Bates, Simon, Hardy, Judy, Hill, Jon and McKain, David (2008) How Design of Online Learning Materials can Accommodate the Heterogeneity in Student Abilities, Aptitudes and Aspirations. Learning and Teaching in Higher Education (2). pp. 3-25.

How Design of Online Learning Materials can Accommodate the Heterogeneity in Student Abilities, Aptitudes and Aspirations

SIMON BATES, JUDY HARDY, JON HILL & DAVID MCKAIN

University of Edinburgh, UK

ABSTRACT

We describe the challenges facing higher education in terms of the heterogeneity of the cohort of students that arrive at university. The reasons why such diversity exists are many: students differ widely in terms of their preparedness for study at university, their degree choice aspirations and the issue of motivation for study of a particular subject. We illustrate how well-designed e-learning course materials can support many of the particular facets of heterogeneity by offering an inherently non-linear pathway through a collection of materials, so as to offer a degree of personalisation of the learning experience.

Drawing on our own experience of several years' development of extensive online materials to support the traditional teaching methods of a large first year physics course at the University of Edinburgh, we highlight three aspects of the design of e-learning materials that facilitate this personalisation. These are: a highly granular source of individual learning objects; online constructions ('one-downs' and 'popups') that provide additional depth and breadth of material; and the ability to import external resources adapted to the local context.

Introduction

In recent years, the number of students entering higher education has increased dramatically. In Scotland, over 50% of 18-19 year olds now proceed to some form of further study and have done for a number of years (Universities UK, 2004), with figures for England around 43% in 2003/04 (DfES, 2004). Whilst some subjects have seen a decline in numbers applying for study (the physical sciences and mathematics amongst them) the overall trend is that our class sizes at Edinburgh are larger than ever before. The context of this paper is a first year mainstream undergraduate course in classical physics taught at the University of Edinburgh (‘Physics 1A: Foundations’).

We have in some respects bucked the national trend of declining applicant numbers in the physical sciences, and have been fortunate to enjoy modest but steady growth in entrant numbers. Not only have numbers increased, but so has the heterogeneity of the student cohort. Approximately 24% of students on the course are female and all entrants arrive with a wide range of prior qualifications in the subject; Scottish Highers (27 and 31% of the cohorts in 2004-05 and 2005-06, respectively), Advanced Highers (45 and 41%) and A-levels (28 and 25%), as well as a small number of other qualifications such as International Baccalaureate and Irish Leaving Certificate. Coupled with the heterogeneity in prior study is a broadening range of preparedness for university study of the subject. The most often-cited example of this is ‘the Maths problem’, the increasingly wide range (and overall decline) of mathematical ability of entrant undergraduates (Smith, 2004). A final dimension of heterogeneity is that the course is taken by over a hundred students each year who will not pursue a physics degree, but have selected this as an ‘outside’ course. Different student abilities, aptitudes and aspirations on the course present complex challenges to us as educators, as does the added ingredient of this being a first year course and the issues of orientation to the mode of studying and learning (and life) at university.

Student activity on the course is equally heterogeneous. An example of this is found in lecture participation. We have correlated participation in lectures, as evidenced by use of an electronic voting system that we have used extensively throughout the course (for a description of use, see for example Bates *et al.* 2006), with end of course assessment. This naively assumes that students who faithfully attended all lectures, participating in all the interactive engagement exercises, would generally fare better in end-of-course assessment than those who did not. The data are shown in Figure 1.

Whilst there is a slight positive correlation, it is not significant and the spread of data indicates widely varying practice. If we assume that all students participate if they attend (reasonable, although some may forget handsets) and that participation confers some sort of learning experience (we would certainly hope so, but it is impossible to measure in isolation) then it is clear that it is possible to gain a very good mark in the end of course assessment without attending all the lectures. Concomitant with this increased size and diversity of student cohorts have been the changes brought about by continuing, perhaps even accelerating, adoption and integration of learning technology

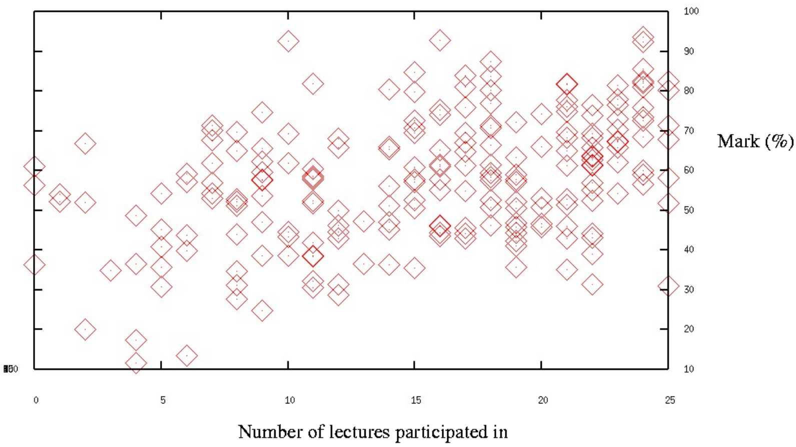

Figure 1: Correlation of lecture participation with end of course assessment mark (as a final percentage).

into all aspects of university life. In the last decade, computers have become much cheaper, vastly faster, increasingly portable and, most recently, wireless. Virtual Learning Environments (VLEs) have become almost ubiquitous at institutional level and 'e-learning' as a phrase is a broad term that can be bandied about for activities as diverse as fully-distance learning courses to simply storing vast quantities of lecture presentations online.

We acknowledge at the outset that this paper reports a modest and practical undertaking, a combination of reasonably affordable (in terms of staff time and effort) measures to move away from a 'silo' model of online learning resources. It impinges on the huge body of work on hypertext: 'computer-based texts that are read in a non-linear fashion and that are organised in multiple dimensions' (Landow, 1992). (Our students inhabit both real and virtual learning environments, which are designed to have strong coherence between them.) Equally, what we report here is a long way from the realm of intelligent tutoring systems and adaptive courseware generation, where multiple instances of an aggregation of learning materials can be provided for students according to their needs before they encounter it, or even dynamically (Brusilovsky & Vassileva, 2003).

This paper attempts to provide an illustration of how careful design of online learning materials can accommodate such heterogeneities, providing different students with a variety of highly non-linear pathways through online materials, offering a degree of personalisation of the learning experience. The paper is structured as follows: in the next section we describe the course context and design philosophy, together with details of how student routes through the material are recorded and analysed; we then present a number of design features of the online materials that facilitate this flexibility. Evidence of different and complex routes through the material are then presented, followed by a discussion and cautionary tale about extensive access or use of online materials being no guarantee of success on the course.

The course, its design and tracking of student activity on the course

The overall design philosophy of the course, together with suggested guiding principles that underpin the design choices, have been reported elsewhere (Bates *et al.* 2005). Here, we present a recap of salient features. The online course materials as they currently exist are the result of over seven years' ongoing development and refinement. The aim of the resources is to support, not supplant, the face-to-face teaching structures employed in an on-campus university (lectures, tutorials, workshops, laboratories). It has been suggested that learning design in the online environment is often not articulated explicitly within the VLE (Vogel & Oliver, 2006). The approach that we have taken is to tightly integrate the online resources with the face-to-face teaching activities and make them available for periods of student self-study. They are an integral part of all activities on the course, thus becoming part of a coherent whole (a blend), rather than something supplementary or an afterthought or add-on. Online resources add another dimension to the range of learning experiences available to a student and therefore make the learning experience more inclusive. Of course they raise a range of potential accessibility issues, many of which are surmountable by observing good practice in design.

The online course material now comprises over 1400 grains of information ('knowledge objects' such as pictures, text, examples, questions, applets etc.), which are aggregated in small clumps to make 'learning objects', which we define as a combination of information and some activity; for example, interactive multiple choice question self-tests, or tutorial-style questions. These learning objects are bundled together to form coherent topics such as the study of linear motion. Seven topics comprise the material covered in the course. The principal components of the online material are electronic course notes, interactive self-test multiple choice questions, tutorial questions (with hints and progressively-revealed solutions) and material to support studio-style group work activities. These are supplemented by various pieces of course information, administration and a discussion forum. Most knowledge objects in the course are reached by, at most, three mouse clicks.

The course is taken in the first semester of the first year of study by approximately 250 students each year, with a diversity of backgrounds. Lectures have recently evolved to be less like a one-way transmission of the course material, and more like a two-way conversation. This conversation is mediated by the technology of electronic voting handsets that are used regularly (at least once per lecture on average) throughout the course to provide interactive engagement exercises and peer instruction (Mazur, 1997). The evidence of the efficacy of such methods to improve student learning is well documented (Hake, 1998; Draper & Brown, 2004; Stuart *et al*. 2004). The course material comprises a tour of classical (Newtonian) physics. It is now accepted that, although elementary, this material embodies complex and often counterintuitive concepts; much of the US Physics Education Research effort over the last decade (see Wieman & Perkins, 2005, for a perspective) has sought to elucidate and provide strategies to dislodge the misconceptions that students intuitively have about the study of forces, kinematics and dynamics. The dislodging process is made more difficult by entrant students taking one look at the syllabus, seeing familiar laws and principles and thinking this course will teach them nothing new. The online materials, being an integral part of the course and its design, play an important role in this process.

Design examples that facilitate personalised routes through material

We start from the assumption that a diverse student cohort will exhibit diverse utilisation strategies for online materials; in other words, one size does not fit all. We illustrate how the course design supports differentiated routes through the materials with three specific examples: the inherent granularity of the materials and the way they are aggregated and cross-linked; the degree of access of on-page interactive elements ('inlines' and 'popups') that afford differentiation in depth and breadth; and the importing of external online resources that can be adapted ('wrapped') to fit the local context. Each of these is described in more detail below.

Granular, cross-linked source

The inherent organisation of the material is itself supportive of differentiated routes through the material. The material is highly cross-linked, allowing multiple navigation possibilities from pages. As an illustrative example, we consider the formative Multiple Choice Questions (MCQs) used throughout the course, a good example because they are not provided for the students via any other medium than online. These resources may be accessed in one of two ways. The first is enabled by tightly coupling them to individual pages of content. A student can work through content and then choose to test their understanding of it by practising any of the questions that directly relate to that content. An alternative route to them is offered by a collected list of all questions for students to browse through, grouped according to the same topics that are used to organise the content itself. The same is true for the course questions, which are either interleaved within relevant content or organised for browsing.

This organisational design represents a balance between maximal reusability (stand-alone learning objects) and contextualised utility for teachers (inter-relationships between discrete objects), as has been previously observed (Duncan, 2003).

On-page interactive elements

The course content is designed with on-page interactive elements that facilitate differentiated routes in depth and breadth, using JavaScript within the page. These constructs are referred to as 'inlines' and 'popups' and have clearly delineated functions. Both do exactly as the names suggest. Inlines allow the students to open up additional pieces of material inline in the page. These constructs contain additional material such as worked examples, proofs, commentaries, help and advice. It can be thought of as the flesh on the skeleton. Some of this material, but by no means all, is discussed and worked through (often in an abbreviated form) in lectures. Students can choose to open these panels if they want access to a greater depth of material. A good example is illustrated in Figure 2, which shows nesting of these elements employed to illustrate and analyse the concepts associated with non-inertial (accelerating) frames of reference.

Inlines are part of the core material of the course. 'Popups' are not, hence the decision to make them distinctly different in operation to inlines. They open up in separate windows and contain material that is deliberately 'off the beaten track'. It might be additional, more challenging material, perhaps containing a greater degree of mathematical rigour. They can also contain media such as applets, simulations or movies and frequently adopt a more discursive tone than the core material. An example is shown in Figure 3, illustrating a more in-depth treatment of Edward Lorentz's famous simulations of weather systems to illustrate extreme sensitivity to initial conditions, one of the signatures of chaos. This material is deliberately included to provide additional challenge and stimulate interest of able students studying towards Physics degrees. Popups and use of resources such as movies or simulations residing externally to the course content raise accessibility issues that should be considered in course design.

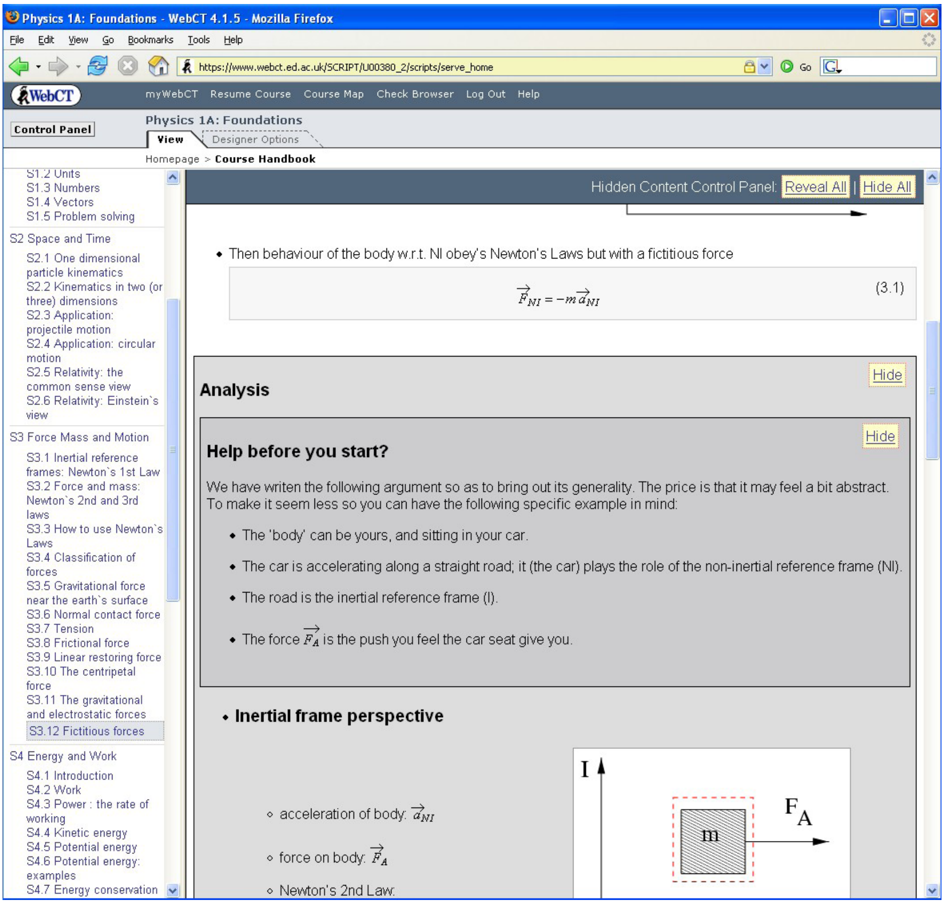

Figure 2: Illustration of nested inlines to provide analysis and description.

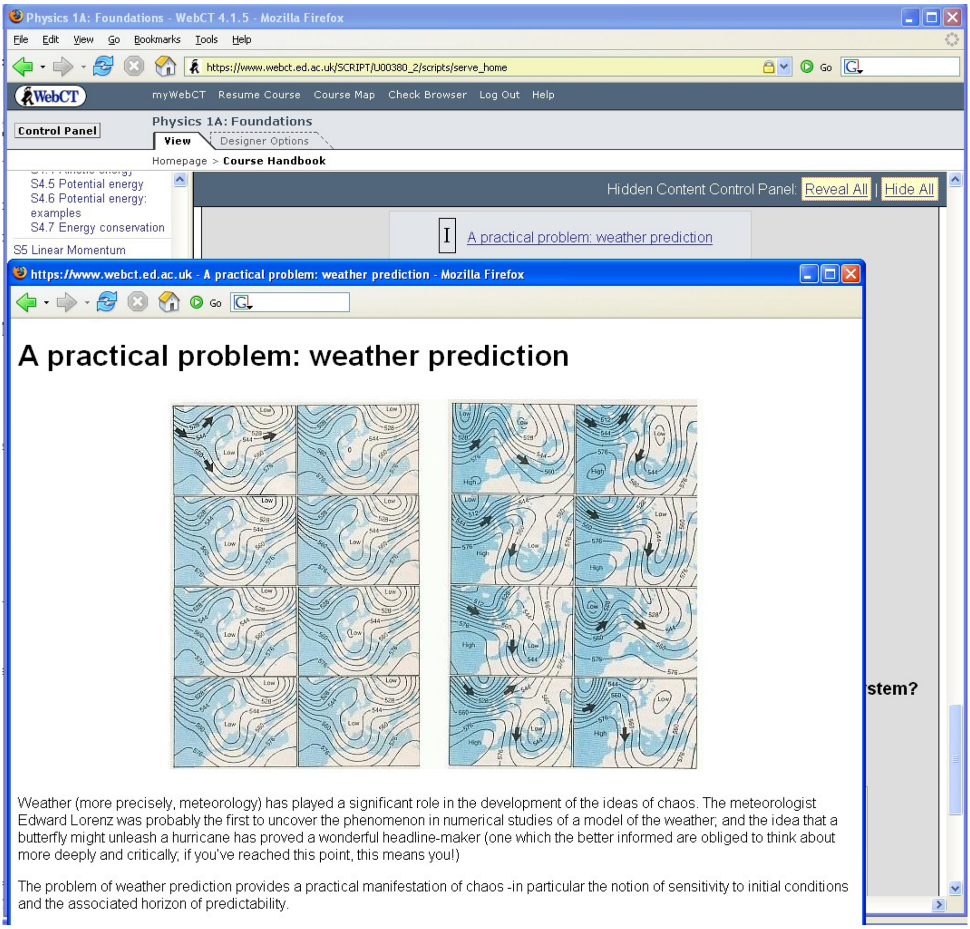

Figure 3: Additional material in a popup, with the referring page behind.

Imported resources, adapted to fit a local context

One of our guiding principles is to use suitable material that is freely available elsewhere on the web, rather than trying to reinvent all these particular wheels. A wide variety of different resources can be offered within the integrated framework of the course provided that they are granular enough to be able to be customised or embedded within the local context. Two good examples of such resources are applets (in plentiful supply on the web) and mathematics support material, in this case video material developed by the Maths for Engineers project (Higher Education Academy Engineering Subject Centre, date unknown). These resources can be wrapped with learning or study activities, to inform and guide student use of these resources.

Example routes through material evidenced from tracking data

For the last two years, student activity and access of the online course materials has been tracked to gain insight into how students make use of the materials, at a far higher level of detail than is provided within commercial VLEs, such as WebCT, through which the students access the material. The technical details of this activity, together with broad patterns of student access (and the implications this has for future course design), have been presented elsewhere (Hardy *et al.* 2005; 2006). A key feature to note here is that the tracking records both the time spent on a particular page and referring page from where the student came. This allows us to construct spatial and temporal ‘maps’ of student sessions of access of the material. Tracking events are not simply restricted to page loads, but also capture the on-page interactive elements.

The access tracking of e-learning content is done via a small JavaScript application, which records data to a database. The information held includes the page currently being visited (via a page ID), the time and date of the visit, the referring page, the page type (one of inline, popup or page), the user ID, a unique session identifier (session ID) and information about the computer platform used to visit the site. The page ID also contains the content type, which is one of MCQ, tutorial, workshop or content (course handbook). For the purposes of this study, the session ID, page ID, the referrer URL and the time fields are the relevant fields.

Overall access patterns

From these data it is possible to calculate an average session, which lasts around 15 minutes and consists of 2.3 page hits. Clearly, this figure is derived from a vast range of possible values ranging from a short 1 page session to a marathon session of nearly 12 hours and over 100 page hits. The length of each session and the number of hits varied throughout the year, culminating in a peak of both metrics just before the end of semester exams (see Figure 4).

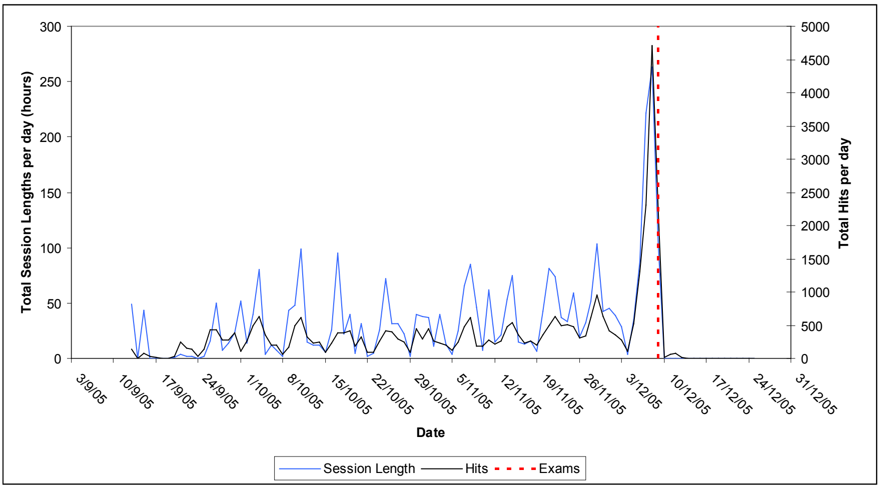

Figure 4: Total session lengths per day and number of hits per day through the course. There is a marked increase in both number of hits and the time spent on the e-learning course as the exams approach (dashed line).

However, more revealing information comes from looking at the type of content accessed through time as stored in the page ID. Not only is there a general increase in the amount of pages being accessed and the length of time being spent on the site, but there is also a clear increase in the access of MCQs as exam time approaches (Figure 5). It is clear that students change how they utilise the site, depending on the demands of the course at the time.

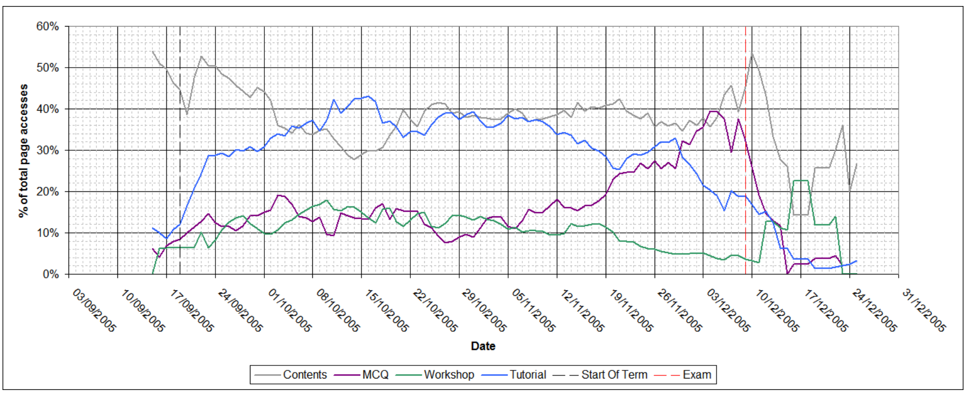

Figure 5: Course access classified by content type. Each line represents the seven-day average of the percentage of page hits within each content type and clearly illustrates a change in the type of content accessed through the semester. The vertical lines represent the formal start and end of the teaching period.

Paths through the course

In an attempt to visualise how a student navigates the vast amount of online material available, three visualisations can be produced for any particular session using the tracking data described above. Each type of visualisation highlights one or more aspect of the session. All visualisations are produced using short Perl scripts which extract data from the database. The graphics are produced by either gnuplot <www.gnuplot.info/> or GraphViz <www.graphviz.org/>.

A path visualisation takes the form of a directed graph (see Figure 6). Each page is represented by a node, which is shaped according to the type of page being accessed and shaded according to the type of content being viewed (Table 1). The nodes are linked by arrows with a numerical label indicating the order of access, starting from zero. The visualisation also includes a student identifier (removed from graphs included in this paper) and the times of the first and last access. This is the only visualisation that includes inline and popup elements.

Like the path visualisation, page access visualisation is in the form of a directed graph, with nodes representing pages and edges representing access between pages. No inline or popup windows are shown so node shape is not relevant and all nodes represent content. The shading of the node represents the content type as described in Table 1. The size of the circle depicts the length of time spent on that page derived from the start and end times recorded in the database. On larger nodes, a label gives the actual time spent on the page. Due to the large ranges of times that occur in the data, the node size is not linearly related to the time spent on the page, but is logarithmic, such that the difference in node size between five and ten seconds is approximately the same as the difference between two and ten minutes. There is a minimum time limit of around 2.5 seconds and a maximum of 600 seconds.

The final visualisation used is access time visualisation, which depicts the time of each page access. The result is a 2D graph with a vertical line representing a page access. The graph looks like a ‘barcode’, with density of lines representing temporal density of page visits.

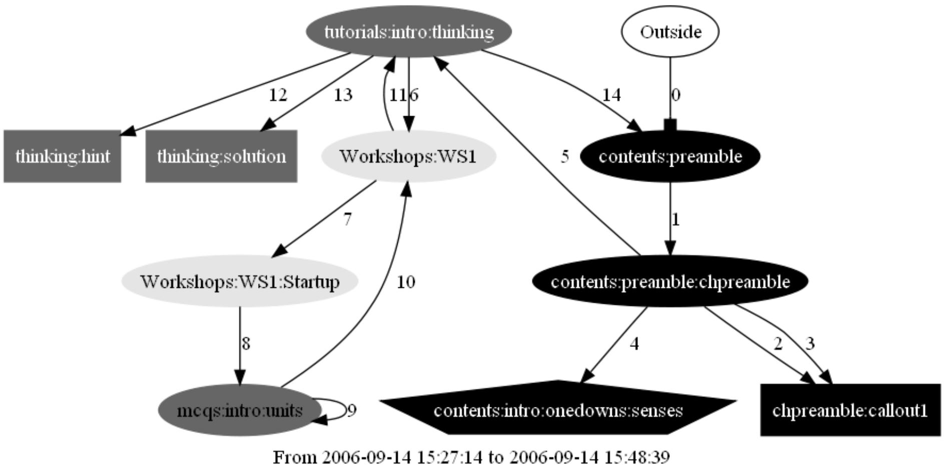

Figure 6: A path through the online course by one of the authors. This path shows access to all four content types and all three types of pages as shown in Table 1 below.

Shape	Page Type
tutorials:intro:thinking	Content Page
thinking:solution	Inline
contents:intro:onedowns:senses	Pop-up
Shading	**Content Type**
mcqs:intro:units	MCQ
tutorials:intro:thinking	Tutorial
contents:preamble	Contents
Workshops:WS1:Startup	Workshop

Table 1: Key for path visualisations. Shapes represent the page type (page, inline or popup). Shading depicts content type (MCQ, tutorial, contents or workshop).

Example sessions

There are a large number of sessions that could be visualised and, rather than attempt to provide some kind of overarching conclusions from these, we shall highlight some interesting examples. Navigation through the site is fundamentally dictated by the available links on any given page and the browser controls available. A single page can have three type of links displayed. The main navigation links occupy a space to the left of the page. These can be used to move around the current content type (e.g. pages of content etc.). Secondly, there are breadcrumb trails and general links on the top of the page. These include the last few sections visited as well as links to the homepage and general WebCT navigation (see Figure 7). Finally, there are links embedded on the page itself, as in any other webpage. Access to the inline elements and popup windows is through this type of link.

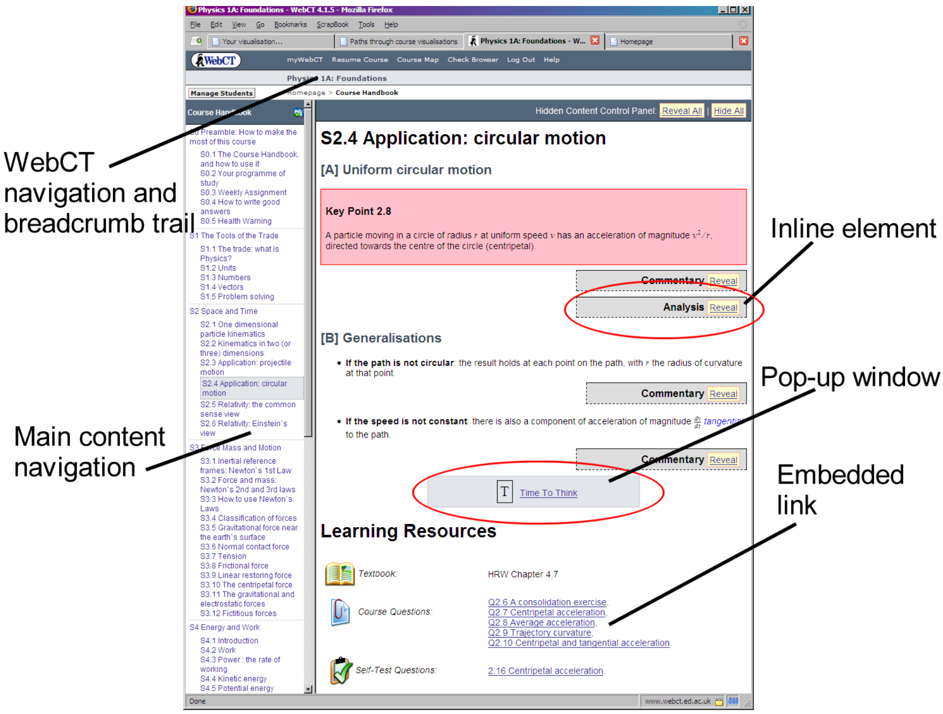

Figure 7: Screenshot of WebCT course pages. The left-hand navigation, breadcrumb trail and embedded links provide non-linear access around the site. Inline elements and popup windows add extra information to the pages.

The complex navigation structure outlined above leads to complex pathways through the course material. An example of such a pathway is shown in Figure 8. The student shows an extremely non-linear access pattern; they did not simply click down the list of a single type of content as displayed in the navigation page, and instead journeyed around the content that they required. This is not to say that they navigated in a random or haphazard way, instead they kept within the same area of the course (oscillatory motion) but accessed different types of content material. From the path visualisation alone, it is not entirely clear what the student was doing. Did they stop and read every page in detail? Were pages simply used as a reference point to gain access to other materials? Using the page access visualisation, more detail can be added and a clearer picture of the session can be constructed.

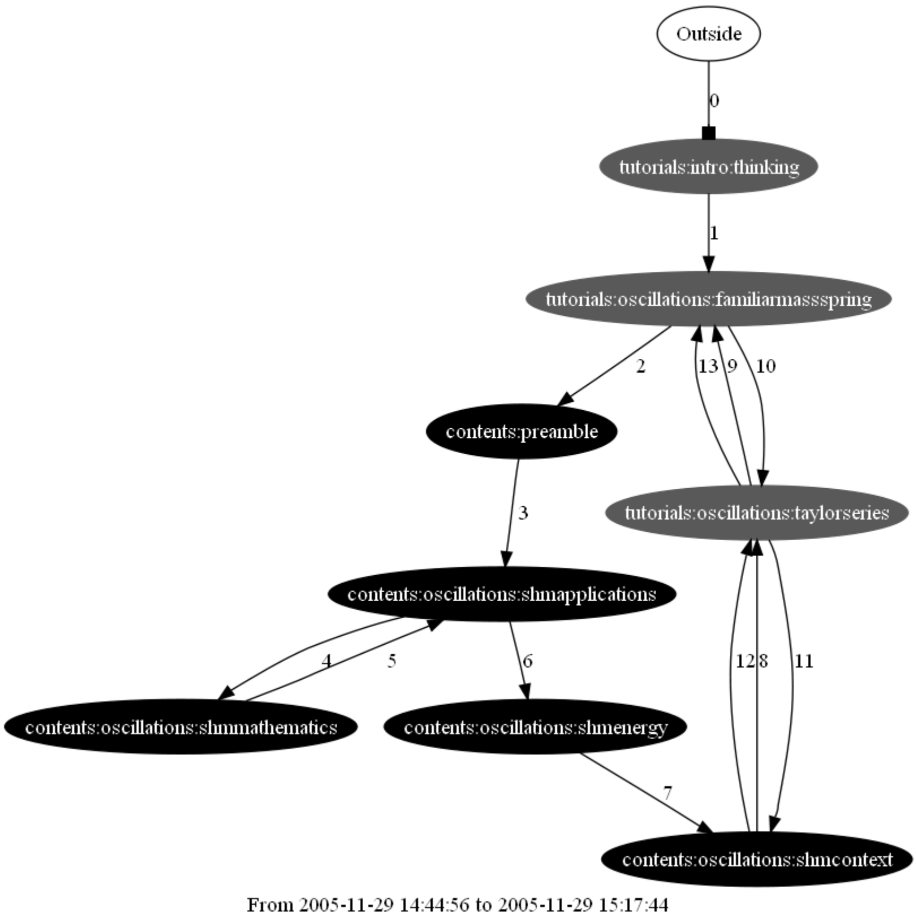

Figure 8 (see also facing page): Path and page access times for a typical student session. A detailed picture of what content was accessed and for how long can be built using these two visualisations.

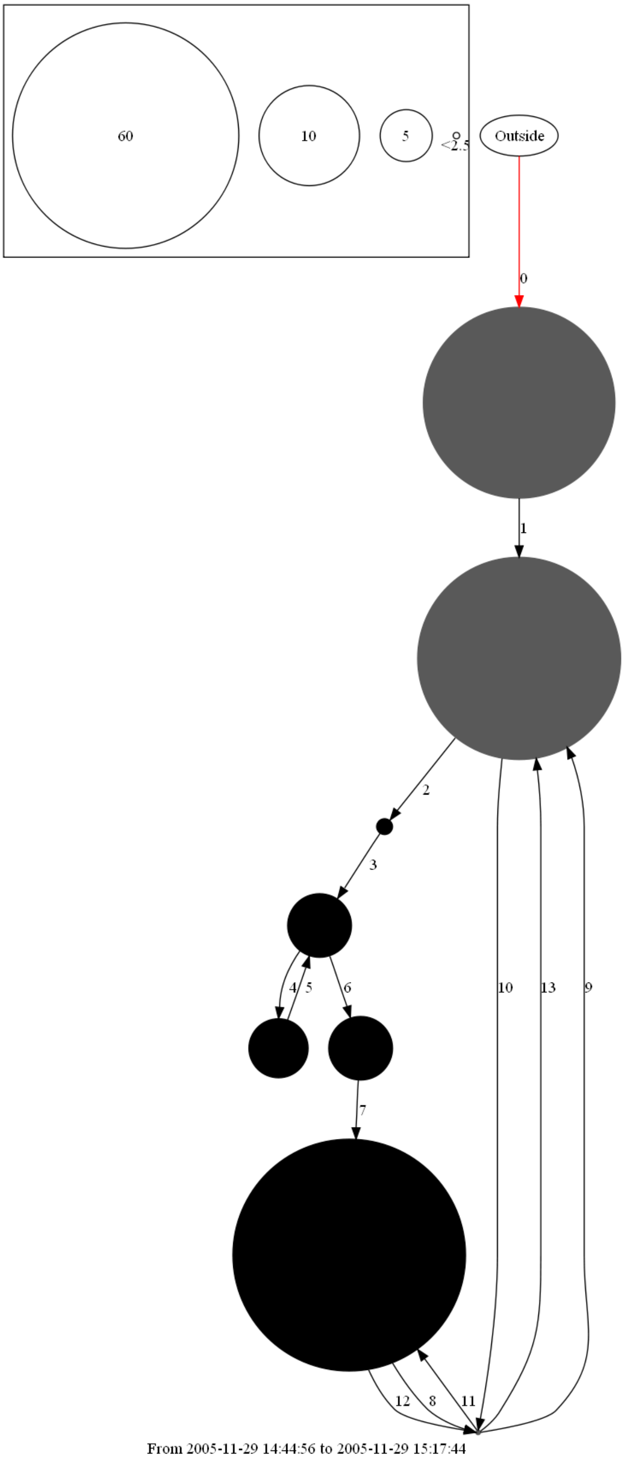

Figure 8 (continued)

A 'story' of the session can be constructed from this information. The student spent a long time on the first two nodes even though the first node (tutorials:intro:thinking) is not in the same subject area as the rest of the pages, but may have been used as a stepping point for content. The student then read the tutorial page on the 'mass on a spring' experiment. They did not access any of the supplementary material on that page, which includes inlines and links to other parts of the course. They then went back to the homepage and looked through the course handbook (content nodes) until settling on the simple harmonic motion context page (from link 7 to 8). They then went back to the original tutorial page via another which is linked on the bottom of the simple harmonic motion context page. They then flicked back and forth between the tutorial section and the contents page, spending some amount of time on each. In other sessions, students have accessed the full range of materials available on a page, including inline elements and popup windows and show repeated visits to the same pages (illustrated Figure 8). Another interesting class of session is one where the student has simply clicked the list of topics in the navigation for a particular type of content in turn resulting in a near-linear pathway. This type of access pattern is particularly common for the MCQs in the days preceding the end of semester exam. These sessions span a large range of page hits and time, from a few pages viewed over a few minutes to nearly a hundred pages in nearly five and a half hours, although examination of the timeline visualisation (Figure 9) shows that they did take a break for a few hours.

Discussion

It is clear that the design of these resources supports a variety of different ways of using and navigating through them. But what, if anything, can we say about the learning that is taking place and how the resources influence or contribute to that? In fact, the non-linear nature of web-based learning spaces has been argued to be disorienting to students (Begoray, 1990), suffering from cognitive overload if the task of navigating through materials distracts from the intended learning goal (Conklin, 1987). This is particularly acute for students with accessibility difficulties (Evans & Sutherland, 2003). Furthermore, it has been observed that one outcome of 'ubiquitous computing' (Crook & Barrowcliff, 2001) is a multi-tasking approach, in which online study and recreation are tightly interleaved. There is not quite the same complexity in the information space represented by these course resources, since there is effectively controlled non-linearity mapped onto a linear progression through the course

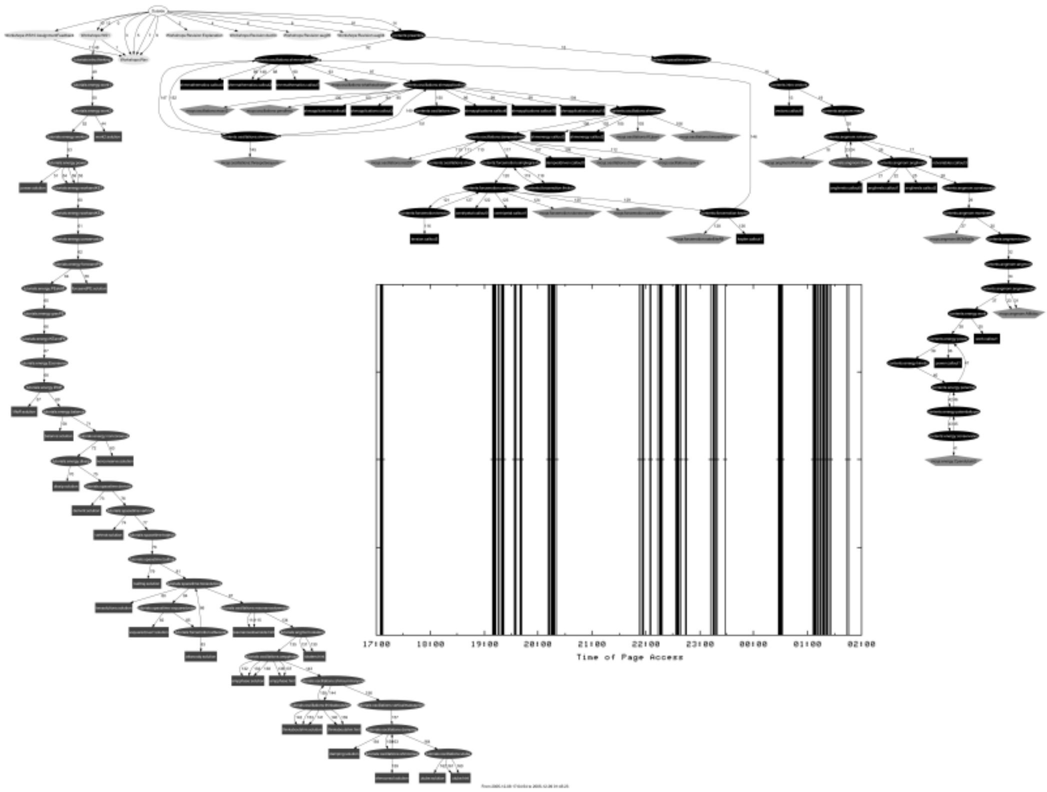

Figure 9: Illustrative example of a long session in which lots of content was accessed, both in terms of type of access and content type.

materials. In addition, the 'real' learning environment of the course maps the same space. However, there is certainly the possibility that flexibility can seduce users, encouraging distraction, so that students 'float' or drift through resources rather aimlessly or get diverted out of the main body of the material by following external links and subsequent follow-on links from these pages. Judging from the path visualisations carried out so far, this does not appear to be the case and students appear to know exactly what they are after when visiting the course material.

Students frequently indicate that the online resources are one of the best things about the course (interestingly, the engagement of staff in the face-to-face teaching episodes fares equally well, suggesting that students really can have the best of both real and virtual learning environments simultaneously). This is perhaps not too surprising; certainly, earlier studies (Crook, 1997; 2002) have drawn similar conclusions.

We have attempted to correlate the level of access of the online resources by students with their end of course assessment. This is admittedly a very crude measure. One might assume that a high volume of access, and therefore use, leads to good learning and understanding and a good mark at the end of course assessment. (Though it is not clear whether this relationship - should it even exist - is one of cause or effect. Are good students well motivated and will they access all available resources heavily, or has using the resources made them 'good' students?) Figure 10 shows the correlation of lecture attendances versus end of course mark shown in Figure 1, this time cross-correlated with the number of page accesses made by a given student on the z-axis. The 3-d grid data points are interpolated and an isosurface plotted (see top panel), with contours of constant number of page hits mapped onto the x-y plane (lower panel). The trends we discuss are relatively insensitive to the choice of measure of online activity; we could have chosen 'number of distinct sessions lasting more than 5 minutes' and the qualitative picture is unchanged.

This graph illustrates the variation in student 'personalities' on the course with respect to attendance at lectures and access of online materials. In a sense, it serves to illustrate the heterogeneity in behaviour on the course, not surprising given our previous statements about heterogeneity of ability, attitude and aspiration. It is abundantly clear that there is no simple correlation between the end of course mark and level of access of online material. Features of note illustrated on the 2-d contour plot (lower panel) include the community of learners (labelled 'A') who achieve first class marks on the course yet participate in very few lectures, whilst at the same time showing a high level of activity on the online materials. Once again, we have difficulty untangling cause and effect. On the one hand, if this group really is making a conscious choice about not attending lectures and deciding to spend study time in a different way, it is an argument for supporting flexible styles of learning. If, however, these students simply don't turn up but get by (rather successfully) on previous knowledge and use of online material, it rather invalidates that argument.

More worrying is the group of students identified by 'B' in the lower panel, who attend few lectures, maintain a high level of activity on the course site and yet do poorly in demonstrating evidence of having achieved the learning outcomes in the end of course assessment. More worrying still are the 'C' group, who faithfully attend all lectures, maintain a high level of activity on the site (though at this resolution we cannot identify what or how effective that activity is) and fail the course. A plot such as this prompts many questions for course designers and teams, and in combination with other analyses and course activities (such as diagnostic testing) will shape the direction of the course in future years of delivery. In ongoing work, we are looking for differences in the online pathways taken by different student 'personalities' that we have just described.

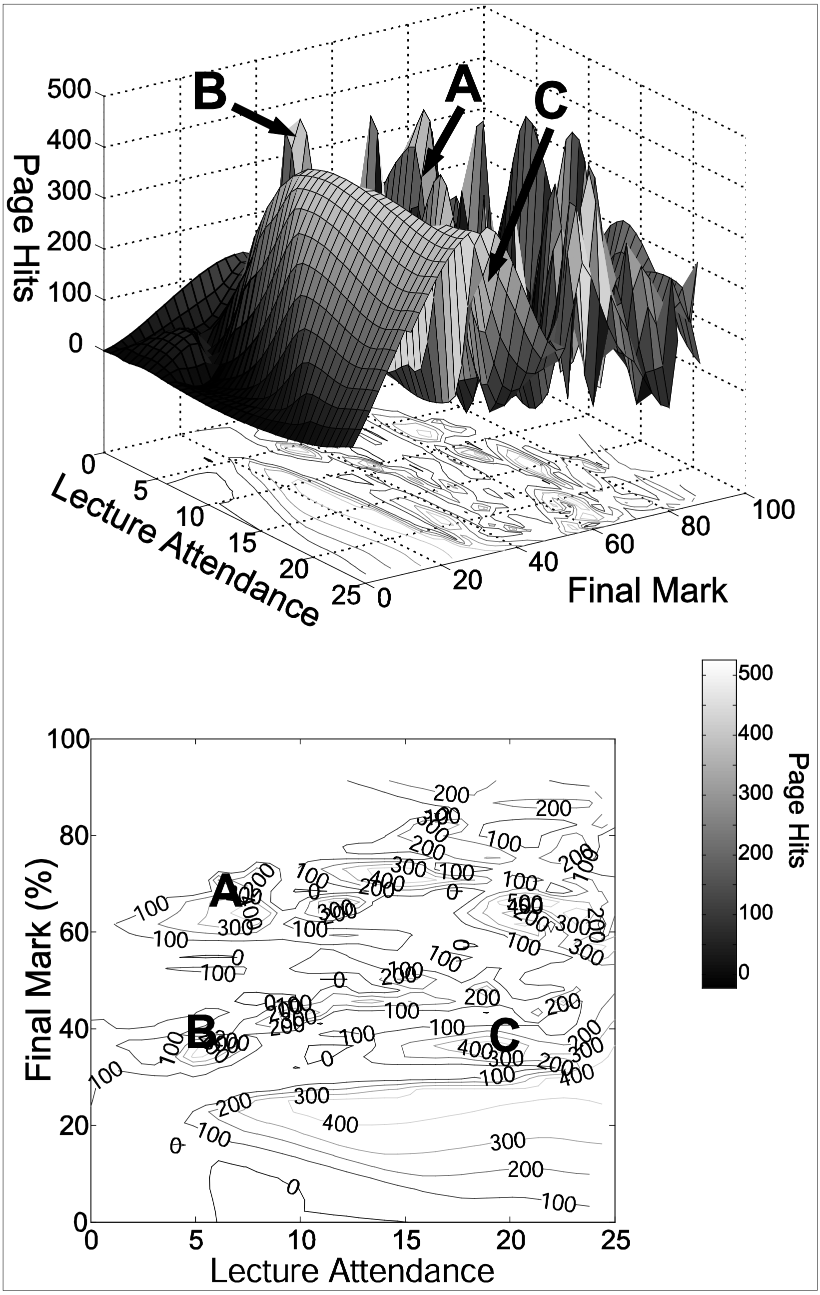

Figure 10: Correlations of lecture attendance, end of course mark and number of hits to online resources. The top panel illustrates a 3-d isosurface, the lower panel a 2-d contour plot showing contours of numbers of hits. See text for descriptions of labels A, B and C.

Conclusions

We have presented details of course design that support non-linear and personalised routes through course material, together with examples from student tracking data to illustrate the complexity and variety of these routes. An analysis of online activity with end of course mark, cross-correlated with lecture participation, leads us to suggest that providing differentiated routes through online material is by no means a guarantee of success on the course or its effectiveness. Online materials may support student learning, but their effectiveness is hard to quantitatively assess in an environment in which a multiplicity of factors contribute in unknown proportions. However, we do conclude that the architecture of online resources for a given course feeds into the entire learning design and how the online resources support the learning outcomes associated with the course. The accessibility challenges that are raised by utilising online materials can be met by adherence to good practice in design. As online materials look to be a continuing and increasing feature of modern learning, much more consideration needs to be given to their design, moving away from a philosophy of 'lecture notes on the web'.

Acknowledgements

The tracking studies in this work have been funded in part from the University of Edinburgh e-Learning Fund.

References

Bates, S.P., Bruce, A.D. & McKain, D. (2005) Integrating e-learning and on-campus teaching I: an overview, in Cook, J. & Whitelock, D. (eds) Exploring the Frontiers of e-learning: borders, outposts and migration, *Research Proceedings of the 12th Association of Learning Technology Conference,* pp.130-139.

Bates, S., Murphy, A.St.J. & Howie, K. (2006) The use of electronic voting systems in large group lectures: challenges and opportunities, *New Directions: the Journal of the Higher Education Academy Physical Sciences Centre,* issue 2, December 2006, pp.1-8, available via <www.heacademy.ac.uk/physsci/publications/newdirections> (date accessed 31 December 2007).

Begoray, J.A. (1990) An introduction to hypermedia issues, systems and application areas, *International Journal of Man-Machine Studies*, vol.33, pp.121-47.

Brusilovsky, P. & Vassileva, J. (2003) Course sequencing techniques for large scale web-based education, *International Journal of Continuing Engineering Education and Life-Long Learning*, vol.13, 1-2, pp.75-94.

Conklin, J. (1987) Hypertext, an introduction and survey, *Computer*, vol.20, pp.17-41.

Crook, C. (1997) Making hypertext lecture notes more interactive: undergraduate reactions, *Journal of Computer Assisted Learning*, vol.13, pp.236-244.

Crook, C. (2002) Deferring to resources: collaborations around traditional vs computer-based notes, *Journal of Computer Assisted Learning*, vol.18, pp.64-76.

Crook, C. & Barrowcliff, D. (2001) Ubiquitous computing on campus: patterns of engagement by university students, *International Journal of Human Computer Interaction*, vol.13, pp.245-256.

DfES (2004), *Participation Rates in Higher Education for Academic Years 1999/2000-2003/2004*, available at <www.dfes.gov.uk/rsgateway/DB/SFR/s000572/index.shtml> (date accessed 13 September 2006).

Draper, S.W. & Brown, M.I. (2004) Increasing interactivity in lectures using an electronic voting system, *Journal of Computer Assisted Learning*, vol.20, pp.81-94.

Duncan, C. (2003) Granularisation, chapter 2 in Littlejohn, A. (ed) *Reusing Online Resources: a sustainable approach to eLearning*, London: Kogan Page.

Evans, S. & Sutherland A. (2003) *VLE User Testing*, Techdis, available at <www.techdis.ac.uk/index.php?p=3_8_26> (date accessed 8 May 2007).

Hake, R.R. (1998) *American Journal of Physics*, vol.66, pp.64-72.

Hardy, J.A., Bates, S.P., Antonioletti, M. & Seed, T. (2005) Integrating e-learning and on-campus teaching II: evaluation of student use, in Cook, J. & Whitelock, D. (eds) Exploring the Frontiers of e-learning: borders, outposts and migration, *Research Proceedings of the 12th Association of Learning Technology Conference*, pp.130-139.

Hardy, J.A., Bates, S.P., McKain, D., Murray, K., Paterson, J., McGonigle, B., Vigentini, L. & Jackson, J. (2006) The modus operandi of the next-generation e-learner: an analysis of tracking usage across the disciplines, in Whitelock, D. & Wheeler, S. (eds) The Next Generation, *Research Proceedings of the 13th Association of Learning Technology Conference*, pp.108-120.

Higher Education Academy Engineering Subject Centre, date unknown, available at <www.engsc.ac.uk/er/engmath/webdisk.asp> and <www.mathsforengineers.net/> (date accessed 12 September 2006).

Landow, G.P. (1992) *Hypertext, the convergence of contemporary literary theory and technology*, Baltimore: John Hopkins University Press.

Mazur, E. (1997) *Peer Instruction: a user's manual*, New Jersey: Prentice Hall.

Smith, A. (2004) Making Mathematics Count, available at <www.dfes.gov.uk/mathsinquiry/Maths_Final.pdf> (date accessed 10 August 2006).

Stuart, S.A., Brown, M.I. & Draper, S.W. (2004) Using an electronic voting system in logic lectures: one practitioner's application, *Journal of Computer Assisted Learning*, vol.20, pp.95-102.

Universities UK (2004) *Higher Education in Facts and Figures*, available at <http://bookshop.universitiesuk.ac.uk/downloads/facts04.pdf> (date accessed 13 September 2006).

Vogel, M. & Oliver, M. (2006) *Designs for Learning in Virtual Learning Environments; insider perspectives*, available at <www.jisc.ac.uk/uploaded_documents/D4L_LVE_report_final.pdf> (date accessed 26 June 2006).

Wieman, C. & Perkins K. (2005) *Physics Today*, pp.36-42.

About the authors

Simon Bates is Director of Teaching in the School of Physics at Edinburgh and course organiser of a large first-year undergraduate course comprising 250 students each year. For this course, the face-to-face delivery is supported by a diverse collection of some 1300 learning objects, including formative self-tests and interactive lecture notes which facilitate a personalised route through the course material. Tracking of students accessing these materials is utilised to inform future course design. Current projects that Simon is working on within the framework of this particular course include the use of Personal Response systems as a tool for interactive engagement in large lectures.

Judy Hardy is a project manager at EPCC, the Edinburgh Parallel Computing Centre, at the University of Edinburgh. She teaches on a number of university courses at both undergraduate and postgraduate level and is responsible for the provision of specialist training in high performance computing. She has been involved in a range of e-learning projects over a number of years; currently these include the tracking of online student activity for the evaluation of e-learning resources and the provision of online training material to support researchers across a number of European academic institutions.

Jon Hill has been an applications consultant at EPCC for just over a year. He completed his degree in Geophysics from the University of Edinburgh in 2001. He then moved on to EPCC to complete an MSc in High Performance Computing, graduating in 2002. Since then he has been carrying out research in geology, involving rocks and computers, for a PhD and working for EPCC. He has been involved in e-learning for the last six months, generating and understanding the paths that students take though e-learning courses. He is also involved with teaching on the HPC (High Performance Computing) Masters course and promoting scientific research using HPC.

David McKain is an e-learning developer in the School of Physics at the University of Edinburgh. He is currently involved in a number of development projects including Aardvark, the School of Physics' system for authoring and managing single-sourced multiple output media learning objects, which powers a number of flagship undergraduate courses and the COSMaP Project, a cross-School collaboration aimed at improving core maths skills by flavouring the material according to students' degree interests. Smaller projects currently underway include developing an interface to a large bank of formative self-test questions that educators can browse, aggregate, package and integrate with their existing e-learning materials with minimal effort.