Energy-Aware RAN x-Haul over ONaaS-Enabled Metro-Access Networks

Energy-Aware RAN x-Haul over ONaaS-Enabled Metro-Access Networks

Ahtisham Ali*†, Ricardo J. B. Pousa‡, Muhammad Umar Masood†, Giang T. Nguyen‡,
Frank H. P. Fitzek‡, Alessandro Galardini*, Vittorio Curri†
*Consorzio TOP-IX, Torino, Italy
†DET, Politecnico di Torino, Torino, Italy
‡ComNets, Technische Universität Dresden, Dresden, Germany
Email: ahtisham.ali@top-ix.org

***Abstract*—Disaggregated radio access network (RAN) architectures require metro-access transport networks that can satisfy stringent fronthaul (FH) and midhaul (MH) constraints while improving energy efficiency under time-varying traffic. This paper evaluates an Optical Network as-a-Service (ONaaS) enabled converged metro-access optical network for flexible RAN *x*-haul provisioning using Digital Subcarrier Multiplexing (DSCM) coherent transmission. A 400G DSCM transceiver (TRx) employing DP-16QAM with 16 digital subcarriers is considered, providing fine granularity for optical resource allocation. The proposed evaluation framework combines a realistic 24-hour mobile traffic profile with physical-layer-aware routing, wavelength, and spectrum assignment (RWSA), accounting for FH/MH latency, quality-of-transmission (QoT), Distributed Unit (DU) capacity, and residual C-band availability. Two DU operation strategies are compared: a dynamic scheme that activates or hibernates DUs according to hourly traffic demand and reachability constraints and a static baseline where all candidate DUs remain continuously active. Results show that dynamic DU operation reduces the active DU set from five DUs at peak load to three DUs during off-peak periods, achieving 13.2% DU-energy savings and 15.4% higher DU power efficiency at equal delivered load. However, the static scheme carries 40.0% more daily traffic, highlighting the trade-off between energy efficiency and service capacity. Blocking remains close to the 1% target and is dominated by spectrum exhaustion, indicating that residual optical spectrum is the main limiting factor in the considered ONaaS scenario.**

***Index Terms*—metro-access networks, energy-aware RAN, 5G/B5G, DSCM, ONaaS, *x*-haul.**

I. Introduction

Future mobile networks are expected to support increasingly dense radio deployments, heterogeneous services, and highly variable traffic demands, ranging from enhanced Mobile Broadband (eMBB) to Ultra-Reliable Low-Latency Communications (URLLC) and massive Machine-Type Communications (mMTC) [1, 2]. This evolution imposes stringent requirements not only on the radio interface, but also on the transport infrastructure interconnecting radio access, edge, and centralized processing functions [3]. Consequently, metro-access optical networks are becoming a key component of 5G and beyond-5G systems, where capacity, latency, reliability, and energy efficiency must be jointly addressed [4].

Modern RAN architectures rely on functional disaggregation, where baseband processing is split across Radio Units (RUs), Distributed Units (DUs), and Central Units (CUs) [5, 6]. This enables flexible function placement and resource pooling, but introduces strict transport requirements over the FH (RU–DU) and MH (DU–CU) segments [7]. In particular, lower-layer splits (LLS) impose tight latency and bitrate constraints, while higher-layer splits (HLS) relax transport requirements at the cost of reduced centralization gains [8, 9]. Therefore, the feasibility of disaggregated RAN operation depends on both computing-resource availability and the ability of the underlying transport network to provide latency-compliant and quality-aware connectivity. Converged metro-access optical networks offer a scalable transport solution to support RAN *x*-haul by leveraging deployed fiber infrastructure and aggregating access traffic toward metro and data-center locations [10]. Dense Wavelength Division Multiplexing (DWDM), coherent transmission, and programmable optical-network abstractions enable shared optical resources to be offered through an ONaaS framework [11]. In this context, QoT estimation becomes essential, since a topologically available path may still be infeasible if accumulated physical (PHY) impairments violate latency, generalized signal-to-noise ratio (GSNR) requirements [12]. Moreover, DSCM coherent optics can improve the granularity of optical resource allocation by mapping moderate-rate RAN traffic flows onto independently assignable subcarriers [13]. This makes DSCM-enabled point-to-multipoint (P2MP) transmission particularly suitable for metro-access RAN aggregation [14].

Beyond capacity and route feasibility, energy efficiency is becoming a major design objective for future RAN transport. Mobile traffic exhibits strong temporal variations, with low-load intervals and peak-hour demand periods. Keeping all DUs continuously active can maximize service availability, but it wastes energy during off-peak hours. Conversely, hibernating selected DUs can reduce energy consumption but may reduce path diversity, concentrate traffic on fewer aggregation points, and increase optical resource blocking. This creates a trade-off between energy saving and service capacity that must be evaluated jointly across the RAN and optical transport layers.

Prior works have addressed RAN function placement, DU deployment, and energy-efficient split selection using optimization or learning-based methods [15–17]. However, these studies often abstract the optical transport layer and do not explicitly account for route-dependent QoT, wavelength/subcarrier occupation, or FH/MH latency constraints. Conversely, optical feasibility studies for converged metro-access RAN transport mainly consider static route evaluation, supported capacity, or PHY-layer feasibility under fixed traffic assumptions [18]. The impact of realistic time-varying RAN traffic on DU activation, optical resource allocation, carried traffic, and energy consumption in a shared ONaaS infrastructure remains less explored.

In this work, we evaluate a converged metro-access optical network supporting RAN x-haul under a realistic 24-hour traffic profile. The hourly load is derived from cellular traffic measurements and mapped onto FH and MH bitrate requirements associated with standardized 5G functional split assumptions [8, 19, 20]. Traffic requests are provisioned over RU–DU–CU paths using PHY-layer-aware routing, wavelength, and spectrum assignment, considering latency, GSNR, bit-error rate (BER), DU capacity, and residual C-band availability. We compare a dynamic DU activation/hibernation strategy against a static all-DUs-active baseline, quantifying the trade-off among DU energy consumption, carried traffic, blocking probability (BP), and optical-link congestion.

II. Enabling Technologies for Optical Network-as-a-Service

A. RAN Functional Splits and x-haul Requirements

The evolution of mobile networks has progressively changed the RAN from a monolithic base-station architecture toward a disaggregated and cloud-native system. In conventional deployments, radio-frequency (RF) and baseband processing are co-located at the cell site: the radio head performs filtering, amplification, frequency conversion, and analogue/digital conversion, while the baseband unit (BBU) executes PHY-layer processing, modulation, coding, scheduling, and upper-layer protocol functions. Although this architecture simplifies deployment, it requires dedicated processing resources at each site and may lead to inefficient utilization under spatially and temporally varying traffic demand.

Modern RAN architectures address this limitation through disaggregation, which enables processing pooling, simplified upgrades, improved coordination, and more flexible deployment but makes the transport network a critical part of the RAN system, since the feasibility of each RU–DU–CU connection depends on the selected function placement and on the capacity, latency, synchronization, and QoT characteristics of the underlying x-haul infrastructure. The placement of RAN functions is governed by the selected functional split, which defines where the radio protocol stack is divided between two processing entities. Moving more functions toward the CU increases centralization gains, such as coordinated scheduling, interference management, and statistical multiplexing of processing resources, but also increases the bitrate and latency constraints imposed on the transport network. Conversely, moving more functions toward the RU or DU relaxes transport requirements but reduces centralization benefits and increases the processing burden at the access site [8, 9]. Therefore, the functional split represents a fundamental trade-off among radio coordination, processing centralization, transport capacity, and latency.

Several standardization and industry bodies have addressed RAN functional disaggregation and x-haul interfaces, including 3rd Generation Partnership Project (3GPP) [21], Common Public Radio Interface (CPRI)/enhanced-CPRI (eCPRI) [22], Next Generation Mobile Network (NGMN) [23], the O-RAN Alliance [24], the Small Cell Forum (SCF) [25], IEEE 1914 Next Generation Fronthaul Interface (NGFI) [26], and the Telecom Infra Project (TIP)

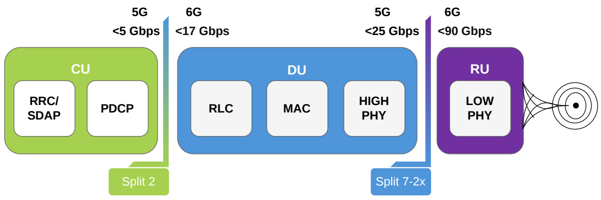

Fig. 1: Distribution of functionalities on the considered CU & DU, with resulting transport network required throughputs

[27]. Among these, 3GPP defines functional split options across the radio protocol stack, while O-RAN specifies an open and interoperable FH (O-FH) interface between the O-DU and O-RU.

In the 3GPP framework, the split options are defined along the radio resource control (RRC), packet data convergence protocol (PDCP), radio link control (RLC), medium access control (MAC), PHY, and RF boundaries of the protocol stack. HLS options place the split above the PHY layer, whereas LLS options divide the PHY layer itself. Option 2 separates PDCP from RLC and has been standardized by 3GPP for the CU–DU interface through F1-C and F1-U control- and user-plane interfaces [5]. O-RAN adopts the O-FH interface based on split 7.2x, which is closely related to 3GPP intra-PHY split options and is designed to enable multi-vendor interoperability between the O-DU and O-RU [7, 28, 29].

In this configuration, the O-CU communicates with the O-DU through the 3GPP F1 interface corresponding to Option 2, while the O-DU communicates with the O-RU through the O-FH interface. This architecture is particularly relevant for flexible RAN deployment, since it allows the DU and CU to be placed at different metro or data-center locations, provided that the transport network can satisfy the corresponding FH and MH constraints. For split 7.2, the FH bitrate depends on radio bandwidth, antenna ports, multiple-input multiple-output (MIMO) layers, modulation format, in-phase and quadrature (IQ) resolution, and occupied resource elements. Using representative 5G reference parameters, the bidirectional bitrate per RU is commonly considered to be approximately 22.2 Gbps [30, 31]. At the same time, stringent latency limits must be respected, typically on the order of $250\,\mu s$ for lower-layer FH and around $1\,ms$ for MH transport in the considered x-haul scenario [8]. These requirements motivate the use of optical metro-access infrastructures capable of providing high capacity, low latency, and flexible resource allocation for disaggregated RAN operation.

As shown in Fig.1, the considered x-haul service follows this disaggregated RAN model: traffic is generated at RU-serving access nodes, transported through an FH segment toward an active DU at some points of presence (PoPs), and then carried over an MH segment toward a CU hosted at a data center (DC). This makes the optical transport layer responsible not only for providing connectivity but also for ensuring latency, spectrum, and QoT constraints under time-varying RAN traffic.

B. DWDM and DSCM as an Enabler for 5G Transport

DWDM is a key technology enabling multiple optical channels to coexist over the same fiber infrastructure, DWDM enables efficient sharing of metro-access resources

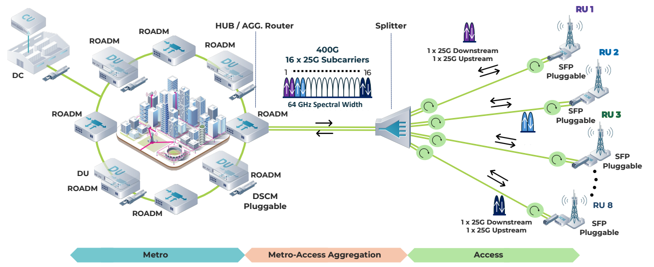

Fig. 2: P2MP converged metro-access network architecture using DSCM pluggable TRxs.

among heterogeneous services, including broadband access, enterprise connectivity, and mobile *x*-haul traffic [11, 32]. However, in RAN *x*-haul scenarios, optical connectivity cannot be evaluated only in terms of graph reachability or shortest-path distance. A candidate route must also satisfy PHY-layer feasibility, since impairments accumulated along the path may degrade the received signal below the required quality threshold. In transparent optical networks, relevant impairments include fiber attenuation, amplified spontaneous emission (ASE) noise from optical amplifiers, nonlinear interference (NLI), and insertion losses introduced by optical nodes such as reconfigurable optical add-drop multiplexers (ROADMs), splitters, and circulators that may reduce the received signal quality below the required threshold [10]. Therefore, QoT-aware routing is essential in an ONaaS framework. Open PHY-layer modeling tools such as GNPy provide a practical way to estimate QoT over multi-span optical routes [12]. This is particularly useful in metro-access networks, where routes can include a mixture of short access links, splitter/circulator losses, metro ROADMs, and amplified fiber spans.

DSCM further improves spectral granularity by dividing a coherent carrier into independently assignable digital subcarriers [13]. This is attractive for RAN *x*-haul because moderate-rate RU flows can be mapped onto subcarrier-level resources instead of consuming a full wavelength. In a P2MP coherent architecture, a hub TRx generates multiple subcarriers that can be distributed toward different leaf nodes, while each leaf processes only its assigned subcarriers [33]. Fig. 2 illustrates the considered ONaaS-enabled metro-access architecture. By combining QoT-aware DWDM routing with DSCM-based subcarrier allocation, the transport network can support fine-grained RAN traffic provisioning while capturing the interaction between DU activation, path diversity, spectrum occupation, and energy consumption.

III. Modeling and Simulation Scenario

The evaluation is performed on the converged metro-access optical topology shown in Fig. 3, whose main parameters are summarized in Table I. The network is modeled as an ONaaS-enabled infrastructure where 40% of the C-band is assumed to be available for RAN *x*-haul services, while the remaining spectrum is considered occupied by pre-existing fixed, enterprise, or cloud traffic. The topology includes metro ROADMs, access aggregation nodes, RU-serving access leaves, and two CU/DC nodes. Access nodes are organized in clusters, each composed of one splitter/MUX node and four RU-serving leaf nodes connected to the metro network through an access–metro interface. DUs are hosted at selected metro PoPs, whereas CUs are located at the DC nodes. Each RAN service request is therefore represented as an end-to-end RU–DU–CU connection composed of an FH segment and an MH segment.

A. DSCM-Based Traffic and Spectrum Model

The optical transmission system is based on an experimentally characterized DSCM coherent pluggable TRx. The characterization follows the back-to-back methodology in [34, 35], providing a key performance model that relates received optical power (ROP) and signal-to-noise ratio (SNR). This model allows the PHY-layer feasibility of each candidate route to account for both line impairments and implementation-related TRx penalties. The considered DSCM signal provides a $400\,\mathrm{Gbps}$ optical carrier divided into 16 digital subcarriers under DP-16QAM configuration, each carrying up to $25\,\mathrm{Gbps}$ and occupying approximately $4\,\mathrm{GHz}$, for an overall spectral occupation of about $64\,\mathrm{GHz}$ [13]. This granularity is well aligned with the considered 5G split 7.2 FH demand of approximately $22.2\,\mathrm{Gbps}$ per RU [31]. Accordingly, RU traffic is mapped onto DSCM subcarrier resources rather than consuming an entire wavelength under time-varying RAN traffic. The offered RAN load follows the normalized 24-hour traffic profile shown in Fig. 4. The profile is derived from cellular traffic measurements reported in [19, 20] and captures the typical diurnal behavior of mobile networks, with reduced traffic during off-peak hours and higher demand during peak periods. For each hour t, the normalized load $L(t)$ scales the peak FH and MH bitrates as

$$r^{\mathrm{FH}}(t) = L(t)R^{\mathrm{FH}}_{\max}, \qquad r^{\mathrm{MH}}(t) = L(t)R^{\mathrm{MH}}_{\max}, \tag{1}$$

where $R^{\mathrm{FH}}_{\max}$ and $R^{\mathrm{MH}}_{\max}$ represent the peak split-dependent FH and MH rates. This enables the network to be evaluated under realistic hourly load variations instead of a fixed peak-load assumption.

B. QoT and Latency Evaluation

For each generated request, candidate FH and MH routes are computed using a k-shortest-path procedure with hop limits of 4 and 8 for FH and MH, respectively. Each candidate route is evaluated in terms of spectrum availability, PHY-layer QoT, and functional-split-aware latency before admission.

The PHY layer is modeled through an impairment-aware abstraction consistent with GNPy-based optical performance evaluation [12]. Fiber spans are modeled as standard single-mode fiber (SSMF) with attenuation $0.25\,\mathrm{dB/km}$. ROADMs introduce $10\,\mathrm{dB}$ insertion loss, circulators introduce $0.81\,\mathrm{dB}$ insertion loss, and splitter loss is computed as $10\log_{10}(N_{\mathrm{out}})$, where N_{out} is the split ratio [10, 36] while EDFAs are modeled with a $6\,\mathrm{dB}$ noise figure, and the channel launch power is set to $0\,\mathrm{dBm}$. A $1\,\mathrm{dB}$ system margin is applied before checking route feasibility. The effective SNR combines line and experimentally characterized TRx contributions as $\mathrm{SNR}^{-1} = \mathrm{GSNR}^{-1} + \mathrm{SNR}^{-1}_{\mathrm{TRx}}$. For DP-16QAM, the BER is estimated as $\mathrm{BER} = \frac{3}{8}\,\mathrm{erfc}\left(\sqrt{\frac{\mathrm{SNR}}{10}}\right)$. A route is considered QoT-feasible only if the resulting BER remains below the adopted pre-FEC threshold of 10^{-2}.

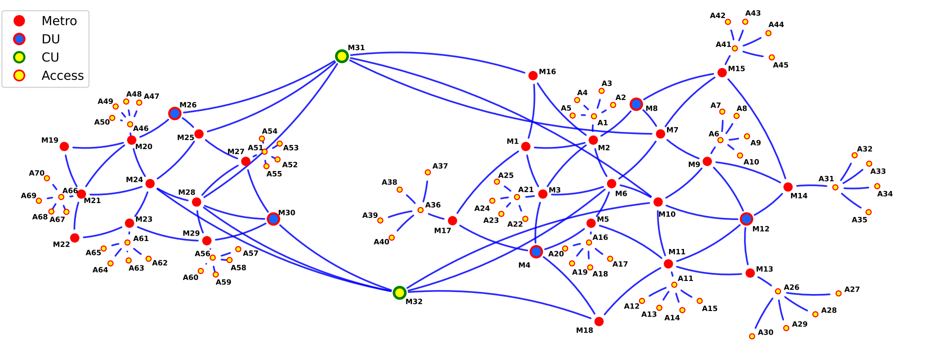

TABLE I: Topology parameters.

Parameter	Value
Optical nodes	102
Metro/ROADM nodes	30
CU/DC nodes	2
Access clusters	14
RU-serving leaves	56
Total fiber links	133
Metro-core links	51
Metro–DC links	12
Access–metro links	14
Splitter–leaf links	56
Avg. metro span	7.25 km
Max. metro span	12 km
Access–metro avg.	1.93 km
Leaf-link avg.	1.71 km
Metro–DC span	20 km to 40 km

Fig. 3: Converged metro-access topology along with parameters considered for RAN *x*-haul provisioning

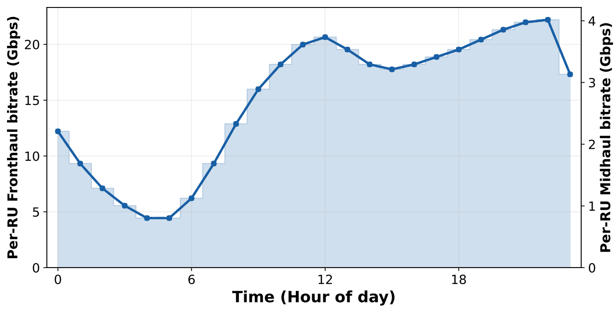

Fig. 4: Hourly RAN traffic profile and corresponding per-RU FH/MH bitrates.

throughout the day. The complete evaluation procedure is summarized in Algorithm 1. First, feasible RU–DU and DU–CU candidates using latency and GSNR constraints are pre-computed. Then, for each hour, FH and MH rates are scaled according to $L(t)$ and requests are admitted through RWSA. A request first attempts its home DU, then another active reachable DU, and finally a hibernated reachable DU in the dynamic scheme. If no candidate satisfies capacity, spectrum, latency, and QoT constraints, the request is blocked. A request r is accepted only if route availability, latency, QoT,

Latency is evaluated separately for FH and MH. For FH, split 7.2/7.2x is considered with a latency budget of 250 µs [8]. The FH latency is modeled as

$$T_{\mathrm{FH}} = T_{\mathrm{RUproc}} + 2T_{\mathrm{TRx}} + T_{\mathrm{prop,FH}}, \tag{2}$$

where $T_{\mathrm{RUproc}} = 50\,\mu\mathrm{s}$ is the RU processing delay [37], $T_{\mathrm{TRx}} = 30.9\,\mu\mathrm{s}$ is the DSCM TRx latency [38], and $T_{\mathrm{prop,FH}} = d_{\mathrm{FH}} \cdot 5\,\mu\mathrm{s\,km}^{-1}$ is the propagation delay over SSMF. For MH, split 2 is considered with a latency budget of 1 ms [8, 29]. The MH latency includes DU processing, TRx latency, and propagation delay, and is modeled as

$$T_{\mathrm{MH}} = T_{\mathrm{DUproc}} + 2T_{\mathrm{TRx}} + T_{\mathrm{prop,MH}}, \tag{3}$$

where $T_{\mathrm{DUproc}} = 150\,\mu\mathrm{s}$ represents the DU processing delay at the DU between CU-UP and DU-High through the F1-U interface [37]. A request is latency-feasible only if both the selected FH and MH routes satisfy their corresponding latency budgets.

C. DU Operation and Request Admission

The energy model accounts for both RU and DU power consumption. Each active RU is assumed to consume $P_{\mathrm{RU}} = 150$ W for outdoor operation, while the DU power depends on the number of distinct RUs served by the DU [39, 40] and is modeled as

$$P_d^{\mathrm{DU}}(n) = P_0 + \delta(n-1), \tag{4}$$

where n is the number of RUs served by DU d, $P_0 = 180$ W is the baseline active-DU power, and $\delta = 20$ W is the incremental power per additional served RU. In the dynamic scheme, hibernated DUs consume zero power, whereas in the static scheme all DUs in the fixed set remain active

Algorithm 1: Diurnal DU steering with integrated reactive overflow and RWSA.

Input : Topology $\mathcal{G} = (\mathcal{N}, \mathcal{L})$; access clusters $\mathcal{C}$; candidate DU set $\mathcal{D}$; CU set $\mathcal{U}$; diurnal profile $L(t)$, $t \in \{0, \ldots, 23\}$; peak rates $R^{\mathrm{FH}}_{\max}, R^{\mathrm{MH}}_{\max}$; DU capacity Γ; QoS limits; blocking threshold $\beta^\star = 0.01$; residual margin $\rho = 0.15$.
Output: Per-hour energy, carried traffic, and blocking for the Dynamic and Static schemes.

```
// Phase I -- physical-layer reachability
```
$\mathcal{R}^{\mathrm{FH}} \leftarrow$ feasible (c, d) pairs, $c \in \mathcal{C}, d \in \mathcal{D}$, where every RU in cluster c reaches d within latency and GSNR $\geq$ GSNR$_{\min}$
$\mathcal{R}^{\mathrm{MH}} \leftarrow$ feasible (d, u) pairs, $d \in \mathcal{D}, u \in \mathcal{U}$, reaching a CU under the same limits
$\forall$ RU: $\Phi(\mathrm{RU}) \leftarrow$ reach-&-MH-eligible DUs, ordered by best GSNR
```
// Phase II -- per-hour evaluation
```
for $t \leftarrow 0$ **to** 23 **do**
  $r^{\mathrm{FH}} \leftarrow L(t) R^{\mathrm{FH}}_{\max}$; $r^{\mathrm{MH}} \leftarrow L(t) R^{\mathrm{MH}}_{\max}$
  if *scheme is* Dynamic **then**
    $\mathcal{A} \leftarrow$ `Steer`$(L(t))$ `// minimum reachable active DU set`
    homing $\pi : \mathcal{C} \to \mathcal{A}$ from `Steer()`
  else // Static
    $\mathcal{A} \leftarrow \mathcal{A}^\star$ `// fixed peak DU set`
    $\pi \leftarrow$ `Rehome`$(\mathcal{C}, \mathcal{A}^\star)$ `// cluster homing over fixed set`
  reset spectrum, DU loads $\ell_d \leftarrow 0$, served-RU sets $S_d \leftarrow \varnothing$
  $A \leftarrow 0$; $B \leftarrow 0$; $\beta \leftarrow 0$
  repeat
    draw request $q = (\mathrm{RU}, c)$ with home DU $d_0 = \pi(c)$
    res $\leftarrow$ `ReactiveOverflow`$(q, d_0, \mathcal{A}, \Phi, \ell, \rho)$
    update A, B, DU loads, spectrum state, and block cause from res
    $\beta \leftarrow B/(A + B)$ `// A, B: served and blocked counts`
  until *target BP is reached*
  $P_t^{\mathrm{DU}} \leftarrow \sum_{d \in \mathcal{A}} \phi(|S_d|)$, with $\phi(n) = P_0 + \delta(n-1)$
  $P_t^{\mathrm{RU}} \leftarrow |\bigcup_d S_d| \cdot P_{\mathrm{RU}}$
  $T_t \leftarrow A\left(r^{\mathrm{FH}} + r^{\mathrm{MH}}\right)$
return per-hour $\{P_t^{\mathrm{DU}}, P_t^{\mathrm{RU}}, T_t, \beta_t\}$ for each scheme

and spectrum constraints are jointly satisfied:

$$\begin{aligned} &N_{\mathrm{FH}}(r) \geq 1, \quad N_{\mathrm{MH}}(r) \geq 1, \\ &T_{\mathrm{FH}}(r) \leq T_{\mathrm{FH}}^{\max}, \quad T_{\mathrm{MH}}(r) \leq T_{\mathrm{MH}}^{\max}, \\ &\mathrm{BER}_p(r) \leq \mathrm{BER}_{\mathrm{th}}, \quad p \in \{\mathcal{P}_{\mathrm{FH}}(r), \mathcal{P}_{\mathrm{MH}}(r)\}, \\ &S_\ell^{\mathrm{res}} \geq S_r^{\mathrm{req}}, \quad \forall \ell \in \mathcal{P}_{\mathrm{FH}}(r) \cup \mathcal{P}_{\mathrm{MH}}(r). \end{aligned} \tag{5}$$

where $N_{\mathrm{FH}}(r)$ and $N_{\mathrm{MH}}(r)$ are the numbers of available FH and MH candidate routes for request r, $T_{\mathrm{FH}}^{\max}$ and $T_{\mathrm{MH}}^{\max}$ are the corresponding latency budgets, $\mathcal{P}_{\mathrm{FH}}(r)$ and $\mathcal{P}_{\mathrm{MH}}(r)$ are the selected FH and MH lightpaths, and $\mathrm{BER}_{\mathrm{th}}$ is the pre-FEC BER threshold. S_ℓ^{res} denotes the number of available DSCM subcarrier slots on link ℓ, while S_r^{req} is the number of subcarrier slots required by request r. The procedure is repeated over all 24 hours for both schemes, recording carried traffic, energy consumption, blocking, and link congestion.

IV. Results and Discussion

This section compares the proposed dynamic DU activation strategy (S1) with the static baseline (S2), where all candidate DUs remain active throughout the day. Both schemes are evaluated under the same 24-hour traffic profile and PHY-layer-aware RWSA constraints, with a target BP of $\beta^\star = 1\%$.

Fig. 5 shows that S1 adapts the active DU set to the hourly traffic load. During off-peak periods, only three DUs are active, namely M8, M12, and M30. During high-load periods, the active set expands to five DUs, namely M4, M8, M12, M26, and M30. The peak DU set is used for only 5 out of 24 hours, showing that the dynamic strategy avoids unnecessary DU activation while preserving additional path diversity near peak demand.

Fig. 6 compares the hourly carried end-to-end traffic achieved by S1 and S2. As expected, S2 carries more traffic because all five DUs remain active, increasing path diversity and the number of feasible RU–DU–CU assignments. Over the full day, S1 carries 885.2 Tbps·h, whereas S2 carries 1239.2 Tbps·h, corresponding to 40.0% higher carried traffic for the static baseline. At the peak hour, $h = 22$ with $L = 1.00$, S2 carries 23.1% more traffic than S1. This confirms the main trade-off: dynamic DU hibernation improves energy efficiency, but the reduced number of active aggregation points can limit carried traffic under spectrum-constrained operation. The blocking analysis shows that S1 operates close to the target BP of approximately 1%. All observed blocking events are caused by spectrum exhaustion, with 495 spectrum-related blocks and no DU-capacity or QoT-related blocking events. This indicates that, for the considered topology and residual C-band allocation, DSCM subcarrier availability on shared optical links is the binding constraint.

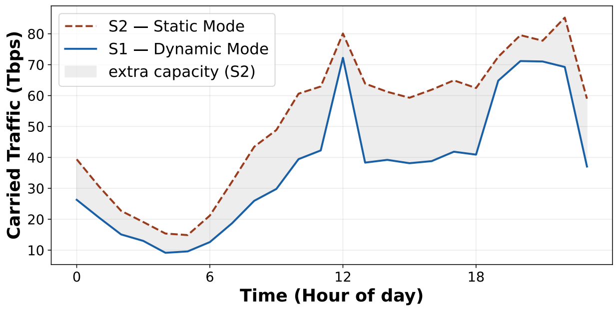

Fig. 6: Hourly carried end-to-end traffic for dynamic and static DU operation.

Fig. 7 reports the DU power per delivered end-to-end bit. The dynamic scheme provides a clear efficiency advantage because fewer DUs are kept active during low- and medium-load periods. In terms of daily energy consumption, S1 requires 40.0 kWh for DUs, whereas S2 requires 46.1 kWh, resulting in a 13.2% DU-energy saving. When RU energy is also included, the total daily energy decreases from 247.7 kWh in S2 to 241.6 kWh in S1, corresponding to a 2.5% total-energy saving. The lower total-energy saving, compared with the DU-only saving, is expected because RU energy remains largely tied to the served access traffic and is not reduced in the same proportion as DU activation. To compare both schemes under the same delivered-load condition, a fair efficiency metric is also considered. Under this normalized comparison, S1 is 15.4% more power-efficient than S2 in terms of DU power per delivered bit. This result shows that the dynamic strategy is more energy efficient when the objective is to serve a given amount of traffic with lower DU power consumption. The static scheme remains preferable when the objective is to maximize carried traffic under a fixed blocking threshold, whereas the dynamic scheme is preferable when energy proportionality is prioritized.

Overall, the results demonstrate that traffic-aware DU activation can reduce DU energy consumption without violating the target BP. The dynamic strategy saves 13.2% of DU energy and improves fair DU power efficiency by 15.4%, while the static strategy carries 40.0% more daily traffic due to the continuous availability of all DU locations. The fact that blocking is dominated by spectrum exhaustion suggests that further improvements should focus on spectrum-aware

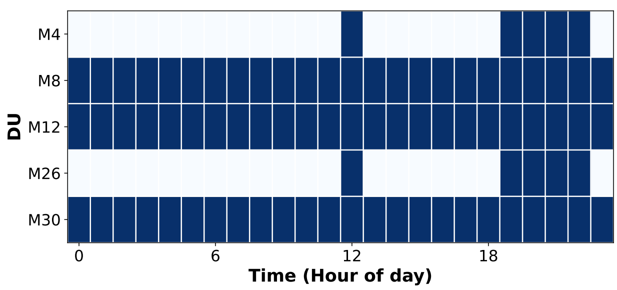

Fig. 5: DU activation pattern under dynamic operation over 24 hours.

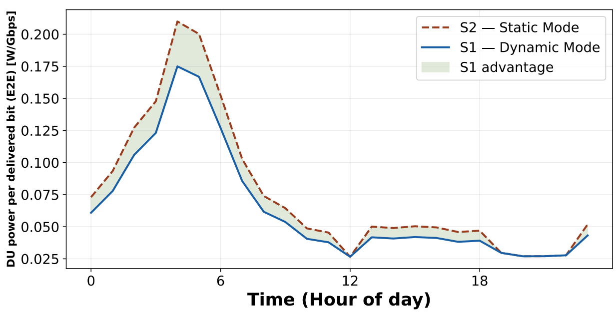

Fig. 7: DU power per delivered end-to-end bit over 24 hours.

DU steering, subcarrier grooming, and congestion-aware assignment of RU traffic to active DUs. These findings confirm that energy-aware RAN x-haul operation must be jointly designed with optical resource allocation, since DU hibernation, path diversity, and residual spectrum availability are tightly coupled in converged metro-access ONaaS infrastructures.

V. CONCLUSION

This paper evaluated energy-aware RAN x-haul provisioning over an ONaaS-enabled converged metro-access optical network using DSCM-based coherent transmission. A realistic 24-hour traffic profile was mapped onto FH and MH bitrate requirements, and each RU–DU–CU request was admitted through PHY-layer-aware RWSA considering latency, GSNR, BER, DU capacity, and residual C-band spectrum constraints. The results show that dynamic DU activation reduces the active DU set from five DUs at peak load to three DUs during off-peak periods, saving 13.2% of DU energy and improving DU power efficiency by 15.4% at equal delivered load. When RU energy is included, the total daily energy saving is 2.5%. In contrast, the static baseline carries 40.0% more daily traffic and 23.1% more traffic at the peak hour due to higher path diversity. Blocking remains close to the 1% target and is dominated by spectrum exhaustion rather than DU-capacity or QoT violations. Therefore, residual optical spectrum is the main limiting factor in the considered ONaaS scenario. Future energy-aware RAN x-haul control should jointly optimize DU activation, optical routing, and DSCM subcarrier allocation to balance energy consumption, BP, and carried traffic.

ACKNOWLEDGMENT

Ahtisham Ali's and Ricardo Pousa's work is fully funded by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie project EWOC under grant agreement No. 101073265. This work has also received support from the EU Horizon Project ALLEGRO (101092766).

REFERENCES

[1] A. Ghosh *et al.*, "5G evolution: A view on 5G cellular technology beyond 3GPP release 15," *IEEE access*, vol. 7, pp. 127 639–127 651, 2019.

[2] International Telecommunication Union (ITU), "IMT vision - framework and overall objectives of the future development of IMT for 2020 and beyond," Tech. Rep. M.2083-0, 2015, [Online]. Available: https://www.itu.int/rec/R-REC-M.2083.

[3] A. Tzanakaki *et al.*, "Optical Networking in Support of 5G and 6G Infrastructures," in *Proc. Int. Conf. Opt. Netw. Design Model. (ONDM)*, 2025, pp. 1–3.

[4] Z. Vujicic *et al.*, "Toward Virtualized Optical-Wireless Heterogeneous Networks," *IEEE Access*, vol. 12, pp. 87 776–87 806, 2024.

[5] 3GPP, "NG-RAN; Architecture Description (Release 17)," 3GPP, Tech. Rep. TR 38.401 V17.4.0 (2023-03), 2023.

[6] J.-i. Kani *et al.*, "Disaggregation and virtualization for future access and metro networks [Invited Tutorial]," *J. Opt. Commun. Netw.*, vol. 17, no. 1, pp. A1–A12, 2024.

[7] O.-R. Alliance, "O-RAN WhitePaper—Building the Next Generation RAN," O-RAN Alliance, Alfter, Germany, Tech. Rep., October 2018.

[8] 3GPP, "Study on CU-DU lower layer split for NR; (Release 15)," 3GPP, Tech. Rep. TR 38.816 V15.0.0 (2017-12), 2017.

[9] L. M. P. Larsen *et al.*, "A Survey of the Functional Splits Proposed for 5G Mobile Crosshaul Networks," *IEEE Commun. Surv. Tutor.*, vol. 21, no. 1, pp. 146–172, 2019.

[10] A. Ali *et al.*, "Statistical Analysis of end-to-end Route Feasibility in Converged Metro–Access Optical Networks," *IEEE Photon. Technol. Lett.*, pp. 1–1, 2026.

[11] S. Shen *et al.*, "Lambda as a service based on dwdm tunable laser in metro access network," in *Proc. Opto-Electron. Commun. Conf. (OECC)*. IEEE, 2023, pp. 1–4.

[12] V. Curri, "GNPy model of the physical layer for open and disaggregated optical networking [Invited]," *J. Opt. Commun. Netw.*, vol. 14, no. 6, pp. C92–C104, 2022.

[13] D. Welch *et al.*, "Digital subcarrier multiplexing: Enabling software-configurable optical networks," *J. Lightwave Technol.*, vol. 41, no. 4, pp. 1175–1191, 2023.

[14] J. Pedro *et al.*, "Extended network applications of coherent pluggable transceivers," *J. Opt. Commun. Netw.*, vol. 17, no. 2, pp. A210–A223, 2025.

[15] H. Hojeij *et al.*, "On Flexible Placement of O-CU and O-DU Functionalities in Open-RAN Architecture," *IEEE Trans. Netw. Serv. Manag.*, vol. 22, no. 1, pp. 660–674, 2025.

[16] N. Sen and A. F. A, "Towards Energy Efficient Functional Split and Baseband Function Placement for 5G RAN," in *Proc. IEEE NetSoft*, 2023, pp. 237–241.

[17] A. Ali *et al.*, "Optimized Distributed Unit (DU) Placement under Physical Constraints for 6G x-Haul," in *Proc. Int. Conf. Transparent Opt. Netw. (ICTON)*, 2026, pp. 1–4.

[18] A. Ali *et al.*, "Evaluating Metro-Access Optical Networks as-a-Service for Flexible RAN X-haul," in *Proc. IEEE Conf. Standards Commun. Netw. (CSCN)*, 2025, pp. 1–6.

[19] F. Xu *et al.*, "Understanding mobile traffic patterns of large scale cellular towers in urban environment," *IEEE/ACM Trans. on Netw.*, vol. 25, no. 2, pp. 1147–1161, 2017.

[20] Y.-J. Cho *et al.*, "Practical load balancing algorithm for 5g small cell networks based on real-world 5g traffic and o-ran architecture," *IEEE Access*, vol. 12, pp. 121 947–121 957, 2024.

[21] "3rd Generation Partnership Project (3GPP)," https://www.3gpp.org/.

[22] CPRI Cooperation, "CPRI Specification V6.1," Common Public Radio Interface, Tech. Rep., 2014, [Online]. Available: https://www.cpri.info/spec.html.

[23] "Next Generation Mobile Networks (NGMN) Alliance," https://www.ngmn.org/.

[24] "Open-RAN (O-RAN) Alliance," https://www.o-ran.org/.

[25] "Small Cell Forum (SCF)," https://www.smallcellforum.org/.

[26] "IEEE NGFI (1914) Working Group," https://sagroups.ieee.org/1914/.

[27] "Telecom Infra Project (TIP)," [Online]. Available: https://telecominfraproject.com.

[28] O-RAN Alliance, "O-RAN Architecture Description," O-RAN Alliance, Germany, Document O-RAN.WG1.O-RAN-Architecture-Description-v07.00, Oct. 2022.

[29] R. J. B. Pousa *et al.*, "Road to Dynamic Functional Split in Radio Access Networks," in *Proc. IEEE Medit. Electrotech. Conf. (MELECON)*, 2024, pp. 844–849.

[30] 3GPP, "Radio access architecture and interfaces (Release 14)," 3GPP, Tech. Rep. TR 38.801 V14.0.0 (2017-03), 2017.

[31] FLEX-SCALE Consortium, "Deliverable D2.1 – 6G network requirements," Tech. Rep., 2023. [Online]. Available: https://6g-flexscale.eu/en/outcomes

[32] ITU-T, "Spectral grids for WDM applications: DWDM frequency grid," International Telecommunication Union, Tech. Rep. G.694.1, 2012.

[33] D. Welch *et al.*, "Point-to-multipoint optical networks using coherent digital subcarriers," *J. Lightwave Technol.*, vol. 39, no. 16, pp. 5232–5247, 2021.

[34] T. Mano *et al.*, "Measuring the Transceiver's Back-to-Back BER-OSNR Characteristic Using Only a Variable Optical Attenuator," in *Proc. Eur. Conf. Opt. Commun. (ECOC)*, 2024, pp. 1499–1502.

[35] R. Schips *et al.*, "Model Validation of DSCM Transceivers from B2B Characterization for Disaggregated Optical Networks," *IEEE Photon. Technol. Lett.*, pp. 1–1, 2026.

[36] R. Pennetta *et al.*, "A low-loss and broadband all-fiber acousto-optic circulator," in *Frontiers in Optics*. Optica Publishing Group, 2024, pp. FW7D–4.

[37] O-RAN Working Group 4 (Open Fronthaul Interfaces WG), "Control, User and Synchronization Plane Specification," O-RAN Alliance, Alfter, Germany, Document O-RAN.WG4.CUS-0-v10.00, Oct. 2022.

[38] G. Simon *et al.*, "200 Gb/s Coherent Point-to-Multipoint Coexistence With 50G-PON for Next-Generation Optical Access," *IEEE Photon. Technol. Lett.*, vol. 36, no. 10, pp. 665–668, 2024.

[39] O-RAN Working Group 7 (White-box Hardware Workgroup), "O-RAN Hardware Reference Design Specification for Outdoor Macrocell with Split Architecture Option 7.2 4.0," O-RAN Alliance, Technical Specification O-RAN.WG7.OMAC-HRD.0-R004-v04.00, Oct. 2025.

[40] Sunwave, "CrossFire M2RU Mid Power: 4T4R Digital Radios, 37 dBm Output Power, Support 5G NR, Outdoor Rated," Sunwave, Product Specification SP_CF_M2RU v1.0.2, Aug. 2021. [Online]. Available: https://www.nexedi.com/RS-Sunware.M2RU.Crossfire.spec/Base_download