



FIG. 6.1. CPU times for applying the Laplace kernel in various cases using LAPACK/ATLAS (LP), the FMM, and recursive skeletonization (RS) as a function of the matrix size  $N$ . For LP and RS, the computation is split into two parts: precomputation (pc), for LP consisting of matrix formation and for RS of matrix compression, and matrix-vector multiplication (mv). The precision of the FMM and RS was set at  $\epsilon = 10^{-9}$ . Dotted lines indicate extrapolated data.