Stable principal component pursuit (SPCP) seeks to decompose noisy signals into their underlying low-rank and sparse components. In this study, we introduce two new methods to parallelize SPCP in order to take advantage of GPU, multiple CPU, and hybridized architectures. Our first development is a parallelizable algorithm for the randomized singular value decomposition (rSVD) that ameliorates the problem of tall-skinny matrix structures that arise often in applications. Our second development is a reformulation of SPCP using Burer-Monteiro splitting, which eliminates the need for an SVD and is particularly well suited for the GPU. We demonstrate, using synthetic data, surveillance video, and data from fMRI brain scans, that both algorithms offer significant speedup over traditional SPCP solvers.