# Ethernet in Core Networks: A Technical and Economical Analysis

Arno Schmid-Egger, *Member, IEEE*, Andreas Kirstädter, *Member, IEEE*

*Abstract*—**Ethernet is an uncompeted success story extending its reach from LAN and metro areas now also into core networks. 100 Gbps Ethernet will be the key enabler for a new generation of true end-to-end carrier grade Ethernet networks.**

**This paper first focuses on functionality and standards required to enable carrier-grade Ethernet-based core networks and possible Ethernet backbone network architectures will be discussed. The second part then evaluates the CAPEX and OPEX performance of Ethernet core networks and competitive network architectures. The results propose that Ethernet will not only soon be mature enough for deployment in backbone networks but also provide huge cost advantages to providers.**

***Index Terms*—Ethernet, core networks, next generation networks, OPEX, CAPEX.**

## I. INTRODUCTION

Backbone networks represent the top of the network hierarchy of carrier networks. They connect networks of different cities, regions, countries or continents, and in the majority of cases comprise SONET/SDH or Packet-over-SONET (PoS) technology. The complexity of these technologies imposes substantial financial burdens on network operators, both in the area of Capital Expenditures (CAPEX) and Operational Expenditures (OPEX).

The Ethernet protocol is a possible enabler of more cost-efficient backbone networks, as it is characterized by simplicity, flexibility, interoperability and low costs. While Ethernet is traditionally a Local Area Network (LAN) technology, continuous developments already enabled its deployment in Metropolitan Area Networks (MANs). Recent research and standardization efforts aim at speeding up Ethernet to 100Gbps, resolving scalability issues and supplying Ethernet with carrier-grade features. For this reason, Ethernet might in the near future become an attractive choice and serious competitor in the market of backbone networks.

The first part of this paper elaborates on requirements and possible architectures of carrier-grade Ethernet-based core networks. Current and future standards in the areas of Quality-of-Service, resilience, network management, and scalability that introduce carrier-grade features into Ethernet will be outlined. Different network architectures and related introduction strategies will also be explained briefly.

The second part of this paper examines the economics of Ethernet networks in comparison to SONET/SDH-based network architectures. Both CAPEX and OPEX are considered in hands-on business cases. The results propose that Ethernet backbone networks have a superior cost performance concerning both Capital and Operational Expenditures.

.

## II. OUTLINE OF CARRIER-GRADE 100G-ETHERNET BACKBONE NETWORKS

Figure II-1 shows two possible protocol stacks of future 100G-Ethernet backbone networks:

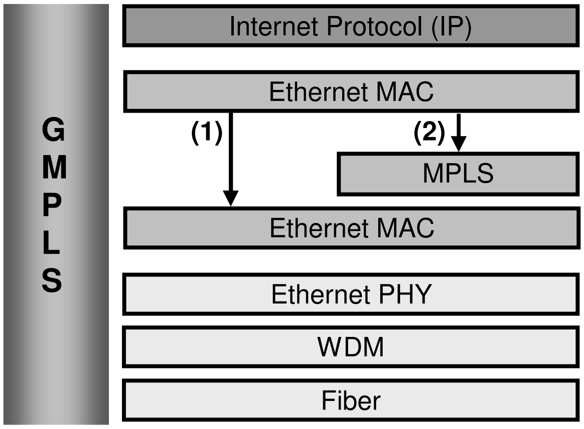

**Figure II-1. Protocol stacks of Ethernet backbone networks.**

Alternative (1) depicts core network transport via native Ethernet that uses MAC-in-MAC encapsulation [1]. Alternative (2) represents a MPLS-based backbone network that uses Ethernet for layer-2 transmission between the MPLS nodes[2]. Both architectures require a control plane to provide for network-wide functionality of signaling, traffic engineering, quality-of-service and protection, established e.g. via Generalized-MPLS (GMPLS) [3].

The mapping of protocol functionality to the requirements of backbone networks is shown in Figure II-2 for both native Ethernet (1) and MPLS Ethernet (2) backbone network

---

Manuscript received December 20, 2005. This work is the result of a diploma thesis conducted at Siemens Corporate Technology and the Institute for Communication Networks of the Technical University Munich.

Arno Schmid-Egger is with Siemens Power Transmission and Distribution, Germany. (e-mail: arno.schmid-egger@siemens.com).

Andreas Kirstädter is with Siemens Corporate Technology, Information and Communications, 81730 Munich, Germany. (e-mail: andreas.kirstaedter@siemens.com).

architectures. Functionality that is provided or controlled by GMPLS is also included in the illustration:

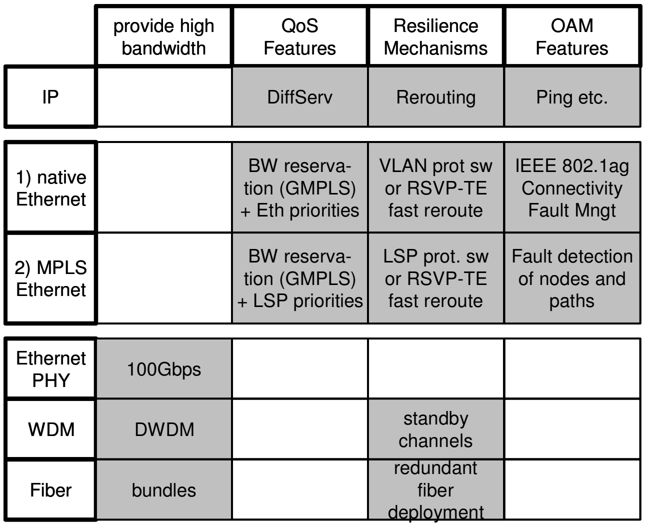

| | provide high bandwidth | QoS Features | Resilience Mechanisms | OAM Features |
|---|---|---|---|---|
| IP | | DiffServ | Rerouting | Ping etc. |
| 1) native Ethernet | | BW reservation (GMPLS) + Eth priorities | VLAN prot sw or RSVP-TE fast reroute | IEEE 802.1ag Connectivity Fault Mngt |
| 2) MPLS Ethernet | | BW reservation (GMPLS) + LSP priorities | LSP prot. sw or RSVP-TE fast reroute | Fault detection of nodes and paths |
| Ethernet PHY | 100Gbps | | | |
| WDM | DWDM | | standby channels | |
| Fiber | bundles | | redundant fiber deployment | |

**Figure II-2: Protocol layer functionality in Ethernet core networks**

In the following, the main mechanisms that enable carrier-grade Ethernet networks are shortly introduced:

### 1) *End-to-End Quality of Service*

Quality of Service (QoS) functionality enables service providers to guarantee and enforce transmission quality parameters (e.g. bandwidth, jitter, delay) according to a specified service-level agreement (SLA) with the customer.

A QoS framework that is currently developed by the Metro Ethernet Forum (MEF) aims at providing hard QoS in Ethernet networks [4]. This framework uses the RSVP-TE protocol to setup end-to-end paths with dedicated bandwidth. In native Ethernet networks, traffic is labeled with Service-VLAN tags that are related to a set of QoS parameters. QoS-conform forwarding in Ethernet switches is controlled by GMPLS [5]. In MPLS Ethernet networks, MPLS packets are labeled with MPLS tags and forwarded along the specific Label Switched Paths. A connection acceptance control, which is also operated by GMPLS, guarantees that the required bandwidth is available along the requested path.

The MEF's definition of generic service-level parameters enables a high flexibility in SLA definitions: The Committed Information Rate (CIR) determines the minimum amount of bandwidth available to the customer, while the Excess Information Rate (EIR) provides additional bandwidth during low network load periods. Maximum burst sizes corresponding to the CIR and EIR are defined accordingly.

For more details on enabling hard QoS in Ethernet networks, please consider [4], [6], and [7].

### 2) *Resilience mechanisms*

The proposed new restoration mechanism of Ethernet, the Rapid Spanning Tree Protocol [8], scales badly with increasing network dimensions, as its convergence time depends on the number of network nodes. For this reason, GMPLS will be used to manage protection of links and paths in carrier-grade Ethernet networks.

For both native Ethernet and MPLS Ethernet networks, GMPLS can pre-provision backup paths and switch over in the case of failure. In native Ethernet networks, multiple spanning trees are set up to accommodate different traffic flows. Every spanning tree corresponds to a certain service-VLAN tag. In the failure case, the connection is switched over to a different S-VLAN that uses a redundant path but connects the same set of nodes. In MPLS Ethernet networks, the GMPLS control plane redirects traffic of affected LSPs to backup LSPs.

Additionally, Link Aggregation Groups [8] may be set up in both scenarios to provide protection for individual links.

Restoration mechanisms like RSVP-TE Fast Reroute might also be used if the protection requirements are less severe. The GMPLS control plane might e.g. pre-establish backup paths for premium traffic, but envision slower restoration mechanisms for best-effort traffic.

The MEF protection framework [9] presents additional, generic resilience concepts that aim at enabling interoperation between future Ethernet devices of different vendors.

### 3) *Operations Administration & Maintenance*

Enhanced OAM functionality is indispensable in carrier networks, as it enables failure detection, localization and performance monitoring.

The IEEE 802.1ag standard extension, named Connectivity Fault Management, defines essential fault management functionality like loopback, continuity check and traceroute. Together, this set of functions will finally introduce path discovery and verification as well as fault detection and isolation into Ethernet. Furthermore, an OAM framework set up by the MEF focuses on providing SLA measurements (connectivity, latency, loss, jitter) for Ethernet networks [10].

### 4) *Core network scalability*

As Ethernet originates from the LAN, it faces several scalability issues on its progress into Wide Area Networks.

Address space limitations will be resolved with the upcoming IEEE Provider Backbone Bridge standard [1], which enables the provider to encapsulate customer Ethernet frames into a second Ethernet header ("MAC-in-MAC"), Figure II-3. Thus, the length of forwarding tables of provider equipment can be reduced tremendously and Ethernet services can be provided across the networks of different carriers.

**Figure II-3: Provider backbone bridge frame format [11]**

In MPLS Ethernet networks, address space is not critical, since MPLS labels are used for switching. These labels have a length of 20bit and enable up to a million of LSPs per link, as they are only locally significant.

Another requirement of native Ethernet backbone networks is a mechanism that enables utilization of meshed network structures. A successor to the spanning tree protocol (STP) that enables new topologies and permits loops would be the most desirable solution. Alternatively, a set of VLANs could be established across the backbone network in a way that all links are covered by one or more VLANs. However, this latter solution comes along with increased management complexity.

The maximum Ethernet frame size of 1500 Byte represents another bottleneck of Ethernet that generates a lot of processing overhead especially with increasing transmission speeds. Therefore, the standardization of “Jumbo frames” that contain 9000 Byte or more of user data is a desirable feature of a new Ethernet standard [12].

The transmission distance of Ethernet signals is another topic that has to be considered carefully upon the deployment of Ethernet in backbone networks, as node distances are usually in the area of a few hundred kilometer or even more. However, the maximum transmission distance of 10G Ethernet is to 70-80km according to vendor specifications, which is already way above the officially declared maximum range of 40km.

Upon speeding Ethernet up to 100Gbit/s, the transmission over long distances will become even more difficult: Second degree (slope) chromatic dispersion has to be exactly compensated, birefringeny effects become grave, and the signal-to-noise ratio of 100Gbit/s signals is generally lower as fewer photons are transmitted per optical impulse. Although several experiments and field trials of the past (e.g. that of Siemens, British Telecom, and the University of Eindhoven in the EU IST project FASHION with 160Gbit/s signals over 275km standard single-mode fiber [13]) showed that long-distance transmissions of high-speed signals are definitely feasible, the complexity and costs of the required optical equipment is at the moment still at a very high level.

On the other hand, the effort spent on extending the signal reach of Ethernet signals is rewarded by equipment savings. Figure II-4 illustrates the possible port count savings in an Ethernet core network where optical grooming can be applied up to the maximum transmission distance avoiding unneccessary electrical processing of transit traffic.

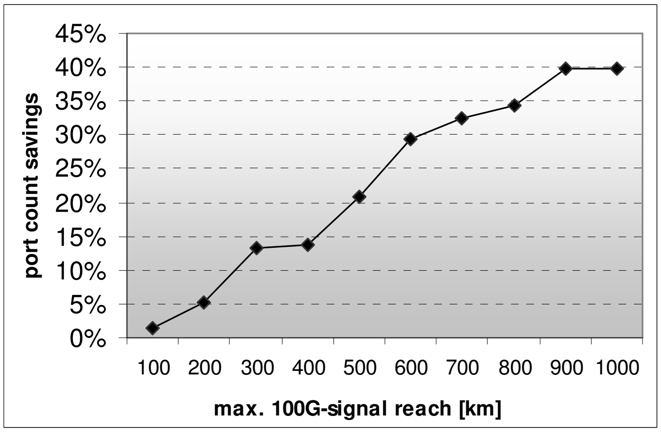

**Figure II-4: Port count savings in grooming-enabled Ethernet networks**

These results were derived for a native Ethernet network follwing the topology of a generic German backbone (Figure II-5) and are also used during the CAPEX and OPEX analyses below [14].

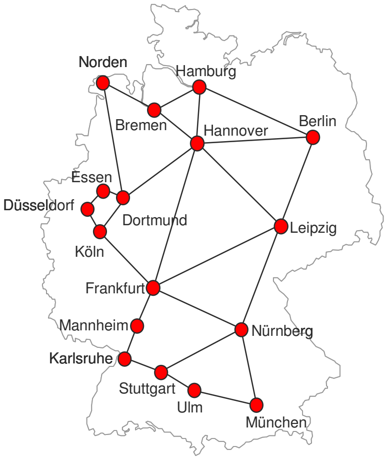

**Figure II-5: Reference network topology.**

The IEEE usually defines a broad set of physical Ethernet interfaces and besides a single 100G-Ethernet signal, a multiplexing of lower bit-rate optical signals into a 100G signal are also conceivable (e.g. 10x10G or 4x25G). Although the maximum transmission range of these multiplexed signals would certainly be longer than that of a pure 100G-signal the multiplexing requires additional WDM equipment and the signal would occupy several wavelength channels on the fiber.

## III. CAPEX COMPARISON OF DIFFERENT BACKBONE NETWORK ARCHITECTURES

In general, the term CAPEX incorporates all expenditures related to the purchase of equipment, infrastructure, buildings or furniture. As we want to compare the impact of different network architectures on CAPEX, we will neglect cost components that are not or barely affected by the choice of the network technology. Therefore, the CAPEX calculation only considers the costs for equipment that is related to the network architecture choice, including switches, routers, linecards and optical modules. However, costs of WDM-equipment and fiber are not included.

In order to calculate the total CAPEX of a specific network architecture, future traffic loads, network device counts and network device prices have to be estimated. The German reference network above was used for the physical topology of the considered backbone and corresponding traffic matrices were extrapolated to determine future link loads. The traffic is assumed to grow homogenously at a rate of 40% per annum, which corresponds to a doubling of traffic every two years as recently predicted by TeleGeography [15]. A shortest-path routing algorithm is applied to determine the single link loads.

From the link loads, the number of switches, routers, and linecard ports can be obtained for any different network architecture. In the SONET/SDH business cases, this device count includes equipment needed for upgrading the existing network in order to accommodate additional traffic (incremental CAPEX scenario). In the Ethernet business cases, the migration from SONET/SDH had to be respected with the consequence that the majority of components have to be acquired at the point of migration and only a few IP routers can be reused (greenfield / migration CAPEX scenario).

Future equipment prices are extrapolated following a careful analysis of market data, past price developments, and price relations between Ethernet and SONET/SDH. The prices of switches and routers in basic configuration (chassis, power supply, backplane, route processor, and switching fabric) are assumed to remain constant. For linecards that make up the largest cost component, a price extrapolation is indispensable, as they experience immense price reductions over time.

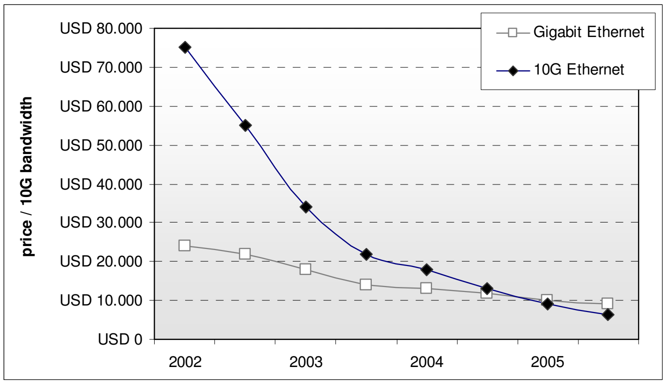

**Figure III-1: Past price development for Ethernet bandwidth (10G port capacity)**

Figure III-1 shows the past development of the prices for 10G Ethernet and Gigabit Ethernet bandwidth in the form of average prices for 10G port capacity (1 x 10GE port or 10 x GbE ports). Evaluated were press releases and official price lists ([16] – [24]) for high-end modular switches with long-reach (10km) optical modules of the following companies: Cisco, Enterasys, Extreme Networks, Force10, Foundry Networks, and Riverstone Networks.

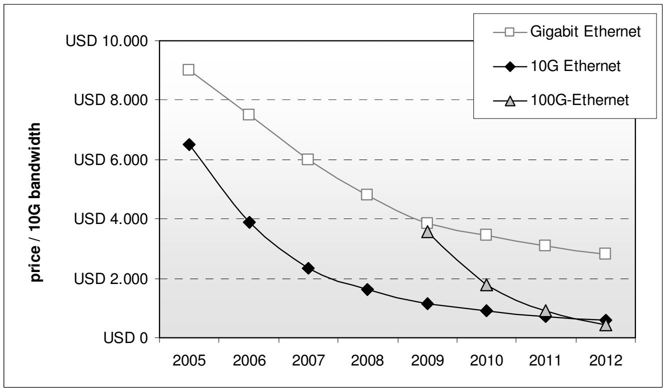

**Figure III-2: Future price development of Ethernet Bandwidth**

As a next step, the price development dynamics (initial price ratios, average price declines dep. on market maturity) of 10GE and GbE were used to estimate future prices for GbE, 10GE, and finally 100GE after its market entry. The market entry of 100GE is expected for the year 2009, as a standardization typically needs 3 years and is expected to start this year (2006). This then determines the curves of Figure III-2.

For SDH and POS linecards, price extrapolations proved to be more difficult than for Ethernet as less market data is publicly available. Therefore, the methodology used to predict future SDH and POS port prices was to consider the current price ratios compared to Ethernet bandwidth and assume that this price ratio remains constant. The prices of POS, SDH, and Ethernet interfaces for routers and switches are illustrated in the Figure III-3 below for different optics [21].

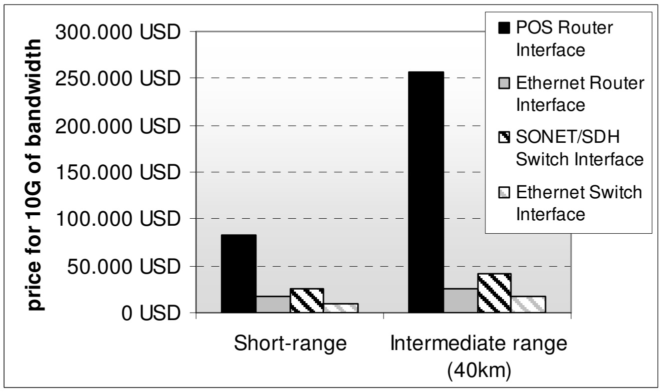

**Figure III-3: Price comparison of router and switch interfaces**

The following generic network architectures were considered in the subsequent business cases:

- IP/POS-over-WDM: The backbone consists of Label Edge Routers (LERs) and Label Switch Routers (LSRs), which are all equipped with POS interfaces. SONET/SDH is only used for transporting the IP packets node to node, directly over WDM or fiber. A 1+1 protection scheme is applied.
- IP/POS-over-SDH-over-WDM: This network scenario considers a backbone where LERs are located at the ingress and egress points of the backbone. The traffic is switched and groomed along SDH add-drop-multiplexers and cross-connects inside the core. A 1+1 protection scheme is applied.
- IP/POS-over-OXC-over-WDM: This case is similar to the previous architecture, however the SDH switches are replaced by optical cross-connects (OXCs). The range of the optical signal is assumed to be large enough to enable end-to-end optical grooming. A 1+1 protection scheme is applied.
- IP/MPLS-over-Ethernet-over-WDM: The backbone consists of LERs at the edge and MPLS-enabled Ethernet switches in the core of the backbone. This architecture corresponds to a MPLS Ethernet architecture as outlined in II (Figure II-1). A 1:1

protection scheme is applied, i.e. all capacity is overprovisioned by 100%.

- Ethernet-over-WDM: The core and outer core are a native Ethernet network with Ethernet switches both at edge and core. A few LERs are deployed to handle a small share of traffic that requires IP routing (share assumed to be 30%), however Ethernet traffic does not have to traverse LERs at the ingress and egress points of the backbone. 1:1 protection is applied.
- Ethernet-over-WDM with service-level protection: The network architecture is identical to the one before except that only premium traffic is protected against failure (share of premium traffic set to 30%).

Figure III-4 illustrates the accumulated CAPEX results for the years 2009 to 2012. The CAPEX is split up into components belonging to the IP layer (LERs, LSRs and interfaces) and components belonging to the transport layer (e.g. Ethernet switches and interfaces). Note: Since the O/E transponders are present in both cases their prices were not counted here.

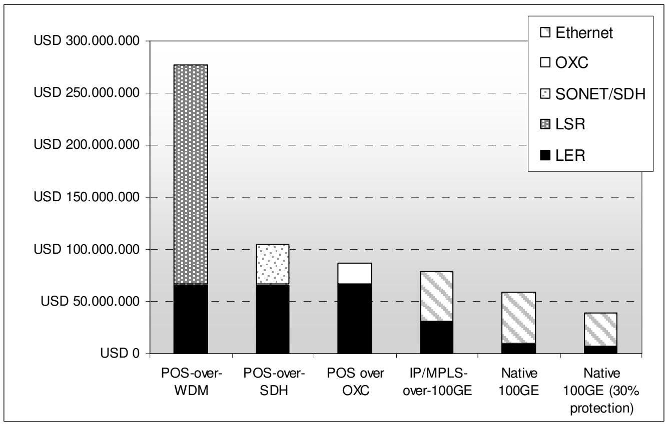

**Figure III-4: Accumulated CAPEX comparison**

The cost component of LERs is high for all SDH-related infrastructures due to the high prices of POS interfaces. On the transport layer, a pure POS-over-WDM network causes the most costs as only expensive router interfaces are used. POS-over-SDH architectures prove to be much cheaper as the SDH network employs SDH switches and interfaces instead of expensive LSR equipment for core switching. A POS-over-OXC network has an even better CAPEX performance due to lower switch and optical transceiver prices of OXC hardware.

Without exception the Ethernet business cases perform all better than SDH architectures. The MPLS Ethernet business case is the most expensive business case among the possible Ethernet architectures, as still a considerable amount of CAPEX is related to LERs and corresponding interfaces. A native 100G-Ethernet network enables higher savings in the LER category. Native 100G-Ethernet networks that employ a service-level differentiated protection scheme, the CAPEX can be reduced even further.

If we consider the the CAPEX performance of the different business cases over time the Ethernet business cases show high initial CAPEX of 100GE networks (new technology) in comparison to SDH network architectures (incremental approach). However, from the year 2009 on the incremental CAPEX of Ethernet networks decrease to a level way below SDH network CAPEX and on an extendend timeframe the CAPEX advantage of Ethernet would even increase.

## IV. OPEX COMPARISON OF DIFFERENT BACKBONE NETWORK ARCHITECTURES

While OPEX generally include a lot more categories, the repair process is selected here since the impact of 100G-Ethernet can be predicted most visibly in this area. Remark: As WDM failures and fiber breaks are not considered the total OPEX repair process values are very low.

The methodology used in these OPEX business cases was to determine the total number and type of equipment for each network architecture and year as done in section III. By using availability figures [25] the average repair time for a given backbone architecture can be estimated. The related costs are derived by multiplying the total repair time with the average salary of a field or point-of-presence technician. The general scenario and assumptions are mostly identical to the CAPEX consideration.

Figure IV-1 shows the results of the repair process OPEX, accumulated over the years 2009 to 2012 (OPEX rising with network growth):

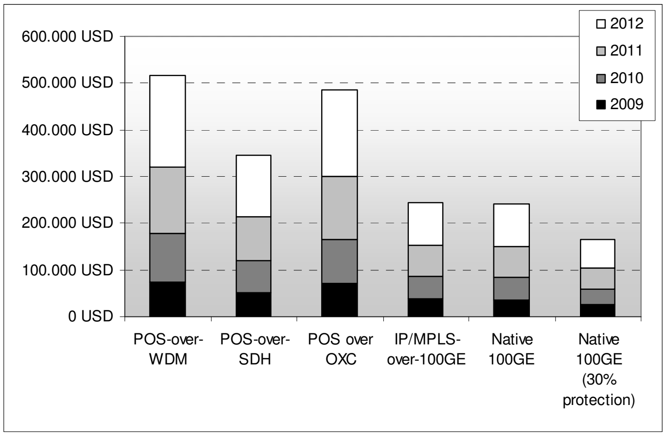

**Figure IV-1: OPEX repair process comparison**

Again, Ethernet networks are more economical than SDH-architectures, which is mostly due to the reduced device count enabled by 100G-Ethernet: Compared to the 40Gbit/s POS/SDH interfaces, Ethernet networks require less switches and linecards due to the higher port bandwidth of 100Gbit/s. The possibility of applying a service-level protection scheme can further reduce overprovisioning and the related OPEX.

While the OPEX of the repair process is very low for Ethernet, further OPEX savings may be enabled via the provisioning of Ethernet services in comparison to legacy services like e.g. leased line or Frame Relay. A study conducted by the MEF indicates total OPEX savings of around 20% in affected areas [26]. Thus, the OPEX advantage of Ethernet is not only limited to a small share of OPEX but Ethernet architectures enable considerable savings in a wide range of OPEX areas.

## V. CONCLUSION

Ethernet evolved from LAN into Metro areas covering speeds from 10 Mbps up to 10 Gbps and the next generation Ethernet speed of 100Gbps will be the enabler of Ethenet-based pure packet core networks. Carrier-gradeness of Ethernet-based packet architectures is the major point. A careful analysis of the required protocol features like network resilience, QoS, and OAM shows many redundancies within the layers of todays network architectures that have to be resolved shaping a new end-to-end Ethernet layer with the required scalability.

A CAPEX and OPEX analysis demonstrates a considerable cost advantage of 100G-Ethernet in comparison to SDH-based solutions. The superior CAPEX performance results from a huge cost advantage of Ethernet devices and their fast price decline. The reduced switch and linecard count in 100G-Ethernet networks and the efficient economics of Ethernet services are responsible for a superior OPEX performance.

Therefore, it can be said that Ethernet has a promising future in core networks, not just as link technology suppporting an upper routing layer, but as a complete, cost-effective, and service-oriented infrastructure layer in the area of core networks. The industry-wide efforts to cover remaining challenges also confirm this outlook.

## ACKNOWLEDGMENT

The authors thank all people contributing to this study for their friendly support, especially Monica Jäger from T-Systems and Andreas Iselt, Sandrine Pasqualini, and Harald Rohde from Siemens Corporate Technology.

## REFERENCES

[1] IEEE, "Provider Backbone Bridges," [Online]. Available: http://www.ieee802.org/1/pages/802.1ah.html
[2] IETF, "MPLS Label Stack Encoding," (RFC 3032) [Online]. Available: http://www.ietf.org/rfc/rfc3032.txt
[3] IEC, "Generalized Multiprotocol Label Switching (GMPLS)" [Online]. Available: http://www.iec.org/online/tutorials/gmpls/
[4] S. Khandekar, "Developing A QoS Framework For Metro Ethernet Networks To Effectively Support Carrier-Class SLAs," [Online]. Available: http://www.metroethernetforum.org/Presentations/IIR_DevelopingQoSFramework.pdf
[5] D. Papadimitriou, J. Choi, "A Framework for Generalized MPLS (GMPLS) Ethernet," IETF Internet Draft [Online]. Available: http://www.ietf.org/internet-drafts/draft-papadimitriou-ccamp-gmpls-ethernet-framework-00.txt
[6] Atrica Corp., "Delivering Hard QoS in Carrier Ethernet Networks," [Online]. Available: http://www.atrica.com/body/products/whitepapers/Delivering_Hard_QoS_in_Carrier_Ethernet_Networks_v12.pdf
[7] MEF, "Bandwidth Profiles for Ethernet Services" [Online]. Available: http://www.metroethernetforum.org/PDFs/WhitePapers/Bandwidth-Profiles-for-Ethernet-Services.pdf
[8] IEEE, "Media Access Control (MAC) Bridges," [Online]. Available: http://standards.ieee.org/getieee802/download/802.1D-2004.pdf
[9] MEF, "Requirements and Framework for Ethernet Service Protection in Metro Ethernet Networks", [Online]. Available: http://www.metroethernetforum.org/PDFs/Standards/MEF2.pdf
[10] M. Squire, "Metro Ethernet Forum OAM," [Online]. Available: http://www.metroethernetforum.org/presentations/MEF%20OAM%202003-12-01.pdf
[11] M. Suzuki, P. Bottorff, M. Chen, "Addressing Issues of Provider Backbone Bridges," [Online]. Available: http://www.ieee802.org/1/files/public/docs2004/ah-suzuki-addressing-issues-1104-01.pdf
[12] Small Tree Communications, "The Importance of Jumbo Frames in Gigabit and 10-Gigabit Networks," [Online]. Available: http://www.small-tree.com/jumbo_white_paper.pdf
[13] G. Lehmann etc., "An 160Gbit/s Network Field Trial Report," [Online]. Available: https://medicongress.be/UploadBroad/Session%2002/Presentation%2002-04.pdf
[14] D. A. Schupke, M. Jäger and R. Hülsermann, "Comparison of Resilience Mechanisms for Dynamic Services in Intelligent Optical Networks", Fourth International Workshop on the Design of Reliable Communication Networks (DRCN), Banff, Alberta, Canada, 2003.
[15] TeleGeography, "Global Internet Geography," [Online]. Available: http://www.telegeography.com/products/gig/
[16] NetworkWorld, "10G or not 10G?," [Online]. Available: http://www.networkworld.com/news/2002/120210gig.html
[17] LightReading, "Force10 Slashes 10-GigE Pricing", [Online]. Available: http://www.lightreading.com/document.asp?site=lightreading&doc_id=27068
[18] Foundry Networks, Press release, [Online]. Available: http://www.foundrynet.com/about/newsevents/releases/pr1_28_03.html
[19] Zach Fierstadt, "10G Matures," [Online]. Available: http://www.arnnet.com.au/index.php/id;2086580617;fp;256;fpid;319049444
[20] Foundry Networks, Press release, [Online]. Available: http://www.hbi.de/clients/Foundry_Networks/pms04/09_02/Modules090204.php
[21] Cisco, Global Pricelist (February 2005), [Online]. Available: http://shop.muk.com.ua/cisco_price/index.html
[22] Foundry Networks, Press release, [Online]. Available: http://www.hbi.de/clients/Foundry_Networks/pms05/18_05/Portpreise_Draft2.php
[23] NetworkWorld, "Force10 aims high with 10G switch", [Online]. Available: http://www.networkworld.com/news/2002/0830force10.html
[24] NetworkWorld, "Bronx hospital leaps to 10G" [Online]. Available: http://www.extremenetworks.com/aboutus/pressroom/news/NetWorld_Bronx10Gig.pdf
[25] NOBEL Project, "Availability Model and OPEX related parameters", to be published .
[26] MEF, "Service Provider Study: Operating Expenditures," [Online]. Available for MEF members: http://www.metroethernetforum.org

**Arno Schmid-Egger** (M'03) studied electrical engineering and information technology at Munich University of Technology (TUM), where he received his Dipl.-Ing. degree in 2005.

He is currently with Siemens AG, Power Transmission and Distribution, Energy Automation in Nuremberg, Germany. Prior to this, he was with Siemens Corporate Technology as a diploma student and worked on the topic of 100G-Ethernet. His main research interests are techno-economical analyses of different network architectures.

**Andreas Kirstädter** (M'97) received the Dipl.-Ing., Dipl.-Wirtsch.-Ing., and Dr.-Ing degrees in Electrical Engineering and Economics from Munich University of Technology (TUM), Germany, in 1990, 1992, and 1997, respectively. From 1991 to 1997 he was with the Institute for Communication Networks at Munich University of Technology, where he worked on research topics in the area of high-speed networking, WDM networks, simulation, analytical modelling and high-speed hardware design.

In 1997 he joined Siemens Corporate Technology in Munich where he was leading the high-speed networks research. Currently he is responsible for the "Optical Networks and Transmission" department within Corporate Technology.

Since 2000 Dr. Kirstädter is lecturing on "Simulation of Communication Networks" at Munich University of Technology. His current research interests include photonic networks, wavelength routing, control plane architectures, hardware implementation of communication protocols, network resilience, mobility, QoS and multimedia concepts for the Internet.