OCP - Operational Curricular Planning: A Visual Decision Support System for Planning Teaching Resources at Universities

1st Raphael Sahann
Visualization and Data Analysis
University of Vienna
Vienna, Austria
raphael.sahann@univie.ac.at

2nd Torsten Möller
Visualization and Data Analysis
University of Vienna
Vienna, Austria
torsten.moeller@univie.ac.at

***Abstract*—We conducted a design study to do an in-depth analysis of the problem of operational planning at universities and designed a decision support tool for that problem, called Operational Curricular Planning (OCP). Based on our observations we abstracted the planning process into separate tasks. Focusing on a subset of tasks that we characterized, we present the OCP tool for visually supporting decision making in the process of planning teaching resources. We show the steps leading to the final design of our visual decision support system and discuss the design decisions made while building the tool. Finally, we present an evaluation with four domain experts in a real-world scenario and talk about lessons learned from building the OCP tool, including the issue of integration and adoption of the system.**

***Index Terms*—design study, university planning, decision support system**

I. Motivation

Planning of teaching resources gets increasingly difficult with the constantly rising number of students at universities. Universities need to offer enough courses to provide places for all students while staying within the budget. Old systems and static reports, which are the current state of the art, complicate planning matters.

We conducted a series of interviews with domain experts of different planning related positions within our university and found out that planning of teaching resources is currently not supported by any tool. It is also only partly covered by existing literature. Room finding, timetabling and scheduling are important parts of teaching planning that are already extensively covered, but predicting student numbers and adapting course capacities accordingly is lacking means of (tool-assisted) support.

The problem of planning teaching resources is based on estimations and knowledge from experts and is heavily influenced by external factors. In order to offer enough groups per course to fit all students, the number of students has to be estimated before registration starts. This task does not rely on a concrete mathematical model, but is just a conglomeration of internal and external factors that change independently from each other. An example for an external factor might be the media naming data scientist the job of the year and causing a rush of new students into the data science study. An internal event could be a far below average success rate in a course that causes many students to re-take that course. Assistance through prediction is possible, but the task cannot be fully automated, because it needs the input of an expert who has knowledge about these events and can factor them in.

Planning related support systems are mainly looking into scheduling and timetabling tasks, which are both algorithmically solvable and do not predict future student numbers. Planning of teaching resources at our university is currently done by manually extracting information out of an event-based database export spanning multiple spreadsheets. Experts need to filter and add up necessary values by hand in order to come to conclusions. The task has not yet been tackled by a visual decision support system.

In this paper we give a detailed domain description of planning teaching resources at universities in section IV. Based on the defined tasks we then present the Operational Curricular Planning (OCP) tool, which facilitates decision making in this domain by preprocessing the available data and visualizing the parts necessary for planning. We show the process of building the tool and give insight into design decisions in sections III and V. A final evaluation in section VII of the tool shows its strengths in helping experts. The tool is currently being adapted to fit into the existing university system. A discussion about its integration and adoption concludes the paper in section VIII.

II. Related Work

Room planning and timetabling are tasks that need to be addressed in many fields. A notable field is the medical domain, namely planning of operation rooms for surgeries. Cardoen et al. [6] provide a broad overview of the literature regarding the problems faced and techniques used in operating room planning. The general task described as scheduling is to fit as many interdependent subtasks in the shortest possible execution time [15], which is relevant in almost all construction and manufacturing processes.

Timetabling and scheduling has also been tackled in planning teaching resources at universities. Boronico [3] shows a mathematical model in conjunction with a discrete event simulation model that projects student enrollment for courses to suggest faculty schedules. Bonutti et al. [2] present a tool that can visualize all viable combinations of rooms and courses to minimize overlap and maximize room utilization. Thomas et al. [14] present a visual framework for the timetabling problem, which focuses on displaying a graph based overview for different options. These tasks are centered around the concept of solvers and automation to reduce human effort. A more general effort to understanding timetabling and room planning is made by Beyrouthy et al. [1]. By analyzing room utilization they show the interaction between timetabling and space planning and provide a foundation to utilize available teaching space more efficiently.

Tauböck et al. [13] advance the capabilities of their toolkit for room planning at universities even further by introducing agent-based modeling to simulate walking distances and congestion of hallways between courses for optimal space usage.

To get insight into the paths students take through their courses Raji et al. [10] visualize the course choices and implications on majors and minors that students took at the University of Tennessee. This helps them uncover internal structures of different majors which were not apparent by only reading the curricula. Their tool is only used to visualize historical data of students and has only very limited predictive qualities and no planning is included.

Since none of these systems is capable of addressing the task of operational planning we designed a decision support system using knowledge of common actions of insight gathering shown by Guo et al. [7]. Daniel Power presents a framework for classifying decision support systems [8] as well as guidelines to building a decision support system [9], which we followed and argue about in section V.

III. Methodology

For the design of the OCP tool we conducted a multi-staged design study [12]. A set of four experts, who are in charge of operational semester planning in their respective faculties, took part in this process. All experts were individually consulted after each step of the design and implementation process and their feedback was included in the further process.

Initially we started with a set of low fidelity paper prototypes (see figure 1) — based on the Five Design-Sheet method by Roberts [11] — which we presented to each expert individually. The initial design was based on the six semester schema a typical bachelor's program is based upon and visually similar to the study guide students are given. After the first round we gained insight in the mental model of the experts and their way of tackling the task. We found out that they were working with a faculty-wide semester overview, which aggregated all courses of the six semesters based on their association to the summer or winter semester.

The second iteration (figure 2) reduced the number of shown semesters from six to four and added filters, details

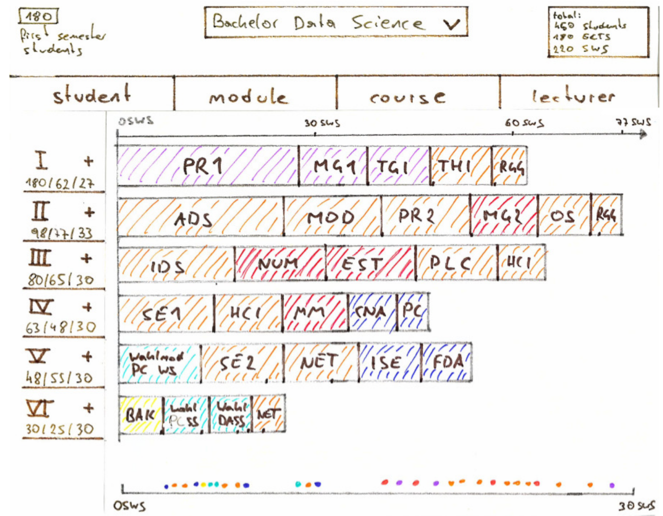

Fig. 1. The paper prototype from the first iteration showed the full six semesters for a bachelors program individually and did not have a dedicated space for showing details about the courses. On the top is a selection of different programs followed by tabs which were later replaced by an axis selection drop-down.

and a separate planning semester. The semesters shown are the past three semesters as well as the next semester which needs to be planned. The planning view showed details about historical values of the courses and added the capability to add courses into the planning semester, adapt or remove them. It also introduced highlighting for courses that were added to the planning semester and need attention as well as courses that were already completely planned. After finishing that prototype we handed it to the experts to evaluate our tool. Since the evaluation period overlapped with the planning period in that semester, the experts could use the tool in their actual planning situation. In the second round of feedback we discovered that teaching contents were missing from the tool — a crucial piece of information for operational planning. A teaching content can consist of one to many different courses that focus on the same topic. The curriculum is divided into teaching contents and the law demands that each one stated in the curriculum has to be covered. In order to finish their studies, students are required to pass a specified amount of credit points (ECTS) for each teaching content, but are able to choose which courses they want to take in order to fulfill this.

The granularity of the visualization was therefore changed from courses to teaching contents and the interface was overhauled. Unnecessary colors were removed to reduce visual clutter and the layout of the tool was enhanced. Information about the teaching contents as well as courses is now shown next to the historical details. The details were also reduced by one data-feature that showed how many students attended the first lecture, which the experts reported they were not using. This design can be seen in figure 4. We also changed some of the detail statistics and aggregation before the tool was evaluated again at the end of the semester in another actual

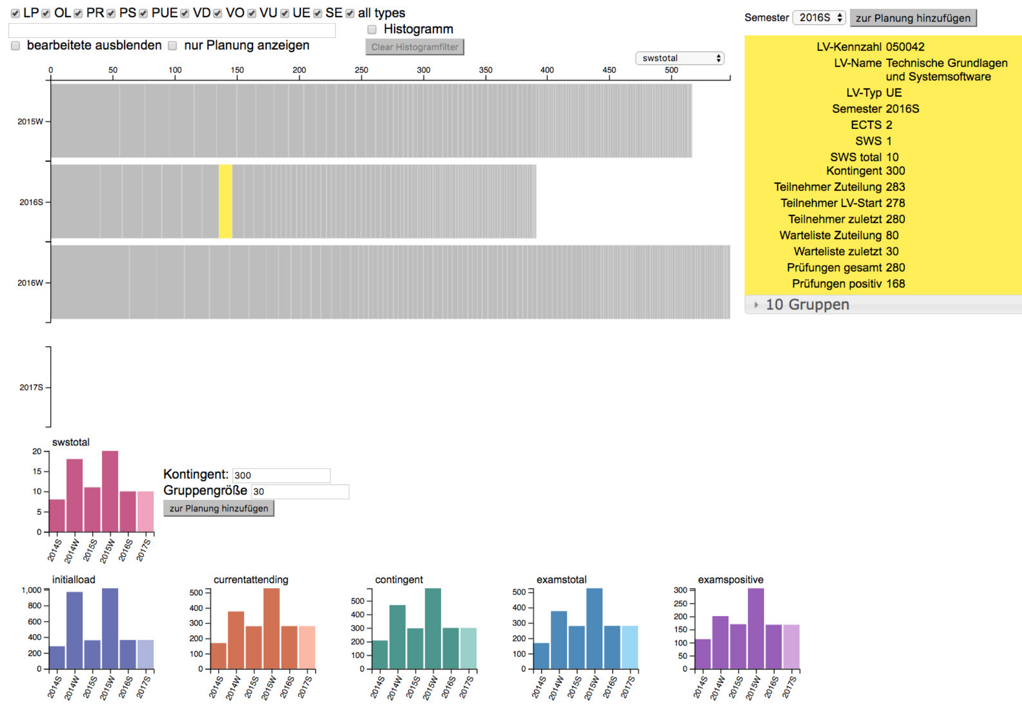

Fig. 2. The first high-fidelity prototype written in D3.js [4]. Its design was still rough and the bar segments showed individual courses instead of teaching contents. It already features the three semester view, as well as the separate planning semester. A filter selection is present at the top of the tool and course details are shown in the yellow box on the right. The planning related historical data is shown as colored bar charts at the bottom of the screen. The number of students at the beginning of the semester is still shown here, which has later been discarded because it was irrelevant for planning.

planning scenario.

The complete design study from the first paper prototypes to the evaluation of the final version of the tool took more than a year, during which the experts had two phases where they could use the tool in their actual semester planning periods. A summative evaluation concluded the design study, its results are discussed in section VII.

IV. Problem Description

A. Requirement Analysis

The general domain of student-related data offers a wide range of subfields and their related tasks and problems. After a general inquiry of open problems we settled for the subfield of operational planning.

Operational planning at universities is generally looking to make sure that all courses that need to be offered will take place, and that all students can attend the courses they need in their studies. At our university operational planning is done by each faculty individually and — depending on the faculty structure — done by either the studies service center or the director of studies. To analyze the process of operational planning we interviewed two directors of studies, several employees as well as two heads of studies service centers. The operational planning routine can be broken down into the following series of questions:

1) Which courses need to be offered in the coming semester?
2) a) How many students are going to attend those courses?
   b) How many teachers are needed to cater for all students?
3) Are the needed teaching hours within the limit of the available budget?
4) Do we have rooms for all classes?
5) Can we provide classes in a fashion that there is little to no overlap?

In the most cases the first question is simple to answer — since the curriculum does not change very frequently courses offered in the previous year will be offered again. Changes that happen to a curriculum need to be addressed manually, because new courses need to be added. Questions two, three and four are directly related to each other. Courses usually have a maximum capacity of students they can support, if more students want to attend a course another instance of that course needs to be offered (called a *group*). Every group needs a teacher, and teachers need to be paid by the amount of hours they teach. An increase in the number of students above the maximum capacity of a certain course also increases the number of teachers needed as well as the teaching hours that need to be paid.

After all courses and groups are planned it is important to find rooms to accommodate them. Some courses have restrictions that they need a certain infrastructure (e.g. chemistry labs) or require large lecture halls to fit all students. The second constraint is that courses which students usually take in the same semester do not overlap to make it possible to attend. Finding rooms and scheduling is therefore also closely related. Classroom restrictions may also effect planning, but the experts told us that this is only rarely the case. We therefore split this routine into two parts, the first one will be referred to as *(operational) planning* and consists of questions 1) to 3), the second part will be referred to as *scheduling* and includes questions 4) and 5).

At our university room data are handled by every faculty internally which makes them hard to gather and the structures of the room data are very different — certain faculties only have seminar rooms and lecture halls, where others have many different types of rooms in their repertoire. Each faculty developed their own strategy to find rooms, and this is already well covered in other works [1], [6], [13]. The problem of timetabling is mainly related to automating the timetabling process by algorithmically optimizing available time slots based on a set of constraints. A plethora of literature tackling that problem algorithmically and visually [2], [14], [15], and commercial solutions are already available as well.

The experts told us they needed assistance in planning teaching resources, since all information they had to rely on were static spreadsheets automatically created by the university reporting system. We therefore chose to focus on *operational planning* and to not address *scheduling* at all. *Operational planning* is not a trivial task, because it has to be finished before the registrations for the next semester starts. This directly implies that the number of students that will register is not known at the time.

A mistake in *operational planning* can have two outcomes:

too few groups: not all students that want to register for a course can do so, therefore hindering the timely progress of the students

too many groups: groups with too few students must be closed due to efficiency reasons, students in those groups need to reschedule and teachers of those groups are losing teaching hours

To tackle question 1 from above the experts need to copy all courses from a previous semester into the currently planned semester. It might also be necessary to copy courses from other semesters or add completely new courses to the currently planned semester. We therefore call this **Task 1: Copying**.

Question 2a is an estimation of how many students will be in each particular course. Experts can easily get insight in the growth or decline in student numbers by looking at past instances of that course and recognizing the trend in historical student numbers. Because of the group size of each course question 2b is inextricably linked to the number of students in that course, thus we combine changing the number of students and changing the size of the groups into **Task 2: Adapting**.

Question 3 can only be addressed after all courses are planned with the right amount of students. It is looking into whether the planned courses stay within the budget by summing up all planned costs and comparing them to the actual budget. If the plan exceeds the budget it is necessary to find courses that can be adapted to lower their teaching hours by either limiting the number of students or increasing the size of the groups. We call this **Task 3: Budgeting**.

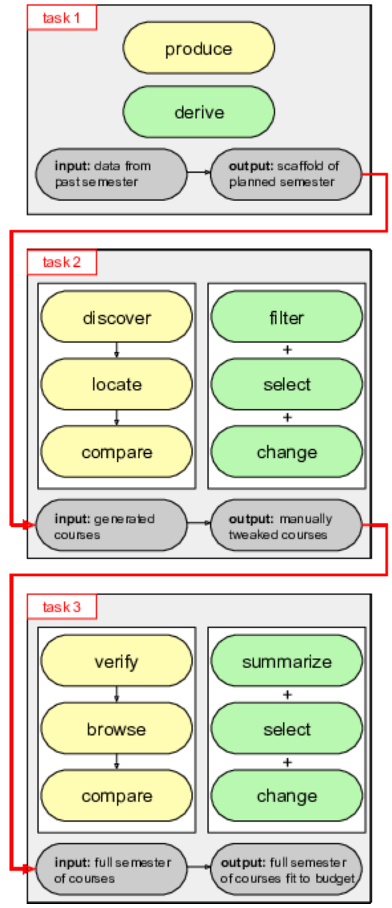

Fig. 3. Task typology of the operational planning tasks based on Brehmer and Munzner [5]. Task one creates a new semester based on the data from the previous semester. Task 2 adapts the automatically created courses that were copied. The third task then summarizes all planned courses to see if the costs of teaching are within the limits of the budget. We used the same colorscheme as in the original paper [5], where yellow encodes *Why?*, green encodes *How?* and gray the input and output (*What?*).

B. Task Abstraction

Figure 3 shows the workflow of operational planning in abstract tasks that have been modeled after the task typology by Brehmer and Munzner [5]. Each task is split into the three questions: *Why?*, *How?* and *What?*. We classified the three tasks identified above as follows:

Task 1: Copying *Why?* We want to *produce* new courses that we can use for planning. *How?* By selecting all courses from a previous semester and copying them as a starting point for planning. The courses are *derived* from the values of their predecessors. *What?* The input data are courses that already happened and have concrete student numbers and group sizes. The output is a copy of all selected courses as a scaffold for planning the next semester.

Task 2: Adapting *Why?* Student numbers are constantly changing and courses in the next semester need to keep up with this change. We therefore need to *locate* courses that need to be adapted and *compare* their capacity to the trend from the historical data. *How?* Individual courses need to be *selected* to be *changed* by the experts. It is also helpful to *filter* courses that have already been changed to avoid redundancy. *What?* The copied courses from task 1 are used as an input, and the manually modified courses are the output from this task.

Task 3: Budgeting *Why?* It is necessary to *verify* that the planned courses fit into the given budget. *Browsing* and *comparing* is used to identify the courses that can be adapted to not exceed planned costs. *How?* All costs of courses in the planning semester are *summarized* to make them comparable to the actual budget. *Selecting* and *changing* courses is used in the same fashion as in task 2 to change them according to the budgetary restrictions. *What?* This task starts with all planned courses as the input and outputs a new iteration of planned courses that now also complies with the available monetary resources.

C. Data Description

The dataset we are working with is event-based, consisting of student-triggered events. Every exam a student took is an event that is linked to the course in which the exam was taken, a grade, the student ID and the time the event took place. We aggregated this data by course and semester. Every course is assigned to a specific teaching content, which is specified in the curriculum. Courses are further split into groups and each course needs to have at least one group. Each group is assigned a teacher and a maximum number of participants.

Before we started our work these data were available to experts as a summary spreadsheet for each semester that aggregated the individual registrations and drop-outs per course and group.

V. Visualization Design

The goal of our tool is to help the experts predict the number of students for each course and to make sure all courses are available in the next semester. Our users need to know how many courses have to change and how big those changes need to be in order to determinate the best trade-off between personnel costs and student acceptance rate. In general the tool focuses on capacity estimation based on historical changes.

The interface of the tool is split into four regions as shown in figure 4. Part (a) is the filter bar, filtering courses is possible based on their types or their current planning status (figure 6), as well as by their name or the name of the teacher using the search bar (figure 5). Filtering the data immediately updates the main view of the tool, after each keystroke within the search field.

The main view (figure 4 (b)) consists of three gray stacked bar charts that represent the past three semesters. Each segment of the stacked bars is one teaching content and the width of the segment encodes the chosen value above the x-axis. Possible values are teaching hours, ECTS, planned contingent, participants, number of exams and number of positive exams – the default view shows teaching hours, since it is the most relevant for planning (note: teaching hours are called *SWS* in the interface, because this is the internal name which is used at our university). The currently selected teaching content is darker than the rest, gray topics happened in past semesters and light gray topics indicate that they have been added to the planning semester. Mousing over a bar segment blurs all non-related segments to show only the occurrences of the chosen segment in all semesters. Also a tooltip is shown that contains the name, the teachers and the width of the segment in the currently selected unit. The fourth stacked bar chart (figure 4 (c)) represents the planning semester where all courses are shown that are currently being planned. Blue and orange color only appear in the planning semester and indicate planning completed and attention required respectively. Courses can be added to the planning semester either individually by selecting them and pressing the *add* button in the bottom right section, or a whole semester can be added by selecting it from the drop-down menu above the planning semester and pressing *add semester to planning*.

Section (d) is the detail view, which is only visible if a segment from above is selected. It is split into three different columns that show different aspects of the currently selected topic:

- left column: a summary of the currently selected teaching content in the respective semester
- middle column: all courses grouped together in the selected teaching content, a course summary can be shown by clicking on the course title (summary is currently shown in figure 4 (d))
- right column: the top left part shows the planning interface where the maximum number of students per course (contingent) and the group size can be set. It also features buttons to add/remove a course from planning and to mark that planning the course is complete. Five bar charts are showing the historical data of the course (light gray) and the prediction based on the currently set values for the planning semester (dark gray).

The historical values we chose to show are *SWS total*, which shows the total teaching hours the course used; participants shows how many students attended the course; contingent is the maximum number of students that has been planned; exams total and exams positive are self explanatory and are used to estimate how many students failed a class and may take it again. Mousing over a bar shows its values and draws a line at the height of the bar into all small multiples. Clicking the bar freezes the line and allows easy comparison of two values (see figure 8).

Since the tool can be classified as a decision support system (DSS) we integrated the following Data-Driven DSS [8] features defined by Daniel Power [9]:

Filtering and Retrieval The filter bar at the top of the application allows the user to filter based on attributes, on planning status or search for items with the search box.

Data Summarization Preprocessing and summarizing the data and aggregating multiple years into one view is automatically done by the tool before visualizing the data.

Fig. 4. Final design of the OCP tool. (a) The filter bar is used to search and filter the main view. The search bar has an auto-completion feature and the filters are based on courses and planning status. (b) Each bar segment in the main view represents one topic of the shown curriculum. The last stacked bar (c) below the three gray (past) semesters is the planning section. Each topic selected for planning is shown using the planned student numbers as basis for the bar widths. Orange signals attention needed, blue states that planning is already completed. (d) Details of the selected topic are shown, the left column is the summary of the topic, all corresponding courses are listed in the middle column and the right column shows planning-relevant historical data of the individual course as multiple small histograms. The historical values shown are teaching credits, number of students, maximum attendees, total number of grades, and number of passing grades. (Note: Names have been redacted.)

Data import/export The tool works with the standard university database format to make importing new data as simple as possible. Planning data can be exported in .csv format.

Reusability By simply importing the dataset from the new semester and changing the planning semester the tool can be used for all semesters.

Not included are 1) alerts and triggers 2) metadata creation and retrieval and 3) reporting capabilities (creation and viewing) [9]. These features did not seem necessary during the iterations, because they did not fit the requirements of this project. The OCP tool also follows the guidelines of targeting capabilities to user needs, ease of use, ease of installation (web-based, automated import) and performance. Cost and support by vendor are left out due to the academic nature of the project.

A. *Tackling Operational Planning*

The OCP tool is capable of dealing with all tasks we identified in section IV-B and therefore able to fully support domain experts in operational planning.

Task 1: Copying is achieved in the OCP tool by selecting the previous semester in the drop-down menu of the planning section (figure 4 (c)) or selecting individual courses and adding them to the planning semester by clicking the *add* button in the detail section.

All courses start colored orange in the planning semester to indicate that a manual confirmation is still needed. For **Task 2: Adapting** each orange colored section in the planning semester needs to be selected to show the details of all courses. By comparing the historical values and seeing trends in student numbers the expert can estimate the contingent needed (figure 7). Sometimes information from other courses

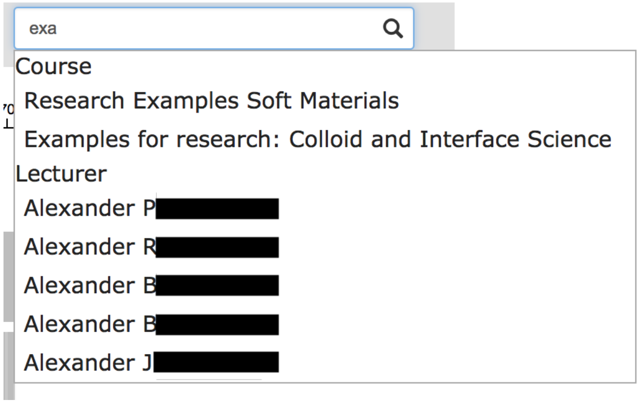

Fig. 5. The search bar automatically suggests entries based on the current input text and updates the main view accordingly. (Note: Names have been redacted.)

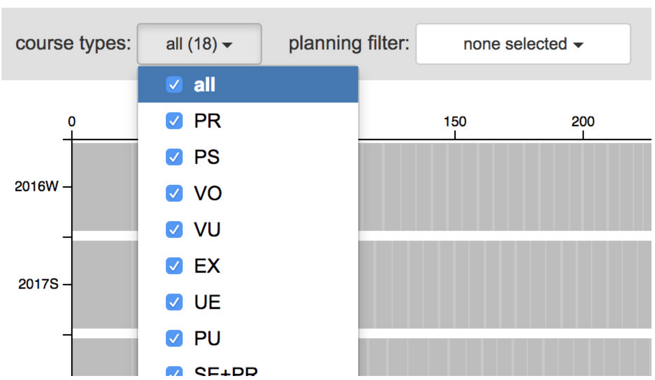

Fig. 6. The main view can be filtered by selecting either course types, planning restrictions (show only courses in planning, hide already planned) or a combination of both.

is needed, which can be easily found by using the search bar (figure 5). After the expert is content with the number of planned students the course gets marked done and turns blue. Task 2 is fully finished once all courses in the planning semester are blue.

Task 3: Budgeting starts with the overview of the planning semester (figure 4 (c)), where teaching hours are selected on the x-axis. The width of the bar shows the total amount of teaching hours needed. If that amount exceeds the budget the expert needs to browse through all planned courses to select and change some of them to fit into the budget limitations.

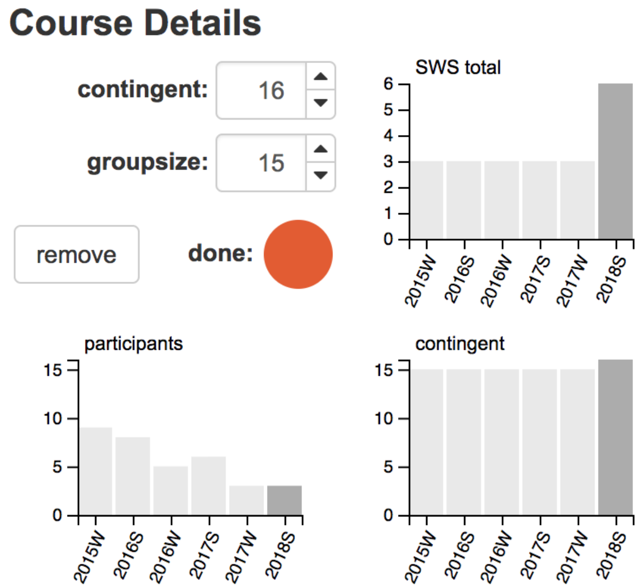

Fig. 7. Changing the values of the contingent and group size updates the dark gray bar in the histograms that represents the currently planned semester. A group size of zero indicates no limitations in terms of student numbers per class (useful for some lectures). In this example the contingent was increased by one compared to the previous semesters. Since the number of students exceeds the capacity of one single group, a second group was automatically planned — which doubles the amount of teaching hours (SWS) needed.

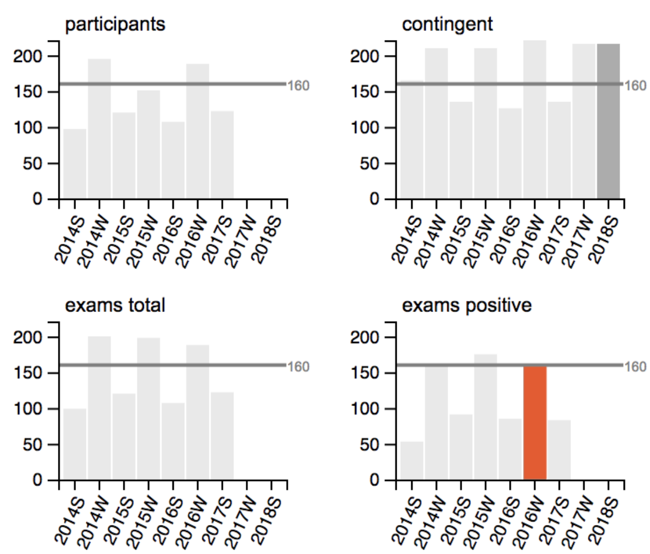

Fig. 8. Mouseover interaction of the bar charts: hovering over a bar highlights it and draws a line at the height of the bar into all small multiples. Clicking the bar freezes the line to make it easy to compare two values.

VI. Use Case

The testing phases of the tool during the design study were chosen in such a fashion that the experts could use the tool in their actual planning of teaching resources at the end of the semester. We were not present during these times to monitor their usage of the tool, we only heard their experience reports in the interviews afterwards.

To see how the experts used the OCP tool we prepared a specific use case that the experts had to solve during the last interview session of our design study. They then provided feedback on this experience as well as the general usage. During this interview they had to plan a new semester with the same values as the previous one, but with a special condition. We selected a set of three courses for each expert, which already required a considerable amount of teaching hours. We then told them due to an external event the amount of students that are going to show up in these courses will double, and they still need to stay within the budget from the previous

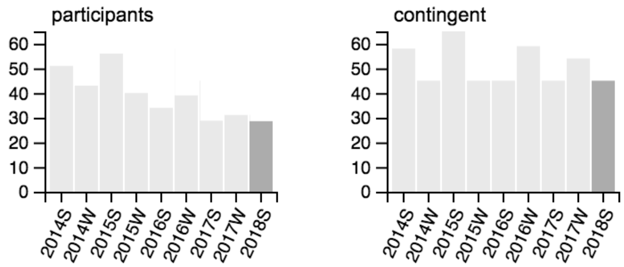

Fig. 9. A downwards trend in participants is easily detectable in the historical course data. Thus the expert reduces the contingent and therefore also the number of teaching hours needed.

semester.

The following example describes the steps taken to tackle that problem as done by one of the experts. The expert starts by copying the previous semester into the planning semester by selecting it from the drop-down menu (figure 4 (c)). He then locates the three courses individually by using the search box (figure 5) and doubles their contingent in their detail section (figure 7). In one of the courses we chose he immediately notices that he can increase the group size to reduce the number of groups needed. He explained that this particular course does not need a specialized room and therefore the number of students can be increased without worrying about laboratory capacities. Looking at the summary of all planned courses he noticed that he was still ten teaching hours above the budget. Browsing through the leftmost part of the planning semester (figure 4 (c)) and selecting various courses one by one he identifies another course where he can increase the group size as well as one course that had a downward trend in attendants where he can reduce the contingent (see figure 9). Because after these changes he is still two teaching hours above the budget, he then decides to double the group size for a small course that uses four teaching hours and has two groups. He told us that the course could handle the increase, but he would still need to talk to the lecturer before actually applying that specific change.

VII. Evaluation

A continuous feedback and evaluation loop was present throughout the whole design study. We met with each expert individually to discuss the current state of the tool. After each programming iteration the experts had at least a week to test the prototype which already integrated their own data. In about one hour long meetings following each testing period we observed how the experts used the tool, interviewed them on their experience using it and discussed improvements and changes. This feedback was iteratively included into the following versions. The prototype shown in figure 2 as well as the final OCP tool (figure 4) were used during actual planning periods by the experts. The feedback for using that prototype in planning showed that the interface was too cluttered and the colors were distracting. This was the reason for the redesign of the interface for the final tool, alongside the restructuring of the depicted data.

The conducted interviews showed a significant perceived gain in both accuracy of predictions and reduction of time invested by all experts. The interface was described as easy to use, but sometimes too cluttered, especially when dealing with the smaller sections of the stacked bar charts. A suggested improvement would be increasing the size of the bar under the cursor. All four experts reported using the tool to help them in their actual planning period. They used it to investigate the historical data of courses they were unsure about while planning. It was also used for retrieving information about courses and teachers, because it was easy and fast to find data in the OCP tool. One expert also used the tool as basis for discussions to show lecturers that they did not need to teach more than other lecturers by simply filtering for courses taught by each individually and comparing them. No expert used the actual planning component to plan the next semester in the tool, because it was not integrated with the proprietary system for course management used by the university. When asked, the experts stated that they were eager to adopt the integrated planning feature into their current workflow, as soon as it would be possible to directly use the planned results in the course management system.

VIII. Lessons Learned and Discussion

We created a tool as a replacement of an established workflow that was using a database dump aggregated into multiple spreadsheets. This task was mutually identified by all experts to be tedious and time-consuming. Even though the task was simplified by the tool and good overall results have been achieved, the adoption of our tool has not taken place so far. A general response for the lack of adoption was the unwillingness to switch from existing patterns to a new system that only improves parts of the experience. Full adoption of the system needs integration into the complex (and proprietary) university tool eco-system — a task whose complexity and time-scale is beyond the scope of this research project.

After showing the OCP tool as well as our results to the university's IT department they told us that an operational planning tool was one of the most requested tools throughout all the university. We are currently working together with the IT department to integrate the OCP tool into the tool eco-system of the university.

Evaluation without adoption of the tool has proven to be very hard since a lot of feedback went towards integrating the OCP tool into the existing system. This made the experts lose focus of reporting about the functionality of the tool. Future projects should therefore evaluate the existing infrastructure first and check for possibility and feasibility of integrating a tool from the start.

In the first iteration of our design study we found out that our mental model (from the perspective of the visualization experts) and the mental model of the domain experts were quite different, even though our expertise comes from creating a new curriculum ourselves. Because building a visual decision support system heavily depends on the mental model of the target users, we urge visualization experts to spend effort on capturing and depicting the mental model of their target users in addition to identifying data and tasks. Low-fidelity prototyping is a key mechanism to tackle this, because changes can be made inexpensively and flexibly and should already be done during the first interviews to create a common ground between the visualization experts and the domain users.

The stacked bar charts have proven to effectively handle the selecting of single courses and adapting them, and easily compare courses over multiple semesters, as well as comparing whole semesters. Because of the simple perception of bar charts and their meaning the experts had no initial problems to understand the visualization.

IX. Future Work and Conclusion

We presented the OCP tool, a domain specific data-driven decision support system, which reduced the time taken by experts during operational planning while increasing the quality of their decisions, because obscured information was made easily accessible. The problem of operational planning at universities was depicted and we identified the abstract tasks of that domain. Furthermore we showed that a tool fully capable of simplifying a task — even depicting the mental model of the domain experts — can be very hard to be adopted in an existing workflow because of a seemingly small integration hurdle.

The next step for our tool after it finds its way into the production system will be the analysis of data gathered from complete semester planning using the tool. It could be viable for a predictive system to use this data and create an approximation of the next semester automatically. This would enhance the creation of the semester scaffold when copying the topics from the previous semesters. Instead of using the same value for the contingent and group size, the predictive system could intelligently adapt the values. The expert then only needs to make a few manual adjustments in the end to correct the automated suggestions.

References

[1] C. Beyrouthy, E. K. Burke, D. Landa-Silva, B. McCollum, P. McMullan, and A. J. Parkes. Towards improving the utilization of university teaching space. *Journal of the Operational Research Society*, 60(1):130–143, 2009.

[2] Alex Bonutti, Fabio De Cesco, Luca Di Gaspero, and Andrea Schaerf. Benchmarking curriculum-based course timetabling: Formulations, data formats, instances, validation, visualization, and results. *Annals of Operations Research*, 194(1):59–70, 2012.

[3] Jess Boronico. Quantitative modeling and technology driven departmental course scheduling. *Omega*, 28(3):327–346, 2000.

[4] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. D3 data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, 2011.

[5] Matthew Brehmer and Tamara Munzner. A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2376–2385, 2013.

[6] Brecht Cardoen, Erik Demeulemeester, and Jeroen Beliën. Operating room planning and scheduling: A literature review. *European Journal of Operational Research*, 201(3):921 – 932, 2010.

[7] Hua Guo, Steven Gomez, Caroline Ziemkiewicz, and David Laidlaw. A Case Study Using Visualization Interaction Logs and Insight. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):1–1, 2015.

[8] Daniel Power. Supporting Decision-Makers: An Expanded Framework. In *Informing Science eBook*, number June, pages 1–6, 2001.

[9] Daniel J Power. Understanding data-driven decision support systems. *Information Systems Management*, 25(2):149–154, 2008.

[10] Mohammad Raji, John Duggan, Blaise DeCotes, Jian Huang, and Bradley T. Vander Zanden. Visual Progression Analysis of Student Records Data. *CoRR*, abs/1710.0, oct 2017.

[11] Jonathan C. Roberts, Chris Headleand, and Panagiotis D. Ritsos. Sketching Designs Using the Five Design-Sheet Methodology. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):419–428, 2016.

[12] Michael Sedlmair, Miriah Meyer, and Tamara Munzner. Design Study Methodology: Reflections from the Trenches and the Stacks. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2431–2440, dec 2012.

[13] Shabnam Tauböck, Felix Breitenecker, Dietmar Wiegand, and Nikolas Popper. The <morespace> Project: Modelling and Simulation of Room Management and Schedule Planning at University by Combining DEVS and Agent-based Approaches Introduction. *SNE Simulation Notes Europe*, 21(1):11–20, 2011.

[14] J. Joshua Thomas, Ahamad Tajudin Khader, and Bahari Belaton. A visual analytics framework for the examination timetabling problem. *Proceedings - Computer Graphics, Imaging and Visualisation, Modern Techniques and Applications, CGIV*, pages 305–310, 2008.

[15] Melanie Tory, Sheryl Staub-French, Dandan Huang, Yu Ling Chang, Colin Swindells, and Rachel Pottinger. Comparative visualization of construction schedules. *Automation in Construction*, 29:68–82, 2013.