

desired one, whereas "bubbles" of liquid metal form and collapse inside the object in the first experiment, which results in a casting of poor quality.




Fig. 3




Fig. 4



Fig. 5



Fig. 6




Fig. 7




Fig. 8

Acknowledgements This work was supported by the Program for Fundamental Research of Presidium of RAS P17, by RFBR (No. 11-01-12136-ofi-m-2011) and by the Program Leading Scientific Schools (NSh-4096.2010.1).

References

1. Albu, A.F., Zubov, V.I.: Mathematical Modeling and Study of the Process of Solidification in Metal Casting. *Computational Mathematics and Mathematical Physics* **47**, 843–862 (2007)
2. Albu, A.F., Zubov, V.I.: Functional Gradient Evaluation in an Optimal Control Problem Related to Metal Solidification. *Computational Mathematics and Mathematical Physics* **49**, 47–70 (2009)
3. Albu, A.V., Albu, A.F., Zubov, V.I.: Functional Gradient Evaluation in the Optimal Control of a Complex Dynamical System. *Computational Mathematics and Mathematical Physics* **51**, 762–780 (2011)
4. Albu, A.V., Zubov, V.I.: Choosing a Cost Functional and a Difference Scheme in the Optimal Control of Metal Solidification. *Computational Mathematics and Mathematical Physics* **51**, 21–34 (2011)
5. Evtushenko, Y.G.: Computation of Exact Gradients in Distributed Dynamic Systems. *Optimizat. Methods and Software* **9**, 45–75 (1998)

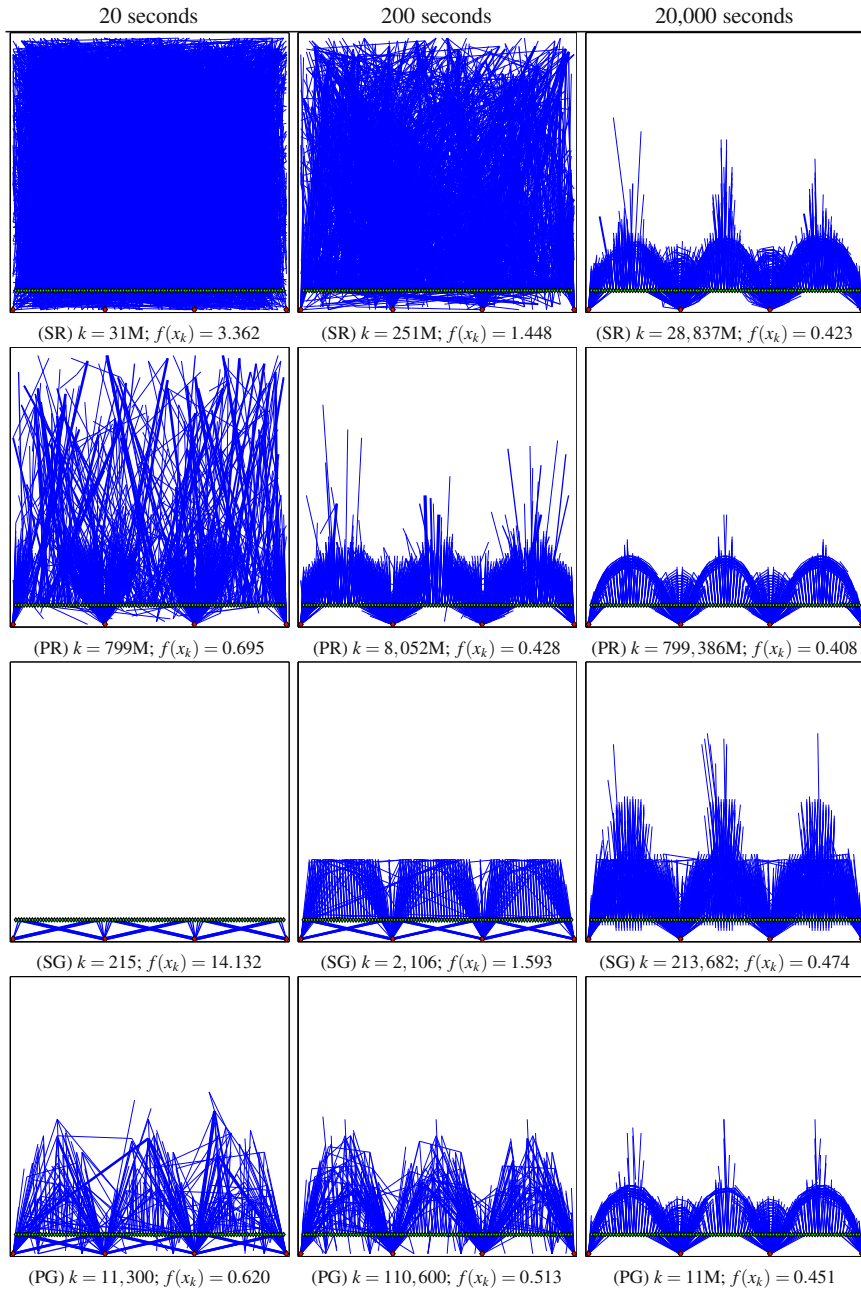


Fig. 1 “Bridge” truss after 20s, 200s (≈ 3.34 minutes) and 20,000s (≈ 5.56 hours) of computation time for algorithms SR, PR, SG and PG (rows in this order). Number of iterations (“M” = millions) and objective value is shown under each plot.

3 Extreme Points Crossover


Each component \tilde{x}^k of the solution \tilde{x} yielded by solving the master problem (2) is a convex combination $\tilde{x}^k = \sum_{p \in P_k} \tilde{\lambda}_{kp} x^{kp}$, with P_k being the (index) set of extreme points obtained by solving the k -th pricing problem. These are integer feasible and satisfy the constraints $D^k x^k \geq d^k$, while \tilde{x}^k satisfies all linear constraints but is usually not integer feasible. This leads to the question: Can we exploit the integer feasibility of the x^{kp} to obtain a solution \tilde{x} satisfying *all* constraints? This is what extreme point based heuristics try to achieve. The extreme points we are particularly interested in are those with $\tilde{\lambda}_{kp} > 0$.

It may happen that there is a number of coordinates in which these extreme points are identical. Especially when their values are also shared by \tilde{x} , it seems worth exploring whether there are more integer feasible or feasible points taking the same values. Therefore, *Extreme Points Crossover* performs a neighborhood search on the extreme points by *crossing* them: For each block, it takes those x^{vp} with the highest $\tilde{\lambda}_{kp}$, fixes the variables in which the selected points are identical and solves a sub-MIP.

As its name indicates, the heuristic can be compared to the Crossover heuristic discussed in [3, 12] and implemented in SCIP, which considers a number of already known feasible solution for crossing. The main advantage of our heuristic against this one is that it does not need any previously found feasible solution.

Fig. 1 Extreme Points

Crossover: An LP feasible solution \tilde{x} is a convex combination of $x^{p1}, x^{p2} \in \text{conv}(X)$. All the three points have a coordinate in common which is also shared by the encircled feasible point.



4 Restricted Master

When we solve the master problem, we dynamically add variables that improve the current solution. As the total number of master variables may be exponentially high compared to the number of original variables, this means that the master problem may soon reach a size where searching for feasible solutions would become very time-consuming. This is what the *Restricted Master* heuristic (suggested by Joncour et al. [9]) tries to overcome. It searches for an integer feasible master solution by restricting the formulation again to a subset of promising variables, hence regarding a problem which is of considerably smaller size than the current master formulation.




Fig. 1 The solving process depicted for the instance `n3seq24`

better is the expected quality of the incumbent solution if we stop the solver at an arbitrary point in time.

We used the benchmark set of the MIPLIB 2010 [5] as test set for our experiments. Since we are interested in the primal part of the solution process, we excluded the four infeasible instances `triptim1`, `enlight14`, `ns1766074`, and `ash608gpia-3col`; further, `mspp16` was excluded since it terminated for memory reasons during presolving for all our tests. Thus, 82 test instances remained. We performed four different runs: SCIP without any primal heuristics (Heur_OFF), SCIP using only the rounding and propagation heuristics which are described in this paper (R&P), SCIP using all default heuristics except the ones described in this paper (NoR&P), and SCIP with default settings (DEF). All experiments were conducted with a time limit of one hour, a memory limit of 4 GB on a 3.00 GHz Intel® Core™2 Extreme CPU X9650 with 6144 KB Cache and 8 GB RAM.

Figure 1 exemplarily shows the primal gap function $p(t)$ for the four settings applied to the instance `n3seq24`. A square shows when a new primal solution is found and its quality. It can be seen that the two settings that use rounding and propagation heuristics find solutions earlier and hence have a smaller primal integral $P(t_{\max})$. For all four settings, SCIP found the optimal solution within an hour, but timed out without proving optimality. Interestingly, with disabled heuristics (Heur_OFF), the optimal solution was found in the smallest amount of time. Nevertheless, the behavior of the default settings (DEF) seems favorable since primal solutions of reasonable quality are found much earlier.

Table 1 shows aggregated results of our experiments. The first row shows the evaluation of the *normalized* primal integral over the entire testset. For each instance, the corresponding integral $P(t_{\max})$ is divided by the integral obtained with the setting Heur_OFF. This reference value is then used to compute the geometric mean $\phi(P)$ for each setting. The remainder of the table shows the geometric means

an evaluation of the incurred trenching cost in overall cost-optimized networks as computed by the FTTx-PLAN software.

To this end, we use an extended Steiner tree model to determine lower bounds on the trenching costs for a given FTTx instance. This model is presented in Section 3, after a high-level description of the FTTx network planning problem is given in Section 2. Finally, Section 4 compares trenching costs in solutions obtained from FTTx-PLAN with the lower bounds obtained by the Steiner tree model on various realistic test instances.

2 Problem formulation

The FTTx network planning problem can be described as the task to connect a number of given customers to central offices of the telecommunication carrier, using optical fibers and various active and passive components. Fibers can be laid out in different types of cables which themselves are direct-buried or embedded into different types of ducts, which are eventually buried into the ground along specified trails of the deployment area.

The topological structure of the input is represented by the *trail network*, an undirected graph whose edges designate the trails along which connections (fibers, cables, ducts) can be laid. To each edge is assigned a cost value that determines the trenching costs for the trail in question. The trails can also have existing infrastructure, such as dark fibers or ducts, which can be used for planning.

Some nodes in the trail network can be of a special type – each such node is associated with a setup cost which is incurred if it is included in the final network. [Figure 1](#) shows an example of a trail network with special nodes BTPs and COs.

Fig. 1 Trail network of an instance, projected onto a satellite image of the deployment area. Yellow dots indicate BTPs, red diamonds possible CO locations; trails are colored according to their trenching costs – green for low, red for high costs.



4 Results and conclusion

We assessed trenching costs in 7 instances, given in Table 1. The first three trail networks were artificially generated using GIS information from OpenStreetMap [1]. The last four are modified networks provided by different German city carriers. Figure 2 shows an example of an FTTX-PLAN solution and the trenched trails attaining the lower bound for the instance a2.



Fig. 2 Solution network (left) to instance a2 and trenched trails attaining the lower bound (right).

Table 1 shows the results from the computations. Besides the size of the instances, we have given the trenching costs incurred by the solution networks, the lower bound obtained with the model in Section 3, and the relative gap between these two values. Additionally, we list the average and maximal connection lengths for the networks and the lower bound.

Instance:	a1	a2	a3	c1	c2	c3	c4
# nodes	637	1229	4110	1051	1151	2264	6532
# edges	826	1356	4350	1079	1199	2380	7350
# customers	39	238	1670	345	315	475	1947
# potential COs	4	5	6	4	5	1	1
network trenching cost	235640	598750	2114690	322252	1073784	2788439	4408460
average connection length	717.1	1186.6	733.4	589.6	1969.2	995.3	2094.1
maximal connection length	1320	2257	2334	1180	4369	2452	4809
lower bound	224750	575110	2066190	312399	1063896	2743952	4323196
relative gap	4.8%	4.1%	2.3%	3.2%	0.9%	1.6%	2.0%
average connection length	299.2	887.5	504.6	363.7	1531.0	590.2	1092.9
maximal connection length	1049	3142	3215	1394	7315	2586	4786

Table 1 Comparison of trenching costs for different instances.

It can be observed that trenching costs in the solution networks are in an acceptable range with respect to the obtained lower bounds (up to 5 percent). Furthermore,


		individual contribution provider			individual contribution customer		group weighting	normalized group weighting	
	Provider			Customer					
Profit provider	Max	0.65	0.33	Min	0.60	-0.30	0.03	Max	0.06
Energy consumption	Min	0.10	-0.05	Min	0.25	-0.13	-0.18	Min	0.44
Cooperation	Max	0.25	0.13	Max	0.15	0.08	0.20	Max	0.50

Table 2 Weighting

Results

Figure 1 depicts the net flows and the rankings of the single decision makers and the group. There are no differences in the results between PROMETHEE I and PROMETHEE II, therefore, the results of the partial ranking are not shown separately. Considering the results of the single decision makers, the provider prefers PSS 3 most. After this follows PSS 2 and on the last position is PSS 1. For the customer the ranking is the other way round. He prefers PSS 1 most and PSS 3 least. Regarding the results of group decision, PSS 2 is the best compromise for both the decision makers. This is ranked on the second position for both decision makers, so no one has to cope with the worst choice. According to the group weighting PSS 3 is at second position with only little distance to the first one. The last-ranked business concept is PSS 1, where the negative flows predominate, what already can be seen in the decision matrix, as this alternative is worse than PSS 3 in all criteria, as long as the profit for the provider is maximized.

Fig. 1 Results single decision maker and group



Discussion

If the criteria with conflicting targets are highly ranked by all decision makers they are losing importance as shown above. Therefore, the decision is mainly based on criteria with equal targets of the decision makers whereby the veto function helps to find a compromise concerning criteria with conflicting targets. In case that only one decision maker considers the conflicting criteria as important and the other decision maker has rather low weights for these special criteria, the calculated difference is

2 Approach

For answering these questions we are using a linear programming model to optimize production³ and investment decisions⁴ of the energy sector. Thereby the objective is to maximize the profit, calculated with a contribution margin accounting, and the decision variables are running times of plants and capacity enlargements. For reliable results, we first have to consider the political, economic and technical framework conditions, what gives us the input parameter, constraints and appropriate model structure. Accordingly, we do not simplify but respect the difference between the ETS and taxes in taking the EA price as uncertain and endogenous variable. This results in a calculation of the EA price as intersection of the supply and demand function. The first is given by the CAP. However, the second is deduced by the optimal quantity of carbon emissions depending on the adaption of the merit order to variations in the EA price. Since the switch load effect leads to a discontinuous slope, we calculate the demand function for EA with a sensitivity analysis.

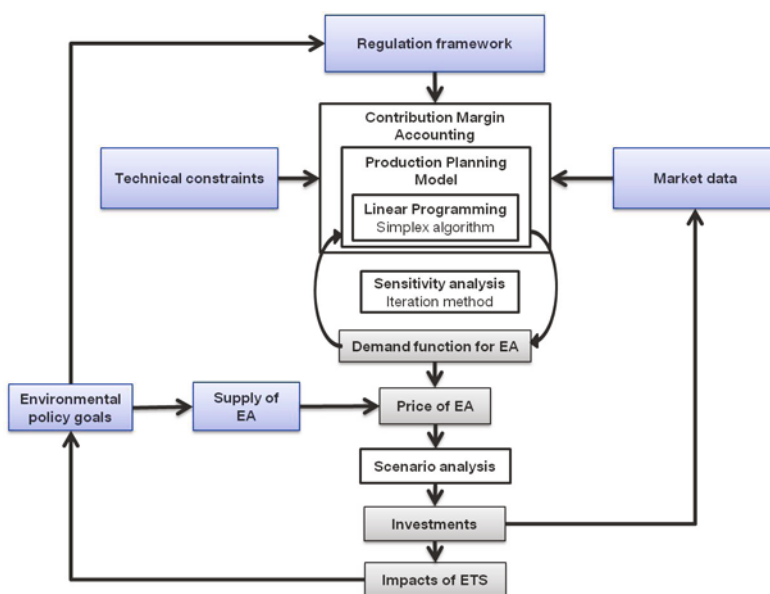


Fig. 1 Approach of the production and investment planning model

In the long run not only running times of plants are variable, but also changes in the plant portfolio are possible due to construction and wear. Since in the long-term perspective the framework conditions may change additionally, we consider the endogenous modifications of the input parameter in a feedback loop on the one hand and use a scenario analysis for the exogenous changes on the other hand. In

³ Cf. Fichtner (2005).

⁴ Cf. Roth (2008).

$$0 \leq \sum_{i=1}^I (P_{el,i,a} + Z_{el,i,a}) \leq \bar{P}_{el,i} \quad (8)$$

$$0 \leq \sum_{h=1}^H (P_{th,h,a} + Z_{th,h,a}) \leq \bar{P}_{th,h} \quad (9)$$

$$0 \leq (P_{el,i,a} + Z_{el,i,a}) \quad \forall i \quad (10)$$

$$0 \leq (P_{th,h,a} + Z_{th,h,a}) \quad \forall h \quad (11)$$

For cross-border grid capacities, the EU commission is planning network investments, which are allocated according to demand and international trading. Dismantling is excluded in order to build an European wide electricity market.

$$0 \leq \sum_{b=1}^B Z_{el,b,a} \leq \bar{Z}_{el,b,a} \quad (12)$$

$$0 \leq Z_{el,b,a} \quad \forall b \quad (13)$$

4 Results

After solving the optimization model with a simplex algorithm, we receive similar but too low investment levels in scenario 0 to 3. This underinvestment problem leads to cutbacks in total capacity. However scenario 4 with lower capital costs for renewables due to subsidized interest rates shows that countermeasures are possible.

In scenarios 1 to 4 (all with ETS) is the enlargement of power production by biomass and the decline of coal remarkable. Additionally, in scenario 1 and 2 the increasing use of gas is significant both in power and heat production. Investments in wind, water and geothermy plants play only in scenario 4 a special role, which is the only scenario which reaches the EU and national goals of 22 % respectively 30 % of renewables in the energy mix.

Furthermore the total amount of power production decreases in all scenarios, but especially in 3 with the expansion of the EU internal electricity market, since in balance there will be an electricity import instead of an export as before.




Fig. 2 Investments in capacities and development of the plant portfolio




Fig. 3 Power and heat production per technology

But what are the results for the ecological effectiveness and economic efficiency of the ETS? The emission allowance price shows a big range between 0 and 85 Euro in 2020 and so the emission costs with 0 to 24 b. Euro. This means specific emission costs of 0 to 4.01 ct/kWh_{el}. Since emission costs are reflected in the power price and the ETS triggers a fuel switch from lignite and coal to more expensive but low emission gas, in 2020 the power price can be 6.16 ct/kWh_{el} higher in comparison to the scenario 0 without ETS. Renewables, nuclear plants and an EU internal electricity market can reduce this price increase effect, but not the emission amount, since the CAP is always used. Nevertheless is the emission abatement ascribed to the ETS 96 to 119 m. t CO₂ or 0.146 to 0.201 t CO₂/MWh_{el}. Other reasons for emission abatement can be a technological efficiency increase (about 2 to 15 % points on average over the total plant portfolio) and so a decline in fuel consumption. However, an increase of the primary energy import quota of up to 7 % points in scenario 1 and a decrease of the supply security of up to 8 % points in scenario 4 can also be observed.

5 Discussion and conclusion

Since the ETS has strong impacts on production decisions, it is expedient for emission control. But low interest rates (scenario 4) influence more effective long-term decisions like plant investments. So effective the ETS is in achieving emission goals, so costly it can be on the other hand. Emission costs can go up to 24 b. Euro in 2020 and can noticeably drive up power prices. But choosing the right regulatory framework can lower this costs.

References

1. Fichtner, W.: Emissionsrechte, Energie und Produktion Verknappung der Umweltnutzung und produktionswirtschaftliche Planung. Berlin (2005).
2. Geiger, B., Hardi, M., Brückl, O., Roth, H., Tzscheutschler, P.: CO₂-Vermeidungskosten im Kraftwerksbereich, bei den erneuerbaren Energien sowie bei nachfrageseitigen Energieeffizienzmaßnahmen. IfE Schriftenreihe Heft 47, München (2004).
3. Roth, H.: Modellentwicklung zur Kraftwerksparkoptimierung mit Hilfe von Evolutionsstrategien. München (2008).




Fig. 1 Development of European power plant capacities and power production

4.2 Production

In Fig. 2 typical production patterns for European power production in the year 2050 (climate-market scenario) are presented. In the upper figures a typical weekend day in summer is depicted. Here, load is low compared to other days of the year and especially solar infeeds are high. On the right side infeed from renewables on a summer day are low while they are high on the left side. The adjustment of the




Fig. 2 European power production at the twelfth hour of a typical summer weekend day with a) high and b) low RES infeeds and an evening of a typical winter workingday with c) high and d) low RES infeeds

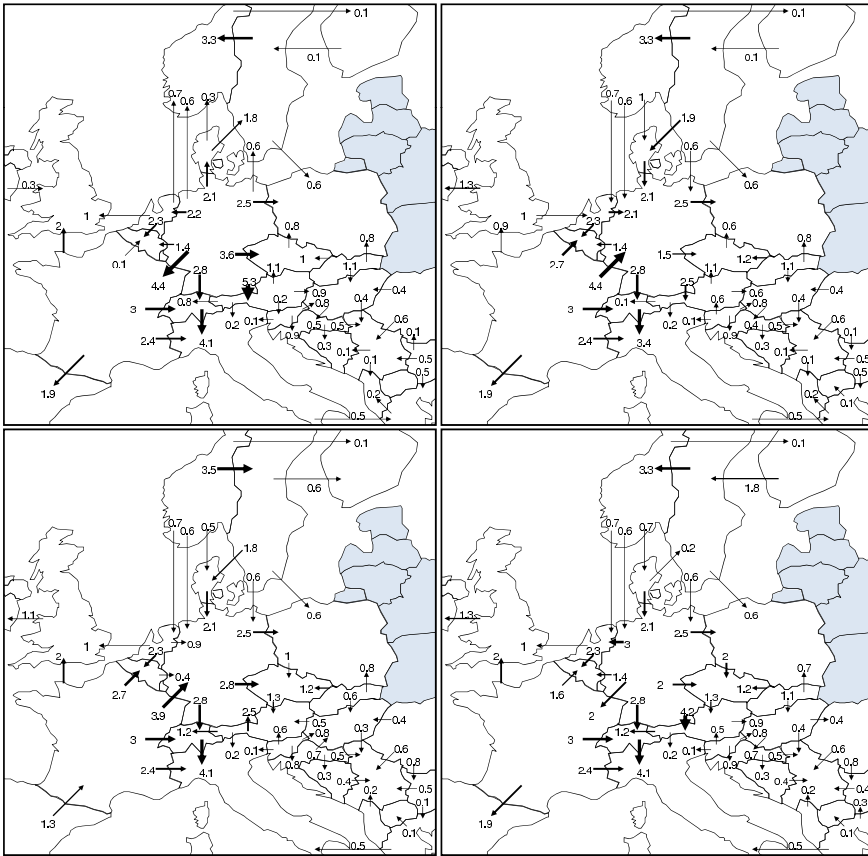


Fig. 3 Transmission on a typical summer weekend day with a) high and b) low RES infeeds and a typical winter workingday with c) low and d) high RES infeeds

References

1. Boiteux, M.: Peak load-pricing. *Journal of Business*. **33**, 157–179 (1960)
2. Kchler, C., Vigerske, S.: Decomposition of Multistage Stochastic Programs with Recombining Scenario Trees. *Stochastic Programming E-Print Series* 9, (2007)
3. Mst, D., Fichtner, W.: Renewable energy sources in European energy supply and interactions with emission trading. *Energy Policy*. **38**, 2898–2910 (2010)
4. Spiecker, S., Weber, C.: Integration of Fluctuating Renewable Energy - a German Case Study. EWL Working Paper (2010)
5. Swider, D., Weber, C.: The costs of wind's intermittency in Germany: application of a stochastic electricity market model. *Euro. Trans. Electr. Power*. **17**, 151–172 (2007)
6. Tuohy, A., Meibom, P., Denny, E., O'Malley, M.: Unit commitment for systems with significant wind penetration. *IEEE Trans. Power Syst.* **24**, 592–601 (2009)

ple [1, 5–9]. Many concepts include still a manual operation of the devices by the user. This leads to a significant loss of comfort and requires the presence of the user. In addition, the user must actively inform about the current and future prices.

Dynamically priced electricity rates are much more complex than constantly priced electricity rates: A dynamically priced electricity rate model must define possible prices, steps as well as upper and lower limits, and must also determine how far in advance consumers will receive the price information. The model must not be too complex, so that it remains understandable to consumers and doesn't violate regulatory restrictions [2]. The price model planned for the field experiment and used in the simulation prototype is kept very simple: The base price is set to 20Ct/kWh and there is a range between a defined minimum (lower than actual price) of 15Ct/kWh and maximum (higher than actual price) of 25Ct/kWh (11 1Ct/kWh steps). The price is defined for each hour and will be constant for at least two hours. The price is announced and fixed for six hours (preview time) in advance. This enables a fast response to changes in wind and solar power and still offers the consumer a reasonable certainty about the costs incurred for its electricity consumption.

3 Simulation of Smart Devices

As smart devices for consumer households are not widely available yet, we simulate their behavior as preparation for the field experiment with 500 households. Smart devices are able to react to changing price signals. A rational behavior for a smart




Fig. 1 Simulation result of devices running once (like dishwashers, washing machines, dryers)




Fig. 2 Simulation result of heating devices (like water boilers with storage)

device is to run when electricity rates are low. The exact optimization strategies used by real smart devices are still unknown, but simulation allows to test several strategies and parameters. The realized load is heavily dependent on the optimization strategy used by the devices. We implement two classes of smart devices:

- Devices running once (for a few minutes to several hours) in a defined limited time slice (of several hours, typically much longer than running time), as simplified model of real devices like dish washer, washing machine, or dryer. These devices start at once as soon as the current time is the unique cheapest time within the preview time. Otherwise the devices wait for a cheaper time. It is ensured that the device execution is complete at the end of the limited time slice. [Figure 1](#) shows exemplary results from simulation of this device class.
- Devices without time restriction and a given daily runtime, as a simplified model of water heaters with large storage or heat pumps. The daily running time is based on running times of real water heaters. These devices have a storage capacity of several hours. The storage is filled completely while the current time is the cheapest time within the preview time. Otherwise the device will only run (for minimal required time) when the storage runs empty in the beginning hour.

The results (exemplarily see [figure 1](#) and [2](#)) of simulation show that the price significantly influences the load. The two implemented device classes result in different delta-loads (difference between uninfluenced and influenced load). The load change is asymmetric: There are high positive peaks and smaller negative troughs. The very high peaks are the result of the cheapest solution for the device owners, but are not in the interest of the energy provider.




Fig. 1 Inconsistency between price and recovery rate for a Greek government bond.

*hazard rate model*¹. This suggested dynamic calibration is in contrast to usual usage of static CDS models, where a deterministic hazard rate function is derived from market prices for a given recovery rate. Instead, the market-implied recovery rate now becomes an output of the model rather than an input. In this way we aim to make practical the recent more expressive models [2, 5, 6], however, we limit the scope of this exposition to market-consistent calibration of static recovery rates for brevity of presentation.

2 Modeling

Let us start by recalling the most common definitions of the *recovery rate* δ of a bond, see for example [4], Section 5.7 or [1]:

- *Recovery of face value*²: A fraction $\delta \in [0; 1]$ of the outstanding notional N is paid at the *time of default* τ .
- *Recovery of treasury*: A fraction $\delta \in [0; 1]$ of the outstanding notional N is paid at the *bond maturity* T .
- *Recovery of market value*: A fraction $\delta \in [0; 1]$ of the value of the claim just prior to default (denoted by $PV_{Bond}(\tau^-)$) is paid at time of default τ .

¹ See for example [4], Section 8.4 or <http://www.cdsmodel.com/cdsmodel/>

² Recovery of face value is the standard assumption used in most CDS pricing models, for example in the already mentioned hazard rate model which is also the standard ISDA model, see <http://www.cdsmodel.com/cdsmodel/>

$$CF_{CDS}^{Prot}(t_i) = \Delta_i \cdot (f_i + s_{Bond}) \cdot \mathbf{1}_{\{\tau \in (t_{i-1}, t_i]\}} + (1 - \delta) \cdot \mathbf{1}_{\{\tau \in (t_{i-1}, t_i]\}}.$$

- *Premium leg*: In return for the protection, a fixed premium s_{CDS} has to be paid by the protection buyer until a default has been observed:

$$CF_{CDS}^{Prem}(t_i) = \Delta_i \cdot s_{CDS} \cdot \mathbf{1}_{\{\tau \geq t_i\}}.$$

Then the present value of the CDS can be obtained similarly to the price of the bond

$$PV_{CDS}(t_0) = \mathbf{E}_{\mathbf{Q}} \left[\sum_{i=1}^n P(t_0, t_i) \cdot CF_{CDS}^{Prot}(t_i) \right] - \mathbf{E}_{\mathbf{Q}} \left[\sum_{i=1}^n P(t_0, t_i) \cdot CF_{CDS}^{Prem}(t_i) \mid \mathcal{F}_{t_0} \right],$$

and it can be shown (by comparison of the cash flows to a default-free floating rate bond) that it holds

$$\forall t : PV_{Bond}(t) + PV_{CDS}(t) = 1 \iff s_{CDS} = s_{Bond},$$

independent of the specific recovery value. More generally, for a given recovery rate δ , the bond price can be deduced from the CDS price and vice versa, as both prices are only depending on the hazard rate (i.e. default probabilities). Thus it is obvious that if the bond trades below the recovery rate, no reasonable CDS price or par spread can be obtained any more for the corresponding CDS, see Figure 2.




Fig. 2 Inconsistency between price and recovery rate leads to non-existing CDS par spread.

In practice, if the bond price drops towards the current recovery assumption, the assumption is reviewed and the recovery rate is adjusted downwards. If the bond

Fig. 1 Distribution of the angles of the social value orientation as measured by the Ring Measure of Social Values




Table 1 Classification of the players into prosocial and prosself according to their social value orientation

	Number of cases	%
Proself	62	71.3
Prosocial	25	28.7

3 Experimental Design

We conducted a computerized PD experiment using the payoffs in Table 2. Players interacted repeatedly with the same partner and were rematched with a new partner two or three times. In 9 cases players are rematched once, in 14 cases the players are rematched twice. Thus we had two or three matchings per session. A matching we call a number of successive rounds in which one player interacts with the same partner. A player is matched with a different partner after a number of rounds not known to the participants in advance. One matching lasts between 10 and 35 rounds.

The experiment is conducted in the Experimental Economics Laboratory of the Department of Statistics and Operations Research at the University of Graz. The participants are recruited in classes and by campus advertisements. Potential participants are informed that they can take part in a paid experiment consisting of an decision making task. We arranged the experiments in a way that 8 or 12 participants take part at the same time. 92 subjects took part in the experiments.²

These participants were then randomly assigned in groups of four that were led to computer workstations. The workstations are seperated from each other so participants cannot directly interact with each other. Subjects do not know which one of

² 5 Players were excluded from the analysis because either they either were outliers in the social preferences distribution or their comments suggested that they did not understand the game properly.

$$d_1 = (.40, .20, .30, .10)$$


$$d_2 = (.60, .10, .25, .05)$$

$$d_3 = (.80, .00, .15, .05)$$

We clearly see an increase in the initial probability of joint cooperation from matching to matching. We like to denote this as a 'restart-effect', because it seems likely that a new partner is credited with an increasing level of trust, once the individual learns that mutual cooperation has a tendency to increase returns. On the other hand, within one matching we find a decreasing probability of joint cooperation over time. This tendency can be exemplified by the steady state distribution generated by M : $(.16, .14, .14, .56)$. So in the long run there seems to develop a dominating probability for joint defection.


This simple model seems to contribute to a better understanding of experimental observations in several ways. Firstly, it mimics the high probabilities of rather long runs of either CC or DD states, which are, e.g., about 60% for a run length of 4, given an initial state of either CC or DD . Secondly, the model can be used to simulate hypothetical experimental results which can be compared with the actual ones. This allows for a whole range of advanced analysis.

The following figure shows the probability of joint cooperation (CC) found in 20 simulated experiments, using the matching-specific transition matrices in connection with the initial distribution for each matching, together with the actual experimental result (thick black line). The simulations seem to capture the experimental observations reasonably well, although we have a rather short length of the matching, and the first two matchings had equal length of 6 rounds. It is quite obvious that players assumed that the third matching would also have 6 rounds, which led them to defect more often because of expected end effects, which could not be captured with transition matrices estimated for whole matches. For the first two matchings, the experimental observations are well covered by the range of the simulated observations. The restart-effect is also prominent.



The simulated experiments show rather wide fluctuations. This indicates that we should expect similar fluctuations in the results of actual experiments if we repeat them. This should make us careful in interpreting certain patterns in experimental data, as they could prove as the mere result of stochastic fluctuations.

The next figure shows these probabilities as inferred directly from the estimated transitions matrices, again together with the actual experimental result (thick black line). The apparent deviation is likely to have two reasons. The fitted model is not flexible enough to reflect learning effects properly. We have to live with that because of data limitations. More important is the insight, that it could well be possible that the typical behavior that we might find based on large data sets might well be represented by Markov matrices, but that just the realization of the actually conducted experiment deviates from this average, although its stochastic behavior is well represented.



To clarify these issues we need to develop test criteria for the accuracy of Markov models, and it is likely advantageous to use time dependent models if enough data is available. Tests need then to be performed out-of-sample.

4 Conclusion

We conducted a series of IPD experiments with unknown length and re-matching of players and observed the patterns of cooperation over time. This pattern shows apparent stochastic behavior, which calls for an attempt of stochastic modeling. We present descriptive as well as simulation results for a very simple Markov model derived under the assumption of stochastically equivalent behavior of all players. They indicate that various patterns in the data may well be the results of an in fact simple memoryless stochastic behavior. With more data, likely learning behavior could also be investigated in more detail.

References

1. Dal Bo P.: Cooperation under the shadow of the future. Experimental evidence from infinitely repeated games, *Am. Econ. Rev.* **95**, 1591-1604 (2005)
2. Hennig-Schmid, H., Leopold-Wildburger U.: How the shadow of the past affects the future - An iterated prisoners dilemma experiment, working paper, University of Graz, (2011)

2.1 Data

The L.A. Jackman et al. (1997)⁵ study was based on the idea that the achievement of maximal calcium retention during adolescence may influence the magnitude of peak bone mass and subsequently lower risk of osteoporosis. Calcium retention is generally considered to reach a plateau at a certain calcium intake. To test this idea, subjects were given two controlled diets with a different amount of calcium in each. However, some individuals completed the study and provided two observations while others provided only one observation due to drop outs.

In our data, the total number of observations is 202 which includes 44 subjects with a single observation, and 79 subjects with two observations. Besides, there are no significant differences between these subjects in the following areas: medical problems, use of medications, pregnancy, abortion, eating disorder, or current use of tobacco. The distribution of the response variable (calcium retention) has clear center, and is slightly symmetric. On the other hand, the distribution of the independent variable (calcium intake) is not normal. Since some subjects have two observations for calcium intake levels (low/high), the distribution for this variable has two peaks. Moreover, there is a nonlinear relationship between the dependent and independent variable.

Figure 1 shows that there are positive correlations between high/low calcium intake levels and high/low retention levels for most of the subjects. These correlations demonstrate the dependence between these two levels of variables for the same subject. This dependence affected the re-sampling part of the study.




Fig. 1 Two levels of calcium intake versus calcium retention for subjects with double observations

and mBIC2 compared with mBIC is larger than in case of extreme sparsity, leading to misclassification rates of the FDR controlling rules which are now visibly smaller than those of mBIC. It is noticeable that BIC is doing slightly better in the denser case. Its FDR is now converging towards 0 for large p , and its misclassification rate approaches that of mBIC.

Effect of correlations: For the second part of the simulation study a design matrix with $n = p = 256$ was obtained by simulating from a multivariate normal distribution $X \sim \mathcal{N}_p(0, \Sigma)$. Figure 2 shows the block structure of Σ . Correlation within blocks was set constant, varying between $\rho = 0$ and $\rho = 0.8$ for different matrices. Trait values of 2000 replicates were simulated similarly to the previous section, where effect sizes were chosen slightly larger with $\tau = 0.5$. The positions of the $k = 24$ non-zero regressors are shown in Figure 2.

For orthogonal designs the minimization of selection criteria is particularly simple, because the estimate of a regression coefficient for a given explanatory variable does not depend on other regressors. It is enough to find the minimum among nested models according to the order obtained from p-values of one-factor regression models. In the non-orthogonal case one has to consider instead more sophisticated search procedures, leading to more time consuming procedures. Here we applied backward elimination starting from the top 40 regressors of one-factor models, followed by forward selection. This heuristic procedure is not guaranteed to find the model which minimizes a given selection criterion, but more sophisticated search procedures would become prohibitively time consuming for a simulation study. The focus of this article is on comparing different criteria, and to this end the relatively simple search strategy serves well.

Figure 3 illustrates the effects of correlation on the performance of different selection procedures. The larger the correlation, the more often regression models select a strongly correlated regressor instead of the correct one, resulting in a loss of power. However, this effect only becomes recognizable for $\rho \geq 0.3$, and really severe for $\rho \geq 0.5$. In the literature on GWAS thresholds on correlation have been used to distinguish between correct detections and wrong detections [8,9]. The presented results indicate the usefulness of a threshold somewhere between $\rho = \sqrt{0.05}$ of [9] and $\rho = 0.7$ of [8].




Fig. 2 The correlation structure of regressors consists of several blocks of different sizes. Within each block correlation was ρ , otherwise 0. Stars indicate the position of non-zero regressors.




Fig. 1 Basic Scenario

their goods from an outside region via the terminal. The flow of a full container from the terminal to the receiver is called inbound full (IF) container. It is obvious that a shipper requires an empty container before he can fill the freight in it. Additionally, an empty container remains at the receiver's location after the IF container is emptied. We define two time windows at each customer location: During the first time window the full/empty container has to be delivered to the customer location. After the container is emptied/filled, it can be picked up by a vehicle during the second time window. Moreover, trucking companies have to consider two additional transportation requests. Due to the imbalance between import- and export-dominated areas, they need to take care of outbound empty (OE) or inbound empty (IE) containers which either have to be moved to the terminal or derive from it. For these requests only the terminal as the destination or as the origin is given in advance. Hence, the locations that can provide empty containers for the OE requests and the destinations for the imported empty containers need to be determined during the solution process. This scenario was defined by Zhang et al. (2009) as the multi-depot container truck transportation problem (CTTP) [5]. The objective is to minimize the carriers' total fulfillment costs consisting of fixed and variable costs. Hence, in a first step the number of used vehicles should be minimized while in the second step the optimization of the vehicles' total operating time symbolizing the transportation costs should be pursued [3]. In this paper the first objective is formulated as a constraint and the minimization of the operating time is chosen as the objective function of our models. Thereby, the number of used vehicles within the employed models is raised iteratively until a feasible solution is found. The resulting solution approach leads to an integrated model which does not only consider vehicle routing but also empty container repositioning. Thus, solving the problem determines: a) where to deliver the empty containers released after inbound full/empty loads, b) where to pick up the empty containers for outbound full/empty loads, and c) in which order and by which truck the loads should be carried out.

In this paper, two scenarios for the hinterland transportation are analyzed. In the first scenario empty containers are uniquely assigned to trucking companies, i.e. they can solely be used by the company they are assigned to. For instance, an empty container obtained at a receiver location served by a certain trucking company can exclusively be used for transportation requests of this company. This scenario is given by the above CTTP. In the second scenario companies can use empty containers of cooper-

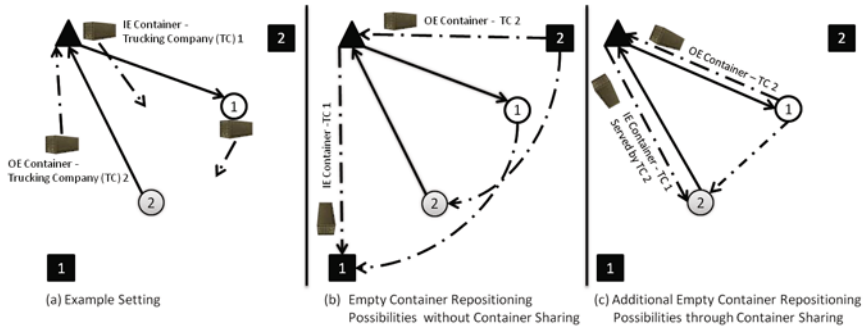


Fig. 2 Possible Benefits through Container Sharing

ating trucking companies. I.e. companies share their information at which locations empty containers are currently stacked and they agree with the mutual exchange of these containers. Thus, the cooperating companies can improve their routes and increase their profit by decreasing transportation costs in return. The permission of sharing empty containers among trucking companies leads to the multi-depot container truck transportation problem with container sharing (CTTP-CS).

Figure 2 demonstrates the rising possibilities to reposition empty containers of two trucking companies cooperating with each other. Trucking company 1 has to serve a receiver and needs to move an IE container to the hinterland. Trucking company 2 has to serve a shipper and needs to move an OE container to the terminal (see Figure 2 (a)). In the non-cooperative case, as illustrated in Figure 2 (b), the only opportunity for both companies to reposition empty containers is moving them either to (trucking company 1) or from (trucking company 2) their own depot. If the exchange of containers is permitted, both companies could benefit through the emerging additional flexibility to allocate empty containers to a vehicle's tour (see Figure 2 (c)). The amount of the benefit of a cooperation highly depends on time and place conditions given by the time windows for pickup and delivery and by the locations of the terminal and customers.

2 Simultaneous Solution Approach

Let V denote the nodes of a directed graph, consisting of customer node set V_C , terminal node set V_T and depot node set V_D . There are two types of customers ($V_C = V_S \cup V_R$), the node sets $V_S = V_{Si} \cup V_{So}$ and $V_R = V_{Ri} \cup V_{Ro}$ describe the shipper and the receiver node sets. V_{Si} and V_{Ri} refer to the first time window of the shipper/receiver, in which an empty/full container has to be made available. After the container has been completely filled or emptied, respectively, container $c \in C$ can be picked up by a vehicle $k \in K$ during the second time window (V_{So} and V_{Ro}). The terminal node set V_T consists of V_{Ti} and V_{To} which correspond to the amount of ingoing and outgoing




Fig. 1 Construction of the tour planning graph.

i.e., $v = (s, t) \in S \times Z := \{0, \Delta, \dots, T - \Delta, T\} \subset [0, T]$, or they represent artificial start and end nodes d_s and d_t for the vehicle paths (depot nodes). Directed arcs connect either adjacent time intervals of the same section, or they connect adjacent sections, i.e., $\forall s \in S$ there is $((s, t_1), (s, t_2)) \in A$ with $t_2 = t_1 + \Delta$, starting at $t_1 = 0$ until $t_2 = T$, and if $(s_1, s_2) \in N$ it holds that $((s_1, t), (s_2, t + \Delta)) \in A \forall t \in Z \setminus \{T\}$. [Figure 1](#) illustrates this construction. In addition, arcs are inserted from the start depot to all other non-depot nodes and from all non-depot nodes to the end depot node. Finally, a profit value is associated with each node $v = (s, t)$. We consider the problem of finding a feasible (d_s, d_t) -path in D for each vehicle f on each day, that respects a tour length restriction of 8h30. Each control tour corresponds to such a *control path*. The profit w_p of a control path p is the sum of the profit of its visited nodes. This approach could be seen as a vehicle routing problem with profits under some additional constraints. Vehicle routing is an well-established research area, see [6] for an overview. For the case of dealing with profits Feillet et al. [3] give a literature survey.

There are restrictions on the feasible starting and ending times of the control tours; the feasible times are defined by the *(Working) Time Windows*. Let P be set of all control paths in D and $P_{f,j} \subset P$ the set of all paths that are feasible for vehicle $f \in F$ and start at day $j \in J$. Furthermore for a given section $s \in S$, the set of all paths $p \in P$ that visit a node $v = (s, t_i) \in V$ is denoted by P_s and the minimum control quota is named by κ_s . The *Tour Planning Problem (TPP)* is then formulated as a 0/1 multi-commodity flow problem in D , where vehicles f represent the commodities. We introduce 0/1-variables z_p , $p \in P$, that indicate if tour p is chosen or not. Then the TPP can be modeled by the following integer program:

$$\max \sum_{p \in P} w_p z_p \quad (1)$$

$$\sum_{p \in P_{f,j}} z_p \leq 1, \quad \forall (f, j) \in F \times J \quad (2)$$

$$\sum_{p \in P_s} z_p \geq \kappa_s, \quad \forall s \in S \quad (3)$$

$$z_p \in \{0, 1\}, \quad \forall p \in P. \quad (4)$$

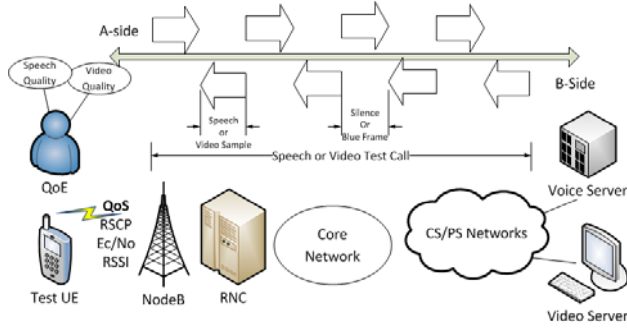


Fig. 1 UMTS architecture and Measurement Campaign configuration.

$\{s_1(\tau_1), s_2(\tau_2), \dots, s_n(\tau_N)\}$, with the corresponding degraded received samples, $\tilde{S}(\tau) = \{\tilde{s}_1(\tau_1), \tilde{s}_2(\tau_2), \dots, \tilde{s}_N(\tau_N)\}$, at the called party, *B-side*, by applying the objective assessment algorithm, ϕ . The objective QoE indicator, $MOS_{i \rightarrow j}$, varies between 1 (bad/very annoying) and 5 (excellent/imperceptible) and $(i, j) \in \{(A, B), (B, A)\}$:

$$MOS_{i \rightarrow j} = \frac{1}{N} \sum_{v=1}^N \phi(s_{v,i}(\tau_v), \tilde{s}_{v,j}(\tau_v)) \quad (1)$$

[Fig.1](#) presents a measurement campaign configuration for E2E QoE assessment with a “Diversity Benchmark” [12]. Besides, we conducted post-processing analysis with *NetQual NQDI* on a measurements database managed by *MS SQL Server* and we used MATLAB for regression and optimization modelling.

3 Robust Optimization in Regression Analysis

In order to predict QoE from QoS, a non-linear regression method can be performed for speech and video quality estimation from E_c/N_0 , $QoE = \exp(\alpha \cdot E_c/N_0 + \beta)$. A log-linear model can be derived, $\ln(QoE) = \alpha \cdot E_c/N_0 + \beta$. The estimates for α and β can be computed through non-linear regression or linear regression and are almost the same. The estimates of linear regression are the optimal solution to

$$\min_{\alpha, \beta} \sqrt{\sum_{i=1}^n (\ln(QoE_i) - \alpha \cdot (E_c/N_0)_i - \beta)^2}, \quad (2)$$

where QoE_i , $(E_c/N_0)_i$ denotes the i -th samples of QoE, E_c/N_0 respectively.

The accuracy of the data used to compute them is limited. To deal with this problem, we propose to follow a robust optimization approach. We assume that that data $(E_c/N_0)_i$ incorporate errors and that their true value resides in a p -sized interval centered at $(E_c/N_0)_i$, $\rho \geq 0$. The robust estimates are the optimal solution to the min-max problem




Fig. 2 Robust optimization in non-linear regression for QoE prediction models.

$$\min_{\alpha, \beta} \max_{\forall i: |x_i - (E_c/N_0)_i| \leq \rho} \sqrt{\sum_{i=1}^n (\ln(QoE_i) - \alpha x_i - \beta)^2}. \quad (3)$$

The constraint $\forall i: |x_i - (E_c/N_0)_i| \leq \rho$ is equivalent to $\|\mathbf{x} - (\mathbf{E}_c/\mathbf{N}_0)\|_\infty \leq \rho$, where $\|\cdot\|_\infty$ is the infinity norm and \mathbf{x} , $(\mathbf{E}_c/\mathbf{N}_0)$ are the vectors which contain the x_i , $(E_c/N_0)_i$ respectively. Problem (3) is equivalent [15] to

$$\min_{\alpha, \beta} \sqrt{\sum_{i=1}^n (\ln(QoE_i) - \alpha(E_c/N_0)_i - \beta)^2 + \rho|\alpha|}, \quad (4)$$

which defines an l_1 -regularized regression estimator [13]. Using our data and SeDuMi [9], we computed the robust estimates for α and β for various values of the size ρ of the uncertainty set. The results can be seen in Fig. 2.

We observe that as the size of uncertainty ρ increases, α drops, namely the robust estimates suggests a weaker dependence between QoE and E_c/N_0 . This weaker dependence is compensated by a smaller fixed term β , as seen in Fig. 3. Thus, in the presence of errors, one should be more conservative in the prediction of QoE for a given E_c/N_0 . The rate of change for QoE with respect to E_c/N_0 estimated by linear or non-linear regression is too optimistic in the presence of errors. The regularized estimator of Eq. (4) takes this phenomenon into consideration by adding the trade-off term $\rho|\alpha|$. Our method connects the impact ρ of this trade-off with the size of the uncertainty sets for the data $(E_c/N_0)_i$. In this way, we can use the information




Fig. 3 Robust estimates of α and β parameters for various values of the size ρ of the uncertainty set.




Fig. 1 Drawing of a continuous roller furnace

multiple plates at the same time. The plates are passing the furnace continuously and one by one. In addition, they have to leave the furnace in the same order as they entered it because inside the furnace overtaking is not possible. Due to technical, physical and logistical restrictions, a plate cannot always enter the furnace directly after its predecessor. The minimal time a plate has to wait is equal to the duration its predecessor needs to completely enter into the furnace. In the following, we assume this duration to be constant ($t^{\min} = 8$ minutes).

One reason for additional waiting times is the finite furnace capacity. A plate cannot enter the furnace if there is not enough space. It has to wait until one or more plates leave the furnace, so that the free furnace capacity is sufficiently large. A further complication results from a special kind of setup procedure in cases when two consecutive plates need different heat treatment temperatures. Depending on the temperature deviation, these plates have to keep a specific distance from each other in the furnace. This distance is necessary in order to avoid an overheating or inhomogeneous heating of the plates. Apart from this, a minimal safety distance $s^{rc,\min}$ of two meters has always to be kept between two consecutive plates.

Another additional waiting time occurs if the processing time of the next plate needed to enter the furnace is smaller than the processing time of its predecessor. We already mentioned that overtaking is not possible inside the furnace. In addition, each plate has to keep its predefined holding time, so that processing times are never exceeded. As a consequence, the waiting time of a plate whose processing time is smaller than that of its predecessor is at least as long as the processing time deviation.

To the best of our knowledge, there is no study on this problem in the literature up to now. However, there are some studies having some parts in common with the HTFSPD (see [3], [10], [9]).

2 Problem formulation

Koné et al. [7] demonstrate different MILP models for the *Resource-Constrained Project Scheduling Problem* (RCPSP), two discrete-time formulations, a flow-based

$$x_{0n'0} + \sum_{j=0}^{n'} y_{0j0} = 2 \quad (14)$$

$$\sum_{j=0}^{n'} x_{n'jk} + y_{n'0k} = 2 \quad (15)$$

$$\sum_{e=0}^k x_{ije} + \sum_{e=0}^k x_{jie} \leq 1 \quad \forall i, j = 0, \dots, n' \quad (16)$$

$$\sum_{e=0}^k y_{ije} + \sum_{e=0}^k y_{jie} \leq 1 \quad \forall i, j = 0, \dots, n' \quad (17)$$

$$\sum_{e'=e}^k \sum_{j=0}^{n'} x_{ije'} + \sum_{e''=0}^{e-1} \sum_{j=0}^{n'} y_{ije''} \leq 1 \quad \forall i = 0, \dots, n', \quad e = 0, \dots, k \quad (18)$$

$$\sum_{e=0}^k y_{ije} \cdot e \leq \sum_{e=0}^k \sum_{j'=0}^{n'} y_{jj'e'} \cdot e \quad \forall i = 0, \dots, n, \quad j = 1, \dots, n' \quad (19)$$

$$\sum_{e=0}^k x_{ije} = \sum_{e=0}^k y_{jie} \quad \forall i, j = 0, \dots, n' \quad (20)$$

$$\sum_{e=0}^k x_{ije} \cdot e \geq \sum_{e=0}^k \sum_{j'=0}^{n'} x_{jj'e} \cdot e + \left(\sum_{e=0}^k x_{ije} - 1 \right) \cdot M \quad \forall i = 1, \dots, n', \quad j = 0, \dots, n' \quad (21)$$

$$r_0 = L - \sum_{i=0}^{n'} \sum_{j=0}^{n'} x_{ij0} \cdot (l_i + s_{ij}^{rc}) \quad (22)$$

$$r_e = r_{e-1} - \sum_{i=0}^{n'} \sum_{j=0}^{n'} (l_i + s_{ij}^{rc}) \cdot x_{ije} + \sum_{i=0}^{n'} \sum_{j=0}^{n'} (l_i + s_{ji}^{rc}) \cdot y_{ije} \quad \forall e = 1, \dots, k \quad (23)$$

$$t_e, r_e \geq 0 \quad \forall e = 1, \dots, k \quad (24)$$

$$x_{ije}, y_{ije} \in \{0, 1\} \quad \forall i, j = 0, \dots, n', \quad e = 0, \dots, k \quad (25)$$

In the above formulation, the objective function (1) minimizes the makespan. Constraints (2)-(21) assure that a feasible solution conforms with the continuous process described in Section 1. Constraints (22)-(23) make sure that the furnace capacity is never exceeded.

3 A genetic algorithm

A genetic algorithm (GA) is developed in order to solve real world problems. Its procedure is shown in [Figure 2](#).

The most individuals of the initial population are generated randomly. In order to increase the performance of the algorithm, the others are generated by some simple sorting methods. One of them corresponds to the method the *Dillinger Hütte GTS* applies in practice. First, the plates are sorted by their furnace temperatures. Then,

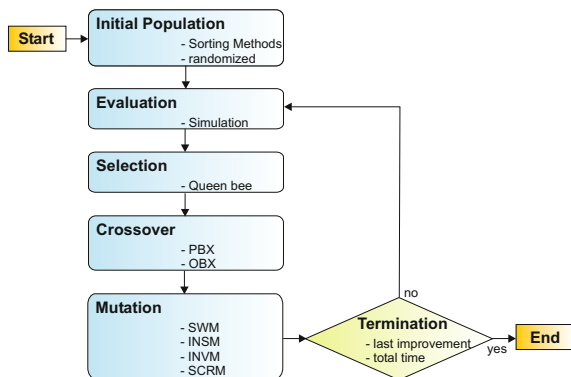


Fig. 2 A genetic algorithm

all plates with identical furnace temperatures are sorted by its processing time in ascending order. We refer to this heuristic in the following as the *rule-based sorting* (RBS).

The next step of the algorithm is to evaluate the individuals. A simulation of the furnace process serves as evaluation function.

The Selection scheme is similar to a proposal made by Karci [5]. It is based on the intelligence in bee swarms. The individual with the best fitness value of the current generation is the queen bee. All the other parents (selected by tournament selection) are crossed over with this fixed parent. For a survey about algorithms simulating bee swarm intelligence, we refer to Karaboga and Akay [4].

For crossover, we use the *Position based crossover operator* (PBX) with probability p_{pbx} and the *Order based crossover operator* (OBX) with probability p_{obx} . Kellegöz et al. [6] compare eleven crossover operators for the one machine n -jobs problem. Experimental results demonstrate the effectiveness of the OBX and PBX crossover operators. With probability p_{nox} , the selected individual does not cross with the queen bee.

As last step of each iteration, the individuals are mutated by the use of the *swap operator* (SWM) with probability p_{swm} , the *insert operator* (INSM) with probability p_{insm} , the *inverse operator* (INVM) with probability p_{invm} and the *scramble operator* (SCRM) with probability p_{scrm} . These operators can be found in [2].


There are two stopping criteria to terminate the algorithm. The first is if there has been no improvement since $stop_{limp}$ iterations whereas the second is a total run time limit of $stop_{time}$ seconds.

4 Computational results

ILOG CPLEX is used to solve the MILP formulation presented in Section 2. The genetic algorithm described in Section 3 has been implemented in C++. All compu-

time and personnel needed to serve these passengers is highly volatile, both over the day and over the week. Figure 1 shows a typical arrival profile with 4 peaks over the day and 2 different daily patterns over the week. In order to reduce costs it is thus desirable to smooth the demand by making use of the allowed service level.

Fig. 1 Two days of the same week. Check-In demand over the day (minutes 0 – 1440) for a smaller carrier on a major European airport. 4 large peaks, but different layout: strong morning peak over the week, strong evening peak on Sunday



Our goal is to smooth the irregular demand by making some passengers wait, i.e. we try to move peaks into valleys. We present a network flow formulation for this problem and discuss additional constraints relevant for real-world scenarios. We show how to solve the model and present numerical results on real-world data.

The model is part of the Inform GroundStar suite. GroundStar is an integrated resources management system to optimize all planning and control processes in aircraft handling at passenger airports. GroundStar is in successful use in various handling areas of more than 200 airports of every size world-wide.

There is not much literature on check-in counter planning (consisting of calculating the demand and allocating the counters). Check-In counter allocation problems have been studied by [5]. They use a simulation for generating the appropriated input. [6] work with a number of passengers for a flight without assumptions on their distribution or on waiting times. To our knowledge, this problem has also not been considered in queuing theory.

2 A Network Model

We consider a planning horizon of T minutes. Typically, we plan for 10 days ($T = 14400$). Let n_t denote the demand at time t . The queue length, Q , is the maximum time a demand is allowed to wait before being served, i.e. demand at time t should be served at time t' with $t \leq t' \leq t + Q$. As this is not always possible we allow violating queue length Q . However, demand at time t must being served before $t + L$, where $L \geq Q$ is a limit of e.g. 6 hours.

possible, etc. A better estimation of $h(t)$ is used if the estimation above yields large numbers. And there are additional rules, e.g. as soon as demand is there at least one check-in counter must be opened.

3 Numerical Results

The above model is implemented in C++ using the COIN-OR LP solver Coin CLP 1.11 [2].¹ Run time is measured on an Intel Core2, 2.8GHz, 3GB Ram on Windows XP using Visual Studio 2008. We start with an artificial test case, Fig. 3. The demand is given as a double sinus wave over 1001 minutes ($n_t := 12 + 10 \sin(t \frac{4\pi}{1000})$, $t = 0 \dots 1000$). The two lines show the resource requirement for a maximum waiting time of 20 and 100 minutes, respectively². No departure flights are assumed. As can be seen in the statistics (Table 1), enlarging the queue length increases the number of edges between I and O and thus the number of LP variables. Nevertheless, run times of less than 2 seconds are no problem for the overall planning process. Quality-wise, the result is good as the number of resources only required for short time intervals is small and only up to 1.6% additional demand is induced by rounding (i.e. the demand of $\approx 12,000$ minutes is only increased by 172 or 192 minutes, resp. when rounding).

The figure on the right shows one day of a real-world scenario originally consisting of 9 days, 149 departure flights. The queue length is $Q = 24$. Whereas the number of variables scales with the number of days when comparing artificial and real model, the latter introduces new edges between C and Z . These variables ensure that all demand is served before departure. Despite their large number, they do barely impact running time here.




Fig. 3 Artificial example (right), 1 day from a real-world scenario (left)

¹ The model above could be adequately solved by a pure network solvers (e.g. [3, 4]). We prefer the more generic LP solver as that gives us some more freedom for future enhancements and constraints.

² waiting time is typically ≤ 1 hour on real airports. $Q = 100$ is chosen for illustration here.

and constraint relations shown on the spreadsheet are for display purposes only, and are not used by Solver.

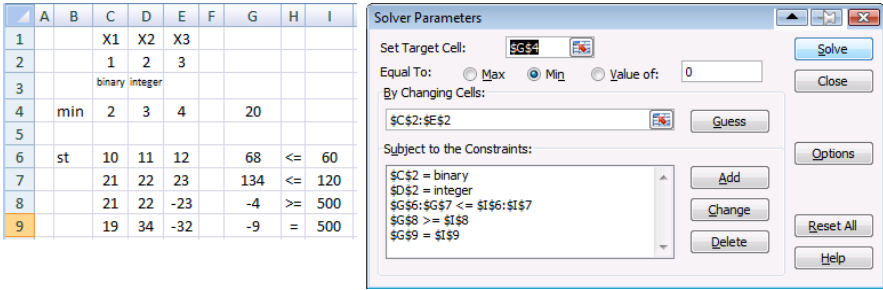


Fig. 1 A typical spreadsheet optimization model (left) and the Solver entry for this. Many models do not follow this layout, but instead ‘hide’ the model inside the spreadsheet formulae.

2 Constructing a Mathematical Model from a Solver Model

The means by which Solver stores a model does not appear to be documented. However, using the Excel Name Manager add-in [7], developed by Jan Karel Pieterse of Decision Models UK, shows that Solver uses hidden ‘names’ to contain all the model’s details. OpenSolver reads these values to determine the cells that define the model and Solver options.

Excel’s representation of the optimisation model is given in terms of cells that contain constants and formulae. Because OpenSolver restricts itself to linear models, we wish to analyse the spreadsheet data to build an optimisation model with equations of the form:

$$\begin{aligned} \text{Min/max } & c_1x_1 + c_2x_2 + \dots + c_nx_n \\ \text{Subject to } & a_{i1}x_1 + a_{i2}x_2 + \dots + a_{in}x_n \geq/\leq b_i, \quad i = 1, 2, \dots, m \end{aligned}$$

where \geq/\leq denotes either \geq , $=$ or \leq . Assuming the model is linear, then the Excel data can be thought of as defining an objective function given by

$$\text{Obj}(\mathbf{x}) = c_0 + c_1x_1 + c_2x_2 + \dots c_nx_n$$

where $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is the vector of n decision variables, $\text{Obj}(\mathbf{x})$ is the objective function cell value, and c_0 is a constant. Similarly, each constraint equation i is defined in Excel by

$$\text{LHS}_i(\mathbf{x}) \geq/\leq \text{RHS}_i(\mathbf{x}) \Rightarrow \text{LHS}_i(\mathbf{x}) - \text{RHS}_i(\mathbf{x}) \geq/\leq 0, \quad i = 1, 2, \dots, m$$




Fig. 2 OpenSolver's buttons and menu appear in Excel's Data ribbon.

OpenSolver and CBC files can be copied to the appropriate Microsoft Office folder to ensure OpenSolver is available every time Excel is launched.

Users can construct their models either using the standard Solver interface or using a new OpenSolver dialog. The new dialog provides a number of advantages, including highlighting of selected constraints on the sheet, and easier editing of constraints.

We have found OpenSolver's performance to be similar or better than Solver's. CBC appears to be a more modern optimizer than Solver's, and so gives much improved performance on some difficult problems. For example, large knapsack problems which take hours with the Excel 2007 Solver are solved instantly using OpenSolver, thanks to the newer techniques such as problem strengthening and pre-processing used by CBC [1].

To review an optimisation model developed using the built-in Solver, the user needs to check both the equations on the spreadsheet and the model formulation as entered into Solver. This separation between the equations and the model form makes checking and debugging difficult. OpenSolver provides a novel solution to this in the form of direct model visualisation on the spreadsheet. As [Figure 3](#) shows, OpenSolver can annotate a spreadsheet to display a model in which the objective cell is highlighted and labeled min or max, the adjustable cells are shaded with any binary and integer decision variable cells being labeled 'b' and 'i' respectively, and each constraint is highlighted and its sense shown. We have found this model visualisation to be very useful for checking large models.

4 Automatic Model Construction

We have developed additional functionality that allows OpenSolver to build Solver-compatible models itself without requiring the usual step-by-step construction process. Our approach builds on the philosophy that the model should be fully documented on the spreadsheet. Thus, we require that the spreadsheet identifies the objective sense (using the keyword 'min' or 'max' or variants of these), and gives the sense (\leq , $=$, or \geq) of each constraint. Our example spreadsheet shown in [Figure 1](#) satisfies these layout requirements.

To identify the model, OpenSolver starts by searching for a cell containing the text 'min' or 'max' (or variants of these terms). It then searches the cells in the

vicinity of this min/max cell to find a cell containing a formula (giving preference to any cell containing a ‘sumproduct’ formula); if one is found, this is assumed to define the objective function. The left and/or right hand side formulae for the constraints are then located in a similar fashion by searching for occurrences of \leq (or ‘<’), $=$ and, \geq (or ‘>’). The predecessor cells of all these formulae are then found, and the decision variables are then taken as those predecessor cells that have as successors either (1) at least two constraints or (2) the objective and at least one constraint.

The final step is to identify any binary or integer restrictions on the decision variables. These are assumed to be indicated in the spreadsheet by the text ‘binary’ or ‘integer’ (and variants of these) entered in the cells beneath any restricted decision variables of these types.

5 Advanced Features

OpenSolver offers a number of features for advanced users, including:

- The ability to easily solve an LP relaxation with a single menu click,
- Interaction with the COIN-OR CBC solver via the command line,
- Faster running using ‘Quick Solve’ when repeatedly solving the same problem with a succession of different right hand sides,
- Viewing of the CBC .lp input file showing the model’s equations, and,
- Detection and display of non-linearities in the model.

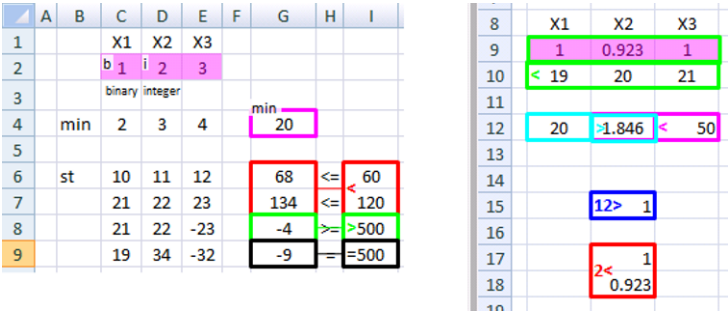


Fig. 3 OpenSolver can display an optimisation model directly on the spreadsheet. The screenshot on the left shows OpenSolver’s highlighting for the model given earlier, while the screenshot on the right illustrates OpenSolver’s highlighting for several other common model representations.

The buyer chooses the fraction β of used items to be collected and remanufactured (collection rate, $0 \leq \beta \leq 1$) as well as the order size q .

The relevant costs of the supplier and the buyer consist of fixed costs per order, stock holding costs and, in the buyer's case, disposal costs. The evolution of the inventory levels at the buyer is shown in Figure 1, where $t^c = q/D$ is the order cycle length. It is straightforward to see that the buyer's cost function expresses then as

$$TC_b(q, \beta) = s_b \cdot \frac{D}{q} + \frac{q}{2} \cdot (h_b + \beta \cdot u_b) + d_b \cdot (1 - \beta) \cdot D.$$

The structure of the inventory levels at the supplier depends on what operation is run first within each order cycle. Figure 2 illustrates the inventory levels during an order cycle with manufacturing executed prior to remanufacturing; t_2 and t_3 represent there the respective durations, and $t_1 = t^c - t_2 - t_3$ is the slack time. Due to the assumption $P_M, P_R > D$, it holds $t_1 > 0$. The expression of the supplier's costs accordingly depends on the sequence of operations during an order cycle [4]:

$$TC_v(q, \beta) = \begin{cases} s_v \cdot \frac{D}{q} + \frac{q}{2} \cdot h_{MR}(\beta), & \text{if manufacturing is run first} \\ s_v \cdot \frac{D}{q} + \frac{q}{2} \cdot h_{RM}(\beta), & \text{otherwise} \end{cases}$$

where

$$h_{MR}(\beta) = \beta^2 \cdot \left[h_v \cdot \left(\frac{D}{P_M} - \frac{D}{P_R} \right) - u_v \frac{D}{P_R} \right] - 2\beta \cdot \left[h_v \cdot \left(\frac{D}{P_M} - \frac{D}{P_R} \right) - u_v \right] + h_v \frac{D}{P_M},$$

$$h_{RM}(\beta) = \beta^2 \cdot \left[(h_v - u_v) \cdot \left(\frac{D}{P_R} - \frac{D}{P_M} \right) + u_v \frac{D}{P_M} \right] + 2\beta u_v \cdot \left(1 - \frac{D}{P_M} \right) + h_v \frac{D}{P_M}.$$




Fig. 1 Inventory levels of serviceables (left) and nonserviceables (right) at the buyer over time




Fig. 2 Inventory levels of serviceables (left) and nonserviceables (right) at the supplier; manufacturing precedes remanufacturing

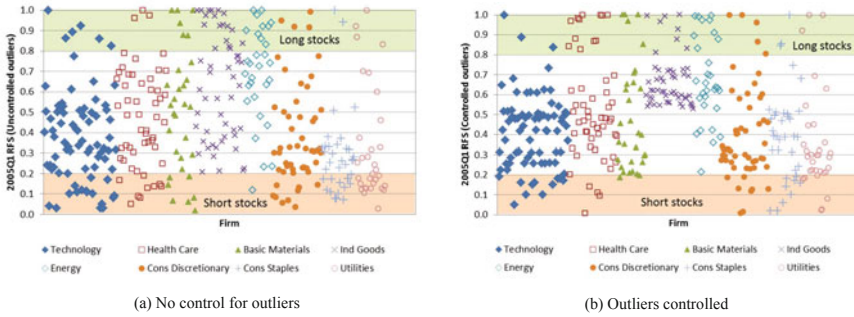


Fig. 1 2005Q1 expected firm-RFS scores




Fig. 2 Out-of-sample performance of RFS-based and ETF portfolios

References

1. Banker, R.D.: Maximum likelihood, consistency and DEA: Statistical foundations. *Management Science* **39** 1265–1273 (1993)
2. Banker, R.D., Natarajan, R.: Evaluating contextual variables affecting productivity using Data Envelopment Analysis. *Operations Research* **56** 48–58 (2008)
3. Charnes, A., Cooper, W.W., Rhodes, E.: Measuring the efficiency of decision-making units. *European Journal of Operational Research* **2**, 429–444 (1978)
4. Dusansky, R., Wilson, P.W.: On the relative efficiency of alternative modes of producing public sector output: The case of the developmentally disabled. *European Journal of Operational Research* **80** 608–628 (1995)
5. Edirisinghe, N.C.P., Zhang, X.: Portfolio selection under DEA-based relative financial strength indicators: Case of US industries. *Journal of the Operational Research Society* **59** 842–856 (2008)
6. Fama, E.F.: Efficient capital markets: A review of theory and empirical work. *Journal of Finance* **25** 383–417 (1970)
7. Huang, M., Li, S.X.: Stochastic DEA models with different types of input-output disturbances. *Journal of Productivity Analysis* **15** 95–113 (2001)
8. Olesen, O.B., Petersen, N.C.: Chance constrained efficiency evaluation. *Management Science* **41** 442–457 (1995)
9. Post, T.: Performance evaluation in stochastic environments using mean-variance Data Envelopment Analysis. *Operations Research* **49** 281–292 (2001)
10. Sengupta, J.K.: Efficiency measurement in stochastic input-output systems. *International Journal of Systems Science* **13** 273–287 (1982)
11. Wilson, G.W., Jadow, J.M.: Competition, profit incentives, and technical efficiency in the provision of nuclear medicine services. *The Bell Journal of Economics* **13** 472–482 (1982)

ture is that it allows multi step forecasts. Additionally, by training an ensemble of HCNNs, we get a distribution of possible outcomes. In this paper we will *not* primarily evaluate the merits of HCNNs as a forecasting model. See, e.g., [5, 6, 8] for detailed performance evaluations. Instead we focus on supporting the decision maker in interpreting the considerable amount of data that the forecast produces.

2 Basic Visualization Techniques for Ensemble Forecasts

Typically, we visualize time series ensembles by plotting summary data, see for example [1–4]. As a first step, we can simply plot the mean or median, see [figure 1](#). The figure shows a 20-day ahead forecast for the price of natural gas in US-Dollar. The exact forecast asset is not central to the following discussion. We continue using the same 20-day ahead natural gas price forecast throughout the paper. Mean and median can, of course, differ. The general visual impression of [figure 1](#) is that of strong downtrend. The figure does not convey any distributional information.

This changes slightly when we add representative percentiles, like the quartiles, see [figure 2](#). Note an interesting feature: future uncertainty, as measured by the difference of maximum and minimum, does not necessarily increase. In fact, uncertainty decreases during the last five forecast days. The figure now only conveys the visual impression of a slight downtrend, due to the width of the distribution. Nevertheless, we still have no idea of the distribution of individual paths.

An additional step is to plot every path of individual ensemble members, see also [7]. This leads to [figure 3](#). For clarity, the figure also shows mean and median in bold lines. It becomes apparent that we do not gain much by plotting every path. To the contrary, the information becomes less clear because the paths overlap. All we can see is that the distribution is dense in the middle and less dense to its borders. We might therefore visually conclude that the distribution is unimodal — a possibly dangerous conclusion as we will see later on. The approach of plotting every ensemble member does not scale well. It becomes unwieldy when the ensemble includes several hundred members. The figure is, e.g., a plot of 190 individual paths.

Fig. 1 Average and median of neural network ensemble forecast.




Fig. 2 Neural network ensemble forecast. The figure shows from the top: maximum, 75 percent quartile, mean and median, 25 percent percentile, minimum.





Fig. 3 All paths of neural network ensemble forecast. Plot of 190 ensemble members.



3 Prototypical Visualization Environment with Heatmaps

The question arises: how can we present the forecast information to a decision maker in an intuitive way but still exploit all distribution information? We propose a heatmap style visualization. A heatmap allows to differentiate between more active and less active regions of the forecast space by color coding, see [figure 4](#). We obtain the color values as follows:

- At every timestep compute the maximum of ensemble members within an (adjustable) ε -environment.
- Optionally divide the number of ensemble members within an ε -environment by the maximum for the specific timestep.
- Assign the color red to a ratio of 1 (the maximum) and light blue to a ratio of almost 0.
- Values with no ensemble members at all in the given ε -environment are color coded in white.

Therefore, red regions indicate regions of high activity according to the forecast. For clarity, [figure 4](#) also shows mean, median, maximum and minimum in thick lines. We note that often, but not always, the mean coincides with red regions. However,

we also note that approximately from day 5 to day 15 the red region splits into several paths and becomes quite large. This becomes especially clear in a smoothed version of the heatmap, see [figure 5](#). The conclusion is that according to the forecast the mean does not represent the distribution well at all, because (especially for days 6 to 8) the distribution is bimodal or multi modal. This significantly changes the interpretation of the decision maker. Looking only at [figures 1–3](#) it was apparent that the forecast for days 6 to 8 is a slight downtrend. Looking at the heatmap we see that the correct answer actually is: the model doesn't know! This is a warning to the decision maker. On the other hand, during the last few forecast days (days 16–20) the uncertainty becomes smaller and the model clearly forecasts a downtrend. Remember, that we are not dealing with realized forecast accuracy. We are just exploiting the forecast information a priori.

We implement heatmap style visualization in a prototype, see the screenshot in [figure 6](#). For plotting the open source software Gnuplot (version 4.2 patchlevel 6) is used. The right part of the program window controls several plot parameters. Especially, vertical and horizontal plot resolution can be configured. Higher parameters produce smoother heatmaps, see 4 and 5 for two extremes. The screenshot 6 shows an intermediate setting. At first sight no additional information seems to be gained by interpolating between different forecast steps. However, the smoothed result allows to follow intersecting paths. See, for example, the split path from timestep 8 to 9 in [figures 4 and 5](#). Another important parameter is the class radius divider which allows to set the width of the mentioned ε -environment. Finally, different checkboxes allow the plot additional information, like extrema, percentiles, etc.




Fig. 4 Heatmap style forecast visualization.




Fig. 5 Smoothed heatmap.

4 Conclusions and Outlook

This paper presents a visualization approach for distributional forecasts. Individual distributional forecasts often arise in the context of neural network ensembles. However, the same visualization style could be useful for other forecasts. We see, that mean and median do not necessarily confer the right information because the distribution may split. This makes mean and median bad representatives. We also see, that forecast uncertainty does not uniformly increase with future timesteps. A split distribution may actually become unimodal again.

An apparent alternative to heatmap style visualization would be a three dimensional plot. This would replace the color coded information by height information. However, the disadvantage of three dimensional plots is that often parts of the plot may hide other parts. Generally, getting adequate information out of a three dimensional plot is more difficult. Heatmap style visualization provides a good compromise between information density and visual interpretability.

The paper only focuses on *visual* representation. It would be interesting to analyze, if we could also *quantify* the advantage of better using the forecast information. This involves, e.g., identifying peaks of the distribution and benchmarking a tri-state model forecast (increases, decreases, don't know) against the realization. Other means of better quantifying the forecast information can, of course, be devised.




Fig. 6 Screenshot of the prototypical visualization environment (alternative heatmap style).

References

1. Andrienko, G., Andrienko, N.: Visual Exploration of the Spatial Distribution of Temporal Behaviors. In: Proceedings of the Ninth International Conference on Information Visualisation, pp. 799–806 (2005) doi: 10.1109/IV.2005.135
2. Andrienko, G., Andrienko, N., Mladenov, M., Mock, M., Poelitz, C.: Extracting Events from Spatial Time Series. In: Proceedings of the 14th International Conference on Information Visualisation, pp. 48–53 (2010) doi: 10.1109/IV.2010.17
3. Buono, P., Plaisant, C., Simeone, A., Aris, A., Shneiderman, B., Shmueli, G., Jank, W.: Similarity-Based Forecasting with Simultaneous Previews: A River Plot Interface for Time Series Forecasting. In: Proceedings of the 11th International Conference Information Visualization (2007) doi: 10.1109/IV.2007.101
4. Feng, D., Kwok, L., Lee, Y., Taylor II, R.M.: Matching Visual Saliency to Confidence in Plots of Uncertain Data. IEEE Transactions on Visualization and Computer Graphics (2010) doi: 10.1109/TVCG.2010.176
5. von Mettenheim, H.-J. Advanced Neural Networks: Finance, Forecast, and other Applications. Leibniz Universität Hannover (2009)
6. von Mettenheim, H.-J., Breitner, M.H.: Robust Decision Support Systems with Matrix Forecasts and Shared Layer Perceptrons for Finance And other Applications. In: ICIS 2010 Proceedings (2010)
7. Potter, K., Wilson, A., Bremer, P.-T., Williams, D., Doutriaux, C., Pascucci, V., Johnson, C.R.: Ensemble-Vis: A Framework for the Statistical Visualization of Ensemble Data. In: International Conference on Data Mining Workshops, pp. 233–240 (2009) doi: 10.1109/ICDMW.2009.55
8. Zimmermann, H.G., Grothmann, R., Tietz, C., von Jouanne-Diedrich, H.: Market Modeling, Forecasting and Risk Analysis with Historical Consistent Neural Networks. In: Selected Papers of the Annual International Conference of the German Operations Research Society, pp. 531–536 (2010) doi: 10.1007/978-3-642-20009-0_84

2 Causal-Retro-Causal Recurrent Neural Networks

2.1 Causal Recurrent Neural Networks

To derive the CRCNN, we let us first introduce causal RNN for closed dynamical systems, a NN for modeling the dynamics of observables Y_τ in a causal formulation [ZGTJ10]. We formulate the RNN as a state space model in discrete time τ that describes the observables Y_τ by using a state transition and output equation [ZGTJ10]:

$$\text{state transition } s_\tau = f(s_{\tau-1}) = \tanh(As_{\tau-1}), s_0 \quad (1)$$

$$\text{output equation } Y_\tau = g(s_\tau) = [Id, 0]s_\tau \quad (2)$$

$$\text{system identification } \text{Error} = \sum_{\tau=t-m}^t (Y_\tau - Y_\tau^d)^2 \rightarrow \min_A \quad (3)$$

The time-recurring state transition equation $s_\tau = f(s_{\tau-1})$ describes the current state s_τ solely dependent on the previous system state $s_{\tau-1}$. The observables Y_τ are derived from the state s_τ using the output equation g . Without loss of generality we can approximate the functions f and g of the state space model with a NN, i. e. $\tanh(As_{\tau-1})$ and $[Id, 0]s_\tau$ [SZ06, ZGN02].

We use the technique of finite unfolding in time to transform the temporal equations into a spatial architecture. The RNN is **unfolded** across the entire time path, i. e. we learn the unique history of the system. Fig. 1 depicts the resulting RNN. The




Fig. 1 Architecture of the Causal Recurrent Neural Network

RNN is trained using error back propagation through time (BPTT) [Wer74] together with an architecture based formulation of teacher forcing (TF) [ZGTJ10].

We applied the causal RNN to forecast prices of energy and copper futures traded on the European Energy Exchange (EEX) and the London Metal Exchange (LME) over 20 days. Fig. 2 compares the RNN forecasts with the actual prices for two forecast horizons in **sequence**.

As depicted in Fig. 2 the causal RNN gives a fairly good forecast of the EEX (LME) market dynamics in case of the first (second) forecast horizon, while it fails in case of the second (first) time horizon. Now the question arises, how we can improve the forecasting accuracy of the causal RNN.

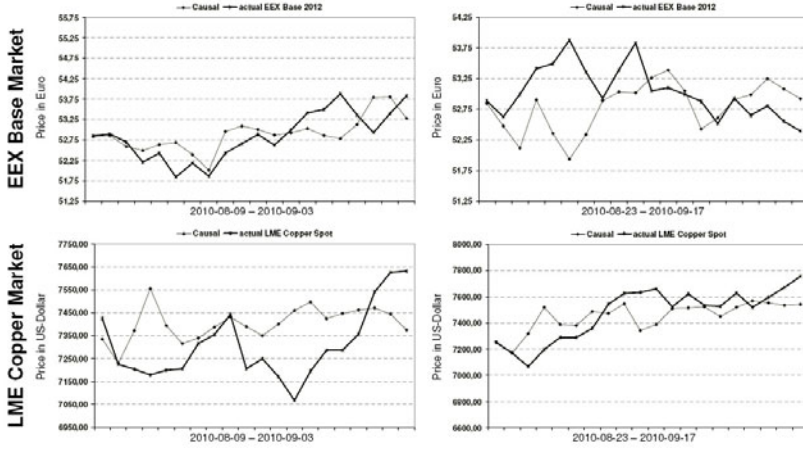


Fig. 2 Forecasting EEX and LME future contracts over 20 days with causal RNNs

2.2 Retro-Causal Recurrent Neural Networks

Causal models might fail due to the impact of utility maximization, which implies an information flow from the future backwards to make an optimal decision at present time. Thus we formulate a model with a retro-causal information flow:

$$\text{state transition } s'_\tau = \tanh(A's'_{\tau+1}), s'_T \quad (4)$$

$$\text{output equation } Y_\tau = [Id, 0]s'_\tau, \quad (5)$$

In the RCNN the internal state s'_τ depends on $s'_{\tau+1}$. The observables Y_τ are part of s'_τ . The final state s'_T and the transition matrix A' are trained such that the historical system behavior is described. Fig. 3 depicts the architecture of the RCNN.




Fig. 3 Architecture of the Retro-Causal Recurrent Neural Network (RCNN)

Likewise to the causal RNN, the RCNN in Fig. 3 is also trained using the BPTT algorithm and an architecture based formulation of TF [ZGTJ10]. We applied the RCNN model on the same time periods as the causal RNN (Fig. 4).

In cases the causal RNN fails, the RCNN is able to give an accurate forecast and vice versa. However, *ex ante* we do not if we are in a causal or retro-causal regime.

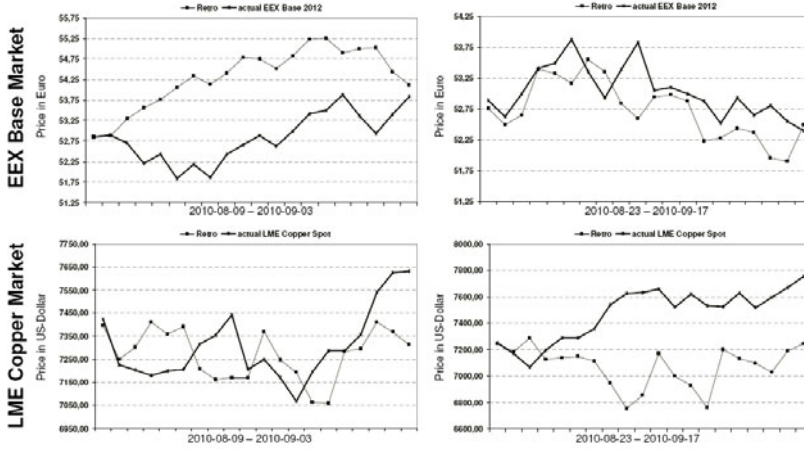


Fig. 4 Forecasting EEX and LME future contracts over 20 days with RCNNs

2.3 Causal-Retro-Causal Recurrent Neural Networks

Causal-retro-causal neural networks (CRCNN) combine a causal and a retro-causal information flow in an integrated NN model given by

$$\text{causal state transition } s_\tau = \tanh(As_{\tau-1}), s_0 \quad (6)$$

$$\text{retro-causal state transition } s'_\tau = \tanh(A's'_{\tau+1}), s'_T \quad (7)$$

$$\text{output equation } Y_\tau = [Id, 0]s_\tau + [Id, 0]s'_\tau. \quad (8)$$

The dynamics Y_τ is explained in the CRCNN by a sequence of causal s_τ and retro-causal states s'_τ using a transition matrices A and A' for the causal and retro-causal information flow (Eq. 6 and Eq. 7). Fig. 5 depicts the architecture of the CRCNN.




Fig. 5 Architecture of the Causal-Retro-Causal Recurrent Neural Network (CRCNN)

To train the CRCNN we use a combined teacher forcing approach on both branches, resulting in an extended CRCNN architecture (Fig. 6).




Fig. 6 Extended CRCNN incorporating a Teacher Forcing (TF) mechanism

The causal as well as the retro-causal branch have to explain only the incompleteness of the opposite side. Thus, we have a moving target problem, where the TF has to take into account the behavior of the opposite side. If the error approaches zero, TF is inactive and thus, the extended architecture converges to the original CRCNN (Fig. 5). The advantage of extended CRCNN (Fig. 6) is that it allows a fully dynamical superposition of the causal and retro-causal information flows. We have to state that the numerical solution is difficult, because the extended CRCNN contains many closed loops. Mathematically this means to identify a dynamical system on a manifold, where we have to find the dynamics and the manifold in parallel.

Fig. 7 depicts the application of the CRCNN to the two consecutive time periods.

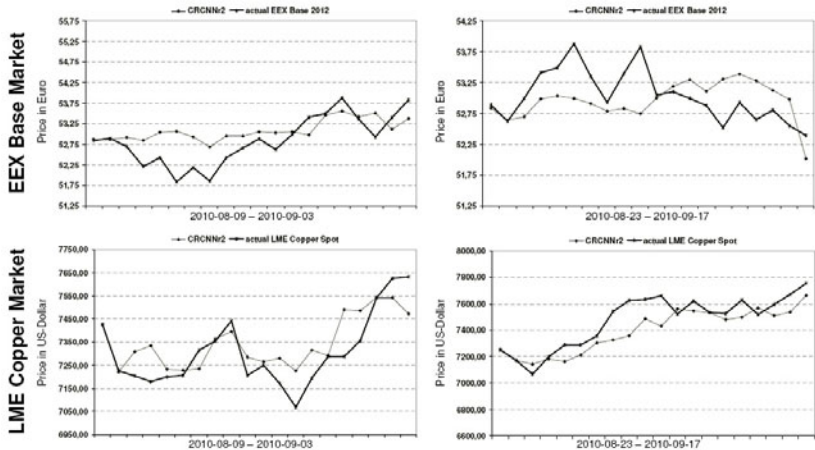


Fig. 7 Forecasting EEX and LME future contracts over 20 days with CRCNNs

The combination of the causal and retro-causal information flow within an integrated NN is able to give an accurate prediction for the future contracts under consideration. Remarkably, the CRCNN is able to detect beforehand if the market is in an causal or retro-causal regime as well as a mixture thereof, since the causal and retro-causal information flows are dynamically combined.

When parameterizing the adapted MNS procedure, we found a **contrary** behavior between the two **objectives** *minimization of delay* and *minimization of empty travel time*. If more relative weight is put on one of both objective criteria, this produces the desired improvement; however, simultaneously the solution quality of the other objective decreases by nearly the same magnitude.




Fig. 1 Pareto curves for different simulation speeds

Figure 1 illustrates this finding: The black square indicates the benchmark value in terms of delay (x-axis) and empty travel time (y-axis). The triangles show the results of different penalty cost parameter settings and different simulation speeds. For a simulation speed of $s = 5$ (black triangles), it can be observed that cost parameters (30,10) and (20,5) result in reduced delay, while the cost parameters (40,8), (30,5) and (40,5) result in a reduced empty travel time. However, there is no black triangle solution that reaches the *preferred area* with simultaneous improvement of both objective criteria. The *preferred area* is only achieved with a simulation speed of $s = 1$ (real-time simulation, grey triangles).

It is interesting to observe the general **impact of simulation speed** on overall solution quality: The slower the simulation is run, the better the solution quality. This finding can be exemplarily investigated for cost parameter (40,8): From simulation speed $s = 120$ (blank triangles) to simulation speed $s = 5$ (black triangles) empty travel time is reduced by 5.1% and delay is reduced by 23.4%. From $s = 5$ (black triangles) to $s = 1$ (grey triangles) empty travel time and delay are reduced by an additional 2.2% and 6.6%, respectively. This finding is due to the MNS procedure's improvement framework: additional time is efficiently used to perform extra calculations resulting in an improved overall solution quality.

Figure 2 shows a typical planning result generated by the adapted MNS procedure. Such type of **graphical illustrations** were used to discuss and verify the planning results with the cooperating freight forwarding company. The *first loaded trip* directs the vehicle from Reutlingen (Germany) at point A, to Miskolc (Hungary) at point B. Afterwards, the vehicle has to perform an empty trip from Miskolc (Hungary) at point B to Mosonszolnok (Hungary) at point C. Here, the vehicle gets its *second loaded trip* from Mosonszolnok (Hungary) at point C, to Vienna (Austria) at point D. In Vienna (Austria) at point D, the *third loaded trip* is directly available: Vienna (Austria) at point D, to Mannheim (Germany) at point E. And so on...



Fig. 2 Exemplary five-week vehicle tour – geographical illustration

Due to the contrary objectives, it is not possible to report *one* best result. Instead, the **best results** of three penalty cost combinations are presented in Figure 3:

For penalty cost combination (40,8) empty travel time is reduced by 3.9% and delay by 6.0%. More relative weight is put on the reduction of empty travel time with penalty cost combination (30,5): Consequently, a reduction of 7.3% in empty travel time is now achieved. However, the reduction in delay decreases to only 1.7%. The highest reduction of empty travel time is achieved with penalty cost combination (40,5): 11.9%. However, at the same time, delay increases by 16.6%.

Our cooperating freight forwarding company chose the results of penalty cost combination (30,5) as their “preferred ones”.




Fig. 1 Distance vs. game theoretical prices

tance and our game theoretical prices. Obviously, it is impossible to consider all of the $2^{85} - 1$ coalitions. The best thing that we can do is to sample and create a pool of some “essential coalitions” among them. Based on this pool one can show graphically a representative comparison between the two price vectors. The picture on the left side represents the relative profits of 7084 essential coalitions in the pool with the two price vectors, while the picture on the right side is just a zoom of the first one. They show that our prices for all of the 100 coalitions which have the worst relative profit with the distance price vector increase significantly. Many of them have even a relative profit of more than 10% with our prices. This example shows clearly that using our game theoretical approach one can come up with price systems that constitute a good compromise between fairness and enforceability and are better (fairer) than the ad-hoc allocations.

Acknowledgements I would like to express my sincere gratitude to my advisors, Prof. Dr. Dr. h.c. mult. Martin Grötschel and Dr. habil. Ralf Borndörfer, for the interesting research theme and for their valuable supports and suggestions. I am grateful to the Zuse Institute Berlin for providing me a Konrad-Zuse Scholarship.

References

1. Bussieck, M. R.: Optimal Lines in Public Rail Transport. PhD Thesis, TU Braunschweig (1998)
2. Dantzig, G. B.: On the Significance of Solving Linear Programming Problems with Some Integer Variables. *Econometrica* **28**, 30–44 (1960)
3. Faigle, U., Kern, W., Kuipers, J.: On the computation of the nucleolus of a cooperative game. *Internat. J. Game Theory* **30**, 79–98 (2001)
4. Faigle, U., Kern, W., Paulusma, D.: Note on the Computational Complexity of Least Core Concepts for min-cost Spanning Tree Games. *Math. Methods of Operations Research* **52**, 23–38 (2000)
5. Grötschel, M., Lovász, L., Schrijver, A.: Geometric algorithms and combinatorial optimization, 2. corr. ed., Springer (1993)
6. Hallefjord, A., Helming, R., Jørnsten, K.: Computing the Nucleolus when the Characteristic Function is Given Implicitly: A Constraint Generation Approach. *International Journal of Game Theory* **24**, 357–372 (1995)
7. Hoang, N. D.: Algorithmic Cost Allocation Games: Theory and Applications. PhD Thesis, TU Berlin (2010)

Operations Research Proceedings 2011
Selected Papers of the International Conference on
Operations Research (OR 2011), August 30 -
September 2, 2011, Zurich, Switzerland
Klatte, D.; Lüthi, H.-J.; Schmedders, K. (Eds.)
2012, XV, 607 p., Softcover
ISBN: 978-3-642-29209-5