

Assessing Student Learning in a Virtual Laboratory Environment

Tilman Wolf, *Senior Member, IEEE*

Abstract—Laboratory experience is a key factor in technical and scientific education. Virtual laboratories have been proposed to reduce cost and simplify maintenance of lab facilities while still providing students with access to real systems. It is important to determine if such virtual labs are still effective for student learning. In the assessment of a graduate computer networks course, the author quantifies the amount of learning that is observed in lectures and labs. The results not only show that learning indeed occurs during lab sessions, but almost equally as much (45.9%) as in lectures (54.1%). Also, it is observed that even students who have prior experience in networking benefit from virtual labs.

Index Terms—Assessment, computer networks education, educational technology, learning events, virtual laboratory.

I. INTRODUCTION

IN ADVANCED technical education, laboratories are important environments for student learning. High setup and maintenance costs make it challenging to provide laboratory facilities at institutions where budgets and technical expertise are limited. In response, “virtual laboratories” have been developed to address these challenges. Virtual laboratories allow users to perform *experiments on real systems* via remote access. Thereby, laboratory resources can be shared among a large community to geographically distributed users while limiting setup and operational costs to a single facility.

While virtual laboratories are beneficial from an economic and organizational point of view, it is not clear that they can achieve the most important goal—student learning. In this paper, an assessment study is presented that explores this question in the context of a graduate course in computer networks. In particular, the contributions of this paper are:

- an assessment methodology that can be used to quantify student learning in lectures and labs,
- detailed assessment results from a graduate networking course that uses a virtual laboratory for computer networks experiments, and
- the observation that learning is indeed observed when using the virtual lab (to an almost equal degree as in lectures) and that even students with prior experience can benefit.

Manuscript received January 12, 2008; revised December 18, 2008. First published June 30, 2009; current version published May 05, 2010. This work was supported by the National Science Foundation under Grant DRL-0632586.

The author is with the Department of Electrical and Computer Engineering, University of Massachusetts Amherst, Amherst, MA 01003 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TE.2008.2012114

The remainder of the paper is organized as follows. Section II discusses related work. Section III introduces the virtual network laboratory and how it is used in the context of the evaluated course. The assessment methodology is described in Section IV. Assessment data and their interpretation are presented in Section V. Section VI summarizes and concludes this paper.

II. RELATED WORK

The importance of laboratory experience in engineering education (and other fields) has long been recognized [1]. One of the main goals of laboratory education is for students to gain experience in the theory and practice of experimentation [2]. Such experimental skills are considered crucial in the sciences as well as computer science [3]. Therefore, it is important that laboratories teach students how to deal with real or realistic environments. This exposure includes potential problems and failures (e.g., as they occur in the context of network security [4]). In this context, simulators may not be sufficiently realistic unless they have been designed specifically with realism (not just accuracy) in mind [5].

Some virtual laboratories, including the Open Network Laboratory (ONL) [6] that is used in this study, solve this tension between realism and easy use and access by providing remote access to a experimental facility that houses real systems. Using an easy-to-use interface, these virtual laboratories provide access to a real system with all or most of its operational characteristics. The use of remote access to educational resources has been explored in the distance education domain [7], [8]. Other virtual laboratories, which aim for realism but do not use a central shared resource, have been designed for various application domains (e.g., thermodynamics [9], civil engineering [10], geotechnical engineering [11], and cell biology [12]).

The virtual lab that is used in this study is the ONL that has been specifically designed for computer networks education [6]. While there are other virtual and remotely accessible lab facilities for computer networks (e.g., Emulab [13] and Planetlab [14]), they are mainly designed for research and are thus more complex to use. The general teaching experience with ONL has been reported previously [15], and therefore the discussion of ONL features and functionality in this paper is limited to a very brief overview. Instead, the focus of this paper is on the assessment methodology, which is roughly modeled after National Science Foundation (NSF) recommendations for evaluation processes [16], and the assessment results.

While virtual laboratories have been used in several domains, the number of assessment studies is small. Yarbrough and Gilbert have studied the effectiveness of multimedia-based

laboratory experiments for civil engineers [10]. They found the virtual lab to be beneficial, but their assessment relies on self-reported perception of learning by students. Wyatt *et al.* have assessed student use of a virtual geotechnical laboratory [11], but they were not able to relate this information to student learning. Kinnari *et al.* have evaluated a networking laboratory qualitatively with student surveys [17]. In the work presented here, the aim is to quantify the amount of learning that is contributed by a virtual laboratory, especially in comparison to lecture learning.

III. VIRTUAL NETWORK LABORATORY

To provide the necessary background, this section first takes a brief look at why laboratory components are important in computer networks education and how ONL fits into this picture.

A. Computer Networks Education

Computer networks are an important infrastructure component in today's world. Society increasingly relies on communication, information, and services that are delivered via computer networks (most notably the Internet). This importance is reflected in computer science and computer engineering curricula, where a basic computer networks course is practically always a requirement for undergraduate students. Even for students who want to focus on other areas, computer networks are important since they tie in with many other domains (e.g., distributed computing, operating systems, embedded systems, wireless communication, etc.).

A typical introductory networking course covers the principles of network design and operation: protocols that define interactions on links, network nodes, and end-system applications; algorithms that enable shortest-path routing, efficient prefix lookups, and packet scheduling; system designs that provide performance guarantees under worst-case traffic scenarios; security issues and solutions. Often, theoretical concepts are illustrated with examples from real networks. This aspect is particularly important since Internet technology is dominated by a set of well-established protocols (i.e., TCP/IP). Without knowledge of this particular technology, a student would have a difficult time applying networking concepts in industry or even in research.

B. Virtual Laboratories

Due to the need for practical knowledge about protocols and systems, networking courses often include laboratories. Typically, these labs either consist of 1) hardware in a lab room for hands-on exercises or 2) software to simulate networks and to complete programming assignments on a desktop computer. The pros and cons of both of these approaches are shown in Table I ('+' indicates that a particular type of laboratory is rated positive at achieving the quality listed in the table, '-' indicates negative rating, and 'o' indicates neutral rating).

While hardware labs are completely realistic, they are difficult to use, costly to set up, and time-consuming to maintain. Software labs are easier to set up, use, and manage, but lack realism. Considering that realism is probably the most important quality of a laboratory when it comes to student learning, the shortcomings of software labs are significant.

TABLE I
COMPARISON OF LABORATORY TYPES USED
IN COMPUTER NETWORKS EDUCATION

	Hardware Lab	Software Lab	Virtual Lab
High realism	+	-	+
Ease of use	-	+	+
Low cost	-	+	+
Ease of maintenance	-	+	o

To address these challenges, "virtual laboratories" have been proposed. In the context of computer networks, a virtual laboratory uses hardware that is located in one centralized place. Users can access this hardware and perform experiments remotely through a software interface. The experiments performed by the user are executed on actual hardware, while results are monitored via the software interface. The separation of hardware and physical presence by the user makes it possible to provide a realistic laboratory experience to many off-site users while limiting the cost and maintenance overhead to a single site.

C. Open Network Laboratory

This paper focuses on one particular instance of a virtual laboratory, the ONL. This laboratory has been designed to hide the complexities of configuring high-performance network systems and solely to expose the features that are important for a good learning experience by students.

The hardware of ONL includes several 8-port gigabit routers with programmable port processors, dozens of end-systems, and gigabit-per-second network connections. The user interface for configuring and monitoring ONL is shown in Fig. 1. The systems shown in the user interface correspond to actual hardware in ONL. Users can remotely log into all end-systems used in the experiment. These machines are full-blown workstations, and any software can be run on them.

To illustrate the operation on ONL, consider an experiment where a student is to set up correct route table entries in the network. To change route table entries on the routers, a user simply updates the values in the user interface (as shown in Fig. 1). It is not necessary to understand the router-specific sequence of commands that updates the route table. Instead, the user can focus on the higher level question of determining correct route tables entries. Then, the user can log into any of the end-systems and use tools (e.g., ping or traceroute) to determine if the network behaves as desired. Traffic that is sent across the routers can be monitored and visualized in real-time. More details on the teaching experience with ONL can be found in [15].

D. Assessed Course

In this study of learning with a virtual laboratory, student learning is assessed throughout a course that uses ONL in labs. This course is "ECE697AA – Computer Networks" and was taught by the author in Spring 2007 at the University of Massachusetts Amherst. ECE697AA is a graduate course that is required for the doctoral program and is also often taken by Master's degree students. Thus, students have a broad range

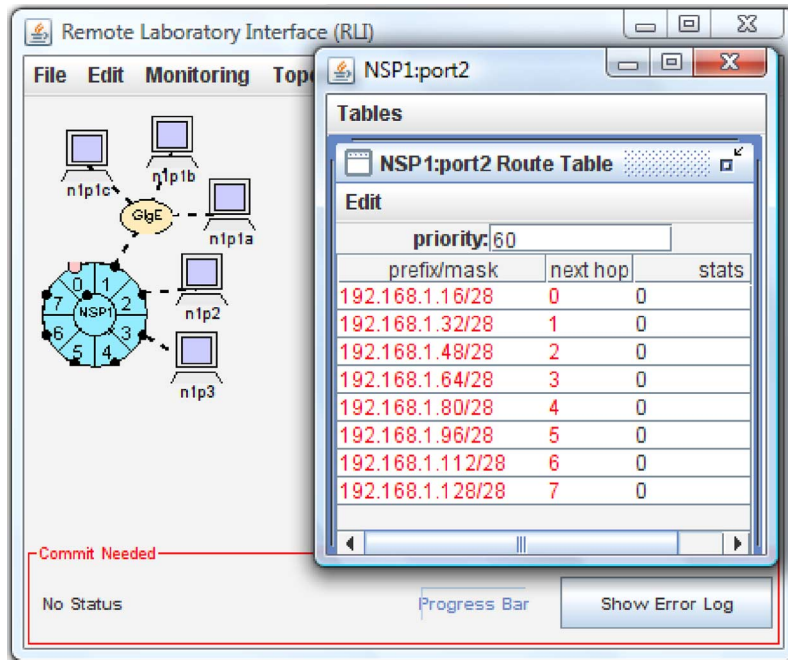


Fig. 1. Open Network Laboratory interface. The figure shows a virtual network topology and the interface to editing route table information.

TABLE II
LABORATORY ASSIGNMENTS

Lab	Networking Concepts	Lab Assignment	ONL Use
1	Routing	Basic connectivity, routing configuration	50%
2	Queuing theory and network delays	Bandwidth and queueing delay measurement	90%
3	Scheduling and congestion control	TCP and UDP scheduling, hypothesis testing	90%
4	Router design, lookups, and network measurement	Space-time diagram analysis, passive network measurement	40%

of background knowledge ranging from no prior exposure to networks to being researchers in the area.

The course consisted of 26 lectures, seven homework assignments, and four lab assignments. The lab assignment topics and related networking concepts are shown in Table II. In each lab, ONL was used to set up and perform experiments and to obtain measurement data. Some parts of each lab assignment required offline data analysis that did not involve ONL. The percentage of ONL use (i.e., the fraction of the assignment that uses ONL) is also shown in Table II. All lab topics were covered to some extent in lectures.

IV. ASSESSMENT METHODOLOGY

The main question that is addressed in the assessment study is if students benefit from use of a virtual laboratory. That is, do students learn anything while using the virtual laboratory (as compared to learning solely from lectures)? To answer this question, an assessment methodology was developed that distinguishes between learning in lectures and learning in the lab.

A. Assessing Lecture and Lab Learning

In a typical course, it can be expected that a new topic is first covered in lecture and then studied in more detail in the lab. Rarely, students are expected to complete lab assignments without having been taught the foundations in lecture. Thus, it is assumed that lectures precede lab exercises. In such a scenario, students may learn during the lecture and/or may learn during the lab. To assess when this learning takes place, the proposed assessment methodology evaluates student knowledge three times.

- First assessment: prelecture assessment of prior knowledge.
- Second assessment: post-lecture and prelab assessment of learning during lecture.
- Third assessment: post-lab assessment of learning during lab (and retention from lecture learning).

In each assessment, students are asked the same set of questions related to the particular topic shown in Table II. Students did not know what questions they would be asked or that there would be overlap between assessments. Based on the answers, it is

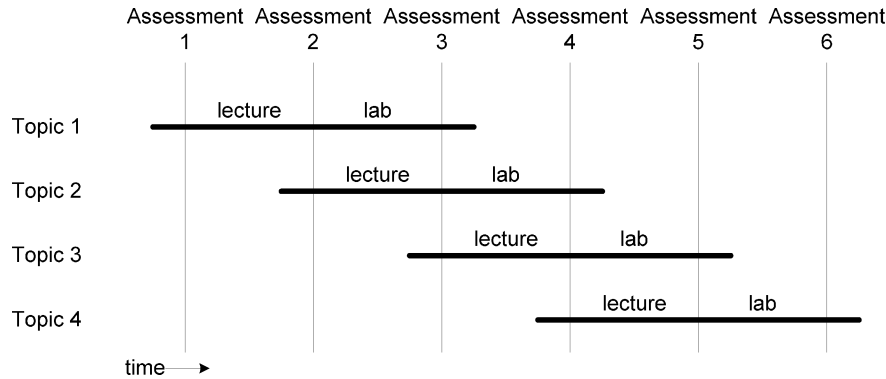


Fig. 2. Schedule of lectures, laboratory assignments, and assessments.

Q1-4: Avoiding Loops: How do switches avoid loops in LANs? (check all that apply)

- (N) by tracking the paths of packet
- (Y) by exchanging information between switches
- (Y) by turning on and off interfaces on the switch
- (N) by running OSPF within the autonomous system

Q2-11: On-Off User Model: Assume user behavior that follows the (discrete time) Markov model below. If A indicates that the user is transmitting data, and B indicates that the user is silent, which of the following statements are correct?

Figure 3 shows a Markov model with two states, A and B. State A is represented by a circle on the left, and state B is a circle on the right. From state A, there is a self-loop arrow labeled $3/4$ and an arrow pointing to state B labeled $1/4$. From state B, there is a self-loop arrow labeled $1/2$ and an arrow pointing back to state A labeled $1/2$.

- (Y) On average, the user is more often transmitting than silent.
- (N) On average, the user is more often silent than transmitting.
- (N) On average, the user is equally often silent and transmitting.
- (N) It cannot be determined if the user is more often silent than transmitting.

Fig. 3. Example of two assessment questions.

inferred if and when students increase their knowledge and thus exhibit learning.

Clearly, one major challenge in this assessment methodology is the sheer number of assessment events (three per lab). In order to reduce the number of assessment questionnaires that students have to fill out, it is possible to pipeline the process as shown in Fig. 2. For example, when performing the second assessment for the first lab, the first assessment for the second lab can be scheduled at the same time. Thus, for the four lab assignments used in this study, only six assessment events are required (instead of 12).

For each lab topic, the author created a set of 10–14 questions that covers the most important concepts. Two examples of such questions are shown in Fig. 3. All questions used multiple-choice answers and were submitted electronically to allow automatic evaluation. Students were permitted to submit their answers at their convenience within a certain time frame (e.g.,

after the end of a lecture and before the posting of a lab assignment).

B. Inferring Learning Events

With the assessment process in place, the process of inferring learning events needs to be considered. Since students are asked the same questions before the lecture, between lecture and lab, and after the lab, a progression of student performance over time can be observed. For example, if a student does not answer a question correctly in the first assessment but then answers it correctly in the second and third assessment, it is inferred that learning took place during the lecture (or during activities related to the lecture, e.g., reading). Of course, there are numerous sequences of events that can take place. Also, students may skip assessments and thus only provide incomplete data.

To formalize the sequence of events, the proposed assessment methodology proposes a “triple” that describes the performance




Fig. 4. Learning events during assessment.

TABLE III
INTERPRETATION OF TRIPLE ASSESSMENT RESULTS

		learning		
		yes	unknown	no
location	lecture	011	01x, x11, x1x	000, 010, 0x0, 100, 101,
	lab	001	00x, x01, x0x	10x, 110, 111, 11x, 1x0,
	unkwn.	0x1	0xx, xx1, xxx	1x1, 1xx, x00, x10, xx0

of a single student on a single assessment question. The triple contains three indicators of the performance of the student on each of the three assessment steps:

- 1: a “1” indicates that the question was answered correctly.
- 0: a “0” indicates that the questions was not answered correctly.
- x: a “x” indicates that the student did not participate in this particular assessment.

There are $3^3 = 27$ possible triples. Table III shows the interpretation of these triples with regard to learning. For example, “011” shows learning during lecture, while “001” shows learning during the lab. Cases where there is no retention (e.g., “010”) are not considered as learning. The case of knowledge loss (e.g., “100”) is considered as no learning. The lack of assessment data (i.e., “x”) contributes toward unknown events in which it may be unknown if learning took place and/or where it took place.

V. RESULTS

The results of the presented study are based on data collected during one semester. A total of 29 students participated in the study. Each student provided informed consent to collecting this data. Also, all data are anonymized and/or presented in aggregate form to avoid exposure of any personally identifiable information.

A. Participation

The participation in the assessment was voluntary, but all students enrolled in the course chose to participate. Lecture attendance was not tracked for this study, but all participants completed all four lab assignments. Even if students agreed to participate in the assessment study, it was not possible to enforce their filling out all assessment questionnaires. The participation results are shown in Table IV. While the participation was high

TABLE IV
ASSESSMENT PARTICIPATION

Assessment	Number of Question	Participation	Burden for Quiz (min/avg/max)
1	14	21/29 (72.4%)	5/17/37 min.
2	14+12=26	27/29 (93.1%)	10/22/45 min.
3	14+12+11=37	24/29 (82.8%)	10/37/100 min.
4	12+11+10=33	25/29 (86.5%)	10/25/60 min.
5	11+10=21	17/29 (58.6%)	8/18/30 min.
6	10	10/29 (31.0%)	— ^a

^aNo accurate information on the burden for Assessment 6 is available since several other questions were asked to study retention. Students included the time to answer these questions in their reported burden for this assessment.

for the first four assessment events, it dropped considerably for the last two. This effect is taken into account when presenting the data, and particular attention is paid to data where assessment results are (nearly) complete. One reason for the drop in participation may be the overall time burden for the quiz (shown in Table IV).

B. Learning Events

Based on the assessment results, a total of $29 \cdot 47 = 1363$ triples are obtained, with each indicating how each student performed on each question. A summary result is shown in Fig. 4. For each student, it is shown if and where learning occurred (according to Table III). The columns are ordered by assessment participation. Students who completed all assessment are shown on the left. Each column is a stacked bar, which shows learning in lecture on the bottom, learning in lab in the middle, and learning that took place either in lecture or lab on top. Solid fills indicate that learning definitely happened (i.e., complete triple), and lighter shading indicates that learning may have happened (i.e., triple with one or more ‘x’).

For example, the student represented by the left-most entry in Fig. 4 showed learning in lecture for 8% of the questions and showed learning in lab for 14% of the questions. For the remaining 78% of the questions, no learning was observed. As the number of assessments goes below six, the results become more ambiguous. This is illustrated by lighter shades that indicate possible learning and light gray bars that indicate learning in either lecture or lab.




Fig. 5. Learning in relation to prior experience. Each graph shows where learning happened depending on number of prior networking courses taken. Point size is proportional to final grade.

The first observation is that definite learning can be observed in up to 27.7% of the assessment questions. The average learning is 10.5% (for students who participated in three or more assessments). As expected, the amount of definite learning that can be observed decreases with a decreasing level of assessment participation.

The second observation is that for definite learning, the amount of observed lecture learning (average of 54.1% of definite learning) and lab learning (average of 45.9% of definite learning) is roughly the same. This result underlines the importance of lab components in courses like the one studied in this paper. Lectures alone can only partially contribute to student learning; lab components are almost equally as important.

C. Learning and Prior Experience

Another question with respect to learning in a virtual laboratory is how much previous experience students need for them to benefit. To answer this question, the observed learning events are related to the number of prior networking courses taken by students. Fig. 5 shows this data in four graphs (one for zero, one, or two or more prior networking courses and one for students where this information is not available). Black stars show the data for definite learning events (i.e., no 'x' in the triple), and gray squares show definite plus possible learning events. The size of the star or square is proportional to the final grade that was obtained by that student.

From Fig. 5, it can be seen that most learning occurred for students with no or one prior course. In all cases, there is a high probability that additional learning occurred during the lab, but this is not certain due to the lack of assessment data. Nevertheless, it can be observed that the virtual laboratory enables learning even for students who have previous experience with computer networks. When considering the final grade achieved by students, there is no apparent correlation to learning events in lectures versus labs.

VI. SUMMARY AND CONCLUSION

In this paper, the impact of a virtual laboratory on student learning in a computer networks course is studied. An assessment methodology is introduced that can distinguish learning events occurring in lectures and labs. Assessment results are presented that show learning can indeed be observed and that

the amount of learning in the virtual lab is approximately equal to learning in the lecture.

For future studies in this area, the author is considering several extensions that may provide further insights. It may be desirable to compare the amount of learning achieved in a virtual lab with that in other types of labs (physical lab or software simulation). Also, the uncertainties observed in the data due to voluntary participation in assessments may be overcome by requiring students to fill out all questionnaires. Finally, the assessment methodology may be calibrated by studying how much students learn by simply repeating assessment questions (without lectures or labs). Despite these open questions, the author believes that this study provides important initial insights into the effectiveness of virtual laboratories and their impact on student learning.

ACKNOWLEDGMENT

The author would like to thank C. Kotori from the School of Education at the University of Massachusetts Amherst for her insightful comments while developing the assessment methodology described in this paper. The author would also like to thank K. Wong and C. Wiseman from the Department of Computer Science and Engineering at Washington University in St. Louis (supported by NSF Grant DRL-0632580), who kept ONL up and running and answered many of the author's and his students' questions.

REFERENCES

- [1] "New directions in laboratory instruction for engineering students," *J. Eng. Educ.*, vol. 58, pp. 191–195, Nov. 1967.
- [2] E. W. Ernst, "A new role for the undergraduate engineering laboratory," *IEEE Trans. Educ.*, vol. E-26, no. 2, pp. 49–51, May 1983.
- [3] W. F. Tichy, "Should computer scientists experiment more?," *Computer*, vol. 31, no. 5, pp. 32–40, May 1998.
- [4] J. M. D. Hill, C. A. Carver Jr., J. W. Humphries, and U. W. Pooch, "Using an isolated network laboratory to teach advanced networks and security," *SIGCSE Bulletin*, vol. 33, no. 1, pp. 36–40, Mar. 2001.
- [5] A. Francis and M. Couture, "Realism in the design process and credibility of a simulation-based virtual laboratory," *J. Comput. Assisted Learning*, vol. 20, no. 1, pp. 40–49, Feb. 2004.
- [6] J. DeHart, F. Kuhns, J. Parwatikar, J. Turner, C. Wiseman, and K. Wong, "The Open Network Laboratory: A resource for networking research and education," *Comput. Commun. Rev.*, vol. 35, no. 5, pp. 75–78, Oct. 2005.
- [7] B. Aktan, C. A. Bohus, L. A. Crowl, and M. H. Shor, "Distance learning applied to control engineering laboratories," *IEEE Trans. Educ.*, vol. 39, no. 3, pp. 320–326, Aug. 1996.

- [8] A. Bagnasco, P. Buschiazzo, D. Ponta, and M. Scapolla, "A learning resources centre for simulation and remote experimentation in electronics," in *Proc. PETRA*, Athens, Greece, Jul. 2008, pp. 1–7.
- [9] K. D. Forbus, S. E. Kuehne, P. B. Whalley, J. O. Everett, L. Ureel, M. Brokowski, and J. Baher, "CyclePad: An articulate virtual laboratory for engineering thermodynamics," *Artif. Intell.*, vol. 114, no. 1-2, pp. 297–347, Oct. 1999.
- [10] S. E. Yarbrough and R. B. Gilbert, "Development, implementation, and preliminary assessment of virtual laboratory," *J. Professional Issues in Eng. Educ. Practice*, vol. 125, no. 4, pp. 147–151, Oct. 1999.
- [11] T. R. Wyatt, P. Arduino, and E. J. Macari, "Assessment of a virtual laboratory for geotechnical engineering education," *Comput. Educ. J.*, vol. 10, no. 2, pp. 27–35, Apr. 2000.
- [12] A. R. White, P. E. McClean, and B. M. Slator, "The virtual cell: An interactive, virtual environment for cell biology," in *Proc. of World Conf. Educ. Multimedia, Hypermedia and Telecommun. (ED-MEDIA)*, Seattle, WA, Jun. 1999, pp. 1442–1443.
- [13] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb, and A. Joglekar, "An integrated experimental environment for distributed systems and networks," in *Proc. 5th Symp. Operat. Syst. Design and Implement.*, Boston, MA, Dec. 2002, pp. 255–270.
- [14] L. Peterson, T. Anderson, D. Culler, and T. Roscoe, "A blueprint for introducing disruptive technology into the Internet," in *Proc. HotNets-I*, Princeton, NJ, Oct. 2002.
- [15] K. Wong, T. Wolf, S. Gorinsky, and J. Turner, "Teaching experiences with a virtual network laboratory," in *Proc. 38th SIGCSE Tech. Symp. Comput. Sci. Educ.*, Covington, KY, Mar. 2007, pp. 481–485.
- [16] J. Frechtling, "The 2002 user-friendly handbook for project evaluation," Nat. Sci. Found., Arlington, VA, Tech. Rep. NSF 02-057, 2002.
- [17] S. Kinnari, A. Suhonen, J. Harju, and K.-J. Rähä, "Evaluating a remote access networking laboratory as a learning environment," in *Proc. WBED*, Chamonix, France, Mar. 2007, pp. 73–79.

Tilman Wolf (SM'07) received the Diplom. in informatics from the University of Stuttgart, Stuttgart, Germany, in 1998, and the M.S. degree in computer science, the M.S. degree in computer engineering, and the D.Sc. degree in computer science from Washington University in St. Louis, St. Louis, MO, in 1998, 2000, and 2002, respectively.

He is an Associate Professor in the Department of Electrical and Computer Engineering at the University of Massachusetts Amherst. He is engaged in research and teaching in the areas of computer networks, computer architecture, and embedded systems. His research interests include network processors and their application in next-generation Internet architectures, embedded system security, and how to leverage technology for teaching.

Dr. Wolf is a Senior Member of the Association for Computing Machinery (ACM). He has been active as a program committee member and organizing committee member of several professional conferences, including IEEE INFOCOM and ACM SIGCOMM. He is currently serving as Treasurer for the ACM SIGCOMM Society. At the University of Massachusetts, he received the College of Engineering Outstanding Teacher Award in 2008.