

Homework Exercise #1

Due: September 29, 2008

1. In this question, you will consider two models of distributed systems. Both models have asynchronous processes that are reliable (*i.e.*, there are no process failures). Each process has a unique identifier.

In both models, processes communicate by sending messages to one another and every process is connected to every other process by a channel. The channels are also asynchronous, so the time for a message to cross a channel from sender to receiver is unpredictable. Each process knows the ids of the processes at the other end of each of its channels, so it can send a message to a particular process, and it knows the identifier of the sender of each message it receives.

In **Model A**, the message channels can carry only 8-bit messages, and messages may get reordered in the channels. (That is, if several messages are en route from process 1 to process 2, they may arrive at process 2 in a different order from the order that process 1 put them into the channel.) Each message sent across a channel is guaranteed to be delivered eventually.

In **Model B**, the message channels can drop some messages. However, if process 1 sends the *same* message to process 2 infinitely many times, at least one copy of it will eventually be delivered to process 2. Aside from the dropped messages, each channel is FIFO (first-in, first-out): messages that are delivered are delivered in the order that they were sent. There is no limit on the length of messages that can be sent in Model B.

- (a) Is Model A at least as strong as Model B? Show your answer is correct.
- (b) Is Model B at least as strong as Model A? Show your answer is correct.