

A network model for routing-fault-free wavelength selection in WRONoCs design

Maddalena Nonato^{a,1} Davide Bertozzi^{a,1} Marco Gavanelli^{a,1}
Andrea Peano^{a,1}

^a *Dipartimento di Ingegneria
Università degli Studi di Ferrara
Ferrara, Italy*

Abstract

Emerging technologies in on-chip communication domain bring about new combinatorial optimization problems at design automation. We address the *Wavelength Selection Problem* in Wavelength-Routed Optical Networks-on-Chip (WRONoCs), where wavelengths act as signal carriers for initiator-to-target communication, so that signals are the least interfering and routing faults are prevented. We present this novel engineering problem and model it as a constrained shortest path on acyclic networks, propose a graph-based mathematical formulation and an iterative procedure on incremental graphs to solve the model on realistic data.

Keywords: WRONoC, constrained shortest path, node forbidden pairs.

1 Introduction

The increasing computer systems performance based on integrated circuits downscaling has fulfilled *Moore's Law* predictions for the last 30 years. This trend has now reached a saturation point in communication, the so called *interconnect bottleneck*. To overcome this flaw disruptive interconnect technologies are being investigated among which Optical Networks-on-Chip (ONoCs), where communication among core components at chip level exploits the potentials of silicon photonics to carry information by light. This technology

¹ Email: maddalena.nonato,davide.bertozzi,marco.gavanelli,andrea.peano@unife.it

has several advantages with respect to its electronic counterparts, including ultra-high communication bandwidth, lower power dissipation, and lower latency [4], beside increased flexibility in resource usage [3]. Despite optical technologies for long-distance communication are well assessed, their application at chip level still presents several challenges. During the synthesis process leading from abstract system specifications to manufacturing, system designers face several decision problems for which there are no reliable standardized methodologies, despite a few having nice combinatorial structures. Here we address one of such problems arising at the fifth of the seven steps of the automatic synthesis process for Wavelength-selective Routing ONoCs (WRONoCs) proposed in [2]. In WRONoCs, communication among cores is based on wavelength selective routing protocols providing all-to-all contention-free connectivity: a master sends information to each receiver via a different wavelength channel and the same set of channels is reused across masters with no need for arbitration. Given the network topology, information delivery to target relies on the accurate selection of the network parameters, such as the signal-carrying wavelengths and the radii of the optical switches WRONoCs are made of, so that interferences as well as potential routing faults are prevented, which is the topic of this paper. In fact, the path taken by an optical signal in the topology depends on its wavelength, and second-order effects of optical switching may route signals to unintended destinations. In Sect.2 we model this problem as a shortest path on acyclic networks with cardinality constraints and node forbidden pairs, and provide an Integer Linear Programming (ILP) model embedded into an iterative procedure on incremental graphs; in Sect.3 we report computational results for realistic fabrication data applied to a well known WRONoC topology, and ongoing research is sketched.

2 A network model for Carrier Wavelength Selection

Carrier Wavelength Selection in WRONoC

In WRONoCs, m initiators I send bits over optical channels to m targets T using a tuple of n_λ wavelengths $\tilde{\Lambda}_{i,t} = \langle \tilde{\lambda}_{i,t}^1 \dots \tilde{\lambda}_{i,t}^{n_\lambda} \rangle$ for each $\langle i, t \rangle$ pair, $i \in I, t \in T$. Due to the Wavelength-selective Routing (WR) protocol, the same n_λ -tuple $\tilde{\Lambda}_{i,t}$ can be reused as signal carriers by other pairs $\langle i', t' \rangle$ as long as $i \neq i', t \neq t'$. Then, mn_λ different wavelength channels are sufficient to support an m to m communication at n_λ bit rate parallelism. At this stage of the network synthesis process, though, the tuples $\tilde{\Lambda}_{i,t}$ are just abstract symbols and optimally mapping them to feasible values in the actual bandwidth $\bar{\Lambda} = [\lambda_{min}, \lambda_{max}]$ is the topic of this paper. This step requires the understand-

ing of the fundamental mechanisms of the WR protocol which guarantee a unique path from i to t for each n_λ -tuple $\tilde{\Lambda}_{i,t}$ on the interconnect network. Each initiator i issues a Wavelength-Division Multiplexed (WDM) signal made of m n_λ distinct carriers with wavelength $\tilde{\Lambda}_{i,t_1}, \dots, \tilde{\Lambda}_{i,t_m}$, each n_λ -tuple assigned to its target $t \in T$. The WR protocol ensures signal delivery to target exploiting the network topology and the ability of optical switches to deviate on a different waveguide all and only those wavelengths that are resonant peaks of the transmission characteristic of the Micro Ring Resonators (MRRs) an optical switch is made of [6]. These peaks depend on the radius of the MRR, and the length of the ring radius is the first parameter to be set. Then, for each radius, n_λ among its resonating wavelengths are chosen to be used as signal carriers. The network topology determines how many switches of how many different radii are necessary to provide an m to m communication: usually $n_R=m$ radii are required. Formally, we face a combinatorial problem that we call *Wavelength Selection Problem for ONoC* (WSP-ONoC) whose decisions are the following: given the set R of n radii available from production, select a subset $R^* \subset R$ of n_R radii, and for each $r \in R^*$ pick n_λ wavelengths in the set $\Lambda(r)$ of the *resonating* wavelengths of r . Solving WSP-ONoC provides physical values to the abstract symbols $\tilde{\Lambda}_{i,t}$. The WDM signal issued by each initiator is thus made of the $n_\lambda n_R$ wavelengths in $\Lambda_{R^*}^* = \bigcup_{r \in R^*} \Lambda^*(r)$, where $\Lambda^*(r) \subset \Lambda(r)$ denotes the n_λ selected resonances for radius r . The aim is that such wavelengths are evenly distributed in $\bar{\Lambda}$, provided that interferences and routing faults are prevented, i.e., all wavelengths in $\Lambda^*(r)$ are routed along the same it -path on the interconnect ONoC. Interference avoidance is a common concern to most wavelength based communication paradigms and motivates the objective function. Indeed, wavelengths that are very close to each other generate optical cross-talks that degrade the signal-to-noise ratio. To avoid it, we forbid wavelengths closer than a given threshold ε and aim at keeping wavelengths as apart as possible, which corresponds to equitably spacing the $n_\lambda n_R$ selected wavelengths in the bandwidth. Interfering (closer than ε) wavelengths are mutually exclusive, one at most can be selected. On the contrary, routing fault prevention is a peculiar feature of WR-ONoCs due to the WR protocol. A *routing fault* occurs when a wavelength carrier is unduly deviated by an MRR. Formally, let $\lambda_r^i \in \Lambda(r)$ denote the i^{th} resonating peak in $\bar{\Lambda}$ of the transmission characteristic of the MRRs with radius r . Fabrication errors during the radii manufacturing may cause two peaks of different radii λ_q^j and $\lambda_r^i : |\lambda_r^i - \lambda_q^j| < \epsilon$ to overlap, ϵ depending on fabrication. Let $\Gamma(\lambda_r^i)$ be the set of wavelengths in $\bar{\Lambda}$ potentially overlapping with λ_r^i , let $\Theta(\lambda_r^i) = \{q \in R : \Gamma(\lambda_r^i) \cap \Lambda(q) \neq \emptyset\}$ be the set of the radii of the poten-

tially overlapping wavelengths, and let $\Omega(\lambda_r^i) = \bigcup_{q \in \Theta(\lambda_r^i)} \Lambda(q)$ be the set of resonances of such radii. We call the wavelengths in $\Omega(\lambda_r^i)$ the *conflicting resonances* of λ_r^i . The following constraint arises: if λ_r^i is selected no wavelength in $\Omega(\lambda_r^i)$ can, and vice-versa. Note that when $\varepsilon \leq \epsilon$ routing fault prevention avoids interference, however the two constraints must be handled differently: interference is a one-to-one relation, while routing fault is one-to-many. Now we show how WSP-ONoC can be modeled as a shortest path on acyclic networks with cardinality constraints and node forbidden pairs.

A graph based Integer Linear Programming model

Given $\Lambda_R = \bigcup_{r \in R} \Lambda(r)$ the set of resonating wavelengths of the available radii in R , consider a digraph $G = (N, A)$ where N is made of the wavelengths in Λ_R plus two dummy nodes s and t associated to λ_{min} and λ_{max} that act as origin and destination of the path, respectively, i.e., $N = \{s, t\} \cup \{i_r : \lambda_r^i \in \Lambda_R\}$. For each pair of wavelengths λ_r^i and λ_q^j that are potentially adjacent in Λ_{R^*} when sequenced in non decreasing order, the arc $a_{i_r j_q}$ belongs to A . G is acyclic but quite dense since the forward star of i_r includes any arc $a_{i_r j_q}$ such that $\lambda_q^j > \lambda_r^i + \varepsilon$ (no arc connects interfering wavelengths) and $\lambda_q^j \notin \Omega(\lambda_r^i)$. Regarding the objective function, as the cardinality of Λ_{R^*} is given, maximally spacing selected adjacent wavelengths corresponds to equally distributing the $n_\lambda n_R$ wavelengths in the frequency range $\bar{\Lambda}$. Thus the ideal gap between two successive values in Λ_{R^*} is $\delta = (\lambda_{max} - \lambda_{min}) / (n_\lambda \cdot n_R - 1)$, and the cost $c_{i_r j_q}$ of arc $a_{i_r j_q}$ is the absolute value of the difference between $\lambda_q^j - \lambda_r^i$ and δ . Additional arcs connect s to each node i_r and each i_r to t . Their cost is the gap between λ_{min} and λ_r^i , and from λ_r^i to λ_{max} , respectively. A solution is an $(n_\lambda n_R + 1)$ arcs path from s to t .

Let $z_{i_r j_q}$ be the binary variable associated to arc $a_{i_r j_q}$, let x_{i_r} be the one associated to node i_r , and let y_r be the selection variable for radius r . Model (1-12) provides an ILP formulation for WSP-ONoC.

$$\min \sum_{a_{i_r j_q} \in A} c_{i_r j_q} z_{i_r j_q} \quad \text{subject to:} \quad (1)$$

$$\sum_{a_{i_r j_q} \in FS(i_r)} z_{i_r j_q} - \sum_{a_{j_q i_r} \in BS(i_r)} z_{j_q i_r} = 0 \quad \forall i_r \in N \quad (2)$$

$$\sum_{a_{i_r t} \in BS(t)} z_{i_r t} = 1 \quad (3)$$

$$\sum_{a_{s i_r} \in FS(s)} z_{s i_r} = 1 \quad (4)$$

$$\sum_{a_{i_r j_q} \in FS(i_r)} z_{i_r j_q} = x_{i_r} \quad \forall i_r \in N \quad (5)$$

$$\sum_{i_r: \lambda_r^i \in \Lambda(r)} x_{i_r} = n_\lambda y_r \quad \forall r \in R \quad (6)$$

$$\sum_{r \in R} y_r = n_R \quad (7)$$

$$x_{i_r} \leq y_r \quad \forall \lambda_r^i \in \Lambda(r), r \in R \quad (8)$$

$$x_{i_r} + y_q \leq 1 \quad \forall q \in \Theta(\lambda_r^i), \forall \lambda_r^i \in \Lambda(r), r \in R \quad (9)$$

$$z_{i_r j_q} \in \{0, 1\} \quad \forall a_{i_r j_q} \in A \quad (10)$$

$$x_{i_r} \in \{0, 1\} \quad \forall i_r \in N \quad (11)$$

$$y_r \in \{0, 1\} \quad \forall r \in R \quad (12)$$

Equations (2-3) are classical flow balance constraints describing the route of a unit flow from s to t . Equations (5) define node variables x_{i_r} as a function of arc variables $z_{i_r j_q}$. Feasibility with respect to cardinality constraints is ensured by (6) and (7): (7) guarantees that n_R radii are chosen and (6) guarantees that either n_λ resonances are selected if r is chosen or none if r is discarded. Constraints (8) are the logical implication $x_{i_r} \rightarrow y_r$ and provide a tighter version of (6). Constraints (9) are the compact version of the *node forbidden pairs* constraints linking each λ_r^i to each of its conflicting wavelengths, which ensure routing faults prevention. In fact, for each λ_r^i each inequality $x_{i_r} + y_q \leq 1$, $q \in \Theta(\lambda_r^i)$, stands for the family $x_{i_r} + x_{j_q} \leq 1 \quad \forall \lambda_q^j \in \Lambda(q)$, casting this problem as a special case of the Path Problem with Forbidden Pairs (PPFP) discussed in [5]. In that work, complexity results are provided for some special classes of forbidden pairs, none of which holds in our problem. For the proof of NP-completeness of PPFP for acyclic graphs (by reduction from 3SAT) see [7]. Note that, while routing fault prevention must be explicitly enforced, interference is prevented by graph topology.

3 Computational results

Testbeds

The experimental campaign aims to stress the efficacy of the ILP model and test whether realistic instances can be solved as is by commercial solvers, i.e. Gurobi 6.05. The experiments run on a two machines cluster equipped with Intel Xeon $E5 - 2640@2.5$ GHz (6 cores and 24GB of RAM each), providing each job with 8GB of RAM and time out $t_{max}=3 \cdot 10^4$ seconds. We report computational results for realistic data generated according to current




Figure 1. Graph $G(\tau)$ for different values of τ . On top, the radii footprints.

industrial practice and referring to the λ -router topology, which we tested for different values of n_R and n_λ . In detail, bandwidth is $\bar{\Lambda} = [1491, 1611] \text{ nm}$; the set of radii is computed as $R = \{R_m + i * R_s, \forall i \in \{0, \dots, \lceil (R_M - R_m) / R_s \rceil\}\}$ with respect to $R_m = 5 \mu\text{m}$, $R_s = 0.25 \mu\text{m}$, $R_M = 30 \mu\text{m}$ which yields $n = 101$ radii as in [1]. The resonant wavelengths are generated exploiting an analytical electromagnetic model [8], which generates $|\Lambda_R| = 1619$ resonances in the range $\bar{\Lambda}$. We set $\epsilon = \varepsilon = 0.1 \text{ nm}$. In this scenario, we analyze the cases $m = 4$ and $m = 8$ cores with $n_R = m$ for different degrees of parallelism $n_\lambda \in \{1, 4, 8\}$, with δ spanning from ~ 1 to ~ 40 : namely, we tackle 6 instances denoted by the pairs $(n_R \times n_\lambda)$ in $\{(4 \times 1), (4 \times 4), (4 \times 8), (8 \times 1), (8 \times 4), (8 \times 8)\}$. These are well assessed real life WRONoC parameters for $n_\lambda = 1$ while the frontier for higher parallelism is still an open research topic [1]. Finally, note that the size of G depends on Λ_R while realistic values for ϵ and ε only marginally decrease the number of arcs with respect to a complete digraph, therefore its size (1621 nodes and $\sim 13 \cdot 10^6$ arcs) is basically the same for all instances.

Exploiting incremental graphs

To cope with the large size of the graph, we present an iterative procedure working on incremental graphs of increasing size obtained by filtering the arcs based on a rising threshold on their cost, exploiting the following results.

Proposition 3.1 *Let p be a cardinality feasible path from s to t in G . If arc $a_{i_r j_q} \in p$ then p costs at least $2c_{i_r j_q}$.*

Proposition 3.2 *Let $G(\tau)$ denote the subgraph obtained by deleting all arcs whose cost is greater than or equal to τ , and $z^*(\tau)$ the optimal solution value of WSP-ONoC on $G(\tau)$. If $z^*(\tau) \leq 2\tau$ then $z^*(\tau)$ is optimal. Otherwise, the optimal solution is a path in $G(z^*(\tau)/2)$.*

Proof 3.2 relies on 3.1, which in turn follows from the cost structure and the fixed cardinality of $\Lambda_{R^*}^*$. In fact, $z^*(\tau)$ is optimal if discarded arcs cannot improve the solution. However, such arcs cost at least τ therefore - because of 3.1 - any feasible path they belong to costs no less than $2\tau \geq z^*(\tau)$, proving $z^*(\tau)$ optimality. If $z^*(\tau)$ is not optimal ($z^*(\tau) > 2\tau$), it still provides an upper bound implying that the optimal path p^* cannot have arcs whose cost is greater than $z^*(\tau)/2$, therefore p^* must be a path in $G(z^*(\tau)/2)$. \square

Figure 3 shows an example with three radii and their resonating wavelengths, overlapping ones circled in yellow. We compare $G(\tau)$ for two different values of τ when solving the problem for $n_R=3$ and parallelism 2 (two nodes per radius). The smaller value τ_1 yields only the solid arcs and $G(\tau)$ admits no feasible solution (overlapping nodes are not allowed in the path). For $\tau_2 > \tau_1$ dashed arcs are added and $s - \lambda_2^1 - \lambda_1^2 - \lambda_3^3 - \lambda_3^5 - \lambda_1^5 - \lambda_2^6 - t$ is a feasible path in $G(\tau_2)$.

This suggests to compute z^* by iteratively solving the problem on a sequence of sub-networks, according to the following procedure.

```

Ensure: Optimal solution
or Unsolvable
 $\tau \leftarrow \tau_0$  {Initial value of  $\tau$ }
repeat
   $\langle \text{Solvable}, z^*(\tau) \rangle \leftarrow \text{solve}(\tau)$ 
  if !Solvable then
    Increase  $\tau$ 
until Solvable or  $\tau > \bar{\Lambda}$ 
if Solvable then
  if  $z^*(\tau) \leq 2\tau$  then
    return  $z^*(\tau)$ 
  else
    return  $\text{solve}(z^*(\tau)/2)$ 
else
  return Unsolvable

```

The model is solved on $G(\tau)$ for increasing values of τ starting from $\tau = 0, 1$ with steps of $\tau_{step} = 0, 1$. Either a feasible solution is found for some $\tau^* \leq \tau_{max} = \lambda_{max} - \lambda_{min} - \delta$ or τ reaches τ_{max} with no solution and the problem is proved infeasible. In the first case, let $z(\tau^*)$ be the solution value for the sub-network $G(\tau^*)$. If $z(\tau^*) \leq 2\tau^*$ the solution is optimal also for G . Otherwise, the optimal solution can be found by solving the model on $G(\tau') = G(z(\tau^*)/2)$, since $z(\tau') \leq z(\tau^*) = 2\tau'$. In that case we enforce no time limit.

Computational results are summarized in Table 1. Columns (left to right) report: instance size $n_R \times n_\lambda$; ideal distance δ ; τ (∞ stands for τ_{max}); number of arcs $|A|$; running time in seconds; linear programming relaxation at the root node, if any; solution value, if any. In the last column “INFEAS” means that the instance proved infeasible, “?” means that no feasible solution could be found within t_{max} ; if a feasible solution is found but optimality is not proven due to either time out (TO) or out of memory (OoM), column 5 reports the time at which the incumbent was found and the termination

status. For each instance we report information on (at most) 5 rows. Row 1 summarizes the values for the infeasible iterations, i.e., when $G(\tau)$ does not admit a feasible path: the running time is the sum for all iterations and column 3 reports minimum and maximum τ values for which the instance is infeasible. Row 2 regards the first feasible instance. Row 3 concerns the first instance such that $G(\tau)$ admits not only a feasible path, but an optimal one w.r.t. G . Row 4 (labeled ∞) reports data if running the model straight on G , while row 5 concerns the results provided by the modified Answer Set Programming (ASP) logic program. Whereas we do not have a benchmark to compare with, in [1] we searched for a feasible set of resonances maximizing n_λ for a given n_R under a pessimistic scenario about fabrication errors. That problem was coded in ASP [10] and solved by Clasp [11]. Here we provide that ASP model with a given n_λ and set fabrication errors tolerance equal to ϵ so that it returns the first feasible solution computed in the search. The rationale is to verify the need for optimization at this WRONoC design step. At the same time, ASP proved to be a useful tool to verify the existence of a feasible solution even when the ILP solver struggles to find one.

Results confirm our incremental procedure efficacy. For instances with no parallelism (4×1 and 8×1) the search converges at very low values of τ , as the first feasible solution is proved to be a global optimum. For 4×4 the overall time required by the procedure is one third of what it takes to solve the problem on G . For both 4×8 and 8×4 instances, the solver cannot prove the optimality of the incumbent on G , which is indeed quite far from optimality, within t_{max} while our procedure converges to optimality. Comparing the computational effort required for these two instances that look for the same number of resonances (and have the same δ as well), we observe that the difficulty grows with the parallelism n_λ . This comes at no surprise as the problem with highest n_λ is more constrained: indeed, optimality is reached at a higher τ for denser networks. Finally, for the largest instance 8×8 we could not prove infeasibility at iteration 9 ($\tau = 0,9$), but, since $z^{ASP} = 93,9$, we know that G admits a feasible solution so we tackled the τ_{max} case and got a feasible solution with cost $z^H(\tau_{max}) = 18,12$. According to the same rationale we applied our procedure backwards, solving the case $\tau = 9,06 = (18,12/2)$ for which a better solution with cost $z^H(9,06) = 16,05$ was found within t_{max} .

Conclusions and work in progress

We introduced an ILP model for the resonant wavelength selection problem arising in WRONoC design. The problem is formalized as a constrained shortest path on acyclic networks and solved by an iterative procedure on incremental graphs which exploits the structure of the arcs cost function. Re-

Size	δ	τ	$ A $	Time	LP	z
4×1	39.93	0.1	5835	0		0.10
		∞	1292057	99	0.1	0.10
		ASP		0		100.23
4×4	7.99	0.1 - 0.4	7070-18494	29.72		INFEAS
		0.5	22712	325.12	2.846	3.71
		1.8533	74015	2172.89	2.548	3.46
		∞	1306464	6297.25	2.546	3.46
		ASP		5		126.40
4×8	3.86	0.1 - 0.8	7069 - 36231	39567.4		INFEAS
		0.9	40553	25093.04	4.062	12.16
		6.0822	204261	97634.82	4.037	10.35
		∞	1306464	9544/OoM	4.037	18.22
		ASP		4.68		117.04
8×1	17.11	0.1 - 0.2	6906 - 10789	3.3		INFEAS
		0.3	14357	1.93		0.40
		∞	1292057	130.47		0.40
		ASP		2		81.94
8×4	3.86	0.1 - 0.4	7069 - 18852	31.64		INFEAS
		0.5	23318	4900.32	3.236	4.61
		2.30	98312	7262.91	3.238	4.58
		∞	1306464	9236/OoM	3.218	8.19
		ASP		4.7		103.56
8×8	1.90	0.1 - 0.8	7047 - 34242	3226.93		INFEAS
		0.9	38171	30000/TO	34.6	?
		9.06	224232	27834/TO	9.5061	16.05
		∞	1306464	19535/OoM	9.5058	18.12
		ASP		4.7		93.90

Table 1
Experimental results for six realistic instances

sults show that realistic instances can be solved by commercial solvers, but for the highest level of parallelism even finding a feasible solution can be challenging. While we believe to be offering practical decision support tools to designers, we are aware that the combinatorial structure of the problem can be further exploited: decomposition methods based on the network structure are currently under investigation as well as valid inequalities that could speed up the search for feasibility and tighten the weak linear relaxation lower bound. ASP confirmed its ability to find feasible solutions very quickly, however selected adjacent resonances are too close for industrial production, which confirms that any wavelength selection process must involve an optimization capability. An on going research activity concerns developing the optimization version of the ASP model and compare vs ILP. Whichever framework proves more efficient, our network based model can support decomposition

approaches that become necessary when scaling to very large size instances. At the same time, though, the ASP based approach could be integrated as an efficient tool to quickly assess feasibility on sub-networks.

References

- [1] Peano, A., L. Ramini, M. Gavanelli, M. Nonato, and D. Bertozzi, *Design Technology for Fault-Free and Maximally-Parallel Wavelength-Routed Optical Networks-on-Chip*, Proceedings of ICCAD 2016, Austin, Texas, 3:1–3:8, <http://doi.acm.org/10.1145/2966986.2967023>
- [2] Bertozzi, D, M. Ortín-Obón, A. Peano, M. Talay, M. Balboni, L. Ramini, M. Nonato, and V. Vignals-Yuferax, *Design Automation Beyond its Electronic Roots: Toward a Synthesis Methodology for Wavelength-Routed Optical Networks-on-Chip*, in “Optical Interconnects for Computing Systems”, Nicolescu G. et al. (Ed.) River Publishers, to be published. http://riverpublishers.com/book_details.php?book_id=331
- [3] Ortín, M., L. Ramini, M. Balboni, L. Zuolo, M. Nonato, V. Vials, and D. Bertozzi, *Partitioning strategies of wavelength-routed optical networks-on-chip for laser power minimization*, Proceedings of SiPhotonics 2015, 17–24.
- [4] D. Miller, *Rationale and challenges for optical interconnects to electronic chips*, Proceedings of the IEEE 88(6) 2000, 728–749.
- [5] P. Kolman, and O. Pangrc, *On the complexity of paths avoiding forbidden pairs*, Discrete Applied Mathematics 157(13) 2009, 2871–2876.
- [6] D.H. Geuzebroek, and Al. Driessen, *Ring-Resonator-Based Wavelength Filters*, in “Wavelength filters in fibre optics”, Venghaus, H.(Ed.) Springer Verlag 2006.
- [7] H.N. Gabow, S.N. Maheshwari, and L.J. Osterweil *On two problems in the generation of test paths*, IEEE Trans. Soft. Eng., SE-2 (3) 1976, 227-231.
- [8] A. Parini, L. Ramini, G. Bellanca, and D. Bertozzi, *Abstract modelling of switching elements for optical networks-on-chip with technology platform awareness*, Proceedings of INA-OCMC’11, 2011, 31–34
- [9] Boris V. Cherkassky, A. Goldberg, and T. Radzik, *Shortest paths algorithms: Theory and experimental evaluation*, Math. Prog. 73(2), 1996, 129–174.
- [10] M. Gelfond, *Answer sets* In F. van Harmelen et al. (Ed.s) “Handbook of Knowledge Representation”, Elsevier Science, 2008, 285–316.
- [11] M. Gebser, B. Kaufmann, and T. Schaub. *Conflict-driven answer set solving: From theory to practice*, Artificial Intelligence 187-188, 2012, 52–89.