Although high-performance computing traditionally focuses on the efficient execution of large-scale applications, both energy and power have become critical concerns when approaching exascale. Drastic increases in the power consumption of supercomputers affect significantly their operating costs and failure rates. In modern microprocessor architectures, equipped with dynamic voltage and frequency scaling (DVFS) and CPU clock modulation (throttling), the power consumption may be controlled in software. Additionally, network interconnect, such as Infiniband, may be exploited to maximize energy savings while the application performance loss and frequency switching overheads must be carefully balanced. This work first studies two important collective communication operations, all-to-all and allgather and proposes energy saving strategies on the per-call basis. Next, it targets point-to-point communications to group them into phases and apply frequency scaling to them to save energy by exploiting the architectural and communication stalls. Finally, it proposes an automatic runtime system which combines both collective and point-to-point communications into phases, and applies throttling to them apart from DVFS to maximize energy savings. The experimental results are presented for NAS parallel benchmark problems as well as for the realistic parallel electronic structure calculations performed by the widely used quantum chemistry package GAMESS. Close to the maximum energy savings were obtained with a substantially low performance loss on the given platform.