Cross-layer optimization of wireless networks (CLOWN): Final report

Aleksi Penttinen
Deparment of Communications and Networking
TKK

August 21, 2008

Contents

1 Introduction

Cross-layer optimization of wireless network (CLOWN) is a Tekes-funded research project at the Department of Communications and Networking / TKK that was carried out between 1.5.2006-31.8.2008. The scientific results of the project are reported as follows:

- 4 journal articles ([24, 21, 18, 12])
- 6 articles in international conferences ([28, 19, 16, 20, 11, 17])

In addition, a state-of-the-art report on mesh networks [9] was produced jointly with the ABI project and a contribution was made to a final report of COST 290 action (to be published by Springer). This report gives an overview of the scientific results with examples of their application.

2 Background

The traditional design of data communications systems is based on a layered architecture. Protocols are stacked on each other, each hiding the implementation of the underlying layers from the protocols above. The architecture is extremely modular; each layer can be implemented independently as long as the interfaces to the adjacent layers remain unchanged. This modularity is one of the driving forces in the rapid development of the communications networks we observe today.

However, the traditional layered design has come under serious criticism in the context of wireless networks, e.g., [26]. Radio transmissions are susceptible to interference and the resulting packet losses and varying channel capacities lead to serious inefficiency with the current protocol design. In a classic example packet losses and timeouts of a wireless link are mis-interpreted as congestion by TCP (Transmission Control Protocol, responsible for congestion control and file integrity in the Internet) which then slows down the data transfer rate and the link becomes underutilized. Although there are numerous proposals for mitigating the effect of this particular phenomenon by tuning the existing protocols, the underlying problem remains with the design.

Cross-layer optimization, cf. [27], is a general approach that attempts to improve the performance of wireless networks by exchanging information and setting parameters directly across several layers of the architecture. Cross-layer optimization has inspired numerous research efforts. However, the simple approaches of adding new cross-layer interfaces into the current protocols do not usually provide enough benefits to counter the incurred loss of modularity [13].

Recently, the interest in cross-layer optimization has taken a more fundamental form. Layering as optimization decomposition [8] is an analytical perspective to cross-layer optimization. The network is modeled as an utility optimization problem and the layering corresponds to the mathematical decomposition of the problem. This view enables quantitative analysis of layering costs and guidelines "how to" or "how not to" layer and provides a solid foundation for design and analysis of novel protocol structures for future wireless communications systems.

However, the most important aspect in this development of improving the performance of wireless networks received only a little attention. Most of the traffic in the networks is data traffic which is elastic by nature. These file transfers, generally referred to as flows, can adapt their transmission rates to share the network resources. Thus the interesting performance measures for a flow, e.g., how long it takes to download a file, depend essentially on other concurrent flows the number of which is dynamically changing. Indeed, the performance of data traffic can only be described by utilizing a dynamic traffic model, where the file transfers are initiated (randomly) and also leave the network upon completion. The utility optimization approach discussed above represents a static problem, a snapshot of a dynamic system, and cannot be thus applied to evaluate the real benefits of cross-layer optimization for data communications in wireless networks.

Performance analysis of data networks in the dynamic traffic model has become available during the past few years with the development of the balanced fairness resource sharing concept. Balanced fairness, invented by Bonald and Proutière [6] in the context of wireline networks, allows one to define the stationary distribution of the flows in the system which, in addition, takes an elegant form depending only on the traffic load in the network. The distribution can then be utilized to evaluate the average system performance. Our research group has been actively contributing to the development of this scheme and its application to wireless networks [25].

The CLOWN project addressed the cross-layer optimization of data traffic performance in wireless networks utilizing the dynamic traffic model. Two particular areas were considered in the context: Performance analysis and dimensioning. In addition, the project addressed the MAC-layer subproblem of link scheduling and studied different optimization strategies in a load balancing problem of wireless networks. In the following we summarize the key scientific results of the project and provide examples of their application.

3 Flow-level performance analysis of wireless networks

Balanced fairness (developed by Bonald and Proutiere [6]) is a resource sharing notion that was introduced in fixed networks to provide a tractable model for dynamic data traffic. Roughly, it can be seen as an extension of a processor sharing queue into a network context. Utilizing balanced fairness as the resource sharing principle between the dynamic flows (file transfers that are able to adapt to the available capacity and also share this capacity among each other), one is able to compute performance measures such as the average throughput on a given route.

When trying to evaluate the potential performance of a wireless network, one can utilize a strong time scale separation that exists between the access layer and the flow layer. The access layer is responsible to resolve potential conflicts between active flows. This access resolution happens on the time scale measured in milliseconds. On the other hand, the typical file transfer duration are measured in seconds. As a result, a file transfer sees the links as "fixed" with a capacity that is the time average of the link capacity under given access mechanism. In cross-layer optimization this time scale separation provides a natural decomposition of the performance problem in wireless networks.

In [25] the concept of balanced fairness was extended to wireless networks where the link capacities are no longer fixed but depend on link scheduling. This results in a cross-layer optimization problem where the resource sharing among active flows is determined on the higher layer and the link scheduling on a lower layer. However, the resulting optimization problem becomes soon intractable as the number of routes are increased in the model. To alleviate this problem, one of the first research tasks in the CLOWN project was to extend the approximate methods developed for mainly fixed networks [5] into the computationally challenging case of wireless networks where link scheduling must be taken into account. The results of this work were presented in the article [19].

Whereas the previous works considered only the case where a centralized scheduler assigns which links are active in which time slots, the project applied the concept also on random access wireless multi-hop networks. In such a network each link has a certain activation probability in each time slot and the link is going to transmit with this probability irrespective of other links. However, the activation probability of the link remains a parameter that can be subject to cross-layer optimization. This was addressed in the conference paper [20], which was later on extended into a journal article [21].

3.1 Networks with centralized link scheduling

Performance analysis of large networks is a computationally difficult problem. In [5] the following approximative method for analysis was proposed: By computing exactly the throughput curve and its derivatives at very low and at very high loads one can easily sketch (by an interpolating function) the throughput as a function of load over the whole load range.

In the case of wireless networks we introduced detailed derivative formulae as well as derived computational methods to find the required derivatives in [19]. This work was a co-operation with the Fancy project funded by the Academy of Finland.

Figure 1 gives a rough example of the performance analysis in the wireless case: Consider a mesh network shown in the figure, with two access points (AP), two relays (R) and 18 traffic classes. We assume that no node can participate in more than one transmission at a time. We assume that each link has a unit capacity when active. The traffic pattern is such that a route i (cf. the routes in the figure) has the load proportion $\rho_i = i$. We study the throughput of class 9 (marked with thick line in the left figure). Figure 1 (right) shows the throughput behavior of class 9 as a function of normalized load r. The throughput reflects the average bandwidth the file transfers observe on the route during their transfer. The value $r = 0$ corresponds to an empty network. The throughput is 0.5 as only every second link can be active simultaneously due to the interference. The value $r = 1$ corresponds to the capacity limit where the network becomes saturated with the traffic. At this point the load of class 9 is $\hat{\rho}_9 = 3/40$ (not shown). The throughput between the extremes is sketched using a cubic polynomial fitted to the end points and to the first two derivatives at $r = 0$.

3.2 Networks with random access

Early wireless computer networks, most notably ALOHAnet [1], applied a very simple channel access principle: transmit a packet at a random time instant (or with some probability in a time slot). The packet may then be successfully received if it is not interfered by another random transmission. Again, from the flow-level perspective a slotted random access link can be

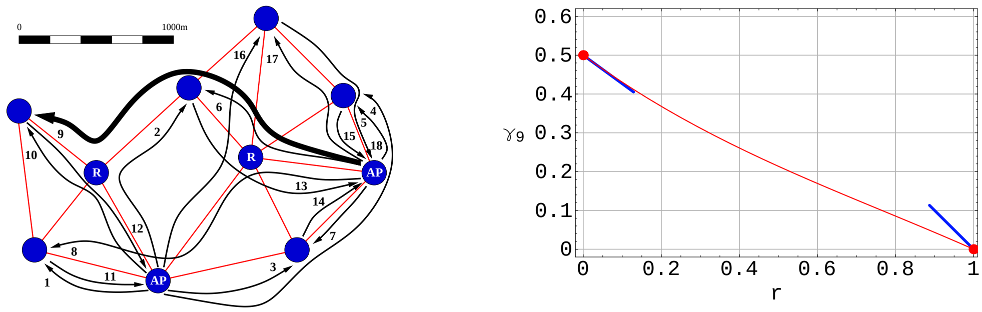

Figure 1: Left: Wireless mesh network and the flow classes. Right: Throughput asymptotics of the class denoted by the thick line (right).

seen as a "fixed" link, the capacity of which is the fraction of the nominal link capacity that corresponds to the probability of successful transmission on that link in a randomly chosen time slot.

If such a network is small, its flow-level performance can be made surprisingly good by adjusting the transmission probability. For an example in a transmission over a single link one can set the transmission probability to 1, which achieves the maximum throughput obtainable by any access scheme. The same phenomenon extends itself also to short linear networks, but with a large number of active interfering links the unavoidable collisions decrease the performance of random access. The classic result by Kleinrock [15] on single resource Aloha system states that the total optimal throughput of the channel tends to the fraction $1/e$ of the channel capacity as the number of stations accessing the channel grows.

In the article [20] and its extended version [21], we studied the performance of file transfers in such a random access, or slotted Aloha, network when the transmission probabilities are optimally adjusted every time a file transfer arrives or departs from the network. We quantified the throughput behavior of flow-optimized random access and compared it against the throughput obtained by optimal scheduling (see previous section). In particular, we derived the flow throughput analytically in certain specific scenarios. For arbitrary flow-optimized random access networks we provided a general scheme that enables evaluating the throughput. The scheme entailed also a novel algorithm to determine the maximum link capacities with given proportions achievable by the considered random access mechanism. Also this work was a co-operation with the Fancy project funded by the Academy of Finland.

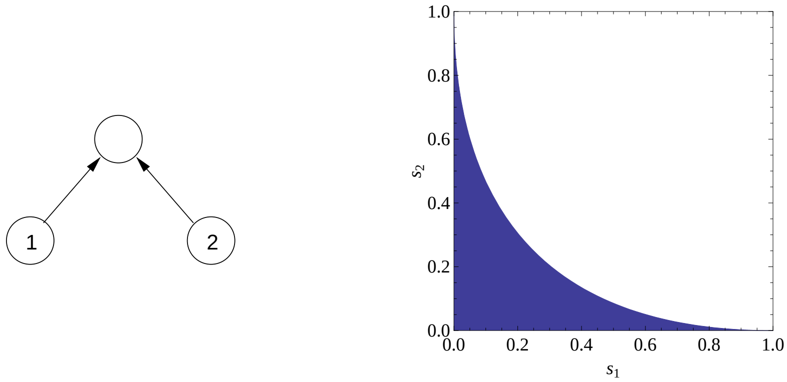

Figure 2: Two nodes accessing a common channel and the capacity set of the system under slotted random access.

3.2.1 Capacity of a wireless random access network

Computation of the flow throughput under balanced fairness starts from defining the capacity set of the network so that for any given flow rate proportions one is able to determine the maximum sustainable rates in the network. In other words, the task is to find which flow-level link capacities are attainable by adjusting the access probabilities. Although the problem is not new, only a few solutions to certain special cases existed in the literature, cf. [15].

For an example of a capacity set consider the simple network in Figure 2 (left), where two nodes are trying access a single unit-capacity channel using the Aloha method. In Figure 2 (right) the darkened area shows the flow-level link capacities that are attainable by adjusting the access probabilities, i.e., the capacity set. In the extreme cases, one of the links has the access probability 1 whereas the other has the probability 0. In other cases the attained capacity is the rate of successful transmissions. For instance, when both the access probabilities are set to 1/2 the flow-level capacity of the links equals 1/4 (probability that the link transmits in a time slot, but the other link remains silent).

In general case we do not have to know the whole capacity set at once. It suffices to be able to determine the boundary point in any given direction, i.e., what is the maximum attainable capacity for given link capacity proportions. For this purpose an algorithm was proposed in [19], which finds the boundary point by iterating two phases. The idea of the algo-

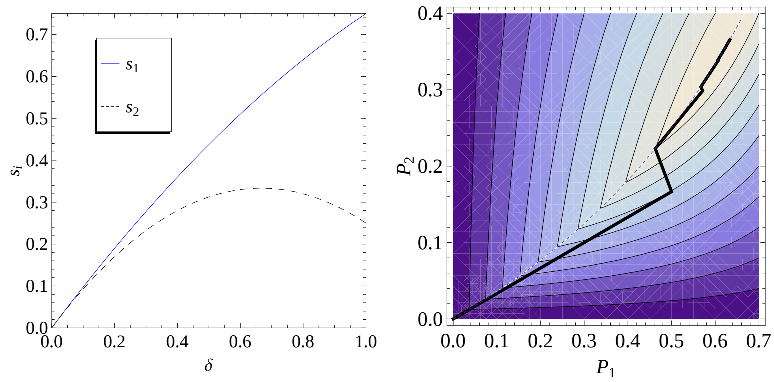

Figure 3: Illustration of the algorithm for a 2-class Aloha system with capacity proportions $\frac{3}{4}$ and $\frac{1}{4}$.

rithm is illustrated in Figure 3, which studies the algorithm in the scenario of two interfering links with capacity proportions $\frac{3}{4}$ and $\frac{1}{4}$ for links 1 and 2, respectively. As the capacity proportions are fixed, we need to find the largest multiplier that can be supported by the link transmission probabilities. For each link we may compute its own multiplier locally, but the joint multiplier needs to be found iteratively. The diagram on the left shows the link specific multipliers in the first iteration as a function of a step size to the direction (defined in the transmission probability space) where the multipliers grow at the same pace. The two phases of the algorithm are clearly visible from the figure on the right. Every first step is a step towards increasing all the multipliers (cf., the figure on the left) and every second step sets equalizes the multipliers (the boundary where the multipliers are equal is shown as a dashed line in the figure).

3.2.2 Flow-level throughput analysis

The flow-level throughput analysis requires cross-layer optimization of the link transmission probabilities. In certain simple cases, the balanced fairness analysis leads to an analytical formula for the throughput as a function of traffic load. This happens for example in the single resource example discussed earlier, see Figure 4, where the throughput of class 1, i.e., the link from node 1 to "base station" in Figure 2 (left), is illustrated.

Other systems that lead to analytical results are linear networks with end-to-end flows and a single resource system that has a large number of

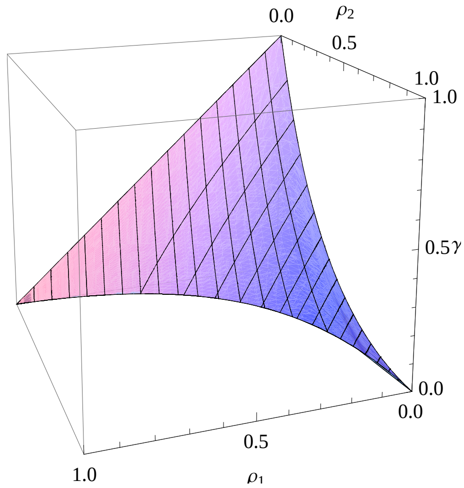

Figure 4: Throughput γ of class 1 as a function of class-1 and class-2 loads, ρ_1 and ρ_2, respectively.

users. In general, one has to resort to the numerical application of the algorithm described in the previous section. This allows performance analysis of arbitrary wireless networks. In order to demonstrate the asymptotic analysis, we study the network illustrated in Figure 5. The network has 30 nodes and 184 links. We assume 6 traffic classes and denote the source and destination nodes of class-i by s_i and d_i, respectively.

Assuming equal load ρ on all traffic classes, asymptotic analysis is used to approximate the flow level performance. Figure 6 illustrates the throughputs of the first three classes together with their performance if the networks was using centralized link scheduling instead of random access.

4 Data network dimensioning

The performance analysis research optimized the flow level performance for given traffic load by optimizing the access layer parameters. The dimensioning setting takes one step further, it attempts to define the nominal capacity of the links, i.e., the data transfer rate when the link is scheduled to be active.

In the simplest case a wireless communication link uses only one transmission rate. What this transmission rate should be that the flow level file

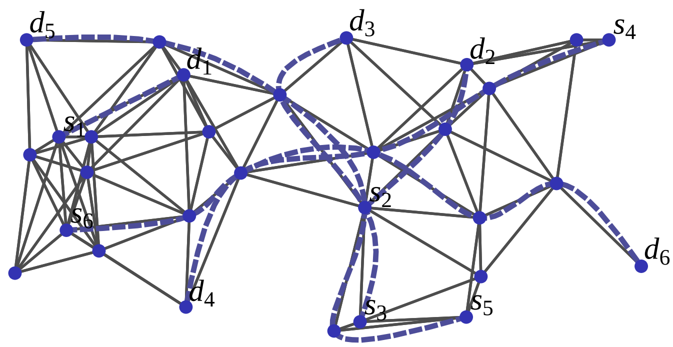

Figure 5: Wireless network with 30 nodes.

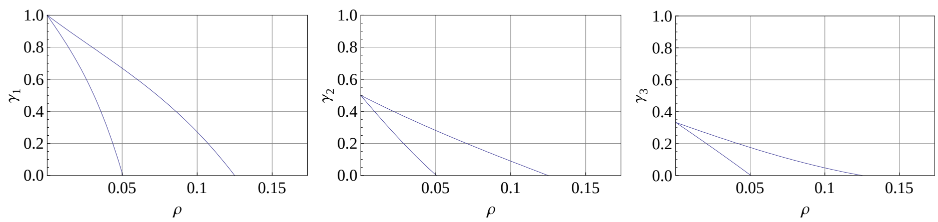

(a) Class 1 (b) Class 2 (c) Class 3

Figure 6: Asymptotic analysis of the first three traffic classes. The upper curves correspond to coordinated MAC and the lower curves to Aloha.

transfer rates would meet given performance requirements was addressed in [16].

However, the wireless network dimensioning problem spans over too many layers to make it directly tractable. It has the physical layer, where the transmission rate was defined, then there is the MAC-layer, where the links are scheduled, and, finally, there are the flow-level sharing effects of the realized capacity. Hence, approximate dimensioning methods were developed in the project. In particular, one of the best approximation method was based on the idea that the dimensioning problem for given traffic was solved as if the network was fixed. So each link was allocated a capacity. In the second step the scheduling problem was solved to determine the required transmission rate. This approach put the fixed network problem in focus, which is also an important problem by itself. The problem was addressed in detail in [17].

Finally, the dimensioning work was collected in to a journal article [18] with certain extensions, e.g., to include access rate limitations. In this work utilized also ns-2 [23] simulations to validate the proposed methods. The proposed methods provide a simple-to-use approach to obtain first estimates of the required capacities to obtain sufficient file transfer performance in the network.

As a spin-off result from the dimensioning work, an algorithm that was developed for dimensioning was applied in a fair resource allocation problem. The results were reported in [24].

The dimensioning work was carried out in co-operation with the Tekes-funded project Algorithms for Broadband Infrastructure (ABI).

4.1 Wireless networks

In wireless networks the dimensioning task is to determine a minimizing nominal capacity of the radio that is sufficient to have a pre-defined performance level under a given traffic load. This problem can be viewed as a first approximation to dimensioning, where the objective is to get a rough idea of how much traffic handling capacity the network needs.

The problem is challenging as it requires simultaneous optimization over MAC-layer scheduling and physical layer radio parameters. In the project several approximative methods were developed. These included:

- Wireless LB method which utilized a upper bound expression for the throughput under balanced fairness. The method readily defines the

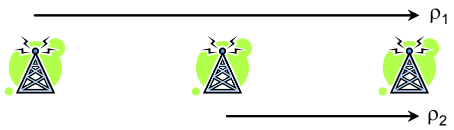

Method		Value (Mbps)
WLB	b_{LP}	19.0
WSF	b_{FS}	21.7
	b_{FC}	21.0

Figure 7: Illustration of the two-link WMN example and the radio transmissions rate determined using different approximation methods, when route 1 has the load $\rho_1 = 5$ Mbps and route 2 has the load $\rho_2 = 7$ Mbps and both the classes have the throughput requirement 1 Mbps

required link capacities with renders the solution of the radio parameter to a linear programming problem.

- Wireless SF method which utilized a conservative lower bound for the throughput under balanced fairness. As the bound is more complex than the one in the wireless LB method, we proposed two approximative method for this approach.
  - *Fixed schedule approach* solves the scheduling problem first which results in certain link capacities. Now the link capacities can be scaled up by increasing the radio rate parameter until the throughput attains the desired level.
  - *Fixed capacity approach* solves the dimensioning problem first *as if* the wireless network was fixed. This results in fixed link capacities and the radio rate parameter can be solved by an LP problem the same way as in the wireless LB method.

Figure 7 illustrates the dimensioning setting in a simple wireless mesh network. The radios are assumed to interfere so that the two links cannot be active simultaneously. The dimensioning task is to determine the sufficient radio capacity when we want that the file transfers are carried out (on average) with the rate 1 Mbps.

4.2 Wireline networks

Dimensioning of fixed networks is an essential subproblem in dimensioning of wireless networks (the fixed capacity approach), but it is an important problem also by its own right. In this case the dimensioning problem for data traffic is to determine the link capacities of the given network topology with a given traffic load and routing so that the average per-flow throughput experienced by users remains at acceptable levels. The link capacities are assigned different costs the sum of which are minimized. This topic has been addressed in [17] and [18].

As in the wireless case we utilize the throughput bounds of balanced fairness derived by Bonald and Proutière [7] to establish the connection between link capacities, throughput and traffic load and separate two methods for

- LB method is based on the upper bound on the throughput (under balanced fairness) and it gives a lower bound for required link capacities by stating that the capacity of a link must be at least equal to the load plus the maximum throughput requirement on that link.

- SF method is based on the so called store-and-forward bound that gives a lower bound on the throughput and hence an optimistic estimate on the required link capacity.

Whereas the two methods above address dimensioning when a throughput demand is specified for each route separately, there is also an alternative way to allocate capacity in the network by minimizing the capacity cost subject to a constraint on *average* file throughput [17].

For the SF method we developed a simple iterative algorithm that solves the problem. This method can also be extended to an improved lower bound derived in [3].

Next we consider a larger network example with the flow classes (routes) as shown in Figure 8. The network consists of 2 gateway nodes and there are 18 flow classes (routes) in the network. We assume that the offered load $\rho = 1$ Mbps for all routes.

The effects of access rate limitations are illustrated in Figure 9 where the access rate limit is 1 Mbps for all classes. The sum of excess capacities (allocated capacity minus the load on a link) is shown here as a function of the throughput requirement γ. Here we also show the results for the ISF method that is based on an improved lower bound for throughput.

4.3 Resource allocation

The algorithm which was developed for the SF dimensioning problem was a novel approach to certain convex optimization problems where the variables are bounded away from zero. This inspired applying the algorithm in the basic resource allocation problem where a number of resources are shared by contending users and the task is to allocate the resources fairly.

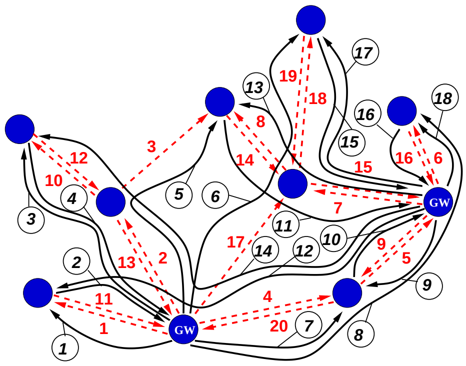

Link	LB	SF	Link	LB	SF
1	3.0	3.9	11	2.0	2.0
2	4.0	7.0	12	2.0	3.0
3	2.0	2.3	13	2.0	3.0
4	3.0	5.0	14	2.0	2.7
5	3.0	5.0	15	3.0	4.4
6	3.0	5.0	16	2.0	2.0
7	3.0	4.0	17	2.0	3.0
8	2.0	3.0	18	3.0	4.0
9	4.0	7.2	19	2.0	2.7
10	3.0	5.7	20	3.0	6.2

Figure 8: Example topology with two gateways and 18 flow classes (shown as arrows) and the resulting link capacities (in Mbps) from the different dimensioning methods

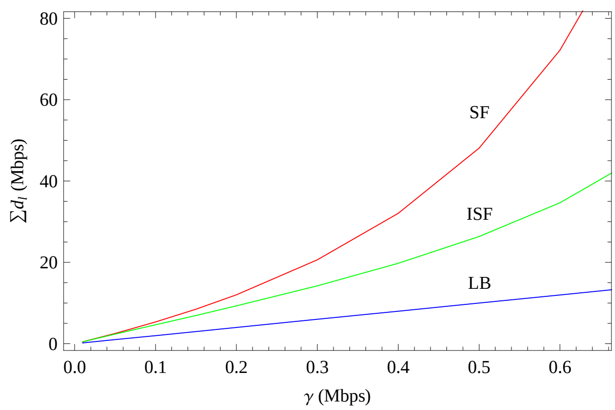

Figure 9: Sum of excess link capacities (the difference between the dimensioned capacity and the traffic load on each link), $\sum_l d_l$, as a function of γ with per-flow access rate limit $\Delta_i = 1$ Mbps, for all i. From top to bottom the curves are: SF method, ISF method and LB method.

Mo and Walrand [22] have presented a generalized fairness notion, so-called α-fairness, which represents a large family of fair allocations parameterized by the parameter α. As a special cases α-fairness includes several well-known allocations:

- With $\alpha = 0$ the allocation corresponds to maximization of the sum of the shares, i.e., utilizing "no" fairness.
- With $\alpha = 1$ the allocation corresponds to maximizing the product of the shares, i.e., proportional fairness [14].
- With $\alpha \to \infty$ the allocation approaches max-min fairness, which represents the most fair allocation.

Naturally also the parameter values between those well-known allocations are possible.

In [24] we derived a simple algorithm for determining the α-fair allocation for $\alpha > 1/2$. The optimization algorithm was presented in the context of fixed networks, where the resources are links with certain capacities and the users are end-to-end flows. However, the algorithm is readily applicable in any networks where the capacity set can be represented by linear constraints. This model can be directly applied also to a variety of wireless networks [4].

The idea of the algorithm is to use a simple iteration for the dual variables that will converge to values that satisfy the KKT optimality conditions [2]. Figure 10 presents an example with three links with capacities 3 Mbps, 2 Mbps and 1 Mbps, respectively. Flow "a" uses all the links, flow "b" uses only the link with capacity 3 Mbps, and flow "c" uses only the link with capacity 2 Mbps.

5 Theoretical work on scheduling problems

Many of the above problems incorporate a subproblem of centralized link scheduling. In co-operation with the Helsinki Institute for Information Technology HIIT the link scheduling problem was addressed from a theoretical viewpoint. By studying the conflict graph of the network, where each link is represented a vertex and edges correspond to conflict between the links, we can solve the associated scheduling problem if we are able to identify the maximum weight independent set of the conflict graph. However, this problem is difficult and [10] have put forth the question of whether there is a family of conflict graphs which arises in realistic network

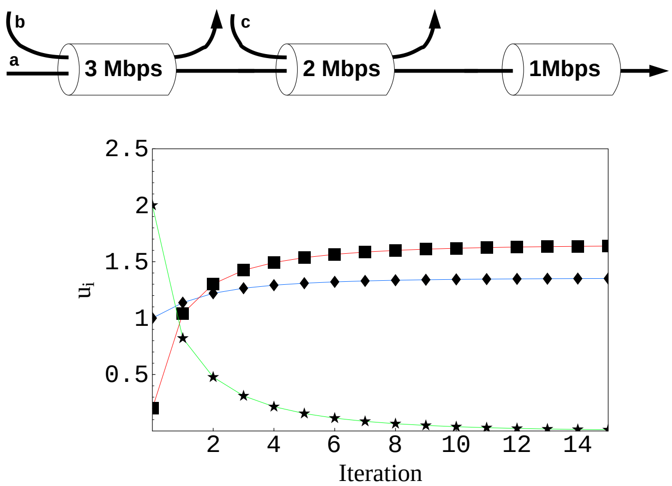

Figure 10: Determining the proportionally fair allocation ($\alpha = 1$) in the example network. Values of the dual variables u_1 (diamonds), u_2 (boxes), and u_3 (stars) in the iteration which starts from an arbitrary starting values, but quickly converges in $u^* = (1.354, 1.646, 0)$. This corresponds the resource allocation 0.785 Mbps, 2.215 Mbps, 1.215 Mbps, for flows "a", "b" and "c", respectively, which is proportionally fair.

deployments and which makes the problem of finding an independent set easier.

In [11] and its extended journal version [12] we study what assumptions must be made on the underlying network in order approximate the above problems in polynomial time. We show that if we assume a bounded density of the devices and a bounded range of interference we are able to obtain a constant factor approximation ratio to the problems, but they are not enough to achieve arbitrarily small approximation ratio unless, e.g., a bounded range of radio transmission is assumed.

6 Load balancing with link scheduling

Familiar to fixed networks, the term load balancing refers to optimization of usage of network resources by moving traffic from congested links to less loaded parts of the network based on knowledge of network state. By this approach QoS experienced by the users, such as transmission delay, is improved. There are many load balancing algorithms proposed for IP networks [29], especially along development of new tunnelling technique MPLS. However, these algorithms are not directly usable in wireless networks as the scheduling must also be addressed.

In [28] the load balancing problem was considered both with unconstrained paths and with predefined paths. The resulting cross-layer optimization task attempts to minimize the maximum link utilization in the network when one jointly optimizes the traffic allocation and the scheduling. It was shown that even the joint optimization task is still a linear programming problem and can be solved for relatively large problem instances.

For a simple example, consider a 2x2 grid network with 4 transmission modes as illustrated in Figure 11. There are two units of traffic; one from node 1 to node 4 and one from node 4 to node 3. Three cross-layer optimization approaches are considered which are also shown in the figure (from left to right):

- Traffic is routed via shortest paths but the scheduling is optimized. The maximum link load is 1.
- Traffic is routed using a load balancing routing. Scheduling is optimized. The maximum link load is 0.78.
- Routing and scheduling are jointly optimized to minimize the maximum link load. The maximum link load is 0.75.

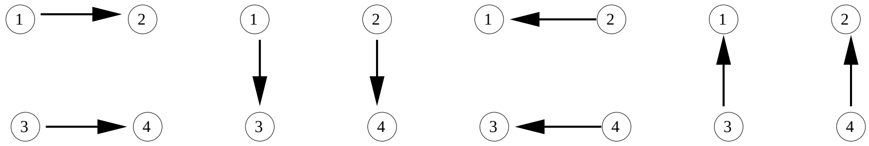

(a) Transmission modes

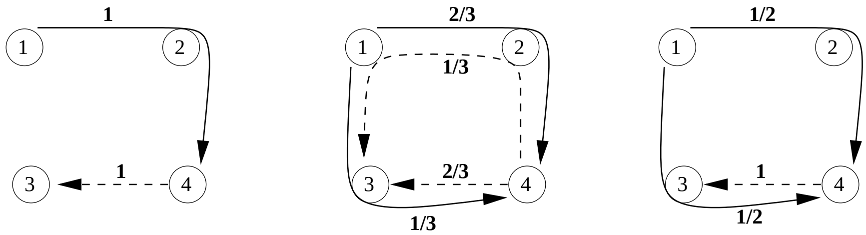

(b) Traffic allocations: Shortest paths with optimized scheduling, load balancing routing with optimized scheduling, joint optimization of routing and scheduling

Figure 11: Load balancing in a 2×2 grid network.

7 Results and impact

The CLOWN project aimed at (a) expanding the understanding of cross-layer optimization in wireless networks, and (b) apply the flow-level modeling approach to the context, and (c) provide algorithmic solution methods to the problems. In achieving these objectives the project was very successful. The scientific output of the project was significant both in quality and in quantity. The project contributed to 4 journal articles, 6 international conference articles and two other documents. The project also contributed to one PhD thesis.

As a practical results the project produced several algorithms that can be directly used for performance analysis and dimensioning of wireless multi-hop networks such as mesh networks.

The project fostered domestic scientific collaboration. Joint meetings and articles were prepared for example with the ABI project (of Tekes GIGA programme) and Fancy (funded by the Academy of Finland). In addition, research co-operation was carried out with the Helsinki Institute of Information Technology HIIT. Industrial and governmental contacts were maintained through the steering group which had members from Nokia Siemens Networks and Finnish Defence Forces Technical Research Centre.

In international collaboration the project personnel participated in the

COST 290 action entitled Traffic and QoS Management in Wireless Multimedia Networks, which had participating universities from 24 different European countries.

In addition to the contributions in methodology cross-layer optimization in wireless networks and the related algorithm development, the results of the project provide a good starting point of further development in wireless communications. Relevant issues have arisen for example in the ICT SHOK Future Internet Program.

References

[1] N. Abramson. The Aloha system - another alternative for computer communications. In Proceedings of Fall Joint Computer Conference, AFIPS Conference, page 37, 1970.

[2] M. Bazaraa, H. Sherali, and C.M. Shetty. Nonlinear programming: theory and algorithms. Wiley, second edition, 1993.

[3] T. Bonald. Throughput performance in networks with linear constraints. In *Proceedings of 40th Annual Conference on Information Sciences and Systems (CISS)*, March 2006.

[4] T. Bonald, L. Massoulié, A. Proutière, and J. Virtamo. A queueing analysis of max-min fairness, proportional fairness and balanced fairness. Queueing Systems, 53:65–84, 2006.

[5] T. Bonald, A. Penttinen and J. Virtamo. On Light and Heavy Traffic Approximations of Balanced Fairness. In Proceedings of SIGMETRICS/Performance 2006, pp. 109-120, Jun. 2006.

[6] T. Bonald, and A. Proutière. Insensitive bandwidth sharing in data networks, Queuing Systems, vol. 44, pp. 69-100, 2003.

[7] T. Bonald and A. Proutière. On performance bounds for balanced fairness. *Performance Evaluation*, 55:25–50, 2004.

[8] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle. Layering as optimization decomposition: A mathematical theory of network architectures. Proceedings of the IEEE, vol. 95, no. 1, pp. 255-312, Jan. 2007.

[9] T. Huovila, P. Lassila, J. Manner and A. Penttinen. State of the Art Analysis of Wireless Mesh Technologies. Helsinki University of Technology, 2006.

[10] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. Wireless Networks, 11(4):471487, 2005.

[11] P. Kaski, A. Penttinen and J. Suomela. Coordinating Concurrent Transmissions: A Constant-Factor Approximation of Maximum Weight Independent Set in Local Conflict Graphs. In Proceedings of the 6th International Conference on AD-HOC Networks & Wireless, AdHoc-Now 2007, pp. 74-86, 2007.

[12] P. Kaski, A. Penttinen and J. Suomela. Coordinating Concurrent Transmissions: A Constant-Factor Approximation of Maximum Weight Independent Set in Local Conflict Graphs. Ad Hoc & Sensor Wireless Networks: An International Journal, 2008, to appear.

[13] V. Kawadia and P.R. Kumar. A cautionary perspective on cross-layer design. IEEE Wireless Communications, 12(1):311, Feb. 2005.

[14] F. Kelly. Charging and rate control for elastic traffic. European Transactions on Telecommunications, 8:33–37, 1997.

[15] L. Kleinrock. Queueing systems, vol II: computer applications. John Wiley & Sons, first edition, 1976.

[16] P. Lassila, A. Penttinen and J. Virtamo. Dimensioning of wireless mesh networks with flow-level QoS requirements. In The 3rd ACM International Workshop on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN 2006), 2006.

[17] P. Lassila, A. Penttinen and J. Virtamo. Dimensioning methods for data networks with flow-level QoS requirements. In Proceedings of 32nd IEEE Conference on Local Computer Networks (LCN), pp. 353-359, 2007.

[18] P. Lassila, A. Penttinen and J. Virtamo. Dimensioning of data networks: a flow-level perspective. European Transactions on Telecommunications, 2008, to appear.

[19] J. Leino, A. Penttinen and J. Virtamo. Approximating Flow Throughput in Complex Data Networks. In Proceedings of 20th International Teletraffic Congress (ITC-20), pp. 422-433, 2007.

[20] J. Leino, A. Penttinen and J. Virtamo. Flow-optimized random access for wireless multihop networks. In Proceedings of ACM MSWiM 2007, pp. 387-394, 2007.

[21] J. Leino, A. Penttinen and J. Virtamo. Flow-optimized random access for wireless multihop networks. Performance Evaluation, 2008, to appear.

[22] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. IEEE/ACM Transactions on Networking, 8(5):556–567, October 2000.

[23] `http://nsnam.isi.edu/nsnam/index.php/Main_Page`

[24] A. Penttinen. Computation of fair flow rates in data networks. IEEE Communications Letters, vol. 12, no. 1, pp. 53-55, 2008

[25] A. Penttinen, J. Virtamo and R. Jäntti. Performance analysis in multihop radio networks with balanced fair resource sharing. Telecommunications Systems, vol. 31, no. 4, pp. 315-336, 2006.

[26] S. Shakkottai, T.S. Rappaport, and P.C. Karlsson. Cross-layer design for wireless networks. IEEE Communications Magazine, 41(10):7480, Oct. 2003.

[27] V. Srivastava and M. Motani. Cross-layer design: A survey and the road ahead. IEEE Communications Magazine, 43(12):112119, Dec. 2005.

[28] R. Susitaival. Load Balancing by Joint Optimization of Routing and Scheduling in Wireless Mesh Networks. In Proceedings of ITC-20, 2007

[29] Y. Wang, Z. Wang, and L. Zhang. Internet Traffic Engineering without Full Mesh Overlaying. In Proceeding of IEEE INFOCOM, 2001.