Encouraging knowledge exchange in discussion forums by market-oriented mechanisms

Manuel Núñez, Ismael Rodríguez and Fernando Rubio

Dept. Sistemas Informáticos y Programación
Universidad Complutense de Madrid, E-28040 Madrid. Spain.
e-mail: {mn,isrodrig,fernando}@sip.ucm.es

Abstract. Discussion forums are one of the simplest, while successful, Internet-based systems to promote the spread of knowledge. Unfortunately, they sometimes lack appropriate incentives to encourage users participation. In this paper we discuss a simple market-oriented mechanism to promote the knowledge exchange activity in discussion forums. The proposed solution departs from other market-oriented approaches since the scarce resource we deal with, the *effort* to create information, has important peculiarities which must be taken into account. In addition to the basic framework, we have also developed a method to dynamically improve the performance of the system. This method uses statistics about users behavior in order to minimize the response time in which questions are satisfactorily answered.

1 Introduction

Knowledge is one of the most valuable goods of our society. Fortunately, the new technologies are currently playing a key aspect in the efficient dissemination of knowledge. On the one hand, Internet has allowed the spread of knowledge without frontiers. On the other hand, a new generation of interactive intelligent tutors (see e.g. [3, 1, 8, 9, 2]) provide personalized learning assistance to users. In this context, following the definition of the eLearning Action Plan of the European Union (March 2001), e-learning is "*the use of new multimedia technologies and the Internet to improve the quality of learning by facilitating access to resources and services as well as remote exchanges and collaboration.*"

One of the simplest and most classical, but also extremely successful, electronic mechanisms used for e-learning are *discussion forums* and *news groups*. These applications allow the holders of knowledge about a specific topic to share their knowledge with the rest of interested people, so that a bigger community can profit. Thus, thanks to the collaborative action of many users, doubts and questions can be solved by their mates, as the more experienced users help beginners by transmitting them their knowledge.

However, this simple mechanism of cooperation presents two main problems. First, the incentives to share information by answering questions of other users can be weak. In most cases, the altruism or the desire of prestige can be the main motivation to help beginners. Unfortunately, such motives could be not shared by many users. Thus, the collaboration degree can be quite asymmetric. Second, experts can lose their motivation to help others as they can get saturated by a huge amount of questions. Even though most of the problems could be solved by not so skilled users, more experience ones receive all the questions.

In fact, the previous problems appear due to the primitive structure underlying current discussion forums. Actually, they do not usually promote the efficient distribution of information. However, the key to understand the problem is to consider information as a *scarce resource*. More precisely, if we consider the cost of the generation of information in terms of the time needed to create it, then it becomes clear that such activity must be considered a scarce resource. Obviously, the costs to copy, store or transmit information are negligible compared to the cost of generating it. So, we will only consider the aspects related to the generation of information.

As the efficiency in the generation and use of information in a discussion forum can be defined in terms of optimizing the access to a scarce resource, there exists a science that can help us dealing with it. Quoting from [5], "*Economy is the science that deals with distributions of scarce resources which can be valued by having different utilities by different users*." In our concrete problem, the effort for generating information is clearly a scarce resource, and the utility that each user of the system gives to each piece of information is also clearly different. As an example, let us consider two different questions performed by two different users. The first of them is more complex than the second one. In fact, only experts can answer such a question, while the second one can be answered by a user with only some command in the topic. In that situation, it is clearly inefficient to ask an expert to answer the easiest question and to assign the hardest question to a user with only some command on the corresponding topic.

The most usual method in economic systems to promote the efficient utilization of resources consists in assigning *prices* to them. These prices mark the effort level that an interested agent should apply to obtain a certain resource. In case a resource is very scarce (supply much lower than demand) the price will raise, so that only the most interested agents will obtain the desired resource. Moreover, a high price in a resource encourages producers of the resource to increase its production. Thus, the difference in prices among the different resources represents the real necessity of each product as well as promotes the production of the most estimated items. The most common unit to measure the effort level to be paid is *money*. However, in order to promote an efficient access to resources it is not necessary to use *real* money: It can be virtual money provided that its quantity is scarce. In fact, several market oriented solutions have already been used to regulate the access to information in databases or the processing time in distributed systems (see e.g. [7, 6, 4]).

We claim that a market-oriented approach could also be applied to the control of information generation in a discussion forum. In this way, each user should have to somehow *pay* other users in order to get her doubts solved. However, a direct application of the market oriented concepts could not be completely satisfactory, as e-learning environments need to consider additional factors not appearing in other computational systems. In this paper we deal with these important differences and propose a suitable market-oriented methodology to optimize the use of the educational effort in discussion forums. Besides, and taking as starting point performance information collected from using [2], we propose a dynamic mechanism to improve the overall performance of the system according to the statistics collected from the behavior of the users. Our experience indicates that discussion forums may benefit from the technique developed in this paper, yielding systems where the overall time in which questions are satisfactorily answered will be reduced just by increasing the efficiency of the educational efforts of users.

2 A Market-oriented Methodology for Discussion Forums

In this section we describe how a market-oriented methodology can be applied to optimize the performance in a discussion forum. As we said in the introduction, e-learning environments present relevant differences with respect to other computational environments. These peculiarities must be taken into account in order to successfully apply a market-oriented approach.

2.1 Rewarding good Answers

First of all, let us remark that the system will not be able to decide, in general, which user is the best choice to answer a concrete question. To overcome this difficulty we can *classify* users based on the *quality* of their effort generating answers. By using classes of users at different levels, the system can determine the suitability of a given user to answer a question of another given user. Let us remark that the quality of the effort of the generation of information does not only depend on the knowledge of the users, but also on their willingness to transmit their knowledge.

It is important to point out that as we are dealing with knowledge, it is impossible for a user to know with full certainty whether her question requires the help of an expert, or only that of a not so skilled user. Thus, not much improvement of the efficiency could be obtained by allowing the questioner to choose the type of user that should answer her doubt. On the contrary, the system must provide a questions distribution mechanism to decides which type of users should answer each type of questions.

Another important point is the fact that the quality of the resources, that is the quality of the received answers, is not known a priori. In other words, a user asking a question does not know how good the answer is until she obtains it. In fact, she can get disappointed with an incomplete/erroneous/bad answer. Analogously, an user answering a question cannot know the quality of her answer until the other user studies it. Thus, it is not possible to set the price of an enquiry a priori. Actually, the receiver of the resource should be the one who fixes the price she will pay. Obviously, if money is the scarce resource that enables access to information then the buyer of the resource will try to pay as few as possible, arguing that she was not satisfied with the answer. This problem can be overcome if the paid amount of *money* does not perceptibly modify the future capacity to *buy* new knowledge. In order to do so, the units the users make use of will be potentially *inexhaustible*, as we will comment below.

The main reason for including prices in our approach is to give incentives to the providers of knowledge. Thus, when a user gets satisfied with an answer then the provider of the answer should receive an incentive. As the only resource available in the system is educational effort, such incentive should be given in terms of it. Due to the users will not be allowed to choose the type of users to answer their doubts, the system will have to appropriately award incentives to users, both by providing and by restricting the access to educational effort. Thus, if a user provides answers that are positively perceived then:

- The system will try to help that user to get her future doubts solved. As she has proved that she has valuable knowledge, her future questions will be shown to more expert users. Thus, she will subsequently obtain satisfactory answers with a higher probability.
- As she has proved that she has valuable knowledge, the system will also try not to show her easy questions in the future. By doing so, she will be able to save her effort for those questions that really require her skills.

In order to obtain such behavior each time a user posts a new question, it will be initially shown only to the users belonging to her own class. If after a period of time (depending on the class) the question is still unanswered then it will be forwarded to the next level. If the question is not answered there, it will be forwarded step by step until reaching the highest level. So, easy questions will be usually answered before they reach the expert level. This is so because, as we remarked before, the user providing an answer to a question may improve her *reputation* by getting into a *better* group. By using this mechanism we encourage that the questions are answered by those users that are the most appropriate to do it: Users with medium command answer the questions of beginners while experts answer both the questions of the medium users and the ones from beginners which could not be solved by medium users after the corresponding delay. Therefore, the scarce resources of the system (that is, the effort to create valuable information) are exploited in an efficient way, keeping the most valuable resources for the specific questions that really need them.

Let us remark that the main underlying idea consists in restricting the access to educational effort while not restricting the access to the information itself. In fact, anybody will be able to access the full repository of questions and answers. Actually, the only way to gain points in the system will not be to answer specific questions of other users. In addition, any other medium to spread knowledge will be available to be used and acknowledged. For example, if a user creates a tutorial and posts it publicly in the forum, then any user who considers that it is valuable will be able to give points to the author in order to acknowledge her work.

Finally, let us comment that we have already identified some situations that may harm the utility of our methodology. Given the space limitations, we are not able to include our *risk assessment* in the bulk of the paper. However, it may be found as an appendix so that reviewers can evaluate it. In the case of acceptance, the final version of the paper will contain a url where this information can be found.

2.2 Structuring users in classes

Users will be placed in different classes. Each group will be a set of sorted users (according to the amount of gained points). So, the set of *guru* users is made of the u_1 best users of the ranking, the set of *expert* users is made of the next u_2 users, and so on. As it usually happens in knowledge communities, the amounts of users in each class should follow a pyramidal structure, so that the condition $u_1 < u_2 < \ldots < u_n$ is surely a must. Besides, let us remark that by structuring classes through the ranking system a user can, after a while, either improve or fall in the hierarchy. For instance, in the case that a user reduces her activity, she will be overtaken by other more active users. Therefore, even if the points each user owns to acknowledge other users are inexhaustible, these points can actually affect her. Let us remark that, however, the effect of acknowledging others has a minimal repercussion if we consider them *individually*. In a system with thousands of users, giving points to some other specific user will not have, in general, a perceptible effect on the ranking of the specific user who gives them. Nevertheless, the result of the simultaneous delivering of points of all the users of the system yields a competitive environment as a whole. Let us remark that by using this method the incentive users can have to reduce the points they give to others is very low. The reason is that what can actually constrain her access to the resources is mainly the activity of *others*. This is specially clear if we compare this method with a method where the units each user delivers are *exhaustible* and each user is *forced* to pay fixed amounts of units before accessing the information provided by another user.

3 Dynamic Behavior of the System

In this section we study how our system can be dynamically adapted to improve its efficiency by taking into account the behavior of the users. More precisely, the system will automatically adjust certain parameters in order to minimize the mean time required to satisfactorily solve the questions. Moreover, it will also try to minimize the distances between the response time of the users of different classes, so that the performance is equitable among all of them.

The main parameters to be dynamically adjusted are the corresponding *waiting time*, that is, the time questions need to stay unsolved in a level before forwarding them to an upper level. We will denote by t_{ij} the time that questions of users of level i remain in level j before forwarding them to level $j + 1$. Let us remark that these values should be neither *large* nor *short*. On the one hand, if they are very large the response time will unnecessarily increase: If a question is not solved in a reasonable time, it will probably not be solved in any time.

On the other hand, if the time is very short then the probability to receive an answer is very small. Thus, upper levels will receive more questions that needed. In fact, the frequency of questions can saturate the upper levels, reducing the overall efficiency of the system. Let us remark that in such case, the values for response time will be also dramatically increased. In the rest of the paper we show how appropriate *intermediate* values can be formally computed.

3.1 An Algorithm to adapt Waiting Time Values

The factors we will use to optimize the values for appropriate waiting times will be controlled by users statistics. These data can be automatically obtained by the system without affecting the normal performance of the application. Let us suppose that there are n levels of users in the system. Then, for $1 \leq i \leq n$ we denote by f_i the frequency of generation of questions of the users belonging to level i. Besides, we will use a set of random variables to codify the time needed for questions to be correctly answered in a given level. Thus, the discrete random variable ξ_{ijk} describes the time required for a question coming from users of level i to be solved by users of level j, assuming that the frequency of questions arriving at level j (both from level j and other previous levels) is k. Let us remark that the frequency of incoming questions influences the response time. This is so because a high value implies that class j users will waste a lot of time just reading the upcoming questions. So, their efficiency answering questions will decay. The random variable ξ_{ijk} will not take into account the time elapsed from the arrival of the question to the moment in which it is forwarded to the level $j+1$ (that is, t_{ij}). On the contrary, ξ_{ijk} will contain the response time assuming that questions are not forwarded, that is, supposing that questions stay at level j indefinitely until they are answered. The system can infer these random variables by recording the amount of questions of level i users that level j users answer for each possible time $t \in \mathbb{N}$, assuming that the frequency of input questions is k. Obviously, the system will not be able to measure all possible times. In particular, it will not be able to measure times larger than t_{ij}, because in that moment the questions will be moved to level $j+1$. However, the probabilities for times larger than t_{ij} can be easily extrapolated. Taking into account that this kind of events can be accurately approximated by using exponential distributions, or its discrete counterpart Poisson distributions, the extrapolation can be done as follows: We use the recorded samples of the system to calculate the probability that a question coming from level i is transferred from level j to level $j+1$ because time t_{ij} runs out. Then, we distribute this probability among times higher than t_{ij} so that they extend the shape of the exponential distribution among samples taken for times lower than t_{ij}. Let us note that we also need to perform another kind of extrapolation. This is so because we need to generate random variables for different frequencies of arrival of questions. Let us take into account that the variability of frequencies during the performance of the system will allow us to observe (part of) the influence of the frequencies on the values for response time. Therefore, those random variables that have not been sampled will be extrapolated from those that have been actually observed.

In this case, we can also assume that the effect of the frequency on the response times is also exponential, so extrapolations will be easily developed. We will denote by p_{ijkt} the probability of the random variable ξ_{ijk} to take a value less than or equal to t, that is, p_{ijkt} will represent the probability that t units of time have passed. So, the following equality holds:

$$p_{ijkt} = Prob(\xi_{ijk} \leq t) = \int_0^t f(x) \cdot dx$$

where f is the probability density function associated with ξ_{ijk}. Let us note that if we consider Poisson distributions then the previous integral has to be substituted by the appropriate sum. In the following we will denote by p_{ijk} the probability of answering in level j a question initially placed at level i, assuming an input frequency k. That is p_{ijk} stands for $p_{ijkt_{ij}}$.

By considering the previous random variables, the real frequency of questions arriving at each level can be computed. Basically, the level i will receive those questions generated at that level as well as those questions coming from lower levels that have not been successfully answered in the previous levels. Thus, if we denote by g_i the frequency of questions arriving at level i then we have

$$g_i = \sum_{1 \leq j \leq i} f_j \cdot \prod_{j \leq k < i} (1 - p_{jkg_k})$$

Let us note that g_i depends on other values g_k, so that it is actually defined in a recursive fashion. However, the recursion is obviously well founded because when g_i depends on g_k, then $k < i$. From now on, the special value p_{jkg_k} will be simply denoted by p_{jk}.

In order to be able to compute the average waiting time for questions coming from level i users, we must take into account that some questions may not be answered even in the highest level. Let us remark that we have assumed that the random variables describing the response times are exponential, so the probability that a question is answered is greater than 0 for any time $t > 0$. Nevertheless, it is more practical to consider that there is a maximal waiting time so that after it runs out the question will be supposed to be unanswered. When the mean response time is calculated, the fact that a question has not been satisfactorily answered must affect it in terms of a high penalty time. This penalty will be added to the time already consumed by the question to range all the levels until the last one (that is, the time the user needs to know that her question will not be answered). The penalty will depend on the relevance we consider the unanswered questions must have in the overall performance of the system. We will denote this penalty time by t_{pen}. For the sake of clarity, we will codify the effect of this penalty in an homogeneous way by adding an artificial highest level $n+1$. Then, we will have that $f_{n+1} = 0$, $t_{i\ (n+1)}$ is irrelevant for any i, and the random variables of response time of level $n+1$ are defined such that the time consumed is always equal to t_{pen}, that is,

$$\forall i, k, t \neq t_{pen} : Prob(\xi_{i\ (n+1)\ k} = t) = 0 \ \wedge \ \forall i, k : Prob(\xi_{i\ (n+1)\ k} = t_{pen}) = 1$$

Then, *obtaining* the answer to a question in level $n+1$ represents that the maximal timeout run out, meaning that the question remains unanswered.

We will denote the mean response time for level i users by T_i. This measure can be calculated by adding the mean times we would have if the answers were obtained at each level from i to $n+1$, weighting this time with the probability that the question is satisfactorily answered at this level. Thus,

$$T_i = \sum_{i \leq j \leq n+1} \left(p_{ij} \cdot \prod_{i \leq k < j} (1 - p_{ik}) \cdot (\mu_{ij} + \sum_{i \leq k < j} t_{ik}) \right)$$

where $\mu_{ij} = \frac{1}{p_{ij}} \cdot \int_0^{t_{ij}} x \cdot f(x) \cdot dx$, being f the probability density function associated with the random variable ξ_{ijg_j}. Again, if Poisson distributions are considered then the integral appearing in the previous formula has to be replaced by a sum. Assuming that the number of level i users is equal to u_i, we can easily calculate the mean response time, denoted by μ, and the variance for time existing among the users of the different levels, denoted by σ^2, as

$$\mu = \frac{\sum_{1 \leq i \leq n} u_i \cdot T_i}{\sum_{1 \leq i \leq n} u_i} \qquad \sigma^2 = \frac{\sum_{1 \leq i \leq n} u_i \cdot (T_i - \mu)^2}{\sum_{1 \leq i \leq n} u_i}$$

Therefore, a measure that takes into account both the mean response time and the equity among the response times of the different classes is:

$$K = \alpha_1 \cdot \mu + \alpha_2 \cdot \sigma^2$$

where α_1 y α_2 are set to represent, respectively, the relative relevance we give to the mean time and to variability in our system.

As a conclusion, let us remark that the problem of optimizing the discussion forum can be formalized as the search of the values t_{ij} such that K is minimal.

4 Conclusions

In this paper we have presented a market-oriented methodology to improve the performance of discussion forums. The idea is that questions should be answered by the less qualified users who are still able to solve the problem. Thus, experienced users will be answering only the really difficult questions. In order to achieve this objective, users are placed in groups according to the capabilities that they have shown by answering other users questions. If a user has a new question then she initially places it in her group. If after a certain amount of time nobody answers the question then the question is moved to a higher group. Let us remind that the membership to a given group may dynamically vary. Users will change their group according to subsequent performance by taking into account the perceived quality of their answers. Finally, in order to guarantee a good performance of the system, the before mentioned maximum waiting time values have to be adjusted. We have provided an algorithm that, by using previously observed performance of the system, computes optimal values for these delays.

References

1. P. Brusilovsky, E. Schwarz, and G. Weber. ELM-ART: An intelligent tutoring system on the World Wide Web. In *3rd Int. Conf. on Intelligent Tutoring Systems, LNCS 1086*, pages 261–269. Springer, 1996.
2. N. López, M. Núñez, I. Rodríguez, and F. Rubio. Including malicious agents into a collaborative learning environment. In *6th Int. Conf. on Intelligent Tutoring Systems, LNCS 2363*, pages 51–60. Springer, 2002.
3. G.I. McCalla. The search for adaptability, flexibility, and individualization: Approaches to curriculum in intelligent tutoring systems. In M. Jones and P.H. Winne, editors, *Foundations and Frontiers of Adaptive Learning Environments*, pages 91–122. Springer, 1992.
4. M.S. Miller and K.E. Drexler. Markets and computation: Agoric open systems. Available at: `http://www.agorics.com/agoricpapers.html`, 2000.
5. L.C. Robbins. *An Essay on the Nature and Significance of Economic Science*. Macmillan, 1984, originally published in 1932.
6. M. Stonebraker, P.M. Aoki, W. Litwin, A. Pfeffer, A. Sah, J. Sidell, C. Staelin, and A. Yu. MARIPOSA: a wide-area distributed database system. *The VLDB Journal*, 5:48–63, 1996.
7. C.A. Waldspurger, T. Hogg, B. Huberman, J.O. Kephart, and S. Stornetta. Spawn: A distributed computational economy. *IEEE Trans. Soft. Eng.*, 18:103–117, 1992.
8. G. Weber. Adaptive learning systems in the World Wide Web. In *7th Int. Conf. on User Modelling, UM'99*, pages 371–378. Springer, 1999.
9. B.P. Woolf, J. Beck, C. Elliot, and M. Stern. Growth and maturity of intelligent tutoring systems: A status report. In K.D. Forbus and P.J. Feltovich, editors, *Smart Machines in Education*, chapter 4. AAAI Press/The MIT Press, 2001.

Appendix: Risk Assessment

The relative freedom users have to deliver their acknowledgement points yields some specific risks that must be commented. First, even though the points a user gives to others do not strongly affect her hierarchical status, a user could systematically give low marks to the answers she receives. Second, a set of users could make a deal so that they give themselves good marks just to grow up in the hierarchy. In the extreme case, a user could even use two different nicks in the system so that each one gives points to the other. Third, the preference of a user to answer the questions of some other specific user could dramatically depend on the points given by this user in previous questions. These problems that can potentially ruin the application of our methodology can be solved in the following way:

- Let us suppose that a given user a receives good marks from other users. Besides a user b, who gives bad marks to all the users, gives a also a bad mark. Then, b is supposed to be underestimating all the marks she gives. In this case, all the marks of b be can be raised automatically by the system. The case is opposite if a uses to receive low marks but b gives always high marks.

– In the long run, it is highly improbable that the questions of a user are *always* solved either with extremely high or extremely low quality. So, in the case that these situations are detected, all the marks of the user can be automatically lowered or raised, respectively.
– Let us suppose that it is detected a set of users such that points flow intensively inside it but rarely outside. Then, this set of users will be supposed to be interested in a *subtopic* inside the topic of the discussion forum. In this case, the system will create an alternative hierarchy for their members, so that the points gained inside the set will influence both hierarchies. However, the effect of these *internal* points in the overall hierarchy will be reduced. In fact, the reduction will be higher as the flow of points outside of the set is reduced. In the case that there is no flow at all outside the set, the forum will be split into two independent forums. As a side effect, the system will encourage that users enlarge the sets of users to whom they answer questions. So, variety will be required by the system to avoid being considered part of an isolated interest group.
– When a user answers a question, the amount of points received will not be reported. Actually, the gained points will only be added to her score after she has answered a fix number of questions. In this way, a user will not be able to detect how many points she received from each specific user.

Another possible source of wrong behavior may appear because, according to the rules of the system, users could benefit from acting as *intermediaries*. Let us remark that, as we will briefly describe, this behavior does not necessarily harm the system. Let us suppose that a beginner formulates a question. As she is not familiar with the specific words and the notation of the topic, her question is difficult to read and understand. This question reaches the medium level and then a medium user reads the question. After some effort, she understands the question but she does not know the answer. Then, she rewrites it by using a correct and advanced notation. Afterwards, she posts the question in the forum as her *own* question. The question reaches the expert level where an expert answers it by using an advanced notation. Then the medium user, with some effort, rewrites the answer in an easy notation and sends it to the beginner. The expert is paid by the medium-level user while this one is paid by the beginner. Clearly, the role of the intermediate is useful, as the educational effort has been distributed in an efficient way, avoiding the expert user to waste her time in *translating* the question of the beginner. So, it is fair that the medium user gains points in the process. Let us remark that in the case that an intermediary does nothing and forwards the question as it is, then the beginner will not pay her, as she can access the answers of all levels and so she can notice that the intermediary did nothing to help her.