This work examines the possibilities of employing highly parallel architectures in database systems, which are based on the similarity search paradigm. The main objective of our research is utilizing the computational power of current GPU devices for similarity search in the databases of images. Despite leaping progress made in the past few years, the similarity search problems remain very expensive from a computational point of view, which limits the scope of their applicability. GPU devices have a tremendous computational power at their disposal; however, the usability of this power for particular problems is often complicated due to the specific properties of this architecture. Therefore, the existing algorithms and data structures require extensive modifications if they are to be adapted for the GPUs. We have addressed all the aspects of this domain, such as efficient utilization of the GPU hardware for generic computations, parallelization of similarity search process, and acceleration of image indexing techniques. In most cases, employing the GPU devices brought a speedup of two orders of magnitude with respect to single-core CPUs and approximately one order of magnitude with respect to multiprocessor NUMA servers. This thesis summarizes our experience and discoveries from several years of research, related algorithms adopted for the specific conditions of GPU architectures, and the results of empirical experiments performed in order to verify our claims.