Interactive Learning of Algorithms

Jarmo Rantakokko

Scientific Computing, Department of Information Technology
Uppsala University, Uppsala, Sweden

May 28, 2004

Abstract

Visualization is believed to be an effective technique for learning and understanding algorithms in traditional computer science. In this paper, we focus on parallel computing and algorithms. An inherent difficulty with parallel programming is that it requires synchronization and coordination of the concurrent activities. We want to use visualization to help students to understand how the processors work together in an algorithm and how they interact through communication. To conceptualize this we have used two different visualization techniques, computer animations and role plays. As the students can *see* how the processors run simultaneously in parallel, it illustrates important concepts such as processor load balance, serialization bottlenecks, synchronization and communication. The results show that both animations and role plays are better for learning and understanding algorithms than the textbook.

1 Introduction

The motivation and expectations for using visualization of algorithms in the education is that the students will grasp and understand the abstract algorithms and concepts used in computer science courses easier, i.e. the algorithms will be more concrete for the students. Visualization is a useful technique for learning in any computer science course. In this paper, we focus at parallel computing and in particular at parallel sorting algorithms [2]. Parallel computing is intrinsically more difficult than sequential computing, it requires coordination of the parallel processes which can be very problematic to conceptualize. We want the students to understand how the processors work together in an algorithm and how they interact. As the students can *see* how the processors run simultaneously in parallel this will illustrate important concepts such as processor load balance, synchronization, and communication.

The use of pictures and visualizations as educational aids is accepted practice. Textbooks are filled with pictures and teachers often diagram concepts on the blackboard to assist an explanation. Animation includes one more dimension, explaining and illustrating the dynamic properties of an algorithm. Role plays goes even further by actively involving the learners within the algorithm. Furthermore, animations and role plays bring a variety into teaching that is usually missing. Different students have different learning styles. The learning styles can be categorized into *Visual (V), Aural (A), Read/Write (R)*, and *Kinesthetic (K)* [1]. While traditional teaching and lecturing are more directed towards the Aural and Read/Write learners, animations are more focused towards the Visual learners and role plays towards Kinesthetic learners. Moreover, computer systems students have some preference to the visual learning style [1].

In this paper we discuss how students perceive these two visualization techniques, animations and role plays, for understanding abstract algorithms in parallel computing. We have used both techniques in the same course and then let students answer questions anonymously in writing. The students have also taken the VARK-test [1] to correlate the answers with their learning style.

Computer animations are commonly used in traditional computer science courses, see e.g. [3, 4], but little efforts have been made in parallel computing. Most visualization techniques are used by researchers for analyzing the parallel performance of their programs [5]. For educational purposes a group at Georgia Institute of Technology lead by John Stasko has developed a software package, Polka, for creating animations of parallel algorithms [9, 10]. We have used this package to develop our animations. Another system that supports visualization of concurrent events is the Pavane system [8]. A historically interesting note of using role plays is the experiment set up by Lewis F. Richardson in 1922 for computing a weather prediction [7], long before the existence of parallel computers. Richardson simulates a parallel algorithm by letting individual workers do partial numerical calculations in parallel, as in todays parallel computers.

2 Method

2.1 Computer Animations

We have developed two different animations for parallel sorting, using the Polka package [9, 10]. The animations are based on two parallelization strategies that are completely different and with different properties regarding communication, load balance, memory requirements, etc. The purpose is to let the students compare the algorithms, learn how they work and understand their inherent problems, i.e. the communication pattern, computation complexities and load balance. The students can work individually with the animations in a computer room exercise.

Learning aspects

Before starting designing animations one must consider the learning aspects. It is important to activate the students in the animations. Previous studies show that there is no significant learning gain in passive viewing [3]. The studies show that; *learners who are actively engaged with the visualization technology have consistently outperformed learners who passively view visualizations*. The learners can be engaged in a number of ways, such as:

- Constructing their own input data sets.
- Making predictions regarding future visualization states.
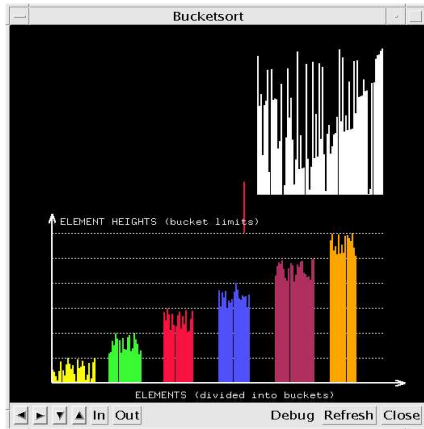- Programming the target algorithm.

- Answering strategic questions about the visualization.
- Constructing their own visualization.

Our goal when construction the animations has been to make them *interactive*, i.e. to engage the students in the animations. We let the students create their own data sets and input parameters. This allows them to actively investigate the algorithms by trying different scenarios. As a leading guideline for experimentation we pose questions pointing to the interesting features of the algorithms in case. The students are then free to use the animations and create their own data sets to find answers.

In the design phase we have consulted students that have taken the course in the previous year. These students have provided valuable input on how to design and improve the animations. They have been an important part of the design process by actively working with and evaluating our animation prototypes.

Parallel bucket sort

The first animation shows a parallel bucket sort algorithm for sorting a sequence of real numbers in the interval [1, 100]. The numbers are represented with bars of corresponding height. We define k buckets by dividing the interval length uniformly into k subintervals, e.g. [1,25], [26,50], [51,75], [76,100] using four buckets on the interval [1,100]. The elements (numbers) are filtered in a sequential step into the equally spaced buckets. The buckets are then assigned to and sorted with quicksort on different processors in parallel. When running the animation, the user is prompt for the number of buckets, the number of processors, and how to distribute the buckets to processors. In addition, the animation is independent of the random number sequence. The user can choose from different random number sequences as input or even provide her own data file. In the animation we have two windows, one dynamic view of how the elements are sorted (Figure 1a), and one static view of the processor load balance (Figure 1b).

(a) Physical sorting view (b) Processor load balance

Figure 1: The *bucketsort* approach for parallel sorting.

In a first step, the elements are moved from a single stack to the different buckets. The buckets are coded with different colors depending on which processors they belong. Next, the elements are sorted simultaneously in the different buckets creating an illusion of parallelism. The processor with the highest work load will continue sorting the longest time, illustrating the effect of the load imbalance. Different random number sequences will give different work loads in the buckets. The workload can then be balanced with, e.g., using more buckets than processors and distributing the buckets cyclicly to the processors or by providing a file with pre-described non-uniform bucket widths.

Parallel Quicksort

The second animation illustrates a memory efficient and fully parallel algorithm. The data is kept distributed all the time and there are no large serial sections as the filtering operation above. In this approach we divide the elements equally among the processors and perform a local sort within each processor. Then we exchange data pairwise between processors to get a global sorting order. The algorithm is outlined in Figure 2.1.

Parallel quick-sort

```
1.  Divide the data into p equal parts, one per processor.
2.  Sort the data locally for each processor.
3.  Perform global sort.
   3.1 Select pivot element within each processor set.
   3.2 Locally in each processor, divide the data into two
       sets according to the pivot ('≤' or '>').
   3.3 Split the processors into two groups and exchange data
       pairwise between them so that all processors in one group
       get data less than the pivot and the others get data larger
       than the pivot.
   3.4 Merge the two lists in each processor into one sorted list.
4.  Repeat 3.1 - 3.4 recursively for each half (log₂ p steps).
```

Figure 2: Parallel Quicksort. The elements are divided equally among the processors and sorted locally. Then, in a number of steps, the processors split their data into two parts according to a pivot and exchange data pairwise with a merging step to get a global sorting order.

(a) Snapshot during global sorting   (b) Execution profile
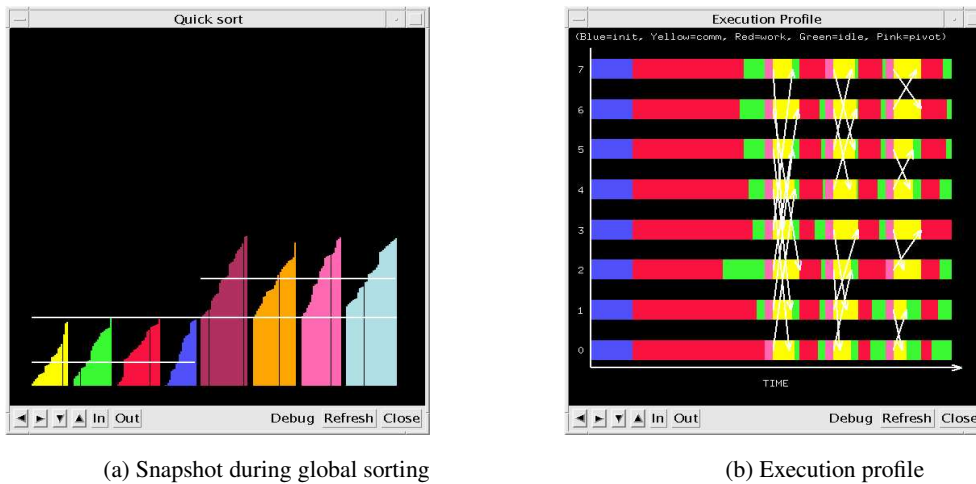
Figure 3: Views of the parallel quicksort animation.

In this animation the user can also choose the number of processors and the random number sequence as input. Here, the processor load balance depends on how the pivot elements are chosen in step *3.1*, Figure 2.1. There are different pivot selections strategies available for the user to choose as input. The animation supports two different views of the algorithm. The dynamic view shows how the elements are physically sorted and moved between different processors (Figure 3a). All parallel activities are animated simultaneously, creating an illusion of parallelism. In the second view, an execution profile is grown simultaneously as the animation continues (Figure 3b). Here, different activities are color coded in bars for respective processor. Arrows represent the communication. The arrows are drawn from the sender to the receiver processor and their stopping point indicates how much data is communicated. From the execution profile it is possible to extract both computation and communication load imbalances.

5

2.2 Role Plays

Role plays are well suited for visualization of parallel algorithms. Each student can take the role of a processor in the algorithm. All bottlenecks, e.g. communication and load imbalance, in an algorithm become very obvious and concrete. The role plays can be performed within class.

Again, we have used different sorting algorithms for demonstration. The first algorithm is a pipelined version of *bubblesort*. The students are given different parts of an unsorted deck of cards. They lay out their cards on the table. The first student starts its first iteration moving the largest card by pairwise compare-exchange to her right while the other are idle. When she comes to the end of her cards in the first iteration she must compare and exchange cards with her neighbor and can then start the next iteration. The other students proceed in the same way but have to wait awhile in the first iteration before they can start. Similarly, the processors/students to the right will finish sorting first and will then be idle. This algorithm requires frequent communication and synchronization but it also has a large load imbalance.

Next, we let the students perform the *odd-even transposition* algorithm. Here the students are also given parts of the deck of cards but now they first sort their own cards (locally). Then, they exchange all their cards alternately with their left and right neighbor. In each step one student in a pair merge its cards with its neighbors cards and gives back half of the cards, keeping the other half. The algorithm proceeds in p-steps (where p is the number of processors/students). This algorithm is much faster and requires less communication than the pipelined bubblesort algorithm. Still, there are some load imbalance, half of the processors doing the merging while the other half of the processors are idle.

Finally, we let the students simulate the parallel *quick-sort* algorithm, see Figure 2.1. This algorithm minimizes the communication compared to the two previous ones and terminates in $log_2(p)$ steps. The load imbalance comes obvious from bad pivot choices giving some students more cards than the others.

2.3 Students

The target course, *Programming of parallel computers*, is a C-level course given in the third year. The students come from different programs, Master of Science, Computer Science, Natural Science, and exchange students. The students are mostly males (85%). The pre-requirements for the course are that they have at least taken two programming courses. This is their first course in concurrent programming techniques. The course is given with traditional lectures and has mandatory computer room exercises.

3 Results

In addition to a course evaluation we have asked the students to fill out a survey with the following questions: Grade 1-5 (where 1=not useful and 5=very useful) traditional lectures (L), textbook (B), animations (A) and role plays (R) with respect to:

1. Understanding parallelism and algorithms
2. Motivation for learning
3. Help for programming assignments
4. Which combination of teaching methods would you benefit most of, e.g. (L)+(B)+(A)

The results from the survey are summarized in Figure 4. The scores for the fourth question are calculated by counting the number of (L), (B), (A), and (R), respectively, then dividing these numbers with the number of students and finally multiplying with five to normalize with the other scores.

The results show that traditional lectures outperform the other teaching methods in all aspects but it is interesting to see that both animations and role plays give better understanding for the algorithms than the textbook and that the students prefer animations before the textbook for learning (in combination with lectures). Comparing the two visualization techniques clearly shows that computer animations are experienced by the students as having higher impact on learning than role plays. The results of the survey depend very much on factors such as quality of lectures, book, animations, enthusiasm of lecturer, preference to learning style, etc.

The students did also take a VARK-test [1]. The results from the test, Figure 5, show that this particular group of students had some preference to the Read/Write and Kinesthetic learning styles. This does not correlate with the results of the questionary where the textbook got low scores while lectures got high scores.
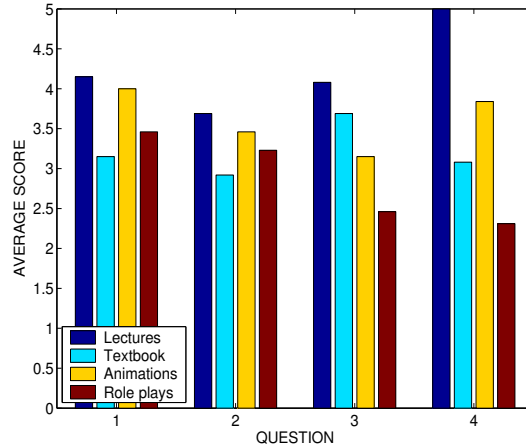
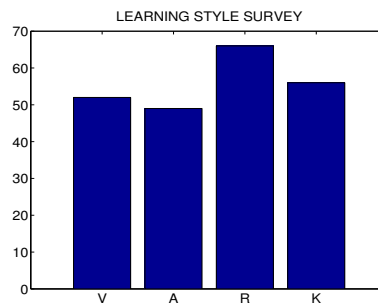Figure 4: Average scores from the survey, 13 students answered the questionary.

Figure 5: Total scores from the VARK-test, 15 students answered the questionary.

The course evaluation, which the students also answered anonymously, was very positive giving an average score 4.3 (out of 5.0) to the course as whole, 4.6 to the lectures and 4.1 to the laboratory exercises including the animations. The textbook got an average score 3.6 which still can be considered good.

4 Discussion

We have used both computer animations and student role plays to visualize abstract algorithms in parallel programming. Both techniques increase and facilitate learning and understanding of the specific algorithms but they also illustrate general concepts, such as load imbalance, synchronization and communication, in a very concrete way. In this particular study, they are experienced by the students as giving better understanding for the parallel algorithms than the textbook.

There are though some fundamental differences between these two visualization techniques. Computer animations allow the students to work in their own pace and let them re-run the animations as many times they want to. They can also experiment with different input parameters and data testing different scenarios and behavior of the algorithms. Role plays on the other hand make the inherent concepts more concrete and real, .e.g. a load imbalance forces one player to work more than the others. Role plays don't require any technical equipment or software development, maintenance, etc. Developing computer animations requires a lot of work and it can be difficult and time consuming to learn how to do this [3]. A problem with role plays is that it can be difficult to perform in a large group of students. It can be impossible to let all participate and then some of the students may not be able to follow or they can feel left out and be unengaged.

In the student survey animations got better scores than role plays. But, the best results were given to the lectures. One explanation is that the impact of the other factors, e.g. quality and enthusiasm, has higher impact than the individual learning styles. Moreover, the lectures are a mixture of different teaching methods, including discussions, writings and visualization with pictures on the blackboard.

An important conclusion is that the visualization techniques cannot replace traditional lectures but can

serve as a valuable complement increasing learning and motivation. This is also how the results of the project will be used, i.e. the developed animations and role plays will be a part of the pedagogy in our target course, Programming of parallel computers.

References

[1] N. Fleming, *VARK, a guide to learning styles*, `http://www.vark-learn.com/`, 2001.

[2] A. Grama, A. Gupta, G. Karypis, V. Kumar, *Introduction to Parallel Computing*, Second Edition, Pearson Education Limited, 2003.

[3] T. Naps, G. Roessling, et. al., *Exploring the Role of Visualization and Engagement in Computer Science Education*, report of the working group on Visualization at the ITiCSE conference in Århus, Denmark, 2002.

[4] T. L. Naps, G. Roessling, J. Anderson, S. Cooper, W. Dann, R. Fleischer, B. Koldehofe, A. Korhonen, M. Kuittinen, C. Leska, L. Malmi, M. McNally, J. Rantakokko, R. Ross, *Evaluating the Educational Impact of Visualization*, inroads - Paving the Way Towards Excellence in Computing Education, volume 35, Number 4. pp. 124-136, ACM Press, New York, 2003.

[5] C. Pancake, *Exploiting Visualization and Direct Manipulation to Make Parallel Tools More Communicative*, in proceedings of Applied Parallel Computing, PARA98, Lecture Notes in computer Science 1541, Springer, 1998.

[6] J. Rantakokko, *Interactive Animation as a Learning Aid for Algorithms*, in proceedings of Utvecklingskonferensen för högre utbildning, Gävle, 26-28 November, 2003.

[7] L.F. Richardson, *Weather Prediction by Numerical Process*, Cambridge University, 1922 (re-published 1965).

[8] G. Roman, K. Cox, C. Wilcox and J. Plune, *Pavane: A System for Declarative Visualization of Concurrent Computations*, Journal of Visual Languages and Computing, Vol 3, No 1, pp 161–193, 1992.

[9] J. Stasko, W. Appelbe, K. Eileen, *Utilizing Program Visualization Techniques to Aid Parallel and Distributed Program Development*, Graphics, Visualization, and Usability Center, Georgia Institute of Technology, Atlanta, GA, GIT-GVU-91/08, June 1991.

[10] J. Stasko, J. Domingue, M. Brown, B. Price, (editors), *Software Visualization: Programming as a Multimedia Experience*, MIT Press, Cambridge, MA, 1998.

Author note