Autonomous Control of Multi-Granular Optical Networks with Pluggable Transceivers for Adaptive and High-Capacity 6G X-Haul

Raül Muñoz*, W. Akbar*, Ll. Gifre-Renom*, V. Lohani *, A. Larrañaga-Zumeta *, A. Sgambelluri†, M. Radovic †, N. Sambo†, R. Vilalta*, S. Lagen*, J.M. Fàbrega*, P. Armingol‡, O. Gonzalez de Dios‡, J. P. Fernandez-Palacios‡, N. Psaromanolakis §, G. Katsikas §, R. Martínez*, R. Casellas*, and I. Tomkos¶
*Centre Tecnològic de Telecomunicacions de Catalunya (CTTC-CERCA), Castelldefells, Spain. Email: raul.munoz@cttc.es
†Scuola Superiore Sant'Anna, Pisa, Italy
‡ Telefónica Innovación Digital, Madrid, Spain
§ UBITECH, Athens, Greece
¶ University of Patras, Patras, Greece

***Abstract*—This paper presents recent advances in multi-granular optical networking and next-generation optical pluggable transceivers that enable adaptive, high-capacity 6G X-haul networks. Building on results from the FLEX-SCALE and PROTEUS-6G projects, we introduce packet-optical architectures that dynamically exploit digital subcarriers, wavelengths, wavebands, and spatial domains to provide terabit-scale transport with fine-grained switching and scalable resource allocation across access, metro, and core segments. These technologies allow transport networks to rapidly adapt to changes in RAN functional splits, offering bandwidth elasticity and support for latency-critical 6G services. We demonstrate how programmable pluggables, multi-granular switching, and intelligent orchestration enable real-time adaptation of fronthaul, midhaul, and backhaul connectivity. Finally, we outline the integration of these capabilities into the cloud-native ETSI TeraFlowSDN (TFS) controller, which incorporates advanced telemetry, analytics, closed-loop automation, multilayer orchestration, and optical-pluggable control for end-to-end management of 6G transport networks.**

***Index Terms*—6G, optical network, MG-ON, orchestration, pluggable transceivers**

I. Introduction

In the coming decade, society, industry, and the global economy are expected to undergo a profound digital transformation driven by advancements in the Sixth Generation (6G) of mobile networks. As emerging applications such as immersive extended reality (XR), large-scale digital twins, pervasive artificial intelligence (AI), autonomous mobility, and massive Internet of Things (IoT) ecosystems place unprecedented demands on connectivity, future infrastructures will need to deliver extreme capacity, sub-millisecond latency, and deterministic end-to-end performance [1]. Meeting these requirements will push beyond the limits of traditional transport solutions and motivate a shift toward highly programmable, AI-driven optical communication systems capable of adapting in real-time to dynamic service conditions. To sustain traffic volumes anticipated to reach multi-terabit-per-second scales per cell site, 6G transport networks must adopt advanced optical architectures that exploit multiple granularities (wavelengths, wavebands, and spatial domains) to overcome spectral saturation and support massive parallelism [2]. These innovations reinforce the tight coupling between wireless and optical domains, making the transport network a critical enabler of 6G by providing the agility, scalability, and efficiency required across fronthaul, midhaul, and backhaul segments [3]

A major driver of this increasing complexity is the growing flexibility of the Radio Access Network (RAN), enabled by the adoption of functional splits that distribute processing tasks between Centralized Units (CUs) and Distributed Units (DUs) according to service requirements and network conditions. By selecting among standardized split options, from lower-layer splits that impose stringent timing and high-capacity fronthaul demands to higher-layer splits that relax transport constraints but reduce coordination gains, operators can dynamically balance performance, cost, and scalability. This adaptability is essential in 6G environments where traffic patterns, mobility behaviors, and service profiles evolve rapidly. However, the choice of functional split has a direct impact on the transport infrastructure, affecting bandwidth and latency [4]. This creates a need for optical systems that can reconfigure capacity with the same temporal granularity as RAN behavior. Digital subcarrier technologies play a central role by enabling elastic, sliceable optical channels whose bandwidth and modulation formats can be rapidly adjusted to follow real-time changes in Flexible Functional Split (FFS) configurations. Through fine-grained spectrum allocation, fast retuning, and per-subcarrier quality adaptation, these technologies allow fronthaul, midhaul, and backhaul connections to expand, contract, or reorganize on demand.

As a result, transport infrastructures must evolve toward highly elastic, multi-granular, and autonomously controlled architectures, leveraging digital subcarriers, wave-

This work was partially supported by: the European Union's Horizon Europe Program under grant FLEX-SCALE (101096909) and PROTEUS-6G (101139134); SHINE (PID2024-159781OB-I00) project funded by MICIU/AEI/10.13039/501100011033 and by ERDF/EU.

lengths, wavebands, and spatial domains, to support continuous re-optimization across the entire X-haul [5] [6]. This tight interplay between FFS-driven RAN dynamics and optical-layer behavior further reinforces the need for intelligent orchestration and automation to ensure that a flexible 6G RAN can operate efficiently and reliably at scale.

The PROTEUS-6G project [7] addresses these challenges by developing two key technologies: (i) a spatially diverse point-to-multi-point (SDPtMP) optical fronthaul distribution network that enables cost-efficient and scalable delivery of multi-Gb/s fronthaul streams, and (ii) next-generation Digital Subcarrier Multiplexing (DSCM) transceivers that offer fine-grained, rapidly reconfigurable optical capacity to accommodate real-time fluctuations in functional split selections. Complementing this, the FLEX-SCALE project [8] introduces a multi-granular optical node (MG-ON) capable of exceeding 10 Pb/s by exploiting spectral and spatial switching to deliver unprecedented capacity scaling and resource flexibility for dense urban and regional 6G transport networks.

To operationalize the multi-granular elasticity and device programmability enabled by PROTEUS-6G and FLEX-SCALE, the cloud-native ETSI TeraFlowSDN (TFS) controller [9] acts as the autonomous transport-network intelligence layer. TFS integrates advanced telemetry, analytics, multi-layer optimization, and closed-loop automation, together with native support for DSCM pluggables, SDPtMP architectures, and MG-ON functionalities. This unified control and orchestration framework allows the X-haul transport network to adapt automatically and intelligently to evolving 6G service demands, ensuring end-to-end performance, scalability, and resilience across fronthaul, midhaul, and backhaul domains.

II. 6G X-HAUL TRANSPORT NETWORK REQUIREMENTS: BANDWIDTH AND LATENCY

The bandwidth requirements of 6G X-haul networks are fundamentally shaped by the expected escalation in radio interface parameters for both basic and advanced 6G RAN configurations. Increased carrier bandwidths that range from 400 MHz to 1 GHz per component carrier, extended use of multi-band aggregation with up to 16 component carriers for basic 6G and 48 for advanced 6G, higher order modulation formats, and an expansion of spatial degrees of freedom with up to 16 Multiple Input Multiple Output (MIMO) streams and 256 antenna ports collectively push transport capacities into the multi-hundred gigabit and multi-terabit regimes depending on the selected functional split [4]. For low layer splits such as options 7.2 and 7.3, capacity demands vary from 50 to 430 Gb/s for single band operation and scale linearly with the number of aggregated component carriers. Under multi-band configurations, these requirements readily exceed multiple terabits per second. Even more demanding are antenna port centric splits such as options 7.1 and 8, where in phase and quadrature (IQ) sample transmission yields fronthaul requirements between 1.25 and 12.6 Tb/s per sector in advanced 6G scenarios. These extreme rates highlight the necessity of multi-granular optical transport that includes digital subcarriers, waveband switching, and spatial parallelism, together with next generation pluggable transceivers and optical nodes capable of ≥ 10 Pb/s to ensure that X-haul capacity evolves in step with the capabilities of the 6G radio interface.

Latency requirements in 6G X-haul are equally critical and are largely dictated by interaction timings between protocol layers distributed across CUs, DUs, and RUs. High layer functional splits such as options 1 and 2 impose relatively relaxed constraints on the order of milliseconds. In contrast, mid layer and low layer functional splits introduce stringent latency budgets because of Medium Access Control (MAC) to Radio Link Control (RLC) coordination in split 4 and Hybrid Automatic Repeat Request (HARQ) timing in splits 6 to 8. For example, split 4 requires one-way transport latency below 100 µs for a 1 ms Transmission Time Interval (TTI), which tightens to 25 µs for higher subcarrier spacings such as 60 kHz. Even more restrictive are splits 6, 7.x, and 8, where centralized MAC and HARQ control impose end-to-end transport latency below 250 µs, shared between fronthaul and midhaul when both segments implement low layer splits. Emerging 6G features such as multi-connectivity, coordinated multi-Transmission Reception Point (TRP) operation and relay-based architectures further increase sensitivity to transport-induced latency since inter-RU coordination must occur within tight HARQ timing loops and under potentially variable multi-hop propagation conditions. Meeting these latency constraints requires not only optical paths with ultra-low latency and fast packet optical switching, but also real-time elastic bandwidth adaptation enabled by digital subcarriers and intelligent orchestration. This ensures that transport behavior can evolve in synchrony with dynamic functional split selections and time-critical 6G radio processes.

III. MULTI-BAND AND MULTI-FIBER PACKET OPTICAL NETWORK ARCHITECTURE COMBINING DSCM, WAVELENGTH DIVISION MULTIPLEXING (WDM), WAVEBAND DIVISION MULTIPLEXING (WBDM), AND SPACE DIVISION MULTIPLEXING (SDM)

A. *FLEX-SCALE Network Architecture*

The FLEX-SCALE project introduces a multigranular Multi Degree Reconfigurable Optical Add Drop Multiplexer (MD ROADM) architecture (Fig. 1) designed to scale optical transport networks beyond the C band and support Ultra Wideband (UWB) and space division multiplexed transmission for 6G [8]. As traffic increases across multiple spectral bands (S, C, L) and multiple spatial lanes per fiber, traditional Wavelength Selective Switch (WSS) based Reconfigurable Optical Add/Drop Multiplexer (ROADM) architectures face limitations in port count, scalability, and energy consumption. FLEX-SCALE addresses these constraints through a multi-granular optical switching node (MG ON) structured around a novel Waveband Selective Switch (WBSS), which enables flexible routing of wavebands carved from the full UWB spectrum. These wavebands act as spectral super channels that can be switched across ingress and egress fibers through

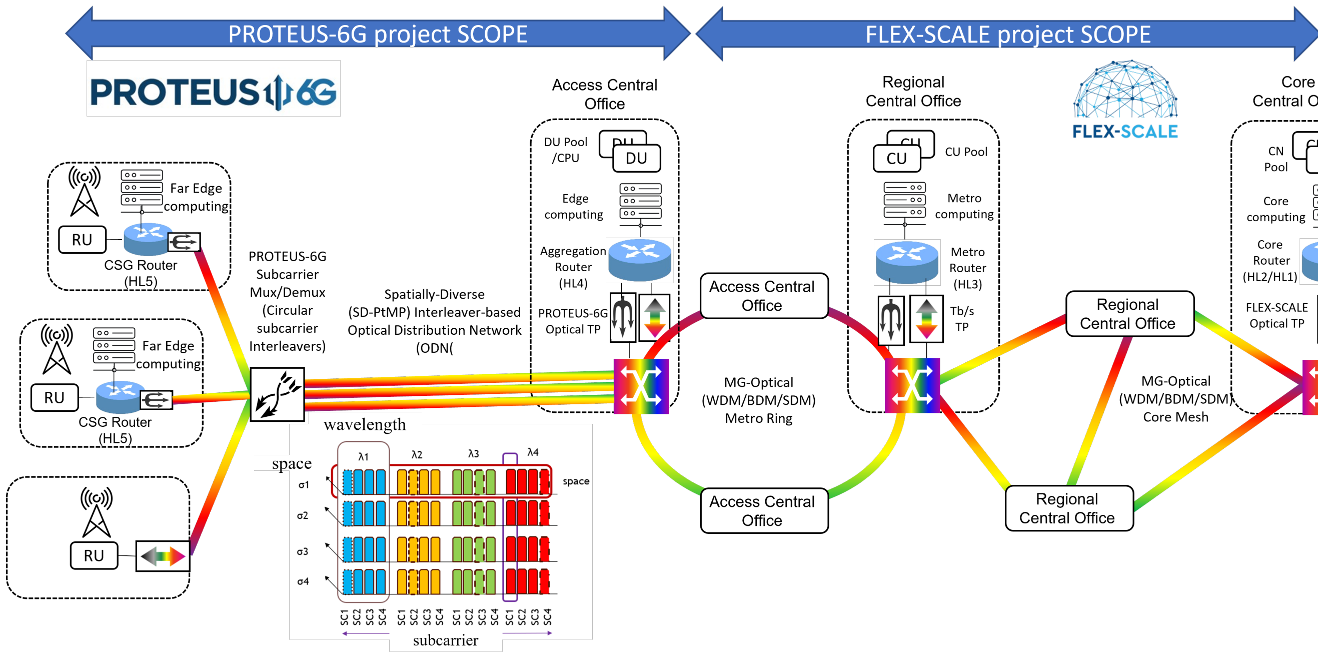

Fig. 1. Multi-granular (DSCM, WDM, WBDM, SDM) network architecture developed in PROTEUS 6G and FLEX-SCALE

a route and select topology capable of switching at fiber, waveband, and wavelength levels.

The proposed MG-ON integrates three switching layers. Layer 1, implemented by WBSS modules, performs routing at the fiber and waveband levels across multiple spatial lanes. Layer 2, implemented by waveband level inter port cross connects, enables flexible add drop and directional routing of entire wavebands across ROADM degrees. Layer 3, implemented through reconfigurable Wavelength Selective Switches interconnected by an intra band cross connect, provides wavelength level manipulation within each waveband. Together, these layers form a fully multigranular ROADM that supports fiber level pass through, waveband switching, and fine wavelength level routing. This architecture is conceived as a backbone building block offering unprecedented scalability, spectral efficiency, and operational flexibility for high capacity 6G transport.

B. PROTEUS-6G Network Architecture

The PROTEUS-6G project proposes a spatially diverse point to multipoint (PtMP) optical fronthaul and metro aggregation architecture that delivers beyond terabit capacities to 6G cell sites (Fig. 1). The architecture combines Space Division Multiplexing (SDM) and Wavelength Division Multiplexing (WDM) with Digital Subcarrier Multiplexing (DSCM) to enable efficient and scalable distribution of optical capacity across multiple Remote Units (RUs) and Distributed Units (DUs) [7]. By assigning groups of digital subcarriers to each RU and DU, multiple cell sites can share the same wavelength without interference, removing the need for power splitters typically used in passive optical network based fronthaul and reducing insertion loss while increasing capacity.

At the core of the PtMP node is a Subcarrier Group (SCG) interlacer, an optical element that performs subcarrier level routing. Each DSCM spectrum is divided into SCGs, which are then distributed among output ports so that multiple RUs can be served from the same wavelength. This enables deterministic spatial spectral routing and scalable PtMP connectivity in the access segment. When additional capacity is needed, Wavelength Selective Switches assign new wavelengths, ensuring seamless integration with metro ROADM rings for midhaul transport. The result is a unified fronthaul to midhaul optical distribution network that supports wavelength, subcarrier, and spatial routing, offering a scalable solution for high capacity 6G X-haul.

IV. Management of MG-Optical nodes (WDM/WBDM/SDM)

In 6G transport networks, the dynamic creation and reconfiguration of IP Virtual Network Topologies (VNTs) over multi-granular optical infrastructures is essential to support highly variable bandwidth demands and latency-sensitive services. Achieving end-to-end orchestration across the Radio Access Network (RAN), Core Network (CN), and Transport Network (TN) requires an integrated control framework capable of coordinating heterogeneous resources across domains. FLEX-SCALE adopts a cloud-native, microservice-based Transport Software Defined Networking (SDN) architecture, implemented using the TeraFlowSDN (TFS) platform [9], in which a packet SDN controller and an optical SDN controller operate under the supervision of an End-to-End (E2E) SDN orchestrator. The orchestrator ensures policy compliance and service-life-cycle consistency across layers, enabling the transport network to satisfy stringent 6G service requirements.

The underlying transport infrastructure is provided by the FLEX-SCALE multi-granular optical node (MG-ON), which integrates waveband-level and wavelength-level switching. Each MG-ON also includes a pool of multi-band transponders that expose Virtual Client Ports (VCPs) to the IP layer. This design enables flexible rearrangement of transport resources across multiple granularities and provides a scalable substrate for constructing high-capacity IP virtual topologies.

A central concept in controlling MG-ONs is the hierarchical connection model, where provisioning a connection at a lower optical layer automatically induces a virtual link at the next upper layer [10]. Provisioning a waveband channel induces a virtual WDM link between a pair of ROADMs; provisioning a wavelength channel (i.e., an optical channel) between a pair of transponders induces a virtual Optical Tributary Signal (OTSi) link; and provisioning a digital service connection between a pair of routers induces a virtual IP link. This abstraction decouples upper-layer control from the underlying optical complexity while allowing efficient reuse of dynamically-provisioned transport resources.

The workflow for establishing a new IP virtual link is as follows: The process begins with a request from the Operator to the E2E orchestrator to provision a digital service between two routers. The orchestrator validates the request and instructs the optical SDN controller to provision the required optical channel. As part of this workflow, the optical controller invokes the Routing, Fiber, Waveband, and Spectrum Assignment (RFWBSA) procedure to identify an appropriate optical path and determine the layer at which resources must be allocated. The E2E SDN orchestrator is responsible for End-to-End (E2E) network orchestration. RFWBSA determines whether suitable wavelength, waveband, or fiber resources exist and triggers the configuration of the corresponding MG-ON switching elements along the selected route.

If a waveband channel must be provisioned, this action automatically creates a virtual WDM link in the optical VNT, with the available spectrum corresponding to the allocated waveband. Once the wavelength-level optical channel is successfully established, the optical controller creates the virtual OTSi link with the required capacity. The orchestrator is then notified, and it instructs the packet SDN controller to instantiate the virtual IP link. The packet SDN controller activates the relevant VCPs on the routers and updates their routing tables. Subsequently, the new IP link is incorporated into the IP VNT, enabling direct, high-capacity router-to-router connectivity while reducing hop count and energy consumption. This multi-layer workflow provides the automation and flexibility required to support dynamic, multi-Tb/s services in emerging 6G backhaul transport networks.

A. Dynamic routing and resource allocation algorithms for MG-ON

Dynamic routing and resource allocation in MG-Optical Networks is inherently more complex than in conventional WDM systems because multiple resource domains (fiber, waveband, and wavelength) must be jointly managed. In the FLEX-SCALE architecture, routing must account for the hierarchical relationship between these layers, in which provisioning a connection at a lower layer automatically induces a virtual link at the next upper layer. To efficiently exploit this structure, MG-ONs rely on the RFWBSA algorithm, which determines the optimal path and resource allocation for each incoming digital service request.

RFWBSA operates over a two-layer optical graph model that reflects the hierarchical structure of MG-ONs [11]. At the *waveband layer*, the network is represented by a physical graph $G_b(V_b, E_b)$, where each pair of physically adjacent nodes is connected by one or more fibers offering multiple fixed wavebands derived from the UWB optical spectrum (S, C, and L bands). This graph captures all physical connectivity in the network and defines the domain in which waveband channels can be established. When a waveband is continuously available along an end-to-end route, RFWBSA provisions a *waveband channel* (WBCh) between the endpoints. Provisioning a waveband channel induces a virtual link in the overlying *wavelength-layer graph* $G_w(V_w, E_w)$. Initially, G_w contains no edges ($E_w = \varnothing$), but it grows dynamically as waveband channels are created, enriching the virtual topology with new wavelength-layer links. These induced virtual links enable subsequent attempts at wavelength-layer resource allocation and significantly improve spectral availability and network adaptability.

When a connection request arrives, RFWBSA first determines the number of frequency slot units (FSUs) required to support the requested bit rate. At the wavelength layer, resources follow the ITU-T G.694.1 flexible-grid specification, using a slot width granularity of 12.5 GHz. The algorithm attempts to assign a contiguous set of FSUs along a feasible route that satisfies *(i)* spectrum contiguity (adjacent FSUs on each link) and *(ii)* spectrum continuity (the same spectral interval across all hops) constraints . If such a route exists in G_w, the optical channel is provisioned, and the request is accepted.

If insufficient wavelength-layer resources are found, the algorithm escalates the request to the *waveband layer*. At this layer, RFWBSA searches for a waveband that is continuously available along the entire path. If such a waveband exists, a waveband channel is provisioned, and a corresponding virtual link is added to G_w between end nodes, enabling a new attempt at wavelength-layer spectrum assignment. If no feasible wavelength or waveband-level resources exist, the connection request is blocked. RFWBSA employs a first-fit spectrum assignment strategy, which provides computational simplicity and efficiency in dynamic traffic scenarios. For waveband assignment, priority is given first to the C-band, followed by the L-band, and finally the S-band.

V. Management of DSCM coherent pluggable transceivers

The introduction of DSCM coherent pluggable transceivers into disaggregated IP-over-DWDM (IPoDWDM) transport networks is a key enabler for highly programmable and bandwidth-elastic 6G X-haul [6]. These pluggables embed coherent optical transmission inside IP routers, replacing traditional transponders and simplifying the end-to-end transport architecture. At the same time, their tight integration of packet and optical functions introduces new control-plane requirements, as pluggable configuration must be consistent with optical-path computation and the operational state of the Optical Line System (OLS).

In the PROTEUS-6G scenario, DSCM pluggables support dynamic FFS adaptation by providing fine-grained, subcarrier-level bandwidth elasticity. Each pluggable generates multiple low-rate subcarriers using Nyquist frequency-division multiplexing (NFDM), which can be routed independently over the Spatially Diverse Point-to-Multipoint (SD-PtMP) optical distribution network. This allows the 6G X-haul transport capacity to scale according to changing RAN centralization requirements, making the pluggables a key interface between packet and optical domains.

Managing these capabilities requires defining how pluggable configuration responsibilities are divided among the IP controller, the optical controller, and the transport orchestrator. In current disaggregated architectures, three control plane models have been identified. In the first, *dual SBI management with full read/write access*, both the IP and optical controllers can directly configure the pluggables. This provides maximum flexibility, since each controller may tune parameters such as central frequency, power, or operational mode, but also introduces the risk of conflicting or out-of-sync updates. Strong coordination by the orchestrator is needed to avoid inconsistencies.

A second approach, *single SBI management with control exclusively in the IP domain*, restricts write access to the IP controller. The optical controller configures only the OLS, while pluggable parameters discovered by the IP controller are forwarded to the orchestrator, which incorporates them into OTSi service computation. This avoids configuration conflicts but increases the orchestrator's responsibility, as it must translate optical-layer planning decisions into pluggable configuration operations applied via the IP controller.

A third model, *dual SBI management with read-only access for the optical controller*, seeks a balanced compromise. The IP controller retains exclusive write access to the pluggables, ensuring configuration safety, while the optical controller obtains read-only visibility into their capabilities via NETCONF/OpenConfig. This allows the optical controller to compute OTSi media channels with full awareness of transceiver constraints but without modifying the device. After computing the optical service using Open Networking Foundation (ONF) Transport API (TAPI), the optical controller provides the orchestrator with the required pluggable settings, which are then applied through the IP controller.

Across all three strategies, accurate and consistent capability discovery is fundamental. Depending on the model, either the IP controller alone, or both controllers (with asymmetric permissions), retrieve the pluggable's supported transmission modes, Optical Signal-to-Noise Ratio (OSNR) requirements, power tolerances, and DSCM-specific parameters. These capabilities are essential inputs to OTSi computation and must be shared with the orchestrator to ensure alignment of pluggable configuration and optical-path planning.

VI. ETSI TeraFlowSDN (TFS) controller

A. Architecture Overview

ETSI TeraFlowSDN (TFS) is a cloud-native, microservice-based SDN controller for multi-layer, multi-technology transport network. Within FLEX-SCALE and PROTEUS-6G, TFS has been extended to support digital-subcarrier, multi-band and space-division multiplexed optical networks for real-time adaptation of 6G X-haul. The baseline architecture includes a *Context* state store, a plugin-based *SouthBound Interface (SBI)* for device drivers (NETCONF/gNMI/OpenConfig), *Service* handlers, *PathComp* for path computation, a *NorthBound Interface (NBI)* exposing IETF RESTCONF-based models, and a *WebUI*. Microservices interact through gRPC for direct calls and Apache Kafka for pub/sub eventing.

To support hierarchical control across IP and optical domains, TFS has been extended with: (i) an *End-to-End Orchestrator* for cross-domain service decomposition and topology composition; (ii) a *Virtual Network Topology Manager* (VNTM) to manage and recommend IP VNT updates; and (iii) an *extended Optical Controller* to provision both wavelength and waveband resources in multi-band MG-ONs. The NBI also incorporates an asynchronous event channel (Socket.IO) for near real-time broadcasting of recommendations and updates.

B. Multi-Layer Dynamic Adaptation of IP VNT over MG-Optical Networks

For FLEX-SCALE, TFS was extended to natively manage MG-ON nodes that support *waveband* and *wavelength* switching across S/C/L bands [13]. The Optical Controller integrates a *Flex-RSA (Flexible Routing and Spectrum Assignment) Engine* that (a) computes end-to-end routes, (b) checks for pre-existing wavebands, (c) creates/expands wavebands if needed, and (d) provisions the requested optical channels inside those wavebands. This "*waveband-first, wavelength-inside*" workflow reduces control-plane overhead and WSS port pressure while enabling scalable reconfigurations.

On the IP side, the VNTM emits *virtual link* recommendations (e.g., new or expanded links) toward the E2E Orchestrator, which maps router ports to optical endpoints and triggers optical service setup. After successful optical provisioning, the IP controller activates the new virtual link and updates the VNT. This hierarchical loop enables dynamic VNT adaptation driven by IP demand while minimizing optical reconfigurations.

A live demonstration [12] presented TFS-based multi-layer control that dynamically adapts IP virtual links over a multi-granular (waveband/wavelength) optical substrate spanning S/C/L bands.

C. Closed-Loop Automation

TFS implements a transport-native *observe–analyze–decide–act* loop to supervise optical services/devices and autonomously enforce corrective or optimizing actions at waveband/wavelength granularity, coordinated with IP VNT when policies require it. Functionally, the loop comprises:

- *Telemetry ingestion*: Real-time, model-driven streaming of device/service Key Performance Indicators (KPIs) (OpenConfig via NETCONF/gNMI), normalized and keyed to services and resources.
- *Analytics & correlation*: near–real-time detection of threshold breaches/trends and mapping symptoms to the smallest actionable scope (band, channel, endpoints).
- *Policy/decision*: triggers and guardrails (hysteresis, dwell times, concurrency limits) with action selection that prefers localized changes before global recomputation.
- *Actuation*: idempotent optical reconfiguration primitives (create/expand bands, (re)provision channels, retune endpoints) with pre/post checks and rollback; integrated with the waveband-first provisioning flow.
- *Multi-layer coordination*: optional emission of VNT recommendations to keep the IP view aligned after optical actions.

D. Management of Coherent Pluggables

TeraFlowSDN provides vendor-neutral, model-driven control of router-hosted coherent pluggables and integrates this control coherently with optical service computation and provisioning. Capabilities and state are handled through OpenConfig over NETCONF/gNMI, while optical paths and media channels are computed and provisioned using a RESTful Transport API; this ensures that decisions on central frequency, output power, operational mode, and spectral placement are applied consistently at device level and remain aligned with the optical line system view. TFS operates these functions through three cooperating controller instances: an IP (packet) controller that discovers capabilities and retains write authority on the pluggable; an optical controller that plans services with read-only visibility of pluggable status under a dual-SBI model; and an end-to-end (E2E) orchestrator that correlates multi-layer topology, sequences requests, and keeps state synchronized (including optional VNT updates)..

A live demo [14] showcased TFS managing router-hosted coherent pluggables, supporting digital sub-carrier multiplexing and point-to-multipoint operation, within an IP-over-DWDM control framework.

VII. Conclusions

This work has presented a unified vision for highly elastic, multi-granular 6G X-haul transport, integrating advances in wavelength-, waveband-, and spatial-domain switching and programmable DSCM pluggables with cloud-native multi-layer control. The FLEX-SCALE and PROTEUS-6G architectures demonstrate that combining MG-ON flexibility, pluggable-based capacity scaling, and the autonomous capabilities of ETSI TeraFlowSDN enables transport networks to dynamically align with rapidly changing RAN functional splits and extreme bandwidth and latency demands. Experimental demonstrations confirm that the proposed multi-layer workflows, closed-loop automation, and pluggable-aware orchestration provide the real-time adaptability, scalability, and operational safety required for future 6G services. Together, these results highlight a clear path toward fully programmable, AI-assisted optical transport infrastructures capable of supporting the end-to-end performance requirements of next-generation mobile networks.

References

[1] D. Uzunidis, *et al.*, "A Vision of 6th Generation of Fixed Networks (F6G): Challenges and Proposed Directions," *Telecom*, vol. 4(4), pp. 758-815, 2023. https://doi.org/10.3390/telecom4040035.

[2] C. Papapavlou, *et al.*, "Forthcoming optical x-haul infrastructure supporting 6G mobile network requirements," in Journal of Optical Communications and Networking, vol. 17, no. 11, pp. E82-E93, 2025.

[3] I. Tomkos, *et al.*,"The "X-Factor" of 6G Networks: Optical Transport Empowering 6G Innovations," *IEEE IT Professional*, vol. 26, pp. 32-39, 2024.

[4] A. Larrañaga, *et al.*, "Fronthaul/Midhaul Networks: Capacity and Latency Requirements Imposed by 6G Disaggregated RANs," in IEEE Communications Magazine, vol. 63, no. 5, pp. 86-93, 2025.

[5] R. Muñoz, *et al.*, "Control of packet over multi-granular optical networks combining wavelength, waveband and spatial switching for 6G transport," *in Proceedings of Optical Fiber Communications Conference and Exhibition (OFC)*, San Diego, California, USA, 2024.

[6] R. Muñoz, *et al.*, "Dynamic SDN Control of Pluggable Transceivers in disaggregated packet-optical networks for 6G X-haul", *In Proceedings of 25th Anniversary International Conference on Transparent Optical Networks (ICTON)*, Barcelona, Spain, 2025.

[7] C. Christofidis, *et al.*, "Programmable Reconfigurable Packet-Optical 6G Front-/Mid-Haul Infrastructure," *in Proceedings of 25th Anniversary International Conference on Transparent Optical Networks (ICTON)*, Barcelona, Spain, 2025.

[8] J.M. Fabrega, *et al.*, "FLEX-SCALE: A Disruptive Approach Toward Flexibly Scalable Energy Efficient Networking," *IEEE Communications Magazine*, vol. 63, no. 9, pp. 1-7, Sept. 2025.

[9] R. Vilalta et al., "TeraFlow: Secured Autonomic Traffic Management for a Tera of SDN flows," *2021 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, Porto, Portugal, 2021, pp. 377-382, doi: 10.1109/EuCNC/6GSummit51104.2021.9482469.

[10] V. Lohani, *et al.*, "Dynamic control, routing, and resource assignment in multi-granular optical node topologies combining wavelength, waveband, and spatial switching for 6G transport," *Journal of Optical Communications and Networking*, vol. 17, no. 1, pp. A59-A70, Dec. 2024.

[11] V. Lohani, *et al.*, "Dynamic Routing, Waveband, and Spectrum Assignment for Optical Superchannels in IP over Multi-Granular Optical Nodes," *Computer Networks*, vol. 9, p. 111731, Oct. 2025.

[12] Ll. Gifre, *et al.*, "Dynamic adaptation of IP virtual network topologies over multi-granular (wavelength and waveband) optical networks supported by ETSI TeraFlowSDN multi-layer control," *in Proceedings of Optical Fiber Conference (OFC)*, 2025.

[13] Ll. Gifre, *et al.*, "Multi-Layer Dynamic Adaptation of IP Virtual Network Topologies over Multi-Granular Optical Networks with ETSI TeraFlowSDN," *in Proceedings of International Conference on Transparent Optical Networks (ICTON)*, 2025.

[14] W. Akbar, *et al.*, "Management of Point-to-Multipoint Coherent Pluggable Transceivers to Provision IP Virtual Network Slice over DWDM Networks using ETSI TeraFlowSDN Multi-layer SDN Controller," *in Proceedings of European Conference on Optical Communication (ECOC)*, 2025.