CORPORATE PARTICIPANTS

Lilian Veras
Moderator

Nadav Yafe
Silicom – VP Product and Product Marketing

Oren Benisty
Silicom – VP Sales

PRESENTATION

Lilian Veras

Welcome, everyone, to the Intel Network Builders webinar program. Thank you for taking the time to join us today for a presentation titled: “Optimizing vRAN with Intel FPGA Platform”. Before we get started, I want to point out some of the features of the BrightTALK tool that may improve your experience.

There's a Questions tab below your viewer. I encourage our live audience to please ask questions at any time. Our presenters will hold answering them until the end of the presentation. Below your viewing screen, you will also find an Attachments tab with additional documentation and reference materials, including a number of websites and documents mentioned in this presentation. Finally, at the end of the presentation, please take the time to provide feedback using the Rating tab. We value your thoughts and we’ll use the information to improve our future webinars.

Intel Network Builders webinar series takes place live twice a month, so check the channel to see what's upcoming, and access our growing library of recorded content. In addition to the resources you see here from our partners, we also offer a comprehensive NFV and SDN training program through Intel Network Builders University. You can find the link to this program in the Attachments tab, as well as a link to the Intel Network Builders Newsletter.

Intel Network Builders’ partners have been working to accelerate network innovation by optimizing their solutions on Intel technologies. These industry leaders are recognized in our Winners’ Circle program, and Silicom is a Gold partner. Learn more about our INB Winners’ Circle program by clicking on the link in the Attachments tab.

Today we're pleased to welcome Nadav Yafe and Oren Benisty from Silicom. Nadav is vice president of product and product marketing at Silicom. He leads Silicom product line that spans into the DU, CU, switches, and PCIe add-in, architecture in vRAN, and other market segments. He owns Silicom engagements with customers and partners in telco, cloud, and data center markets. Nadav’s responsibility includes product development for the full portfolio of Silicom products and solutions based on FPGA, x86, and other ASIC. He joined Silicom in 2010. Nadav holds a Bachelor of Electrical Engineering from the Ben Gurion University and an EMBA from Tel Aviv University.

Oren is EVP strategic sales at Silicom, covering the 5G market. Oren holds a Bachelor's in Electrical Engineering and a Master of Business Administration. Oren has over 25 years of experience in the telecom and networking markets, working with major network equipment providers, including ECI Telecom and Rad-Bynet Group, as well as 13 years at Intel embedded in telecom teams.

Welcome, Nadav and Oren. Thank you, again, for joining us today, and I will hand over to Oren to start off. Thank you.

Oren Benisty

Thank you very much, and thank you for Intel for hosting us today, and today we will cover optimizing vRAN with Intel FPGA, and I will start with the agenda. The first slide will cover the Silicom overview. After that, we will look at the ORAN offering by Silicom, and details about N6010 and N6011 product line, and we will finalize with the Silicom Time Synchronization that is an embedded part of all our new NIC and FPGA cards.

Silicom has over 30 years in the market. We trade on the NASDAQ under the name SILC. Silicom's got three R&D centers in Israel, in Denmark, and in the US. Headquarter is based in Israel. We've got over 300 employees. We are collaborating with strategic partners. The main one of them is Intel, who are developing both SOCs and FPGA, both cards and appliances. Silicom is part of TIP, ORAN, OCP, and many standards organizations. Today, we are heavily investing in the TIP, ORAN, and OCP, contributing our designs both in the Edge, FPGA, and RAN solutions.

Looking at Silicom, you can expect from us to see many product lines, and our key value proposition is our rapid innovation. When Silicom is looking at a specific market, we are looking to see how we can create, let's say, a derivative of a product, we would like to create added value. So, when we're looking at a specific market, we are developing our product line to meet the customer needs. We don't charge any NRE. We believe in the business and in the partnership. So, no NREs, the moment that the customer is placing a PO it's sufficient for us to move forward and design with our partners, new products.

Key value proposition for us is the NPI, the process, including a big engineering team that is covering the product from top to bottom, and define all the manufacturing capabilities, and quality assurance to make sure that those are telco grade and high quality products that can sustain availability for at least 10, and even more, years in the field, while doing full certification of the product line. And we're offering a flexible manufacturing location. Today, we're able to manufacture in Israel, China, Philippines, in the US, and India, and if there will be more requirements, we are able to change our manufacturing locations.

Moving forward, the product line that Silicom has includes two main product lines. On the top, we can see we have appliances. Appliances are ranging from ATOM product line for the edge, including fixed wireless access capabilities. We support both LTE and 5G uplink. We have 1U and 2U appliances. The 1U including Xeon SP 3rd Generation that holds a capability to function as a DU and a CU. We have integrated Forward Error Correction on the motherboard based on Intel ACC100 and a Time Sync capability. A range of ports, starting with eight and 12 port, and 25-gig, and we have to add-in slots to cover more ports.

The middle platform, you can see, Silicom Vigo, this is a Xeon D platform. Again, with two slots, we have integration with a Broadcom switch to offer a higher port count.

The Silicom Palma is a 2U platform, similar to the Silicom Barcelona. This one is Xeon SP with Ice Lake processor 3rd Generation Intel processor. It's including also a Time Sync capability with SyncE and 1588, and Forward Correction on the motherboard.

At the bottom, we can see Silicom add-in cards. We're collaborating with major COTS platform suppliers and we provide a STS product family. The STS product family offers a range of product, STS1 with four ports of 10/25, STS2 with eight ports of 10-gig, STS3 with 10/25, eight ports, and STS4 with 12 ports. The range of products is also including Genesis receivers, so our cards can be also a Grandmaster. N6010 is the topic of this presentation today and we'll show much more information. N6010 and N6011, those are FPGA cards, can also support a Time Sync, and we have the FPGA that we can create a full functionality fronthaul connectivity.

We have the Silicom Lisbon, which is an ASIC-based chip from Intel, and we support both 4G and 5G Turbo and LDPC. And at the bottom, we can see the high-end card N5014 for UPF acceleration.

So, the product line, in all, can create opportunity for vendors to adopt our white box solutions and create your own platform, with differentiation using FPGAs and our hardware.

Looking deeper into the ORAN market, on the left-hand side, we can see the opportunity in orange. We're looking at the traditional market, marked in blue, which is decreasing along the time, and we're seeing the Open vRAN market growing across the years. So, currently, we are just in the beginning of this trend, and on the right-hand, we can see the chart that is showing that today, we believe that we are in the early majority. There are plenty of vRAN POCs around the world. I believe today the number is over 30 operators already doing POCs, and some of them are already migrating to a production stage, and we believe that moving forward into 2022, we will move higher in the product, and better integration, and will create much more traction in the market

Silicom's offering for the add-in cards, we can see three types of deployment models. On the left-hand side, we can see the software-only. Software-only, it means that all the Layer 1 is done. On the x86, it means Intel Xeon SP 3rd Generation. Fronthaul connectivity is

done using Silicom STS card. STS, by the way, stands for Silicom Time Sync. So, the STS cards provide connectivity for 10/25 gig, which supports both 1588/SyncE. Of course, with the Grandmaster. We can be ePRTC and we can provide a full connectivity to the radio. We support both C1, C2, C3. It means a direct connection between the DU and the RU, connectivity through a switch for C2 that will provide the Grandmaster, and C3 that the network provides the connectivity between the DU and RU.

Moving to the middle one, we're looking at the eASIC ACC100 from Intel, functions as Look Aside model. This model today is the highest runner in the market. Silicom Time Sync is providing both fronthaul and backhaul connectivity, of course with a Time Sync. It means that we can provide clock directly to the radio and we can provide also clock to the transport network in the back.

On the right-hand side, we are seeing the Silicom N6010. N6010 is an FPGA card. It's including Intel Agilex FPGA and this card can provide Layer 1 acceleration, and those will be the next slides.

I forgot to press, but you can see there all those cards in the right-hand side. The ACC100 does not have any ports because it's a Look Aside model. The Silicom Time Sync, including this one, is STS3 with eight ports of 10/25, and over here we have 1PPS 10MHz and GNSS receiver to function as a Grandmaster. On the bottom we can see the N6010 that includes both the Time Sync and an FPGA.

Moving to the next slide, we can see a full network architecture. We can see on the left-hand side, the Silicom appliance including three types of cards, both the FPGA, the N6011. We can see the ACC100 and the STS cards. Those provide the fronthaul connectivity and acceleration.

On the middle appliance we can see the Silicom CU that has the C5010X. This card provides connectivity and provides also acceleration for different types of workloads, including IPsec and OVS offloading.

On the right-hand side, we can see the network core. Over here, we provide the C5010X, and this one function is a UPF acceleration. It means that we're able to accelerate the user plane function within the 5G core.

Moving to the next slide, I will hand over to Nadav to cover the N6010 and N6011. Nadav, over to you.

Nadav Yafe

Hello everyone. So, why are we here today? We are here to introduce you to the new SmartNIC platform coming up from Silicom, what we call, codenamed N6010/N6011. This is a SmartNIC built for NFVi, VNFs, vRAN, with a focus on Open RAN interface.

This is a high performance acceleration card that enables DPDK and OPAE software packages coming with accelerated workloads that we'll discuss later on. The card architecture is around Intel Agilex FPGA and Intel E810 Columbiaville NIC. The card offers high speed ethernet interface in the ranges of 100, 25, and 10 gigabit ethernet with SyncE support, and ability to support CPRI and eCPRI, which are the interfaces in LTE and 5G networks.

This Agilex FPGA provides PCIe Gen 4 x16 and 16 gigabyte of memory DDR4 for both the logic and the hard processing system. This Agilex provides a quad core ARM A53 to provide several MEC functionality.

Other than that, Silicom incorporated a very state-of-the-art BMC based on Intel Max 10 that is accessed through an SMBus with a PLDM protocol. Time Synchronization is a key factor in vRAN world today and we provide a precision clocking system with 1588 and SyncE to enable Open RAN, low level splits, C1, C2, and C3 timing configuration. As we understand, DU today requires a small form factor to replace the existing base station, that is why we created this card as a full height, half length form factor to meet all the standards in an edge appliance. The card is also optimized for 75 to 100-watt TDP, again, depending on the workload.

Silicom will offer two main SKUs of this Smart NIC, an FPGA-only and an FPGA with Intel NIC.

So, let's discuss a little bit about Agilex. Agilex is the new and the latest and greatest FPGA coming out from Intel fabricated in a 10-nanometer process to enable double the performance over its previous generation with the same feature set that we have seen in the past, DDR4 and DDR5, as well as high bandwidth memory and support in Intel Optane. The hard IP that this chipset provides is a 400 gig ethernet and PCIe Gen 4 today, and later on, PCIe Gen 5 and higher transceiver data rates.

Agilex family that Silicom selected for N6010 is the F-Series. On the screen, you can see the next generations, which is the I-Series and the M-Series, as I mentioned, will offer higher connectivity rate for both fronthaul and backhaul, meaning both PCIe and ethernet.

The Agilex F014 flavor that Silicom selected represents a sweet spot between FPGA logic resources and power and price. This Agilex supports PCIe Gen 4 x16, on the one side, and 100-gigabit ethernet on the other side.

Another major building block in the N6011 with Intel E810 ethernet controller. Intel Ethernet Network Controller E810 delivers up to 200-gigabyte of total bandwidth in PCIe Gen 4 compliant systems. Each port supports up to 100-gigabit and it can be bifurcated to modes of four times 25-gig, eight times 10-gig, and two times 50-gig ethernet.

On here, let's switch over to see what's under the hood.

Silicom provides, alongside the hardware, a flexible FPGA stack that we call OFS. OFS is an FPGA hardware and software infrastructure representing 2nd generation of Intel FPGA platform software. This new software and hardware infrastructure simplifies the developer's experience, making it easier for our customers and partners to develop highly tailored optimized solutions. An already existing ecosystem of Intel OFS enabled boards, workloads, and OS distributors enables faster time to market. A collection of open source accessible hardware and software code is delivered via an open source methodology. The code will be fully open source and upstreamed to the Linux kernel. Whether you're a hardware, or software, or application developer, you can leverage Intel OFS to form the foundation of your development, saving time to innovate and differentiate of your unique value-add.

Intel OFS is a hardware and software infrastructure. All the hardware is open sourced on GitHub. On the hardware side, Intel OFS includes what we call AFU, which is the Acceleration Functional Unit Region. In this region, IP developers can implement their workloads in this, what we call, sandbox, and can contain both static and dynamic regions. Silicom provides a sample AFU to jumpstart your workload development. FPGA Interface Manager provides platform management, functionality, clock and reset standard interface to the host and to the AFU. The FIM is a static region of the FPGA that contains the control logic and the I/O ring. The BMC in this board provides a route of trust and telemetry monitoring support.

From a software side, Silicom provides open source kernel drivers that comprise the first layer of FPGA software stack. OPAE libraries, tools, and APIs which facilitates the development of software application of this accelerator by allowing the host to manage the FPGA accelerator in an abstracted way.

Here, you can see a block diagram of the elements included with OFS. The top half shows the software infrastructure on the Linux driver to user space in application code, and even including the simulation and debug code. Let's zoom in to the lower part.

The bottom half illustrates two examples of FIMs, which is the FPGA Interface Management, both targeting vRAN workloads with slightly different interface requirements. The design on the left supports the interface to the E810 NIC, whereas the one on the right does not. These designs represent typical FIM configuration for the N6010 on the left, and N6011 on the right – sorry, the opposite, for the board variants shown earlier. It's important to note that OFS supports partial reconfiguration. In mode where acceleration workloads, highlighted in purple, can be swapped out dynamically under software control to change system accelerator needs.

Silicom enables solution providers and customers with a variety of workloads. The one we will focus today is vRAN.

The N6010/N6011 supports Silicom Time Sync technology. Through its I/O interfaces, enabling support from all Time Synchronization configurations, C1, C2, and C3 mentioned earlier, making Silicom's N6010 and Open RAN platform, custom designed to meet the real-time processing needs for 5G DUs. Time Sync as well as it's High-PHY Offload IP makes N6010 uniquely positioned to offer a platform that can lower DU costs by reducing the number of CPU cores in memory. With N6010, the TCO, total cost of ownerships, for mobile network operators is up to 50% savings. Silicom's N6010 is a fully programmable and customizable platform that enables lowering cell site power and cost footprint, and easy migration path to eASIC for higher power and cost savings.

On the left hand side, you can see the feature set that I mentioned before, where this picture illustrates connectivity to the cell site where the card is hosted in a DU.

Another use case we see here, and has high demand, is a CSR site router. Using the N6011, we can take any COTS, custom off-the-shelf server and transform it into a virtual CSR or fronthaul gateway. This card provides an optimized edge site integration into standalone 5G deployments. In legacy 4G, we saw requirement to support CPRI to eCPRI conversion in appliances known as fronthaul gateways.

Fronthaul gateways and CSR for routing integration and F2 aggregation towards vDU is a requirement we see today in known – in brownfield deployments.

So, let's talk about 5G vRAN use cases and OFS. This rather busy slide shows how the N6010 is used to provide both FEC and fronthaul/backhaul processing. If you look at the bottom portion of this block diagram, going from left to right, on the left side, you can see fronthaul interfaces used to connecting to the radio units or fronthaul gateway. As mentioned before, the N6010 supports both Split 8 and Split 7.2 configuration. As the packets stream into the card through the QSFP connectors, they arrive into the fronthaul block, which performs inline processing such as packet classification with precise time protocols. All low layer processing for Split 8 configuration. The data flows through, then flows through the Columbiaville NIC across the PCI interface to the host where the Xeon CPU performs Layer 1 processing using FlexRAN software. Data can be sent to the FPGA to the FEC Block for higher Layer 1 processing and will return to the CPU for Layer 2 processing. This offloads the computationally-intense FEC operations, freeing cores for more productive revenue-generating tasks.

If midhaul connection is required, say, to connect to other CUs or DUs, data can be forwarded using the QSFP ports or to a separate SDS NIC in the system. Depending on the traffic model, this system can support 24 to 36 cells, 4T4R 20 megahertz narrow band traffic, or up to three to four cells of 100 megahertz massive MIMO traffic.

Silicom, an industry-leading provider of high performance networking and infrastructure solutions designed N6010 being a flexible platform. It was created to answer a need for a platform that can lower O-DU costs with a need to answer Time Synchronization Open RAN definition. The combination of Silicom's IP that's shown here as Layer 1 High-PHY Offload makes it much easier for vRAN solution providers to implement a complete High-PHY, saving them the need to integrate and test all its modules separately. The most common deployment scenario in today's vRAN environment is 7.2x split. The combination of Silicom's IP and the N6010 Smart NIC platform makes it much easier for a vRAN solution providers to implement a High-PHY solution in their O-DU.

From here, I would like to spend a few minutes and discuss Silicom Time Synchronization IP.

As mentioned before, Silicom identified a problem that was introduced by the Open RAN specification. The Open RAN specification required the DU to have a PTP and SyncE support. What Silicom did, we developed our own unique IP that supports both 1588, SyncE and other Time Synchronization standards to support low level split C1, C2, and C3 configuration.

What you are seeing here on your screen, on the right hand side, taking from the Open RAN specification is C1 and C2 configuration where you can see that the DU is – requires to support SyncE and PTP in a Grandmaster mode where – and actually, it means the DU is synchronizing the radio. C2 configuration means that the DU still synchronizes the radio, but this time, it's for a PTP-aware switch.

In C3 configuration, the DU requires to have transparent or slave clock capabilities, and this is also a mode that Silicom supports today.

Here, on screen, you can see a list of the Time Synchronization standards defined by the Open RAN specification that Silicom supports. Among them is PTP and SyncE, as well as Boundary Clock, Grandmaster, and the ability to be a PRTC Class B Grandmaster.

Other than that, Silicom provides our own unique PTP stack and server software that can be implemented in both x86 or the Agilex ARM.

Thank you everyone for listening to us today. We're very happy that we were given the chance to introduce this very unique and interesting platform and I'll hand over to you back, Lilian. Thank you.

Lilian Veras

Thank you, Nadav, and thank you, Oren, as well for sharing such great information with us all. While you were presenting, a few questions have come in, so let's start our Q&A.

The first question we have here is, "What is the key use case Silicom is targeting?"

Nadav Yafe

OK, so, as I mentioned earlier in the slides, we are targeting several use cases, vRAN being one of the key ones that we'll start, but there are other ones like NFV and OVS and other – and cell site or virtual CSR router workloads.

For vRAN implementation, it's important for our audience to understand that what, basically, Silicom is doing with this platform, we're using this platform for a slot saver, but also optimizing power and cost in the DU. We know that one of the key factors for mobile network operators today when they come to deploy a cell site is the footprint of not only the cost of a cell site, but also power and maintenance footprint later on. Using this FPGA-based platform, we enable exactly that.

Oren Benisty

There are a few more use cases being targeted for this type of card. We engage with the OpenBNG to introduce this card as a small-scale BNG. We have got some demand from the industrial market that are requiring Time Synchronization within FPGA capabilities. And as Nadav showed before, also, we can function as a cell site gateway within a DU.

Lilian Veras

Great, thank you both. Another question we have here, "Can you clarify the difference between N6010 and N6011?"

Nadav Yafe

Of course. We created two SKUs, again, N6011 is FPGA with the E810 NIC. The ability of the NIC, that I did not mention before, is that it comes with DPDK drivers, standard Intel drivers that enables seamless integration of a data path from the FPGA platform to the host.

That being said, there are customers that has their own IP or DNA engines or Virtio interfaces to the host and do not want to use the NIC because of that. So, we created another SKU that creates a higher bandwidth path to the host, a x16 instead of a x8 connection to the host exactly for those types of customers.

Lilian Veras

Thank you, Nadav. Another question we have here, "Can this platform fit into more use cases other than vRAN?"

Oren Benisty

OK, I think that we have covered this one already in some of the slides. As we discussed before, we can function as a cell site router or a cell site gateway, we can function also for a BNG fronthaul solution, industrial automation. We have discussed with the medical companies that are looking to use this type of card beyond the DU market or the DU/CU market. And we are seeing demand also from the financial market for timestamping capabilities.

So, this type of card can function as a forwarding engine with Time Sync capabilities in many markets.

Lilian Veras

Great, thank you. Another question we have here, "Can Silicom offer the N6010/N6011 in a DU platform for POCs?"

Nadav Yafe

The question is yes, of course. We already are in the process of integrating that into several OEM platforms. Among them are Supermicro, Dell, and others, and we can offer a list of already qualified, or partially qualified OEM service.

Lilian Veras

Thank you. A question from the audience, “Are the DU products designed for option two, or do they support 7.2x as well?”

Oren Benisty

Yes, of course. The platform is an x86, of course it was designed for 7.2x fronthaul connectivity, and we support also option two, it means the PDPC will be part of the DU. We have x86, so we can support the Layer 2 and Layer 3 within the same platform. Actually, we are targeting some of the private networks to host both DU and CU within the same platforms.

Lilian Veras

All right, thank you. We do have one last question here from the audience. “Will there be a different performance in connectivity challenges for option two, for example, DU onsite?”

Nadav Yafe

So, the answer is yes. The panel was created with the – to support extended temp and very harsh environments. So, basically yes, the answer is yes.

Lilian Veras

That is great.

Nadav Yafe

So, no performance differences, because the platform is validated to work in very harsh conditions that are used in a DU on the site.

Lilian Veras

All right, that’s great, Nadav. I'd like to thank you both for sharing such great information with us. I would like to ask our audience to please do not forget to give our team a rating for the live recording so that we may continuously improve the quality of our webinars.

Thank you again for joining us today, and this concludes our webcast. Thank you.