A cost-benefit analysis of remote collaborative tutorial teaching

by
Stephen W. Draper
Department of Psychology
University of Glasgow
Glasgow G12 8QQ U.K.
email: steve@psy.gla.ac.uk
WWW URL: http://www.psy.gla.ac.uk/~steve

and
Sandra P. Foubister
Department of Computing
Napier University

Abstract

Near the end of the MANTCHI project, an interview study of 10 participating university teachers was carried out. This was a teacher-centered evaluation to complement the student-centered evaluation of the use of the materials. This chapter reports on a cost-benefit analysis based on this of the approach to collaborative teaching developed in the project. The analysis suggests that there are gains in authoring effort, a gain in quality because more frequent delivery means more revisions based on experience, a gain in curriculum quality, a balance in delivery time costs or else a small extra cost in return for staff development gains. Besides the analysis of this case, the study identifies classes of cost and benefit for use in future, more detailed, studies.

Introduction

A strong tendency in the evaluation of learning innovations and learning technology is towards learner-centred evaluation. There are at least two important reasons for this: if the aim is to establish claims about educational effectiveness, this is most directly done by measuring learning outcomes (i.e. studying the performance of learners) rather than only the opinions of teachers or other experts; and if the aim is to discover the often unexpected bottlenecks determining performance, which often vary widely between different cases, then again observing real learners in their actual situation is crucial. Our own work on Integrative Evaluation (Draper et al., 1996) has been in this direction. However there are other important issues which cannot be addressed in this way, but instead require some form of teacher-centred study. That is particularly the case with innovations whose main benefit is likely to be in saving costs of some kind.

As with studying learning benefits, the issues are likely to be complex and the literature is much less developed (but see, for example, Doughty, 1979; Doughty, 1996a, 1996b; Reeves, 1990; and Hawkridge, 1993). Identifying the kinds of costs and benefits for teachers, some of them unanticipated, is at least as important as taking measurements of those kinds that were expected. This chapter reports an attempt at this, based on interviewing 10 teachers involved in an innovative project on remote collaborative tutorial teaching. It addressed the general question of whether this approach to collaborative teaching was of net benefit to them, in terms of the cost headings which they would themselves accept as personally relevant (mainly time spent on various component activities).

The MANTCHI project

The MANTCHI project — Metropolitan Area Network Tutoring in Computer-Human Interaction — (MANTCHI; 1998), explored remote collaborative task-based tutorial teaching over the Internet. The project involved four universities in central Scotland for about a year and a half, and worked in the subject area of Human Computer Interaction (HCI). The material was delivered in existing university courses, where students were working for a qualification.

"Tutorial" was broadly defined to mean anything other than primary exposition (such as lectures). A unit of this material is called an "ATOM" (Autonomous Teaching Object in MANTCHI), and is typically designed as one week's work on a module for a student i.e. eight to ten hours, including contact time. Ten ATOMs were written (listed in the table below), of which eight were delivered and evaluated during the project. Typically this material was an exercise created at one site and available to all sites as web pages. Sometimes it was designed to use a remote expert (usually the author) as part of each delivery, who might give a video conference tutorial or give feedback on student work submitted and returned over the Internet. In some cases students on different courses, as well as the teachers, interacted. More details of each are given below, and through the project website (MANTCHI; 1998).

Responsibility for the courses and assessment remained ultimately with the local deliverer. In this chapter, these deliverers and also the authors and remote experts are referred to by their role as "teachers" (though their job titles might be lecturer, professor, teaching assistant etc.) in contrast to the "learners", who in this study were all university students. There were three basic roles: author (also sometimes active as a remote expert), local deliverer, and learner.

A key emergent feature of the project was its organisation around true reciprocal collaborative teaching. All of the four sites have authored material, and all four have delivered material authored at other sites. Although originally planned simply as a fair way of dividing up the work, it has kept all project members crucially aware not just of the problems of authoring, but of what it is like to be delivering to one's own students (in real, for-credit, courses) material that others have authored: a true users' perspective. This may be a unique feature. MANTCHI has in effect built users of teaching material further into the design team by having each authoring site deliver material "not created here". It is also a system of collaborative teaching based on barter. This goes a long way to avoiding organisational issues of paying for services. However the details may not always be straightforward, and the future will depend upon whether costs and benefits balance favourably for everyone: the core reason for the present study.

Evaluation of learning in MANTCHI

There was extensive evaluation work within MANTCHI (reported in a companion chapter "Evaluating remote collaborative tutorial teaching in MANTCHI") of the educational effectiveness of this material in actual classroom use, based on the method of Integrative Evaluation (Draper et al. 1996). The evaluation studies provided formative information for improving the overall delivery of the material, and some summative evidence on learning outcomes and quality that suggests that the ATOM-ised materials are of at least as high a quality as other material, although by no means always preferred by students. One telling piece of favourable evidence was that on one course many students preferred the ATOM authored by their teacher to other parts of the course also developed by him. These studies suggested however that the main gains might be in improved curriculum content and in staff development (expanding the range of topics a teacher is confident of delivering), something that learner-centred evaluation could not directly demonstrate, but that teacher-centred evaluation could.

Another issue those learning evaluation studies could not address was whether teachers could be expected to use the ATOM materials beyond the end of the project. The educational benefits seem to be sufficient to warrant this, but not enough to provide an overwhelming reason by themselves regardless of costs and other issues. This would depend upon whether teachers found them of overall benefit. Another kind of study was needed to investigate these issues.

The purpose of the study

Whether teachers decide to adopt an ATOM, at least outside the project and its special motivations for the participants, will depend upon their perception of the costs and benefits such adoption will afford. This is crucial to the longer term success of the work, but requires a quite different kind of evaluation: centred on teachers not learners, and requiring an investigation into what they count as benefits, and the development of methods for measuring costs as experienced by teachers. As we began to recognise the need for this, we launched a small interview study of the teachers involved, asking how much work (time) it had taken to author ATOMs, to deliver them, whether they thought they would continue to use the ATOMs beyond the project, and why. This yielded information on estimated quantity of work, on the kinds of cost and benefit perceived, and also on where the ultimate significance of this ATOM approach (and hence of the MANTCHI project) might lie.

Method

Accordingly a short study was undertaken in the final month of the project, consisting of retrospective interviews with the participating teachers (authors, remote experts, local deliverers). Each interview lasted about an hour. The agenda for the interviews was to ask how much time and effort had gone into activities related to creating and delivering the ATOMs, whether the teachers expected to use the ATOMs beyond the end of the project, and to identify what they thought the costs and benefits were. Actual times were obtained (and are shown in the table below), although the accuracy of these estimates is probably quite poor due to the retrospective nature of the measure: this discussion took place well after the actual delivery of the ATOMs, in some cases over a year later. However at least as important a result was the insight afforded into how these teachers think about the pros and cons, the costs and benefits, of these materials i.e. identifying the kinds, rather than the quantity, of costs and benefits found to be relevant by the participants. Also noted were comments about how to use ATOMs, the relative advantages of each type, the best way to use a remote expert, etc., although not all are discussed here.

The variety of ATOMs

The ATOMs varied considerably in their nature, which of course affects the costs associated with each. This section briefly describes those features likely to affect time costs for teachers, and hence the values given in the table. (The ATOMs themselves are available via the project website: MANTCHI, 1998.)

From the point of view of time costs, there were three types of ATOM. In the first group described below, a remote expert was actively involved in the delivery and this was a cost not incurred in other ATOMs. In the second group, students interacted between institutions, incurring extra coordination costs, and implying that deliverers at both institutions were involved simultaneously, but other than that no remote expert or author was involved in delivery. In the third group, there was no dependency during a delivery on people at a remote site (although there were dependencies on remote resources such as web sites). This grouping is only for the purpose of bringing out types of cost: in a classification in terms of pedagogical method, for instance, the groupings would be different e.g. the CSCLN and remote presentation ATOMs (described below) would be grouped together as being based on teachback.

The CBL (computer based learning) evaluation ATOM concerned teaching students how to perform an educational evaluation of a piece of CBL. The students had to design and execute such an evaluation, and write a report that was assessed by the local deliverers. The interaction with the remote expert was by two video conferences, plus some discussion over the Internet (email and a web based discussion tools). The UAN (User Action Notation), ERMIA (Entity-Relationship Models of Information Artefacts), and Statecharts ATOMs each concerned a different notation for specifying the design of a user interface. These ATOMs each revolved around an exercise where students, working in groups, had to specify an interface in that notation. These solutions were submitted over the Internet, marked by the remote expert, and returned again over the Internet.

The CSCW (Computer Supported Cooperative Work) ATOM was a group exercise, which involves students working in teams assembled across two institutions to evaluate an asynchronous collaboration tool (BSCW is suggested). They first worked together using the tool to produce a short report related to their courses, and then produced evaluation reports on the effectiveness of the tool. The formative evaluation ATOM took advantage of the fact that students at one university are required to produce a multi-media project as part of their course. In this ATOM, students from a second university were each assigned to one such project student, and performed an evaluation of that project, interacting with its author remotely through email, NetMeeting, and if possible video conferencing. The ATOM on remote student presentations was a seminar based course, where students took turns to make a presentation based on their reading to other students, and were assessed on this. In this ATOM, these presentations were both to the rest of their class and, via video conference, to another class at another university.

The CSCLN (Computer Supported Cooperative Lecture Notes) ATOM required students to create lecture notes on the web that accumulate into a shared resource for the whole class, with one team assigned to each lecture. There was no role for a remote expert. The website evaluation ATOM involves the study and evaluation of three web sites on the basis of HCI and content. Students complete the exercise on their own, over the course of a week, and submit and discuss their evaluations via Hypernews. In the website design ATOM students work in groups to produce a web site; there is no remote collaboration. In this ATOM, as in fact in all the ATOMs in this group, the exercise could be reorganised to involve groups split across sites, as in the previous group.

Result table

Here is a table that allows comparison between the time estimates of the various elements in an ATOM's lifecycle. Since many of the ATOMs were elaborations on pre-existing tutorial exercises, the presence or absence of an existing version is also noted.

ATOM	**Pre-existed**	**Finding resources**	**Authoring**	**Preparation - *Author***	**Preparation - *Deliverer***	**Delivery**	**Marking**	**Revision**
CBL	YES	-	3 - 4 hrs	2 hrs	3 hrs	2 X 2hrs video-conference 1 hr photo-copying	1 hr per group	0
UAN	YES	-	4 hrs	0.5 hr	**Napier:** 2 days **Heriot-Watt:** 1 hr	0	0.5 hr per group + 0.5 hr feedback on Web	0.5 hr
ERMIA	YES	24 hrs		4 - 7 hrs	0	Normal tutorial time	1 hr per group	4 - 7 hrs
Statecharts	NO	16 hrs	4 hrs	-	**First delivery:** 2 days **Subsequently:** 0	3 hrs	20 mins/ group 0.5 hr printing + 1 hr put feedback on Web	4 hrs + 8 RA hrs
CSCW	YES	-	2 hrs	1 - 1.5 hrs	0	1.5 hrs	Normal marking time	0
Formative Eval	YES	-	1 hr	13 - 14 hrs		2.5 hrs	3.5 hr	0
Remote Present	YES	-	(No ATOM description)	**Napier:** 4.5 hrs + programmer time **Glasgow Caledonian:** 1 hr + programmer time		4 X 1hr	0	0
CSCLN	NO	-	3 - 4 hrs	-	4 hrs	0	0	4 hrs
Website Eval	YES	-	0.5 hr	0.5 hr	0	(ATOM not used)		0.5 hr
Website Design	NO	16 hrs	2 hrs	0	0	(ATOM not used)		8 hrs

N.B. in the case of the ERMIA ATOM, the interviewee gave a single time (24 hours) for finding resources and authoring combined.

Accuracy of times

The times are probably underestimates (one subject said at the end of the interview "I bet these are all underestimates"). One of the authors, having been interviewed as a subject, later found a rough time diary he had forgotten about. Comparing the times recorded in that diary with the estimates he gave in the interview 4 months later shows he had lowered his estimate by about 30% for both the ATOMs in which he was involved. The diary itself may be an underestimate by omitting a few occasions, and because it was at best filled in retrospectively at the end of each day.

A second problem is that of accuracy in the sense of comparability (not systematic underestimation) of the times given. Many respondents noted that it is very hard to estimate "time" in this context. The time it takes to do something, such as physically type in an ATOM description, may not have much relationship with elapsed time — from, say, the original outline of the ATOM to the final, usable, version. People also mentioned "thinking" time and "research" time, as in "Do I count an hour in the bath spent thinking about it? An hour at home? An hour in the office?". Similarly: "I regard that as a free by-product of research thinking!" Nevertheless, every person interviewed felt able to give rough estimates. These vary very widely, e.g. from 0.5 hour to 24 hours for writing an ATOM description, so it may be that the implicit definition of "time" was indeed different between respondents.

The research aspect of MANTCHI will have increased the time costs because of the need to monitor various details, both at the time and in retrospect. This "costing" exercise has added yet a further hour for each of the dozen or so lecturers involved.

Cost categories used in the table

The columns in the table represent the main categories of time cost that emerged from the interviews.

- The "Pre-existed" column records whether there was an existing exercise on which to base the ATOM.
- "Finding resources" refers to collecting material to which the students will be referred e.g. papers for the primary exposition.
- "Authoring" is writing the ATOM pages themselves.
- "Preparation — Author" is time the author or remote expert spent preparing for an interaction (e.g. a video conference).
- "Preparation — Deliverer" is similar work by the local deliverer.
- "Delivery" is contact time with students by local and remote teachers, and any other delivery costs (e.g. collecting student work).
- "Marking" is generating feedback on student work.
- "Revision" is rewriting the ATOM material in the light of earlier deliveries for re-use in a new delivery.

Authoring from scratch often requires little creativity or design time, as people only volunteer for ATOMs that they already know how to write, or perhaps already have written in some form. The "authoring" category therefore is mainly about writing, and not about designing or creative thinking. In general it might be useful to have a separate "design" or "creativity" category, but in this study such a category would have had small quantities of time attributed to it.

Authoring in general often has a significant element of iterative design to it, meaning that first draft authoring is often quite cheap, but the cost of revisions after a delivery or two should be allowed for, as these are really part of the process. Thus the authoring column in the table represents an estimate for a teacher considering joining in an ATOM exchange, but should not be compared directly with authoring times in other media (such as textbook writing) where some revisions would be part of the author's work. Conversely, the "revision" column combines (and so confounds) revision work that increases the quality with revision work simply to adapt the written material for a new occasion (e.g. updating times and places in handouts, modifying URL links). N.B. the times for the two ATOMs not yet delivered are estimates of the latter time for adaptation. N.B. the times for the two ATOMs not yet delivered are estimates of the latter: of specialisation time. Future work should distinguish these two types of revision. In general authoring might have four related categories: creativity or design, collecting resources to be used, actual writing, revision of the content in the light of previous uses of the material.

Time for marking is proportional to the number of solutions marked. The table gives the time per solution, which must be multiplied by the number of solutions in any attempt to predict times for future deliveries. This is one of the biggest issues in agreeing an exchange involving a remote expert: if class sizes are very different, marking loads will be different. There are several possible ways forward:

different group sizes might mean the same number of solutions to mark even with different class sizes; a remote expert might just mark a sample with exemplary feedback, leaving the rest of the marking and commenting to be done by the students and/or local deliverer.

A related point to note is the use of groupwork in most of the ATOMs, since each group only produces one solution to be marked: thus the fewer the groups, the less time needed for marking. However larger groups certainly make it harder for students to agree meeting times, and may be less effective in promoting individual learning. Thus there is probably a strong tradeoff between learning quality and teacher time costs here, which is even more obvious in cases where the student work never gets marked at all or the feedback is of low quality.

Kinds of cost

One of the central results of this study is the identification of categories of cost, which future studies might design special instruments for. Each of the column headings above are one such category. As noted, "revision" should be divided into two, and perhaps "authoring" should be split into creative design vs. writing, besides the related categories of collecting resources and revising to improve quality.

One issue that emerged from the interviews was that time is not a currency with a fixed value. Time spent or saved at a period when other pressures on time are high is more valuable than time at periods of low pressure. Thus being able to serve as a remote expert at a low pressure time in return for getting the services of a remote expert at a high pressure time could be a "profitable" trade; while the converse could be disadvantageous even if the durations involved seem to balance. ATOMs move time costs around the calendar, which may be good or bad for the participants.

Another issue is also important: how difficult an activity feels to the teacher (perhaps measurable in terms of their confidence about performing the activity). The fundamental advantage to the trade behind ATOMs is that a remote expert usually feels it is easy and indeed often interesting to field student questions on the topic or to comment on unusual solutions, while the local deliverer would feel anxious about it. The time spent in each case might be the same, but the subjective effort would be significantly different. (The probable educational quality is less easy to estimate without detailed study because students may sometimes benefit more from a teacher who does not provide the answer but thereby succeeds in getting them to do more work in seeking it, than an expert whose instant provision of an answer inhibits further student thought.)

A kind of cost not visible in this study is the groundwork of understanding what an ATOM is. For project members, this was done at a series of project meetings and perhaps in thinking about the ATOMs they authored, and is probably missing from all the time estimates. If new teachers were to use the ATOM materials, this learning of the background ideas might be a cost. On the other hand, this is comparable to the costs of any purchaser in gaining the information from which to decide whether to buy: a real cost, but often not considered. With ATOMs, a teacher would have to learn about ATOMs before making the decision whether to "buy in" at all, rather than while they were using them.

Finally, there were clearly costs in using some of the technology e.g. setting up video conferences, getting CSCW tools to work. As noted below, much of this can be written off as a cost of giving students personal experience of the technology which is appropriate in the subject area of HCI. This would not apply if the ATOM approach were transferred to another subject area, but retaining those same technical communication methods. However it is also true that such costs will probably rapidly reduce as both equipment and staff familiarity with the technology improves.

Kinds of benefit

A number of kinds of benefit are apparent.

1. Subjective effort (or confidence): donating time at teaching felt by a teacher to be easy for them in exchange for support at teaching felt to be difficult for them.
2. Prime time vs. time at low pressure periods for an individual: this exchange could be a cost or benefit.

3. Using an exercise without having to write it, in exchange for an exercise that had already been written for local use anyway. It may have to be written more carefully for use at multiple sites, so the gain is then not having to write an exercise in exchange for improving one already written.
4. The main value may be in a better curriculum. It is clear that even within the project, local deliverers only "took" exercises they felt would improve their courses (there were many more offers of ATOMs at the planning meetings than takers i.e. the topics developed were more demand-driven than supply-driven). Many deliverers definitely believed they improved their courses (although sometimes this may have been offset to some extent by the costs of moving topics around) by improving the selection of topics or the depth in which they were treated, although this study did not have another measure of curricular quality.
5. Some teachers expressed the idea that after having delivered an ATOM in someone else's speciality a few times, they would probably feel comfortable delivering it without a remote expert. This would mean that ATOMs were serving the function of staff development, rather than remaining permanently as an item of exchange.
6. In some ATOMs, the students were given a personal experience (rather than only an abstract concept) of the topic e.g. experience of CSCW by collaborating with a remote group via software. That is, the ATOMs additionally amounted to practical laboratory exercises i.e. a learning objective in themselves. This additional benefit comes from the synergy between the HCI content area being taught and the project's exploratory use of learning technology. Thus for instance in the CSCLN ATOM, collaborative lecture notes would probably benefit any course, but the students (and teachers) were also particularly keen on this as an occasion to practise web authoring for its own sake.
7. In the ATOMs involving learners at remote sites, this contact was often seen as a benefit for learners in itself: seeing other learners' solutions, and getting a sense of how a subject is taught elsewhere.

Roles for a remote expert

Remote experts were only used as part of the delivery in four ATOMs, forming the first of the three groups. Here their use is seen to have several functions:

- As staff development: "Next time I feel I could do the marking myself"
- As resource: "The ATOM fell at a time when I was very busy, so having X do the marking was very useful"
- To give the primary exposition (e.g. a lecture delivered by video link), though this was not the focus of MANTCHI.
- To give students access to the expert, as postgraduates might have at workshops and conferences, particularly to respond to student questions and suggested solutions to exercises. There was some tendency at undergraduate level, though, to perceive the use of a remote domain expert as an admission of inadequacy by the local deliverer: "Why can't you teach us this yourself?"

This latter issue in fact relates to which view prevails (on the part of learners, teachers, and indeed institutions) of what higher education should be, as described by Perry (1968). This view, which varied within this project, also varies generally and widely with the age of the learner, the institution, but perhaps above all with the discipline. Even at the age of 16, history teaching may be organised around weekly essays based on library reading with the teacher's role one of discussing student work, whereas even at Masters level, you may find chemistry being taught by lecture with the student's role being to reproduce what the teacher says without criticism. Perry's view was that a university's duty is to move students from the latter simplistic view and dependent role to the former. (Nowadays this might be redescribed in terms of acquiring learning or critical thinking skills.) Students however often resist this, and criticising their local teachers for not apparently "knowing" the material is consistent with what Perry regards as the primitive state of seeing teaching as telling, learning as listening, and not having to trouble to examine evidence and alternative views in order to formulate and defend a reasoned view of one's own.

The role of a remote expert will depend upon where on the Perry spectrum a course is pitched. It is however well suited to an approach where students are expected to be able to learn by reading and by attempting to do exercises, but will benefit from (expert) tutors' responses to their questions and examples not directly covered in the primary material. From the teachers' viewpoint, the interviews indicated that the most important function was to deal with those unexpected questions and to comment on student solutions to exercises (which also requires the ability and confidence to judge new variations).

Although local deliverers might well welcome a remote expert who volunteered to give lectures, do all the marking, and so on, their most valuable actions (because hardest to find a substitute for) are probably to give some discussion and to answer questions after students have done the reading or heard a lecture, and to give comments on the good and bad features of student solutions (even if they do not decide numerical marks).

Conclusions about future cost benefit studies

Clearly this was an exploratory study, one of whose main contributions is to suggest how to design improved studies in future. The accuracy of the time measures could be improved, firstly, by having time diaries (logs) kept throughout the period, rather than relying on retrospective recollections; although the burden this places on subjects may often make this impractical. Secondly, some direct observation of the work could be used as a check on the accuracy of the diaries. Thirdly, they would be improved by drawing up and communicating definitions of what time was to be counted. The categories identified in this chapter would be the starting point for this.

Having said that, any future study should still continue to look for, and expect to find, new categories of time and other costs and benefits. Future studies should repeat and extend the interview approach of this study as one way of doing this. Furthermore comparative studies of other teaching situations would be illuminating, as little is known of how higher education teaching work breaks down into component activities.

Finally, the point of cost benefit studies is to support, justify, and explain decision-making: in this case, whether it is rational to join an ATOM scheme for collaborative teaching. The actual measured costs and benefits are just a means to that end. It would therefore be valuable to do some direct studies of such decision-making, for instance by interview or even thinkaloud protocols of teachers making such decisions. That might well bring up more essential categories that need to be studied.

Conclusion: Overall comments on the cost-benefit relationships

The fundamental potential advantages here are that:

1. Teachers volunteer to author a topic they are already expert in and in exchange adopt an exercise in a topic they think is important, but do not feel expert in
2. Material can be re-used in several institutions, thus both saving authoring costs and raising quality by increasing the number of deliveries and hence of iterative refinements to the material.

From the authoring viewpoint of creating exercises and the associated handouts (assuming authors use their own material), the author receives materials they do not have to write in exchange for improving something they have or would have already written. The advantage shows up in large potential time savings (at least 3:1 for a teacher who authors one ATOM and adopts two), in lower subjective effort, and in higher quality topics (i.e. a gain in curriculum quality).

A further potential gain comes from the fact that each exercise will be re-used more often, because it is used at several institutions, than it normally would be. This reduces the authoring cost per delivery, and will often lead to higher quality as feedback leads both to revisions and to the use of stored past solutions and feedback (a feature of MANTCHI not dealt with in this chapter).

From the viewpoint of the local deliverer as a course organiser, adopting an ATOM is less work than creating one's own. It is less stressful because its quality is supported by an expert author, and by having been trialled elsewhere, and because its delivery may be supported by a remote expert. Above all, it gives a higher curriculum quality. Within the project, teachers dropped the topics they valued least on their own courses and requested or selected ATOMs from others that they felt most increased the value of the set of topics on their course. (We must however remember that further steps are involved in converting improved curriculum and resource material quality into improved quality of learning attained in practice.)

From the viewpoint of the work of local delivery itself, there are three cases.

- In ATOMs without any use of a remote site, the workload is the same.
- If a remote expert is used, then local deliverers donate some tutor time on a subject about which they are highly confident in return for the same time received on a topic they have low confidence about. In contact time, this may not be a saving as some local deliverers will attend as facilitators for the occasion. However such "contact time" does not require the preparation it normally would. For marking, there is a negotiated balance, so no time should be lost.
- In ATOMs with remote student interaction, there is an extra cost of coordinating the courses at two institutions, which has to be balanced against the pedagogical gains of this form of peer interaction for students, and any relevant gains due to practice with the technology involved.

Thus in return for savings in authoring time, a gain in curriculum quality, and a gain in the quality of individual materials due to increased use and revision, there is either no penalty in net delivery time, or a small increase in time spent "facilitating". In the case of ATOMs that involve a remote expert, there is the added staff development reward of the deliverer becoming gaining confidence in teaching that particular topic.

In constructing the above arguments, I have implicitly used approaches related to those such as the "equal cost comparison" technique described in Hawkridge (1993). When it is possible, arranging comparisons where costs cancel out avoids the need for measuring the actual size of those costs. This justifies some of the above conclusions despite the inaccuracies of the quantities given in the table of data. However without improvements in quantitative measurements, it will not be possible to answer further questions that could be of practical importance in some cases, such as how much imbalance in marking should be accepted before the exchange ceases to be of overall benefit, or how to trade off reduced subjective effort for teachers in teaching topics they have more confidence about against the effort of changing courses that could have been left unchanged and finding partners for an exchange.

Acknowledgements
The authors acknowledge the support of Scottish Higher Education Funding Council (SHEFC) through the Use of MANs Initiative for the MANTCHI project. Thanks to all the people who contributed to MANTCHI (see the web site for a list of personnel: MANTCHI, 1998) but particularly those who (in addition to the authors) were interviewed: David Benyon, Alison Cawsey, Phil Gray, Alistair Kilgour, Patrick McAndrew, Julian Newman, Michael Smyth, and Alison Varey. Our thanks also go to two referees, whose efforts improved, though they could not perfect, the clarity of this chapter.

References

Doughty, G. (1996a) "Technology in Teaching and Learning: Some Senior Management Issues", "Deciding to invest in IT for teaching" in TLTSN Case Studies (HEFCE)

Doughty, G. (1996b) "Making investment decisions for technology in teaching" (University of Glasgow TLTSN Centre) [WWW document] URL http://www.elec.gla.ac.uk/TLTSN/invest.html

Doughty,P.L. (1979) "Cost-effectiveness analysis tradeoffs and pitfalls for planning and evaluating instructional programs" J. instructional development vol.2 no.4 pp.17,23-25

Draper,S.W., Brown, M.I., Henderson,F.P. & McAteer,E. (1996) "Integrative evaluation: an emerging role for classroom studies of CAL" Computers and Education vol.26 no.1-3, pp.17-32

Hawkridge,D. (1993) Evaluating the cost-effectiveness of advanced IT and learning CITE report no.178 (Open University: Milton Keynes)

MANTCHI (1998) MANTCHI project pages [WWW document] URL http://mantchi.use-of-mans.ac.uk/

Perry, W.G. (1968/70) Forms of intellectual and ethical development in the college years (New York: Holt, Rhinehart and Winston)

Reeves,T.C. (1990). "Redirecting evaluation of interactive video: The case for complexity" Studies in Educational Evaluation, vol.16, 115-131.