LA-SiGMA
Louisiana Alliance for Simulation-Guided Materials Applications

NSF Highlight

LSU
CENTER FOR COMPUTATION & TECHNOLOGY

Interdisciplinary Physics/Computer Engineering Collaboration speeds many-body codes to 30,000 processors

Outcome: An interdisciplinary team of physicists and computer engineers has developed many-body codes that can efficiently use 30,000 or more processors at NSF supercomputing facilities, using techniques that overlap the inter-processor communication with local computations.

Impact/ Benefits: These advanced latency hiding codes will allow scientist to more accurately model competing phases, including magnetism, superconductivity, insulating behavior etc. in strongly-correlated electronic systems.

Background/ explanation: Advanced latency hiding methods, together with new national leadership class machines such as NSF's Kraken at NICS, have brought two-particle many-body theories to the point where they may be efficiently simulated. These more accurate methods will be used to more reliable model competing phases, including magnetism, superconductivity, insulating behavior etc. in strongly correlated electronic systems. The 2007 International Technology Roadmap for Semiconductors stresses that highly correlated electron systems can enable new devices by greatly enhancing their sensitivity to different applied fields. Continuous advances in computing power over the last 30 years are attributable to Moore's Law, according to which the transistor density in modern semiconductor chips doubles every two years. However, CMOS technology, upon which these increases are based, can no longer sustain this growth. New materials, with new functionalities, are required to sustain Moore's law into the future. Devices based on correlated electronic materials can use the electron's spin and orbital degrees of freedom, in addition to the electronic charge, to encode information, yielding an exponential increase in computing power. In addition, the competing phases these materials display offers new ways to control these devices.

Images/videos understandable by general audience, with captions and credits

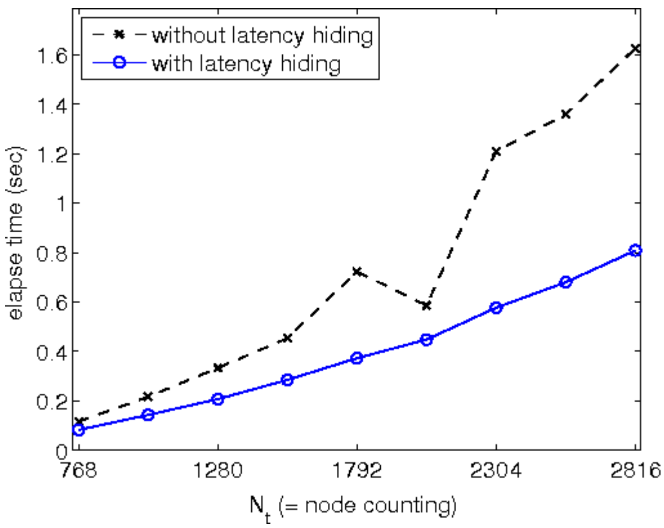

Caption: Time spent in data communication as a function of the number of computer nodes (12 processors per node). For large data sets each process sends messages to all the others, and the communication time scales linearly with the number of processes. Latency hiding techniques, that overlap the interprocessor communication with local computations, yields a factor of two speedup when compared with standard message-passing libraries as the number of processors increases beyond 30,000.

Credit: T.W. Lee (CCT, LSU), J. Ramanujam (ECE/CCT, LSU), M. Jarrell and J. Moreno (Physics/CCT, LSU).

Permission from copyright holder via email or form 1515 to use images publicly