CS586: Distributed Computing
Instructor: Panagiota Fatourou

Exercise 2

***Deadline**: December 19, 2011*

1. Consider the unbounded lock-free queue implementation. Explain (with a counter-example) why the order of lines 29 and 30 (slide 8 - Section 8: Pools) is important for correctness.
2. Consider the non-blocking linked-list implementation of Tim Harris (slides 26-30 - Section 8: Pools).
    a. Is the implementation correct if we delete lines (39) and (40) of Delete()?
    b. Is the implementation correct if we delete lines (17)-(20) of Search()?
    c. Present a simplified version of the implementation where all the parts of the code referring to marking of nodes have been removed. Is the new version of the algorithm correct?
3. Read Section 7.6 of Chapter 7, M. Herlihy and N. Shavit, "The art of multiprocessor programming".
    a. Present C-like pseudo-code for the TOLock described in this section.
    b. Explain, in your own words, the way the algorithm works.
    c. Discuss the liveness properties of the algorithm, the number of cache-misses that it causes, and its space complexity.
    d. Prove that the algorithm is correct.