

# Scaling Workshop Projections

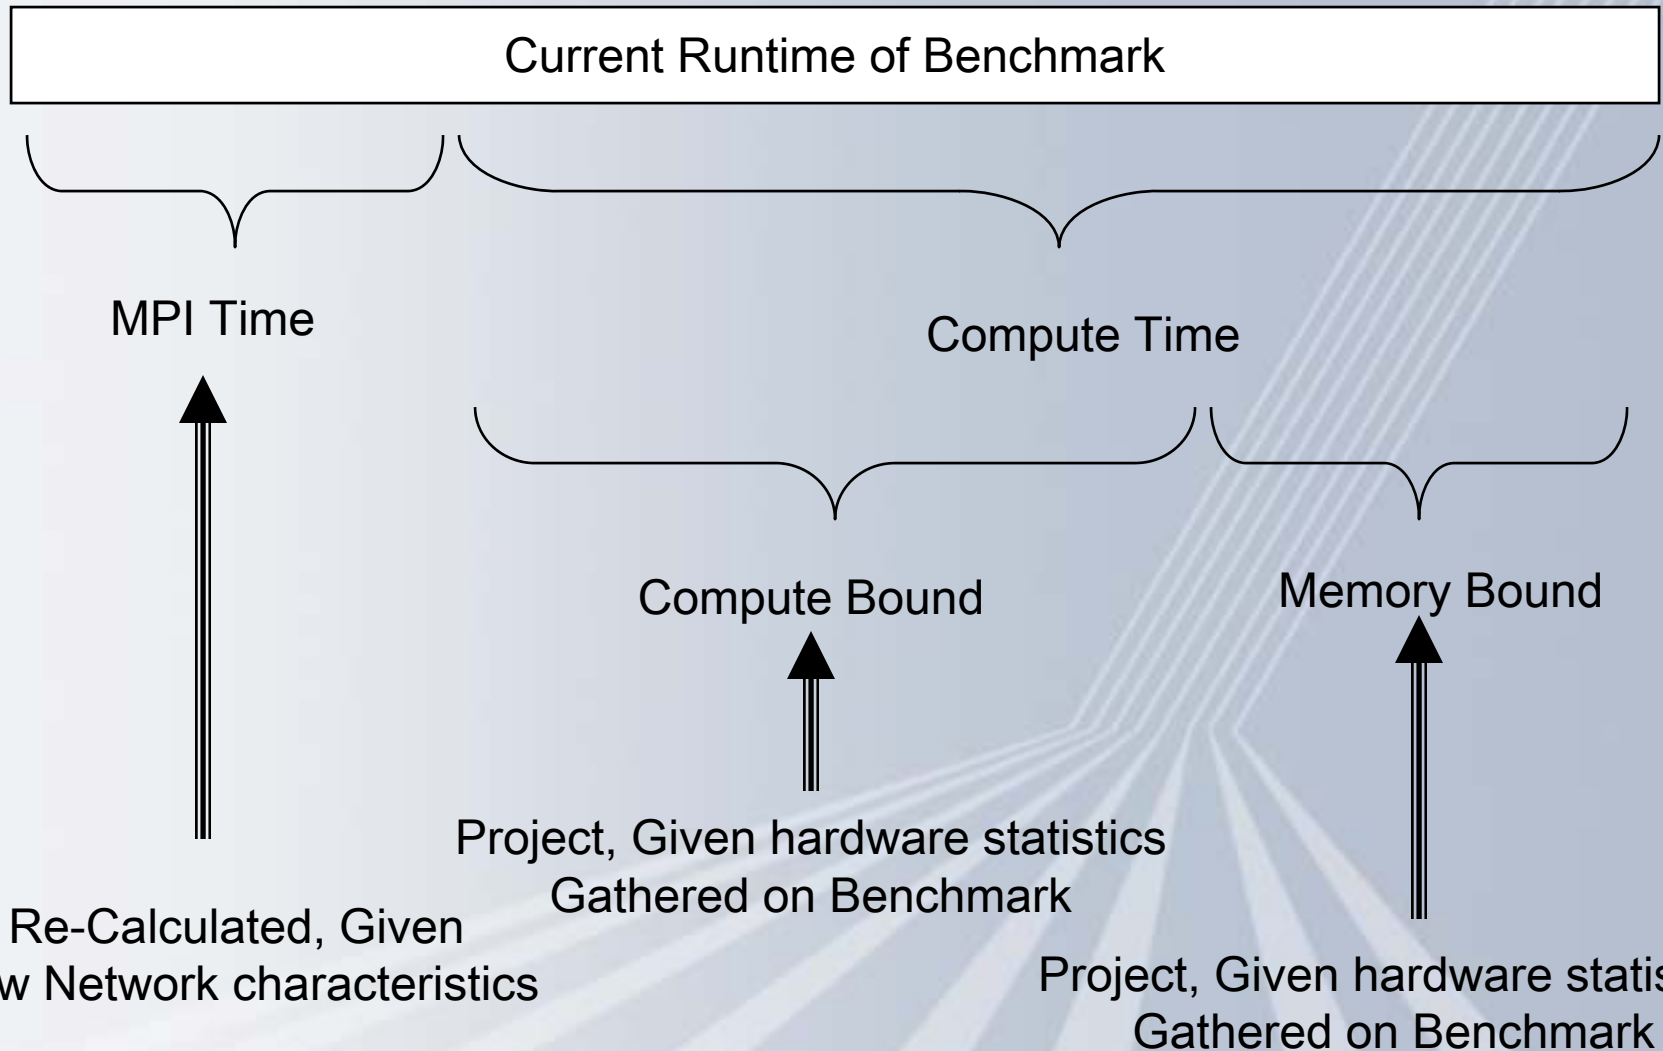
Cray Supercomputing Center of  
Excellence

John M. Levesque  
Director

# Quad Core System – Jaguar System

- Nodes contain 1 sockets of 4 cores each
- Each core can generate 4 Floating point results each cycle
  - 16B wide SSE instructions (4)
  - 2. GHZ
- Memory is is 1channel @ 667 MHZ
  - $1 * 667 = .667$  GB/sec
- Interconnect
  - Latency 3-4 microseconds
  - Bandwidth 4 GB/s
- Nodes contain 1 sockets of 2 cores each
- Each core can generate 2 Floating point results each cycle
  - 8B wide SSE instructions (2)
  - 2.6 GHZ
- Memory is 1channel @ 667 MHZ
  - $1 * 667 = .667$  GB/sec
- Interconnect
  - Latency 3-4 microseconds
  - Bandwidth 4 GB/s


# Projection Methodology




# Projection Methodology

- In Hand
  - Times for Dual and Single core
  - Hardware counters for SSE instructions
- Determine
  - Percentage of code that can use SSE Instructions  
 $1 + 2(FPMULTS + FPADDS + SSE) / SSE$
  - Percentage of code that is memory bound
    - Memory bandwidth bound fraction =  
 $(\text{Dual Time} - \text{Single Time}) / \text{Dual Time}$

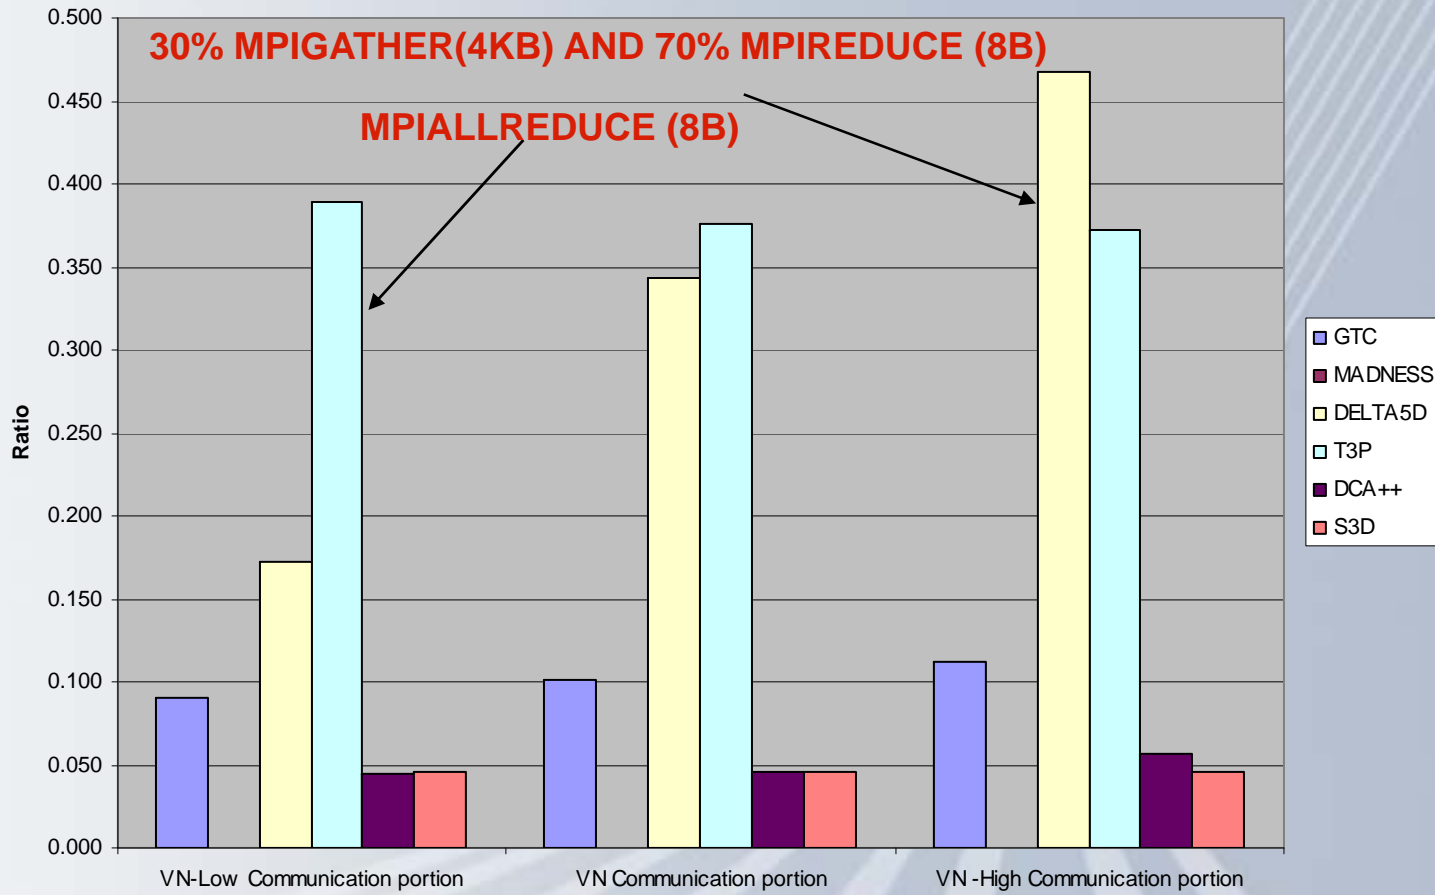
Compute Analysis




Communication Analysis




Communication Analyssi



Projections to 250TF  
Assuming communication the same



Projections to 1000 TF



# Scaling Projections

- Madness, DCA++ and S3D will continue to scale
- GTC may be running into difficulty
  - OpenMP on the node may help on the 250 TF
- DELTA3D and T3P will have significant troubles scaling until Gemini arrives
  - Might try using point to point for collectives on Seastar
- DELTA3D and T3P might benefit from OpenMP on the socket to reduce communication issues

# Compute Recommendations

- DELTA3D, T3P and S3D must be analyzed for cache optimization – they are the most memory bound of the bunch
- DCA++ seems to be well set for the future
- T3P needs to be vectorized badly
- S3D, GTC, DELTA3D and Madness can also use some vectorization