This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE Globecom 2011 proceedings.

Coverage Optimization in DSL Networks by Low-Complexity Discrete Spectrum Balancing

Martin Wolkerstorfer*,
* FTW Telecommunications Research Center Vienna,
Donau-City-Straße 1, A-1220 Vienna, Austria
Emails:{wolkerstorfer, nordstrom}@ftw.at

Tomas Nordström*†
† Centre for Research on Embedded Systems (CERES),
Halmstad University, Box 823, SE-30118 Halmstad, Sweden
Email: Tomas.Nordstrom@hh.se

***Abstract*—Spectrum balancing is an established optimization approach in multi-carrier digital subscriber line (DSL) systems. It has previously been applied to very different performance objectives such as sum-rate, min-rate, or fairness maximization and sum-power minimization. In this work we study the maximization of the service coverage, which will be defined as the number of DSL lines which can be granted an operator-specified high-bandwidth service. The proposed algorithm is based on a previously described mathematical decomposition framework. We extend this framework for our new problem and enhance its scalability by various low-complexity heuristics. Simulations demonstrate the applicability of our algorithm for DSL networks of realistic sizes. More precisely, our results obtained in thousand 25 user near-far DSL scenarios show an average gain in service coverage of more than $13\,\%$ compared to state-of-the-art sum-rate maximizing spectrum balancing algorithms.**

I. Introduction

Dynamic spectrum management (DSM) [1] comprises various techniques for interference mitigation in discrete multi-tone (DMT) based digital subscriber line (DSL) networks. Focussing on spectrum balancing (DSM level 2) we deal with a discrete and non-convex multi-carrier and multi-user power allocation problem. We introduce the maximization of service coverage as a novel objective for DSM and propose a scalable algorithm for its solution. Service coverage will be defined as the number of DSL lines which can be granted an operator-specified high bit-rate service. The network operators think in terms of the number of customers they can offer a particular service, which is better reflected by this objective than by previous DSM formulations. Current objectives for DSM in the literature range from sum-rate [2], [3] and min-rate (or reach) maximization [3], to sum-power minimization [4] and general utility (e.g., fairness [5] or hardware power consumption [6]) optimization. Spectrum balancing algorithms applicable for a large number of users are currently based on complexity reduction techniques such as constructive greedy heuristics [7], [8], continuous and/or (sequential) convex relaxation [2], [6], [9], Lagrange relaxation (LR) [9], [10], and sequential power updates over users [2], [9]. The unique features of the LR-based DSM framework in [10] are its robustness with respect to the initialization and the single-carrier solutions' suboptimality, as well as the separate storage of the single-carrier solutions. In the development of a practical algorithm for coverage maximization (CM) we extend the framework in [10] in two ways: by various single-carrier heuristics which allow the scheme to scale with the number of users, and by a successive variable fixing heuristic which exploits features that are specific to the framework.

We introduce the CM problem in Section II and give background information on the DSM framework in [10]. Our CM heuristic is described in Section III, together with numerous simulation results investigating its suboptimality. We approach the involved combinatorial single-carrier problems in Section IV by first analyzing case studies of the two greedy heuristics in [10], followed by the proposal of two novel randomized local search heuristics and performance evaluations in numerous 30-user very-high-speed DSL (VDSL) scenarios. A warm-start local search heuristic for complexity reduction of the single-carrier problems is given in Section IV-E, where we also show results indicating that investing complexity on a few subcarriers by using randomized local search is sufficient. Results on the performance of our complete CM algorithm including the single-carrier heuristics compared to sum-rate maximization schemes in near-far DSL scenarios are given in Section V. Our conclusions are summarized in Section VI.

II. System and Optimization Model

We assume a far-end crosstalk limited DSL system based on discrete multi-tone (DMT) modulation with the index sets of users and subcarriers denoted by $\mathcal{U}=\{1,\ldots,U\}$ and $\mathcal{C}=\{1,\ldots,C\}$, respectively. We model the users' rates per DMT-symbol $\mathbf{r}^c(\mathbf{p}^c)\in\mathcal{R}^U$ on subcarrier $c\in\mathcal{C}$ by the commonly used gap approximation [11], where $\mathbf{p}^c\in\mathcal{R}^U$, and p_u^c denotes the power spectral density (PSD). The power allocation vector $\mathbf{p}^c(\mathbf{r}^c)$ for rates $\mathbf{r}^c$ is then given as the solution of a linear matrix equation, see [12] for details. To model constraints on the modulation's constellation size and regulatory PSD mask limitations we use the set of feasible per-subcarrier PSDs

$$\mathcal{Q}^c=\{\mathbf{p}^c|r_u^c(\mathbf{p}^c)\in\mathcal{B},p_u^c\in[0,\hat{p}_u^c],u\in\mathcal{U}\},c\in\mathcal{C},\qquad(1)$$

where $\hat{p}_u^c, c\in\mathcal{C}, u\in\mathcal{U}$, indicates the PSD mask, and $\mathcal{B}=\{0,\Delta,2\,\Delta,\ldots,\hat{B}\}$ is the set of discrete per-subcarrier bit allocations with rate-steps of size Δ and a bit cap $\hat{B}$.

This work has been supported in parts by the Austrian Government and the City of Vienna within the competence center program COMET.

A. A Problem Formulation for Coverage Maximization (CM)

Our optimization problem is to guarantee a basic target-rate for all users while obeying all system constraints and maximizing the number of users which can be upgraded to a high-bandwidth service. Denoting the decision variables for service upgrades by $\delta_u \in \{0,1\}, u \in \mathcal{U}$, we formulate this coverage maximization (CM) problem for DSL as

$$\underset{\substack{\delta_u \in \{0,1\},\, u \in \mathcal{U},\\ \mathbf{p}^c \in \mathcal{Q}^c,\, c \in \mathcal{C}}}{\text{maximize}} \quad \sum_{u \in \mathcal{U}} \delta_u \tag{2a}$$

$$\text{s.t.} \quad \sum_{c \in \mathcal{C}} r_u^c(\mathbf{p}^c) \geq R_u + \delta_u(\hat{R} - R_u), \quad \forall u \in \mathcal{U}, \tag{2b}$$

$$\sum_{c \in \mathcal{C}} p_u^c \leq P_u, \quad \forall u \in \mathcal{U}, \tag{2c}$$

where $\hat{R} \in \mathcal{R}_+$ is the target-rate for the service upgrade, and where $\mathbf{R} \in \mathcal{R}_+^U$ and $\mathbf{P} \in \mathcal{R}_+^U$ are the users' minimum target-rates and maximum sum-powers, respectively. As the variables $\delta_u, u \in \mathcal{U}$, are binary the objective in (2a) reflects the number of high-bandwidth users, i.e., the service coverage. We proceed by studying the Lagrange relaxation of the problem in (2) for continuously relaxed binary variables, i.e., $\delta_u \in [0,1], u \in \mathcal{U}$, and will return to a feasible binary solution in Section III.

B. Background Information on Lagrange Relaxation (LR)

LR [13] has emerged as a scalable tool for the design of DSM algorithms [12], [14]. It results in a decomposition of the problem in (2) into separable and non-convex combinatorial subproblems on subcarriers $c \in \mathcal{C}$ given by [10], [12]

$$\underset{\{\mathbf{r}^c | \mathbf{p}^c(\mathbf{r}^c) \in \mathcal{Q}^c\}}{\text{minimize}} \quad f^c(\mathbf{r}^c) = \boldsymbol{\nu}^\intercal \mathbf{p}^c(\mathbf{r}^c) - \boldsymbol{\lambda}^\intercal \mathbf{r}^c, \tag{3}$$

where $\boldsymbol{\lambda}, \boldsymbol{\nu} \in \mathcal{R}_+^U$ are the Lagrange multipliers for the relaxed constraints (2b) and (2c), respectively. Nondifferentiable optimization methods [13] are typically used to adjust the multipliers and enforce the relaxed constraints. The proposed heuristics in Section IV target the subproblems in (3) and are independent of the specific mechanism generating the multipliers. Differently, the CM heuristic in Section III builds upon a column generation [15] based approach for optimizing $\boldsymbol{\lambda}$ and $\boldsymbol{\nu}$. This iterative approach keeps at iteration t a memory of the feasible per-subcarrier solutions $\mathbf{p}^{c,i} \in \mathcal{Q}^c, i \in \mathcal{I}^{c,(t)}$, it has obtained so far, indexed by the set $\mathcal{I}^{c,(t)}, c \in \mathcal{C}$. In each iteration it then computes a convex combination of the corresponding per-subcarrier powers and rates by adjusting weights $\xi^{c,i}$ in a linear problem (LP) of the form, cf. (2),

$$\underset{\substack{\delta_u \in [0,1],\, u \in \mathcal{U},\\ \xi^{c,i} \geq 0,\, i \in \mathcal{I}^{c,(t)},\, c \in \mathcal{C}}}{\text{maximize}} \quad \sum_{u \in \mathcal{U}} \delta_u \tag{4a}$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{I}^{c,(t)}} \xi^{c,i} = 1, \qquad \forall c \in \mathcal{C}, \tag{4b}$$

$$\sum_{c \in \mathcal{C}} \sum_{i \in \mathcal{I}^{c,(t)}} r_u^c(\mathbf{p}^{c,i}) \xi^{c,i} \geq R_u + \delta_u(\hat{R} - R_u), \ \forall u \in \mathcal{U}, \tag{4c}$$

$$\sum_{c \in \mathcal{C}} \sum_{i \in \mathcal{I}^{c,(t)}} p_u^{c,i} \xi^{c,i} \leq P_u, \quad \forall u \in \mathcal{U}. \tag{4d}$$

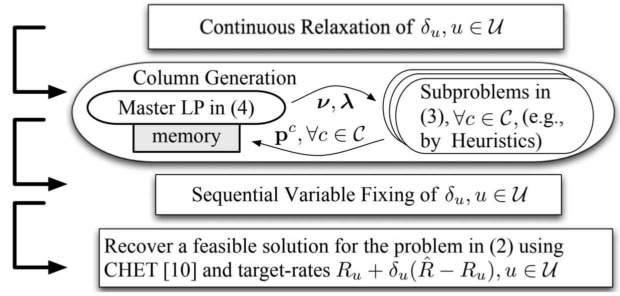

Fig. 1. Illustration of the proposed coverage maximization heuristic.

In column generation the Lagrange multipliers $\boldsymbol{\lambda}$ and $\boldsymbol{\nu}$ are then computed as the optimal dual variables associated with the constraints in (4c) and (4d), respectively. The CM heuristic presented next exploits two features of column generation: the information in the relaxed variables $\delta_u, u \in \mathcal{U}$, and the memory of the produced allocations $\mathbf{p}^{c,i}, i \in \mathcal{I}^{c,(t)}, c \in \mathcal{C}$.

III. A Heuristic For Coverage Maximization

In the above sections we considered the column generation based relaxation in (4) of the original problem in (2) with continuously relaxed variables $\delta_u \in [0,1], u \in \mathcal{U}$. Our next step is a heuristic for sequentially fixing these variables to either zero or one: We simply try setting the lower-bound of the largest fractional variable $\delta_u, u = \text{argmax}_{i \in \mathcal{U}}\{\delta_i\}$, to one and re-solve the master LP in (4) from the final iteration $t = \hat{t}$. If the LP is feasible we keep the lower-bound of δ_u at one or reset it to zero otherwise, cf. [10] for a simple feasibility detection mechanism. This process is repeated until all variables $\delta_u, u \in \mathcal{U}$, are fixed at integer values. Next we solve yet another, similar LP with the constraints as above but with lower-bounds of binary variables now set to the found integer values and the objective being the minimization of the total transmit power. This is to avoid unnecessarily wasting power while maintaining the found level of service coverage. Finally we apply the column combination heuristic for time-sharing (CHET) in [10] to recover from the column generation based relaxation and to find a feasible solution for our original problem in (2), cf. the schematic overview in Figure 1. Summarizing, the key ideas of our CM heuristic are a) to exploit the information in the optimized continuous variables $\delta_u, u \in \mathcal{U}$, in order to prioritize users in the sequential fixing process, and b) to use the memory in column generation based schemes for checking feasibility of specific 0-1 fixings. As a result we obtain a low-complexity method for the sequential fixing of continuous variables $\delta_u, u \in \mathcal{U}$, where we only solve in total $U + 1$ LP's as in (4).

We will now study the suboptimality of our CM heuristic by comparison to an exhaustive search over all 0-1 combinations of binary variables $\delta_u, u \in \mathcal{U}$. We randomly generate 1000 4-user DSL scenarios as described in Section V, with 3 and 1 lines deployed from central office and cabinet, respectively, with simulation parameters as detailed in Section IV-D, and using the mentioned column generation framework with optimal solutions [12] of the per-subcarrier problems in (3). We

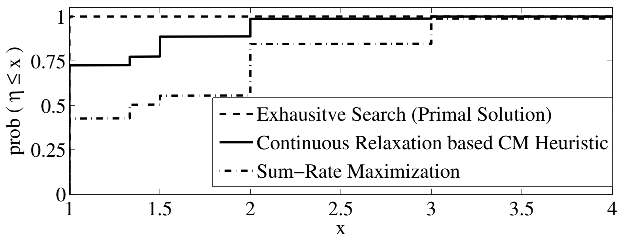

Fig. 2. Cumulative distribution function of the performance ratio η in (5).

define an algorithm's performance ratio η between the optimal objective P^* in (2) and the algorithm's coverage P^{ALG} through a feasible solution in (2) as

$$\eta = \frac{P^*}{P^{\mathrm{ALG}}}. \tag{5}$$

Figure 2 shows cumulative distributions of the ratio η for our CM heuristic, the primal solution obtained after the exhaustive search through CHET [10], and the result of a sum-rate maximization. Regarding the dashed curve we see that after the exhaustive search a feasible primal solution was obtained in all cases (using CHET), assuring that we compare to the *primal* optimum P^*. We find that sum-rate maximization and the CM heuristic are optimal in less than $47\,\%$ and more than $68\,\%$, respectively, with $99\,\%$ confidence according to a t-test. Further results and an intuition behind the good performance of the applied continuous relaxation will be developed in Section V. We recognize that other relaxations such as a semidefinite programming relaxation [16] could be used instead and we plan related studies in future work.

In the following sections we study the solution of the subproblem in (3) and therefore omit subcarrier indices c for ease of notation. We search over the set of discrete rate allocations $\times_{u\in\mathcal{U}}\mathcal{B}$, considering the constraints in (3) only implicitly through $f(\mathbf{r}) = \infty, \forall \mathbf{r} \notin \{\tilde{\mathbf{r}} | \mathbf{p}(\tilde{\mathbf{r}}) \in \mathcal{Q}\}$. We will use the number of solved linear matrix equations [12] needed for the evaluation of the objective $f(\mathbf{r})$ as a reproducible complexity metric, cf. Section IV-D.

IV. Selected Single-Carrier Heuristics

The number of subproblems C as in (3) which has to be solved is in the order of thousands in the latest DSL technologies [17]. Moreover, these C problems have to be solved in every iteration t of the column generation scheme. Finding the optimum of the problem in (3), although of polynomial complexity [10], was in general found to be intractable for practical values of U [10], motivating the application of two fast heuristics discussed in the next section. Both heuristics are extended by local search (LS) techniques, cf. [18] for details.

A. Constructive Base-Heuristics

Before proposing our novel heuristics for the subproblem in (3) we review two greedy schemes we proposed in [10] and analyze their performance. The constructive, joint greedy optimization scheme (JOGO) iteratively updates the rates $\mathbf{r}$, starting from $\mathbf{r} = \mathbf{0}$ and computing in each iteration the

TABLE I
Worst-case scenarios for the greedy heuristic JOGO in [10].

# Users with Line-Lengths of ...				Suboptimality [%] of sum-objective over subcarriers
200 m	400 m	600 m	800 m	
1	1	0	4	13.1
2	0	0	4	10.0
0	1	1	4	9.3
1	0	5	0	9.0
0	2	0	4	8.9

additional cost $f(\tilde{\mathbf{r}}) - f(\mathbf{r})$ of loading another Δ bits, i.e., $\tilde{r}_u = r_u + \Delta, \tilde{r}_i = r_i, \forall i \in \mathcal{U}\setminus\{u\}$, for any user $u \in \mathcal{U}$. Next, r_u is increased by Δ bits for the user u which minimizes this additional cost. In case of a draw we take the allocation with minimum sum-power. These steps are repeated until no update with negative additional cost can be found, cf. [10] for a complete description. In [10] a second, less complex sequential greedy optimization (SEGO) algorithm was proposed where users greedily set their own rates r_u one after the other, each minimizing the global objective $f(\mathbf{r})$. For the analysis of these two base heuristics we study a set of 84 6-user VDSL scenarios with possible line-lengths of $\{200, 400, 600, 800\}$ m, respectively, and a subset $\tilde{\mathcal{C}} \subset \mathcal{C}$ of all subcarriers, cf. Section IV-D for details on the simulation parameters. The maximum and average suboptimality of JOGO over subcarriers $\tilde{\mathcal{C}}$ and scenarios is $39.3\,\%$ and $3.7\,\%$, respectively. While the suboptimality values of JOGO were found to be zero for all collocated scenarios, the highest values appear in classical near-far type of scenarios, cf. Table I. As one example we pick the first scenario in Table I and a specific subcarrier at approximately $9.26\,$MHz. The greedy base heuristic JOGO assigns the bits $\mathbf{r} = [9, 4, 0, 0, 0, 0]$, i.e., only the shortest lines are transmitting. Compared to an optimal allocation $\mathbf{r}^* = [6, 4, 3, 2, 2, 2]$ JOGO has a suboptimality of more than $31\,\%$. Note that in this example SEGO gives the result $\mathbf{r} = [16, 0, 0, 0, 0, 0]$ when the user sequence starts with the shortest line, corresponding to a suboptimality of more than $15\,\%$. Also, having the two greedy algorithms followed by the local search in [18] does not improve the solutions beyond $15\,\%$ of suboptimality. We noted that any user sequence in SEGO which starts with one of the longer lines would have lead to an optimal result after a following local search initialized at the solution of SEGO. This motivates the extension of SEGO in Section IV-C where the user-sequence is optimized together with the bit-allocation. In Section IV-B we present a modification of JOGO which uses randomized greedy decisions and local search. These two heuristics presented next were seen to outperform those in this section as well as various other meta-heuristics [18].

B. A Greedy Randomized Adaptive Search Procedure (Grasp)

As the name indicates, Grasp [19] is a randomized extensions of greedy search schemes, cf. Algorithm 1. More precisely, Lines 3–11 mimic the joint greedy bit-loading of JOGO in [10] with the important difference that in Line 10 a randomized, potentially non-greedy decision is taken. The parameter β allows for a trade-off between standard greedy

Algorithm 1 Grasp based Bit-Loading

1: Initialize $\mathbf{r} = \mathbf{0}$, $\mathbf{r}^*$, $\delta^* = 0$, $f^{\text{prev}} = f(\mathbf{r})$, $\boldsymbol{\beta} \in \mathcal{R}^M$, $\beta_i \in [0,1], g_i = \check{f}, \forall i \in \mathcal{M} = \{1,\ldots,M\}$, $k = 1$, K
2: **for** $k = 1,\ldots,K$ **do**
3: **while** $\delta^{\min} \leq 0$ **do**
4: Set $\beta = \beta_{i^*}$, $\tilde{g}_i = \max_{i\in\mathcal{M}}\{g_i\} - g_i, \forall i \in \mathcal{M}$, where $i^* \in \mathcal{M}$ is sampled from the distribution

$$P_i = \frac{\tilde{g}_i}{\sum_{j\in\mathcal{M}} \tilde{g}_j},\ \forall i \in \mathcal{M}, \tag{6}$$

5: **for** $u = 1,\ldots,U$ **do**
6: **if** $\exists \mathbf{p} \in \mathcal{Q} | r_u(\mathbf{p}) = r_u + \Delta, r_i(\mathbf{p}) = r_i, i \in \mathcal{U}\setminus\{u\}$,
7: **then** $\delta_u = f(\mathbf{p}, \hat{\mathbf{w}} + \boldsymbol{\nu}, \check{\mathbf{w}} + \boldsymbol{\lambda}) - f^{\text{prev}}$
8: **else** $\delta_u = \infty$
9: $\delta^{\min} = \min_{u\in\mathcal{U}}\{\delta_u\}$, $\delta^{\max} = \max_{u\in\mathcal{U}}\{\delta_u\}$
10: Uniformly sample a user u^* from the set $\{u \in \mathcal{U} \mid \delta_u \leq (1-\beta)\delta^{\min} + \beta\delta^{\max}\}$
11: **if** $\delta^{\min} \leq 0$ **then** $r_{u^*} = r_{u^*} + \Delta$, $f^{\text{prev}} = f^{\text{prev}} + \delta_{u^*}$
12: Search a local optimum $\tilde{\mathbf{r}}$ starting at $\mathbf{r}$ as in [18]
13: Update incumbent $\mathbf{r}^* = \tilde{\mathbf{r}}$ if $f(\tilde{\mathbf{r}}) < f(\mathbf{r}^*)$
14: Update the average g_{i^*} with the newest value f^{prev}

Algorithm 2 rSEGO

1: Initialize $\mathbf{r}^*$, K, $\bar{K}$, values $\tilde{\tau}_u(i), \forall u, i \in \mathcal{U}$, q_0, β, ρ
2: **for** $k = 1,\ldots,K$ **do** $\{/*\text{iterations}*/\}$
3: **for** $\bar{k} = 1,\ldots,\bar{K}$ **do** $\{/*\text{ant runs}*/\}$
4: Set $\mathbf{r}^{(\bar{k})} = \mathbf{0}$, $\tilde{\mathcal{U}} = \mathcal{U}$, Uniformly sample $s_1^{(\bar{k})} \in \mathcal{U}$
5: **for** $i = 1,\ldots,U$ **do** $\{/*\text{sequential decisions}*/\}$
6: Set $u = s_i^{(\bar{k})}$, $f^{\min} = \min_{\{r_u^{(\bar{k})}\in\mathcal{B}|\mathbf{p}(\mathbf{r}^{(\bar{k})})\in\mathcal{Q}\}} \{f(\mathbf{r}^{(\bar{k})})\}$, $f^{\max} = \max_{\{r_u^{(\bar{k})}\in\mathcal{B}|\mathbf{p}(\mathbf{r}^{(\bar{k})})\in\mathcal{Q}\}} \{f(\mathbf{r}^{(\bar{k})})\}$
7: Uniformly sample $r_u^{(\bar{k})}$ from the set $\{r_u^{(\bar{k})} \in \mathcal{B} \mid \mathbf{p}(\mathbf{r}^{(\bar{k})}) \in \mathcal{Q}, f(\mathbf{r}^{(\bar{k})}) \leq (1-\beta)f^{\min} + \beta f^{\max}\}$
8: **if** $i < U$ **then**
9: Exclude u from $\tilde{\mathcal{U}}$, Uniformly sample $q \in [0,1]$
10: **if** $q < q_0$ **then** $\{s_{i+1}^{(\bar{k})} = \text{argmax}_{j\in\tilde{\mathcal{U}}}\{\tilde{\tau}_u(j)\}\}$
11: **else**
12: Sample $s_{i+1}^{(\bar{k})} = j \in \mathcal{U}$ from $P_u(j \mid \tilde{\mathcal{U}}) = 0, j \in \mathcal{U}\setminus\tilde{\mathcal{U}}$, $P_u(j \mid \tilde{\mathcal{U}}) = \frac{\tilde{\tau}_u(j)}{\sum_{l\in\tilde{\mathcal{U}}}\tilde{\tau}_u(l)}, j \in \tilde{\mathcal{U}}$,
13: Update $\tilde{\tau}_u(s_{i+1}^{(\bar{k})}) = (1-\rho)\cdot\tilde{\tau}_u(s_{i+1}^{(\bar{k})})$
14: Update $\mathbf{r}^{(\bar{k})}$ by local search [18] starting at $\mathbf{r}^{(\bar{k})}$
15: $\bar{k}^* = \underset{\bar{k}=1,\ldots,\bar{K}}{\text{argmin}}\{f(\mathbf{r}^{(\bar{k})})\}$, Update $\tilde{\tau}_{s_i^{(\bar{k}^*)}}(s_{i+1}^{(\bar{k}^*)})$ as in (7), $\forall i \in \{1,\ldots,U-1\}$, Set $\mathbf{r}^* = \mathbf{r}^{(\bar{k}^*)}$ if $f(\mathbf{r}^{(\bar{k}^*)}) < f(\mathbf{r}^*)$

schemes ($\beta = 0$) and purely random schemes ($\beta = 1$). However, just as in JOGO the greedy bit-loading process is considered complete when even a purely greedy decision would increase the solution objective, i.e., $\delta^{\min} > 0$, cf. Line 3. The free parameters in the presented basic implementation of Grasp are the local search strategy and neighborhood definition, cf. Line 12, the number of restarts K (where an alternative stopping criterion will be used in our simulations), and the randomization parameters $\beta_i, i \in \{1,\ldots,M\}$, where in each iteration of Grasp a single parameter β is chosen depending on the average objective value experienced under all parameters $\boldsymbol{\beta}$, cf. Lines 4 and 14 of Algorithm 1, respectively.

C. Randomized SEGO (rSEGO)

Next we present an improvement over SEGO [10] in the sense that the sequence with which the users are loaded is adaptive, cf. Algorithm 2. It combines elements of SEGO, Grasp, as well as ant colony system heuristics [18], [20]. More precisely, in Lines 9–13 of Algorithm 2 the currently active user u decides which user should allocate his bits next, where each possible decision $i \in \tilde{\mathcal{U}}$ is assigned a sequence decision value $\tilde{\tau}_u(i)$. The decision on the bit-rate r_u is made greedily in Line 7 similarly as in SEGO, where similarly as in Grasp the decision is randomized based on the objective values of all possible rate decisions. Finally, after $\bar{K}$ solutions have been constructed in this way, defining $\hat{f} = \sum_{u\in\mathcal{U}} \nu_u \hat{p}_u$ and using any $\rho \in [0,1]$, a global update of the decision values $\tilde{\tau}_u(j)$ is done in Line 15 as

$$\tilde{\tau}_u(j) = (1-\rho)\cdot\tilde{\tau}_u(j) + \rho\cdot\left(\hat{f} - f(\mathbf{r}^{(k^*)})\right). \tag{7}$$

D. Simulation Results for the Single-Carrier Heuristics

The parameters for our xDSL simulator available in [21] were selected according to the ETSI VDSL standard [17], with $\Gamma = 12.8\,\text{dB}$, $\hat{B} = 16$, $\Delta = 1$, upstream band plan 997-M1x-M, and noise N_c^u comprising ETSI VDSL noise A added to a flat background noise at $-140\,\text{dBm/Hz}$. We make the practical assumption that there is a restriction in simulation time for solving the subproblems in (3). Furthermore, in order to make our results reproducible for future research we will use the number of power evaluations $\mathbf{p}(\mathbf{r})$ by solving a linear matrix equation [12] as the stopping criterion for Grasp and rSEGO. This is further motivated by the almost identical simulation times of these heuristics as a function of the number of power evaluations. We will initialize the incumbent solution (but not the initial starting-point $\mathbf{r} = \mathbf{0}$) of all methods by the result of the base heuristic JOGO. As this heuristic is guaranteed to give a solution with negative objective value, we have that also all other heuristics will produce a solution $\mathbf{r}^*$ with negative objective value $f(\mathbf{r}^*)$. For comparing our single-carrier heuristics we choose a subset of subcarriers $\tilde{\mathcal{C}} = 1, 51, \ldots, 1601$, fixed Lagrange multipliers $\nu_u = 1/U, \lambda_u = 1, u \in \mathcal{U}$, and construct our network scenarios using a set of specified line lengths $\{200, 400, 600, 800\}\,\text{m}$, considering all U-combinations with repetitions [18]. The choice of parameters[1] is based on Monte-Carlo simulation in six-user networks where optimal solutions are known [10], cf. [18] for details. Comparing the two greedy randomized schemes rSEGO and Grasp, both, their performance metrics as well as the scenarios where their performance was suboptimal were found to be similar. Grasp is based on JOGO and we found that it takes a parameter $\beta = 1$

[1]For all randomized heuristics we chose a first-improving local search with two-step neighborhood as proposed in [18]. Parameters for rSEGO were $\beta = 0.75$, $\bar{K} = 5$, $\rho = 0.99$, and $q_0 = 0.2$, while for Grasp we use $\boldsymbol{\beta} = [0.75, 1]$.

TABLE II

IMPROVEMENTS BY RANDOMIZED HEURISTICS COMPARED TO JOGO.

Alg.	Mean obj. improv. [%] (per subcarr.)	[Min.,Max.] improv. [%] of sum-obj. over subcarr.	Mean op.-count [$\times 10^3$] per subcarr.[3]
JOGO	0	$[0, 0]$	1.79±0.06
JOGO+LS	8.03±0.78	$[0.91, 32.13]$	3.92±0.07
rSEGO$^{(10*)}$	8.90±0.00	$[0, 29.024]$	10*
rSEGO$^{(20*)}$	9.76±0.00	$[1.36, 21.89]$	20*
Grasp$^{(10*)}$	9.05±0.00	$[0, 29.02]$	10*
Grasp$^{(20*)}$	9.86±0.00	$[1.48, 21.89]$	20*

to remedy the shortcomings of JOGO in the example made in Section IV-A.[2] We chose to have two parameters $\boldsymbol{\beta} \in \mathcal{R}^2$ among which we select, cf. Line 4 in Algorithm 1.

Due to the large number of 30-user scenarios for the four selected line-lengths we uniformly sample 200 and 100 scenarios (not cable lengths!) for the deterministic and randomized algorithms, respectively, make 50 repetitions for randomized algorithms, and present mean values with 95 % confidence intervals according to a t-test. In order to reduce the variance of our results we ran all heuristics on identical, sampled scenarios. Furthermore, as we lack an optimal solution we use the greedy algorithm JOGO as our reference and show the improvements in objective value. Table II reports our average simulation results, where the values marked by a star reflect the applied stopping criterion (complexity budget) and exclude the incumbent initialization using JOGO. We find that using a combination of JOGO and local search (with a negligible difference whether we initialized LS at the solution of JOGO or at $\mathbf{r} = \mathbf{0}$, using JOGO solely to provide an incumbent solution) already gives improvements compared to the objective under JOGO of around 8 %, while the average improvements of our randomized algorithms are notably almost 10 % under the limit of $2 \cdot 10^3$ matrix inversions.

E. Warm-Start Local Search (WS-LS)

Despite the presented low-complexity heuristics, solving the single-carrier problem in (3) on all subcarriers $c \in \mathcal{C}$ leads to a large total complexity. We therefore propose a novel scheme where only a subset of all subcarriers $\tilde{\mathcal{C}} \subseteq \mathcal{C}$ is solved using one of the single-carrier heuristics. The approach WS-LS takes to obtain a reasonable starting point $\bar{\mathbf{r}}$ for local search on all other subcarriers $c \in \mathcal{C} \setminus \tilde{\mathcal{C}}$ is to use the heuristic solutions $\mathbf{r}^{\hat{c}}$ and $\mathbf{r}^{\check{c}}$ on the "neighboring" subcarriers $\hat{c}, \check{c} \in \tilde{\mathcal{C}}$. The conservative strategy we suggest to use is $\bar{r}_u = \min\{r_u^{\hat{c}}, r_u^{\check{c}}\}, \forall u \in \mathcal{U}$. Note that WS-LS is different to the idea of subcarrier grouping in wireless networks or the complexity reduction technique in [6] as we do not approximate or interpolate solutions on different subcarriers. The following results are based on a maximum of 10^3 power evaluations per subcarrier, a selection of subcarriers $\tilde{\mathcal{C}}$ in regular intervals, and other settings as specified in Section IV-D. As a base-line we added a randomized local search (rLS) scheme where the LS algorithm is reinitialized at random starting points uniformly drawn from the set of allocations feasible w.r.t. single-user bit cap and power mask constraints. Figure 3 depicts the average improvements in sum-objective by WS-LS and its complexity in all 6-user scenarios as described in Section IV-D compared to using JOGO on all subcarriers. Figure 3(a) shows that the performance gain by increasing $|\tilde{\mathcal{C}}|$ actually flattens out, with this effect happening for lower values of $|\tilde{\mathcal{C}}|$ using Grasp or rSEGO compared to rLS. The initial drop in complexity in Figure 3(b) is explained by the dependence of the local search complexity on its starting point. As an example, we find that using WS-LS with $|\tilde{\mathcal{C}}| = 21$ we can achieve a complexity reduction of more than 90 % compared to using rSEGO with given complexity limitation.

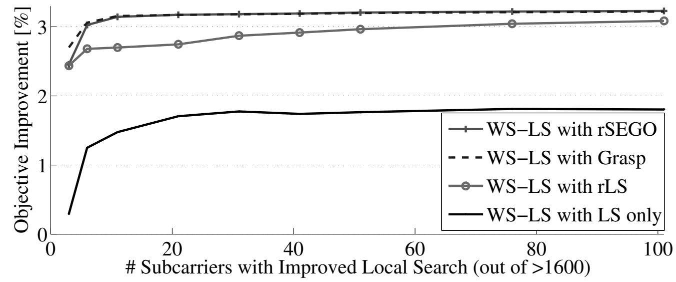

(a)

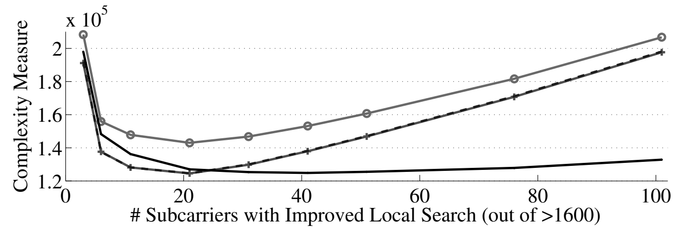

(b)

Fig. 3. Trade-off for heuristic WS-LS between average objective improvement compared to JOGO (a) and sum-complexity over all C subcarriers (b).

V. PERFORMANCE EVALUATION IN LARGE NETWORKS

We apply the coverage maximization (CM) algorithm in Figure 1 to the problem in (2), as well as the sum-rate maximization (RM) scheme in [10], both using the heuristics rSEGO with a complexity limit of $3 \cdot 10^4$ power evaluations per subcarrier, WS-LS with 40 regular subcarrier intervals shifted in every iteration t by a uniformly drawn random number of subcarriers between zero and one interval length, and other simulation parameters as specified in Section IV-D. We generated 1000 downstream VDSL [17] 25 user network topologies by uniformly sampling the lengths $L_{CO-Cab} \in [100, 1400]$ and $L_{Cab-u} \in [50, 500], \forall u \in \mathcal{U}$, cf. Figure 4, and assigning 10 users to the cabinet and 15 users to the central office. We compare our CM and RM schemes to optimal single-user bit-loading [22] under worst-case crosstalk noise $\sum_{i \in \mathcal{U} \setminus u} H_{ui}^c \hat{p}_i^c, c \in \mathcal{C}$, (MaskBL) and the sum-rate maximizing

[2]Lines are not identical in their transmission parameters in practice. Still, the example suggests β to be set to a large value in near-far scenarios.

[3]The numbers marked by * reflect the set complexity budget.

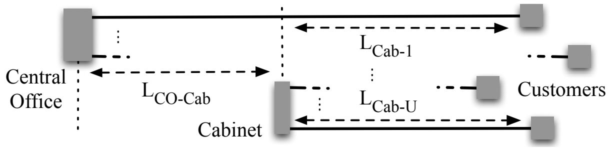

Fig. 4. Structure of the randomly generated DSL networks.

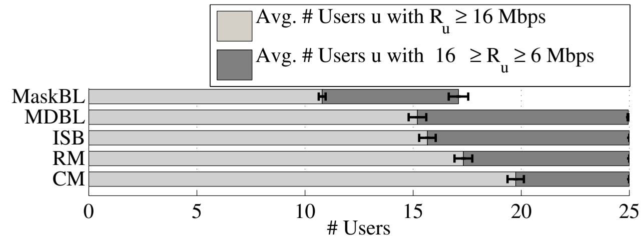

Fig. 5. Average service coverage in 1000 DSL scenarios.

algorithms in [7] (MDBL) and a scheme closely related to [14] (ISB).[4] Under RM and CM we assume minimum target-rates $R_u = 6\,\text{Mbps}, \forall u \in \mathcal{U}$. The upgrade target-rate is set to $\hat{R} = 16\,\text{Mbps}$. Our average simulation results are shown in Figure 5 where we observe a large gap in terms of the number of upgraded users between the four tested DSM algorithms and MaskBL. For example, CM improves upon MaskBL in this respect by on average more than $82.9\,\%$ (8.95 ± 0.41). Comparing CM to the best-performing rate-maximizing DSM algorithm RM we find that our CM scheme yields a $13.9\,\%$ $(2.42 \pm 0.32$ users) increase in service coverage. Intuitively, the improvements by our CM heuristic under the applied continuous relaxation of binary variables in Section II-A upon the sum-rate maximizing algorithms can be explained by the reduction in rates of the lines which would over-achieve the upgrade target-rate under sum-rate maximization. This results in lower crosstalk into the longer, e.g., central office deployed lines, and consequently in a higher service coverage.

VI. Conclusions

We propose a spectrum balancing algorithm for a novel coverage maximization problem in multi-carrier digital subscriber line (DSL) networks. It is enabled by scalable heuristics for the combinatorial single-carrier power allocation subproblems which show a favorable trade-off between their complexity and suboptimality. Evaluated in a large number of randomly varied DSL scenarios our spectrum balancing algorithm increased the average service coverage by more than $13\,\%$ compared to state-of-the-art sum-rate maximization algorithms. This is a significant increase in service coverage that should be attractive for DSL network operators.

References

[1] K. B. Song, S. T. Chung, G. Ginis, and J. Cioffi, "Dynamic spectrum management for next-generation DSL systems," *IEEE Communications Magazine*, vol. 40, no. 10, pp. 101–109, October 2002.
[2] W. Yu, G. Ginis, and J. M. Cioffi, "Distributed multiuser power control for digital subscriber lines," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 5, pp. 1105 – 1115, June 2002.
[3] W. Yu, R. Lui, and R. Cendrillon, "Dual optimization methods for multiuser orthogonal frequency division multiplex systems," in *IEEE Global Telecommunications Conference 2004 (GLOBECOM '04)*, vol. 1, Dallas, Texas, USA, 29 November – 3 December 2004, pp. 225–229.

[4]We perform iterative spectrum balancing (ISB) similarly as in [14] but differently using a line-search over bit-rates and a bisection search over Lagrange multipliers. Feasible power allocations under ISB are obtained by the heuristic in [8]. The used convergence criterion for ISB is a maximum number of 50 user sweeps not improving the Lagrangian or a total of 200 iterations, while we set bounds on the number of iterations for CM and RM as $30 \leq \hat{t} \leq 60$ and stop if the primal improvement over iterations falls below $0.1\,\text{ppm}$. In CM we initially increase $\hat{R}$ and $R_u, u \in \mathcal{U}$, by $3\,\%$ but use the original values for the final heuristic CHET, cf. Figure 1.

[4] M. Wolkerstorfer, D. Statovci, and T. Nordström, "Dynamic spectrum management for energy-efficient transmission in DSL," in *IEEE International Conference on Communications Systems 2008 (ICCS '08)*, Guangzhou, China, 19–21 November 2008.
[5] P. Tsiaflakis, Y. Yi, M. Chiang, and M. Moonen, "Fair greening for DSL broadband access," *ACM SIGMETRICS Performance Evaluation Review*, vol. 37, no. 4, pp. 74–78, March 2010.
[6] M. Guenach, C. Nuzman, K. Hooghe, J. Maes, and M. Peeters, "Reduced dimensional power optimization using class AB and G line drivers in DSL," in *International Workshop on Green Communications, IEEE GLOBECOM 2010*, Miami, USA, 6–10 December 2010, pp. 1443–1447.
[7] J. Lee, R. Sonalkar, and J. Cioffi, "Multi-user discrete bit-loading for DMT-based DSL systems," in *IEEE Global Telecommunications Conference 2002 (GLOBECOM '02)*, vol. 2, Taipei, Taiwan, China, 17–21 November 2002, pp. 1259–1263.
[8] D. Yu, K. Seong, and J. Cioffi, "Multiuser discrete bit-loading for digital subscriber lines," in *IEEE International Conference on Communications 2007 (ICC '07)*, Glasgow, Scotland, 24-28 June 2007, pp. 2755–2760.
[9] P. Tsiaflakis, M. Diehl, and M. Moonen, "Distributed spectrum management algorithms for multiuser DSL networks," *IEEE Transactions on Signal Processing*, vol. 56, no. 10, pp. 4825–4843, October 2008.
[10] M. Wolkerstorfer, J. Jaldén, and T. Nordström, "Column generation for discrete-rate multi-user and multi-carrier power control," submitted to *IEEE Transactions on Communications*, July 2011.
[11] P. Golden, H. Dedieu, and K. Jacobsen, Eds., *Fundamentals of DSL technology*. Auerbach Publications, 2006.
[12] R. Cendrillon, W. Yu, M. Moonen, J. Verlinden, and T. Bostoen, "Optimal multiuser spectrum balancing for digital subscriber lines," *IEEE Transactions on Communications*, vol. 54, no. 5, pp. 922–933, May 2006.
[13] D. P. Bertsekas, *Nonlinear Programming*. Athena Scientific, 1999.
[14] R. Cendrillon and M. Moonen, "Iterative spectrum balancing for digital subscriber lines," in *IEEE International Conference on Communications 2005 (ICC '05)*, vol. 3, Seoul, Korea, 16-20 May 2005, pp. 1937–1941.
[15] J. Desrosiers and M. Lübbecke, "A primer in column generation," in *Column Generation*, G. Desaulniers, J. Desrosiers, and M. Solomon, Eds. Springer, 2005, ch. 1, pp. 1–32.
[16] C. Helmberg and F. Rendl, "Solving quadratic $(0,1)$-problems by semidefinite programs and cutting planes," *Mathematical Programming*, vol. 82, no. 3, pp. 291–315, August 1998.
[17] ETSI, "Transmission and Multiplexing (TM); Access transmission systems on metallic access cables; Very high speed Digital Subscriber Line (VDSL); Part 1: Functional requirements," ETSI, Tech. Rep. TM6 TS 101 270-1, Version 1.3.1, July 2003.
[18] M. Wolkerstorfer and T. Nordström, "Heuristics for discrete power control: A case-study in multi-carrier DSL networks," in *ALIO/EURO Workshop on Applied Combinatorial Optimization 2011*, Porto, Portugal, 4–6 May 2011.
[19] M. Resende and C. Ribeiro, "Greedy randomized adaptive search procedures," in *Handbook of Metaheuristics*, F. Glover and G. Kochenberger, Eds. Kluwer Academic Publishers, 2003, ch. 8, pp. 219–249.
[20] M. Dorigo and L. Gambardella, "Ant colony system: A cooperative learning approach to the traveling salesman problem," *IEEE Transactions on Evolutionary Computation*, vol. 1, no. 1, pp. 53–66, April 1997.
[21] (2008, October) xDSL simulator v3.1. [Online]. Available: xdsl.ftw.at
[22] J. Campello, "Optimal discrete bit loading for multicarrier modulation systems," in *IEEE International Symposium on Information Theory 1998 (ISIT'98)*, Cambridge, MA, USA, 16-21 August 1998, p. 193.