC market is moving very fast in terms of performances. Of course there is always room for continuous improvement in clock-chips while trying to take the best of any new high resolution high speed ADC. The evolution of jitter in ADC is more like in saturation in order of 100fs but we can oversee improvement in terms of frequency and extended delay control ranges. One of the limitations we are always challenging with is board space and power dissipation which are limited in standard low cost package. Integration and trigger time resolution is another challenge that we also decided to respond to in our team as it will be covered in the next post. And as mentioned in the [previous post](#), another key area where we can foresee huge improvement, is the ADC driver or signal amplifier with the venue and availability of new State of the Art IC technology.

Trigger circuitry: 12ps rms of precision



Nicolas

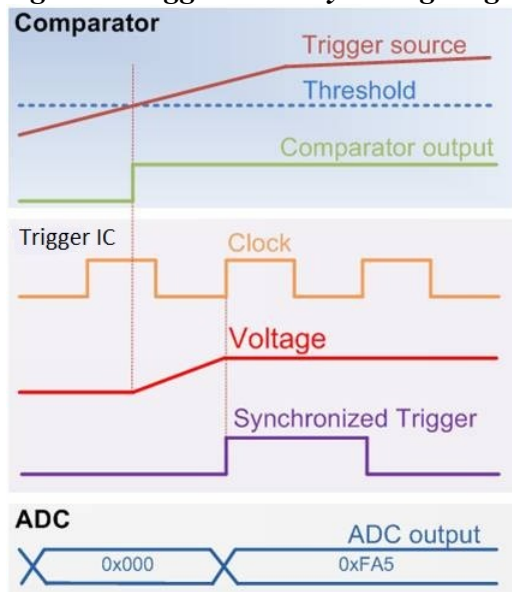
Today, we receive Nicolas who has been concentrating on the development of the Trigger IC that has been implemented on the M9703A. This circuit was developed as part of a trigger circuitry which includes others elements as you will discover in this interview.

Hello Nicolas, tell us a bit more about yourself. Have you always been working in the development of IC?

Actually, I started my career in an automotive company, mainly doing software coding as well as CPLD firmware but this was not my cup of tee. I changed job and worked for a DSO company and was mainly involved in the design of Front-end, multi-chip-modules and also GaAs IC design. After this I completed a PhD at the Swiss federal institute of technology in the field of IC design in the area of RF power amplifier efficiency improvement. Since then, I have been working for Agilent mainly on IC design for time base, trigger circuit and vertical signal conditioning. Incidentally, I am also in charge of the full time base circuitry development, implementing our own IC. Having to implement a circuit in a real application always gives new ideas for improvement.

The trigger circuit you developed is part of a trigger circuitry which is also made of several different blocks. Tell us a bit more about this circuitry. What is the influence of this dedicated trigger IC on the timing precision of a trace positioning (a precision of 12ps rms as a reminder) ? This target of some pico seconds would most probably be very difficult to reach with a full discrete design. As mentioned in your question, the trigger system is made of several blocks, based on ultra-fast analog comparators, commercial Analog-to-Digital Converter (ADC) and a Field Programmable Gate Array (FPGA) (cf figure 1). The trigger chip integrates a lot of functionalities including fast and low jitter logic and a time-to-analog converter (TAC) to precisely place the trigger event. The advantage of using this dedicated trigger chip is to dramatically reduce the occupied PCB area and above everything power dissipation compared to a full discrete design.

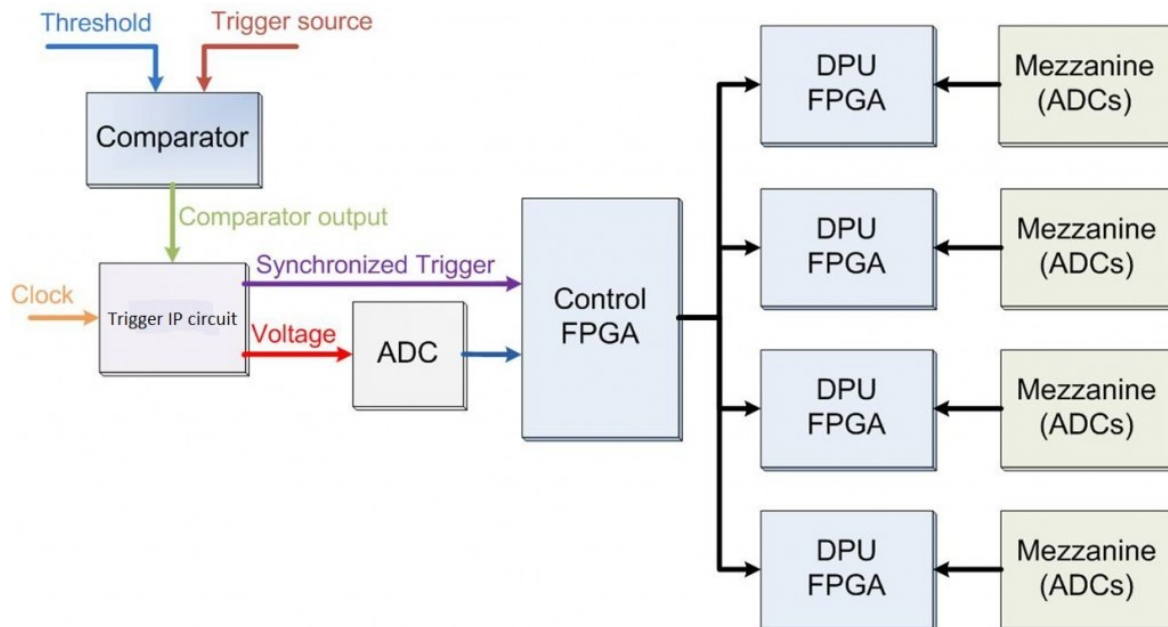
Figure 1 Trigger circuitry timing diagram



What is the added-value of your in-house design trigger circuit?

The added value of the trigger chip is quite straightforward. Definitely, without this trigger chip, we could not propose our customers good trigger capability and precision because we would be limited by the available PCB area and the dissipated power. The edge trigger provides a marker which determines the exact moment the trigger source signal crosses the chosen threshold. Furthermore, the edge trigger information is propagated in real-time to a configurable area of the Digital Processing Unit (DPU) FPGA where customers could potentially implement their own processing algorithms.

Note from the editor: The next figure summarizes the flexibility and synchronization offered by this trigger technology:



Nicolas as a last word: Do you still see room for improvement in trigger circuit design? What is your feeling about it?

The trace positioning precision depends on the slope of the signal one wants to trigger on. When the slope is very steep, the ultimate precision limitation is the TAC repeatability. The room of improvement would be mainly based on finding new architecture for the TAC with better repeatability.

What do we mean by synchronization?



Eduardo

Let's meet with Eduardo who will take us on a tour around synchronization. Eduardo, can you briefly describe your background before you joined Agilent?

I have spent most of my career working as digital designer on FPGAs and ASICs. After finishing my Masters, I joined CERN (European Organization for Nuclear Research). Firstly, I was in charge of the FPGA optimization of SRAM controls in an existing data acquisition board based on VME for high energy physics applications. Secondly, I was responsible for the design and development of the full digital circuitry (data signal processing and digital interface) in a mixed mode ASIC to scope Time Projection Chamber detectors. The work was defended as part of my PhD at Universidad Politecnica de Valencia (Spain). I then joined Agilent as part of the FPGA design team and I have been involved in different projects; for instance, the U1084 zero suppression scheme, the trigger circuitry, high-speed digitizer architecture, multi-board synchronization described here... among others.

How would you define synchronization and in what type of applications, synchronization is a key attribute for a digitizer?

The term synchronization is quite general; we usually say that two systems are synchronous when they run at the same reference clock. The M9703A presents a multi-channel multi-board scalable architecture that provides not only synchronous sampling over multiple channels located on several boards but also synchronous processing over multiple boards.

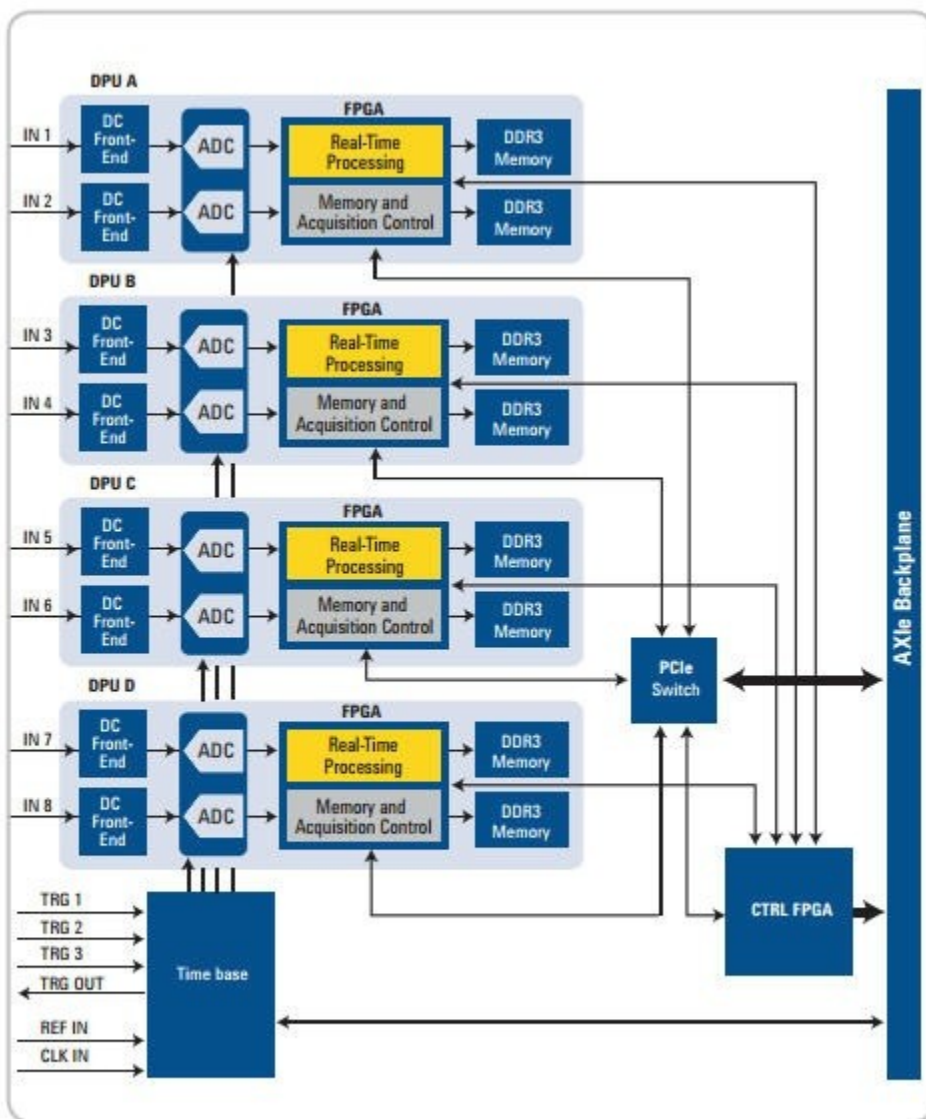
The synchronous sampling assures that each acquisition channel is sampled at the same absolute time. While the synchronous processing guarantees that all the processing algorithms are running simultaneously on the same set of samples (trigger information is included in real time). The M9703A synchronous sampling is possible due to in-house ASICs while the synchronous processing over multiple boards is possible due to the firmware flexibility and precise calibration algorithms.

Currently, we are able to scale-up seamlessly multiple boards to create 40 synchronous processing channels (High-speed digitizers, DDCs or any other processing algorithm) using the

Agilent's M9505A chassis. This is a tremendous gain for a diverse range of applications as: phased array radars, radiographic facilities, fusion research, particle physics...

What are the key R&D challenges that are lying behind the channels' synchronization of a digitizer like the M9703A?

It's not possible to synchronously distribute GHz range sampling clocks from one board to another without affecting either its quality or its relative phase matching. So, the challenge is to provide all the necessary hardware and software elements to assure synchronous sampling and synchronous processing in multiple "asynchronous" boards. It is also important that these elements are part of our M9703A and the chassis to avoid using any additional hardware. For instance, this solution employs some AXIe backplane resources, such as the common AXIe 100MHz reference clock and AXIe Local Bus.



Simplified block diagram of the M9703A AXIe Digitizer.

Do you think that the synchronization across channels on a digitizer can still be improved? As explained by Nicolas (Best in test #4), the trace positioning standard deviation in all channels individually is 12 ps rms. If we consider all channels in a single board, the standard deviation is ~13.5 ps rms. And the first tests have shown that, if we consider all channels in multiple boards, the standard deviation is ~15 ps rms. These are great results for such a scalable system!!

The only limitation (in decreasing the standard deviation) is the accuracy of the reference clock distributed through the backplane which introduce a fixed inter-channel skew. Nevertheless, this is not an issue in the M9703A, because we provide channel deskewing controls. These controls can further improve the inter-channel skew by deploying a system calibration including all mismatches (even including external cables if required).
What would it bring to the big physics experiments?

We know that physics experiments require more and more high-speed synchronized channels with fast data acquisition to accurately capture experimental events. We therefore present a real solution to facilitate the tedious and complex task of synchronizing multiple channels. On top of that, the M9703 provides the possibility of including custom user-designed processing algorithms in these multiple channels (for more info contact us). So, the client will “only” focus on the processing algorithms and will profit from the Agilent Technologies multi-channel infrastructure.

And how do you see it in the future?

In the near future, we would like to test our solution in a chassis with greater number of slots and potentially with lower backplane skew. We expect to have around 100 channels with synchronous sampling and synchronous processing in a single chassis. But, this is not the end... a new challenge will come with the synchronization across several chassis. This is a new challenge that personally I am quite excited about!!

Processing with Xilinx Virtex-6 FPGA's



Giovanni

Hello Giovanni, we are pleased to receive you today. You are FPGA developer and along with what has been exposed lately on this blog, some serious innovations were brought in and around the on-board data processing units (DPU) or more commonly called FPGA's that the M9703A carries (the M9703A counts no less than 4 Xilinx Virtex-6 FPGA's each unit supporting the on-board signal processing of 2 channels).

Giovanni, before we start could you please briefly describe your background?

Well, I have a Master of Science in Digital Electronic from University of Pisa in Italy. I have since the beginning worked with FPGAs and embedded systems at first in Italy in a start-up company where I was in charge of the development of sensor interfaces, then in the Netherlands at European Space Agency (ESA) mainly focusing on designs for space applications. I joined Agilent in Geneva in 2007 working on our first generation PCIe (PCI express) board and then on the new PXIe and AXIe platforms. Today I am mainly working on the deployment of new on-board processing functions looking to systematize this process across our high-speed digitizers. Could you explain to us why Agilent ADC team considers the role of the DPU's more and more important?

The DPUs are responsible for controlling the module functionality, data flow, and obviously real-time signal processing. On-the-fly processing is becoming more and more important, as requirements on both ADC resolution and sampling rate are continually increasing, whilst the available data readout bandwidth remains limited. To address this problem, techniques may be employed at the digitizer level to reduce processing time and optimize data transfer.

How has Agilent then addressed the problem?

Real-time data processing is reached by either getting of-the-shelf (OTS) commercially available firmwares, or by implementing custom algorithms. Agilent provides a variety of IP firmware real-time processing, such as fast Fourier transform (FFT), averaging, peak detect, or digital down converter (DDC) option on the M9703A.

So it seems that Agilent has actually developed quite a few IPs. Why is that? Wouldn't be easier to buy an existing signal processing IP's?

Well the answer is simple. It's impossible to find a commercially available IP which contains the right set of functionalities which address the requirements of our customers when using our digitizers and which is easily adaptable and customizable to fit our systems and the frequencies involved. If we want to guarantee the best real-time performance at the multi-GS/s data rates our digitizers platforms are capable and at the same time fully leverage their exceptional hardware capabilities in terms of trigger positioning or signal fidelity this is the only way to go.

We heard that you are really proud of the latest firmware Agilent developed. What does it do and what are the main advantages of implementing it?

The DDC allows tuning and zooming on the signals to be analyzed, improving the dynamic range, reducing the noise floor, extending the capture time, and accelerating the measurement

speed. This exclusive IP algorithm provides very powerful and flexible digital downconversion on all 8 channels. The filters and local oscillators (LO) are synchronized to maintain constant phase and timing relationships to allow phase-coherent post processing. The DDC allows isolating the signal of interest from other signals in crowded spectrum, and improving the dynamic range as the integrated noise is reduced, increasing the SNR and effective number of bits (ENOB).

Does Agilent consider to open their digitizers' FPGA's to customers signal processing algorithms?

Of course, Agilent high-speed digitizers' team does. Any customers who would have interest in doing so should contact us. As I mentioned above, we believe that the on-board processing is the future of high-speed digitizers. Now we need to make sure that the deployment of real-time processing algorithms does not deteriorate the hardware performances. That's where I spend most my energy and the successful DDC deployment on the M9703A AXIe digitizer is the proof that we can do it.

Thanks a lot for your time Giovanni, we are looking forward to hear about the next FPGA's developments on high-speed digitizers with you in a near future.

[High Speed Digitizer - Agilent](#)