Exercise problems that facilitate deep approaches to learning

Martin Hell & Jonas Johansson

August 19, 2010

1 Introduction

University courses use several teaching and learning activities. The most common is probably the lecture, but it is often complemented, at least in the faculty of engineering, by problem solving classes, laboratory practicals, projects and sometimes also other activities, such as home assignments or essays. This project will focus on exercises and problem solving classes and how they can be improved to support better learning.

So what does a good problem or exercise look like? This seems at first glance to be a very difficult question to answer since learning styles vary greatly between students and problems are designed for a particular course. Two courses that have roughly the same content, but are given in different universities, might have different examination forms. As a result, problems that are suitable in one course might not be suitable in the other.

The goal of this project is not to develop a universal algorithm for problem construction. That would clearly not be possible. Instead, we will look at well known concepts and beliefs that are results from research in teaching and learning. We look at the following concepts

- Motivation - High motivation influences and stimulates learning.
- Transfer of Learning - Previous knowledge is used when acquiring new knowledge.
- Constructive Alignment - There should be a clear link between course goals, teaching activities and examination.

With this as background we aim to formulate a few ideas that can be kept in mind when designing new problems. The goal is to improve the problems used in two courses at LTH.

2 Courses

In this section we give an overview of the two courses that we plan to work on. We also identify and motivate the need for better exercises and problems in the courses.

2.1 Web Security

Web security is a course given at the department of Electrical and Information Technology. It is an obligatory course in the third year of the Information and Communication Engineering Technologies (C) program, but also optional for students in the forth year of the Computer Science and Engineering (D) and Electrical Engineering (E) programs. Approximately 60 students participate in the course, which is rather small as it gives the students 4 hp.

The web security course was introduced in 2008 and has thus been available for two years. There are seven lectures and one project and the examination is in the form of a written exam based on problem solving and understanding of the course content. So far there has not been any scheduled problem solving sessions. In 2008, the students were not even given any problems or exercises to solve, apart from one practice exam, even though the exam was based on problem solving. Despite this, the course received a good evaluation on the CEQ. The only criticism from students was that it was difficult to know what was expected from them. This criticism is in line with the lack of constructive alignment which is evident when comparing the exam and the course content. As a result, in 2009, the students were given a set of exercises and problems (without solutions) to solve during the course. However, there were no scheduled problem solving classes. Instead, students had the possibility to discuss the problems with the teachers during office hours. It is not clear to which extent the student actually tried to solve the problems. The goal for the course instance of 2010 is to improve these exercises. Moreover, in 2010 there are 3 scheduled problem solving classes. For more information, the course description can be found in [1].

2.2 Processing and Device Technology

Processing and Device Technology is a course given at the division of Solid State Physics. It is compulsory in the third year of the Engineering Nanoscience (N) program, but optional for students in the fourth year of the E and Engineering Physics (F) programs. The course is also taken by several international master and exchange students. Often, a couple of students from the natural sciences faculty also take the course. Thus, the course is given in English for a very mixed group of students. Each year around 60 students take the course, which gives 7.5 hp (ECTS credits).

Just as the Web Security course, the Processing and Device Technology is a fairly new course. It was first given 2004 as an optional course but as a compulsory course since 2005. There are 12 lectures, 5 problem solving classes, and an extended laboratory practical consisting of 5 occasions, where each occasion is 4 hours. The course has a written examination and to be able to pass the course the laboratory report must be approved. The problem solving classes are organized as follows: The students are encouraged to try to solve the problems in advance. During the class the students discuss and solve problems (usually) in groups. The teacher helps the students and gives feedback. After some time, the teacher solves most of the problem on the whiteboard.

Processing and Device Technology generally gets positive responses on the CEQ. The most common criticism though, is that there are too long waiting times during the practicals. We have improved this by using more laboratory supervisors, allowing the students to work in smaller groups.

The goal for the course in the fall of 2010 is to further improve the laboratory practicals but also to improve the set of problems that the students practice on during the problem solving classes. For more information, the course description can be found in [2].

3 Theory & Previous work

In this section we look at well known and established concepts from the theory of teaching and learning. These will later be used as a basis for understanding how to improve problems and exercises.

Surface and Deep Approaches to Learning

A student's approach to learning a particular subject can be divided into two categories, deep and surface approach. Basically, students that focus on just memorizing as much as possible are said to having a surface approach to learning. Putting the different parts together is not seen as important. Learning how to solve a problem does not imply that other, similar, problems can be solved. At the other end of the scale are the students that are trying to *understand* the concepts and the message in the material. They relate the knowledge to previously gained knowledge and put everything together in a bigger picture. Learning how to solve a problem will allow these students also to solve other problems since they can use the knowledge in other situations. This, of course, does not mean that these students never have to memorize information. Memorization is still an important part of the deep approach to learning, but it is important to emphasize that it is just a part of it. Students that take a deep approach also remember the information for a longer time.

It is important to remember that deep and surface approaches to learning is not a constant defining characteristic of a student. One student can take a surface approach in one situation or course and a deep approach in another. Also, this may not be an active choice, but can be a subconscious choice. In this case, teaching activities can play an important part in this choice. Thus, teachers should actively develop problems and exercises that motivate the students to take on a deep approach. However, even if the importance of good problems can easily be motivated, since we want the students to take a deep approach, this still does not tell us *how* to construct these problems.

The terminology is related to atomistic and holistic views of knowledge. An atomistic view, being more traditional, is a view that sees the knowledge as small units of facts that together build up a larger understanding. As learning proceeds, new pieces of knowledge are added. In this view, the focus is on the teacher and how the teacher transfers knowledge to the students. In the holistic view, the parts are instead seen as different aspects of the whole. The gathered knowledge is more than just the sum of the parts [3].

Motivational Aspects

There are several factors that influence and stimulate learning. One important factor is *motivation*. Motivation is usually divided into *external* and *internal* motivation. Externally motivated students are motivated by outer factors, e.g., getting a degree or getting a good job. Internally motivated students are driven by interest in the subject, curiosity, or that they believe that the knowledge is actually useful for them. Students that are internally motivated learn better than students that have only external motivation [4]. It is possible to influence the internal motivation. If teaching activities can be formed such that they show that what the students learn is actually useful for them, then by definition the internal motivation is increased, and learning will benefit. Relating this to problem solving, the problems should be stated such that they support this. This could e.g., mean that problems should be taken from real situations that students can identify with. We base this on a belief that knowledge is easier to find useful if it relates to real situations. Thus, we conclude that a real life scenario is much better than a made up ad hoc scenario. The problem has to emphasize that what is learned is really relevant. Other ways to stimulate motivation is to give students challenges that are neither too hard nor too easy [5]. This is of course difficult when there are many students with different background and initial knowledge. More motivational factors can be found in [5], but most are difficult to apply in our context.

The Transfer of Learning

One important goal of teaching is that students should be able to use knowledge from one course in a later course. This often means that they must be able to transfer what they have learned in one context into another context. The two contexts can be two different courses but they can also be two contexts within the same course. The problem of transferring knowledge from one context into another has been a subject of extensive research and can also be looked at from different points of view. We restrict ourselves to four key characteristics of learning and transfers given in [6].

- Initial learning is important. The transfer of knowledge to a new context will not work unless the student has good knowledge in the first context. Motivation, as discussed above, can help improve learning in the first context.

- Knowledge should not be too contextualized, but should be learned in a more abstract setting in order to promote transfer. Problems and exercises supporting this can involve finding connections between several contexts and to work out general solutions that apply to all these contexts.

- Transfer is an active and dynamic process. A transfer from a first context to a second may not require the same strategies and activities as a transfer from a second to a third context. Teacher interaction is often very important to help transfer, but the exact type of interaction is very individual.

- All learning is transfer. When we learn something new this builds upon something that we know since before. Thus transfer is present in all types of learning. In teaching this can be taken advantage of by relating new knowledge to something students already know. This can e.g., be done by relating problems and exercises to other well-known facts or events.

Having these characteristics in mind, it is possible to construct problems and exercises that can help transfer learning between contexts. At first glance, the second bullet, "knowledge should not be too contextualized", seems to contradict the conclusion that problems should be based on real life scenarios from the section on motivational aspects. However, we believe that using real life scenarios instead of made up ad hoc scenarios in problems based on actual situations does not contradict the fact that abstract problems should also be used, since these are a different type of problems. Moreover, real life problems can include and encourage abstract reasoning.

Constructive Alignment

There should be a clear link between course goals, teaching activities and examination. This is the concept of constructive alignment [7]. Consider a course which consists of lectures, problem solving classes, and a few labs or projects. Lectures will provide necessary information and hopefully help the students structure the information [8]. The problem solving will allow the students to practice just that, solving problems related to the course content. The laboratory practicals or projects will give the students the opportunity to apply the knowledge in practice. Approved lab practicals/projects are typically required to pass the course, but there is often also a written exam based on problem solving. This is the case on the two courses that we focus on in this project, but also in many other courses in LTH. Often, problems are taken from the course book, where they are located at the end of each chapter. These problems test the material covered in the corresponding chapter. However, the exam may include problems that connect the material in different chapters and require the students to combine information from several parts of the course. This would not be in agreement with the constructive alignment as students do not get the opportunity to practice on these types of problems, which might show up on the exam. Thus, it seems favorable to include in the list of problems, also problems that combine the material from several parts of the course if these types of problems will be used on the exam. Often, this is solved by giving the students a set of all exams. Then they can practice also on these problems. Still, in the view of constructive alignment we consider this solution to be suboptimal. These problems will be considered by students as typical exam problems, while problems solved during the course will be considered to be another type of problems. Hence, the problem with constructive alignment remains in this case.

4 Applying Theory to Problems and Exercises

Considering the previous section, we summarize the aspects of problems and exercises that should be considered. It must be noted that this is not a definitive and exhaustive list, but it can be used as starting point when constructing new problems aimed at helping students take on a deep approach to learning.

- Problems based on real life situations or facts can help the students to increase the motivation.
- Be careful not to make too easy or too difficult questions.
- Problems based on finding connections between different contexts can help transferring the learning.
- Relating problems to something that the students already know, e.g., previous courses or earlier parts of the same course, can also help transfer.
- Give students problems that are comparable to the problems given on the exam. It should not be completely different types of problems on exam and during the course.

The following two examples are related to the Web Security course. Their purpose is to use examples from real life that the students are already familiar with. Moreover, the examples incorporate several parts of the course so that students more easily can see the full picture instead of just one small part at a time. Hopefully, this can both motivate and help students to take a deep approach to learning.

Example 1:
In February 2008, a Danish court ordered Tele2, a Danish Internet Service Provider, to block all access to the popular website "The Pirate Bay". The website was blocked by using a DNS redirect. This means that, when the client tries to look up the IP to the website, the DNS server owned by Tele2 responds with another name.

a) Give a few different ways to easily bypass this block.

b) It is easy to conclude that this kind of web site blocking is rather useless. The discussion whether it is right or wrong is outside the scope of this *course. The implementation of the blocking directed users to another website with information that the pirate bay was blocked. Compare this to the DNS cache poisoning attack. What similarities and differences can you see?*

This problem not only relates to a real life example, but something that practically everyone has heard of and knows something about from the beginning. The second part of the question relates the DNS block to something that at first glance is completely different, namely an attack on DNS servers performed by someone that has no (administrative) access to the server. However, there might be more similarities than expected as the result is exactly the same. End users are directed to a "fake" web server, without knowing it (except visually of course).

Example 2:
Home DSL routers are commonly used to provide Internet access for home users. They also allow several users on a home network to access Internet using one connection. The fact that many home routers could be compromised received much attention in late 2008. A typical default security setting is to disable WAN administration, i.e., administration from outside the home network. Only LAN administration is allowed. Since only computers inside the home network can access the configuration menu of the router, another typical default setting is to allow administration without a password. In other words, computers on the home network are considered trusted. It is possible to enable WAN administration and set a password for accessing the router from outside the home network. The setting can be set in one of the configuration menus. The setting is enabled by the HTTP request:

```
POST /setremote HTTP/1.1
Host:  192.168.1.1

RemotePassword=abc&WANAdminAccess=on&timeout=10&Enable=Enable
```

In some cases it turned out that GET requests could be used as well. The attacker's goal is to enable remote administration and choose the password. After that, the attacker can control all network traffic leaving the router by e.g., changing the DNS used.

- *a) Describe a CSRF attack on this setting. Under what circumstances will the attack work?*
- *b) To improve security, the default password can be used for LAN access. Describe a CSRF attack on this setting. Under what circumstances will the attack work?*
- *c) To improve security further, basic or digest access authentication with a well-chosen password can be used to protect the router. Describe a CSRF attack on this setting. Under what circumstances will the attack work?*

d) *An alternative to HTTP authentication is to use web based login. Describe a CSRF attack on this setting. Under what circumstances will the attack work?*

Similar to the first example, this relates to something that most people are familiar with. The fact that simple attacks can take control of your Internet connection even though the router is password protected should be motivating. The problem also incorporates other parts of the course (HTTP authentication) and require the students to look at the problem from several points of view. This will broaden their understanding of the CSRF attack.

The following example is related to the Processing and Device Technology course. The purpose is to connect different areas in the course into one realistic problem, which should favor a holistic view, necessary to fully grasp the subject.

Example 3:
You want to fabricate a simple pn-junction diode, according to Figure 1, or according to your own design. All feature sizes can be in the micron range, so UV lithography will work. Write the sequence of process steps you need to make this device. Give a thorough description of each step. Explain why it is important that the process steps are executed in the order you have chosen.

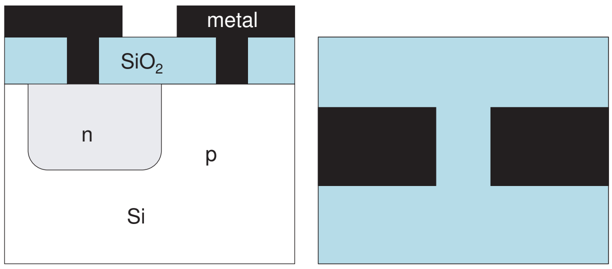

Figure 1: Sketch of a diode. The left panel shows a cross-section view and the right one shows the top view.

This problem is about the fabrication of a real electronic device, which all of the students should be familiar with from previous courses. The core of the problem is process integration, which is about executing the different process steps (oxidation, lithography, etching, diffusion of dopants, etc) in the correct order. Text books on the subject are always organized with one chapter per process. This has the effect that the typical end-of-the-chapter problems only deal with single processes. This is without a doubt important, but clearly not enough for the students to fully understand the subject. To be able to fabricate devices and circuits it is of highest importance to understand how the processes are related and which order the individual processes steps should be carried out.

In the problem above, Example 3, the description of each step is a subtask that resembles typical end-of-the-chapter problems and belongs to the lower end of the SOLO-taxonomy [9]. On the other hand, the construction of the process sequence and the explanatory part of the problem gives the students a chance to synthesize and analyze. These parts of the problem can thus be classified with higher order verbs from the SOLO-taxonomy.

This problem also facilitates constructive alignment as it is more based on understanding than memorization, questions that are common on an exam.

Another interesting problem related to process integration is to show the students a principal sketch of a device and pictures of the lithographic masks that are needed to fabricate the device. The problem for the students is to organize the masks in the correct order. This is a smaller problem, which is well suited as a discussion problem during a lecture.

5 Evaluating the Changes

We intend to make the set of problems in our respective courses (Web Security and Processing and Device Technology) more complete by including a few problems of the type discussed above. Typical end-of-the-chapter problems will still be used so the students also can practice details.

We will use questionnaires to monitor whether the students felt that these new problems have increased their level of understanding and in that case how. The questionnaires will have only a few but open questions enabling the students to freely express their thoughts.

Casual discussions together with the students can also help us to understand their impression of the problems.

6 Conclusion and Discussion

In this report we propose ideas on how to construct problems that should facilitate a deep and holistic approach to learning. While the typical end-of-the-chapter problems only contain one aspect it is a risk that the students adopt an atomistic view of the subject. This is not beneficial, since most realistic situations require that the students can combine several aspects of the course content. Thus, the set of problems should be completed with some, possibly larger, problems that include several parts of the course, i.e., combines the contents of several chapters of the course text book.

Further on, we discuss how to construct good problems that will facilitate deep and holistic learning. Our discussions are based on motivational aspects, the transfer of learning, and constructive alignment.

Finally, we give three concrete examples of new problems from our respective courses and we also propose how to evaluate the new kind of problems.

Since the main goal is to facilitate better learning, it could be interesting to let the student give input to how problems should be constructed, or even let them construct problems by themselves. Some teachers have taken this even further by letting the student construct problems on their own, and then include the best problems on the exam. The respective student will obviously benefit from this since the solutions should be clear to him or her.

We conclude this report by reminding that learning is very complex and the problem of helping the students to take a deep approach to learning requires more than good problems and exercises. Other teaching activities are as, or sometimes even more, important. We must not forget that all students are individuals and that the best way to learn is also unique to each student.

References

[1] URL: http://www.ka.lth.se/kursplaner/arets/eitf05.html

[2] URL: http://www.ka.lth.se/kursplaner/arets/fff110.html

[3] M. Elmgren and A-S. Henriksson. Universitetspedagogik. Norstedts, 2010. ISBN 978-91-1-302297-0.

[4] M. A. Eppler and B. L. Harju. Achievement motivation goals in relation to academic performance in traditional and nontraditional college students. Research in Higher education, 38(5), pp. 557–573, 1997.

[5] J. Turner and S. G. Paris. How literacy tasks influence children's motivation for literacy. The reading teacher, 48(8), pp. 662–673, 1995.

[6] J. D. Bransford, A. L. Brown and R. R. Cocking. How People Learn: Brain, Mind, Experience, and School. Commission on Behavioral and Social Sciences and Education, 1999. Url:*http://www.nap.edu/openbook.php?record_id=6160&page=R1.*

[7] J. Biggs. What thte student does: teaching for enhanced learning. Higher Education Research & Development, 18(1), pp. 57–75, 1999.

[8] D. Bligh. What's the use of lectures?, Intellect Books, 5th revised edition, 1998.

[9] J. B. Biggs, K. F. Collins, Evaluating the quality of learning: The SOLO Taxonomy (Structure of the Observed Learning Outcome), Academic Press, New York 1982.