

APPLIED RESEARCH

Accelerating Innovation in 6G Research: Real-Time Capable SDR System Architecture for Rapid Prototyping

MAXIMILIAN ENGELHARDT¹, SEBASTIAN GIEHL², MICHAEL SCHUBERT¹,
ALEXANDER IHLOW^{1,2}, CHRISTIAN SCHNEIDER^{1,2}, ALEXANDER EBERT^{1,2},
MARKUS LANDMANN¹, GIOVANNI DEL GALDO^{1,2}, (Member, IEEE),
AND CARSTEN ANDRICH²

¹Fraunhofer Institute for Integrated Circuits IIS, 91058 Erlangen, Germany

²Institute of Information Technology, Technische Universität Ilmenau, 98693 Ilmenau, Germany

Corresponding author: Sebastian Giehl (sebastian-wilhelm.giehl@tu-ilmenau.de)

This work was supported in part by Bavarian Ministry of Economic Affairs, Regional Development and Energy in the project “DSAI;” in part by the Federal Ministry of Education and Research of Germany in the project “6G-ICAS4Mobility” under Grant 16KISK241; in part by the Federal State of Thuringia, Germany; and in part by European Social Fund (ESF) in the projects “ML4ASP” and “MOTA” under Grant 2019 FGI 0031 and Grant 2018 FGI 0041.

ABSTRACT The upcoming 3GPP global mobile communication standard 6G strives to push the technological limits of radio frequency (RF) communication even further than its predecessors: Sum data rates beyond 100 Gbit/s, RF bandwidths above 1 GHz per link, and sub-millisecond latency necessitate very high performance development tools. We propose a new SDR firmware and software architecture designed explicitly to meet these challenging requirements. It relies on Ethernet and commercial off-the-shelf network and server components to maximize flexibility and to reduce costs. We analyze state-of-the-art solutions (USRP X440 and other RFSoc-based systems), derive architectural design goals, explain resulting design decision in detail, and exemplify our architecture’s implementation on the XCZU48DR RFSoc. Finally, we validate its performance via measurements and outline how the architecture surpasses the state of the art with respect to sustained RF recording, while maintaining high Ethernet bandwidth efficiency. Building a 6G integrated sensing and communication (ISAC) example, we demonstrate its real-time and rapid application development capabilities.

INDEX TERMS 6G, integrated sensing and communication (ISAC), rapid prototyping, RFSoc, software-defined radio (SDR), system architecture, wideband streaming and processing.

I. INTRODUCTION

Working towards the upcoming Release-19, the 3GPP has introduced integrated sensing and communication (ISAC) into the work plan for 5G NR [1]. A feasibility study on the chances and requirements of ISAC has been conducted beforehand [2] and another one, targeting the relevance of channel modeling for ISAC, is ongoing [3].

The ITU, contributing research organizations, and leading hardware vendors do agree, that ISAC will also be part of the emerging 6G standard [4], [5], [6], [7], [8], [9],

The associate editor coordinating the review of this manuscript and approving it for publication was Francisco Rafael Marques Lima¹.

[10], [11], [12], [13]. Their use case visions define 6G’s technological cornerstones: Sum data rates in excess of 100 Gbit/s, radio frequency (RF) bandwidths in excess of 1 GHz, sub-millisecond latency, and possibly full-duplex radio communication. The increased complexity and scope of 6G mandates a highly efficient research and development process to sustainably deliver the required results in time [14], [15], [16].

In contrast to application-specific monolithic architectures, the software-defined radio (SDR) concept [17] provides the flexibility to cover a wide range of requirements, making it a popular choice in research and development. However, the ease of use of existing ready-to-run SDRs in ecosystems such

as GNU Radio or MATLAB is counterbalanced by limitations in terms of throughput and latency, which would, e.g., require individual adaptation in low-level system programming with C/C++ or even field-programmable gate array (FPGA) design. The performance requirements of 6G, in particular the high data rates and bandwidths paired with low latency, certainly necessitate the low-level programming approach. Our own research endeavors of the last two decades [18], [19], [20], [21] reinforce our conviction that SDR represents the best approach to mobile communication research and development.

In this contribution, we propose a novel SDR system architecture that bridges the gap between support for rapid application development and sustained high performance, i.e., high throughput and low latency. To achieve the required performance, we still have to rely on a highly parallel low-level FPGA design and heavily optimized C/C++ software where necessary. We establish zero-overhead interfaces for real-time signal processing modules. All non-performance-critical parts, especially application-specific control code, are implemented as high-level scripting using a concurrent interface that provides deterministically timed access to the hardware functions, e.g., RF frontends, sample acquisition, and waveform generation.

This architecture offers an unparalleled set of features compared to other state-of-the-art (SotA) solutions [22], [23], [24], [25], [26], [27], [28]:

- *Strong isolation between scripted application logic and the hardware abstraction layer* resolves performance issues of SotA solutions where high-level scripts interfere with low-level realization. Additionally, realizing application logic in high-level scripts reduces the initial hurdle for developers and enables rapid application development.
- *Full exploitation of the hardware's capabilities.* The high-speed serial link between FPGA and host server is the system bottleneck. Therefore, the streaming efficiency in receive (Rx) and transmit (Tx) is optimized towards the link's data rate limit. Beyond that, a separate solution (comparable to an arbitrary waveform generator) for periodic Tx signals allows simultaneous use of all DACs at the maximum sample rate.
- *Continuous sample recording only limited by the size of the SSDs.* Using hardware comparable to the SotA, the architecture realizes uninterrupted recording of a 10 GB/s sample stream via a single 100 Gbit/s Ethernet link.¹

The remainder of the paper is structured as follows: Starting from the state of the art in Sec. II, we derive the design goals that are required to establish such an architecture in Sec. III. Afterwards, we point out our design decisions and the resulting SDR system architecture in Sec. IV to VIII. Finally, we portray our implementation on a Xilinx RFSoc and commercial off-the-shelf (COTS) server hardware in Sec. IX and discuss the

¹This corresponds to a real-valued 16 bit sample signal at 5 GSa/s or an equivalent complex 2x16 bit sample baseband signal at 2.5 GSa/s. Due to analog and digital filter roll-off, the modulated bandwidth of either signal is limited to approximately 2 GHz.

achieved performance based on system benchmarks in Sec. X and an exemplary 6G ISAC measurement in Sec. XI.

II. STATE OF THE ART

Table 1 enumerates a selection of architecture solutions and their realization. The existing approaches can be split into two groups, namely dedicated monolithic architectures (1. and 2.) and classic SDR systems (3. through 5.).

Monolithic solutions are meant to be composed of dedicated hardware framework components. Architecture approaches like these usually – due to the fine-tuned component selection – reach outstanding performance, but lack flexibility due to hardware dependencies and proprietary data interfaces. Adapting the system to different applications therefore requires changing its setup in components and low-level programming. This results in prolonged development cycles and increased cost.

On the other hand, classic SDR approaches gain flexibility by realizing a split architecture design [17]. They define a protocol layer, which realizes three planes: Control, Rx, and Tx. The transceiver hardware must be compatible with the protocol; aside from that, it is interchangeable. Using an appropriate data link, it is connected to a host personal computer (PC), where a software driver interacts with the transceiver hardware to realize generic functionality. The driver offers a high-level interface to the user. The actual application-specific code can be built on this interface and is therefore largely independent of the utilized transceiver hardware and vice versa. In general, this enables designing a powerful and hardware independent architecture, which can utilize popular standard high-speed interfaces and COTS hardware. As trade-off against the increased flexibility, many SDR architecture realizations suffer from performance limitations.

Solutions 3 to 5 utilize the Xilinx RFSoc chipset, which combines direct RF-sampling data converters with FPGA hardware. It offers eight analog-to-digital converters (ADCs) and digital-to-analog converters (DACs) with up to 5 GSa/s each and two 100 Gbit/s Ethernet interfaces. Comparing the realized figures of merit, the impact of architecture design can easily be seen: Solution 3 utilizes direct RF-sampling, but also digital down conversion. This allows for a possibly high channel count, but cuts down the realizable instantaneous bandwidth. Furthermore, the underlying architecture only implements receiver operation.

In contrast, solutions 4 and 5 represent the two RFSoc-based SDRs available as COTS devices from Ettus Research. Both are ready to run solutions and include the necessary peripheral components and interconnects. The Ettus open source software suite USRP Hardware Driver (UHD) can be utilized for both devices. The UHD is accompanied by the RF Network on Chip (RFNoC) architecture, which provides partially run-time reconfigurable hardware acceleration for signal processing by exploiting the computational capabilities of the FPGA. Although it simplifies the FPGA design flow through its modular structure, it increases the overhead in

TABLE 1. Comparison of the figures of merit of various wideband sampling solutions.

Index	System	Max. Instantaneous Bandwidth / MHz	Max. Rx Channels	Data Rate Gbit/s	Interface
1.	Keysight Digitizer [22]	12 500	4	160	Optical data interface (ODI)
2.	Teledyne ADQ7DC [23]	3000	2	56	PCI Express
3.	Universal RFSoc-based Signal Recorder [24]	256	8	10	10G Ethernet
4.	Ettus USRP X410 [25]	400	4	200	2x100G Ethernet
5.	Ettus USRP X440 [26]	1600	8	200	2x100G Ethernet
6.	KIT MIMO testbed [27], [28]	2000	8	200	2x100G Ethernet

terms of resource usage and data rate. Design decisions required at compile-time limit the flexibility that can be achieved. The X410 integrates an analog frontend for mixing and filtering which limits the bandwidth per channel. In contrast, the X440 directly connects the balun coupled chassis inputs to the converters' pins and therefore supports enhanced bandwidths. Ettus specifies the reference system to be only capable of utilizing the available two 100 Gbit/s Ethernet links to up to 61%. The worst case configuration even reaches only 34% of link utilization [29], which severely degrades performance. Furthermore, as shown in [30], although Ettus' driver allows to specify the execution time of commands, the architecture does not support parallel synchronous switching operations. Moreover, the driver's interface is not designed for burst Tx streaming, massively increasing the control overhead for the high-level script in applications such as multiple input multiple output (MIMO) channel sounding.

To overcome the existing SDR systems' data rate limitations, solution 6 is being developed in [27] and [28]. The RFSoc-based MIMO testbed presented there achieves outstanding performance, utilizing the links between host and converter device to over 90%. However, depending on the application, further optimization is desirable to reduce the number of samples transferred over these links: In the Rx path, burst sampling is not supported, and in the Tx path, periodic signals have to be streamed continuously through the host. Due to the targeted application, this system is not designed for continuous streaming and storage on SSDs, but buffers Rx samples in RAM during the acquisition, which limits the measurement duration to a maximum of 10 s [28].

III. DESIGN GOALS

In order to design a platform as versatile as possible, the basic functionality should be isolated from the interchangeable application logic. Similar to the UHD, we aimed for a high-level application programming interface which allows to control all low-level functionalities from a programming language of the users' choice, such as Python. The application logic can thus be modified without recompiling FPGA design or C++ code, also obviating the need to flash and to reinitialize the device. This saves from a few minutes up to several hours per development iteration and is one central reason for the widespread success of SDRs in research and development. Beyond this, we aimed to achieve this isolation not only functionally, but also in terms of performance: Decoupling

of software components prevents blocking of critical tasks like sample recording by concurrent non-critical processes, e.g., real-time data analysis. Asynchronous procedure calls allow commands to be issued without waiting for previous operations to complete.

We designed our architecture from the ground up for real-time capability: This allows the application to react to measurement data with low latency and to precisely control the timing of all hardware interactions, which is essential in highly parallel MIMO and multi-node setups. One potential application is the implementation of an automatic gain control in software.

As a result, the following design objectives guided the development of our architecture:

- 1) Optimize the data rate both over the host-to-device link and to the SSD storage to enable continuous recording
- 2) Isolation between application logic and abstracted basic functionalities
- 3) Shifting the complexity from the FPGA to the host, maximizing the use of Python and minimizing the use of C++
- 4) Low latency in both control and data planes
- 5) Deterministic timing of hardware actions independent of software timing jitter
- 6) Flexibility in hardware, number of nodes/channels, and applications.

These goals allow to create one versatile platform that can be adapted via high-level user scripts to a multitude of research and development wideband RF use cases, particularly in the context of 6G design, rapid prototyping, and implementation. This encompasses applications such as multi-node MIMO RF propagation measurements, radar target characterization, ISAC, and real-time algorithm testing, which also incorporates AI-driven features.

IV. SYSTEM ARCHITECTURE OVERVIEW

Fig. 1 outlines the basic structure of our proposed architecture, which consists of the PC-based host and the converter device, with an FPGA at its core, connected via a high-speed Ethernet link. Due to its generic design, the hardware components as well as the high-speed data interface can be exchanged to match individual requirements without incurring changes to the architecture.²

Ethernet provides the most promising solution to connect host and converter device, as suitable network interface

²Minor changes to the implementation might be required.

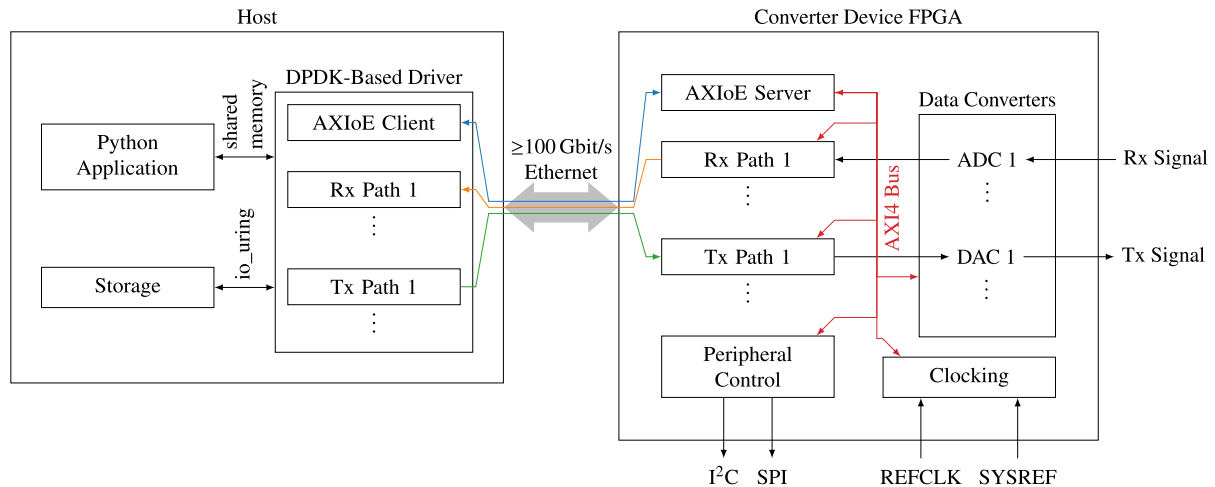


FIGURE 1. Overview of our system architecture, consisting of the converter device, which is operated by a driver software on the host computer via one or multiple high-speed Ethernet links. Its generic design allows for adaption and exchange of hardware components and high-speed data interface. This allows to achieve the individually desired performance parameters without requiring changes to the architecture. Encapsulation of base features in an abstraction layer on the host allows for interfacing with the system via the high-level scripting language Python.

controllers (NICs) are widely available for COTS server hardware. It supports high data rates and allows for a wide range of connection types: From short direct connections via copper cable to kilometer-long fiber optic cables and complex switched networks. Using Ethernet as an interchangeable, generic, and flexible interface, our architecture overcomes the limitations of monolithic SotA setups and proprietary interfaces.

The Ethernet interface is used for both the communication to control the device (control plane) and the streaming of samples (data plane). For the latter, there is one data path per channel, whereby a distinction is made between Rx (ADC channels, sample flow from device to host) and Tx (DAC channels, reverse sample flow). The planes' individual communication protocols were each designed with the goals of efficiency and shifting complexity from the FPGA to the host to simplify implementation. Unlike UHD, our architecture natively supports burst Tx streaming, eliminating any control overhead. Moreover, it is optimized for ultra-fast and overhead-free error recovery in both Rx and Tx streaming operation.

To link the control plane modules within the FPGA, we employ the de facto standard AMBA advanced eXtensible interface 4 (AXI4). All modules provide their configuration via registers, which can be read and written through this memory-mapped interface. As AXI4 is widely used, a variety of ready-made infrastructure cores and functional IP cores with compliant control ports are available. Bridges from the AXI4 network to other protocols, e.g., serial peripheral interface (SPI) or inter-integrated circuit (I²C), integrate peripheral components like analog RF frontends. The AXI over Ethernet (AXIoE) protocol [31], [32] tunnels register accesses through Ethernet and enables direct access to all components via straightforward Python³ programming on the

³Without loss of generality, we utilize Python to implement the application logic due to its widespread use, readability, and extensive library support.

host. On the FPGA, a timed command infrastructure enables the precise timing of actions independent of software timing jitter. Combined, this facilitates rapid prototyping with analog components, e.g., development of advanced 6G features like full-duplex RF transceivers.

For all data streams within the FPGA, e.g., ADC and DAC samples, the unidirectional variant AXI4-Stream is used.

The link to the host is the bottleneck of the architecture as the combined data rate of the utilized converter channels may exceed its capacity. In that case, not all streams can be continuously operated with maximum RF bandwidth and 100% duty cycle.

To realize maximum flexibility, each channel features an independent data path on the FPGA, which allows time-based sampling control per channel, supporting both burst and continuous sampling modes. Samples from individual channels are independently packed into Ethernet packets, which are then transmitted on the Ethernet interface in a round-robin arbitration scheme. All Ethernet packet headers include virtual local area network (VLAN) tags to distinguish between channels and to enable building distributed setups with multiple converter and host nodes connected to a switched network. To ensure only intact packets are processed, the FPGA monitors all inbound packets' cyclic redundancy check (CRC) and Reed-Solomon forward error correction (RS-FEC) status. Faulty packets will be discarded.

The host-side software is separated into two components:

- 1) A driver which encloses the base functionality in low-level code and the high-level application-specific scripting. The driver, written in C++, handles high-rate and latency-critical communication. This is supported by the Data Plane Development Kit (DPDK) framework, which is a widely used solution for performance-optimized networking in the data center industry [33]. Like the FPGA image, the driver is developed to be generic and reusable.

- 2) A Python program governs the application-specific operations. This program interacts with the driver via a shared memory interface, implementing mechanisms for both asynchronous remote procedure call (RPC) and sample transfer. This separation allows the application logic to be modified without recompiling C++ code or even the FPGA design, also obviating the need to flash and reinitialize the SDR device.

We realize continuous streaming and recording of a 10 GB/s sample stream via a single 100 Gbit/s Ethernet link. Thus, we overcome the data rate and duty cycle limitations of UHD [26] and the MIMO testbed presented in [27] and [28].

V. SYNCHRONIZATION AND TIMING

In line with our design goals, we realize synchronous sampling and deterministically timed command execution in distributed setups. To allow for synchronous sampling, a reference clock (REFCLK) is distributed to each converter device. Absolute synchronizability in time is gained by additionally introducing a system synchronization reference signal (SYSREF) and using a simplified synchronization scheme similar to JESD204 [34]. It is generated and distributed with a defined timing relation to the REFCLK. All clocks required internally are derived from the REFCLK on each individual converter device. Ensuring integer clock relations yields deterministic synchronizability as well as synchronous sampling for local and distributed converter device setups. We plan to publish a more detailed description of the clock generation, distribution, and synchronization scheme for single and multi-device setups in a dedicated article.

In order to arrange commands and data in a chronological sequence, two time bases are employed. Firstly, we define a sampling-oriented time base: The data converters aggregate samples into beats resulting in a configuration-dependent beatclock. From this, we derive the **beatcounter** which provides a direct relation between sample beats and time. It is used for all sampling-related purposes. Secondly, as low sample rates unavoidably incur insufficient resolution of the beatcounter, we derive an invariant **real time clock (RTC)** from the FPGA's REFCLK. It is used for all other timing needs, e.g., interaction with peripheral components.

A. FPGA SYNC AND TIME

The FPGA implements individual counters for both time bases. Since the RTC is clocked by the REFCLK, it can be directly synchronized using the external SYSREF signal, which has a known timing relation to the REFCLK. The synchronization mechanism is armed via AXIOE and will subsequently trigger on the SYSREF's next rising edge. Based on the previously synchronized RTC, the beatcounter is initialized via a timed command. The synchronization algorithm predicts both clocks' common edges and thereby ensures deterministic operation for all possible (integer) REFCLK to beatclock relations no matter which possesses the highest frequency.

The implementation of these fast counters is a major challenge. In fact, their adder's possibly long carry chain might impair the design's timing closure and therefore requires designing ultra-fast and lightweight counter modules. Instead of implementing a monolithic adder structure, the logic is split into two synchronized counters as proposed by [35]. While the first counter increments the lower part of the output word in every cycle, the second one precalculates the next upper part and propagates it to the output as soon as the fast counter overflows. This way, the long carry chain can be split into two parts. While the first counter's short chain still has to fulfill the tight requirements, the second counter's long chain's timing can be eased by applying multicycle path constraints.⁴

B. HOST-SIDE PACKET PACING

Most FPGAs have very limited internal buffer capabilities, e.g., in the order of a few megabits. Without incurring further hardware dependencies like utilization of external memory, the SDR may only buffer a few microseconds of Tx sample data. To prevent overflows, the host must not send its samples to the device too early. Similarly, on the control plane, the host must adhere to queue size limitations of the FPGA. Therefore, a flow control mechanism is required.

In the *condensed hierarchical datagram for RFNoC (CHDR)* protocol used by the universal software radio peripherals (USRPs), the device gives clearances to the host to transmit data up to a given stream position (transmit window) [36]. As it relies on the host to react in time, the buffers in the FPGA must be significantly larger than the worst-case software latency. Therefore, this solution is not viable here. In the newer USRP X440, Ethernet pause frames are used for flow control [29]. These are packets that request the remote device to suspend data transmission for a specified period of time. However, the disadvantage of this method is that it is not channel specific and, in particular, slows down time-critical control communication (head-of-line blocking).

Instead, we take a different approach: For both timed commands and Tx samples, it is known when the device will consume them from its buffers. This allows the host to send the packets accurately timed so that they arrive at the device with a fixed lead time. To achieve this precisely, we employ the *send on timestamp* feature of the NIC hardware, which allows the software to schedule packets ahead of time [37]. The NIC puts the packets on the wire at the predetermined transmission time. Thus, the arrival time at the FPGA is decoupled from the relatively high software timing jitter, rendering significantly smaller buffers possible.

To achieve this, the host has to know the relationship between its local and the device's remote time bases. In fixed intervals, it requests the device to read out the current time stamp counter value. The response then tells the time at which the request packet arrived at the device. This is used as input to a clock model to adjust the transmission timestamps of future packets.

⁴The generic design allows for compile-time reconfiguration of the counter's split.




FIGURE 2. AXIoE packet processing FPGA design: AXI4-Stream-based command processing and AXI4 bus manager interface. Separate AXIoE preprocessing enables fast error recovery.

VI. DEVICE CONTROL

Full-featured remote control of the FPGA's internal modules and peripheral components by the host is required to enable an architecture independent from both hardware and application. In the FPGA design, we realize all configuration and control functionality via memory-mapped AXI4. We utilize AXIoE as generic and reliable access protocol for tunneling the AXI4 accesses from the host to the FPGA over Ethernet. This allows for easy integration of commercial and custom modules and peripheral components into the architecture.

A. AXI OVER ETHERNET

For the remote control of the different blocks on the FPGA, we employ the protocol AXIoE, specified in [32], which allows to tunnel memory-mapped AXI4 accesses via Ethernet [31].

The protocol realizes an ordered and reliable command stream, but also allows for independent commands to be transmitted unreliably and out of sequence. It is specifically designed for the constraints of the communication between a host and an FPGA: Its asymmetric design requires the AXIoE server on the FPGA side to only realize a simple request-response mechanism. In contrast, the error detection and recovery process is to be realized by the client running on the host. This allows for a lightweight FPGA implementation, shifting complexity to the host software.

The host can send request packets, which may contain one or multiple transactions. Each transaction consists of one atomic command, either a read from or write to a specified memory-mapped address range. Upon packet loss, the host performs a resynchronization: It either repeats the lost request packets or asks the FPGA to repeat lost responses.

On the FPGA side, we choose a split implementation of the server functionality as visualized in Fig. 2. The preprocessing module checks incoming packets' AXIoE headers. Whereas protocol compliant packets are forwarded into the inbound command first in, first out (FIFO) buffer, faulty packets are discarded and error tickets – containing the information required for generating the error response – are inserted into the FIFO instead.⁵ The AXIoE state machine processes the incoming requests from the command FIFO, executes transactions on the AXI4 interface, and generates

⁵This prevents faulty packets from entering the FIFO and therefore fully eliminates the time required to read them out of the command FIFO in error case.




FIGURE 3. Timed command network: FPGA design topology. AXI4 bus to AXI4-Stream conversion, AXI4-Stream routing, and timestamp distribution. Standard interfaces and lightweight design allow for simple and fast functional extension. Either beatcounter or RTC may be used as timestamp.

response packets. For proper operation it requires utilizing two additional buffers: The response FIFO stores individual AXI4 transaction response data until its status header can be generated. The replay memory is addressable and stores full response packets for potential replay requests.

B. FPGA COMMAND TIMING

Fig. 3 depicts the timed command network topology implemented on the FPGA. The respective time bases' counter output is distributed to all related blocks. The timed command dispatcher module is accessible via the FPGA's AXI4 network and converts memory-mapped accesses to AXI4-Stream packets containing the timed commands. It checks the command time margin, generates the routing information, and inserts it alongside the actual packet data into the timed command AXI4-Stream network. Each timed command actor features its own FIFO buffer. This allows queuing timed commands independently for each actor and avoids head-of-line blocking. An actor module loads a command and executes it as soon as the execution target time is reached. Due to the timing mechanism's time base independent design, each actor module may utilize either the beatcounter or the RTC.

To further relax critical timing paths, the execution target time check is uncoupled from the actor module's actual command logic by setting a start bit and delaying the execution by one cycle instead of performing the check and the command in the same clock cycle.

VII. HIGH-RATE ADC STREAMING

When host and SDR are connected via one 100 Gbit/s Ethernet link, theoretically up to 99.623 Gbit/s net data rate can be achieved when using jumbo frames of 9000 bytes payload. We were able to demonstrate a data rate of 95.885 Gbit/s for the maximum standard-compliant packet size of 1500 bytes. [38]

To allow for any meaningful error recovery, the converter device would need to buffer the outbound data stream for a significant amount of time. The internal memory resources typically found on FPGAs do not suffice to realize this. Using two or more DDR4 memory banks would provide the bandwidth required for prolonged buffering, but would introduce unwanted hardware dependencies on the system architecture.

The same reason also opposes the use of the RDMA over converged Ethernet (RoCE) protocol, which is a common




FIGURE 4. Rx FPGA path: AXI4-Stream-based, beatcounter-timed triggering, MTU compliant packing, efficient buffering, parallel header generation, and data unit relation based merging. The design enables burst and continuous sample streaming.

solution for implementing reliable high-rate data streams: On the host side, it is directly supported by many NICs, allowing zero-copy interaction with continuous memory regions without CPU intervention. However, it demands significant resources on the FPGA side [39]. Beyond this, RoCE's real-time capabilities are limited as the software on the host does not have control over the timing of individual packets.

Instead, we designed an ultralight Rx protocol without retransmission capabilities. Similarly to the control plane protocol, it is constructed asymmetrically to attain minimal implementation complexity on the FPGA. Thus, we centered our design around the sample beats generated by the ADCs: They form atomic units, which will not be fragmented on the FPGA and their generation frequency is the basis for their timestamping (see beatcounter in Sec. V).

In addition to the actual sample data, each Rx packet contains the following meta information:

- The beatcounter based timestamp of the packet's first sample, allowing the host to infer the timestamp of all samples within the packet
- Packet sequence number, independent per channel
- ADC status bits providing overrange, overvoltage, and threshold information, allowing the host to detect analog faults and implement an automatic gain control.

A. FPGA-DESIGN REALIZATION

Fig. 4 illustrates the FPGA side ADC path design. The trigger module executes beatcounter-timed commands, allowing for both burst or continuous sampling. This module forwards an ADC channel's sample beats as one AXI4-Stream packet per trigger event. The first beat's beatcounter value is attached to the packet. The subsequent module, i.e., the data packer, fragments the original packet according to the high-speed interface's maximum transmission unit (MTU) size. To maintain the stream's embedded timing information, intermediate beatcounter values are calculated internally and attached to each fragment. The data packer directly feeds the sample data into the sample FIFO for subsequent merging with the associated packet headers: In parallel, the header generator produces one packet header per fragment from the accumulated metadata. The header merge module combines the sample fragments with their respective headers and forwards them to the Ethernet interface.

In burst mode, the resulting data rates can exceed the capacity of the Ethernet link in the short term. Therefore, buffering is necessary to cover the occurring backpressure.

Due to the limited resources of the FPGA, an efficient design is essential. This is achieved by splitting the data path into multiple, parallel sub-paths, e.g., for sample data and packet headers. Each individual sub-path handles either data units of AXI4-Stream packets or single AXI4-Stream beats. This relation empowers implicit synchronization by AXI4-Stream flow control and path merging by a lightweight mechanism based on the data units.

The data relation between the sub-paths enables a pipelineable, modular, and expandable path design and individual buffering or processing per path without explicitly implementing any synchronization mechanism. To preserve this inter-path data relation, all modules within the paths must be designed to be fully flow control compliant. A recovery mechanism handles backpressure conditions impacting the trigger module⁶ and prevents the sub-paths from losing synchronization.

Special considerations are necessary to allow continuous operation of the ADC paths. Backpressure occurring in paths with critical load will eventually cause FIFO overflows. To prevent systematic backpressure, all modules used in paths with critical load must be implemented with an initiation interval⁷ of 1. Stochastic backpressure must not occur either, since it can never be caught up with. This requires an adequate path design with respect to data width conversions and clock domain crossings.

B. HOST-SIDE IMPLEMENTATION

A central use-case of our SDR architecture is recording the received samples for offline processing, requiring particularly fast access to large amounts of storage. In order to achieve the maximum possible write rates, we use multiple SSDs and combine their individual write speeds using a software RAID0 and the XFS filesystem. The host software interacts with the storage using *io_uring*, an asynchronous, high-performance interface to the Linux kernel [40].

For maximum performance, we use the *O_DIRECT* access mode of the Linux kernel. It imposes a fixed block size, in our case 512 bytes [41]. This is in conflict with the requirements of the communication protocol, which handles sample beats as atomic units. Therefore, we decided not to strive for a zero-copy implementation in the driver, but rather to copy the samples from the individual packets into a large ring buffer. This provides maximum flexibility. In fact, samples can be selected arbitrarily in this step without being bound to beat limits. If samples are missing in the output stream, e.g., due to a lost packet, they are zero-padded to ensure that the subsequent samples are found at the expected position in the stream. The ring buffer allows for simultaneous online processing of samples, e.g., to provide a live view of the data or to implement an automatic gain control, without disrupting the real-time recording in any way.

⁶To maintain the timing relation to the continuous stream of ADC samples, it must not support backpressure.

⁷The delay between processing of successive input data in units of clock cycles.

It would be a desirable feature to be able to configure the network card not to drop packets with erroneous checksums: Since there is no retransmission option in our application, we would rather accept bit errors in the sample stream than lose entire packets. Unfortunately, DPDK currently does not provide a way to configure the NIC accordingly.

VIII. ANALOG SIGNAL GENERATION: DAC-PATH

In common SDR solutions, only streaming is supported in Tx mode, making the Ethernet link the bottleneck of the system. To overcome this, we have implemented approaches to efficiently support both static and dynamic Tx sequences:

The first approach operates in a **loading-looping** manner – like an arbitrary waveform generator – realizing dynamic exchangeability of sequences at runtime via AXIoE. Since a sequence only needs to be transmitted to the converter device once, this mode of operation does not require high performance host hardware to setup and control the waveform generation. It enables the synchronized playback of periodic sequences on multiple channels, vastly exceeding the Ethernet interface data rate.

The second approach implements the most flexible solution realizable for an SDR platform: **Real-time Tx sample streaming**. The ability to transmit arbitrary RF signals enables the platform to implement a full-featured Tx. In this setup, the host server transmits the DAC samples and their associated timing information to the converter device, which buffers the data until the specified playback time. To ensure proper sample handling, both ends implement mechanisms to compensate for network jitter. If there are no samples to be played, zero-samples are automatically passed to the respective DACs by the FPGA (auto-zero on idle). In contrast to UHD, this eliminates the need for explicit start and stop commands when transmitting burst signals.

A particular challenge with streaming Tx is handling errors: In the CHDR protocol used in the USRPs, this is done by monitoring the sequence numbers of incoming packets [36]. Errors are reported to the host and Tx operation is only resumed after explicit acknowledgement. The start and end of burst must be explicitly marked to allow discrimination between an intentional interruption and packet loss. Correctly handling the loss of these delimiter packets is particularly complex [42]. In our application, due to limited resources and high data rates, the buffers on the FPGA cannot be realized large enough to have sufficient slack to allow for the retransmission of lost packets.

We have therefore opted for a different approach, significantly simplifying the protocol design: The header of each sample packet indicates its desired playback time expressed as beatcounter value. The FPGA offers multiple statistics counters, capturing the number of successfully transmitted packets as well as the number of packets discarded due to late arrival. The host regularly reads these counters via AXIoE, whereby the FPGA attaches the timestamp when exactly the counter values were last changed. Since the software knows how many sample packets should have been consumed by the




FIGURE 5. Tx loading-looping FPGA path: Conversion of AXI4 writes to AXI4-Stream packets. Timed command-based loading-looping of sequences. Lightweight approach allows for playback of periodic signals, eliminating the bottleneck of the link to the host.




FIGURE 6. Tx streaming FPGA path: Beatcounter-based payback of inbound sample packets based on AXI4-Stream. Includes protocol-based error detection and handling by buffer reset. Fire-and-forget protocol design reduces complexity and resource consumption of the FPGA realization.

DAC at any given time, it can determine whether packets have been lost.

This solution also keeps the gap in the Tx signal as small as possible: When a packet is lost, only the samples in the affected packet are missing and the output auto-zero on idle feature implicitly handles the error on the FPGA side. Consecutive packets are played as intended by their beatcounter target. In addition, there is no delay caused by waiting for explicit acknowledgments.

A. FPGA-DESIGN REALIZATION

1) LOADING-LOOPING-APPROACH

Fig. 5 visualizes the architecture of the loading-looping DAC data path: Using AXIoE, the host writes the desired sequence into the loader module, which passes the sequence as an AXI4-Stream packet to the load FIFO. After loading is finished, the host may control the playback of the sequence using timed commands. The playback module outputs the sequence as sample stream to the DAC, but also feeds it back to its input via the loop FIFO for repetitive playback.

2) REAL-TIME TX STREAMING

The Tx streaming path design is shown in Fig. 6. The inbound module performs a sequence check and inserts samples and timestamps into individually buffered sub-paths. The outbound module manages time-controlled playback and outputs zero-sample words when no sample data are available. Inbound and outbound module provide timestamped statistics counters, which can be read out by the host via AXIoE.

The architecture implements a global packet status check, which ensures packet bit integrity. On protocol level, four types of errors may occur and are handled appropriately:

- 1) Packet loss is implicitly handled by the auto-zero on idle feature.
- 2) An out-of-order packet is handled by the inbound module by discarding the late packet.

- 3) A sample FIFO underrun occurs when packets arrive too late. The outbound module recognizes a missed target time and triggers a reset of both FIFOs. The inbound module monitors the reset condition.
- 4) A sample FIFO overflow occurs when too many packets arrive too early. The inbound module detects and safely resolves backpressure conditions at both FIFO inputs. This ensures sub-path synchronization.

After any error condition occurred and has been resolved, the inbound module discards inbound data until it resyncs its input interface to the next Ethernet packet. This design guarantees the shortest possible interruption of the output sample stream.

B. HOST-SIDE IMPLEMENTATION

The loading-looping approach places no special demands on the host – the sequence to be played back is loaded into the device by a driver call from the user application (e.g., a Python script) using the AXIoE control path.

Tx streaming is more complex: Here, the host driver ensures that the relatively small buffers in the FPGA neither overflow nor underflow. As described in Sec. V-B, we use the NIC's send on timestamp feature for this purpose. To detect errors, the host reads the previously described status counters periodically via AXIoE. Based on the timestamp attached to the result, the host driver calculates how many sample packets are expected to have been consumed by the DAC. From this, the host determines how many errors of which type have occurred during the readout interval.

IX. IMPLEMENTATION EXAMPLE

The proposed SDR architecture, described in the previous sections, is not tied to any particular hardware. Starting from this section, we introduce a specific implementation (based on the Xilinx RFSoc). This allows us to discuss relevant implementation details and validate our architecture with real measurements.

A. XILINX ZYNQ ULRASCALE RFSOC XCZU48DR

Besides the actual FPGA, the Xilinx RFSoc XCZU48DR monolithically integrates specialized blocks, e.g., 100 Gbit/s Ethernet, ADC, and DAC:

The RF data converter (RFdc) hard-IP offers direct RF-sampling ADCs and DACs. Table 2 and Table 3 list their parameters. Both converter types integrate digital signal processing features like digital down conversion (DDC) and digital up conversion (DUC). The ADCs and DACs are organized in tiles. Their power up sequence is neither synchronized nor timed, which results in an undeterministic timing relationship between tiles of a single as well as multiple devices. The converter multi-tile synchronization (MTS) procedure ensures a consistent and deterministic timing across all ADC and DAC tiles. This requires specific external clocks and sync signals, which have to be reconfigured multiple times. All converter features are accessible through a single IP core, which can be customized for the design and provides data as well as control ports.

TABLE 2. Xilinx XCZU48DR: ADC characteristics. [43], [44].

Number of channels	8
Interleaved sub-ADCs per channel	8
Resolution	14 bit
Sample rate	1–5 GSa/s
Analog bandwidth (-3 dB)	6 GHz

TABLE 3. Xilinx XCZU48DR: DAC characteristics. [43].

Number of channels	8
Resolution	14 bit
Sample rate	0.5–9.85 GSa/s
Analog bandwidth (-3 dB)	6 GHz

Xilinx provides a 100 Gbit/s Ethernet interface by combining a CMAC hard IP with a GTY serial transceiver quad. The RFSoc features two⁸ of these interfaces.

Beside the FPGA as user-programmable logic (PL), the RFSoc integrates a central processing unit (CPU) core as processing system (PS), which in our case runs a Linux system and initializes the platform by configuring peripheral clock generation components⁹ and loading the bitstream onto the FPGA. AXI4 interfaces between PL and PS allow communication between both sides. By connecting AXIoE to the PS-PL interface and utilizing a driver in the PS, the host is empowered to remotely reconfigure all peripheral components that are accessible via the PS.

B. HOST

On the host side, enterprise COTS hardware is used. Particularly noteworthy is the SSD array, which allows continuous recording of one channel's data at its full sample rate of 5 GSa/s: It consists of 4x *Samsung SSD PM9A3* with 4 GB/s write rate each, so that a RAID0 reaches a sustained write rate of 16 GB/s. The 100 Gbit/s Ethernet NIC used to communicate with the FPGA is an *Nvidia Mellanox MCX623106AN-CDAT*. It offers the required feature send-on-timestamp (also called packet pacing), allowing for the exact timing of packet transmission [37]. Both the SSD array and the NIC are connected via PCI Express 4.0 directly to the CPU for maximum transfer rate.

C. DEVICE CLOCKING

As REFCLK, we use a 100 MHz square wave instead of the common 10 MHz sine wave to achieve a better phase noise performance. A 1 PPS reference signal is fed to the converter device. Based on these, each device derives its clocks and SYSREF. This enables synchronous sampling and absolute synchronization in multi-device setups.

D. RFDC-OVER-ETHERNET

In addition to the sample transfer interface, the RFdc IP core in the FPGA design provides an AXI4 control interface. It is

⁸Currently, we are limited to a single CMAC due to evaluation board layout constraints.

⁹The clock generation network's communication interface is hardwired to the PS package pins on the used evaluation board.

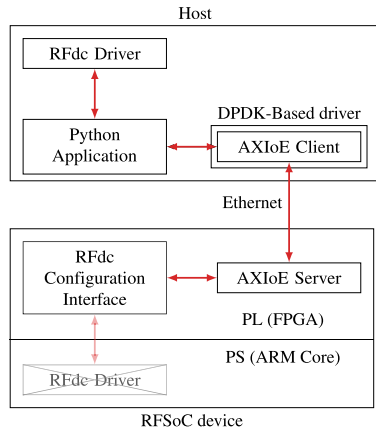


FIGURE 7. The configuration interface of the RF data converters is typically operated out of the PS, an integrated ARM core. Instead, we connect it to the host via our AXI over Ethernet infrastructure. There, the driver is linked to the Python application, allowing it to configure the hardware with maximum flexibility.

used, e.g., for initialization, calibration, and synchronization. To access this interface, Xilinx provides a driver which is intended to run on the PS. This contradicts the goal of our SDR architecture to have the host control the device as comprehensively as possible.

Therefore, we decided to give the host direct control over the RFdc by attaching its configuration interface to the AXIoE AXI4 network as depicted in Fig. 7. However, the control registers of the RFdc are not documented publicly, so the original driver must be utilized to interact with them. Fortunately, it is published in C source code, so it is possible to compile it for the host architecture. We adapted the functions that access the memory-mapped control register in the PS to trigger AXIoE transactions instead. This is transparent to the driver, as the underlying AXIoE layer handles any problems that may occur, such as packet loss.

The only remaining hurdle is that the driver is designed for synchronous register accesses, as from the PS this is a fast local access, but the Python application follows the paradigm of asynchronous programming. The solution here is provided by the library *ucontext*, which makes it possible to create a separate execution context (especially stack) for the driver. Its execution is suspended for every AXIoE access and continued after the asynchronous arrival of the result.

The concept of tunneling accesses and commands through the reliable and ordered AXIoE interface is not limited to the RFdc and its driver. Instead, it is generically applicable for most commercial IPs, custom modules, and even hardware peripherals, which are controlled by software drivers via register accesses. Beyond that, AXIoE can easily be utilized as reliable tunnel for non-AXI4 interfaces via wrapper modules and on top benefits from the timed command functionality already included in the design.

X. MEASUREMENT RESULTS

After implementing the proposed system architecture, we carried out a multitude of tests to verify its performance.

First, we examined the stability of the synchronization and command timing: To do this, we power cycled the device numerous times and performed all synchronization steps, which include the MTS, repeatedly. In each cycle, we measured the system latency in an analog loopback. Using timed commands, we set up the DAC path to produce a periodic test signal. The generated analog Tx signal was looped back into the Rx channels, the ADC samples were recorded, and the delay between transmitted and received signal was determined. Within each DAC/ADC combination, the measured latency was constant across multiple reboots. This not only confirms that the synchronization, both within the FPGA design and between the data converters, works deterministically, but also validates proper command timing.

Next, we evaluated our control plane implementation in high-rate switching scenarios similar to, e.g., MIMO channel sounding [30, Fig. 2]. The software not only has to produce packets fast enough, but also pace them precisely so as not to overload the command queues in the converter device. We verified that commands were reliably executed even when sending AXIoE packets with the high pace of 5 μ s. This proves that even fast switching tasks do not have to be implemented explicitly in the FPGA but can be handled by the host in software.

To verify the performance of the send-on-timestamp mechanism, we measured the arrival time jitter. We did this by regularly inserting AXIoE requests into the data stream, which command the FPGA to read out the current value of the time stamp counter. As shown in Fig. 8, the range in which the real arrival time differs from the planned arrival time (jitter) spans 3.34 μ s. This proves the precision of both the hardware send-on-timestamp feature and clock modeling, confirming our design decision to keep Tx streaming and control plane buffers on the device as small as possible.

Next, we evaluated the sustained performance of the Rx path: Over several hours of continuous sample streaming at the full rate of 5 GSa/s, no packet loss occurred. Also, continuously storing the incoming data stream of 10 GB/s to the SSD array has been proven to work reliably, only limited by the storage capacity.

We successfully verified the timing and synchronization of both Tx path variants. Tx streaming was demonstrated to work reliably at 5 GSa/s, i.e., a net data rate of 10 GB/s. Hereby, 112 μ s of sample buffer were required on the device, which is more than we expected. This is not a matter of the architecture itself, but most likely caused by a problem with the utilized host hardware, which is currently under investigation [45]. Nevertheless, at a sample rate of 2.5 GSa/s, this effect disappeared and only 14 μ s of samples had to be buffered on the device for reliable Tx streaming.

Finally, to verify the real-time capabilities of our platform, we measured the round-trip latency in the following test scenario: For both Rx and Tx, we set up a single continuous stream at a rate of 2.5 GSa/s. To ensure that each Tx sample packet arrives in time, the software must enqueue it into the NIC transmit queue with a fixed lead time. We identified




FIGURE 8. Histogram of packet arrival time deviation, positive values indicate late packet arrival. Using hardware send-on-timestamp, the time at which packets arrive at the device can be planned with a jitter of only 3.34 μs . This accuracy allows implementing Tx streaming with very small buffers in the FPGA, conserving device resources.

54 μs as the minimum viable value for reliable operation. Connecting two converters in an analog loop and analyzing the respective timestamps, we determined a latency of 0.12 μs from the DAC's digital input to the ADC's digital output.¹⁰ In the Rx path, a sample reaches the software at most 17.3 μs after it was output by the ADC. The individual values add up to the system's worst-case **round trip latency of 71.4 μs** .

This makes our system ideally suited to meet the sub-millisecond latency requirements of state-of-the-art 6G applications [8], [9], [10], [12], [13]. Therefore, it enables hardware-in-the-loop evaluation of each element of the communication system in a rapid prototyping environment: From communication protocol components like scheduling, channel estimation, and beam steering to novel applications such as ISAC.

XI. APPLICATION EXAMPLE: 6G-ISAC-RADAR

Integrated sensing and communication (ISAC) is a proposed feature of the upcoming 6G standard [4], [5], [7], [14], [15]. Both infrastructure and sidelink-based communication may integrate radar-like sensing functionality [6], [46]. This promises great benefits for use cases like assistance systems, health monitoring, mobility, public safety, and many more [6], [14], [16].

Owing to its unique features and performance (cf. Table 4), our SDR architecture is well suited for 6G rapid prototyping, e.g., for ISAC [7]. The standardization of ISAC is at an early stage, i.e., comprising a work item [1] and case studies [2], [3]. Therefore, we opted for a straightforward application example: A quasi-monostatic, single-input single-output Doppler radar demonstrator, inspired by the ISAC features envisioned for 6G [6]. The example realizes arbitrary waveform generation and continuous sample acquisition

and storage using our SDR architecture. The emitted signal uses the same orthogonal frequency-division multiplexing (OFDM) modulation scheme that is used by 5G and potentially 6G, representing illumination by a base station or user equipment [4]. The architecture generally supports common mobile communication features, e.g., MIMO and beamforming, which can be implemented if appropriate multi-channel RF frontends are available. The real-time streaming capabilities with low latency and high bandwidth even facilitate the future implementation of advanced ISAC demonstrators that seamlessly integrate into a 6G radio access network, serving as simultaneously sensing and communicating 6G network nodes.

Regarding our simplified ISAC application example, evaluating the Doppler effect over time and range will not only reveal the speed of a radar target as a whole but also the movements of its inner parts, the so-called micro-Doppler. The entirety of a target's inner movements, e.g., a pedestrian's limbs or the individual rotor blades of a UAV, can be detected as characteristic patterns: Its micro-Doppler signature. To resolve the target's inner structure in the range domain, a radar system with high instantaneous bandwidths is necessary. With a bandwidth of 2 GHz, our setup realizes a path distance resolution of 15 cm, which corresponds to a range resolution of 7.5 cm in the monostatic case. In [47], we discuss the promising chances of joint evaluation of static and dynamic target reflectivity for detection, localization, and classification of targets in upcoming ISAC solutions and therefore, future mobility applications.

Our measurement setup is shown in Fig. 9: As only a single converter device is used, no multi-device synchronization is required. It continuously transmits a complex-valued baseband OFDM sequence with a period length of 2500 samples and a bandwidth of 2 GHz. The baseband signal is mixed up into an RF signal with a center frequency of 5 GHz. One antenna radiates the amplified signal while a second one receives the target's reflection. This Rx signal is amplified and mixed down into the baseband before being sampled by the converter device and sent to the host. The resulting stream of 10 GB/s is stored on the SSD array in real-time, while computing delay-Doppler plots in parallel.

As shown in Fig. 10a, we adapted a typical ISAC mobility scenario as measurement example: A pedestrian is walking towards a concentrated radar sensor node, which is realized by our SDR system and utilizes a quasi-monostatic antenna setup. Fig. 10b contains a snapshot of the continuously generated delay-Doppler plots. The target's micro-Doppler signature can easily be identified as three characteristic peaks: The pedestrian's torso appears with its walking speed as the peak in the center. Moving relatively to the torso, the swinging arms appear centered around it as distinct peaks with different distances and speeds. The signature's asymmetric shape, which is shifted towards higher absolute speeds, is caused by the stepping foot moving toward the antennas in the plotted

¹⁰This latency is primarily caused by the RFdc IP core.

TABLE 4. Comparison of our solution with the state-of-the-art USRP X440. [26], [29].

	UHD Ettus X440	This work XCZU48DR
Ecosystem	Mature	Recent ¹¹
Rx (cf. Sec. VII)		
Continuous, high-rate recording	Must be implemented by user	Native support; demonstrated 10 GB/s per 100 Gbit/s link
Real-time sample access	Must be implemented by user	Native support; decoupled, concurrent, zero-overhead access via shared memory
Packet loss handling	Resynchronization must be implemented by user	Automatic zero-padding retains synchronization
Sample distribution	A separate endpoint can be specified per stream	VLANs enable arbitrary distribution of samples to multiple hosts
Tx (cf. Sec. VIII)		
Arbitrary waveform generation	Flexible playback from SDR memory	Loading-looping via SDR memory
Burst streaming	Separate handshake for each burst	Native support enables zero overhead
Flow control mechanism (cf. Sec. V-B)	Pause frames, blocking all traffic on interface (head-of-line blocking)	Software-controlled per-stream packet pacing
Packet loss handling	Resynchronization requires lengthy handshake [42]	Auto-zero on idle inherently retains synchronization
Hardware Control (cf. Sec. VI)		
Synchronous parallel command execution	Limited due to shared command queues [30]	Full support via isolated command queues
Fast automatic gain control (AGC)	Necessitates FPGA implementation [30]	Demonstrated low latency enables implementation in software
Performance Parameters		
Maximum complex baseband sample rate	2 GSa/s [29]	2.5 GSa/s
Maximum instantaneous channel bandwidth	1600 MHz [29]	2000 MHz
Rx streaming rate (1x 100 Gbit/s link)	4 GB/s (1 channel) – 6.4 GB/s (4 channels) [29]	10 GB/s
Tx streaming rate (1x 100 Gbit/s link)	7.2 GB/s [29]	10 GB/s

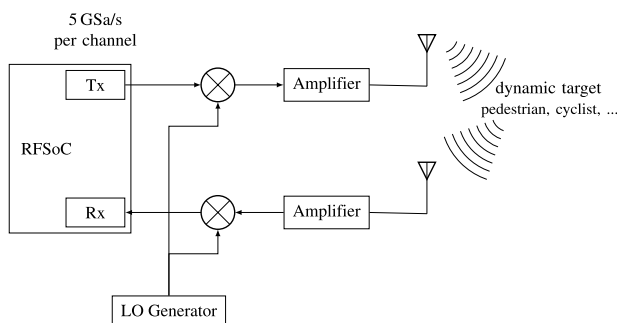


FIGURE 9. Micro-Doppler radar measurement setup with a center frequency of $f_c = 5$ GHz, reaching an analog bandwidth of 2 GHz, corresponding to a range resolution of 7.5 cm in the monostatic case.

time instant. Due to the continuously changing target geometry, the plot’s subsequent snapshots in time would show the arms’ peaks moving around the torso’s on elliptic curves. This leads to the target’s actual time-dependent micro-Doppler signature.

These measurements demonstrate a range resolution superior to that of state-of-the-art SDR solutions like the Ettus USRP X440 [29]. In the context of ISAC-related research, we already used the system to measure micro-Doppler signatures of different kinds of targets relevant for future mobility scenarios [47] as well as to validate a target

¹¹Interfaces to other scripting languages and GNU Radio can be easily developed.

modeling algorithm for generating training data for artificial intelligence [48].


XII. FUTURE WORK

First, as shown in Sec X, a further latency improvement of the real-time Tx streaming at 5 GSa/s is possible. To accomplish this, a limitation associated with the specific host hardware must be overcome. However, this is not a matter of the architecture itself and does not affect any application in which a static Tx sequence is used as we successfully demonstrated.

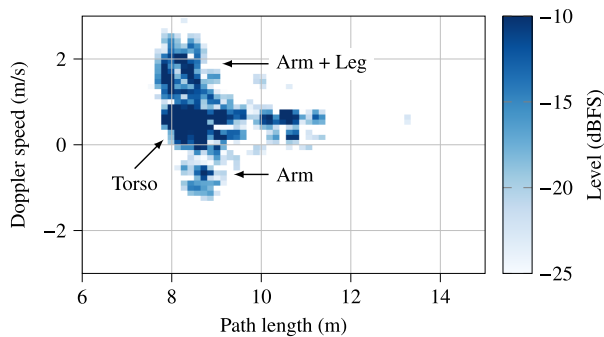
Although already addressed in the architecture, one topic that we have not yet tested is the synchronization of multiple devices: The first step is to synchronize multiple devices in a single node using a distributed clock in a wired setup. Moving on to distributed multi-node measurement arrangements using GNSS as a time source is the next challenge.

We are also working on the integration of multiple RF frontends, which will be used to access a variety of frequency bands. For practical application in channel sounding, very high dynamic ranges are necessary, which is why we are working on an automatic gain control.

With applications demanding bandwidths exceeding the capability of a single converter channel, a transceiver has to extend its instantaneous bandwidth beyond the Nyquist limit of an individual channel. To achieve this, channel bonding can be used. In a test setup with reduced complexity based on the RFSoc, we already investigated and realized multiple approaches [49]. They shall now be implemented



(a) Setup: Pedestrian walking towards our quasi-monostatic radar setup.



(b) Micro-Doppler signature: In the delay-Doppler plot, torso and extremities can be resolved. The proposed SDR system achieves a resolution (size of one pixel) of 7.5 cm by 11.8 cm/s.

FIGURE 10. ISAC scenario: A pedestrian is captured by a radar sensor node implementing the proposed SDR system architecture.

and evaluated with our novel architecture. Channel bonding requires multiple channels to be streamed from one or more converter devices to the host server in parallel. This implies challenges concerning multi-channel and multi-device synchronization, the high-speed serial interface, and the host's real-time processing capabilities, which shall also be addressed in our future research.

Finally, a more precise evaluation of the RFSoc as a measurement system is also desirable: The long-term stability of the RFdc's internal calibration needs to be investigated here. Imperfections such as inter-channel imbalances should also be analyzed for possible correction via pre-distortion or post-processing.

XIII. CONCLUSION

This contribution proposes a novel generic and hardware independent SDR system architecture, which covers functionality for Rx, Tx, and remote control. It enables synchronized and distributed multi-node transceiver setups with high channel counts. Both Rx and Tx implement time-triggered burst and continuous sample-streaming. To enable periodic signal generation beyond the limits of host and data interface, the Tx additionally features runtime loading-looping signal playback. The control infrastructure realizes remote configuration, synchronization, and timed command execution. In summary,

the proposed system architecture covers the base functionality of an SDR allowing for fast and easy high-level adaption of the realized platform to a variety of applications without requiring low-level changes to the FPGA design or C++ code. Therefore, we intrinsically support rapid prototyping. Due to its modular, generic, and reconfigurable design, our architecture is ready for future extensions.


In order to validate the functionality of the design, we have realized our architecture on a Xilinx RFSoc XCZU48DR in combination with COTS server hardware. Rx streaming and recording as well as dynamic Tx streaming were successfully implemented and demonstrated for continuous operation at the converters' maximum rate of 5 GSa/s. For periodic Tx signals, the loading-looping design allows multi-channel high-rate operation exceeding the Ethernet interface data rate limit.

Table 4 compares our solution with the latest Ettus USRP which is also based on Xilinx RFSoc technology. The table shows that our solution achieves superior performance, leveraging on a better utilization of the available link data rate. To the best of our knowledge, it is the only SDR which allows to continuously send, receive, and record a signal with a sample rate of 5 GSa/s over a single 100 Gbit/s Ethernet link. At the same time, a latency of less than 80 μ s is realized, which, to this date, no other software-based system achieves, and paves the way for sub-millisecond latency in 6G development. Therefore, we bridge the gap between the high performance requirements of modern mobile communication and rapid application development. Combining flexibility and performance, our SDR system architecture is ideal for research and applications requiring scalable and distributed RF transceiver solutions with high instantaneous bandwidths. These include antenna measurements, radar target characterization, multi-node MIMO channel sounding, real-time algorithm testing, also incorporating AI-driven features, and ISAC in 6G.

REFERENCES

- [1] *Integrated Sensing and Communication*, document 3GPP SP-230750, 2023. [Online]. Available: https://www.3gpp.org/ftp/Information/WI_Sheet/SP-230750.zip
- [2] *Feasibility Study on Integrated Sensing and Communication*, document 3GPP TR 22.837 V19.4.0, 2024. [Online]. Available: https://www.3gpp.org/ftp/Specs/archive/22_series/22.837/22837-j40.zip
- [3] *Study on Channel Modelling for Integrated Sensing and Communication (ISAC) for NR*, document 3GPP RP-240799, 2024. [Online]. Available: https://www.3gpp.org/ftp/Information/WI_Sheet/RP-240799.zip
- [4] *Future Technology Trends of Terrestrial International Mobile Telecommunications Systems Towards 2030 and Beyond*, International Telecommunication Union, document ITU-R M.2516-0, 2022. [Online]. Available: https://www.itu.int/dms_pub/itu-r/opb/rep/R-REP-M.2516-2022-PDF-E.pdf
- [5] *Framework and Overall Objectives of the Future Development of IMT for 2030 and Beyond*, International Telecommunication Union, document Rec. ITU-R M.2160-0, 2023. [Online]. Available: https://www.itu.int/dms_pubrec/itu-r/rec/m/R-REC-M.2160-0-202311-I!!PDF-E.pdf
- [6] P. Rosemann, S. Partani, M. Miranda, J. Mähni, M. Karrenbauer, W. Meli, R. Hernangomez, M. Lübke, J. Kochems, S. Köpsell, A. Aziz-Koch, J. Beuster, O. Blume, N. Franchi, R. Thomä, S. Stanczak, and H. D. Schotten, "Enabling mobility-oriented JCAS in 6G networks: An architecture proposal," in *Proc. 4th IEEE Symp. Joint Commun. Sens.*, Mar. 2024, pp. 1–6.


- [7] C. de Lima, D. Belot, R. Berkvens, A. Bourdoux, A. Dardari, M. Guillaud, M. Isomursu, E.-S. Lohan, Y. Miao, A. N. Barreto, M. R. K. Aziz, J. Saloranta, T. Sanguanpuak, H. Sameddeen, G. Seco-Granados, J. Suutala, T. Svensson, M. Valkama, H. Wymeersch, and B. van Liempd, "6G white paper on localization and sensing," in *6G Research Visions*, vol. 12, Oulu, Finland: Univ. Oulu, 2020. [Online]. Available: <https://urn.fi/URN:ISBN:9789526226743>
- [8] (2023). *6G—Connecting a Cyber-Physical World*. Ericsson. [Online]. Available: <https://www.ericsson.com/en/reports-and-papers/white-papers/a-research-outlook-towards-6g>
- [9] (2022). *6G: The Next Horizon White Paper*. Huawei. [Online]. Available: <https://www.huawei.com/en/huaweitech/future-technologies/6g-white-paper>
- [10] H. Viswanathan and P. Mogensen. (2023). *Communications in the 6G Era*. Nokia Bell Labs. [Online]. Available: <https://onestore.nokia.com/asset/207766>
- [11] (2024). *Transforming the 6G Vision to Action*. Nokia Bell Labs. [Online]. Available: <https://onestore.nokia.com/asset/214027>
- [12] (2020). *6G: The Next Hyper Connected Experience for All*. Samsung. [Online]. Available: <https://research.samsung.com/next-generation-communications#6gPop>
- [13] (2022). *B5G Technology White Paper*. ZTE. [Online]. Available: <https://www.zte.com.cn/global/about/news/20221108e1.html>
- [14] *The European view on 6G Use Cases*, 6G SNS-ICE, 3GPP Stage 1 Workshop IMT2030 Use Cases, document SWS-240018, 2024. [Online]. Available: https://www.3gpp.org/ftp/workshop/2024-05-08_3GPP_Stage1_IMT2030_UC_WS/Docs/SWS-240018.zip
- [15] *NGA Vision for 6G*, NextG Alliance, 3GPP Stage 1 Workshop IMT2030 Use Cases, document SWS-240017, 2024. [Online]. Available: https://www.3gpp.org/ftp/workshop/2024-05-08_3GPP_Stage1_IMT2030_UC_WS/Docs/SWS-240017.zip
- [16] M. Latva-aho and K. Leppänen, "Key drivers and research challenges for 6G ubiquitous wireless intelligence," in *6G Res. Visions*, vol. 1, Oulu, Finland: Univ. Oulu, 2019. [Online]. Available: <https://urn.fi/URN:ISBN:9789526223544>
- [17] T. Ulversoy, "Software defined radio: Challenges and opportunities," *IEEE Commun. Surveys Tuts.*, vol. 12, no. 4, pp. 531–550, 4th Quart., 2010, doi: [10.1109/SURV.2010.032910.00019](https://doi.org/10.1109/SURV.2010.032910.00019).
- [18] R. S. Thoma, D. Hampicke, A. Richter, G. Sommerkorn, A. Schneider, U. Trautwein, and W. Wirmitzer, "Identification of time-variant directional mobile radio channels," *IEEE Trans. Instrum. Meas.*, vol. 49, no. 2, pp. 357–364, Apr. 2000, doi: [10.1109/19.843078](https://doi.org/10.1109/19.843078).
- [19] D. Stanko, M. Döbereiner, G. Sommerkorn, D. Czaniera, C. Andrich, C. Schneider, S. Semper, A. Ihlow, and M. Landmann, "Time variant directional multi-link channel sounding and estimation for V2X," in *Proc. IEEE 97th Veh. Technol. Conf. (VTC-Spring)*, Jun. 2023, pp. 1–5, doi: [10.1109/VTC2023-Spring57618.2023.10199213](https://doi.org/10.1109/VTC2023-Spring57618.2023.10199213).
- [20] V. Ramireddy, M. Grossmann, M. Landmann, and G. Del Galdo, "Sub-band versus space-delay precoding for wideband mmWave channels," *IEEE Wireless Commun. Lett.*, vol. 8, no. 1, pp. 193–196, Feb. 2019, doi: [10.1109/LWC.2018.2866250](https://doi.org/10.1109/LWC.2018.2866250).
- [21] R. S. Thoma, C. Andrich, G. D. Galdo, M. Döbereiner, M. A. Hein, M. Kaske, G. Schafer, S. Schieler, C. Schneider, A. Schwind, and P. Wendland, "Cooperative passive coherent location: A promising 5G service to support road safety," *IEEE Commun. Mag.*, vol. 57, no. 9, pp. 86–92, Sep. 2019, doi: [10.1109/MCOM.001.1800242](https://doi.org/10.1109/MCOM.001.1800242).
- [22] (Jan. 2022). *M8131A–16/32 GSa/s Digitizer Data Sheet*. Keysight. [Online]. Available: <https://www.keysight.com/de/de/assets/7018-06368/data-sheets/5992-3412.pdf>
- [23] *Teledyne ADQ7DC Manual*. Accessed: Aug. 1, 2024. Teledyne. [Online]. Available: https://www.spdevices.com/en-us/Products/_Documents/ADQ7DC/16-1796-E%20ADQ7DC%20manual.pdf
- [24] F. Michalak, W. Zabolotny, L. Podkalicki, M. Malanowski, M. Piasecki, and K. Kulpa, "Universal RFSoc-based signal recorder for radar applications," in *Proc. 23rd Int. Radar Symp. (IRS)*, Sep. 2022, pp. 136–140, doi: [10.23919/IRS54158.2022.9905032](https://doi.org/10.23919/IRS54158.2022.9905032).
- [25] *Ettus USRP X410 Specifications*. Ettus. Accessed: Aug. 1, 2024. [Online]. Available: <https://www.ni.com/docs/de-DE/bundle/ettus-usrp-x410-specs/page/specs.html>
- [26] *Ettus USRP X440 Product Page*. Ettus. Accessed: Aug. 1, 2024. [Online]. Available: <https://www.ettus.com/all-products/usrp-x440/>
- [27] B. Nuss, P. Groeschel, J. Pfau, J. Becker, M. Vossiek, and T. Zwick, "Broadband MIMO testbed for the development and research on 6G," in *Proc. Eur. Wireless; 27th Eur. Wireless Conf.*, Dresden, Germany, Sep. 2023, pp. 89–91.
- [28] M. Neu, C. Karle, B. Nuss, P. Groeschel, and J. Becker, "A scalable and cost-efficient antenna testbed using FPGA-server compound structures for prototyping 6G applications," in *Proc. 19th Int. Conf. Distrib. Comput. Smart Syst. Internet Things (DCOSS-IoT)*, Jun. 2023, pp. 171–178, doi: [10.1109/DCOSS-IoT58021.2023.00039](https://doi.org/10.1109/DCOSS-IoT58021.2023.00039).
- [29] *X440—Ettus Knowledge Base*. Ettus. Accessed: Aug. 1, 2024. [Online]. Available: <https://kb.ettus.com/X440>
- [30] D. Stanko, G. Sommerkorn, A. Ihlow, and G. D. Galdo, "Enable SDRs for real-time MIMO channel sounding featuring parallel coherent Rx channels," in *Proc. IEEE 95th Veh. Technol. Conf. (VTC-Spring)*, Jun. 2022, pp. 1–5, doi: [10.1109/VTC2022-Spring54318.2022.9860841](https://doi.org/10.1109/VTC2022-Spring54318.2022.9860841).
- [31] W. Kamp, "AXI over Ethernet; a protocol for the monitoring and control of FPGA clusters," in *Proc. Int. Conf. Field Program. Technol. (ICFPT)*, Dec. 2017, pp. 48–55, doi: [10.1109/FPT.2017.8280120](https://doi.org/10.1109/FPT.2017.8280120).
- [32] W. Kamp and S. Wong. (2024). *AXI over Ethernet Base Specification*. [Online]. Available: <https://gitlab.com/axioe/axioe-base-specification/releases/rel-base-0.7.2>
- [33] (2024). *DPDK Homepage*. DPDK Project. [Online]. Available: <https://www.dpdk.org/>
- [34] 2023. *JESD204D*. JEDEC Solid State Technology Association. [Online]. Available: <https://www.jedec.org/system/files/docs/JESD204D.pdf>
- [35] C. W. Wagner, G. Glaeser, G. Kell, and G. Del Galdo, "Every clock counts—41 GHz wide-range integer-N clock divider," in *Proc. Int. Conf. SMACD 16th Conf. (SMACD/PRIME)*, Jul. 2021, pp. 1–4.
- [36] *RF Network-on-Chip (RFNoC) Specification*. Ettus. Accessed: Aug. 1, 2024. [Online]. Available: https://files.ettus.com/app_notes/RFNoC_Specification.pdf
- [37] (2024). *DPDK NVIDIA MLX5 Ethernet Driver*. DPDK Project. [Online]. Available: <https://doc.dpdk.org/guides/nics/mlx5.html>
- [38] S. Giehl, "Evaluation, implementation and benchmarking of high-speed-data-interfaces for FPGA-based measurement applications," B.S. thesis, Inst. Inf. Technol., Technische Universität Ilmenau, Ilmenau, Germany, 2020. [Online]. Available: <https://opac.lbs-ilmenau.gbv.de/DB=1/XMLPRS=N/PPN?PPN=1746330716>
- [39] W. Mansour, N. Janvier, and P. Fajardo, "FPGA implementation of RDMA-based data acquisition system over 100-gb Ethernet," *IEEE Trans. Nucl. Sci.*, vol. 66, no. 7, pp. 1138–1143, Jul. 2019, doi: [10.1109/TNS.2019.2904118](https://doi.org/10.1109/TNS.2019.2904118).
- [40] (2024). *Io_Uring Manpage*. Linux Man-Pages Project. [Online]. Available: https://man7.org/linux/man-pages/man7/io_uring.7.html
- [41] (2024). *Open Manpage*. Linux Man-Pages Project. [Online]. Available: <https://man7.org/linux/man-pages/man2/open.2.html>
- [42] M. Engelhardt, C. Andrich, A. Ihlow, S. Giehl, and G. Del Galdo, "Low-latency analog-to-analog signal processing using PC hardware and USRPs," 2022, *arXiv:2210.06067*.
- [43] (Jun. 2023). *DS889—Zynq UltraScale+ RFSoc Data Sheet: Overview V1.14*. Xilinx. [Online]. Available: <https://docs.xilinx.com/v1/u/en-US/ds889-zynq-usp-rfsoc-overview>
- [44] (Oct. 2023). *PG269—Zynq UltraScale+ RFSoc RFData Converter v2.6 Gen1/2/3/DFE*. Xilinx. [Online]. Available: <https://docs.xilinx.com/v1/en-US/pg269-RF-data-converter>
- [45] M. Engelhardt. *Loss of Packet Pacing Precision Under High Tx Loads*. Accessed: Aug. 1, 2024. [Online]. Available: <https://www.mail-archive.com/users@dpdk.org/msg07395.html>
- [46] Q. Wang, A. Kakkavas, X. Gong, and R. A. Stirling-Gallacher, "Towards integrated sensing and communications for 6G," in *Proc. 2nd IEEE Int. Symp. Joint Commun. Sens. (JC&S)*, Mar. 2022, pp. 1–6, doi: [10.1109/JCS54387.2022.9743516](https://doi.org/10.1109/JCS54387.2022.9743516).
- [47] H. Cesar Alves Costa, S. James Myint, C. Andrich, S. W. Giehl, C. Schneider, and R. S. Thomä, "Bistatic reflectivity and micro-Doppler signatures of drones for integrated communication and sensing," 2024, *arXiv:2401.14448*.
- [48] H. C. A. Costa, S. J. Myint, C. Andrich, S. W. Giehl, C. Schneider, and Reiner S. Thomä, "Modelling micro-Doppler signature of drone propellers in distributed ISAC," in *Proc. IEEE Radar Conf. (RadarConf)*, May 2024, pp. 1–6, doi: [10.1109/RadarConf2458775.2024.10548468](https://doi.org/10.1109/RadarConf2458775.2024.10548468).
- [49] S. Giehl, C. Andrich, M. Schubert, M. Engelhardt, and A. Ihlow, "Receiver bandwidth extension beyond Nyquist using channel bonding," in *Proc. 17th Eur. Conf. Antennas Propag. (EuCAP)*, Mar. 2023, pp. 1–5, doi: [10.23919/EuCAP57121.2023.10133262](https://doi.org/10.23919/EuCAP57121.2023.10133262).



MAXIMILIAN ENGELHARDT received the Bachelor of Science (B.Sc.) and Master of Science (M.Sc.) degrees in computer and systems engineering from Technische Universität Ilmenau, Germany, in 2019 and 2021, respectively. Since 2021, he has been a Research Assistant with Fraunhofer IIS, Germany. His research interests include novel software-define radio solutions and radio-frequency measurement architectures, whereby his focus is on implementing real-time systems on off-the-self hardware.



ALEXANDER EBERT received the B.Sc. and M.Sc. degrees in electrical engineering and information technology from Technische Universität Ilmenau, Germany, in 2012 and 2014, respectively. Since 2021, he has been with the Electronic Measurements and Signal Processing Group, Fraunhofer IIS. His current research interests include microwave and millimeter-wave systems, front-end design, and multiband and multichannel sounding up to THz frequencies.



SEBASTIAN GIEHL received the B.Sc. and M.Sc. degrees in electrical engineering and information technology from Technische Universität Ilmenau, Germany, in 2020 and 2022, respectively. He has been working on parallel computing FPGA systems, since 2017. In 2022, he was with Fraunhofer IIS, Germany. Since 2023, he has been a Research Assistant with Technische Universität Ilmenau on ISAC in mobility applications for the next generation of mobile communication and FPGA-

based, high-bandwidth, and real-time capable SDR systems, for e.g., radar or channel sounding.




MARKUS LANDMANN received the Dipl.-Ing. and Dr.-Ing. degrees in electrical engineering from Technische Universität Ilmenau, Germany, in 2001 and 2008, respectively. Until 2009, he was a Research Assistant with Technische Universität Ilmenau and Tokyo Institute of Technology on wireless propagation, channel modeling, and array signal processing. Since 2010, he has been with Fraunhofer IIS responsible for the Facility for Over the Air Research and Testing, including test and development of satellite and terrestrial-based communication systems (2G-5G). Since 2018, he has been a Chief Scientist of the EMS Department being the Co-Department-Head, since 2022, and responsible for the strategic topics related to 5G standardization.

MICHAEL SCHUBERT received the M.Sc. degree in computer engineering from Technische Universität Ilmenau, in 2018. Since then, he has been working on RFSOC-based measurement systems with the Fraunhofer Institute for Integrated Circuits IIS. His research interest includes novel physical layer algorithms especially for the upcoming 6G mobile communication. In particular, he works on FPGA implementations and related software components for signal processing.



ALEXANDER IHLOW received the Dipl.-Ing. degree in electrical engineering and information technology from Technische Universität Ilmenau, Germany, in 1999, and the Dr.-Ing. degree from Otto von Guericke University Magdeburg, Germany, in 2006. Since 2008, he has been with the Institute of Information Technology, Technische Universität Ilmenau, as a Scientific Staff Member. His current research interests include signal processing, wireless communications, and measurement and testing technology.



GIOVANNI DEL GALDO (Member, IEEE) received the Laurea degree in telecommunications engineering from the Politecnico di Milano and the Dr.-Ing. degree in MIMO channel modeling from Technische Universität Ilmenau, Germany, in 2007. Then, he joined the Fraunhofer Institute for Integrated Circuits IIS, focusing on audio watermarking and spatial sound. Since 2012, he has been leading a joint research group composed of a department with Fraunhofer IIS and as a Full Professor and the Chair of Technische Universität Ilmenau on the research area of electronic measurements and signal processing. His current research interests include the analysis, modeling, manipulation of multidimensional signals, over-the-air testing, and sparsity promoting reconstruction methods.



CHRISTIAN SCHNEIDER received the Diploma degree in electrical engineering from Technische Universität Ilmenau, Germany, in 2001. He is currently the Group Leader of the Electronic Measurements and Signal Processing Department (EMS), Technische Universität Ilmenau and Fraunhofer IIS. His research interests include multi-dimensional channel sounding, radio channel characterization and modeling, and its application to space-time signal processing and ISAC questions. He received

the Best Paper Award at European Wireless Conference, in 2013, and European Conference of Antennas and Propagation, in 2017 and 2019.

CARSTEN ANDRICH received the B.Sc. and M.Sc. degrees in electrical engineering and information technology from Technische Universität Ilmenau, Ilmenau, Germany, in 2014 and 2016, respectively. Subsequently, he joined the Fraunhofer Institute for Integrated Circuits IIS as a Researcher and engineered software-defined radio systems for integrated sensing and communication and radio propagation measurements. Since 2020, he has been with the Institute of Information Technology, Technische Universität Ilmenau, as a Scientific Staff Member. There, he is currently the system architect responsible for the development of distributed, RFSOC-based measurement systems. His area of expertise lies in distributed radio frequency measurements, software-defined radio systems, time and frequency synchronization, and digital signal processing.

...