



UCL Department of Computer Science
CS M038/GZ06: Mobile and Cloud Computing
Spring Term 2014
Kyle Jamieson and Brad Karp

One-pager: MapReduce (Dean and Ghemawat, 2004)

Due: Start of lecture, 7th March 2014

*Instructions: **in your own words**, answer the following question as **succinctly** as possible (in 200–500 words, but shorter answers within this range are encouraged). Quoting figures or text from the assigned reading or from any other source is categorically prohibited.*

Trace the flow of data at each step of a MapReduce computation. For the input files, and every step of the MapReduce computation, describe *a)* where the data is physically, *b)* what filesystem contains the data, and *c)* all the mechanisms MapReduce uses to move data closer to the consumer of the data.