



# Overview of Berkeley UPC

**Kathy Yelick**

Christian Bell, Dan Bonachea,  
Wei Chen, Jason Duell,  
Paul Hargrove, Parry Husbands,  
Costin Iancu, Mike Welcome



# Goals of the Berkeley UPC Project

---



- Make UPC Ubiquitous on
  - Parallel machines
  - Workstations and PCs for development
  - A portable compiler: for future machines too
- Components of research agenda:
  1. Runtime work for Partitioned Global Address Space (PGAS) languages in general
  2. Compiler optimizations for parallel languages
  3. Application demonstrations of UPC



# Where Does Berkeley UPC Run?

---



- Runs on most SMPs, clusters & supercomputers
- Support Operating Systems:
  - Linux, FreeBSD, Tru64, AIX, IRIX, HPUNIX, Solaris, MSWindows(cygwin), MacOSX, Unicos, SuperUX
- Supported CPUs:
  - x86, Itanium, Alpha, Sparc, PowerPC, PA-RISC
- GASNet communication:
  - Myrinet GM, Quadrics Elan, Mellanox Infiniband VAPI, IBM LAPI, Cray X1, SGI Altix, Cray/SGI SHMEM
- Specific supercomputer platforms:
  - Cray T3e, Cray X1, IBM SP, NEC SX-6, Cluster X (Big Mac), SGI Altix 3000



# Recent Progress on Runtime

---



- Runtime portability, interoperability [Jason]
  - New pthread version runs on SGI Altix, SMPs, clusters of SMPs
  - Support for Intrepid, C++, mixed MPI
- GASNet communication layer [Dan]
  - Previously existing ports: IBM LAPI, Myrinet GM, Quadrics Elan-3
  - FY04 Ports: Infiniband, UDP, Shmem, GM+threads, Elan-4
  - Research on support for pinning-based networks such as Infiniband and Myrinet
  - Third party ports: SCI by UFL



# Recent Progress on Compiler



- Enabled optimizations in Open64 base
- Static analyses for parallel code
  - Understand when code motion is legal without changing views from other processors
  - Extended cycle detection to arrays with three different algorithms
- Message Coalescing
  - Replacing small messages with larger ones [Wei]
- Message strip-mining
  - Find optimal message size for pipelining [Costin]
- Experiments with vectorization on the X1
  - [Christian and Wei]



# Recent Progress on Applications



- NAS PB-size problems
  - Berkeley NAS MG avoids most global barriers and relies on UPC relaxed memory model [Parry]
  - Berkeley NAS CG has several versions, including simpler, fine-grained communication
  - Berkeley NAS FT [Christian]
  - Sparse triangular solve [Rajesh]
- Algorithms that are challenging in MPI
  - 2D Delauney Triangulation [SIAM PP '04] [Parry]
  - AMR in UPC: Chombo Poisson solver [Mike]
  - Investigation into AMR potential



# Progress on the Language

---



- Specification of UPC memory model in progress
  - Joint with MTU
  - Behavioral spec [Dagstuhl03]
- UPC IO nearly finalized
  - Joint with GWU and ANL
- UPC Collectives V 1.0 finalized
  - Effort led by MTU
  - Optimized version on GASNet underway [Paul]
  - Investigation of automatic tuning [Rajesh]
- Improvements to UPC Language Spec
  - Led by IDA



# External Activities

---



- Participation in UPC bi-annual consortium meeting
- 4 Tutorials: PSC, SIAM PP04, SC02/SC03, IPDPS03
- UCB Parallel Computing course
  - Assignment using 4 problems in 2 PGAS languages
  - Slides used at elsewhere (UCSB,...)
- 10 Presentations at workshops, conferences, and panels, poster sessions
- 11 Publications
  - 7 in refereed conferences/journals
  - 4 are language or runtime interface specifications





# Presentation Details



1. **Performance Workshop: Programming Models Panel, K. Yelick, April 2004.**
2. **PSC Petamethods Workshop: "UPC" B. Carlson, K. Yelick, April 2004.**
3. **Open64 Workshop: "Berkeley UPC Compiler," C. Iancu, 2004.**
4. **SIAM PP04 Tutorial: "PGAS Languages," T. El-Ghazawi, K. Yelick, B. Carlson, B. Numerich, Feb. 2004.**
5. **SIAM PP04 talk: "Delaunay Triangulation in UPC," P. Husbands, Feb. 2004.**
6. **SIAM PP04 Poster: "Automatic Tuning of Collectives," R. Nishtala et al, Feb. 2004.**
7. **NRC/LBNL Site visit: "Programming Language Issues," K. Yelick, Jan. 2004.**
8. **SC03 Workshop on Petaflops Programming: "What is wrong with MPI for Petaflops Programming?," Nov. 2003.**
9. **SC03 Poster: "GASNet," D. Bonachea, J. Duell, P. Hargrove, Nov. 2003.**
10. **Dagstuhl workshop on Memory Consistency Models: A Behavioral Model for the UPC Language, K. Yelick, Oct. 2003.**



# 2003/4 Publications



1. ***Evaluating Support for Global Address Space Languages on the Cray X1***  
C. Bell, W. Chen, D. Bonachea, K. Yelick. ICS 2004 (to appear).
2. ***Message Strip Mining Heuristics for High Speed Networks***  
C. Iancu, P. Husbands, W. Chen. VECPAR 2004 (to appear).
3. ***Problems with using MPI 1.1 and 2.0 as compilation targets for parallel language implementations***  
D. Bonachea and J. Duell. 2nd Workshop on Hardware/Software Support for High Perf. Scientific and Engineering Computing, SHPSEC-PACT03. (Also to appear in IJHPCN.)
4. ***Polynomial-time Algorithms for Enforcing Sequential Consistency in SPMD Programs with Arrays***  
W. Chen, A. Krishnamurthy, K. Yelick. 16th International Workshop on Languages and Compilers for Parallel Computing (LCPC), 2003.
5. ***A Performance Analysis of the Berkeley UPC Compiler***  
W. Chen, D. Bonachea, J. Duell, P. Husbands, C. Iancu, K. Yelick. 17th Annual International Conference on Supercomputing (ICS), 2003.
6. ***A New DMA Registration Strategy for Pinning-Based High Performance Networks***  
C. Bell and D. Bonachea. Communication Architecture for Clusters (CAC'03), 2003.
7. ***An Evaluation of Current High-Performance Networks***  
C. Bell, D. Bonachea, Y. Cote, J. Duell, P. Hargrove, P. Husbands, C. Iancu, M. Welcome, K. Yelick. 17th Intern'l Parallel and Distributed Processing Symposium (IPDPS), 2003.
8. ***Proposal for Extending the UPC Memory Copy Library Functions, v0.7***  
D. Bonachea, UPC community forum, 2004.
9. ***A Proposal for a UPC Memory Consistency Model, v1.0***  
K. Yelick, D. Bonachea and C. Wallace, LBNL Tech Report LBNL-54983.
10. ***UPC-IO: A Parallel I/O API for UPC, v1.0pre10***  
T. El-Ghazawi, F. Cantonnet, P. Saha, R. Thakur, R. Ross, D. Bonachea, <http://upc.gwu.edu>.
11. ***GASNet Specification, v1.1***  
D. Bonachea. U.C. Berkeley Tech Report CSD-02-1207.



# Schedule



## Introduction

- 8:15 Coffee
- 8:30 Overview of Berkeley UPC

## Runtime Session

- 8:40 Runtime (Duell)
- 9:00 Gasnet (Bonachea)
- 9:40 Collectives (Hargrove, Nishtala)

## 10:00 - Break (15 min)

## Compiler Session

- 10:15 Berkeley UPC on the X1 (Bell, Chen)
- 11:00 Message Coalescing (Chen)
- 11:30 Message Stripmining (Iancu)

## 12:00 - Lunch (45 Minutes)

## 12:00 - Lunch (45 Minutes)

- 12:45 AMR in UPC (Welcome)
- 1:10 FFT (Bell)
- 1:20 Sparse Triangular Solve (Nishtala)
- 1:30 Scaling UPC Applications (Husbands)

## Summary

- 1:45 Future Directions (Kathy Yelick)
- 2:30 Break and move to 50F conference room
- 2:45 Discussion