Semiparametric joint models of longitudinal and survival data are computationally costly and their current implementations often do not scale well to large data. We investigate and address some key computational barriers in a typical EM algorithm for joint models, which include numerical integration, risk set calculation, standard error estimation, and choice of the initial values. We show that efficient algorithms and implementation can generate drastic speedups, reducing the run-time from days to minutes for large data. Illustrations will be presented on both simulated and real world data.