

We consider the problem of distributing data blocks on servers that are heterogeneous in both storage capacity and network connection speed in a formalized context. For this, we present two formal problem definitions that involve distributing data block copies onto different servers to maximize the number of requests the storage system can handle. Using more than one copy per block to optimize performance will be a core aspect of this work. One problem definition will consider a homogeneous data set, the other a heterogeneous one. In the former data blocks will be equal in importance, in the latter there will be different demand on them. We formalize the problems, prove that both are NP-hard and for each present a first method for simplifying the problem instances. We focus on the case of a homogeneous data set, analyze cases of heterogeneity in one of two server dimensions, and show a polynomial 2-approximation as well as a first relationship between the heterogeneity of servers and a guaranteed greedy approximation ratio. Following that we present a greedy approach to the heterogeneous data set case with various different strategies that we later compare in an empirical analysis. Last but not least we conclude our work with an outlook into possible further research.