Satellite Networking in the Context of Green, Flexible and Programmable Networks

(Invited Paper)

Franco Davoli(✉)

Department of Electrical, Electronic, Telecommunications Engineering
and Naval Architecture (DITEN), University of Genoa/CNIT – University
of Genoa Research Unit, Via Opera Pia 13, 16145 Genoa, Italy
`franco.davoli@unige.it`

Abstract. In order to support heterogeneous services, using the information generated by a huge number of communicating devices, the Future Internet should be more energy-efficient, scalable and flexible than today's networking platforms, and it should allow a tighter integration among heterogeneous network segments (fixed, cellular wireless, and satellite). Flexibility and in-network programmability brought forth by Software Defined Networking (SDN) and Network Functions Virtualization (NFV) appear to be promising tools for this evolution, together with architectural choices and techniques aimed at improving the network energy efficiency (Green Networking). As a result, optimal dynamic resource allocation strategies should be readily available to support the current workload generated by applications at the required Quality of Service/Quality of Experience (QoS/QoE) levels, with minimum energy expenditure. In this framework, we briefly explore the above-mentioned paradigms, and describe their potential application in a couple of satellite networking related use cases, regarding traffic routing and gateway selection, and satellite swarms, respectively.

Keywords: Network flexibility · Network programmability · SDN · NFV · OpenFlow · Satellite networking

1 Introduction

Among other types of traffic, the Future Internet should support a very large number of heterogeneous user-led services, increased user mobility, machine-to-machine (M2M) communications, and multimedia flows with a massive presence of video. In order to face the challenges posed by the increased volume and differentiation of user-generated traffic, many Telecom operators believe that next-generation network devices and infrastructures should be more energy-efficient, scalable and flexible than those based on today's telecommunications equipment, along with a tighter integration among heterogeneous networking platforms (fixed, cellular wireless, and satellite). A possible promising solution to this problem seems to rely on extremely virtualized and "vertically" (across layers) optimized networks. At the same time, the interaction between the network and the computing infrastructure (user devices, datacenters and the cloud),

I. Bisio (Ed.): PSATS 2014, LNICST 148, pp. 1–11, 2016.
DOI: 10.1007/978-3-319-47081-8_1

where applications reside, needs to be redesigned and integrated, with the aim of achieving greater use of mass standard Information Technology (IT), ease of programmability, flexibility in resource usage, and energy efficiency (goals actually pursued since a long time in the IT world, also by means of virtualization techniques). In all network segments (access, metro/transport and core), and across different networking infrastructures, this attitude, aiming at leveraging on IT progress, as well as achieving energy consumption proportional to the traffic load, is rapidly being adopted [1–3]. In this perspective, energy efficiency also plays a central role, and can be viewed as an indicator of the "health" of the overall computing and networking ecosystem. It reflects the extent of exploitation of computing, storage, and communications hardware capabilities to the degree needed to support the current workload generated by applications at the required Quality of Service/Quality of Experience (QoS/QoE) levels. Thus, flexibility and programmability of the network itself and of all other physical resources come naturally onto the scene as instruments that allow optimal dynamic resource allocation strategies to be really implemented in practice.

In this short note, we will explore the state of the art in energy-efficiency in various networking platforms, including the satellite segment, and the integration of green technologies in the framework of two emerging paradigms for network programmability and flexibility – Software Defined Networking (SDN) [4, 5] and Network Functions Virtualization (NFV) [6] – as a sustainable path toward the Future Internet.

2 Flexibility and Programmability in the Network

Bottlenecks in the networking infrastructure have been changing over time. Whereas one of the main bottlenecks once used to be bandwidth (still to be administered carefully in some cases, though), the increase in the capacity of transmission resources and processing speed, paralleled by an unprecedented increase in user-generated traffic, has brought forth other factors that were previously concealed. Among others, some relevant aspects are:

- The networking infrastructure makes use of a large variety of hardware appliances, dedicated to specific tasks, which typically are inflexible, energy-inefficient, unsuitable to sustain reduced Time to Market of new services;
- The so-called "ossification" of the TCP/IP architectural paradigm and protocols – implemented most of the time on proprietary components – is hindering the capability to host evolutions/integrations in the standards;
- The efficient (in terms of resource usage) management and control of flows, be they user-generated or stemming from aggregation, has become increasingly complex in a purely packet-oriented transport and routing environment.

Then, as one of the main tasks of the network is allocating resources, a natural question is how to provide architectural frameworks capable of efficiently supporting algorithms and techniques that can make this task more dynamic, performance-optimized and cost-effective. Current keywords in this respect are Flexibility, Programmability, and Energy-Efficiency. SDN and NFV aim at addressing the first two. We do not enter any details here (among others, see [4–6]), but only note some essentials. By decoupling the

Control Plane and the Data (Forwarding) Plane of devices, SDN allows a more centralized vision to set the rules for handling flows in the network, by means of specific protocols for the interaction between the controller and the devices under its supervision. OpenFlow is the most well known and widespread of such protocols and a paradigmatic example. It allows setting up, updating and modifying entries in a flow table on each forwarding device, by establishing matching rules, prescribing actions, managing counters and collecting statistics. On the other hand, NFV leverages "…standard IT virtualization technology to consolidate many network equipment types onto industry standard high volume servers, switches and storage, which could be located in Datacentres, Network Nodes and in the end user premises" [6]. It fosters improved equipment consolidation, reduced time-to-market, single platform for multiple applications, users and tenants, improved scalability, multiple open eco-systems; it exploits economy of scale of the IT industry (approximately 9.5 million servers shipped in 2011 against approximately 1.5 million routers). NFV requires swift I/O performance between the physical network interfaces of the hardware and the software user-plane in the virtual functions, to enable sufficiently fast processing, and a well-integrated network management and cloud orchestration system, to benefit from the advantages of dynamic resource allocation and to ensure a smooth operation of the NFV-enabled networks [3]. SDN is not a requirement for NFV, but NFV can benefit from being deployed in conjunction with SDN. Some examples of this integration are provided in [3], also in relation to energy-efficiency, which will be the subject of the next Section. For instance, an SDN switch could be used to selectively redirect a portion of the production traffic to a server running virtualized network functions. This way the server and functions do not need to cope with all production traffic, but only with the relevant flows. The SDN-enabled virtual switch running inside the server's hypervisor can dynamically redirect traffic flows transparently to an individual network function or to a chain of network functions. This enables a very flexible operation and network management, as functions can be plugged in and out of the service chain at runtime [3]. As the main focus of these notes is on the relevance of these architectural paradigms and techniques in the context of satellite networking, we can remark explicitly that, among functionalities that would lend themselves to such treatment, we might include many of those typically delegated to Performance Enhancing Proxies (PEPs), a kind of middlebox quite frequently encountered in satellite communications.

Essentially, with the adoption of these two paradigms, the premises are there for a – technically and operationally – easier way to more sophisticated and informationally richer network control (quasi-centralized/hierarchical vs. distributed) and network management. The latter may exhibit a tighter integration with control strategies, and closer operational tools, with perhaps the main differentiation coming in terms of time scales of the physical phenomena being addressed. In our opinion, the technological setting brought forth by the new paradigms enables the application of the philosophy that was at the basis of some of the early works on hierarchical multi-level and multi-layer control concepts, both in the industrial control and networking areas [7–9], to an unprecedented extent.

3 Energy Efficiency

How does all this interact with network energy-efficiency? As a matter of fact, making the network energy-efficient (“Green”) cannot ignore QoS/QoE requirements. At the same time, much higher flexibility, as well as enhanced control and management capabilities, are required to effectively deal with the performance/power consumption tradeoff, once the new dimension of energy-awareness is taken into account in all phases of network design and operation. The enhanced control and management capabilities and their tighter integration offered by the application of SDN and NFV concepts go exactly in that direction.

The reasons that drive the efforts toward “greening” the network are well known [10, 11], and the impact of green networking on cutting the power consumption and Operational Expenditure (OPEX) is non negligible [12]. Again without entering too many details, we are particularly interested here in recalling the potential of the group of techniques known as Dynamic Adaptation, where two among the typical control actions that can be applied are Low Power Idle (LPI) and Adaptive Rate (AR), consisting of the modulation of “energy operating states” in the absence and presence of traffic, respectively [11]. Their effect can be summarized in the “power profile” of energy-aware components of network devices, i.e., in the characterization of the power consumption as a function of the traffic load [12]. In terms of QoS, the difference among operating states lies essentially in the wakeup times from “sleeping modes” for LPI (where lower power consumption implies longer wakeup time) and in different operating frequencies and/or applied voltage for AR (which affects processing capacity). Therefore, there is a natural tradeoff between power and performance, which can be optimized for different values of traffic load. Given a certain number of operating states, there are then basically two different kinds of control strategies to perform Dynamic Adaptation: (i) entering a certain LPI configuration when no packet is present to be processed in a specific component of the device and exiting to a certain active configuration upon packet arrivals (which can be “sensed” in different ways); (ii) choosing the idle and operating configurations in order to optimize some long-term figure of merit (e.g., minimize average delay, average energy consumption, or a combination thereof), while at the same time respecting some given constraints on the same quantities. In the first case, control is effected at the packet level; the strategy is dynamic and based on instantaneous local information (presence or absence of packets). In the second case the control can be based on parametric optimization, typically relying on information acquired over a relatively long term (e.g., in time scales of minutes, possibly comparable to flow dynamics – anyway several orders of magnitude greater than the time scales of packet dynamics) and typically related to long-term traffic statistics (average intensity, average burst lengths, etc.). The parametric optimization with respect to energy configurations can be combined with other traffic-load related optimizations, like load balancing in multi-core device architectures [13].

It is worth noting that optimization techniques at different time scales require some form of modeling of the dynamics of the system under control. In this respect, models based on “classical” queuing theory [13, 14] lend themselves to performance analysis or parametric optimization for adaptive control and management policies over the longer time scales (with respect to queueing dynamics). The already cited examples are in packet processing engines at routers' line cards [13] and in Green Ethernet transmission modules [14]. On the other hand, fluid models suitable for real-time control can be derived from the classical queueing equations (we recall here the very interesting approach pursued in [15]), or even from simpler, measurement-based, stochastic continuous fluid approximations [16]. In [15], optimal dynamic control strategies were applied upon fluid models derived from the classical queueing theory approach, but capable of describing the dynamic evolution of average quantities of interest (e.g., queue lengths). In our opinion, it would be worth revisiting the approach in the light of the new power consumption/performance tradeoff.

The above-mentioned models and techniques are suitable for Local Control Policies (LCPs), to be applied at the device level. However, it is also important to be able to establish energy-aware Traffic Engineering and routing policies at a "global" level (i.e., regarding a whole network domain), residing in the Control Plane and typically acting on flows, which we can refer to as Network Control Policies (LCPs). These have been considered in the recent literature, for instance in [17–20], also in relation with SDN capabilities [20]. In this respect, a relevant issue concerns the interaction between LCPs and NCPs, and the way to expose energy-aware capabilities, energy-profiles and energy-related parameters affecting QoS (e.g., wakeup delays) toward the Control Plane. A significant step in this direction has been achieved through the definition of the Green Abstraction Layer (GAL) [21, 22], now an ETSI standard [23], which allows summarizing the essential characteristics that are needed to implement energy-aware NCPs and to possibly modify device-level parameter settings accordingly.

Whereas most of the recent work cited so far was implemented in the framework of the ECONET project [24], which was devoted to energy-efficiency in the fixed network, it is worth pointing out that very similar situations in which Dynamic Adaptation strategies find useful applications are encountered also in the wireless environment [25, 26] and in datacenters [27, 28].

4 Satellites in a Green and Flexible Heterogeneous Networking Environment

A recent survey on energy-efficiency in satellite networking is that of Alagöz and Gür [29]. They discuss aspects related to the device level (terminal/earth station/satellite payload) regarding security and energy efficiency, energy constraints in the airborne platform, integration with the terrestrial segment, mobile terminals, as well as networking aspects, particularly in the context of hybrid heterogeneous networks, with the satellite playing the role of relay between various access networks and the core. They also explore emerging factors such as dynamic spectrum access and cognitive radio, cross-layer design, integration of space/terrestrial networks, Smart Grid support, emergency communications, and the Interplanetary Internet. Among some additional recent works related to energy-efficient satellite communications that appeared after the survey we can cite [30–33]. Reference [33] is related to one of the two exemplary topics we will briefly discuss in the following, and it applies what appears to be a very promising optimal control technique, based on Lyapunov optimization [34].

Here we consider two different satellite environments in their relation with flexible and green networking: (i) High Throughput Satellite (HTS) systems (at Terabit/s capacity) [35]; (ii) Nano-satellite networks (or, more generally, satellite swarms) [36].

4.1 HTS Scenario

HTS systems operate in Ka band to the users, but the scarcity of the available spectrum pushes to the use of the Q/V (40/50 GHz) bands for the gateways [37]. At these high bandwidths, where rain attenuation can produce particularly deep fading, gateway diversity is adopted to ensure the required feeder link availability [38, 39]. In essence, when each user is assigned to a pool of gateways (so-called Smart Gateways), a switching decision must be taken whenever the gateway serving the user experiences deep fading, to reroute the traffic to another unfaded gateway. Apart from the different architectural choices and ways to achieve the goal, gateway cooperation is required to efficiently obtain the desired availability level at a reasonable cost. Handover decisions should be taken at the Network Control Center (NCC), where channel state information from all the gateways should be conveyed.

At the same time, in integrated satellite-terrestrial architectures such as that envisioned by the BATS (Broadband Access via integrated Terrestrial & Satellite systems) project [35], Intelligent Network Gateways (INGs), as well as their user-side counterparts Intelligent User Gateways (IUGs), will be required to take routing decisions on traffic flows, on the basis of QoS/QoE requirements.

Then, let us recall the SDN and energy-aware scenario sketched in Sects. 2 and 3 above, and consider a situation were proper enhancement to OpenFlow allows taking advantage of the information conveyed through the GAL [20]. We can then imagine to have SDN-enabled network nodes (possibly a subset of them [40]), capable of executing power management primitives (e.g., Dynamic Adaptation, Sleeping/Standby) and associated LCPs, and an SDN Control Plane with an Orchestrator/NCC (that can reside in a cloud) in charge of implementing NCPs. SDN network nodes can include Smart Satellite Gateways, either directly or indirectly (through the SDN-enabled upstream router). Each interaction between the NCPs and the LCPs is performed according to the OpenFlow Specification.

Then, we can envisage a situation as depicted in Fig. 1, where incoming traffic is (dynamically at the flow level) directed to terrestrial or satellite paths according to joint Energy Efficiency and QoS/QoE performance indexes, and decisions are taken (dynamically with respect to channel outage conditions) on redirecting flows (or re-adjusting their balance [41]) among satellite gateways. We do not maintain the necessity of SDN for the implementation of such scenario (nor its straightforward feasibility); however, the architectural implications, the possible solutions, the required protocol extensions and the performance evaluation are certainly worth investigating.

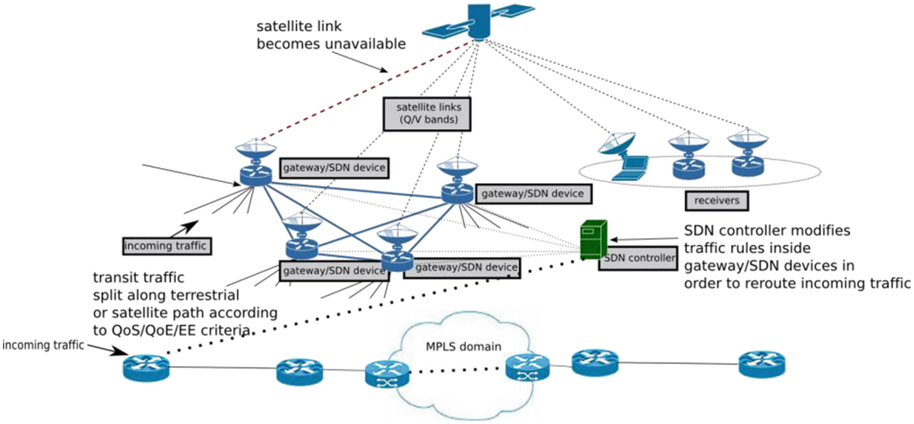

Fig. 1. HTS scenario integrated with SDN.

4.2 Satellite Swarms

There is a recent growing interest in this area, owing to the continuous development of the Internet of Things and to the desire to overcome the digital divide [36, 42], fostered by the relatively low cost of such solutions as compared to the traditional non-geostationary (NGEO) ones. Operating according to a Delay Tolerant Networking (DTN) paradigm [43] is practically a must here, and we should note the "intrinsic" energy-efficient operating characteristics of DTN. By forming a store-and-forward overlay network at the Bundle Layer [44], DTN performs grouping of smaller messages into larger aggregations, which can then be scheduled for transmission opportunities. In terms of exploiting the smart-sleeping techniques that constitute a category of methodologies for green networking, this kind of operating characteristics tends to increase the overall energy-efficiency of the system. Indeed – though operating at the packet level – one of the earliest proposed strategies to exploit smart sleeping and adaptive rate techniques has been the so-called "buffer and burst" [45], and "packet coalescing" has been suggested in connection with the Green Ethernet [46]. Forwarding decisions could then be taken at the bundle layer with attention to link/node availability and delay, but also to energy efficiency.

Recent work in this area [47] has taken into consideration the dynamic "hot spot" selection, where hot spots here play the role of small gateways that upload bundles to the satellites, which will then forward them to "cold spots" connecting users in rural or secluded areas. Here again, providing SDN capabilities to the hot spots and to the central node of the nano-satellite constellation is worth investigating, from the architectural and performance evaluation points of view.

5 Conclusions

We have briefly recalled the potential benefits of introducing flexibility, programmability and energy efficiency in the network, at all segments and levels. In relation to satellite communications, we have considered two specific examples, namely, HTS systems (at Terabit/s capacity) and nano-satellite networks. In both cases, we have tried to highlight the opportunities offered by SDN deployment, extended with energy-efficiency related primitives. In our opinion, this is a very challenging and timely field for further investigation, from the point of view of both protocol architecture and of the effective deployment of sophisticated network management and control strategies.

More specifically, combining SDN, NFV and energy-aware performance optimization can shape the evolution of the Future Internet and contribute to CAPEX and OPEX reduction for network operators and ISPs. Many of the concepts behind this evolution are not new and ideas have been around in many different forms; however, current advances in technology make them feasible. Sophisticated control/management techniques can be realistically deployed – both at the network edge and inside the network – to dynamically shape the allocation of resources and relocate applications and network functionalities, trading off QoS/QoE and energy at multiple granularity levels. Satellite networking does fit in this scenario as a relevant component, by:

- Providing energy efficient by-passes in the backhaul;
- Dynamically diverting flows, while preserving QoS/QoE requirements;
- Benefiting of increased flexibility in resource allocation to compensate fading in Q/V band smart diversity for Terabit/s speeds;
- Integrating with terrestrial networks;
- Adding energy efficient solutions in the access network for rural areas (nano-satellites and DTN);
- Benefiting of virtualization in the flexible implementation of related functionalities (PEP, optimization strategies in the cloud, …);
- Participating in consolidation of flows over a limited number of paths where possible.

Further research activities are needed for the full development of a large spectrum of possibilities.

References

1. Manzalini, A.: Future edge ICT networks. IEEE ComSoc MMTC E-lett. **7**(7), 1–4 (2012)
2. Manzalini, A., Minerva, R., Callegati, F., Cerroni, W., Campi, A.: Clouds of virtual machines in edge networks. IEEE Commun. Mag. **51**(7), 53–70 (2013)
3. Jarschel, M., Hoßfeld, T., Davoli, F., Bolla, R., Bruschi, R., Carrega, A.: SDN-enabled energy-efficient network management. In: Samdanis, K., Rost, P., Maeder, A., Meo, M., Verikoukis, C. (eds.) Green Communications: Principles, Concepts and Practice. Wiley, Chichester (2015)

4. Software-Defined Networking: The New Norm for Networks, Open Networking Foundation (ONF) White Paper (2012)
5. Nunes, B.A.A., Mendonça, M., Nguyen, X.-N., Obraczka, K., Turletti, T.: A survey of software-defined networking: past, present, and future of programmable networks. IEEE Commun. Surv. Tutorials **16**(3), 1617–1634 (2014)
6. Network Functions Virtualisation – Introductory White Paper. SDN and OpenFlow World Congress, Darmstadt, Germany (2012)
7. Lefkowitz, I.: Multilevel approach applied to control system design. J. Basic Eng. Trans. ASME **88**(Ser. B2), 392–398 (1966)
8. Findeisen, W., Bailey, F.N., Brdyś, M., Malinowski, K., Tatjewski, P., Woźniak, A.: Control and Coordination in Hierarchical Systems. Wiley, Chichester (1980)
9. Hui, J.Y.: Resource allocation for broadband networks. IEEE J. Select. Areas Commun. **6**(9), 1598–1608 (1988)
10. Global e-Sustainability Initiative (GeSI) Report. SMARTer2020: The Role of ICT in Driving a Sustainable Future. http://gesi.org/SMARTer2020
11. Bolla, R., Bruschi, R., Davoli, F., Cucchietti, F.: Energy efficiency in the future internet: a survey of existing approaches and trends in energy-aware fixed network infrastructures. IEEE Commun. Surv. Tutorials **13**(2), 223–244 (2011)
12. Bolla, R., Bruschi, R., Carrega, A., Davoli, F., Suino, D., Vassilakis, C., Zafeiropoulos, A.: Cutting the energy bills of internet service providers and telecoms through power management: an impact analysis. Comput. Netw. **56**(10), 2320–2342 (2012)
13. Bolla, R., Bruschi, R., Carrega, A., Davoli, F.: Green networking with packet processing engines: modeling and optimization. IEEE/ACM Trans. Netw. **22**(1), 110–123 (2014)
14. Bolla, R., Bruschi, R., Carrega, A., Davoli, F., Lago, P.: A closed-form model for the IEEE 802.3az network and power performance. IEEE J. Select. Areas Commun. **32**(1), 16–27 (2014)
15. Filipiak, J.: Modelling and Control of Dynamic Flows in Communication Networks. Springer, Berlin (1988)
16. Bruschi, R., Davoli, F., Mongelli, M.: Adaptive frequency control of packet processing engines in telecommunication networks. IEEE Commun. Lett. **18**(7), 1135–1138 (2014)
17. Bolla, R., Bruschi, R., Cianfrani, A., Listanti, M.: Enabling backbone networks to sleep. IEEE Netw. **25**(2), 26–31 (2011)
18. Bianzino, A.P., Chiaraviglio, L., Mellia, M., Rougier, J.-L.: GRiDA: green distributed algorithm for energy-efficient IP backbone networks. Comput. Netw. **56**(14), 3219–3232 (2012)
19. Niewiadomska-Szynkiewicz, E., Sikora, A., Arabas, P., Kołodziej, J.: Control system for reducing energy consumption in backbone computer networks. Concurrency Comput. Pract. Exp. **25**(12), 1738–1754 (2013)
20. Bolla, R., Bruschi, R., Davoli, F., Lombardo, A., Lombardo, C., Morabito, G., Riccobene, V.: Green extension of OpenFlow. In: ITC 26 Workshop on Energy-efficiency, Programmability, Flexibility and Integration in Future Network Architectures (EPFI), Karlskrona, Sweden (2014)
21. Bolla, R., Bruschi, R., Davoli, F., Di Gregorio, L., Donadio, P., Fialho, L., Collier, M., Lombardo, A., Reforgiato Recupero, D., Szemethy, T.: The green abstraction layer: a standard power management interface for next-generation network devices. IEEE Internet Comput. **17**(2), 82–86 (2013)
22. Bolla, R., Bruschi, R., Davoli, F., Donadio, P., Fialho, L., Collier, M., Lombardo, A., Reforgiato, D., Riccobene, V., Szemethy, T.: A northbound interface for power management in next generation network devices. IEEE Commun. Mag. **52**(1), 149–157 (2014)

23. ETSI Standard ES 203 237 V1.1.1, Environmental Engineering (EE); Green Abstraction Layer (GAL); Power Management Capabilities of the Future Energy Telecommunication Fixed Network Nodes, ETSI (2014)
24. http://www.econet-project.eu
25. Correia, L.M., Zeller, D., Blume, O., Ferling, D., Jading, Y., Gódor, I., Auer, G., Van der Perre, L.: Challenges and enabling technologies for energy aware mobile radio networks. IEEE Commun. Mag. **48**(11), 66–72 (2010)
26. Scalia, L., Biermann, T., Choi, C., Kozu, K., Kellerer, W.: Power-efficient mobile backhaul design for CoMP support in future wireless access systems. In: IEEE INFOCOM 2011 Workshop on Green Communications and Networking, Shanghai, China, pp. 253–258 (2011)
27. Baliga, J., Ayre, R.W.A., Hinton, K., Tucker, R.S.: Green cloud computing: balancing energy in processing, storage, and transport. Proc. IEEE **99**(1), 149–167 (2011)
28. Liu, J., Zhao, F., Liu, X., He, W.: Challenges towards elastic power management in internet data centers. In: IEEE International Conference on Distributed Computing Systems Workshops, Los Alamitos, CA, pp. 65–72 (2009)
29. Alagöz, F., Gür, G.: Energy efficiency and satellite networking: a holistic overview. Proc. IEEE **99**(11), 1954–1979 (2011)
30. Kandeepan, S., Rasheed, T., Reisenfeld, S.: Energy efficient cooperative HAP-terrestrial communication systems. In: Giambene, G., Sacchi, C. (eds.) PSATS 2011. LNICST, vol. 71, pp. 151–164. Springer, Heidelberg (2011)
31. Alagöz, F., Gür, G.: Energy-efficient layered content dissemination for multi-mode mobile devices. In: 2013 First International Black Sea Conference on Communications and Networking (BlackSeaCom), Batumi, Georgia, pp. 243–246 (2013)
32. Tang, Z., Wu, C., Feng, Z., Zhao, B., Yu, W.: Improving availability through energy-saving optimization in LEO satellite networks. In: Mahendra, M.S., et al. (eds.) ICT-EurAsia 2014. LNCS, vol. 8407, pp. 680–689. Springer, Heidelberg (2014)
33. An, Y., Li, J., Fang, W., Wang, B., Guo, Q., Li, J., Li, X., Du, X.: EESE: Energy-efficient communication between satellite swarms and earth stations. In: 16th IEEE International Conference on Advanced Communication Technology (ICACT 2014), PyeongChang, Korea, pp. 845–850 (2014)
34. Neely, M.J.: Stochastic Network Optimization with Application to Communication and Queueing Systems. Morgan & Claypool, San Rafael (2010)
35. Pérez-Trufero, J., Watts, S., Peters, G., Evans, B., Fesquet, T., Dervin, M.: Broadband access via integrated terrestrial and satellite systems (BATS). In: 19th Ka and Broadband Communications, Navigation and Earth Observation Conference, Florence, Italy, pp. 27–34 (2013)
36. Burleigh, S.: Nanosatellites for universal network access. In: 2013 ACM MobiCom Workshop on Lowest Cost Denominator Networking for Universal Access (LCDNet 2013), Miami, FL, pp. 33–34 (2013)
37. Kyrgiazos, A., Evans, B., Thompson, P., Mathiopoulos, P.T., Papaharalabos, S.: A Terabit/second satellite system for European broadband access: a feasibility study. Int. J. Satell. Commun. Netw. **32**(2), 63–92 (2014)
38. Jeannin, N., Castanet, L., Radzik, J., Bousquet, M., Evans, B., Thompson, P.: Smart gateways for Terabit/s satellite. Int. J. Satell. Commun. Netw. **32**(2), 93–106 (2014)
39. Gharanjik, A., Rama Rao, B.S.M., Arapoglou, P.-D., Ottersten, B.: Gateway switching in Q/V band satellite feeder links. IEEE Commun. Lett. **17**(7), 1384–1387 (2013)
40. Agarwal, S., Kodialam, M., Lakshman, T.V.: Traffic engineering in software defined networks. In: IEEE INFOCOM 2013, Torino, Italy, pp. 2211–2219 (2013)

41. Kyrgiazos, A., Thompson, P., Evans, B.: Gateway diversity via flexible resource allocation in a multibeam SS-TDMA system. IEEE Commun. Lett. **17**(9), 1762–1765 (2013)
42. https://www.outernet.is
43. http://www.dtnrg.org/wiki/Home
44. Scott, K., Burleigh, S.: Bundle Protocol Specification. RFC 5050, IETF (2007)
45. Nedevschi, S., Popa, L., Iannaccone, G., Ratnasamy, S., Wetherall, D.: Reducing network energy consumption via sleeping and rate-adaptation. In: 5th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2008), San Francisco, CA, pp. 323–336 (2008)
46. Christensen, K., Reviriego, P., Nordman, B., Bennett, M., Mostowfi, M., Maestro, J.A.: IEEE 802.3az: the road to energy efficient ethernet. IEEE Commun. Mag. **48**(11), 50–56 (2010)
47. Cello, M., Marchese, M., Patrone, F.: Hot spot selection in rural access nanosatellite networks. In: ACM CHANTS 2014, Maui, HI (2014)