An Integrated View of Teaching and Learning for a Foundational Course on Computer Networking

Rocky K. C. Chang
Department of Computing
The Hong Kong Polytechnic University
Hung Hom, Kowloon, Hong Kong
Email: csrchang@comp.polyu.edu.hk

***Abstract*— This paper presents an integrated view of teaching and learning to address some of the obstacles to learning Computer Networking effectively. This view specifically addresses teaching and learning activity integration, peer integration, and teaching and assessment integration. As a result of this view, this paper prescribes effective approaches for lecturing, analytical problem-solving, experimental problem-solving, group-based learning, and assessments. These approaches have been applied to a foundational course on Computer Networking offered to junior students of a typical Computer Science or Computer Engineering curriculum. Students' feedbacks on the effectiveness of these approaches are reported for a small class (54 students) and a big class (168 students). The feedbacks also show that students' interests in Computer Networking are more dichotomized after taking the foundational course.**

I. Introduction

Computer networks and the Internet have already become the most critical infrastructure today for information dissemination, business transactions, human communications, computer games, scientific computation, and even national security. Moreover, computer networks are ubiquitously deployed in many other commercial sectors, such as automobiles, smart highways, smart clothing, smart appliances, and coffee shops, just to name a few. As a result, there is a pressing need to equip students on both undergraduate and graduate levels with a solid foundation in the field in order to further foster and flourish its development and applications.

However, in many ways Computer Networking education is still in an exploratory stage. For example, we have seen for the last ten years many different approaches adopted by various textbooks—analytical [1], bottom-up [2], engineering [3], system [4], balanced-view [5], visual [6], and top-down [7]. Whether these approaches can effectively facilitate students' learning are yet to be seen and evaluated.

As the content of the field continues to increase in a very rapid rate, it is also very important to identify a minimal set of core principles to teach undergraduate students, and this has been discussed in the ACM Workshop on Networking Education last year [8]. Moreover, there is increasing effort on providing hands-on experience to students through more traditional laboratory sessions/courses, implementation of networking hardware, simulation tools, and socket programming.

This paper considers a foundational course on Computer Networking for junior students in a typical Computer Science or Computer Engineering program. However, this paper's focus is quite different from previous works which usually considered a particular aspect of teaching activity, such as hands-on work, or curriculum issues. Instead, this paper attempts to give a more holistic treatment by considering all teaching and learning aspects. The main thesis is that all teaching and learning activities, individually and corporately, are essential to effective teaching and learning of the subject. In particular, these activities need to be integrated in a more thoughtful manner to yield the most effective learning outcome.

A. Identifying learning and teaching difficulties

"Why is it (not) so difficult to learn Computer Networking?" is perhaps the first question that we as educators should ask ourselves (and our students). One of the obvious answers is of course to do with the rapid advances in the field. Only ten years ago, one undergraduate subject and one postgraduate subject were perhaps sufficient to provide a reasonable coverage. Today, however, an entire undergraduate program can be devoted to Computer Networking, and its associated subjects, such as economic and social issues. Many Computer Networking topics have also quickly developed into separate subjects, notably Wireless and Mobile Networks, Optical Networks, and Network Security.

From the educators' point of view, it also becomes increasingly difficult to teach the subject effectively, partly because the Internet's success has attracted students with various backgrounds. Thus, it is hard to provide one class for all. Even among students with similar educational backgrounds, say Computer Engineering students, some of them may have taken professional examinations, such as Cisco's CCNA, before attending a first course on Computer Networking, while others are not even able to expand the term TCP/IP.

Besides the issues related to the forever changing nature of the field (at least for now) and heterogeneity in students' backgrounds, there are quite a number of obstacles to teaching and learning Computer Networking effectively, as outlined below.

1) The principles underlying Computer Networking are intrinsically very complex. The layered model helps understand and manage the complexity. But very soon students find out that this layering approach has its own inadequacy. For example, the layers are not really

independent of each other, and a network layer could be degenerated into a data-link layer, e.g., IP over ATM. Therefore, getting the whole picture correctly is already a challenge to many students.

2) Computer Networking concepts and protocols are also very abstract to many students. For one, they cannot see typical networking equipments, and visualize packets and protocols for themselves during lectures. For example, talking about a LAN switch without seeing one is already a hinderance for many students. Examining the kernel code without a prior understanding of how the protocols work helps very little. Computer animation alleviates the problem to some extent, but it may not be able to equip students with the ability of conceptualizing other more abstract concepts on their own.
3) Unlike Computer Programming and Computer Architecture courses, for example, resource provision for hands-on practical experience in Computer Networking is problematic. A Computer Networking laboratory, if available, usually has a ready-to-use computer network on which students can conduct various experiments. However, students should also be expected to know how to set up a computer network from scratch after taking the course, very much like knowing how to write programs after taking a Computer Programming class. But it is not quite possible to provide adequate resources to achieving that in many academic environments.
4) Many terminologies and acronyms are introduced and used in the field, and some of them are very similar, e.g., ARP and ARQ. Worse yet, these terminologies are often not used consistently, especially in the industry, e.g., hubs, switches, switching hubs, port switching hubs, and segment switching hubs. Students, on the other hand, are more used to clear and formal definitions. For example, many students in my class are very perplexed by the term *round-trip time* when it is first mentioned during the subject overview.
5) Some of the networking problems are difficult to comprehend and appreciate due to students' common lack of practical experience. It is generally not difficult for undergraduate students to understand the access network technology, because almost all of them have experience of assessing the Internet via ISPs and LANs. However, it is relatively difficult to comprehend the scalability issues in inter-domain routing and the importance of traffic engineering.
6) The sequence of coverage can also affect students' learning significantly. Computer Science students are perhaps more comfortable with a top-down approach while engineering students may find the pure bottom-up approach more logical. In either case, it is bound to have missing information, which is yet to be revealed later, in order to complete the entire storytelling of the Computer Networking internals. Unfortunately, students may stumble on these missing information. Therefore, a good sequence of coverage may follow a "nonlinear" path which is in contrast to a rigid top-down or bottom-up sequence.

B. *Course objectives*

There must be a set of objectives to achieve in any course, including Computer Networking. These objectives and their order influence how the course is conducted in every aspect. For example, if the objectives are mainly for training researchers in the field, the course should perhaps put more emphasis on the theoretical foundation, such as protocol correctness and performance issues. On the other hand, if the objective is mainly to meet the industry's needs, more practical training, whether socket programming or network setup and diagnosis, should be included to enhance students' marketable skills.

This article considers a first course on Computer Networking offered in the junior year of a typical Computer Science or Computer Engineering curriculum. In what follows, the discussion will be based on the course that I have been teaching for the last ten years in The Hong Kong Polytechnic University [9]. Although there are clear differences between my teaching environment and others, I will attempt to concentrate on issues that, I believe, are common to similar courses offered in other academic institutions.

The overall objective of this course is to equip students with a solid background in the field, so that they may continue to develop their professional career or further study by taking more advanced courses or engaging in academic research. Therefore, mastering of the main Computer Networking principles and acquiring of practical skills are both important. The specific course objectives are:

1) Understand the major problems encountered at the data-link layer (framing, error detection, reliability, multiple access) and the respective solutions.
2) Understand the principles of packet, circuit, and virtual circuit switching methods, and their intended uses.
3) Understand the design issues in a layer-two switched network through the example of switched Ethernet.
4) Understand the internetworking problem and the motivation for a layer-three switching solution.
5) Understand IP forwarding and routing problems, and the respective algorithms and protocols.
6) Understand end-to-end issues, arguments, and protocols through the example of TCP.
7) Understand some application layer issues, such as data formatting, data and network security, and name services.
8) Understand inter-layer issues, such as data encapsulation, address translation, fragmentation, and error detection.
9) Acquire basic socket programming skills.
10) Acquire practical experience in setting up an IP network and using of network diagnosis tools.

In the next section I first describe an integrated view of teaching and learning for the foundational course on Computer Networking. Sections III and IV discuss in details how this

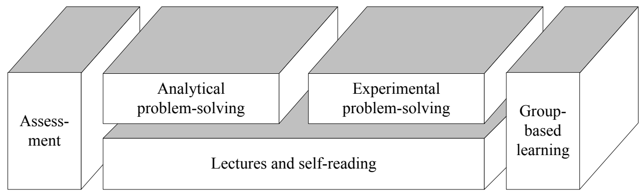

Fig. 1. An integrated view of five major teaching and learning elements.

integrated view affects the lecturing and problem-solving aspects, respectively. Section IV also discusses group-based learning and assessment which have been identified as an important component under the integrated view. Section 5 presents the learning outcomes resulted from applying the approaches discussed in sections III and IV for small and large classes. Finally, section 6 concludes this paper with a future work on adopting the experiential learning approach to this course.

II. An Integrated View of Teaching and Learning

To address the teaching and learning difficulties presented in section I, a promising approach, in my view, is to take an integrated view on all the elements in the entire teaching and learning process. That is, it is inadequate to simply strengthen one element, such as providing more hands-on experience, while neglecting other important elements. To be more specific, there are three types of separations that prevent one from viewing the teaching and learning in a more integrated manner—separation between various teaching and learning activities, separation between students in the class, and separation between teaching and assessments. Pictorially, Fig. 1 depicts an integrated view of the five major elements: lectures and self-reading, analytical problem-solving, experimental problem solving, group-based learning, and assessment.

A. Teaching and learning activity integration

The middle part between assessment and group-based learning in Fig. 1 refers to the teaching and learning activity integration. Although the teaching and learning activities are quite standard, the purpose of each activity and their inter-relationship must be carefully thought out and communicated to students.

The main foundation (the lower layer) is still based on teacher's lectures and students' self-reading. Owing to the limited lecture time, the lectures must concentrate on big issues and motivations, and provide more than a *linear* treatment of the subject material in order to facilitate students' self-reading. As will be elaborated more in section III, there are many do's and don'ts in an effective lecture.

The second layer is the problem-solving layer. This layer intends to deepen the learning in both the conceptual understanding and practical skills. One aspect of deepening is to delve into important details that cannot be treated adequately during lecture hours, such as important properties of CRC and the receiver-side ambiguity problem in sliding window protocols. Another aspect is to deepen students' impression on the topics by letting them visualize packets and how protocols work for themselves. Thus, both analytical and experimental problem-solving activities can enhance students' understanding of the networking concepts. By participating in the experimental problem-solving activities, students also acquire practical skills of building up a computer network and diagnosing network problems.

Across these two layers are group-based learning and assessments, which are discussed next.

B. Peer integration

The best teaching assistants that a teacher has are in fact students themselves. Therefore, one of the teacher's important responsibility is to facilitate student-student interaction, in addition to teacher-student interaction. A simple approach is to form study groups at the beginning of the course. Each group is encouraged to attend lectures and tutorials together, and to work on assignments and class projects together.

Another motivation for peer integration is that we as teachers are often unaware of students' learning difficulties, mostly because we are so familiar with the subject materials already. We have therefore forgotten the difficulties experienced during the learning process. Students, on the other hand, are more capable of understanding their peers' learning problems, and articulating them, to say the least. Moreover, other equally important benefits of the peer integration include mutual encouragement, support, and stimulation.

One of the study groups in my last year's course even came up a very innovative way of "studying together." They divide-and-conquered the course content by having each group member concentrate on some part of the syllabus. By having "experts" for all topics in the syllabus, the whole group was able to understand much more than studying alone. More details on study groups can be found in section IV. In fact, peer integration, or better known as peer learning or collaborative learning in the education sector, has been studied and applied, for example, to introductory computer programming courses [10].

C. Teaching and assessment integration

Teaching and assessment are the two main components in any course. Owing to the common understanding of the two components, teach-and-then-assess becomes the usual practice. That is, a number of topics is first taught (and assignments given), and then followed by quizzes, mid-term tests, and examinations on the covered material. I therefore call this a *teach-then-assess cycle*, and there are usually several such cycles throughout a course.

A strict segregation view of the two components, in my opinion, does not represent the best view to teaching and learning. First of all, the purpose of assessment is not only to discriminate students' understanding on the course materials, so that different grades can be assigned. More importantly, assessment should be used by a teacher to unearth learning difficulties, so that they can be addressed promptly by perhaps giving more examples or assignments. Therefore, assessments must be imbedded in the teaching and learning activities throughout the course, in addition to those formal assessments. One way to accomplishing this type of integration is to interact actively with students during lectures, and another is to work on the written assignments together with students (more details can be found in the next section). These teacher-student interactions allow a teacher to assess students' progress and discover learning problems immediately.

III. Lecturing: an Art of Exposition

Sometimes we have been so used to giving lectures that we do not think deeply about the purposes of lecturing and its relation with other activities. Compared with other teaching and learning activities, lecture is the *only* time where a teacher meets with *all* students, and this time is not very long (2 hours per week in my case). Because of the limited time for teaching as well as for learning during lectures, I suggest that the main objectives of lecturing should be:

1) Define and explain the problems at hand clearly, including why some obvious solutions do not work.
2) Explain the networking concepts and protocols in a more intuitive manner.
3) Register a clear framework in students' minds to relate different networking components.
4) Attempt to make the networking topics and problems relevant to students.
5) Communicate interests and enthusiasm in the subject matter.
6) Assess students' learning outcomes by actively interacting with them.

Of course, anything that prevents us from achieving the above can be classified as don'ts.

A. Defining the problems

According to my own experience, I cannot over-emphasize the importance of taking time to define problems. A major factor responsible for not fully understanding networking protocols, which are essentially solutions to some networking problems, is an incomplete understanding of the problems addressed. A prime example is perhaps the internetworking problem. Why do we need an internetworking solution at layer three? In order to understand this problem, one needs to explain the problems connected to a layer-two solution, i.e., scalability in terms of network size, problems of bridging various data-link networks, and incompatibility between MAC addresses. Therefore, understanding the internetworking problem thoroughly requires an understanding of the layer-two issues.

B. Explaining them intuitively

To lecture effectively, one has to explain problems and solutions intuitively. An effective approach is to teach the unknowns based on the knowns. For example, to explain the multiple access problem, I often ask my students to think about the scenario where more than one student in the class ask questions at the same time. These "messages" are not going to be received correctly by all, and they are "collided". When explaining the CSMA/CD as a solution to the MAC problem, the same analogy can be used—each student can immediately ask his questions only if he "senses" that the channel is not used by anyone else. Of course, two or students may speak at the same time, resulting in collisions again. This kind of intuitive explanations can help students comprehend the nature of the problem and the essence of its solution immediately.

C. Making them relevant

Students are usually more motivated to learn if they can see the relevancy of the subject to their daily lives. There are several ways to accomplishing this during lectures. If possible, at the beginning of each class, we ask students to name one important news related to Computer Networking during the week. One recent example is the root DNS servers' out-of-service during a denial-of-service attack, which serves an excellent example to explain the importance of DNS caching to the Internet stability. Another example took place on a weekday in Hong Kong when all ongoing K12 classes were dismissed due to a heavy rainstorm warning. The mobile phone network (and some fixed phone networks as well) was immediately jammed by parents' calls. This event serves as a real-life example to reveal the scalability limitation of circuit-switched networks, and at the same time demonstrate the versatility of packet switching (all the email facilities were still working at that time).

D. Registering a framework

Whether a top-down or bottom-up approach is used, many students begin to lose the entire picture as we move across the layers. In addition to understanding the concepts relating to a new layer, students also need to know how this layer relates to the layers previously discussed. Providing a simple framework would greatly help students understand the intricate relationship between different networking technologies. I have personally found the framework used in [4] very helpful, which starts first with directly connected hosts, and then move to indirectly connected hosts. In the latter, there are two cases to consider: connecting homogeneous data-link networks and connecting heterogeneous data-link networks. Finally, it moves to the end-to-end layers and issues. I re-iterate this framework again and again, especially when moved to a new chapter. Gradually, students are able to remember this framework without going through the textbook and notes.

E. Making them interesting

One important objective to achieve in any course is to instill interests in students. At the end of the course, we hope to see that most of students, if not all, are more interested in Computer Networking (pursuing it as their careers is another matter). Interests and enthusiasm, like the knowledge, need to be communicated. Therefore, while teaching the subject content, the teacher unknowingly also communicates his interests in the subject. Sometimes, some object lessons would be helpful for making the subject interesting. I have been using LEGO blocks (and other visible objects) to illustrate the concept of network layering and data encapsulation. It was very effective—everybody (around 170 students) watched me as I manipulated the LEGO blocks, and listened intently to my explanation of the concept.

F. Interacting with the learners

Interacting with students during lectures increases students' participation and attentiveness, because they know that they are part of the lecture. Teacher and students interact mainly through questions. Questions initiated from students should be turned around and addressed to the whole class, because most likely other students would have the same question. On the other hand, teacher should always prepare good questions to ask students. Good questions sometimes also challenges a already-know learning attitude. For example, after presenting the IP fragmentation and path MTU discovery approaches to handling heterogeneous MTU values, I asked my students which approach was taken by IP. As predicted, those, who are familiar with the topic, would pick fragmentation. To their amazement, both are actually used by IP: fragmentation by IPv4 and MTU discovery by IPv6.

IV. Problem-solving: a deepening exercise

Lectures and self-reading provide a general understanding to the networking principles, while problem-solving activities deepen the understanding by inviting students to apply the principles learned to tackle specific problems. Therefore, problem-solving learning takes students to another level where the general networking principles can be further explained and expounded. Very often, after going through the problem-solving exercise, many students realize their understanding of the networking concepts is not thorough and in-depth enough.

A. Analytical problem-solving

The first type of problems solving is analytical in the sense that the problems can be analyzed based on the principles presented in lectures and textbook. These problems include *simple* mathematical derivations of the properties of algorithms and protocols, protocol design issues, network configuration problems, etc. The analytical problems are given to students as assignments and mid-term tests. Since the assignments are not meant for testing students' general understanding of the networking principles covered in lectures and textbook, they are not straightforward problems. One class of problems intends to give a more in-depth treatment to some networking concepts and principles, e.g.,

1) Understand the key properties of CRC, e.g., conditions for detecting single-bit errors, double-bit errors, odd number of errors, and burst errors.
2) Understand the differences between layer-two broadcast and layer-three broadcast.
3) Understand the differences between a layer-three's view of the network and a layer-two's.
4) Understand hop-by-hop IP packet forwarding through a proxy ARP example.
5) Understand how IP routing discovers a new path to the Internet when the existing path is not available with the help of ICMP.

Besides, there are other analytical problems that even introduce new problems and issues to students, e.g.,

1) Understand the ambiguity problem at the receiver in a sliding window protocol when the sequence number space is not large enough.
2) Understand why Request-to-Send and Clear-to-Send messages are not sufficient to prevent frame collisions in a wireless LAN.
3) Understand why variable-length subnet masks may cause forwarding problem if the routing protocol does not support subnet masks.
4) Understand the impact of the duplicate IP address problem.
5) Understand why an increase in the initial TCP window size may not help increase the connection's throughput.

Since the assignments serve to deepen students' understanding, they are not left as home-work. Rather, assignments are meant to be tutorial-work; that is, I work on the assignments with students during tutorials. In my academic system, a one-hour tutorial is scheduled every week, which is similar to recitations in the States, and it is conducted in small class size (20-30 students). Therefore, the tutorial sessions are usually very interactive and lively. Students are invited to show and explain their answers. For more difficult problems, I would start by giving some leads and then invite participation from students.

There are many benefits in this *problem-based learning*. First, this is the best way to understand what students do not understand and misunderstand, when they give a wrong answer or give a correct answer based on an incorrect understanding. After each tutorial, it usually dawns on the students that there are so much depth into many networking issues and problems. Second, *process modelling* takes place throughout the teacher-student interaction during tutorials, and students observe first-handedly how I approach a problem, starting from the beginning and then solving the whole problem in logical steps. Honestly, this part is much more interesting than delivering the lectures.

B. Experimental problem-solving

Another type of problem-solving activities consists of practical works, such as socket programming (as assignments), laboratories on application protocols (HTTP, SMTP, and DNS), and a class project. The first two are conducted in the beginning of the course, because they do not require much understanding of the underlying networking principles.

Although the socket programming assignments and laboratory sessions are useful, they do not allow students to interact with the "kernel" of a computer network. Therefore, I initiated a class project three years ago to let students experiment with a "personal" computer network throughout the course—each group is expected to *work on* their own network in parallel to the course' progression. However, the amount of resources required for supporting 170 students is obviously very demanding, in terms of both equipments (PCs, hubs, NICs, etc) and space (space is extremely costly in Hong Kong). Finally, I have come to the conclusion that the only solution to resolving the resource constraint is to "bring the networking laboratory to where students are." That is, each group is responsible for finding their own machines, OSes, networking equipment, and space. Although it sounds like a very difficult task, all groups in the past two years were able to find all the resources and finish the project.

The first part of the group project is to configure a private network (situated in students' homes) connected to the Internet and the systems are usually based on Linux or BSD. A minimum set up, as depicted in Fig. 2, includes two subnets—a shared-medium LAN (Ethernet) and a PPP network, and a Linux system is configured as a router that also performs NAT. Advanced students are able to configure more complex networks and services, such as that in Fig. 3. Starting from this year I have asked students to configure an IPSec tunnel between two such private networks. In the process of setting up the network, students pick up practical skills of diagnosing network problems and finding solutions to them.

The second part is to reinforce the principles learned in lectures by performing various experiments on the network. In other words, the hands-on experiments draw students closer to the Computer Networking principles, because they can now visualize how protocols work for themselves. For example, they will find out from this set of experiments

1) The effect of frame collisions on an Ethernet network's performance,
2) The relationship between MTU and network's throughput,
3) The effect of packet fragmentation on the network performance,
4) TCP's retransmission behavior,
5) The states of a TCP connection not affected by a nonpersistent network outage, and
6) The difference between regular and passive-mode FTP.

V. Group-Based Learning and Assessment

As mentioned in section II, an important teaching and learning element is group-based learning. Starting from last year, I asked students to form study groups at the beginning of the course. All members in a study group attended the same tutorial session, and they were also expected to study together outside classes. Each group had a group leader to plan for study activities outside classes. Each study group also undertook the class project together.

Whether the group members would benefit from the group-based learning mode depends on many factors, such as trust among members, the group leader's leadership, the complementary strength of the members, etc. To put it in another way, putting students in a small group does not automatically make them study the subject cooperatively and complementary. However, knowing the importance of the study groups, students are usefully very careful in selecting their group members. Therefore, most of the study groups turned out quite effective.

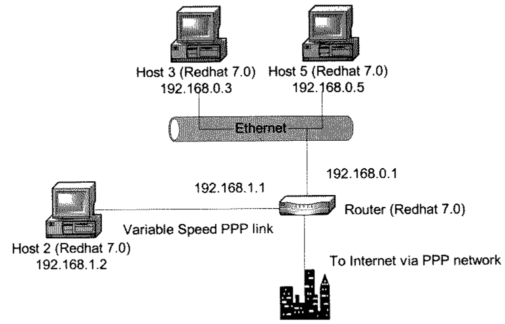

Fig. 2. A minimum network setup for the group-based class project.

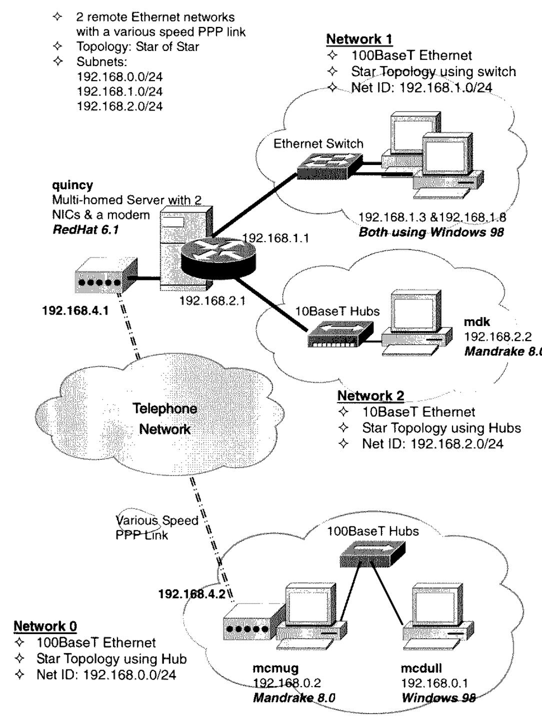

Fig. 3. A more elaborate network setup for the group-based class project.

Moreover, an important factor that could affect the group-based learning outcome is the assessment scheme. For example, if a noncontributing member gets the same grade as other members, members may be discouraged from contributing their very best to the group activities. In my case, students are assessed individually in all components except for the group project. In order to maximize the cooperativeness in the group project, I assess each member's contribution to the project in two steps. The first one is to give an overall grade based on the project report and presentation. In the second step, each member is asked to evaluate the contributions from other group members. Based on the students' evaluations, high and low contributors can be identified, and their individual grades can then be adjusted upward and downward accordingly.

VI. The Learning Outcomes

This section presents students' feedbacks received from the undergraduate course on Computer Networking in the last two years. The class in 2001 was a large class, consisting of 168 students, and the lectures were conducted in a large lecture theater. But each tutorial class was still limited to 20 to 30 students. The class in 2002, on the other hand, was a small class, consisting of 54 students. The learning outcomes were evaluated mainly from three aspects: teaching activities, assessments, and group-based learning. Moreover, I assessed whether students' interests in Computer Networking have been increased after taking the class. The number of responses received were 133 (79%) and 52 (96%) for 2001 and 2002, respectively.

A. Effectiveness of teaching activities

This section evaluates the effectiveness of the lectures, tutorial sessions where assignments were discussed, and self-reading based on the textbook. Therefore, the following three statements were given, and there were five possible responses to each statement: strongly agree, agree, neutral, disagree, and strongly agree.

- The lectures helped me understand the subject materials.
- The tutorial sessions helped me understand the subject materials.
- The textbook helped me understand the subject materials.

The most outstanding result is perhaps the feedback on the effectiveness of tutorials. Almost 30% of the responses from the small class *strongly* agreed that the tutorial sessions helped them understand the subject materials, and the agreed and strongly agreed combined even exceeded 80%. The written comments also indicated that students found the assignments and the discussions on the assignments very useful in clarifying the networking concepts. Similar results were also obtained for the large class in 2001, in spite of less favorable responses as compared with that of the small class—10% on the strongly agreed and 65% on the strongly agreed and agreed combined.

On the other hand, it is not too surprising to find that the lectures were not perceived as effective as the tutorials. In the small class, 75% of responses (strongly) agreed that the lectures helped them understand the subject materials, while this percentage was 47% for the large class. There are many factors contributing to the difference in the effectiveness between lectures and tutorials. The most obvious one is to do with the class size; the small-group tutorial arrangement facilities more teacher-student interactions. Students' learning difficulties can therefore be promptly and specifically addressed in the small-group setting, whereas it is quite impossible to address them during lectures.

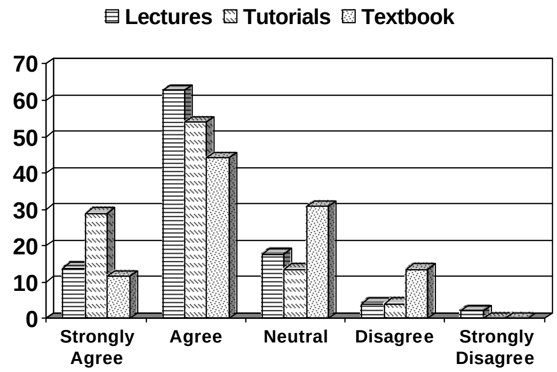

Fig. 4. Effectiveness of teaching activities for a small class.

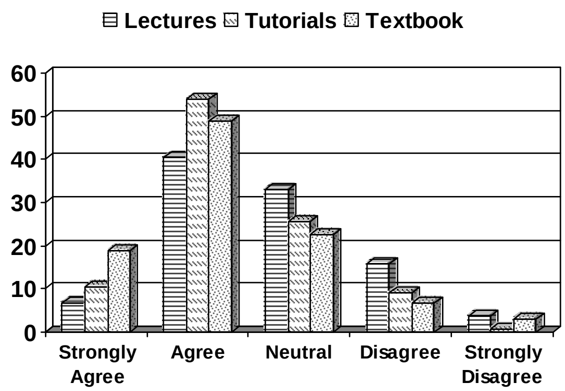

Fig. 5. Effectiveness of teaching activities for a large class.

More importantly, the learning outcomes perceived by students are actually a direct result of the lecture and problem-solving integration. As recalled, lectures serve as a platform to introduce networking principles for the first time, and it is bound to leave many questions unanswered. Analytical problem-solving, on the other hand, deepens students' understanding received from lectures and self-reading. As a result, it is predictable that students generally find that tutorials, where analytical problem-solving takes place, are more effective than lectures. But in fact lectures and tutorials cannot and should not be viewed and evaluated separately from each other. A more accurate interpretation of the data is: the lectures and tutorials combined are responsible for a very effective learning outcome. Moreover, the data shows that the tutorials are very successful in terms of deepening students' understanding of the subject.

Both classes used the same textbook by Peterson and Davie [4]. However, it is interesting to observe the discrepancy between the two classes. Nearly 70% from the large class (strongly) agreed that the textbook helped them understand the subject materials, but this percentage dropped to 56% for the small class.

B. Effectiveness of assessments

This section evaluates the effectiveness of various assessment components, including assignments (mainly analytical problem-solving), class project (experimental problem-solving), and mid-term tests. Therefore, the following four statements were presented to the students.

- The assignments helped me understand the subject materials.
- The class project helped me acquire practical skills in Computer Networking.
- The class project helped me understand the subject materials.
- The tests helped me find out how much I have understood (or have not understood).

Since assignments and tutorials are closely related, the first statement regarding the effectiveness of assignments received highly favorable feedbacks. Almost 90% and 75% of responses (strongly) agreed that the assignments helped them understand the subject materials for the small and large class, respectively. Note that these results are quite consistent with that for tutorials presented in the last section.

The next two statements assessed the effectiveness of the class project in terms of acquiring practical skills and enhancing understanding of the subject materials. The responses to both statements were quite similar in that around 80% of the responses from the small class (strongly) agreed that the class project was helpful in the two aspects with slightly more favorable responses on the practical skill. In the large class, the aspect of practical skills received a more favorable response than the aspect of enhanced understanding. Around 65% of the responses (strongly) agreed that the class project helped them acquire practical skills in Computer Networking, while 51% of responses (strongly) agreed that the class project helped them understand the subject materials.

The responses regarding the class project were quite understandable. Although understanding the networking principle was actually a major objective of the project, students still perceived that acquiring the practical skills throughout the process was more important. Based on students' written comments, many of them felt that acquiring the skills of setting up networks from ground up would help them lend on jobs upon graduation. Another reason is that most of them had never set up a computer network before; therefore the practical work itself seemed to be more rewarding and significant. In any case, the class project was quite effective in both aspects.

One mid-term test was given in the large class and two were given in the small class. The tests were used to serve as feedbacks to the students, so that they can know their own progress in the course. In both classes, more than 60% of the responses (strongly) agreed that the tests helped them find out how much they have understood (or have not understood). Similar to the assignment problems, the mid-term test questions were not straightforward, and students generally did not do very well.

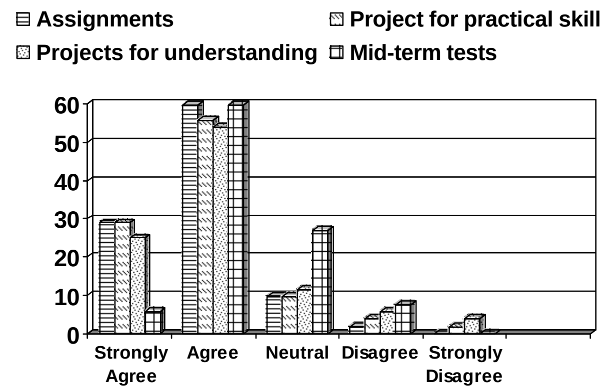

Fig. 6. Effectiveness of assessment components for a small class.

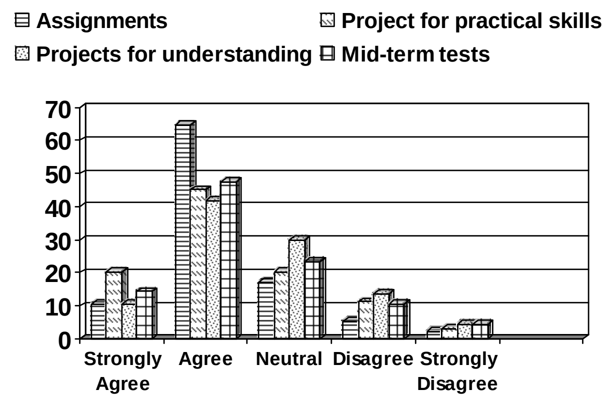

Fig. 7. Effectiveness of assessment components for a large class.

C. Effectiveness of group-based learning

This section attempts to assess whether the group-based learning is effective. The following two statements were presented to students.

- My group members were helping each other to understand the subject materials.
- My group members were dedicated in working together on the class project.

Since the group-based learning was first introduced in 2002, there is only one set of results. The responses to both questions are similar, and around 70% of the responses (strongly) agreed to both statements. The written feedbacks also reflected that the groups were generally effective. Some groups reported that they had encountered many problems in the class project, but each member determined to overcome them by consulting other students and by reading relevant reference books.

D. Students' interests in the subject

The most important students' feedback, in my opinion, is whether students are more interested in Computer Networking after taking the course. Therefore, students in the small class were asked whether they are interested in Computer Networking at *two separate times*: during the first class (indicated by pre-teaching in Fig. 9) and during the last class (indicated by post-teaching in Fig. 9). The reason for doing this is to accurately assess the impact of this course on their interests

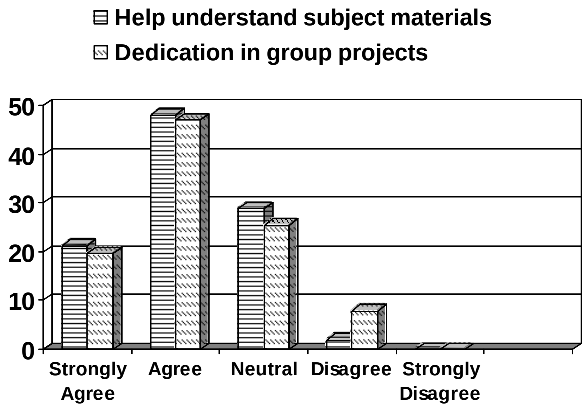

Fig. 8. Effectiveness of study groups for a small class.

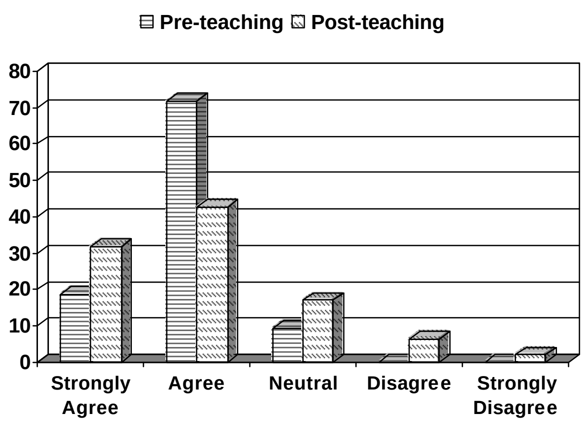

Fig. 9. Students' interests in Computer Networking before and after taking the course for a small class.

in the subject. Previously in the large class, students were asked whether they were *more* interested in the subject after taking the course, but the responses did not clearly indicate the absolute levels of interests.

The feedback results are actually very revealing. During the first class, around 90% of the students (strongly) agreed that they are interested in Computer Networking. This result is also predictable, because of the importance of the Internet. Moreover, some students wrongly believe that Computer Networking is equivalent to WWW, or other network applications. Therefore, after learning the actual Computer Networking principles, many students' perceptions about the subject changed. Some became aware that Computer Networking was not for them, because the content may be a little too difficult to comprehend. However, others found the subject much more interesting than before, because both the extent and depth of the subject far exceeded what they knew before. As a result, we see a higher variance in the post-teaching responses. The percentage of strongly-agreed cases jumped from 19% at the beginning of the course to 32% at the end of the course. On the other hand, the percentage of agreed cases decreased significantly from 72% at the beginning of the course to 43% at the end of the course. The other cases also increased at the end of the course.

VII. CONCLUSIONS AND FUTURE WORK

Doing research and teaching on Computer Networking are two quite different tasks. Although it is true that research underpins teaching, teaching requires an extra set of skills that may not be needed in research. For example, one major difficulty in teaching the undergraduates is to "stoop down" to their level, so that they can comprehend the intricate principles of Computer Networking from what they know. Explaining a complex concept simply is indeed not simple.

In this paper I have presented an integrated view for teaching and learning a foundational Computer Networking course for undergraduate students. Under this view, the objectives of the take-it-for-granted lectures, assignments, projects, etc, and their intertwining relationship need to be thoroughly thought out and planned. Moreover, this integrated view discourages us from paying attention to just a single teaching methodology, such as hands-on work. Lectures, problem-solving, group-based learning, and assessments are all very important to providing a quality and enjoyable teaching and learning experience.

One possible improvement is to further tighten the integration of the class project with the instructional part. Currently, I start giving out the project approximately a month after the beginning of the course. In the next year, I am planning to give out the class project in the first class, and students are expected to "experience" a certain networking topic on their personal computer network as soon as the topic is covered in lectures. This approach is somewhat similar to the experiential learning approach discussed in the education sector. In contrast to cognitive learning, experiential learning attempts to provide incentives to learn, and in my case the incentive is to build a personal computer network after completing the course.

ACKNOWLEDGMENT

This work was partially supported by a Teaching Development Grant from The Hong Kong Polytechnic University. The network setup in Fig. 2 is part of the project report by Swire Lo, Ivan Lam, Antony Wong, and Wesker Chan. The network setup in Fig. 3 is part of the project report by Quincy Cheng and Polly Cheng. I also thank all my students for providing their feedbacks to the various aspects of the course. It is impossible to improve the course over the years without their faithful participation.

REFERENCES

[1] D. Bertsekas and R. Gallager, *Data Networks*, Second Edition, Prentice-Hall, 1992.
[2] A. Tanenbaum, *Computer Networks*, Fourth Edition, Prentice Hall PTR, 2003.
[3] S. Keshav, *An Engineering Approach to Computer Networking*, Addison Wesley Longman, 1997.
[4] L. Peterson and B. Davie, *Computer Networks: A Systems Approach*, Second Edition, Morgan Kaufmann, 2000.
[5] A. Leon-Garcia and I. Widjaja, *Communication Networks: Fundamental Concepts and Key Architectures*, McGraw-Hill, 2000.

[6] B. Forouzan, *Data Communications and Networking*, Second Edition, McGraw-Hill, 2000.

[7] J. Kurose and K. Ross, *Computer Networking: A Top-Down Approach Featuring the Internet*, Addison Wesley Longman, 2001.

[8] *ACM Workshop on Networking Education*, Aug. 2002. Available from http://www-net.cs.umass.edu/sigcomm/education/workshop1.html.

[9] *Homepage of Computer Communications Networks (COMP312)*, Department of Computing, The Hong Kong Polytechnic Univeristy, http://www2.comp.polyu.edu.hk/∼comp312/CCN/CCN.html.

[10] C. Wills, "Application of peer learning to the introductory computer science curriculum," available in http://www.cs.wpi.edu/∼peercs.