

How do students understand computer network protocols?

Anders Berglund

Uppsala Computing Education Research Group, UpCERG
Department of Information Technology
Uppsala University
Sweden

Anders.Berglund@it.uu.se

Raymond Lister

Faculty of Information Technology,
University of Technology, Sydney,
Australia

raymond@it.uts.edu.au

Abstract

When computing education research is performed in a rigorous and theoretically sound way, it can inform and improve teaching. Having conducted a phenomenographic study of students' understandings of computer network protocols, we present our findings and explain how they can inform the teaching of these protocols. We go on to explore what our results can tell us about the discipline of computer science as a whole.

Keywords: Computing Education Research, Computer Science Education, Phenomenography, Networking, Network Protocols.

1 Introduction

This paper has two objectives. First, and more important, it is intended to inspire teachers of computer networks. Second, it illustrates how research in computing education, performed rigorously and founded in theoretically sound pedagogical theory, can offer insights that promote our understanding of teaching and learning computing.

The ultimate aim of the research project reported on in this paper is to influence and improve the teaching of computing. Somewhat more narrowly, the focus of this paper is on how students understand a particular computing concept, namely computer network protocols. The rationale for concentrating on the students and their learning, is explained in the following way by Ramsden (1996):

Teaching and learning in higher education are inextricably and elaborately linked. To teach is to make assumptions about what and how the student learns; therefore, to teach well implies learning about students' learning. (p. 7)

The students in a cohort present markedly different ways of understanding computer networks. This paper presents analyses and descriptions of these different ways, and discusses their value in the sense of applicability to the learning aims of the course. This focus puts aside implications for teaching, tips and techniques, and issues

This quality assured paper appeared at the 20th Annual Conference of the National Advisory Committee on Computing Qualifications (NACCQ 2007), Nelson, New Zealand. Samuel Mann and Noel Bridgeman (Eds). Reproduction for academic, not-for profit purposes permitted provided this text is included. www.naccq.ac.nz

concerning implementations of these ideas. These are left for future research.

1.1 Content of this paper

We begin with an overview of how research in students' learning of computer networking can be studied, and introduce phenomenography, a specific approach to the study of learning. We then discuss the project course in computer systems, from which the empirical material for our study was collected. In section 4, the core of the paper, we describe the different ways in which the students understand network protocols, and the manner in which they shift between these ways. In section 5 we present our analysis, relating it to prior findings within this field of research, and discussing our results in the context of the learning aims. We conclude with a brief discussion.

2 Phenomenography

We have selected phenomenography (Marton & Booth, 1997) to guide the research process in this project, since this research approach offers an opportunity to gain insights into the *students' own experience* of their learning object and of their learning. Being qualitative and explorative, our project aims to produce results in the form of descriptions and insights in how students understand the concepts of networking.

Phenomenography has been used in earlier research on learning in higher education and has proved successful in computing education research. Berglund (2006) offers an overview of the use of phenomenography in computing education research, while Lister, Berglund, Box et al. (2007) goes further, exploring the computing teacher by the means of phenomenographic research.

The object of research in a phenomenographic project is the different ways in which the students *themselves* understand or experience something, in this case the general concept of network protocol. A general description of the approach can be found in, for example, Marton & Booth (1997). The outcome is a description of a limited number of qualitatively different ways in which this phenomenon is understood or perceived by students in a particular cohort. An understanding is, from a phenomenographic perspective, shaped by both the student and the phenomenon. This means that an understanding of something would not be the same if this "something" changed. For example, a student's understandings of the two network protocols TCP and

RMI is (hopefully ☺) different. In the same way, the learner changes by learning: A student is thus “not the same” after understanding TCP in a new or different way. This relation is illustrated by arrow 1 in Figure 1.

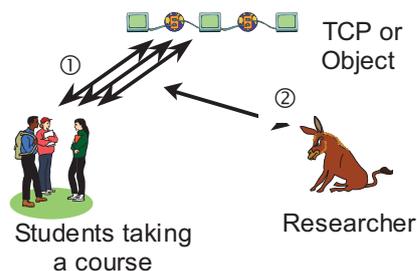


Figure 1. The relationships in phenomenographic research between the learner, the object of their learning and the researcher

Denzin and Lincoln (1994) argue that a qualitative research approach is “multi-method in focus, involving an interpretive approach to its subject matter”. This means that a qualitative researcher “studies things ... attempting to make sense of, or interpret, phenomena in terms of the meanings people bring to them.” (p. 2). The outcome of a phenomenographic research project is thus the researcher’s interpretation (indicated as arrow 2 in figure 1) of the students’ understanding of what they learn about. In this way, it is shaped by both the researcher and the object of his research.

3 The study

The findings discussed in the paper stem from data collected by interviews with computer science students in their third and fourth year, who took an open-ended project-based course in computer science (Daniels, 1999; Last et al., 2000). During the course, the students worked on a large project in teams of six, each team consisting of participants from the two universities concerned, one in Sweden and one in USA. The task of each team was to produce a software system intended to control a motorized toy. This demands good skills in computer systems, particularly in computer communication, since the design of the toy involved several computer communication solutions.

Data was collected through interviews with 15 students from both countries, selected to represent different ages, educational backgrounds, earlier study results etc. A thorough presentation of the research process is available elsewhere (Berglund, 2005).

4 Empirical findings

The findings are based on the analysis of the full interviews; the excerpts shown only serve to illustrate the categories and to give the reader a feel for the situation. To preserve anonymity, the excerpts from Swedish students are assigned names that start with an “S”, while American students are given names starting with “A”.

4.1 How students understand the concept of a network protocol

We have identified four different ways of understanding the general concept of a network protocol.

1. *A network protocol is a way of communicating between two computers*

In this category we encounter an understanding where network protocols are understood as way to communicate between two computers. Anthony articulates such an understanding:

Int: You talked about Java RMI. What is RMI?
 Anthony: I don’t even know. I know it’s a type of protocol used between, um, talking between two machines.

The discussion continues, and UDP and TCP are discussed:

Int: Uhum. What is TCP?
 Anthony: TCP is another type of protocol, used between two machines. There is TCP and there’s UDP that’s one of the things that I actually do remember from ah, networking class. And I believe TCP sends packets to one machine and then there is some sort of response saying that they got the packets or not.

While talking about TCP as another type of protocol, also used for communication between two machines, he spontaneously mentions UDP. By referring to the three protocols in this way, Anthony makes clear that there are properties that are shared between protocols: The three protocols mentioned are for communication or talking, and are experienced in a framework of two computers or machines.

2. *A network protocol is a method of communication over an internet*

This category describes a network protocol as a method of communication. The scope of this communication is wider, compared to the previous category: it is extended to a full network.

Later during the interview, Albert mentions RMI when answering a question about sockets:

Int: There is another word you mentioned there, and that’s socket. What is a socket?
 Albert: A socket is pretty much like a, a port that is opened up on the server, or that is requested by the client and, it’s assigned a number. And it’s just sitting there and listening and um, it’s just an open port and that port is um, designed to use a specific type of protocol, you know whether it be TCP, um, or the RMI. And it’s opened up to listen on that and once it receives that connection you know, it connects on that port. So it’s like an outlet socket, you know, you connect it in, you communicate and then when it’s down it gets turned off and then that port is either closed or it stays open if it’s required by the server.

It is clear that Albert sees UDP, TCP and RMI as protocols that share important properties, that are similar in many ways and that are parts of the same framework. He also describes the individual protocols in such a way as to suggest that he experiences a relationship between the protocols. There are differences between the three, but they are all closely related. Protocols are methods of communication.

3. *A network protocol is a set of rules*

This category discusses an understanding of network protocols as a set of formal rules experienced in a framework of an internet.

Allan discusses TCP as a protocol language used for sending data across a network in the following excerpt:

Int: Um, you've talked about TCP. What is TCP?

Allan: Basic concepts ... it's a protocol language, I guess you can call it, that you just put your data in and it's sent across the network using the different protocols you want to use, like IP or ... I can't think of any other protocols off my head. But it is more or less a packet that you put your data in and you send across and it has some features such as, keeps things in order when you, um, when you get to the, um, when it gets to the server you want to go to.

TCP is a protocol language that is used for sending data across a network. In this way, he clearly indicates his view that TCP is an integrated part of the network. He then explains its main feature, as he sees it: the order of data is preserved when sent to the application program through the TCP sockets, although the data packets might have arrived at the server in any order. Clearly, Allan has a focus on the end-to-end communication. The connection becomes safe, while the timing, routes, etc. of the individual packets vary.

His statements that TCP is "a protocol language", that you "put your data in", and that different protocols might be used for the actual transfer, indicate that protocols are sets of rules. He does not regard his answer as only valid for TCP, since he talks about different protocols, without

wanting to, or without being capable of, mentioning others by name. By mentioning the term IP, Internet Protocol, Allan relates to an internet.

4. *A network protocol is a standard*

The fourth category describes the most abstract understanding of the concept of a network protocol. Here the student abstracts the idea of a standard as well as the background against which it is seen, namely formal decisions.

This is illustrated by a discussion between the interviewer and Adam. Adam argues that the purpose of a protocol is to get computers to communicate. He then points out that TCP is a standard for a protocol, which was created by a committee. A protocol, in its turn, specifies the format for data sent across the network. Seen from this perspective, TCP is one of many protocols. Adam sees a standard as a set of rules that are created by a committee and are a result of human decisions.

4.2 Discussion of the categories

Since the categories offer different perspectives on *the same phenomenon*, they are logically related, according to the phenomenographic theory of learning. The relationships in our categories can be found in the framework of which the protocols are integrated parts. This framework is at its narrowest in category 1 (*A network protocol is a way of communicating between two computers*), consisting of only a limited set of computers. For categories 2 (*A network protocol is a method of communication over an internet*) and 3 (*A network protocol is a set of rules*) the framework is wider, encompassing a full network. In category 4 (*A network protocol is a standard*) it is broader still, encompassing not only computers and a network but also humans, as humans set standards. Table 1 summarizes these relationships between the four category.

A related perspective is given by Salleh et al. (2005), who studied how students understand what a computer network is. They showed two categories, labelled "Two or more communicating entities" and "A system of interconnected devices with associated protocols and structures to enable communication to occur". We see a clear association of the first of these categories with

	What is the general concept of a network protocol experienced as?	What is in focus?	Of what framework is it a part?
1	A protocol is a way of talking/communicating between two machines	Packets, and two computers	Two computers
2	A protocol is a method of communication on an internet	End-to-end-communication	A network
3	A protocol is a set of rules that are used on an internet	The rules governing the communication	A network
4	A protocol is a standard	Definitions	A world stretching outside the world of net-works

Table 1 A summary of the categories of the general concept of computer network protocol

category number 1 in our study, and of the second with categories 2, 3 and 4 of our work.

5 What do the results mean?

Insights concerning the value, or use, of different ways of understanding the protocols pave the way for a discussion concerning a desirable learning outcome. To discuss such issues, the researcher has to step out of his role as someone who only interacts with, analyses and describes data. She or he also has to interpret the results in a wider context, taking into account such considerations as applications and professional aspects.

5.1 What is important to know about TCP and other protocols?

The examples below illustrate that different ways of experiencing a network protocol are pertinent to different tasks. That is, a student must be able to experience a phenomenon in context-dependent ways in order to be able to solve the types of problem that arise during different parts of a project.

1. *Relevance for a programmer*

The description of TCP as communication between two computers (category 1) closely resembles discussions about programming. A quote from Sebastian serves as an illustration. When questioned by the interviewer about UDP, he compares UDP and TCP:

Int: UDP?

Sebastian: UDP ... but that is another form of communication. TCP/IP is set up ... like TCP, in contrast to UDP, TCP sets up communication between two points, and they talk to each other and make sure that they don't drop anything sort of.

TCP and UDP offer procedures, or operations, to a programmer who writes application programs. The procedures for TCP offer services like setting up a connection or sending data. The statements by Sebastian can be related directly to programming issues, for example when using TCP in an application program. As the interview continues, Sebastian returns to UDP, his way of talking still close to the issues of programming:

Sebastian: UDP is ... that the client asks what does this mean. Or what is this, or any question, whatever, and, so the server answers. And the server doesn't care in the end if the answer gets there or not. It is only a question and an answer, and then it is up to the client. If it feels that I didn't get any answer, it gets to ask again.

This line of reasoning is close to the steps that are taken by a program that uses the protocols, and is thus relevant when programming.

2. *Relevance for a program designer*

Experiencing a network protocol as a connection over a network (category 2) is useful when discussing the properties of a certain protocol or deciding which protocol to apply in a given situation. In focus here are

such issues as what circumstances and in what manner to use the protocol. This category is thus fruitful for designing project solutions and selecting between protocols.

The excerpt of an interview with Abraham can serve as an example:

Int: Um, what is RMI? Java/RMI? ...

Abraham: Very nice. It allows two Java virtual machines to talk to each other. They, an object on one machine could instantiate an object that lives on another machine and use that one's methods. That's how RMI is useful.

By discussing the advantages of RMI when asked what it is, Abraham clearly demonstrates an understanding of the purpose of RMI: execution of code on a remote machine. This understanding is useful for deciding when to use RMI in a particular system, or for deciding whether to choose another protocol during the design phase of a project.

3. *Relevance for theoretical development*

Category 3 describes a way to understand network protocols that to a large degree resembles that of category 2. The key difference lies in whether the protocol is seen as the application of a set of rules (category 2) or as the rules themselves (category 3). This distinction has implications for practical work: a focus on a theoretical development is certainly fruitful when developing purpose-made protocols.

4. *Relevance to policy issues*

The discussion about TCP as a standard (category 4) concerns how a protocol is developed, what possible protocols there could be, and what properties a protocol could have. This understanding is useful for policy discussions and for theoretical development, where the rules governing a protocol have to be thoroughly analysed and discussed.

This position is clear in the excerpt presented in section 1, where Alec argues that he is not aware of all features of RMI. RMI is, as he understands it, quick and concise, but it is not as flexible as he thinks it ought to be. His conclusion is that he does not know the features of the protocol well enough. This line of argument is relevant when considering policy questions, such as how to design network protocols.

The kind of reasoning presented here is thus relevant when developing protocols and considering policy questions.

Our conclusions concerning the relevance of different ways of understanding TCP are summarized in Table 2.

5.2 Learning Computer networks

One of the keystones in the phenomenographic theory of learning is that the way a student experiences a specific phenomenon, such as TCP or RMI, changes dynamically as the student reflects on, reads about, or discusses the phenomenon. According to phenomenographic theory, a

student does not have a given limit for his or her capacity to reach an advanced understanding. Rather, as students interact with the phenomenon they experience it in new ways.

5.3 What is desirable learning of network protocols?

Marton and Booth (1997) discuss good learning and argue that the ways in which learning is experienced “differ in richness (different aspects of learning that are discerned and held in focus simultaneously) and situational appropriateness (which particular aspects are held in focus under the prevailing conditions)” (ibid. p. 55).

5.4 Richness and situational appropriateness in ways of experiencing network protocols

To illustrate the value of richness in the understanding of network protocols, we observe that an understanding expressed in a more inclusive category offers the broader perspective needed to inspect and evaluate an understanding expressed in a less inclusive category. For example, to evaluate the solution to a programming problem concerning two interacting computers, that is, the code of an end-user program using TCP, it is necessary to shift to an understanding where the program is experienced in the framework of a network. “Stepping outside” the original reasoning, to look at the problem as a design issue rather than a coding issue, permits the discussion of questions about the efficiency, usefulness and relevance of the solution. As we have shown, discussions based on a more inclusive category in the field of computer networking are also more abstract, a feature that makes such judgements easier.

For solving new or complex problems it is thus necessary to shift between different ways of experiencing a protocol, since problem-solving involves different sub-tasks. But a shift in perspective, whether intra- or inter-contextual, whether triggered or spontaneous, is not in itself sufficient for problem-solving. Shifts must be made in a relevant way, they must be situationally appropriate, in order for the student to be capable of evaluating when and why a specific way of understanding a protocol is fruitful.

5.5 A professional perspective on the learning outcome

A student who can deploy and shift between different categories has a mastery of the concept of a network protocol that is a meaningful subset of the professional view (a view defined, for example, by Feit (1998)). Such an understanding encompasses the protocol both as a set of rules and as a software system. Furthermore, conscious shifts between the categories permit a student to work in different ways with the protocol. Still, differences can be found between the students’ categories and the professional view. Professionals tend to refer to the format of packets, modes of addressing and other protocol features in a manner that is not evident in the students’ statements. Whether these differences between

how professionals and students talk about protocols are due to the context of this study remains an open question for future investigation.

In exploring the students’ understanding of network protocols we have found a number of categories that are linked in a hierarchical structure, but we cannot conclude that any given category is “more functional” than another. In this respect, the results presented here differ from those of most phenomenographic research, where the highest category is frequently judged as the generally most desirable (many examples of such studies are available in Marton & Booth, 1997). However, Booth (1992) presents a similar result when discussing three categories of programming (product orientation, problem orientation, computer orientation), identified within a cohort of novice engineering students. The three categories, she argues, are all relevant for professional programming during the phases of design, prototyping and coding, respectively.

It is thus tempting to search for an explanation for these differences in the character of computer science itself. The structure of computer science, with respect particularly to education, but also to professional life, is discussed in the IEEE/ACM Computing Curricula 2001 (IEEE/ACM, 2001). The report states that the foundations of computer science are drawn “from a wide variety of disciplines” (p. 12) and that “all computer science students must learn to integrate theory and practice, to recognize the importance of abstraction, and to appreciate the value of good engineering design”. Particularly for the area of “net-centric computing” it is argued that “mastery of this subject area involves both theory and practice”. A professional in computer science must then be familiar with these different aspects of the field.

This discussion can be applied to the results concerning network protocols in this study. The statement from the

	What is the general concept of a network protocol experienced as?	What is it relevant for?
1	A protocol is a way of talking/communicating between two machines	Programming
2	A protocol is a method of communication on an internet	Project design
3	A protocol is a set of rules that are used on an internet	Theoretical development
4	A protocol is a standard	Development of new communication solutions

Table 2 The relevance of different ways of understanding computer network protocols

ACM curriculum committee supports the conclusion that a computer networking concept ought to be understood in several ways. In particular, the report stresses the need to understand a computer system in a more advanced way than merely as a program:

Graduates of a computer science program must develop a high-level understanding of systems as a whole. This understanding must transcend the implementation details of the various components to encompass an appreciation for the structure of computer systems and the processes involved in their construction and analysis. (p. 62)

On the other hand, the inclusive structure of the categories indicates that, from a learning perspective, the more advanced categories are the more desirable, since an advanced category necessarily includes the less advanced ones.

6 Conclusions

Problem-solving in computer networks, as within computing in general, normally involves working with new or complex sub-tasks of different character. It is thus necessary to have a repertoire of different ways of understanding the network protocols, and a capacity to shift between these different categories. To shift perspective is not alone sufficient for problem solving. Shifts have to be made in a relevant way, and the student needs to be capable of evaluating when and why a specific way of understanding a network protocol is fruitful.

7 References

- Berglund, A. (2005). Learning computer systems in a distributed project course: The what, why, how and where (Vol. 62). Uppsala, Sweden: Acta Universitatis Upsaliensis.
- Berglund, A. (2006). Phenomenography as a way to research learning in computing. *Bulletin of the National Advisory Committee on Computing Qualifications, BACIT*, 4(1).
- Booth, S. (1992). *Learning to Program. A phenomenographic perspective*. Gothenburg, Sweden: Acta Universitatis Gothoburgensis.
- Daniels, M. (1999). Runestone, an International Student Collaboration Project (No. 11). Linköping: NyIng report.
- Denzin, N. K., & Lincoln, Y. S. (1994). *Handbook of Qualitative Research*: SAGE.
- Feit, S. (1998). *TCP/IP*. New York, NY, USA: McGraw-Hill.
- IEEE/ACM (2001) Computing Curricula 2001, Computer Science Volume. <http://www.sigcse.org/cc2001/cs-principles.html>
- Last, M., Almstrum, V., Daniels, M., Erickson, C., & Klein, B. (2000). An International Student/Faculty Collaboration: The Runestone Project. In the proceedings of the ACM SIGCSE/SIGCUE Conference on Innovations and Technology in Computer Science Education (ITiCSE'00), Helsinki, Finland
- Lister, R., Berglund, A., Box, I., Cope, C., Pears, A., Avram, C., et al. (2007). *Differing Ways that Computing Academics Understand Teaching*. In the proceedings of the Australasian conference on computing education, Ballarat, Australia, 97-106.
- Marton, F., & Booth, S. (1997). *Learning and awareness*. Mahwah, New Jersey, USA: Lawrence Erlbaum Associates.
- Ramsden, P. (1996). *Learning to Teach in Higher Education*. London, UK; New York, USA: Routledge.
- Salleh, E. A., Winslett, G., Bruce, C., & Tickle, A. (2005). *Diagnostic Learning: Using web-based self diagnostic tools for learning abstract concepts in data network education*. In the proceedings of the Online Learning and Teaching Conference (OLT) 2005 conference, Brisbane, Australia, 241-250.