

Developing efficient algorithms to synchronize between different versions of files is an important problem with numerous applications. We consider the interactive synchronization protocol introduced by Yazdi and Dolecek, based on an earlier synchronization algorithm by Venkataramanan et al. Unlike preceding synchronization algorithms, Yazdi and Dolecek's algorithm is specifically designed to handle a number of deletions linear in the length of the file. We extend this algorithm in three ways. First, we handle nonbinary files. Second, these files contain symbols chosen according to nonuniform distributions. Finally, the files are modified by both insertions and deletions. We take into consideration the collision entropy of the source and refine the matching graph developed by Yazdi and Dolecek by appropriately placing weights on the matching graph edges. We compare our protocol with the widely used synchronization software rsync, and with the synchronization protocol by Venkataramanan et al. In addition, we provide tradeoffs between the number of rounds of communication and the total amount of bandwidth required to synchronize the two files under various implementation choices of the baseline algorithm. Finally, we show the robustness of the protocol under imperfect knowledge of the properties of the edit channel, which is the expected scenario in practice.