

Energy Efficiency in integrated IT and Optical Network Infrastructures: The GEYSERS approach

Anna Tzanakaki¹, Markos Anastasopoulos¹, Konstantinos Georgakilas¹, Jens Buysse², Marc De Leenheer², Chris Develder², Shuping Peng³, Reza Nejabati³, Eduard Escalona³, Dimitra Simeonidou³, Nicola Ciulli⁴, Giada Landi⁴, Marc Brogle⁵, Alessandro Manfredi⁵, Ester López⁶, Jordi Ferrer Riera⁶, Joan A. Garcia-Espin⁶, Pasquale Donadio⁷, Giorgio Parladori⁷, Javier Jiménez⁸

¹Athens Information Technology (AIT), Peania, Greece, Email: atza@ait.edu.gr, ²IBBT, Ghent, Belgium, ³University of Essex, Colchester, UK, ⁴Nextworks, Pisa, Italy, ⁵SAP, Zurich, Switzerland, ⁶i2CAT, Barcelona, Spain, ⁷Alcatel-Lucent, Italy, ⁸Telefonica, Madrid, Spain

Abstract—In this paper we propose energy efficient design and operation of infrastructures incorporating integrated optical network and IT resources. For the first time we quantify significant energy savings of a complete solution jointly optimizing the allocation and provisioning of both network and IT resources. Our approach involves virtualization of the infrastructure resources and it is proposed and developed in the framework of the European project GEYSERS - Generalised Architecture for Dynamic Infrastructure Services.

Keywords- Optical fiber networks; WDM networks; virtual infrastructures, energy efficiency, IT resources

I. INTRODUCTION

It is estimated that Information and Communication Technology (ICT) accounts for 4% of the primary energy consumption worldwide [1]. The expansion of the Internet in size and complexity incurs increased energy consumption of both IT and network resources, thus leading to an acute interest in energy efficient networking [2]. Designing, engineering and operating infrastructures comprising IT resources and communication networks in a power-aware manner seems to be instrumental for a more energy-efficient Future Internet. Therefore, it becomes crucial for network and IT resources to be integrated and operated in a coordinated fashion, optimizing them jointly in design and operation.

IT resources require very high levels of power consumption for their operation and their conventional operating window is commonly not optimized for energy efficiency. Hence, allocating IT processing jobs in an energy aware manner through a relatively low energy consuming optical network infrastructure is expected to offer significant energy savings. On the other hand, although optical networking is an energy efficient technology to use, authors in [2] show that the system design itself cannot fulfill the needs for limiting the overall power consumption, thus network design and protocol implementation have to also become power-aware. Authors in [3] follow an energy-minimizing network design approach and identify solutions that enable switching-off of both node and link equipment when unused. In [4] an extensive optimization and simulation framework that minimizes energy consumption in IP-over-WDM networks is presented, showing substantial energy consumption reduction.

Following the increasing interest in energy-efficient infrastructures, in this paper we propose for the first time a complete solution facilitating energy efficient operation of infrastructures incorporating integrated optical network and IT resources. This approach is proposed and developed in the framework of the European project GEYSERS - Generalised Architecture for Dynamic Infrastructure Services [5] that proposes the concept of Virtual Infrastructures (VIs) over one or more interconnected Physical Infrastructures (PIs) comprising both network and IT resources. The introduction of VIs in this context facilitates sharing of physical resources among various virtual operators, introducing a new business model that suits well the nature and characteristics of the Future Internet and enables new exploitation opportunities for the underlying physical infrastructures.

The paper is structured as follows: Section II provides an overview of the GEYSERS approach and architecture, section III describes the energy consumption models used for the optical network and the IT resources. Section IV details the VI planning methodology and models that have been used to optimally design energy efficient VIs considering a specific PI, section V details an energy efficient service provisioning approach and finally section VI summarizes the conclusions.

II. GEYSERS APPROACH

GEYSERS [5] architecture presents an innovative approach by adopting the concepts of the Infrastructure as a Service (IaaS) and service oriented networking to enable infrastructure operators offer new IT and optical network converged services. Figure 1 shows the GEYSERS's layered architecture.

In the GEYSERS layered architecture, devices in the PI layer are abstracted and partitioned or grouped into virtual resources that can be selected to form the Virtual Infrastructures in the Logical Infrastructure Composition Layer (LICL). Within each VI, controllers in the IT-aware network control plane (NCP+) layer configure and manage virtual network resources, while virtual IT node controllers at the Virtual IT Manager (VITM) control the virtual IT resources. The Service Middleware Layer (SML) is responsible for translating the application requests and service level agreements (SLAs) into technology specific requests to trigger the provisioning procedures at the NCP+.




Figure 1: GEYSERS layered architecture

In GEYSERS energy efficiency is targeted, in an integrated IT and optical network infrastructure, taking a multilayer approach. GEYSERS is inherently addressing the issue of energy efficiency, given its focus on optical network technology that has been proven to be significantly more energy efficient than alternative network technologies [6]. However, additional energy efficiency considerations are addressed in both the VI planning and the VI operation/service provisioning phases taking into consideration the combined energy consumption of both optical network and IT resources.

Regarding the VI planning phase the responsible layer is LICL. LICL defines and implements a resource information model to uniquely identify and uniformly abstract physical resources, including energy efficiency properties of physical infrastructure elements. The objective of VI planning is to implement a dynamically reconfigurable virtual network that not only meets customer's specific needs, but also satisfies the virtual infrastructure provider's-driven requirements, while maintaining cost-effectiveness and in our case the specific requirement of energy efficiency. Through this process the least energy consuming VI that can support the required services is identified in terms of both topology and resources. To identify this least energy consuming VI the detailed power consumption models and figures of the underlying physical infrastructure, including joint consideration of optical network and IT resources, are taken into consideration. Mapping the virtual resources to the physical resources is also part of the VI planning phase. Therefore the VI planning phase is also responsible to define the energy consumption parameters of the VI itself. Through the VI planning process, virtual resources will be effectively abstracted from the physical devices and will be marked with "green parameters". Using these parameters the virtual infrastructures can be operated in an energy-efficient manner, through the implementation of energy-aware provisioning mechanisms at the NCP+ layer.

The GEYSERS NCP+, based on the ASON/GMPLS and Path Computation Element (PCE) architectures and protocols, is aiming at playing a key role in the operational energy efficiency of VI infrastructures. This can be achieved as the NCP+ is responsible for the path computation, required as part

of the service provisioning, and therefore the corresponding resource allocation of both network and IT resources. In the effort to maximize the energy efficiency of operating VIs, the NCP+ can compute end-to-end paths with an objective different to these considered in conventional approaches (e.g. length-based, or hop-based shortest path etc.), such as the minimization of the energy consumption associated with the corresponding services. This computation is performed at the PCE(s) and, in case of inter-domain paths, requires the cooperation of multiple PCEs following a hierarchical model. The routing algorithms implemented at the PCEs take into account the "green" parameters, describing the power consumption of the heterogeneous resources involved. To further increase the energy saving during the operation of the VI the option of switching-off of unused resources (network and IT) is also proposed. This approach involves the introduction of a sleep or stand-by mode during which unused resources can be set to, and from which they can be awoken in a fast and seamless manner when required to be used [7], [8].

In addition, in the GEYSERS architecture the LICL dynamically and continuously synchronizes the energy status of physical resources and their associated virtual resources. Therefore, sudden energy changes of the infrastructure can be compensated by initiating VI re-planning.

III. ENERGY MODELS FOR NETWORK AND IT RESOURCES

The estimation of energy consumption of the PI network resources is highly sensitive to the network architecture employed and the network technology used. As GEYSERS is focusing on optical network technologies utilizing wavelength division multiplexing (WDM) the energy consumption of the corresponding technologies is taken into account. The overall network power consumption model is based on the power-dissipating (active) elements of the network and can be classified as switching nodes, referred to as Optical Cross-Connect (OXC) nodes in the context of WDM networks, and transmission line related elements. More specifically the OXCs assumed are based on the Wavelength Selective Switch architecture using Micro-Electrical Mechanical Systems (MEMS), while for the fiber links a model comprising a sequence of alternating single mode fiber and dispersion compensating fiber spans together with optical amplifiers to compensate for the losses is employed. The details of these models are described in [9], with the only difference being that unlike [9] the current work assumes wavelength conversion capability available at the OXC nodes.

The physical IT infrastructures in this paper are considered to be data centers i.e. facilities used to house computer systems and associated components. Data centers are primarily used to process data (servers) and data storage (storage equipment). The collection of this processing and storage equipment is referred to as the "IT equipment". The rest of the IT infrastructure facilities, such as cooling mechanisms etc correspond to "infrastructure equipment". A large part of the energy consumption of a data centre resides in the IT equipment and the cooling [10], making infrastructure equipment (especially cooling) an important factor when trying to reduce the energy consumption in data centers.

In this paper we base our analysis on a linear power consumption model that mainly concentrates on the power consumption associated with the CPU load of IT resources. More specifically p_s is referred to the CPU resources and \mathcal{E}_s to the energy consumption for utilizing a portion $u_s = p_s/p_s^{\max}$ of the maximum CPU capabilities, p_s^{\max} , of server s . For simplicity, the following linear energy consumption model has been adopted [11], [12]

$$\mathcal{E}_s(u_s) = P_{\text{idle}}^s + (P_{\text{busy}}^s - P_{\text{idle}}^s)u_s$$

where P_{idle}^s and P_{busy}^s is the energy consumption of IT server s at idle state and under full load, respectively. For further details regarding the technical specifications of the IT servers the reader is referred to [11].

It should be also noted, that in addition to the power consumption of both network and IT resources, a 100% power overhead due to cooling has been incorporated in the energy consumption model described above.

IV. ENERGY AWARE VI PLANNING

The VI planning is the process of using historical operating data and estimated future virtual resource (VR) requirements, to determine the optimal design of the VI. The objective of VI planning is to identify the topology and determine the virtual resources required to implement a dynamically reconfigurable VI based on both optical network and IT resources. This VI not only will meet customer's specific needs, but will also satisfy the virtual infrastructure provider's (VIP's) requirements for minimum energy consumption in our case. The Energy Aware VI planning problem is formulated through a Mixed Integer Linear Programming (MILP) model that aims at minimizing jointly the total energy that is consumed by optical network components including WDM transponders, amplifiers, switches and IT resources.

A. Problem Formulation

The problem is formulated using a network that is composed of one resource layer that contains the physical infrastructure (1st lower layer) and will produce as an output the virtual infrastructure layer (2nd lower layer) illustrated in Figure 1. The physical infrastructure is described through an eleven-node topology corresponding to the Pan-European optical network in which randomly selected nodes generate demands d ($d = 1, 2, \dots, D$) to be served by a set of IT servers s ($s = 1, 2, \dots, S$). The granularity of demands is the wavelength. The IT locations (demand destinations) at which the services will be handled, are not specified and are of no importance to the services themselves. Therefore, the demand destinations will be identified through the optimization performed by our model. In order to formulate this problem, the binary variable a_{ds} is introduced to indicate whether demand d is assigned to server s or not and it equals 1, if and only if demand d is processed on server s . Moreover, it is assumed that each demand can be assigned only to one server:

$$\sum_s a_{ds} = 1 \quad d = 1, 2, \dots, D \quad (1)$$

Furthermore, for each demand d , its demand volume h_d is realized by means of a number of lightpaths assigned to paths of the VI. Let $p = 1, 2, \dots, P_{ds}$ be the candidate path list in the VI for the lightpaths required to support demand d at server s and x_{dps} the non-negative number of lightpaths allocated to path p . The following demand constraints should be satisfied in the VI:

$$\sum_s \sum_p a_{ds} x_{dps} = h_d \quad d = 1, 2, \dots, D \quad (2)$$

Summing up the lightpaths through each link e ($e = 1, 2, \dots, E$) of the VI we can determine the required link capacity y_e for link e :

$$\sum_d \sum_s \sum_p \delta_{edps} x_{dps} \leq y_e \quad e = 1, 2, \dots, E \quad (3)$$

where δ_{edps} is a binary variable defined as follows

$$\delta_{edps} = \begin{cases} 1, & \text{if link } e \text{ of VI belongs to path } p \\ & \text{realizing demand } d \text{ at server } s \\ 0, & \text{otherwise} \end{cases}$$

Using the same rationale, the capacity of each link e in the VI is allocated by identifying the required lightpaths in the PI. The resulting PI lightpaths z determine the load of each link g ($g = 1, 2, \dots, G$) of the PI, and hence its capacity u_g . Assuming that $q = 1, 2, \dots, Q_e$ is used for denoting the PI's candidate path list realizing link e , then, the following demand constraint for link e should be satisfied:

$$\sum_q z_{eq} = y_e \quad e = 1, 2, \dots, E \quad (4)$$

where the sum is taken over all paths q on the routing list Q_e of link e . Introducing the link-path incidence coefficients for the PI

$$\gamma_{geq} = \begin{cases} 1, & \text{if link } g \text{ of PI belongs to path } q \\ & \text{realizing link } e \text{ of VI} \\ 0, & \text{otherwise} \end{cases}$$

the general formula specifying the PI capacity constraint can be stated as:

$$\sum_e \sum_q \gamma_{geq} z_{eq} \leq u_g \quad g = 1, 2, \dots, G \quad (5)$$

where G is the total number of links in the PI and the summation for each link g is taken over all lightpaths in the PI layer.

Apart from link capacity constraints (3), (5) for the VI and PI, respectively, the total demands that are assigned to each server should not exceed its capacity p_s , $s = 1, 2, \dots, S$. The latter capacity corresponds to the underlying physical resources, such as CPU, memory, disk storage etc. The inequality specifying servers' capacity constraints is given by

$$\sum_d \sum_p a_{ds} c_{ds}(x_{dps}) \leq p_s, \quad s = 1, 2, \dots, S$$

where the summation is taken over all demands that arrive at server s and $c_{ds}(x_{ds})$ is a parameter specifying the computational requirements for demand d on server s . In practice, this parameter is determined by the set of relevant benchmarks for computer systems provided by the Standard Performance Evaluation Corporation (SPEC) [13].

The objective of the current problem formulation is to minimize the total cost of the resulting network configuration as this cost consists of the following components:

- a. k_g that is the cost of the capacity of link g of the PI. It consists of the energy consumed by each lightpath due to transmission and reception of the optical signal, optical amplification at each fiber span and switching according to the model described in [9]
- b. \mathcal{E}_s that is the cost for using capacity p_s of the IT servers. The linear energy consumption model [11], [12] described in section III has been adopted where:

$$\mathcal{E}_s(\mathbf{u}_s) = P_{\text{idle}}^s + (P_{\text{busy}}^s - P_{\text{idle}}^s) \mathbf{u}_s$$

In this context, minimum energy consuming VI is obtained by minimizing the following cost function

$$\text{Minimize } F = \sum_g k_g u_g + \sum_s \mathcal{E}_s(\mathbf{u}_s)$$

The above MILP problem has been solved analytically employing the methods of Lagrangian relaxation and dual decomposition [14].

B. Numerical Results

To investigate the energy efficiency of the proposed VI design scheme, the GEYSERS architecture illustrated in Figure 1 is considered: the lower layer depicts the PI and the layer above depicts the VI. For the PI the COST239 Pan-European

reference topology [15] has been used in which four randomly selected nodes generate demands to be served by two IT servers located in Luxemburg and Milan. Furthermore, we assume a single fiber per link, 40 wavelengths per fiber, and wavelength channels of 10Gb/s each. It is also assumed that each IT server can process up to 2Tb/s, and its power consumption ranges from 6.6 to 13.2KW, under idle and full load, respectively [11], while c_{ds} has been assumed to be equal to 1.

An example of the optimal VI topology design for a scenario in which four source nodes that are located in London, Vienna, Copenhagen and Paris generate demands equal to 50 wavelengths each, is depicted in Figure 2. In this scenario, the generated VI topology consists of 7 virtual links and 6 virtual nodes, while all demands are routed to the IT server in Luxemburg. The capacity of each virtual link along with its mapping to the PI is given in Table I where e.g. it is observed that virtual link Y3 connecting Copenhagen and Luxemburg is realized via physical layer paths u5-u10 and u8-u17-u13, with capacities 25 and 40 wavelengths, respectively.

In Figure 3, the performance of the proposed energy aware VI design is compared to the demand allocation scheme presented in [16] where demands from each source node are assigned to its closest IT server. Note that ‘‘closest’’ refers to the shortest distance between a source node and a data center. Comparing these two schemes, it is observed that the energy aware VI design consumes significantly lower energy for serving the same amount of demands compared to the closest IT scheme in the order of 30%: in the former approach fewer IT servers are activated to serve the same amount of demands. Given that the power consumption required for the operation of the IT servers is dominant in this type of networks, switching-off the unused IT resources achieves significant reduction of energy consumption. Furthermore, it is observed that in both schemes the average power consumption increases almost linearly with the number of demands. However, the relative benefit of the energy aware design decreases slightly with the number of demands, as we get closer to full system load.

TABLE I. SAMPLE VIRTUAL TO PHYSICAL MAPPING

Virtual link	Capacity (wavelengths)	Physical Layer Paths realizing virtual links	Capacity of PI paths (wavelengths)	Average cost / wavelength for Virtual link
Y2	50	Path A: u2-u10	15	21,97
		Path B: u3-u11	35	18,97
Y3	65	Path A: u5-u10	25	20,96
		Path B: u8-u17-u13	40	29,69
Y4	15	u5-u6-u7-u15-u21	15	51,86
Y5	15	u14-u23	15	20,83
Y6	40	u22	40	11,25
Y7	40	u16	40	11,62
Y8	70	Path A: u12	40	9,83
		Path B: u15-u14	25	22,01
		Path C: u7-u11	5	18,68




Figure 2: Example of the virtualization of a physical infrastructure




Figure 3: Comparison of the energy aware scheme with the closest IT server demand allocation scheme

V. ENERGY AWARE SERVICE PROVISIONING OVER VIS

As already discussed the multi-layer GEYSERS architecture allows for energy efficiency considerations both at the VI planning phase as well as the operation phase of the VI through the use of energy efficient service provisioning supported by the NCP+.

In order to facilitate energy efficient operation of the VI we propose an energy efficient routing algorithm. This algorithm aims at provisioning IT services that are originating from specific source sites and need to be executed by suitable IT resources (e.g. data centers). Candidate IT resources reside at different geographical locations and connectivity of the source

site to these IT resources is provided through the underlying optical network as dictated by the already designed VI, available as the output of the VI planning phase. The routing scheme used is anycast routing [17], since the requirement for the IT services is the delivery of results, while the exact location of the execution of the job is of no interest. In this context energy efficiency is achieved by identifying the least energy consuming IT and network resources required to support the services, and switching-off of any unused network resources (such as links, nodes, etc) and IT resources (such as servers) [7].

To evaluate the performance of energy awareness in the service provisioning phase, we have created an Integer Linear Program (ILP) based on the model described in [17] modified appropriately. This ILP assumes a capacitated VI and an energy consumption model where a datacenter's energy increases linearly with processing load, as described in detail in section III. The energy consumption of the optical network is calculated through the VI planning phase described in detail in section IV. Our model takes as inputs the: (a) parameters specifying the energy consumption of the VI IT and network resources, (b) a set of IT resource sites which are able to handle IT requests, (c) a set of lightpath requests per source, which need to end in one of the proposed server sites, and (d) the capacity of the links and the Data Centers. The output of the model includes: (a) the IT resource allocation i.e. which server site serves which source node's light path request and (b) a route to the allocated IT site allowing switching of server sites, OXCs and links.

In order to compare the performance of the proposed energy efficient algorithm, we have also implemented three alternative strategies: (i) one where the objective is to minimize

the number of used wavelengths, which corresponds to conventional shortest path (SP) calculation (ii) one where we only try to minimize the energy consumed by the network elements and (iii) a third one which aims at minimizing the energy consumed by the IT resource sites only. We have run the calculations for 10 random demand vectors for each demand size (15 to 285), so the results which are plotted in Figure 4 represent averages over these 10 demand vectors.




Figure 4: Energy aware service provisioning schemes vs SP routing

Figure 4 illustrates the energy consumption of the VI infrastructure when applying the proposed energy aware algorithm compared to SP. It can be seen that the energy saving achieved through the energy aware scheme reaches 40% for low service loads and reduces with the load. This effect of the energy saving reduction, with the service load, is due to that at high loads, most of the VI resources are utilized and the overall energy consumption reaches its maximum value. Among all approaches under consideration the highest VI energy consumption is obtained as expected, in the case where no energy consideration is taking place. In addition, the approach that focuses on the network power consumption provides very similar performance to that of the SP, with a small difference (3,37% on average). On the other hand, the proposed scheme considering jointly the energy consumption of IT and network resources always provides optimum performance. However, it should be noted that the routing approach taking into consideration only the power consumption of IT resources, demonstrates only a small additional energy penalty compared to the proposed solution (a difference ranging from 2,76% up to 11,05% depending on the offered load). These observations clearly indicate that in the VI under consideration the IT resources have the most dominant contribution in the overall power consumption compared to the optical network resources. Figure 4 shows that when applying the proposed routing approach as well as the approach considering the energy consumption of IT resources, a step like increase in energy consumption is observed at 195 requests. This is due to that a resource site's maximum capacity is 200 requests. Hence, the step like increase indicates the startup cost for switching-on a second server site. In low loading conditions the preferable solution for energy minimization involves the use of only one server site. However, in high load conditions switching-on a second server site is inevitable to accommodate the load.

VI. CONCLUSIONS

This paper focuses on virtualization of infrastructures incorporating integrated optical network and IT resources as they are proposed and developed in the framework of the EU project GEYSERS. In this context an energy efficient approach is proposed addressing both the design and operation of these infrastructures and significant energy savings are quantified when modeling the proposed solution.

ACKNOWLEDGMENT

This work was carried out with the support of the GEYSERS (FP7-ICT-248657) project funded by the European Commission through the 7th ICT Framework Program.

REFERENCES

- [1] M. Pickavet et al., "Worldwide energy needs for ICT: The rise of power-aware networking", in Proc. IEEE ANTS, pp. 1-3, Dec. 2008.
- [2] J. Baliga et al., "Energy Consumption in Optical IP Networks", Journal of Lightwave Technology, Vol. 27, No. 13, July 1, 2009.
- [3] L. Chiaraviglio, M. Mellia, and F. Neri, "Reducing power consumption in backbone networks", in Proc. IEEE ICC, pp. 1-6, June 2009.
- [4] G. Shen and R.S. Tucker, "Energy-Minimized Design for IP Over WDM Networks", in IEEE/OSA JOCN, Vol. 1, No. 1, pp. 176-186, June 2009.
- [5] <http://www.geysers.eu/>
- [6] R. S. Tucker et al., "Evolution of WDM Optical IP networks: A cost and energy perspective", Journal of Lightwave Technology, Vol. 27, No. 3, pp. 243-252, Feb. 2009.
- [7] Energy Star, "Small network equipment", www.energystar.gov/index.cfm?c=new_specs.small_networkequip.
- [8] A. Muhammad, P. Monti, I. Cerutti, L. Wosinska, P. Castoldi, and A. Tzanakaki, "Energy-Efficient WDM Network Planning with Protection Resources in Sleep Mode", in Proc. Globecom 2010, Miami, Florida, USA, 6-10 December 2010.
- [9] A. Tzanakaki et al., "Power Considerations towards a Sustainable Pan-European Network", JWA061, OFC2011.
- [10] Z. Davis, "Power Consumption and cooling in the data center: A survey", 2005.
- [11] Oracle Data Sheet, "Sun Oracle DataBaseMachine", <http://www.oracle.com/us/products/database/database-machine-069034.html>
- [12] Xiaobo Fan, Wolf-Dietrich Weber, and Luiz Andre Barroso, Power provisioning for a warehouse-sized computer. *SIGARCH Comput. Archit. News* 35, 2, 13-23 June 2007.
- [13] Standard Performance Evaluation Corporation (SPEC) (www.spec.org)
- [14] M. P. Anastasopoulos, A. D. Panagopoulos, P. G. Cottis, "A Distributed Routing Protocol for QoS provisioning in Wireless Mesh Networks operating above 10GHz", Wireless Communications and Mobile Computing, Volume 8, Issue 10, pp. 1233-1245, Dec. 2008.
- [15] P. Batchelor et al., "Study on the implementation of optical transparent transport networks in the European environment-Results of the research project COST 239", Photonic Network Communications, Vol. 2, No. 1, pp. 15-32, 2000.
- [16] K. Bouyoucef, I. Limam-Bedhief, and O. Cherkaoui, "Optimal Allocation Approach of Virtual Servers in Cloud Computing", in Proc. of 6th Euro-NF conference on Next Generation Internet, July 2010.
- [17] J. Buysse, M. De Leenheer, B. Dhoedt, C. Develder, "On the impact of relocation on network dimensions in optical grids", in Proc. Int. Conf. Optical Netw. Design and Modeling (ONDM 2010), Kyoto, Japan, 31 Jan.-3 Feb. 2010.