

What is the difference between inline and lookaside accelerators in virtualized distributed units?

According to LightCounting/TÉRAL RESEARCH, the global 5G Open Virtualized RAN (vRAN) market is expected to triple in size by 2027. This rapid growth suggests that many mobile network operators (MNOs) are currently evaluating which options are best for them. This paper will outline the advantages and disadvantages of the two architectural choices open to them.

Unlike traditional RAN, which uses purpose-build appliance hardware for baseband processing, vRAN disaggregates the hardware from software. By separating software from hardware, MNOs have the flexibility to use commercial off the shelf (COTS) hardware from different manufacturers. With increased competition from multiple vendors, MNOs not only get the benefits of better pricing and faster innovation, but they also benefit from improved supply chain resiliency. In addition, by

disaggregating the hardware from software MNOs can not only realize the continuous improvement and continuous development (CI/CD) and containerization benefits of cloud-native solutions, they can also pool hardware in large sites or a centralized RAN configuration to better scale compute capacity with demand, reduce capital expenditures (CAPEX), optimize power consumption and increase redundancy. To recognize all of these benefits, MNOs need to be able to source all hardware from multiple vendors.

vRAN is often (but not always) implemented in conjunction with Open RAN. While vRAN disaggregates hardware and software, Open RAN separates the base station into three functional elements — Radio Unit (RU), Distributed Unit (DU) and Centralized Unit (CU) — and specifies standardized interfaces so that these elements can be purchased from different vendors.



If CU and DU are virtualized, then they are referred to as vCU and vDU. As with the CU and DU, the vCU provides support for the higher layers of the protocol stack, while the vDU provides support for the lower layers such as Radio Link Control (RLC), Medium Access Control (MAC) and physical layer (PHY).

The demand for exceptional computing power and lightning-fast latency is non-negotiable in the dynamic realm of radio access networks, where the vDU Layer 1 (L1) processing orchestrates a symphony of intricate algorithms such as Forward Error Correction (FEC), channel estimation, modulation, and layer mapping. This is precisely where accelerator cards step in, bridging the gap between the complexity of these tasks and the limitations of conventional General-Purpose Processors (GPP) found in off-the-shelf hardware. In the same way that Graphics Processing Units (GPUs) revolutionized the realm of graphics-intensive gaming, accelerator cards offload physical layer processing from the Central Processing Unit (CPU) to elevate the DU's performance.

There are two different types of accelerator architectures, and each architecture has cards that perform different functions. The primary commercially available 'lookaside'

solution uses accelerator cards that are primarily limited to FEC, which we will refer to as the FEC accelerator. On the other hand, 'inline' architectures use accelerators that not only perform FEC but almost all layer 1 processing, which we will refer to as layer 1 or L1 accelerators. In most cases, inline architectures offer the best options for MNOs because they enable the highest performance, lowest cost, lowest power consumption and greatest flexibility.

Comparing inline and lookaside architectures

In the lookaside architecture, the CPU acts as the master controller for layer 1 processing. The accelerator is often a separate peripheral component interconnect express (PCIe) card in the server that is only used for select functions like FEC, which are effectively outsourced by the CPU. Many of the layer 1 real-time computations continue to be processed by the CPU in addition to the layer 2 and layer 3 processing. While the CPU is efficient at layer 2 and layer 3 processing, it is not as effective at layer 1 processing when implemented in a cloud-native architecture. That's because cloud-native software focuses on development efficiency and maintainability, so it is not always suitable for processing that requires high L1 load and low processing latency.

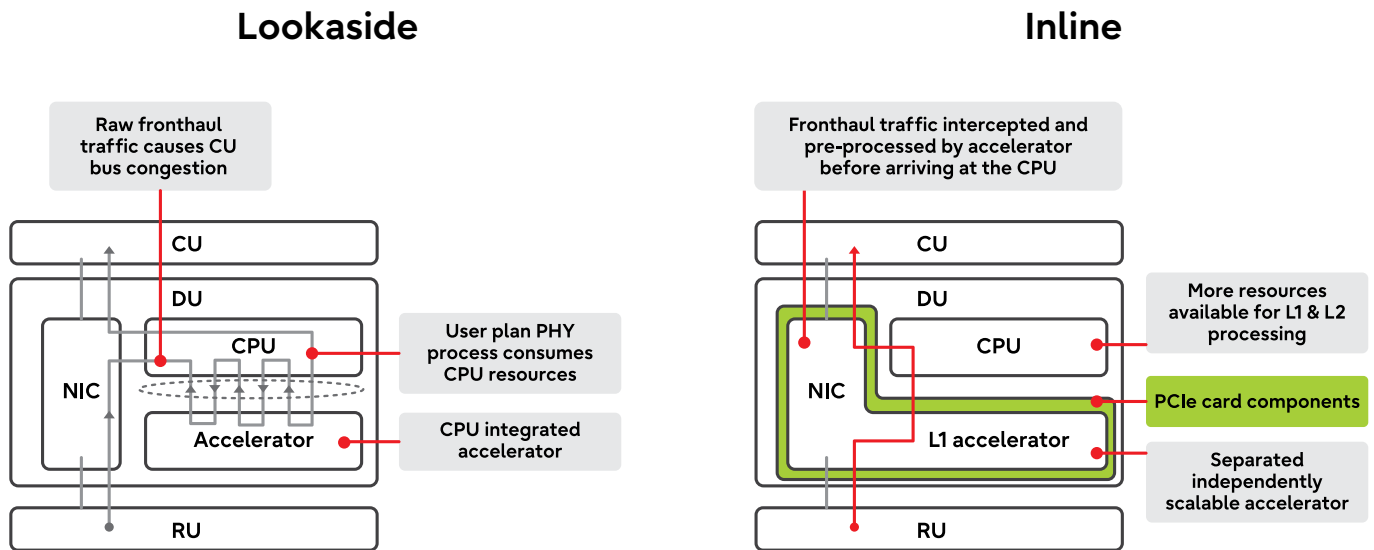


Fig 1: Lookaside and inline accelerator architecture comparison

In the inline architecture, layer 1 processing for the user plane data (which is almost all of the L1 data) is intercepted by the L1 accelerator before reaching the CPU. Inline L1 accelerator cards can perform layer one processing more efficiently than in a general-purpose CPU, and can do so more cost-effectively particularly at high capacity. By

processing almost all layer 1 data in the L1 accelerator, the general-purpose CPU resources are freed up and applied to layer 2 and layer 3 processing. The following table summarizes where the processing for different functions is performed.

Category	Items	Lookaside	Inline vendor A	Inline vendor B
FEC Encoder	LDPC	Accelerator	Accelerator Card	Accelerator Card
	Polar	CPU	Accelerator Card	Accelerator Card
FEC Decoder	LDPC	Accelerator	Accelerator Card	Accelerator Card
	Polar	CPU	Accelerator Card	Accelerator Card
Layer Mapping	Precoder	CPU	Accelerator Card	Accelerator Card
Modulation	DL Modulation	CPU	Accelerator Card	Accelerator Card
	UL Modulation	CPU	Accelerator Card	Accelerator Card
Beam Weighting	Beam Weighting	CPU	Accelerator Card	Accelerator Card
	Weight Calculation	CPU	Accelerator Card	Accelerator Card
Equalizer	Channel Est.	CPU	Accelerator Card	Accelerator Card
	Freq Offset Comp.	CPU	Accelerator Card	Accelerator Card
PRACH	Correlation calc.	CPU	Accelerator Card	Accelerator Card
	Detection	CPU	Accelerator Card	Accelerator Card
Other Layers	L2 and upper, OAM	CPU	CPU	CPU
NIC	Fronthaul	NIC Card	Accelerator Card	Accelerator Card
	Mid haul	NIC Card	Accelerator Card	NIC Card

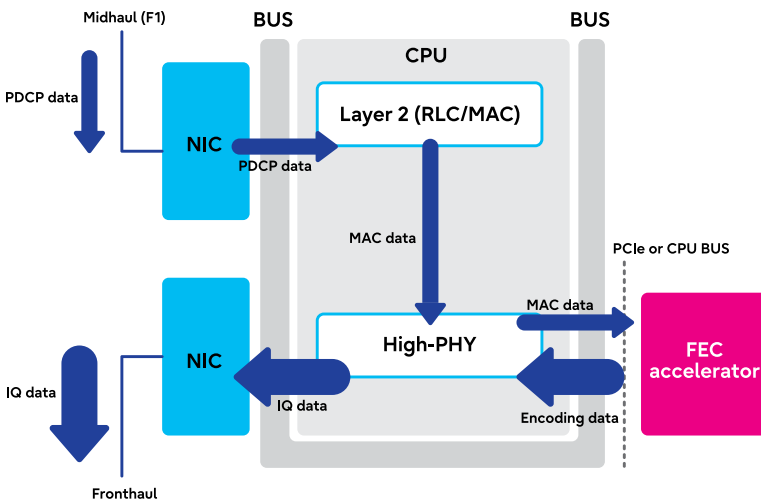
Fig 2: Lookaside vs inline accelerator functions

Since layer 1 processing scales with the air interface, the amount of layer 1 processing increases with bandwidth or the number of antennas connected. Layer 2 and layer 3 processing requirements scale with the number of mobile devices or User Equipment (UEs), number of connections and the volume of traffic. With a lookaside acceleration architecture, all user plane data passes through the CPU,

consuming resources that would be better used for increasing UEs, connections or traffic. With an inline architecture, user plane data remains exclusively on the L1 accelerator card.

Another difference between lookaside and inline accelerators is the way in which they use the bus interconnection inside the CPU.

Lookaside acceleration



Inline acceleration

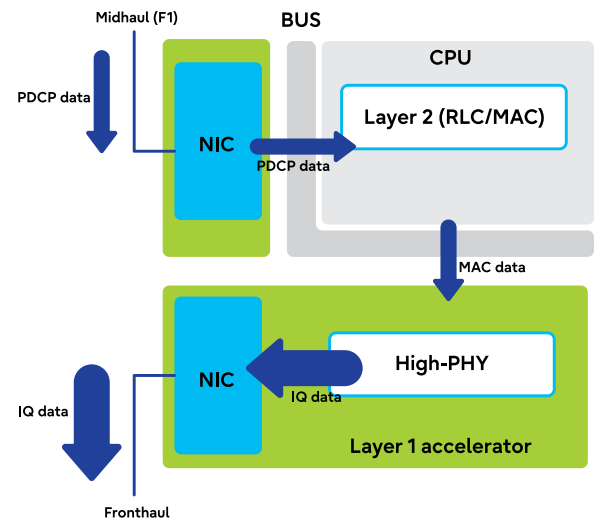


Fig 3: Lookaside and inline architectures with downlink data flow

As the diagram above illustrates, in the downlink for a lookaside architecture the vDU receives Packet Data Convergence Protocol (PDCP) data from the vCU, which is sent from the Network Interface Card (NIC) through the bus to the CPU core. The PDCP data is processed by the RLC, which is responsible for data reliability and flow control, and then the MAC manages and schedules the radio link before passing the data to the layer 1. The High PHY layer function, located in the DU, converts the data into a radio signal, which significantly increases the data capacity by a factor of two. This High PHY data volume increase is due to IQ modulation and the addition of error correction bits in encoding. In the

inline acceleration architecture, encoding happens in the latter stage just before exiting to the fronthaul transport system, and as a result, it does not use the bus inside the CPU.

So, why does the frequency and volume of data traffic on the bus matter? Well, the bus is a shared resource; therefore, the more often it is used, the greater the probability of conflict and delay. The result of this delay can be seen in the following chart, which empirically shows the processing time resulting from two different vendors' inline L1 accelerator cards compared with another vendor's lookaside FEC accelerator.

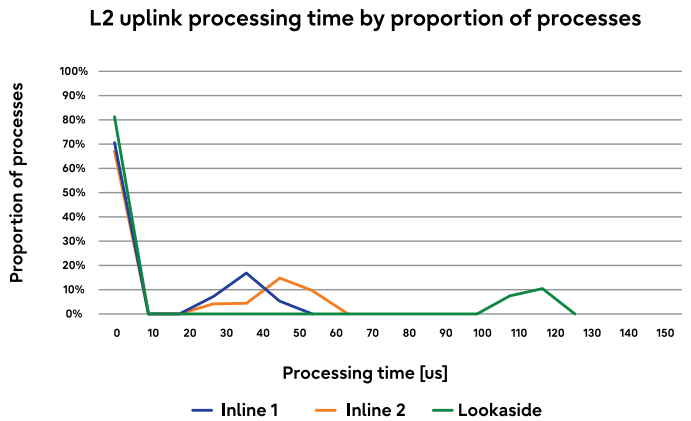
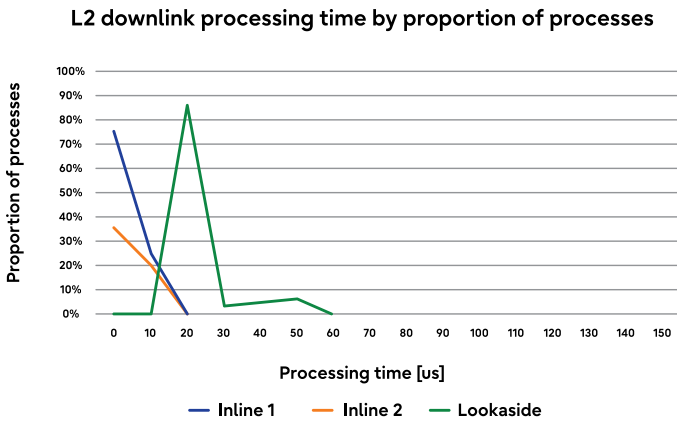


Fig 4: CPU processing time comparison

The graph on the left illustrates the downlink processing delay from lookaside architectures. This is primarily due to delays associated with bus conflicts. The graph on the right shows the processing delay from lookaside architectures that is associated with the uplink.

The delay in the uplink can be much worse for a portion of the processes. Though the delay can be measured empirically, it is difficult to prove the reason(s) for it conclusively. A probable source of the delay relates to the way the cache is consumed in a lookaside architecture. Each CPU core has a cache allocated to it as shown in the diagram below. This primary cache on each core can be accessed very quickly and typically only requires less than 10 CPU clock cycles, but this cache is quite small. If the core is designated to perform the type of intensive processing needed for layer 1, this designated cache can be quickly consumed. Fortunately, additional secondary cache is available; however, the secondary cache typically requires 5-10 times as many CPU clock cycles to access, as compared to the primary cache.

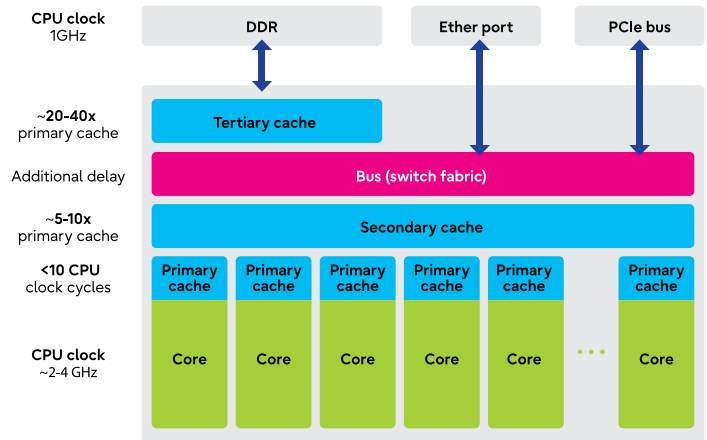


Fig 5: CPU cache

Moreover, since a portion of the cores are performing intensive layer 1 processing, the capacity of the secondary cache can be clogged with data needed for layer 1 processing.

On the plus side, there is a tertiary level of cache. However, accessing the tertiary cache requires transiting a bus, which introduces significant delay and can require up to approximately 20-40 times as many CPU clock cycles to access, compared to the primary cache.

The resulting processing delay translates into lower performance and capacity of the vDU. This delay can inhibit very latency sensitive applications. For example, Fujitsu has developed a vDU feature that optimizes the allocation of CPU core resources to reduce the processing time. By reducing the processing time required in the vDU, the latency budget on the fronthaul can be increased, enabling MNOs to locate RUs up to 50km away from the vDU. Due to its delay, this would not be possible with a lookaside accelerator architecture. To ensure consistent, reliable performance of time-sensitive features under high-capacity loads, these applications are best implemented on inline L1 accelerators that do not suffer the kinds of processing delay described above.

The delays described assume there is contention for cache and/or bus resources. In sites with less traffic and fewer cells with less bandwidth, the frequency of contention is likely lower and perhaps decreases the delay.

Additional Advantages of Lookaside FEC Accelerators

L1 accelerators are designed with the assumption that most layer 1 processing needs to be performed separately so as not to interfere with layer 2 and layer 3 processing being performed on the CPU. However, this is not always the case. vDUs on small de-centralized sites may not need to support as many cells with as much bandwidth or as much traffic. In these instances, merely offloading FEC functions may be adequate. If the accelerator card performs fewer functions on less traffic, it is possible that it might not consume as much power. Under high traffic/high bandwidth circumstances this would be offset by increased power consumption in the CPU. Yet for low traffic, low bandwidth sites with a small number of cells, it is possible that lookaside architectures could offer lower power consumption.

Today the lookaside architecture is only employed by a single vendor, and both CPU and accelerator are provided by them in their implementation. There can be some advantage to having a simple all-in-one solution. For example, there would be fewer parts/inventory to manage and no mismatched development cycles to handle.

Additional Advantages of Inline Layer 1 Accelerators

As mentioned earlier, since inline L1 accelerators pre-process much of the layer 1 processing, they free up valuable CPU resources to perform more layer 2 and layer 3 application processing. This enables the vDU to support more traffic,

more UEs and/or more connections. Alternatively, capacity levels can be maintained using fewer cores and memory resulting in less power consumption. Moreover, by using an inline L1 accelerator for layer 1 processing, MNOs can optimize both performance and expenditure by harnessing the power of less complex and more cost-effective CPUs for layer 2 and layer 3 processing. By separating the layer 1 processing to be processed on the inline L1 accelerator from the layer 2 and layer 3 processing performed on the CPU, the different components can perform what they do best. With its silicon specifically optimized for layer 1 processing, the inline L1 accelerator is best suited to do so, while the CPU is better suited for layer 2 and layer 3 processing.

With inline acceleration, layer 1 capacity can scale independently from layer 2 and layer 3 capacity because there is more disaggregation between layer 1 and layer 2/3 processing. This offers MNOs the flexibility to select server hardware with a CPU capacity better sized for their expected layer 2 and layer 3 capacity needs. With inline L1 accelerators, MNOs have the flexibility to select server hardware that enables them to optimize for a combination of capacity, energy efficiency and/or cost. By separating layer 1 from layer 2 processing and offloading layer 1 processing to the L1 accelerator, MNOs have greater flexibility of COTS server hardware from which to choose. Specifically, they can choose servers with fewer or less powerful cores for low traffic sites with fewer UEs, or select servers with more or more powerful cores for high traffic sites with more UEs. In addition, MNOs have the flexibility to choose either x86 or ARM processors from multiple different vendors.

The same logic can be applied to scaling layer 1 processing. So, if an operator plans to deploy lots of massive MIMO (mMIMO) configurations, that require more layer 1 processing, it makes sense to do that extra processing on a L1 accelerator card which is more cost-effective for layer 1 processing. In this case, the vDU will be limited by the capacity of the L1 accelerator card. To avoid excess capacity in the CPU, MNOs planning to deploy lots of mMIMO can purchase servers with fewer CPU cores to better balance the ratio of layer 1 processing to layer 2/3 processing.

In contrast, lookaside architectures offer two options. The first option is when the FEC accelerator card is independent of the CPU core. In this case, the vDU would be limited by the capacity of the FEC accelerator. However, it would still require a higher quantity of more expensive CPU cores to perform the additional layer 1 processing being performed in the CPU versus an inline model. Thus, an inline model allows MNOs a more economical solution.

Alternatively, a second lookaside option involves the CPU and accelerator card combined in a single unit. In this case, the combined unit would still be limited by its ability to perform FEC processing. However, by combining CPU and accelerator together, there is no longer the flexibility to scale the number of CPU cores needed to accompany the accelerator. An inline model, on the other

hand, allows MNOs more flexibility to balance the relative capacity of layer 1 and layer 2/3 processing capabilities. This is particularly true when the CPU core is merged with the accelerator card.

So, why are scaling and capacity important? Why can't an MNO just add another server when needed? Well, some of the greatest advantages of virtualization are the benefits of pooling. Whether an MNO is deploying in a Centralized RAN (C-RAN) architecture or deploying large sites with several frequency bands in a Distributed RAN (D-RAN) architecture, there are several strong economic reasons to pool vDU resources. In any capacity planning exercise, it can be difficult (and sometimes risky) to design hardware requirements to exactly the right level. As a result, network planners often do budget some excess capacity. The greater the capacity of the vDU, the more they can pool, requiring

less excess capacity overhead. By reducing unused excess capacity, the MNO can not only save CAPEX, by avoiding unnecessary computing resource purchases, but they can also save all the associated operational expenses (OPEX) that accompanies reduced hardware, including power, heating or cooling, space and maintenance.

Although more testing is required to assess the energy efficiency of these two solutions when they support a minimal number of cell sites, inline acceleration architecture is more energy-efficient when supporting many cells due to the advantages described above.

Conclusion

As has been shown, an inline L1 accelerator offers MNOs a lower total cost of ownership (TCO) based on the following factors:

Category	Lookaside FEC acceleration	Inline Layer 1 acceleration
Capacity	Low-Medium	High
Energy Efficiency	Best on low capacity Sites	Best on medium-high capacity sites
Scalability	Less Flexible (CPU needs to support both L1 and L2/L3)	L1 scales independently from L2/L3
Ecosystem	Currently Single Vendor	Multi-vendor
Complexity	Simple All-in-One	Maximum Flexibility with multiple options
Cost Effective L1 Processing	Less effective (Not as relevant in low-capacity sites)	Most effective
Delay	Meets 5G Fronthaul Requirements	Better than 5G Fronthaul Requirements

Fig 6: Lookaside vs inline summary

About Fujitsu

Fujitsu Network Communications is a leading provider of digital transformation solutions for network operators, service providers and content providers worldwide. We design, build and manufacture best-in-class hardware and software to enable cost savings, faster services delivery and improved network performance. With our extensive multivendor service expertise, we also design, build, operate and maintain better networks for the connected world.

Learn more @ us.fujitsu.com/telecom

Contact

Fujitsu Network Communications, Inc.
2801 Telecom Parkway,
Richardson, TX 75082
Tel: 888.362.7763
us.fujitsu.com/telecom

©Copyright 2023 Fujitsu Network Communications, Inc. FUJITSU (and design)® are trademarks of Fujitsu Limited in the United States and other countries. All Rights Reserved. All other trademarks are the property of their respective owners.

Configuration requirements for certain uses are described in the product documentation. Features and specifications subject to change without notice.