



Proteus project closure report (public)

June 2024



Version History

VERSION	DATE	DESCRIPTION OF CHANGE
1.0	May 2024	Initial version
1.1	June 2024	Updates for publication

Table of Contents

Version History 2

Table of Contents3

List of Figures 4

1	Introduction	5
1.1	References.....	6
1.2	Abbreviations	6
2	Proteus architecture	8
2.1	5G Network Overview	8
2.1.1	<i>University of Bristol 5GUK Test Network</i>	8
2.2	Parallel Wireless gNB CU and DU	8
2.3	DU PHY and Stack separation.....	10
3	Proteus physical layer	11
3.1	Guiding principles.....	11
3.2	Proteus physical layer architecture.....	12
3.3	Wireless performance validation framework	13
4	Platforms enabled in the project	14
5	5GUK deployment	15
6	Conclusion	17

List of Figures

Figure 1: VBBU in the 5G Network	8
Figure 2: 5G protocol stacks – control-plane	9
Figure 3: 5G protocol stacks – user-plane	9
Figure 4: SCF 5G FAPI Ports	10
Figure 5: Proteus O-DU-PHY application and its main external dependencies	12
Figure 6: Proteus deployment on the 5GUK outdoor test-bed	15
Figure 7: Cell reselection test on the 5GUK test-bed.....	16

1 INTRODUCTION

Proteus is a project executed under the UK's department for Digital, Culture, Media and Sport (DCMS, now DSIT) *Future of RAN Competition* (FRANC). In this report we provide an overview of the project's stated objectives and of its achievements.

Over the past 15 years the wireless communications industry has begun a profound transformation, triggered by technological improvements in CPU and datacentre technologies and the commoditisation of compute platforms. Intel was first to realise that the wireless communications industry could benefit from these opportunities and started in 2010 what was later to become the FlexRAN programme. In 2016 Facebook, Intel and Nokia founded the Telecom Infra project to federate the efforts of technology vendors, RAN equipment manufacturers, integrators and operators, towards a goal of reducing the cost of communications infrastructure and make these accessible to most. In 2018, 5 operators (AT&T, China Mobile, Deutsche Telekom, NTT DOCOMO and Orange) founded the O-RAN Alliance industry group to define standard interfaces and "organise" the ecosystem to be more open and interoperable.

The benefits of these approaches are clear: Open interfaces enable more players to enter the RAN infrastructure market, increase competition and as a result accelerate innovation and change. Generic COTS platforms enable economies of scale, a reduction in Capex and facilitate reuse of application software through platform generations. Standard tools, virtualisation and cloud-based technologies enable scalability, agility and reduced Opex.

The first part of the network to undergo this change at scale was the core network, but the industry focus has now turned towards the more technically challenging RAN. Within the RAN the DU, and particularly the physical layer, also known as the PHY, layer 1 or L1, presents a difficult challenge because of the computational intensity of the signal processing algorithms within the base-station. Meeting these computational demands had traditionally required the design of purpose-built chips known as ASICs, which demand very large investments, limit access to the market to large players and are typically not interchangeable and not backward compatible through generations. The latest general-purpose processors CPUs (GPPs) have become sufficiently performant to support these workloads due to their ability to operate on arrays of data (vectors) rather than element per element – a feature called SIMD (Single Instruction Multiple Data). But SIMD programming is a specialist task and the language used to do so is CPU-specific. This means that SIMD code written for an Intel GPP will not run on an Arm based GPP, and vice versa.

It has taken 10 or so years for Intel to mature its FlexRAN offering and "seed" the market under a proprietary license to enable use of its processors, something which other GPP vendors including AMD and Arm-based CPU vendors have not done. The lack of portability of SIMD code and the licensing restrictions constitute significant barriers to enabling comprehensive competition and supply chain diversity for GPP-based Open RAN DU products.

The objective of Proteus was to solve this challenge and enable better supply chain diversification by:

1. Developing a 5G DU PHY application deployable on any general-purpose CPU, ideally from a single source of software.
2. Validating the approach by deploying the Proteus 5G DU PHY on differentiated platforms including at least one x86 based and one Arm based GPP, in a representative environment.

In doing so Parallel Wireless and the Proteus project partners BT, The University of Bristol and Real Wireless would enable a more disaggregated supply chain, one of the key objectives of the 5G supply chain diversification strategy [1].

1.1 References

Ref	Doc number and version	Title
[1]	ISBN 978-1-5286-2283-7	UK 5G supply chain diversification strategy
[2]	N/A	5GUK test bed
[3]	SCF 222.10.04	Small Cell Forum 5G FAPI: PHY API Specification
[4]	3GPP TS 38.300 v15.13.0	NR and NG-RAN Overall Description
[5]	3GPP TS 38.401 v15.9.0	NG-RAN; Architecture description
[6]	3GPP TS 38.470 v15.8.0	NG-RAN; F1 general aspects and principles
[7]	3GPP TS 38.801 v14.0.0	Study on NR; Radio access architecture and interfaces
[8]	3GPP TS 38.104 v15.16.0	NR; Base Station radio transmission and reception
[9]	3GPP TS 38.211 v15.10.0	NR; Physical channels and modulation
[10]	3GPP TS 38.212 v15.13.0	NR; Multiplexing and channel coding
[11]	3GPP TS 28.541 v15.8.0	5G Network Resource Model
[12]	O-RAN.WG4.CUS.0 v07.00	O-RAN Fronthaul Working Group; Control, User and Synchronization Plane Specification
[13]	N/A	AttoCore Atto5G
[14]	N/A	Benetel RAN650
[15]	N/A	Genevisio PAC-012
[16]	N/A	gRPC project

1.2 Abbreviations

Term	Definition
ASIC	Application Specific Integrated Circuit
BBDev	Baseband Device (in DPDK)
CN	Core Network
CPU	Central Processing Unit
CU	Central Unit
DCMS	Department for Digital, Culture, Media and Sport
DPDK	Data Plane Development Kit
DU	Distributed Unit
FAPI	Function Application Platform Interface
FEC	Forward Error Correction
FH	Fronthaul (DU-RU interface)
FRANC	Future of RAN Competition
GM	Grand Master (PTP reference clock)
gNB	5G basestation
GPP	General Purpose Processor
NIC	Network Interface Card

Term	Definition
NR	New Radio (5G)
O-RAN	Open RAN
PCIe	Peripheral Component Interconnect express
PTP	Precision Timing Protocol (IEEE-1588)
RAN	Radio Access Network
RRH	Remote Radio Head
RU	Radio Unit
SCF	Small Cell Forum
SDK	Software Development Kit
SIMD	Same Instruction Multiple Data (vectorisation)
SoC	System on a Chip
TRL	Technology Readiness Level
VBBU	Virtualised Baseband Unit

2 PROTEUS ARCHITECTURE

2.1 5G Network Overview

Figure 1 below identifies the key elements in a 5G standalone network, highlighting the VBBU where the software development for the Proteus project was concentrated:

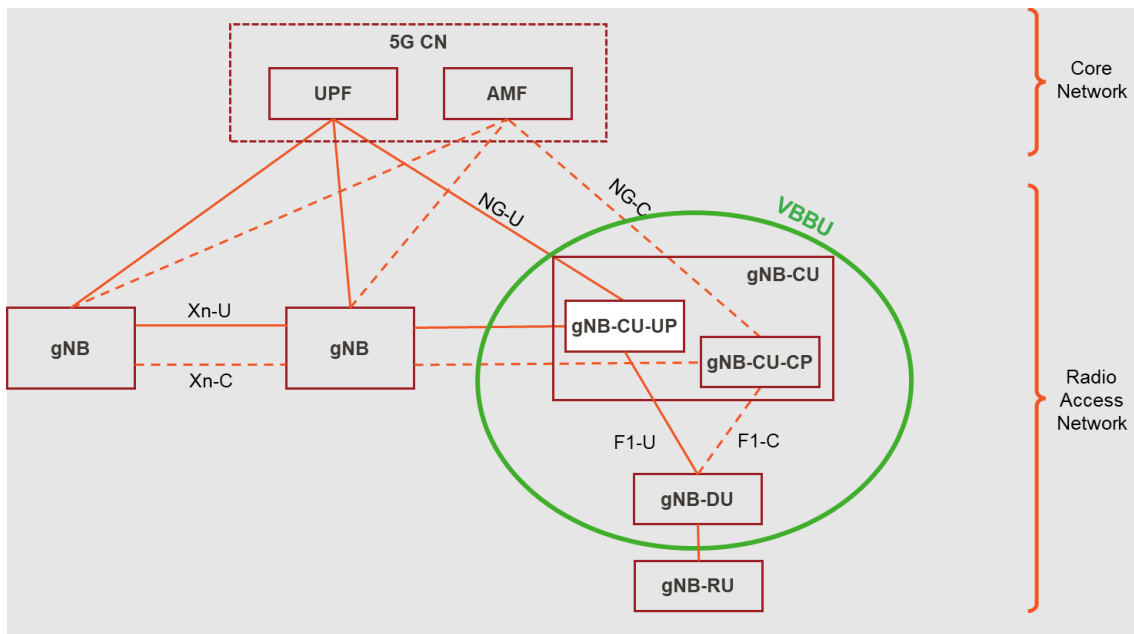


Figure 1: VBBU in the 5G Network

2.1.1 University of Bristol 5GUK Test Network

For the validation phase the Proteus project installed its equipment in the 5GUK [2] test network. In this network the commercial grade *Atto5GC* core from *Attocore* [13] was integrated with the Parallel Wireless 5G RAN solution comprising the VBBU gNB (CU and DU), all located in the datacenter of the University of Bristol’s Smart Internet Lab, and a number of RAN650 O-RU radios from Benetel in band n77 [14] were deployed at locations around the city center. The 5GUK framework provides an environment enabling validation of the Proteus solution at Technical Readiness Level (TRL) 5 to 6.

2.2 Parallel Wireless gNB CU and DU

The Parallel Wireless gNB solution used in Proteus consists of a single VBBU server running a CU and DU as distinct entities. The CU-DU partitioning follows the 3GPP *Split Option 2* definition from [7], elaborated in [5], with the F1 interface between them. The CU is further split between the user-plane, connected back to the UPF in the 5G core,

and the control-plane, connecting to the AMF. The DU-RU partitioning follows the 3GPP *Split Option 7-2* definition from [7], elaborated by O-RAN in [12], with the upper-PHY running in the DU and the lower-PHY located in the RU.

In the Proteus project the CU connects 1:1 with the DU, together forming a vNode and serving a single 5G cell (sector), for simplicity.

Pushing down into more detail, Figure 2 and Figure 3 below capture the protocol stacks for the 5G control-plane and user-plane respectively; refer to [4] and [6] for more details. In the Proteus project we focussed particularly on the NR PHY in the gNB-DU as highlighted – this is where the platform-agnostic functionality is required:

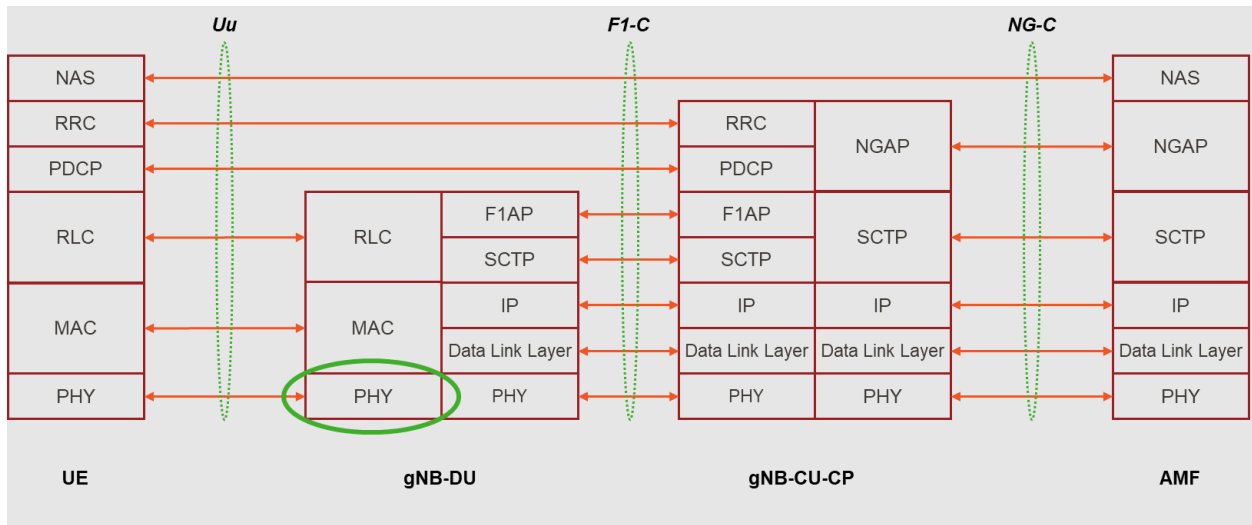


Figure 2: 5G protocol stacks – control-plane

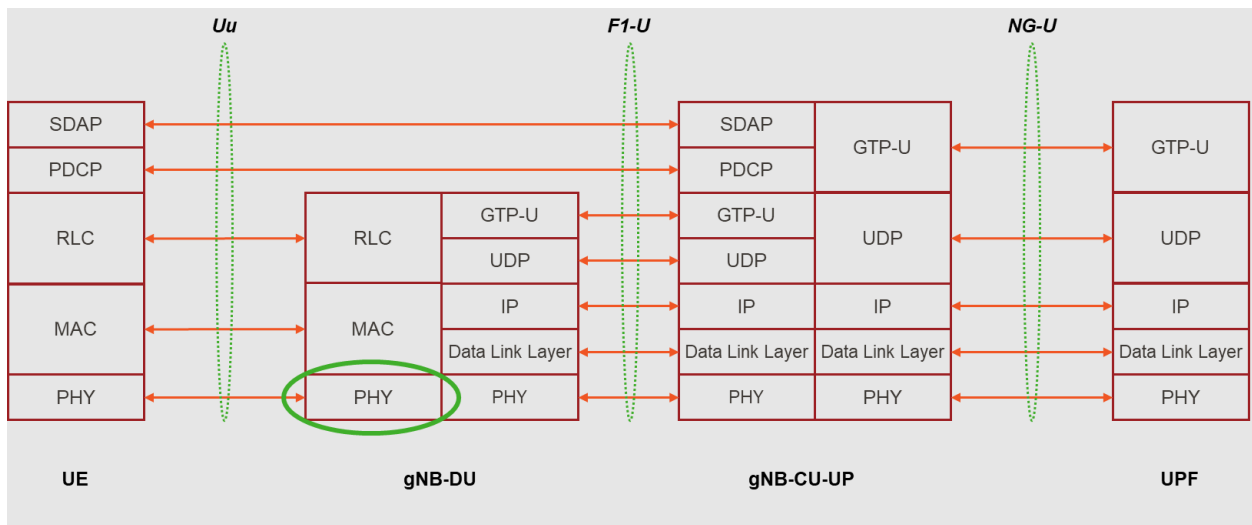


Figure 3: 5G protocol stacks – user-plane

With the move to a generic, portable PHY, it is also necessary to move to a generic L2-L1 API to enable maximum flexibility. In Proteus we decided to adopt the 5G FAPI [3] from the Small Cell Forum, a widely used open-source definition. The SCF also provide ongoing development of the 5G FAPI to align with advances in the 3GPP releases.

Note that the full SCF 5G FAPI¹ defines a set of message ports which include functionality such as SON and RRH configuration, but for the purposes of the Proteus project we only focused at the PHY ports: P5 (PHY control-plane) and P7 (PHY user-plane) as illustrated in Figure 4 below:

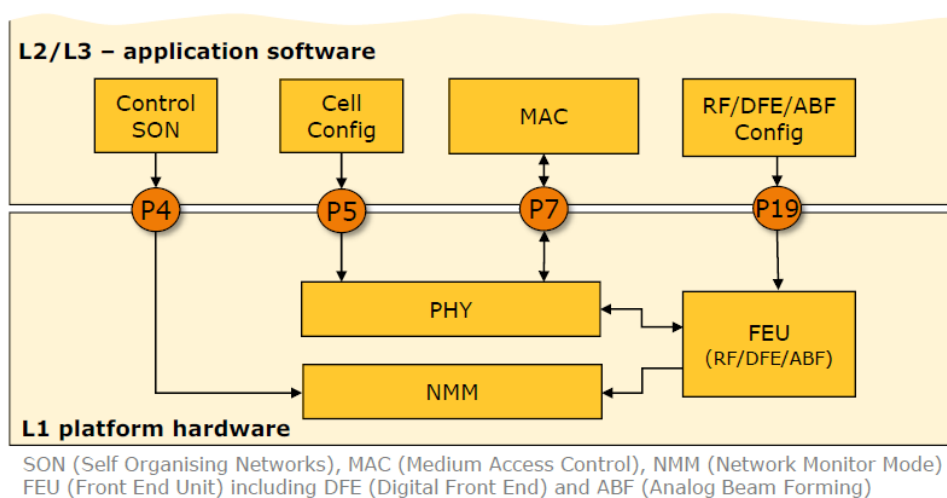


Figure 4: SCF 5G FAPI Ports

2.3 DU PHY and Stack separation

The Proteus project’s main focus is the NR PHY part of the gNB-DU. In order to control scope all other elements of the gNB DU were kept unchanged in the project, based on Parallel Wireless’ standard 5G product. For this reason the L2/L3 elements were pinned to a fixed Intel Xeon based platform, and the PHY + radio interface deployed on a range of target platforms, following a topology known as in-line.

¹ <https://www.smallcellforum.org/5g-fapi-suite/>

3 PROTEUS PHYSICAL LAYER

Traditional RAN solutions, in particular the “core” signal processing elements of the PHY, are tightly coupled with the platform they execute on for reasons of compute efficiency, which in turns translates into cost and power efficiency. This coupling can take place in a number of ways including:

- Tailoring the software architecture of the PHY to the underlying CPU instruction set (e.g. through using CPU-specific SIMD intrinsics) and CPU core characteristics (e.g. through tailored mapping of tasks to cores and utilisation of proprietary thread scheduling frameworks) of the target platform;
- Utilising hardware acceleration and/or offload for compute intensive and/or latency critical functions;
- Utilising platform specific tools – e.g. specialised compilers, libraries.

3.1 Guiding principles

In Proteus the development of the PHY was guided by the following principles:

- **Functional equivalence:** The hardware abstracted Proteus software must be deployable with the same feature set on the range of enabled platforms. Exceptions are only acceptable on hardware capacity grounds.
- **Algorithmic near-equivalence:** Hardware abstraction techniques must enable an algorithm to be implemented with equivalent or near equivalent (wireless) performance on the range of enabled platforms.
- **Code reuse:** The hardware abstraction framework must enable the PHY application to be written and maintained as one code base, and to limit the amount of platform specific code to a well segregated minimum.
- **Compute performance:** The hardware abstraction should enable the compute performance features of the target platforms – specifically SIMD vectorisation – to be exploited optimally or near optimally.
- **Flexibility:** The PHY application solution is expected to be deployable on GPPs ranging from low cost embedded SoCs to high end COTS server cores to cloud platforms. As such the Proteus software architecture shall offer some flexibility e.g. be deployable onto “pools” of cores of varying size.
- **IP portability:** The Proteus PHY must not use any 3rd party tool or library restricting its utilisation onto processors of a specific type or vendor.
- **Future proofing:** Where possible the abstracted code shall be constructed using techniques facilitating enablement on additional platforms or CPU architectures.

3.2 Proteus physical layer architecture

In Proteus the functionality of the PHY has been decomposed into a number of sub-processes, or threads, mappable to a variable number of cores. Within each of the processing threads the PHY needs to be as decoupled as possible from the implementation specificities to the platform, for example its instruction set, memory interface, network interface, or accelerated functions capabilities. The main hardware abstraction interfaces are shown in Figure 5.

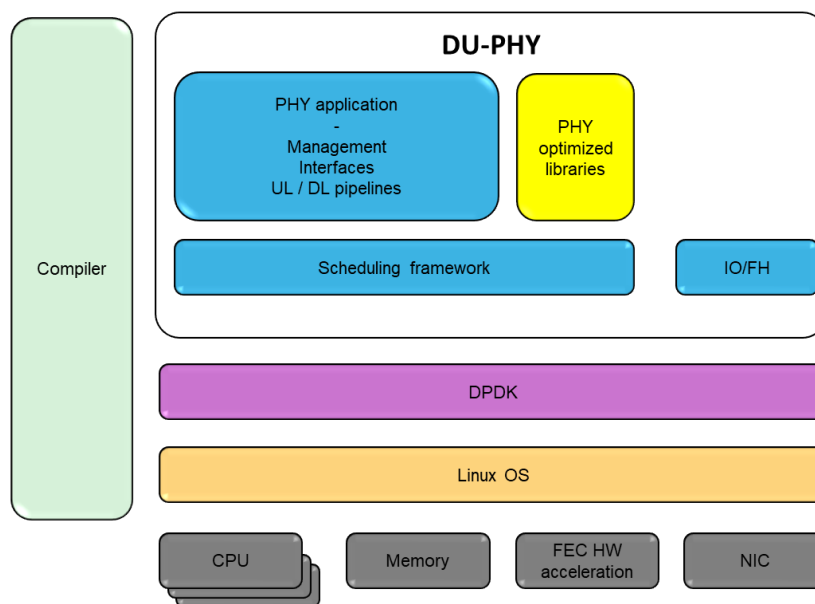


Figure 5: Proteus O-DU-PHY application and its main external dependencies

The key elements adopted by Proteus as the basis for its architecture, are:

- **Linux scheduler:**

Linux is the most widely deployed operating system and is increasingly well suited for real time applications. During the Proteus PHY development we analysed context switching latency times in detail and concluded that the flexibility benefits of a pre-emptive task execution framework outweigh the small performance cost of task switching. Together with suitably partitioned workloads and the ability to parallelise some tasks across pools of cores, this allows great flexibility in the choice of the target CPU.

- **DPDK:**

DPDK is an open-source library available on Intel and AMD x86, Arm, as well as RISC-V platforms. It provides an efficient abstraction framework to interface hardware functions including the platform's memory interface, the network interface hardware and the FEC lookaside accelerators, with a common API. The library also enables a generic interface to other blocks such as GPUs and machine learning engines.

- **SIMD abstraction and substitution libraries:**

Utilisation of the SIMD capabilities of different types of CPUs from a common source code, was an imperative goal at the start of the Proteus project and presented a number of unknowns. We evaluated several SIMD substitution and SIMD abstraction techniques and successfully demonstrated that these techniques, when utilised appropriately, can provide the required portability whilst maintaining an optimal level of performance.

3.3 Wireless performance validation framework

Validation of the wireless performance is an important aspect in a PHY product development. In Proteus we wanted to leverage Parallel Wireless' existing wireless performance validation framework based on Matlab. In legacy x86 based flow we would validate the performance of production code by embedding it within Matlab simulations using the MEX (Matlab Executable) framework. For non x86 target platforms, despite Proteus being a common source code application, it is important to repeat this validation to measure the impact of low-level silicon implementations e.g. in relation to rounding. We faced the problem that Matlab does not run natively on Arm architecture and therefore evolved our validation framework to utilise remote procedure calls through the open source gRPC library [16].

4 PLATFORMS ENABLED IN THE PROJECT

The Proteus project had the initial ambition to deploy the PHY to at least one x86-based and one Arm-based platform. The project selected the Parallel Wireless standard Intel Xeon x86-based vBBU platform and the NXP Layerscape Arm-based Genevisio PAC-012 card [15]. Over the course of the project this objective was exceeded and a further 3 platforms were enabled as shown in **Table 1**, covering the following 3 instruction sets: AVX512, Neon, SVE2.

Form factor	Server / Card OEM	Model	CPU	Frequency	Peripherals	Notes
1u server	Supermicro	SYS-110P-FDWTR	Intel Xeon 6312u 24 Icelake-SP cores	2.7GHz	NIC: Intel FEC HW: Intel ACC100	Proteus initial scope platform
PCIe Gen3 x4	Genevisio	PAC-012	NXP Layerscape LX2162A 16 Arm A72 cores	2GHz	NIC : NXP LX2162A FEC HW : NXP LA1200	Proteus initial scope platform, mounted within the Intel Xeon server
2u server	Gigabyte	Edge	Ampere Altra Q64-30 64 Arm N1 cores	3GHz	-	Additional platform
SBC	Nvidia	Jetson AGX	12 Arm A78 cores	2.2GHz	-	Additional platform
ATX	Marvell	CRB	Marvell Octeon 10 24 Arm N2 cores	2.5GHz	-	Additional platform

Table 1: Proteus deployment platforms

5 5GUK DEPLOYMENT

The deployment of the Proteus PHY around the 5GUK testbed was an important step to validate operation of the Proteus PHY in a representative environment at TRL5/6. The deployment is centralised and based on commercial grade elements:

- Hosted in the centralised datacentre at the University of Bristol Smart Internet lab:
 - Commercial 5G core network from AttoCore
 - Commercial SMO and EMS from Parallel Wireless
 - Commercial CU from Parallel Wireless
 - 5G DU, integrating the Parallel Wireless commercial 5G DU-Stack and O-RAN fronthaul software, with the Proteus PHY – with both fully Intel x86, and x86 + in-line Arm options.
- A network of dark optical fiber connecting the Smart Internet Lab to 1 indoor and 3 outdoor cell sites, carrying the O-RAN 7.2A fronthaul traffic.
- Cell sites equipped with RAN650 n77 O-RU radioheads from Benetel, and Alpha AW3924-T0-F antennas
- Commercial Samsung Galaxy S22 5G handsets, with sim cards matching the deployed network parameters.

Figure 6 presents a high level map of the deployment:



Figure 6: Proteus deployment on the 5GUK outdoor test-bed

The test-bed was used to validate operation of the test network in the following scenarios:

- Establishment of cells on both Proteus-x86-based and Proteus-Arm-based DUs
- Registration of UEs on all cells
- Bi-directional data transfer
- Cell reselection between sites, including from Intel x86 to Arm based DUs, or vice-versa.



Figure 7: Cell reselection test on the 5GUK test-bed

6 CONCLUSION

The Proteus project achieved or exceeded all its intended benefits, in particular:

1. It accelerated the development of a single-source, hardware-agnostic 5G physical layer application and its deployment on a range of x86 and Arm based CPU architectures.
2. It significantly advanced our understanding of Arm-based CPUs and confirmed their suitability (in addition to x86 based CPUs) to deploy RAN functions including the algorithmically demanding PHY.
3. It progressed the telecoms supply chain diversification agenda of the UK and significantly stimulated competition in the CPU and platforms vendors ecosystem. This opens up opportunities for products which are better tailored to each deployment topology, more energy efficient and cheaper.
4. It deepened the ties between Parallel Wireless and our project partners and suppliers BT, The University of Bristol, Real Wireless and Benetel, and with the teams at DSIT.
5. It raised the profile of Parallel Wireless within the O-RAN community and promoted UK PLC and DSIT as an innovation engine.

Since the completion of the project Parallel Wireless has formally adopted a number of the methodologies developed in Proteus in its all-G GreenRAN™ Hardware Agnostic product line and has expanded its scope to a fully portable O-RAN DU solution in collaboration with additional partners Adtran, Accelercomm, Scotland5G and the University of York, with further support from DSIT through the FRANC DU-Volution project.

We are extremely grateful to the DSIT team for their support throughout the Proteus project.