



Market Insight Report Reprint

Rambus buys into CXL interconnect ecosystem with two new deals

June 17 2021

John Abbott

The addition of PLDA and AnalogX will speed up Rambus' initiative to enable memory expansion and pooling in disaggregated infrastructure through the emerging Compute Express Link interconnect ecosystem. CXL 2.0 is a PCIe-based technology intended to make it easier to connect CPUs with memory and specialist accelerators, and to separate memory from physical servers to improve memory bandwidth, capacity and efficiency.

451 Research

S&P Global

Market Intelligence

This report, licensed to Rambus, developed and as provided by S&P Global Market Intelligence (S&P), was published as part of S&P's syndicated market insight subscription service. It shall be owned in its entirety by S&P. This report is solely intended for use by the recipient and may not be reproduced or re-posted, in whole or in part, by the recipient without express permission from S&P.

Introduction

Rambus has acquired two companies, PLDA and AnalogX. The silicon memory IP company is looking to boost its portfolio of new products and IP through contributing to the emerging CXL (Compute Express Link) interconnect ecosystem. CXL 2.0, launched last November, is a PCIe-based technology intended to make it easier to connect CPUs with memory and specialist accelerators, and to separate memory from physical servers to improve memory bandwidth, capacity and efficiency.

Rambus intends to combine its existing serial connection, memory and security skills, and IP with the newly acquired IP and engineers from the two companies to produce CXL buffering chips and memory controllers, powering next-generation PCIe 6.0 and CXL 3.0 devices. Simultaneously, Rambus launched a new research and development effort to explore the use of memory expansion and pooling, techniques that the next generation of datacenters could use to implement disaggregated and composable server architectures.

Snapshot

Acquirer	Rambus
Targets	PLDA, AnalogX
Subsector	Semiconductor IP
Deal value	Not disclosed
Date announced	16-Jun-21
Closing date, expected	Q3 2021
Advisers	None disclosed

THE 451 TAKE

CXL is a future prospect rather than a current one, but Rambus must look three to four years ahead of the curve in order to contribute to the product design phases of current processors' memory subsystems and server/cloud datacenter architectures. The ecosystem has developed enough for Rambus to see an opportunity to get its IP built into system architecture designs underway today, particularly in the cloud and on edge devices. Even together, the two new deals are relatively small, but they add valuable expertise and existing customer relationships at a time when design decisions must be made now as a foundation for future products and services.

Deal details

Rambus has signed agreements to acquire both PLDA and AnalogX, expecting to close the deals by the end of the third calendar quarter. Terms were not disclosed for either deal. They won't have a material effect on Rambus' 2021 results due to timing and acquisition accounting, but they are both expected to be accretive in the following year. Rambus was holding \$529m in cash at the end of its fiscal Q1 (April 30, 2021) and generated \$39.5m in cash through operating activities during Q1. It has initiated an accelerated share-repurchase program for \$100m – four million shares – of its common stock.

Target profiles

Although it's far from a household name, PLDA, founded in 1996 by Arnaud Schleich (CEO) and Staphane Hauradou (CTO) in Aix-en-Provence, France, has been a central developer of PCIe silicon IP, and has also been involved in the gestation periods of related and competitive memory interconnects such as CCIX and Gen-Z. There are 80 employees. PLDA claims to have 3,200 customers and 6,400 licenses in 62 countries, with offices in Silicon Valley, Bulgaria, Taiwan and China, as well as in France.

The PLDA Group is also the parent company of Accelize, the cloud-based distribution platform for FPGA-related software IP, which was spun off from its QuickPlay FPGA development tools business at the end of 2016. Last year, PLDA was one of the endorsers of the memorandum of understanding between the Compute Express Link Consortium and the Gen-Z Consortium promising future collaboration and compatibility. PLDA had already developed Gen-Z Controller IP, and more recently, introduced a CXL equivalent: XpressLink CXL Soft IP Controller, intended to be implemented on either ASICs or FPGAs.

Toronto-based AnalogX, founded in 2017, has specialized in ultra-low-power interconnect IP for AI processors, networking, optical interfaces and datacenters through its multiprotocol SerDes chip IP. It is intended to augment Rambus' existing PCIe 5.0 and 32G multi-protocol PHYs, adding the ultra-low power and low-latency capabilities developed at AnalogX, which has specific expertise on DSP-based designs and PAM4 signaling – both necessary building blocks for CXL.

Robert Wang, AnalogX cofounder and CEO, was previously director of analog IP at Intel, SerDes architect at Toronto's V Semiconductor (acquired by Intel in 2012) and senior project director at another Canadian silicon IP company, Snowbush Microelectronics. Rambus acquired Snowbush (via Semtech) in June 2016 for \$32.5m. His fellow cofounders also come from V-Semiconductor and Snowbush. There are 22 employees.

Acquirer profile

Rambus was founded in 1990 and became known for the invention of RDRAM, a type of synchronous dynamic memory. It was backed by Intel in 1996, but ultimately lost out to an alternative memory standard – DDR SDRAM. There followed a long period of patent battles and litigation. In 2014, a change of strategy saw the settling of many long-running disputes, and a drive to develop a new portfolio of memory IP and products, especially memory buffer interconnects.

Rambus still makes significant revenue from royalties and patents, but its product revenue now makes up the larger portion. In its first quarter of 2021, its total Q1 revenue of \$70.4m was made up of \$28.9m from royalty revenue, \$30.8m from products (mostly buffer chips) and \$10.7m from contract and other revenue (primarily silicon IP).

Rambus says the quarter saw a rebound in memory consumption after some short-term inventory digestion. Demand for DDR4 server memory is expected to remain high, with DDR5 modules for both AMD and Intel platforms currently in end-customer qualification. Supply uncertainties are by no means over, but Rambus says it remains confident in the robustness of its supply chain.

Much of Rambus' new business is datacenter-driven. The shift to the cloud and the demand from AI and ML workloads to feed systems with data on the fly is driving demand for more memory bandwidth. Memory subsystems need to be rethought to achieve this. At the same time, the general industry shift to DDR 5 memory is beginning.

Rambus expects to see its share of chip content per module increase as this transition progresses, because higher-speed memory requires such additional components as temperature sensors, power management ICs, and serial presence detector hubs that are in its portfolio. Specialist AI and ML processors require high-speed interfaces such as HBM (high bandwidth memory) and GDDR6 SDRAM.

Deal rationale

CXL extends the PCI-e bus into a fabric capable of linking heterogeneous CPUs, GPUs, FPGAs and ASIC accelerators, as well as memory. It's a coherent memory interconnect enabling processors and accelerators to recognize CXL memory and utilize it as main memory. The likely result if it becomes established is that memory and storage resources will start to migrate away from the inside of servers, and instead be disaggregated and deployed in racks of just memory.

That has a number of advantages. Some 50% of the cost of current servers is accounted for by memory, and because it is siloed, a lot of it remains unused. Memory bandwidth will increase, and memory resources will become 'composable' – in other words, more fluidly applied to jobs that need lots of memory. It will be easier to manage and upgrade memory. Compute can be placed closer to memory, particularly in edge devices. And because of the memory coherency, programming models between heterogeneous processors should be simplified.

CXL 1.0 came out in 2019. Initially Intel-led, it soon gained support from all the main chipmakers (AMD, ARM, IBM, Microchip, Xilinx), cloud providers (Alibaba, Facebook, Microsoft, Google) and system makers (Dell EMC, Hewlett Packard Enterprise, Huawei). Version 2.0 of the specification came out last year, significantly extending the concept with switching (connecting multiple hosts with multiple resources and enabling pooling), support for distributed persistent memory, and security.

Rambus already has buffer chips for DIMM memory modules, connecting CPU and DRAM. It's now doing the same thing for CXL, developing PHYs and memory controllers (to interface with host processors and other devices), building DDR memory PHYs and controllers (to interface with memory devices) and also adding security IP in the form of cryptographic cores and secure protocol engines, to protect the otherwise exposed links against data tampering and physical attacks.

It doesn't expect to start seeing a significant contribution from CXL revenue until 2023, moving into 2024, generated from the sale of CXL buffers. But designs for CXL chip architectures and memory subsystems are starting now. Aside from the buffer chips themselves, Rambus has memory controller, 32G SerDes PHY (serializer/deserializer, physical layer transceiver) silicon IP, and security IP that other chip makers will require to design their products. Combined with the DDR5 transition, the expectation from Rambus is that the total addressable market for its buffer chip business will have doubled by 2024.

The new deals also build on previous Rambus acquisitions, including the security IP, secure protocols and provisioning from Verimatrix (\$65m, September 2019), the memory controller IP from NorthWest Logic (July 2019) and the memory interconnect unit of Inphi (\$90m, June 2016).

Competitive landscape

CXL is not the only game in town. There is also CCIX (the Cache Coherency Interconnect for Accelerators) led by Xilinx; Gen-Z led by HPE but also supported by Dell, NVIDIA's NV-Link, IBM's Open CAPI, and AMD's Infinity Fabric. Each has its own design points. However, CXL appears to be the most promising. The memorandum of understanding between CXL and Gen-Z advocates indicates that compromises will be made to avoid technology wars. Gen-Z will likely extend CXL across racks and datacenters, supported by optical interconnects using silicon photonics.

Other memory companies are taking notice. Last month, Samsung claimed a first with the development of a CXL-based DDR5 memory module, packaged in an EDSFF solid state drive form factor. Samsung said that the device would scale the memory capacity of servers to terabyte levels, increase memory bandwidth and reduce the latency caused by memory caching. It's been validated using Intel servers, but there's no word on commercial availability.

CONTACTS

The Americas

+1 877 863 1306

market.intelligence@spglobal.com

Europe, Middle East & Africa

+44 20 7176 1234

market.intelligence@spglobal.com

Asia-Pacific

+852 2533 3565

market.intelligence@spglobal.com

www.spglobal.com/marketintelligence

Copyright © 2021 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively, S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers, (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages.

S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global Market Intelligence does not endorse companies, technologies, products, services, or solutions.

S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain non-public information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its websites, www.standardandpoors.com (free of charge) and www.ratingsdirect.com (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at www.standardandpoors.com/usratingsfees.