(a) There is no significant difference in execution times for different pairs of nodes, since all vertices and edges are traversed each time.

(b) Speed-up is most significant for pairs of nodes that are close to each other – the closer they are, the sooner the algorithm will terminate.

There is extra computational cost per node, the one of comparing the current node to the distance. Therefore, there is a tradeoff: early termination decreases the cost by visiting fewer nodes while increasing the cost per node. If most of the nodes are visited, which is the case for distant pairs of nodes, the increase might overpower the decrease, yielding a slightly higher execution time.

(c) If both edge endpoints and the landmark are in the same connected component, the proof is the same as given in lecture. If edge endpoints are in the same component but different from landmark, then they both have potential C and therefore the edge weight does not change. Finally, if edge endpoints are in different components, there is no edge between them, and their potentials are irrelevant. (Note that a component to which destination belongs is irrelevant, since its shortest distance to the landmark is a constant that is subtracted when computing each potential.)

It is more useful to precompute distances to the landmark rather than potentials because potentials are bound to a specific destination, while distances can be reused to compute potentials with respect to multiple destinations with negligible additional cost. Moreover, computing potentials would be meaningless, since it would take one full run of Dijkstra's algorithm to save us some computation on another run of Dijkstra's algorithm.

(d) The speed-up is most significant when a landmark is close to the destination. Potentials method effectively increases the shortest path from the source to a node by the value of node's potential. Destination potential is always 0. It is useful to increase potentials of nodes that are unimportant for finding a route from source to destination (i.e., nodes far from that route), such that their shortest path is now longer than the one to the destination and Dijkstra's algorithm with early termination does not reach them. This is best achieved when a landmark is close to the destination.

If multiple landmarks are given, choose the one that is closest to the destination. Some additional criteria can be thought of if several landmarks are close to the destination.

(e) The combined potential function is always better or equal than any single-landmark potential function in terms of number of visited nodes (assuming that such a single landmark belongs to the set of landmarks used in the combined potentials method). This is true because a method of max potentials can only increase potential of a node and thus increase chances that it is not reached before the destination. (Note that the potential of the destination is always 0.) Therefore, the combined method on a set of landmarks will visit at most the same number of nodes as any single-landmark method for a landmark in that set.

(f) The performance matches the statements above.

(g) (Optional)