Petascale Quantum Simulations of Nano Systems and Biomolecules

Emil Briggs
North Carolina State University

1. Outline of real-space Multigrid (RMG)
2. Scalability and hybrid/threaded models
3. GPU acceleration
4. RmgLib
5. Application: Electron transport in DNA
   Conduction mechanisms
   Counterions, water, base pair matching

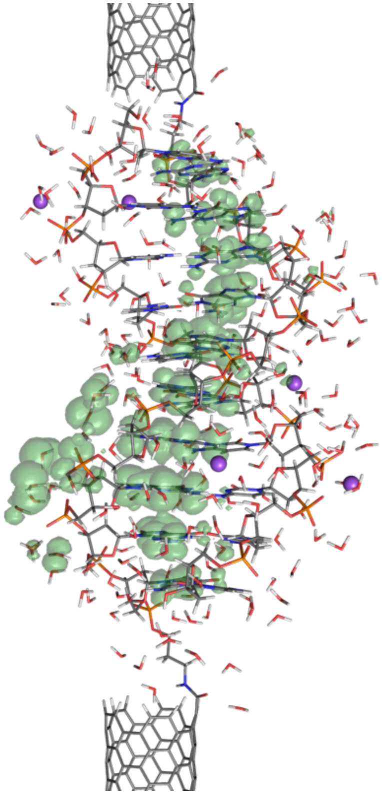

Collaborators: Wenchang Lu, Bikan Tan, Yan Li, Miroslav Hodak, Jerzy Bernholc

Real-space Multi-Grid (RMG)

- **Density functional equations solved directly on the grid instead of with plane waves**
- **Multigrid techniques remove instabilities by working on one length scale at a time**
- **Convergence acceleration and automatic preconditioning on all length scales**
- **Non-periodic boundary conditions are as easy as periodic**
- **Compact "Mehrstellen" discretization**

$$A[\phi_i] + B[(V_{eff} + V_{NL})\phi_i] = \varepsilon_i B[S\phi_i]$$

- **Allows for efficient massively parallel implementation (no FFTs)**

See E. L. Briggs, D. J. Sullivan and J. Bernholc Phys. Rev. B 54, 14362 (96).

Ultrasoft pseudopotentials:

M. Hodak, S. Wang, W. Lu and J. Bernholc, PRB 76, 085108 (07)

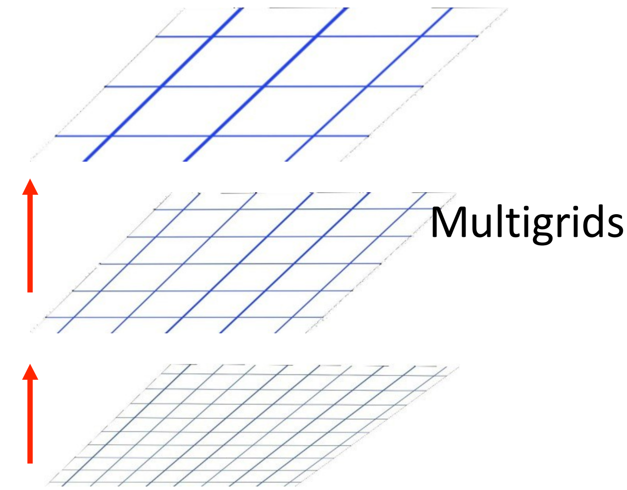

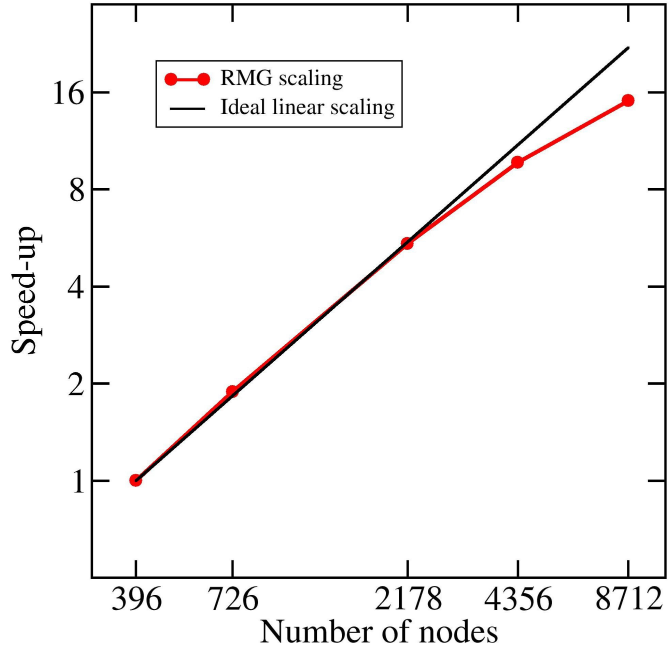

Scalability as HPC architecture changes

Multisocket/multicore SMP nodes

- 1999 – Cray T3E 1-socket and core per node
- 2013 – Cray XK6 2-sockets and 32 cores per node
- 2020 - ?

High speed interconnect between nodes

- Infiniband
- Myrinet
- Cray Gemini

GPU/Accelerator

- Nvidia Fermi/Tesla/Kepler
- AMD Radeon HD
- Intel Xeon PHI

Current machines **400,000** cores
Next generation **1,000,000** cores

Multicore CPU's

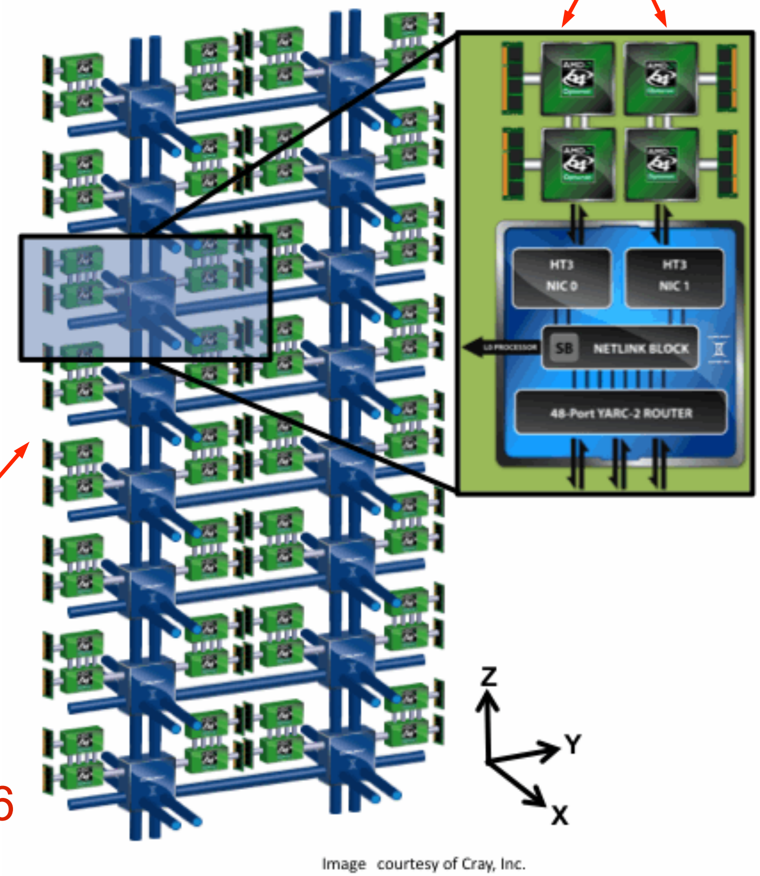

Schematic of Cray XE6

Hybrid MPI/threads/OpenMP to improve scalability

Use 1 MPI process per node rather than 1 process per core
- Inter node parallelization uses traditional MPI
- Intra node parallelization uses shared memory threads
- Extra cores used via Libraries, Pthreads, Openmp

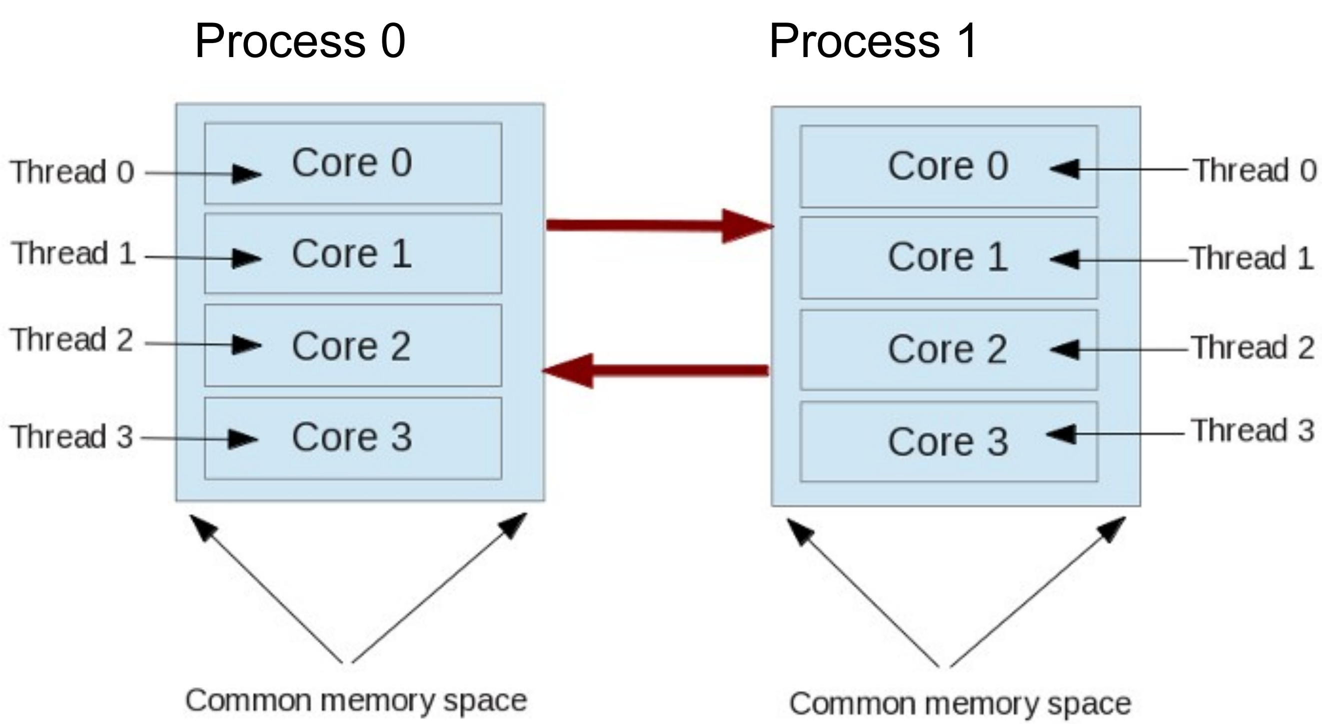

Pthreads for multigrid preconditioner

Consider a typical electronic structure problem with N orbitals
The computer system used for solution has P processing cores per node
Orbitals may be processed independently and N>>P
Natural parallelization method is to assign each orbital to a single core
OpenMP or Pthreads should work equally well

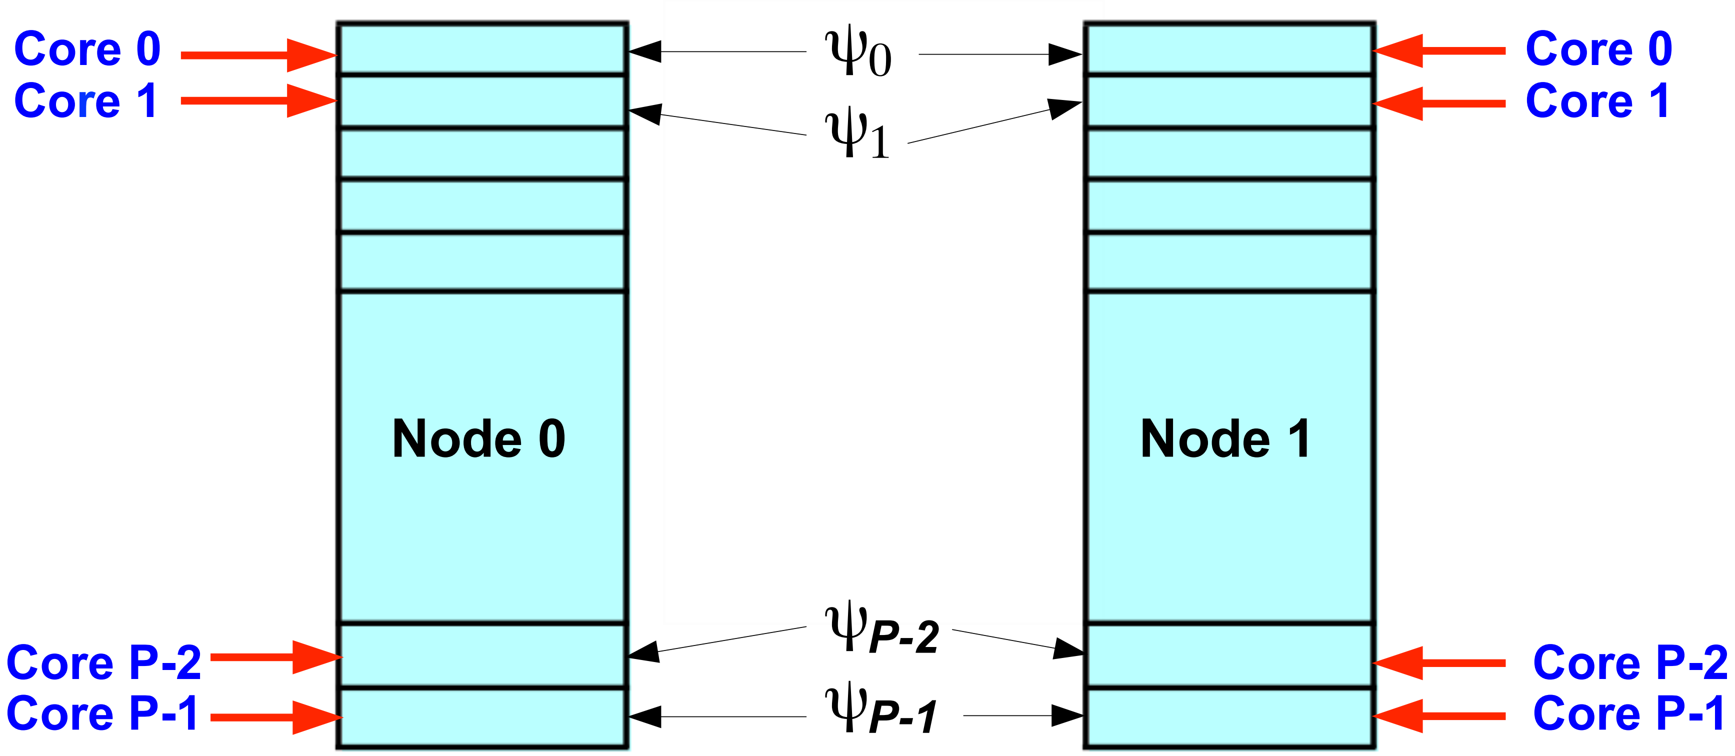

While each thread can operate independently is this the best approach?

Scaling test

Test problem: Gas phase Amyloid Beta 1-42 protein
Test system: Cray XK7
1 node = 16 Opeteron cores + 1 Nvidia K20x GPU accelerator
Strong scaling

Largest run used **139,392** CPU cores and **8712** GPU's.

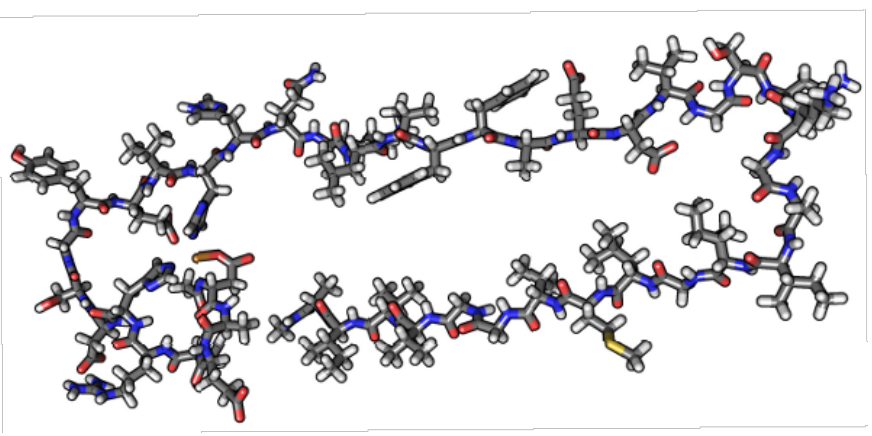

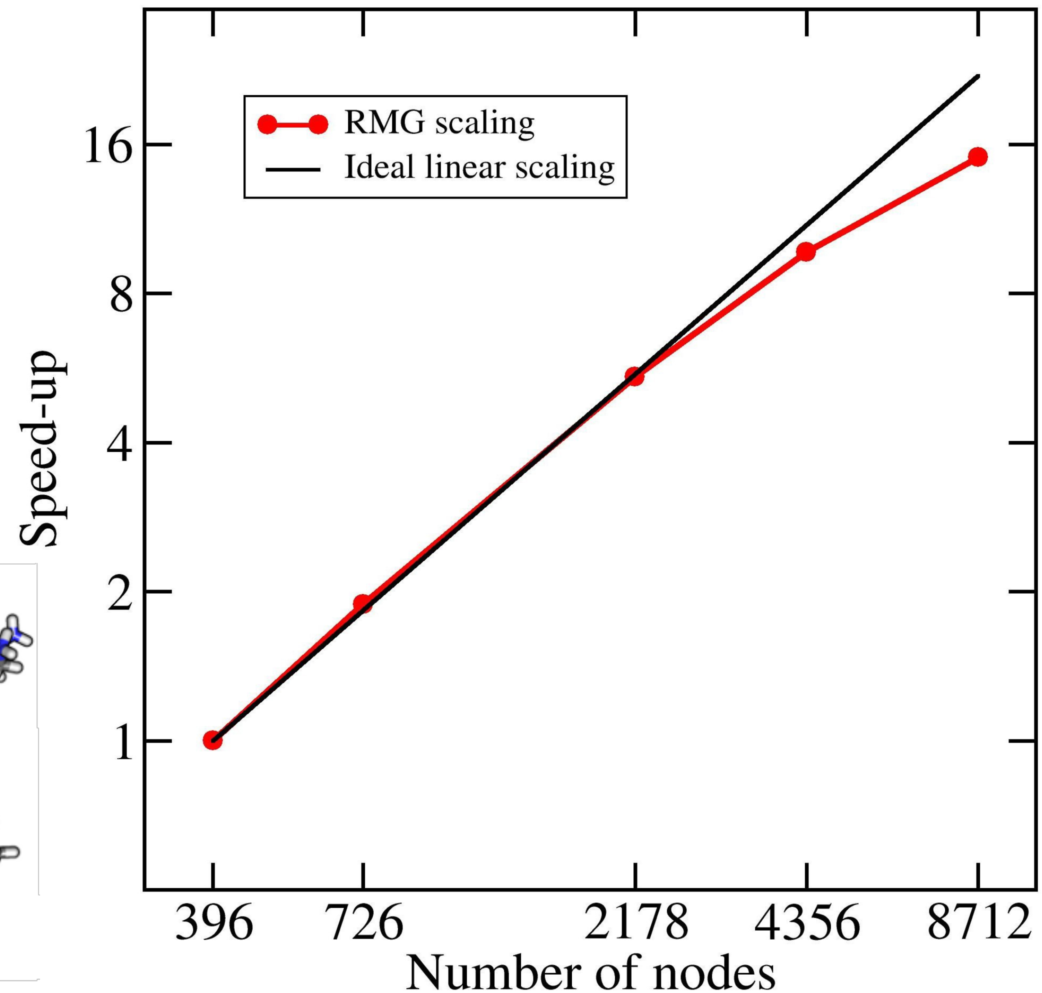

Each node has 16 cores

GPU acceleration for RMG

GPU programming very different from CPU

CPU:

High clock speed, small number of powerful execution units.
Memory latency hidden by caches and out of order execution.
Good single-threaded performance.

GPU:

Low clock speed, large number of weaker execution units.
Memory latency hidden by high thread counts.
Poor single-threaded performance.

Most HPC codes have components that only run well on CPU's

Mixed CPU/GPU model required
Data transfer issues from CPU to GPU (PCI bus latency)

Hints: Avoid writing GPU code as much as possible.
Use vendor supplied libraries.
Data transfer issues still require careful consideration

GPU performance improvements

Small test case: C60 molecule in vacuum

60 atoms: 200 electronic orbitals

CPU only calculation Xeon workstation:12 cores No GPU's
10.32 seconds/SCF step

CPU/GPU calculation
Xeon workstation:12 cores
1 Nvidia K20 GPU
6.72 seconds/SCF step

Speedup of approximately 1.53

Large test case: Solvated amyloid beta protein fragment

3337 atoms: 4672 electronic orbitals

CPU only calculation
2904 nodes: 92,928 Opteron cores
No GPU's
76 seconds/SCF step

CPU/GPU calculation
2904 nodes: 46,464 Opteron cores
2904 Nvidia K20x GPU's
25 seconds/SCF step

Speedup of approximately 3.02

Performance on 3,872 Cray XK7 (K20x GPU) Blue Water nodes: **1.14 PFLOPS**

Introducing RmgLib

Portable library of C++ routines for HPC including

- Grid decomposition across MPI nodes
- Threading on a node
- Finite differencing (Mehrstellen and central operators)
- Communications (asynchronous ghost images)
- Recursive Multigrid solver with support routines

Library supports multiple data types including

- Float, double, complex

Linux/Unix, Windows and OS/X support

- Cmake build system

Open source (BSD type) licensing with initial release in 2nd half 2014

Single DNA Electron Transport

Guo, Barton, Nuckolls et al. Nature Nanotechnol. 2008

- Experiment in 0.05 mol/L saline solution
- Single walled carbon nanotubes as leads
- Alkane linkers -CONH-$(CH_2)_3$ connect B-form DNA to leads
- Very high transmission (T=0.05) gives resistance around 0.5 MΩ
- A single GT or AC mismatch increases the resistance of DNA nearly 300-fold relative to a well-matched one

G-guanine; C-cytosine; A-adenine; T-thymine

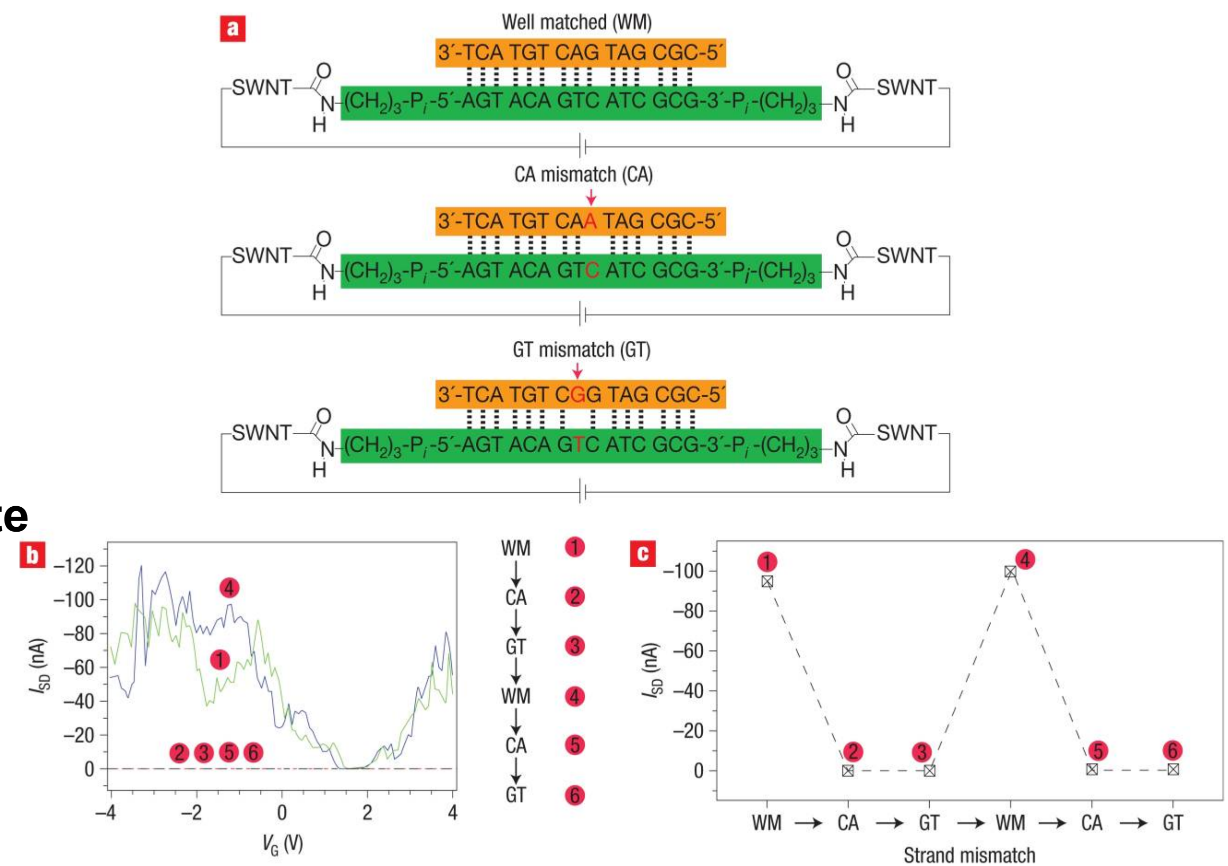

Our goals are to investigate the effects of

- ➢ Linkers
- ➢ Counterions
- ➢ Solvent
- ➢ DNA conformation
- ➢ DNA sequence

Introduction to DNA

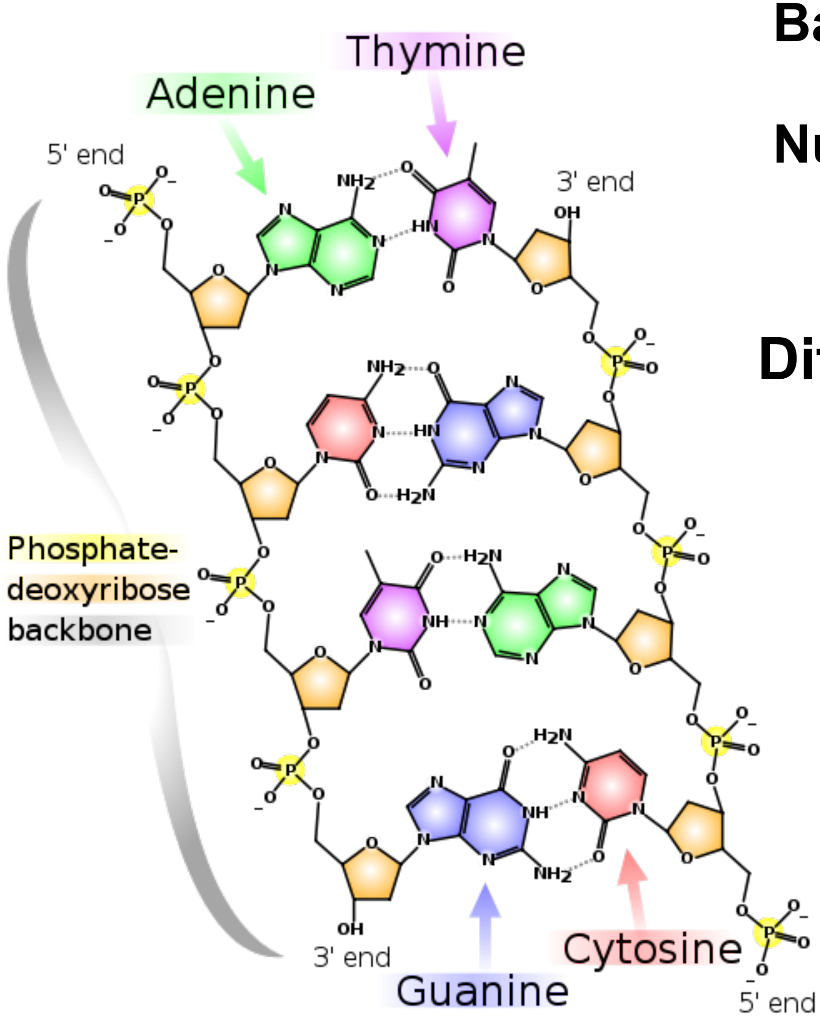

Backbones: sugar rings and phosphate groups

Nucleobases: adenine(A), guanine(G), thymine(T) and cytosine(C). Base pairing can be formed via hydrogen bonds (A-T, G-C)

Different forms of DNA:

	A-form	B-form	Z-form
rise:	2.4 Å	3.6 Å	3.7 Å
twist:	33.6°	35.9°	30.0°

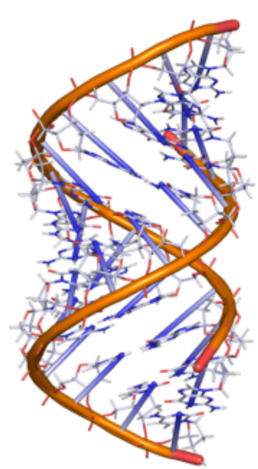
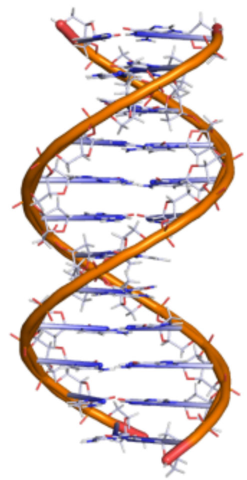
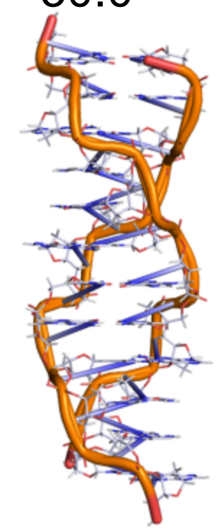

Calculations

- 20 snapshots recorded from 2 ns MD simulation for subsequent quantum calculations
- First solvation shell as well as alkane linkers and CNT leads included in the quantum calculation
- *O(N)* calculation to generate optimized localized orbitals for use in transport calculations
- Iterative non-equilibrium Green function (NEGF) method to calculate transmission self-consistently

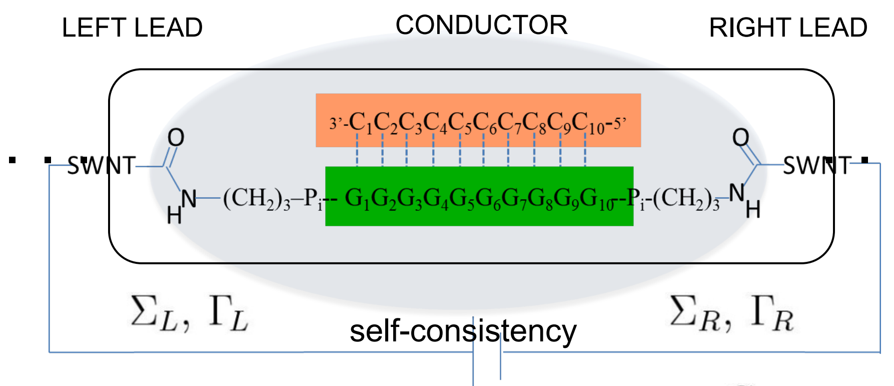

Hamiltonian: H_C orbital overlap: S_C

Self-energy: Σ_L Σ_R

$$G_C = (E\,S_C - H_C - \Sigma_L - \Sigma_R)^{-1}$$

$$T = \mathrm{Tr}\left(\Gamma_L\, G_C^R\, \Gamma_R\, G_C^A\right), \qquad \Gamma = \frac{\mathrm{i}}{2}\left(\Sigma^\dagger - \Sigma\right)$$

$$I(V) = \frac{2e^2}{h}\int_{-\infty}^{\infty} T(E,V)\,[f(E-\mu_L) - f(E-\mu_R)]\,\mathrm{d}E$$

Fattebert, Bernholc, Phys. Rev. B 2000
Nardelli, Fattebert, Bernholc, Phys. Rev. B 2001
Lu, Meunier, Bernholc, Phys. Rev. Lett. 2005

DNA conductivity histogram

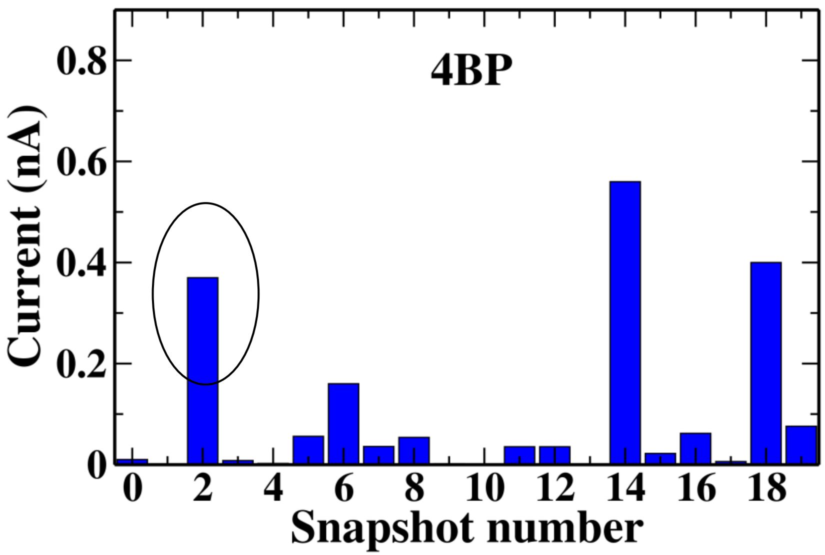

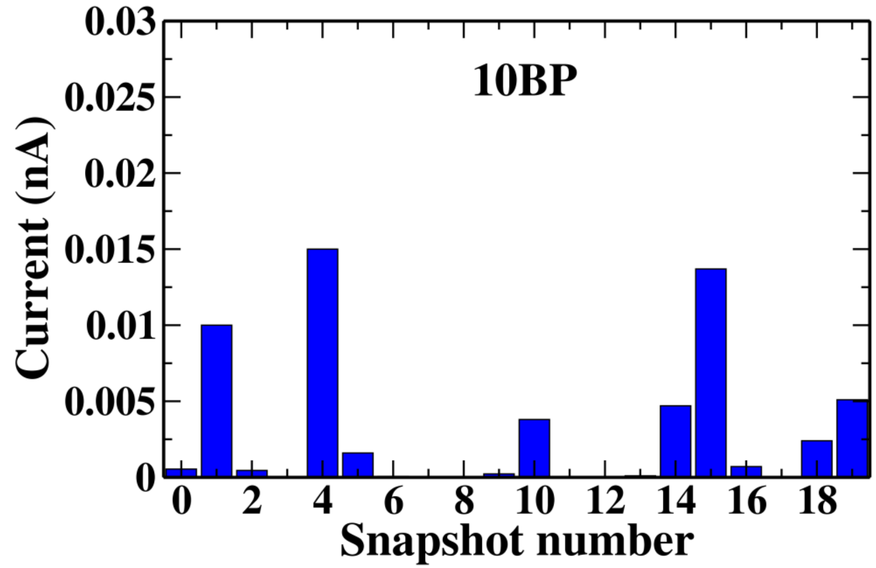

Large variation of conductivity (12 orders of magnitude for 10 BP) among snapshots

- Conformational gating

- Average current for 4 BP poly(G)-poly(C) is 0.1 nA
- Average current for 10 BP poly(G)-poly(C) is 0.0029 nA
- Assume an exponential decay model:

10 base-pair (10BP) configuration

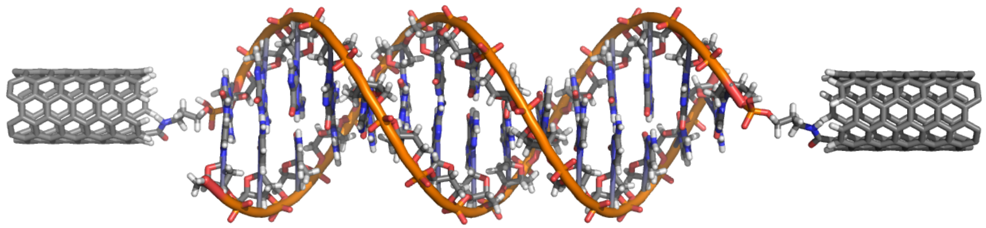

Effect of individual counterions (snapshot 2)

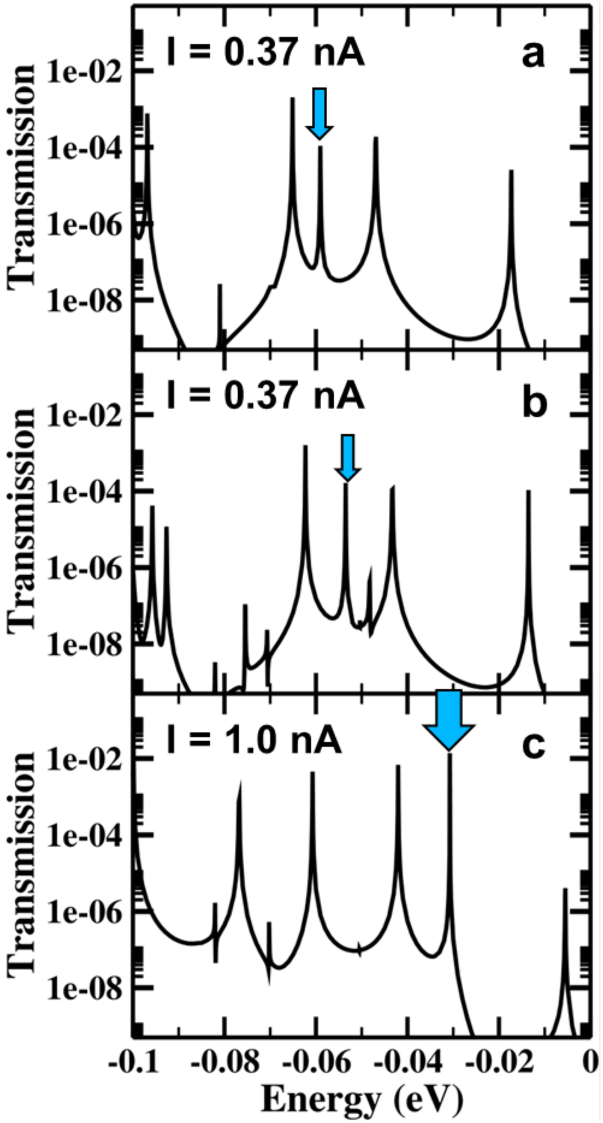

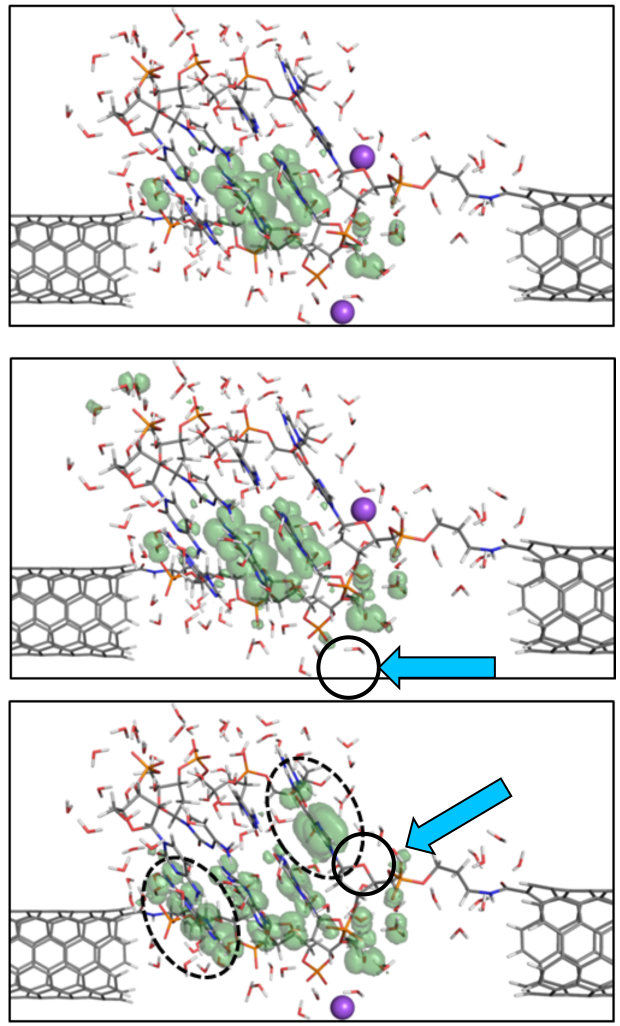

- Two explicit counterions in snapshot 2 (right panel in a)
- The conductivity is unchanged when removing the counterion close to the phosphate group
- After the removal of a counterion near G4, the conductive states become more delocalized and the current increases by a factor of three
- Average current increases by a factor of 2 after removal of ions for the 20 snapshots

Water effects on DNA conductivity

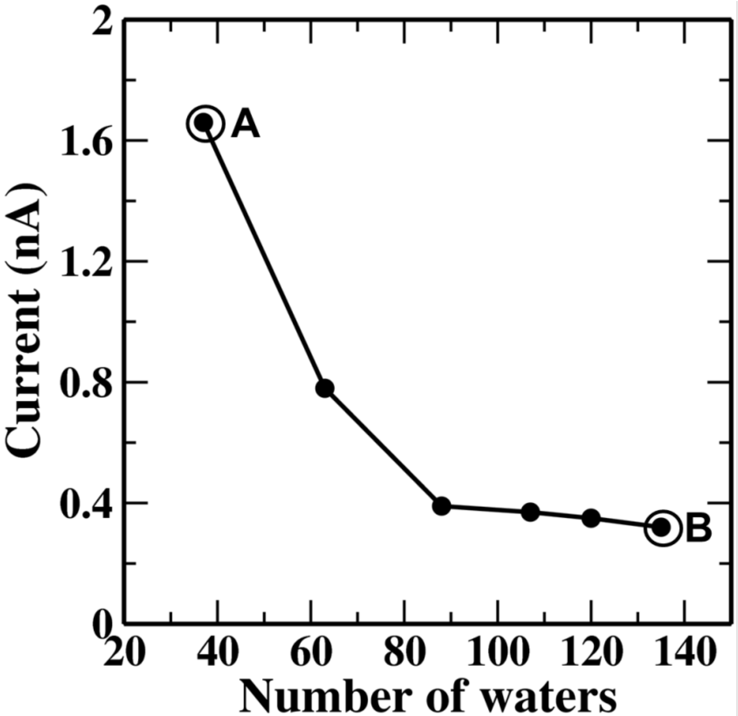

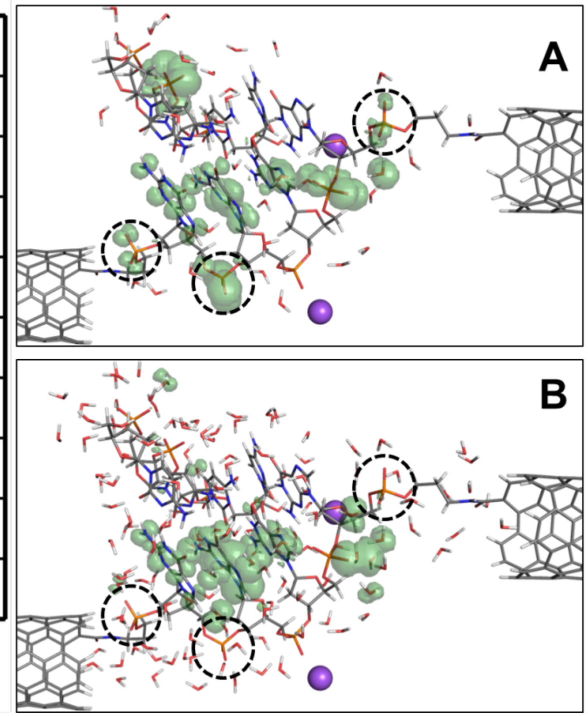

37 waters

135 waters

- Conductivity decreases with the number of water molecules around DNA
- Water dipoles suppress conducting electrons' amplitude on phosphate groups, leading to reduced conductivity
- Presence of water decreases the current by a factor of four
- Presence of ions decreases the current by a factor of two
- Large variation of conductivity persists

Correlation analysis of 10 BP poly(G)-poly(C) DNA

Correlation of conductivity with helicoidal parameters

	mean	minimum	maximum
twist	0.4888	0.1505	-0.0592
roll	0.1571	-0.0128	0.0126
tilt	0.1303	0.181	0.1959
rise	-0.0833	0.2417	0.1569
slide	0.3821	0.383	-0.1662
shift	0.4568	0.6176	0.1275
overlap	0.4208	0.9106	0.051

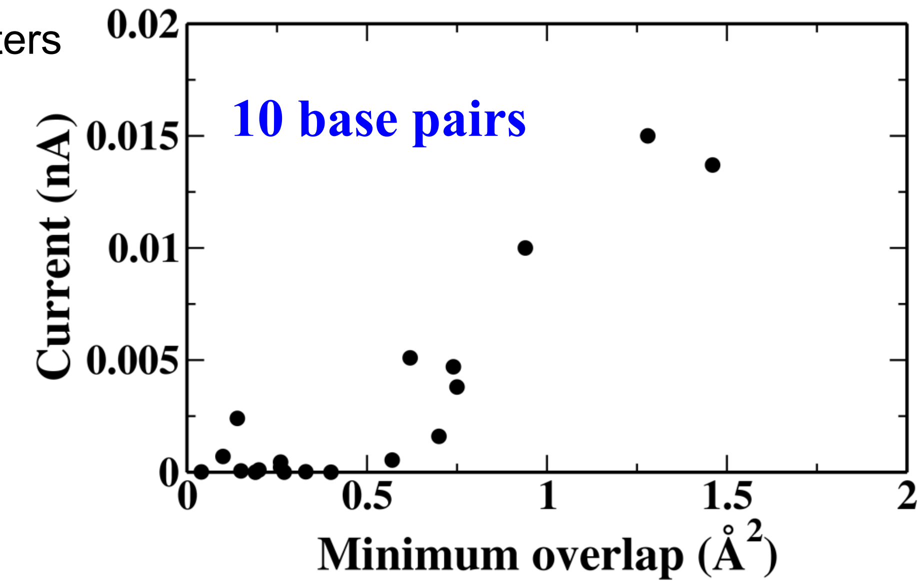

- Helicoidal parameters in DNA calculated using 3DNA
- We consider correlation of conductivity with not only the mean, but also with minimum and maximum values of the parameters
- The minimum *overlap* between successive guanine bases has the dominant correlation with conductivity
- Conductivity is highly sensitive to local distortions
- Snapshot 15: the minimum overlap is 1.46 ⇒ delocalized conducting states with relatively high conductivity (0.0137 nA)
- Snapshot 15’: minimum overlap reduced to 0.73 ⇒ less delocalized conducting states and lower conductivity (0.0045 nA)

Summary

- **RMG code rewrite to enhance scalability and portability**
  - **Real space multigrid method for electronic structure calculations**
  - **Hybrid model using MPI/threads/OpenMP**
  - **GPU acceleration: 1.144 PFLOPS using 3872 Bluewater XK nodes**
  - **Open source RmgLib to be released in 2nd half 2014**
- **Charge transport in DNA**
  - **Charge transport through delocalized hole orbitals**
  - **Highly dependent on instantaneous DNA configuration**
  - **Largest conductivity for largest *minimum overlap* between guanine bases**
  - **Counterions and water molecules significantly reduce conductance**
  - **Unmatched base pairs act as barriers, reduce orbital delocalization.**