The complications of the orchestration clock

Pierre Dillenbourg*[1], Nan Li*[2] and Łukasz Kidziński*[3]

*École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

Introduction

European universities responded to the emergence of MOOCs (massive online open courses) in diverse ways. The majority of the 4,000 universities are observing the phenomenon without actively taking part in it. A few dozen have produced MOOCs, gaining some understanding of the ins and outs, but without placing MOOCs at the core of their strategy. A few universities have engaged more radically, producing dozens of MOOCs. This is the case for our university, EPFL (École Polytechnique Fédérale de Lausanne), which has produced 40 MOOCs, is preparing as many and has reach its millionth participant in September 2015. The rationale behind this engagement is not that MOOCs will become the dominant paradigm of HE (higher education). Instead, our efforts are driven by the conviction that MOOCs contribute to a deeper evolution of universities towards digital entities. Other forms of online education will emerge in the coming years. However, we believe that we learn more about this transformation by being an actor than by being an observer. In this chapter, we have selected one aspect of the digital transformation in HE: the timing of learning activities. Our analysis has been triggered by the current debate between two types of MOOCs: those which impose a rigid time structure and those which allow individual pacing. However, time is not a MOOC-specific issue, as it concerns any form of education. Hence we broaden our analysis to any formal education setting.

In the watch industry, a complication refers to any watch feature beyond the mere display of hour and minutes, such as dates, moon phases, alarms, etc. In a classroom, the clock on the wall simply indicates hours and minutes, but this information triggers multiple wheels in the orchestration of pedagogical scenarios. In a MOOC, the servers run on a central clock, but complications refer to the pedagogical mechanisms that are articulated to the central clock.

[1]Email: pierre.dillenbourg@epfl.ch
[2]Email: nan.li@epfl.ch
[3]Email: lukasz.kidzinski@epfl.ch

Time matters in education

This chapter is neither about philosophy nor about physics. Time is a very pragmatic component of education as illustrated hereinafter by many examples. A curriculum is defined by two values: a set of learning objectives and a time budget. In the European academic system, this time budget is defined in terms of ECTS (European Credit Transfer and Accumulation System) credits, where one credit is often understood to be approximately 28 hours of students' work. Constructivist methods have been criticized by teachers as being too time-consuming, compared with lecturing. The timing of feedback is another old education dilemma: immediate feedback is efficient for creating an association, whereas delayed feedback leaves space for reflection. In empirical studies, the most frequent methodological bias is often the difference in 'time on task' between the control group and the experiment group. In field studies, teachers often expect to devote more of their preparation time or more of the classroom time to a curriculum item than they would to everyday teaching. This explains why many of the great experimental results are unfortunately not reproduced, when the same method is applied in everyday settings, where timing constraints cannot be waived. Is it better to have two lectures of 25 minutes on two days or 50 minutes on one day? Although neuroscience has an answer to this question, instructional sciences emphasized the difficulty of doing anything other than lecturing within 25 minutes. What are the time slots in the day in which we should place the most difficult lesson, because the kids' biological clocks predict the highest level of attention during these time slots? Although universal time results from the segmentation of astronomical cycles (Earth rotation and Earth revolution), instructional time also integrates biological and institutional constraints: energy (e.g. need to break for food), cognition (e.g. the level of attention cannot remain at maximum for a long time), teachers' workload, room availability and external economic constraints (e.g. a bachelor's degree programme should not take 10 years), etc. Looking closer at instructional time, it is indeed a very complex issue. In this chapter, we refer to the orchestration clock rather than more generally to the instructional clock. The term 'classroom orchestration' emerged in the last decade to refer to the management of pedagogical scenarios [1]. Compared with other areas in learning sciences or educational psychology, this approach emphasizes the practical issues in classroom management such as forms of organization (e.g. individual work, teamwork and class-wide activities), discipline, space, flows of data between activities and time. This chapter also considers education and online education from a very practical, organizational and even logistic viewpoint.

Time issues in MOOCs

From their outset, MOOCs introduced innovative ways to manage time. The first salient feature of MOOCs has been to segment the traditional 50 minute lectures into slices of 6–12 minutes. The length of MOOC videos is a trade-off. On one hand, human attention favours short videos: when videos are longer than 12 minutes, the median amount of the video watched is around 60% [2]. Actually,

some schools have evolved in the same way, reorganizing classroom lectures in slices of 15 minutes. On the other hand, the course content constrains this segmentation: it may be pedagogically detrimental to split some segments into different videos. We observed that teachers who transformed their 50 minute lectures into a sequence of MOOC videos had to restructure their content entirely. Segmenting lectures in short videos is a simple idea, but its effect on education is far from negligible.

A second aspect of video time is the relationship between the length of the video and the viewing time. Not only can students skip to the end of the video as mentioned in the previous paragraph, but also they can watch it several times, pause it or play it forwards or backwards. Actually, participants do not use these controls as much as we would have expected. In one of our computer science MOOCs, over 20,000 participants watched a set of 36 videos. On average, they paused only 1.4 times and moved 1.6 times 'forwards' and one time 'backwards'. The low number of pauses is probably for three reasons. First, the videos are short. If, for instance, a participant faced a difficulty when there were 40 minutes of video left, she might interrupt the video. Instead, if she sees that there is only 1 minute left to play, she may choose to let the video play. Secondly, on average, participants watch each video twice (1.93 times), probably because it is easier to play a short video again than to search for a specific point by using navigation buttons. Thirdly, videos include some automatic pauses at the points where the teacher inserted a question.

A third time feature, which received less attention, is the fact that MOOC participants may change the speed of the videos, from 0.75 times to twice the natural speed. In a lecture theatre, students do not have the ability to speed up or slow down their lecturer. Of course, participants may use the standard video controls, such as pause, forwards and backwards. In one of the MOOCs mentioned in the previous paragraph, 29% of the video sessions have been played at a speed other than the natural speed. This percentage includes video sessions during which the participant has changed the speed: on average, they do it 0.6 times per session. It also includes sessions in which the video player continues playing at the same speed as in the previous session [3]. A change of speed is related to a change in the video difficulty, as perceived by the participants. For instance, students watching the videos with the speed at ×2 perceive them on average 0.37 simpler, on a five-point scale, than people watching with the speed ×0.75.

A fourth time feature is that 'a MOOC never sleeps': whenever the number of participants reached several thousands of worldwide participants, it was always daytime for some participants. In MOOCs, the forum activity is generally continuous.

The fifth and most controversial time feature is that MOOCs initially adopted the weekly schedule of campus courses. Although learning technologies emphasized the slogan 'learn anytime anywhere' for more than a decade, the first MOOCs that emerged in 2012 had re-introduced a strict time structure: the videos were only made available on the week they were supposed to be watched, and, every week or so, participants had to upload their assignment.

This rigid time structure was probably not designed purposely, but simply originated from the fact that the first MOOCs were the online version of

existing courses. The appeal of MOOCs was to give access to courses from the best universities in the world, given by renowned teachers, and hence inherited their time structure. The rigid time structure of MOOCs is the heart of many design choices within the MOOC platform, and it has positive and negative aspects.

On the positive side, a rigid time structure constitutes some kind of anti-procrastination prosthesis: to be able to do it 'whenever I want' may end up being 'never'. One may compare this with the difference between the papers to be read some day, and the papers to be reviewed for a given deadline. Secondly, beyond the mere rigidity, this time structure induced some level of synchronicity: all participants are expected to watch the same videos and to do the same assignments during the same week. This is not a fine-grained synchronicity: it can be any time during the week, but it is not 'any time'. This synchronicity probably explains the social dynamics that characterized MOOCs and generated so much hype in the media. Being together induces a social facilitation effect: it is easier to run a marathon with 10,000 people than alone. Similarly, when a MOOC participant struggles to understand a video or to complete an assignment, she may feel better if she knows that thousands of other participants face the same difficulty. The messages in the forum are not posted by participants who did the course two years ago, but by peers currently working on their assignments since the deadline is tomorrow. This weekly wave is illustrated in Figure 1.

Now, a general rule of MOOCs is that universities lose control of who registers, when and why. The social dynamics generated by the rigid time structure are weakened by this loss of control; namely, people may register for a MOOC at any time. As illustrated by Figure 2, between 100 and 500 participants registered on this MOOC each week after the start of the MOOC, and around 100 participants even continued to register every week after it ended. However, we see in Figure 3 that those who register late have a lower rate of success, which may be due to the loss of social dynamics, but also to the difficulty of completing all assignments on time.

On the negative side, this rigid time structure probably also increases the drop-out rate. We do not want to enter into the long debate of attrition rates in MOOCs here. However, a large portion of those who register do not watch

Figure 1
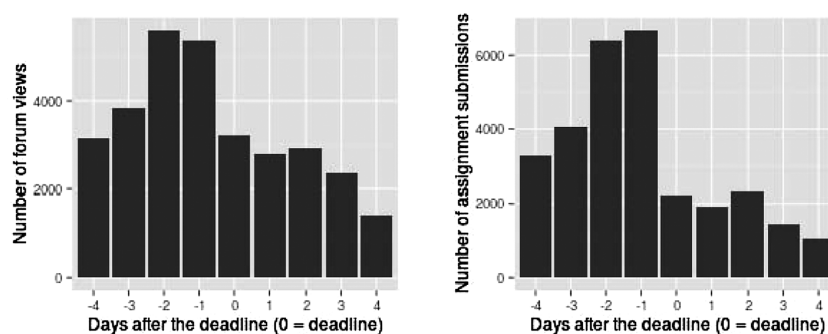
Number of forum views (left) and assignment submission (right) depending on the proximity of the deadline
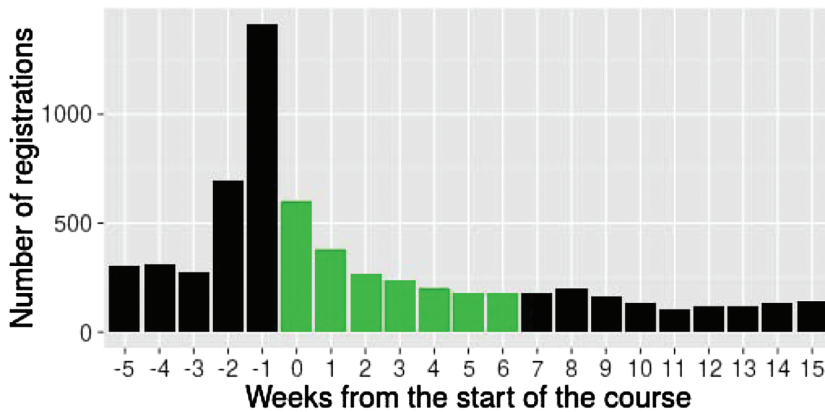*Source: data from the C++ MOOC of EPFL in 2013.*

Time of registration on the C++ course of participants who watched at least one video

The green bars are the weeks in which the MOOC was running.

any video or stop after watching a few videos. Now, among those who intend to complete the MOOC, many participants may fail to keep the pace. A majority is not enrolled in any educational programme and 82% of them have a job [4], a family or other constraints, many reasons for which it is pretty difficult to stick to a weekly schedule. It can be hypothesized that introducing more time flexibility would reduce the drop-out rate after week two. Another drawback of these session-based MOOCs is that when there are only one or two sessions of a MOOC per year, new participants may have to wait several months for the next session to start.

Recently, some MOOC platforms used a more sophisticated clock that aims to account for the problem of rigid timing: MOOC cohorts are launched on a regular time interval, such as every month. A student who fails to adhere to the schedule of his or her cohort can 'catch' the next one. This seems to be
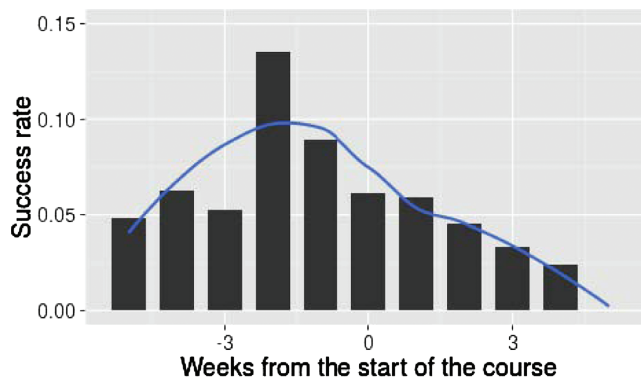
Success rate of students of the C++ course depending on the week they joined from the start of the course

The blue line describes the trend (spline fit).

a good compromise for the participants, but changes the relationship between the universities that produce the MOOCs and the platform that operates the MOOC. As mentioned earlier, initially there was a close relationship between the on-campus course and the MOOC. Typically, the university professor who manages the on-campus class accepts a few thousand free riders on board via the MOOC. The new emerging time structures, such as monthly cohorts, break this link: the MOOC is now independent of the class. Would the university teachers be amenable to a new cohort every month? Can their extraneous workload continue to be provided for nothing by their universities? Would the MOOC platform provide other human resources required to operate the MOOC iterations? These resources could be so-called 'community TAs' (students who followed the MOOC previously and volunteer to be teaching assistants) or online instructors hired by the MOOC platform. In the latter case, will the MOOC platform act as a competitor of the university that produced the MOOC? These questions illustrate that time is not a mere logistic issue; it ends up changing the relationship between universities and MOOC platforms.

The above discussion illustrates the creativity required for designing the clocks that suit both on-campus and online education. These 'clocks' refer to the computational mechanisms that we need to invent in order to efficiently run both online education and on-campus education as well as any blend of both, such as blended learning.

The orchestration clock

The timing issues result from the tension between two points in the instructional design process. The first point is the instructional material that has been designed, which has its intrinsic time requirements, such as 12 videos of 9 minutes or seven assignments of 50 minutes. The second point is the actual timing of learning activities and its variations across individual learners, such as when a video of 9 minutes is watched in 6 minutes, when the assignment due for Wednesday is delivered on Friday, or when a team member waits for her teammate to complete their part of the group assignments. The term 'orchestration' refers to the second point, i.e. the real time management of activities, with its many organizational constraints and unavoidable pitfalls. Therefore the way a MOOC scenario unfolds over time is metaphorically described as an orchestration clock. This metaphor is based on similarities, but these are rather superficial. Like a clock or a watch, the timing of an educational system is a gear train. The largest gear counts cycles (elementary, secondary, Bachelor, Masters, etc.); the next gear counts years; the next one counts quarters or semesters, then weeks, days and hours. Within a lesson, the gear has teeth or time units that are specific to the pedagogical scenario such as 'ten exercises', '40 slides' or 'three cases studies'. The metaphor is easy, but it is not really useful to simply show that instruction and clocks have nested time structures. Instead, it is more interesting if the metaphor reveals how an instructional system differs from a standard clock and hence illustrates the complexity of instructional time.

In a watch, an escapement mechanism allows a wheel to rotate by one tooth at a time, hopefully not any more. In an instruction system, the escapement

may jump several teeth (the teacher skips three exercises or five slides). In a watch, an oscillator, either a pendulum or a balance wheel, determines the beat. In instruction, the wheel may speed up, slow down, stop for a while (when the teacher asks a student to wait until the others complete the series) or even turn the other way around (when a student repeats the year). Mechanical watches are independent of each other, whereas the clocks of MOOC participants are not independent of the clocks of those who will, for instance, participate with them in the next hangout session. In a watch, time is discrete, even though many teeth of multiple gears give a feeling of continuity. In instruction, time does not really stick to the beat: it has some 'elasticity' that we will illustrate in this chapter.

Time is multi-plane

In our theory of orchestration, we describe a pedagogical scenario with graphs [5], as illustrated in Figure 4. The vertical axis of the graph shows learning activities. The horizontal axis represents time. The horizontal position on the left end of an activity is its starting time, and the length of the box that represents an activity corresponds to its duration. The horizontal distance between two successive activities represents the time lag between two activities, for instance a pause or the interval between two weekly classes. The vertical axis is discrete; it is decomposed into several social planes: the lowest plane (π_1) describes individual activities, the second plane (π_2) refers to activities in teams, and the upper plane (π_3) refers to class-wide activities. In a MOOC, class-wide activities are those which engage all active participants. (NB: the modelling language includes three higher planes that are not described here.)

The edges between activities are associated with workflow operators. Let us consider a science lesson based on discovery learning, inspired by Hannie Gijlers and Ton de Jong [6]. After a short collective introduction (a_1, π_3), the teacher asks students to individually write (a_2, π_1) a hypothesis regarding the relationship between pressure and heat. They enter this hypothesis in the system, which
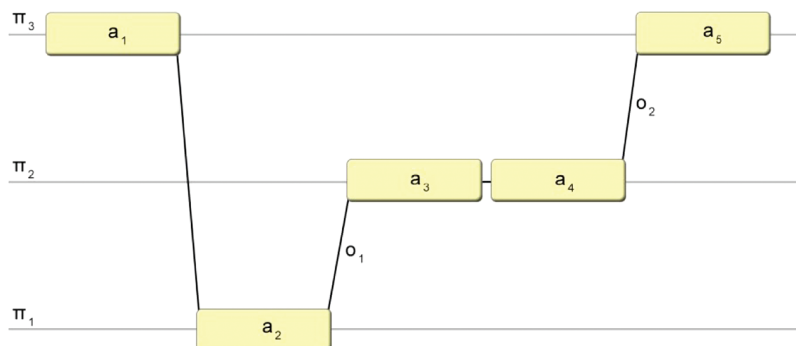

Figure 4

Modelling a pedagogical scenario with an orchestration graph
Horizontal lines respectively correspond to individual ($\pi 1$), team ($\pi 2$) and class ($\pi 3$) planes, boxes correspond to activities and black lines correspond to operators linking the activities.

automatically forms (o_1) pairs of students having expressed divergent hypotheses. The pairs have to design an experiment that could determine which hypothesis is right (a_3, π_2). Thereafter, they use a simulation to run the experiment and save their results (a_4, π_2). Finally, all results are aggregated (o_2) and displayed through means of a visualization tool that the teacher uses in her synthesis lecture (a_5, π_3).

The time structure of this graph includes two synchronization points: all individuals (π_1) must enter their hypothesis before applying o_1 for forming pairs (π_2); all pairs must complete their experiments before applying o_2 for aggregating results and using them at π_3. In a classroom or in cohort-based MOOCs, this synchronicity constraint is managed by setting up deadlines and eventually neglecting those who did not meet this deadline. This is fine if the number is low and if the operators have been implemented robust to missing data. In a self-paced MOOC, where any participant may start the MOOC whenever she wants and do it at the pace she wants, these constraints would require new timing operators such as 'pause the scenario until the participants' set includes someone with opposite opinions' (o_1) or 'pause the scenario until the participants have produced enough results to produce meaningful visualizations' (o_2). Such operators will be necessary to apply peer grading methods in a MOOC, where participants may upload their assignments whenever they want.

In other words, the orchestration of rich pedagogical scenarios requires the development of new clocks in which every plane has its own flexible timing, partly independent of the other planes, but punctuated by co-ordination constraints.

Time is a random variable

The fundamental factor that explains the complexity of the orchestration clock is that the duration of a learning activity is a random variable. In mathematics, this term does not describe a purely random element, but a variable for which the values follow a probability distribution: for instance, the height of a male adult has some probability to be 160 cm, a higher probability to be 175 cm and a low probability to be 225 cm.

As an example, let us consider an experienced designer who believes that an activity a_i, such as reading a text, will last 10 minutes. This actually means that he expects that, for instance, 90% of students will have completed a_i within 10 minutes. If, after 10 minutes, the teacher decides to move to a_{i+1} without letting the late students finish a_i, 10% of the students will be penalized by having to do a_{i+1} without having completed a_i. However, if the teacher decides to give more time and to close a_i after 15 minutes, it would mean that 90% of the students wait for their peers for 5 minutes. This does not only constitute a waste of time, but, as all teachers know, it is not good to let students do nothing for 5 minutes. This dilemma can be reformulated as finding the optimal point in time (t_{opt}) which, on one hand, minimizes the function $U_1(t)$, the number of students who did not finish the activity, and, on the other hand, minimizes $U_2(t)$, the sum of time lost by students who finish early. Often one of these two objectives is more important than the other, which can be expressed by introducing weights $\beta_1 > 0$ and $\beta_2 > 0$ of objectives U_1 and U_2 respectively. In some situations, for example when we teach how to use a parachute, finishing the activity is crucial. Then we will have $\beta_2 > \beta_1$

(to maximize the number of students who succeeded). In other situations, such as when the activity is just a simple short exercise which serves more as an example, we put $\beta_2 < \beta_1$ (to minimize the time lost). In any case, the optimal time is given by

$$t_{opt} = \text{avg min } \beta_1 U_2(t) + \beta_2 U_2(t); t > 0 \qquad (1)$$

    To illustrate this phenomenon, we use an example from the 'Introduction to Visual Computing' course, given at the EPFL in 2015 in a flat lecturing room. Some lectures included activities where students could experience by themselves phenomena such as perceptive biases and cognitive overload. Typically, the lectures were interrupted for n minutes and the students individually performed the online activity on their laptops. The value of n was not fixed, but regulated by the teacher who monitored online the number of students who completed the task. We considered the data from one of these activities and aimed to experience different interaction styles such as language commands, direct manipulation or forms. The 106 students had to do the same task, ordering five train tickets, with each of four interfaces presented in random order.

    Completion times of all students were recorded. Figure 5 presents the number of students who did not finish until t, together with the cumulative time lost by students who already finished. We assume linearity of the lost time; for example, 5 minutes of time of 60 people is equally valuable as 5 hours of one person. Under our assumptions, the optimal time for the activity is 26 minutes.

    As a second example, we do not consider the duration of an activity, but instead focus on the time lag between two activities. Weights introduced in equation (1) correspond to a parameter introduced in the modelling language:

Figure 5
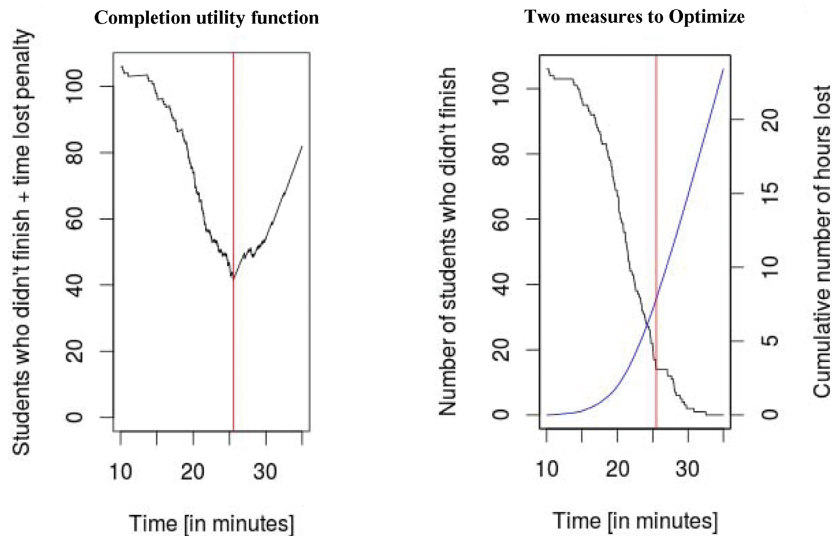
Utility function to optimize (left) and utility decomposed to number of students who did not finish and cumulative number of hours lost (right)
*The minimum of the sum is denoted by the vertical line.*

the strength of the link between two activities. Often this link is time-dependent on its own, which brings us to the concept of the elasticity [5]. For example, the importance of a calculus course as a prerequisite of the probability course may fade out with time, or motivational tricks may be beneficial only within a very short period of time.

Consider two activities: learning neural networks within a course on machine learning and applying these networks for image classification in a computer vision class. Although a student can easily recall the concept just after the machine learning course, they may not be able yet to apply this knowledge in practice. On the other hand, if the lag between the two activities is too long, the student may not be able to recall the method. Ideal timing is somewhere in the middle; so again it is given by t_{opt} = avg min$_{t>0}$U(t), where U is a convex function utility function describing how valuable the second activity is at particular time t. Note that U depends not only on the activities, but also on the personal characteristics of a student, class and other activities; so it should be optimized for each student separately.

The optimal time T is a random variable. The more we know about the student, the better we can assess the optimal time. Using the notion of conditional expectation from the probability theory, we can assess the optimal time better by adding additional information. Suppose T has some unknown distribution with the mean 10 minutes and the standard deviation 3 minutes. If we know that the student is anxious, we may be able to assess the value better. For example, let ($T \mid$ *anxious*), time when the student is anxious, be a Gaussian distribution with mean 12 minutes and standard deviation 1 minute and let ($T \mid$ *confident*), time when the student is not anxious, be a Gaussian distribution with mean 8 minutes and standard deviation 1 minute. Information about the anxiety may allow us to further reduce the variance, and make our estimators more accurate, as illustrated in a simulation depicted in Figure 6.

Figure 6

Distribution of activity time
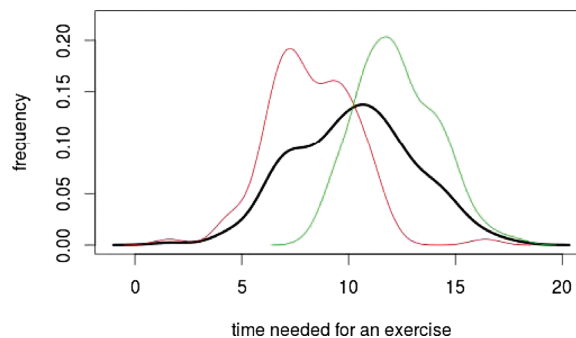
Simulation of the distribution of activity time
Even if the distribution in the entire class may resemble a Gaussian distribution (black), we can assess it more precisely if we know more about the students, for example, if they are confident (red) or anxious (green). The data for this example were simulated by sampling from two Gaussian distributions.

This example requires either very precise assessment from the personality questionnaire in a MOOC context or an attentive teacher who knows each individual in the class well. However, the technique can also be found useful even if we have less granular information. We expect, for example, smaller variance of the required time on activity in a group of professionals learning in a management class than in a primary school class.

Time is relative

We take the measure of time for granted, even though it is partly arbitrary. If the notion of day is related to the Earth's rotation time, its division in hours, minutes and seconds is rather subjective. We are used to taking these arbitrary units as fixed and treat other values, e.g. the learning gain, video duration and course length, as derivatives of time. Although this approach makes perfect sense in the social context where time is fixed between two individuals (teacher devoting one hour of his time to a student), this normalization is not always needed in massive and open educational contexts. What if, instead, we do not consider absolute time, but some relative time? In orchestration graphs, it may be enough to know which activities have been finished. Similarly, to define students' time in a video, we may use the relative position within the video, instead of a number of seconds. This is useful when analysing complex video navigation paths in which the students watch a 6 minute video for 6 minutes, but actually viewed the first two minutes three times.

Therefore, instead of referring to universal time, the orchestration clock defines time as a multi-dimensional variable, depending on the position of the student on different planes. Imagine that a course consists of certain collaborative and individual activities. For each student s we may define his current progress in the course as $(t^{(s)}_i ; t^{(s)}_t ; t^{(s)}_c)$, where coefficients correspond to the progress in individual tasks, team tasks and class tasks.

The above definition simplifies the notion of synchronicity. To define a synchronicity index, we introduce $S(t) = \exp[-\sigma_c(t) - \sigma_i(t) - \sigma_t(t)]$, where each σ is the standard deviation of time used by students in the progress t. Note that each σ is non-negative, so $S \in [0,1]$. We will have $S = 1$ when students are fully synchronized and small values S (still greater than 0) if they differ much on any of the temporal dimensions. A similar idea can be applied to activities: when designing classroom activities, one should try to maximize the synchronicity in order to reduce the time lost by students who finish early. In Figure 7, we present synchronicity for the individual dimension as a function of progress for one of the sub-activities in our train ticket experiment described above. Thereby, our definition of synchronicity allows us to compare four subtasks even with different expected durations.

In a classroom, synchronicity is managed by the teacher and it relates to the wall clock. At universities, synchronicity is imposed by constraints put on students and courses; one cannot take a course without finishing prerequisites. In any context, these synchronicity requirements can be imposed directly on the flow, without referring to the absolute time.

Figure 7
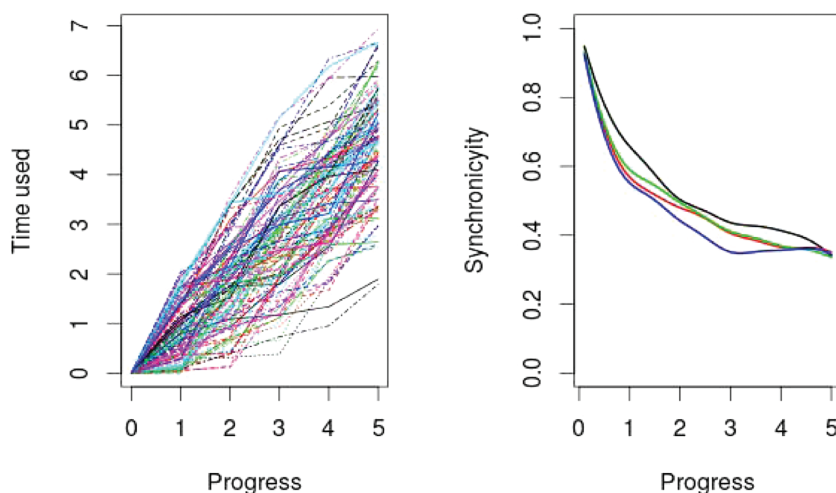
Time spent on ordering virtual tickets in a graphical interface (left) and synchronicity as a function of progress for four different interfaces (right)
Left: time in minutes. Right: drag and drop (black), form (red), command line (green) and graphical (blue).

Conclusions

The timing of learning activities is a key issue in any educational situation, and it is even more critical in MOOCs, where different time structures have affordances and constraints. Any course or curriculum may be structured according to various 'time contracts', from the traditional '2 hours per week' to the total absence of time structure, as would be 'learning from Wikipedia'. Actually, MOOCs may be located at any point of this curriculum. Since physical constraints are partly waived, the choice of a time structure for MOOCs depends on pedagogical constraints. In other words, the time structure of a MOOC can be designed rather than determined by the infrastructure. We have seen that the time for completing a task follows a probabilistic distribution. This observation and common sense suggest maximizing the learner's freedom on time management. Anyway, the teacher has very low control over the way participants do the MOOC activities. However, there are also pedagogical reasons to force some synchronicity, for creating social dynamics, or some synchronization points between the clocks of each social plane. These synchronization points fortunately have some flexibility such as 'the four team members should have submitted their individual work before starting teamwork, but it is acceptable to start if one is missing for a maximum of 3 days'. Some MOOCs indeed discriminate 'soft deadlines' and 'hard deadlines'. Some MOOCs run a new cohort every week, i.e. they combine a rigid clock with the freedom to jump on another clock. We expect that, in the coming years, novel time structures will emerge in online education as well as in blended education. It may be that this creativity will propagate to traditional schools. In this chapter, we have shared some of our preliminary reflections, but our intuition is that some deeper modelling of this clock should receive attention in learning sciences.

Disclaimer

The views expressed in this article are purely those of the authors and should not be regarded as the official position of the European Commission.

References

1. Dillenbourg P. and Jermann, P. (2010) *Technology for Classroom Orchestration: New Science of Learning*. Springer, New York
2. Guo, P.J., Kim, J. and Rubin, R. (2014) How video production affects student engagement: An empirical study of MOOC videos. *Proceedings of the First ACM Conference on Learning @ Scale (L@S '14)*, ACM, New York
3. Li, N., Kidziński, Ł., Jermann, P. and Dillenbourg, P. (2015) How do in-video interactions reflect perceived video difficulty? *Proceedings of the European MOOCs Stakeholder Summit 2015*, PAU Education, Barcelona
4. Jermann, P. and Dillenbourg. P. (2015) *MOOC Report 2012–2014*, EPFL, Lausanne
5. Dillenbourg, P. (2015) *Orchestration Graphs*, EPFL, Lausanne
6. Gijlers, H. and de Jong, T. (2005) The relation between prior knowledge and students' collaborative discovery learning processes. *Journal of Research in Science Teaching* **42**, 264–282