Teaching parallelism in an interdisciplinary scientific computing programme

Rob H. Bisseling

Mathematical Institute, Utrecht University, the Netherlands

February 24, 2010

Outline

Past
- History
- Hardware
- Software

Present
- Master in SC
- Parallel SC

Future

Universiteit Utrecht

Past: parallel computing courses in Utrecht

- History
- Hardware
- Software

Present: Utrecht master programme in scientific computing

- Master in SC
- Parallel SC

Future: revolutionary changes

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Universiteit Utrecht

Parallel algorithms history in Utrecht

- 1990–1992 Linear Algebra for Supercomputers by Henk van der Vorst
- 1993–1995 Linear Algebra for Supercomputers by RB
- 1996–2002 Parallel Algorithms for Supercomputers by RB
- 2003–present Parallel Algorithms by RB
- 10–12 students each year take the course, from mainly maths, but also physics, and computer science
- Level: first year of MSc. Language: English (since 2004)
- 2009/2010: became part of Dutch national mathematics master (http://www.mastermath.nl), now 17 students
- Follow-up course High-performance scientific computing: individual project

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Hardware used over the years

- 1995–1997 Cray T3E (Delft)
- 1998–2000 SGI Origin 2000 (SARA computing centre, Amsterdam)
- 2000–2003 SGI Origin 3800 (SARA)
- 2003–2006 SGI Altix 3700 (SARA)
- 2007–present IBM Power 5/6 (SARA)
- Supercomputer access attracts students
- We use one shared student account for teaching on the machine which is officially a research machine. So far we have behaved well!
- We also used: a cluster of workstations (painful);
  Linux PCs (good for code development);
  and even Linux emulated under Windows (it does work).

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Dutch National Supercomputer Huygens

IBM Power 6 computer with 3328 processing cores named after Christiaan Huygens, Dutch astronomer who in 1655 proposed the form of the rings around Saturn

Outline

Past

History

Hardware

Software

Present

Master in SC

Parallel SC

Future

Universiteit Utrecht

Software used over the years

- Bulk Synchronous Parallel (BSP) programming is used. BSP seems to have caught on, as everybody urges us to abandon it!
- 1995–1997 Oxford BSP library (Richard Miller, 6 primitives)
- 1998–2005 Oxford BSP Toolset (Jon Hill et al., 20 primitives, native implementation of BSPlib)
- 2006–present BSPonMPI (Wijnand Suijlen, BSPlib on top of MPI)
- BSPlib is great for students, as it is easy to learn
- But some researchers of parallel graph algorithms and sparse matrix computations also like BSPlib a lot, especially because of the `bsp_send` primitive, which does a lot of buffering and communication optimisation for you. This primitive came from the Green BSP library (Mark Goudreau et al. 1996)

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

BSPonMPI

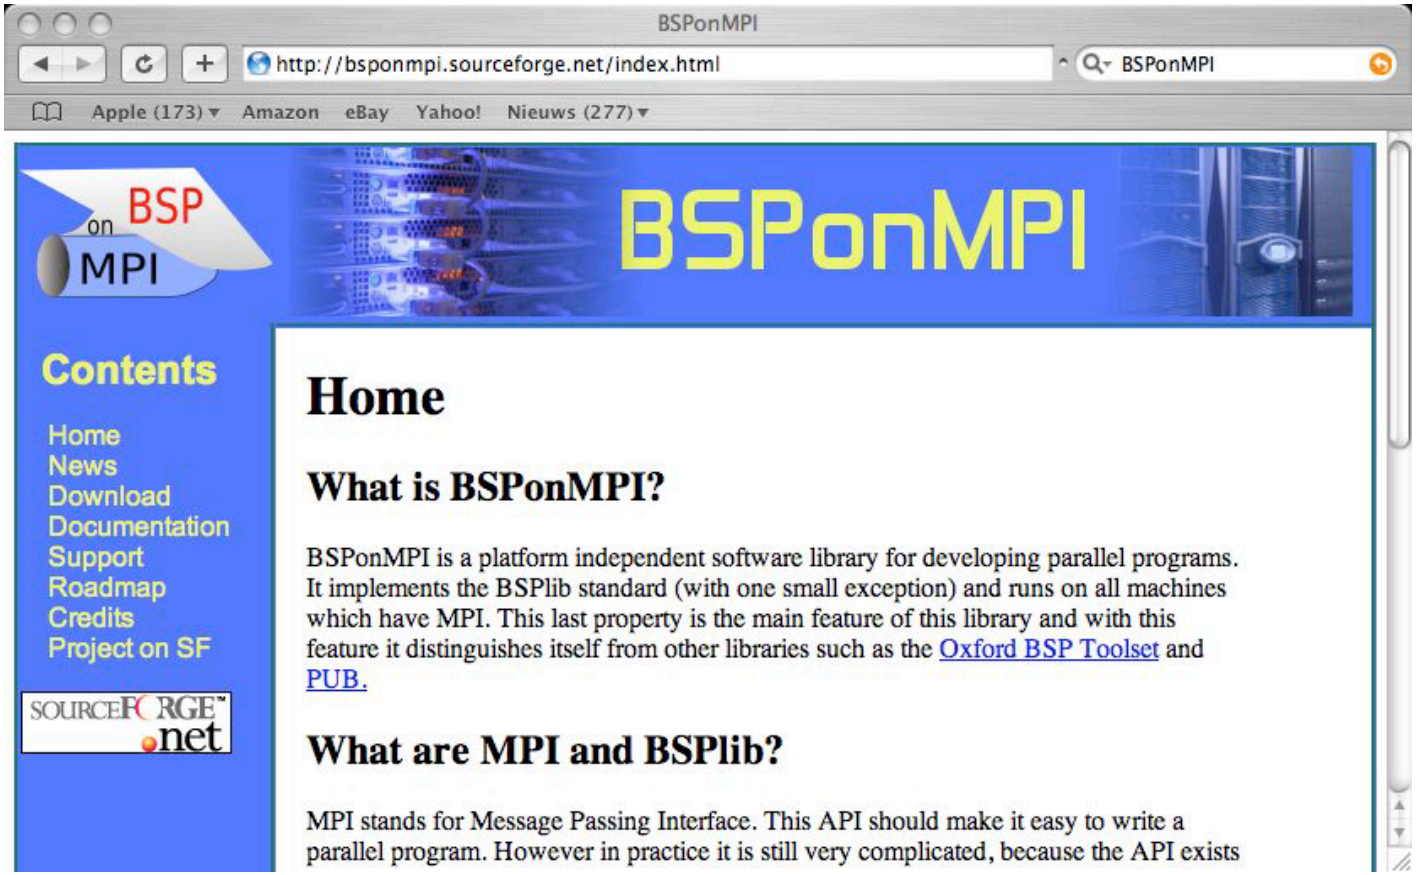

BSPonMPI by Wijnand Suijlen, version 0.3 available since February 20, 2010 from http://www.bsp-worldwide.org/
Developed as a BSc thesis in Utrecht, written in C++

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Universiteit Utrecht

Utrecht MSc programme in scientific computing

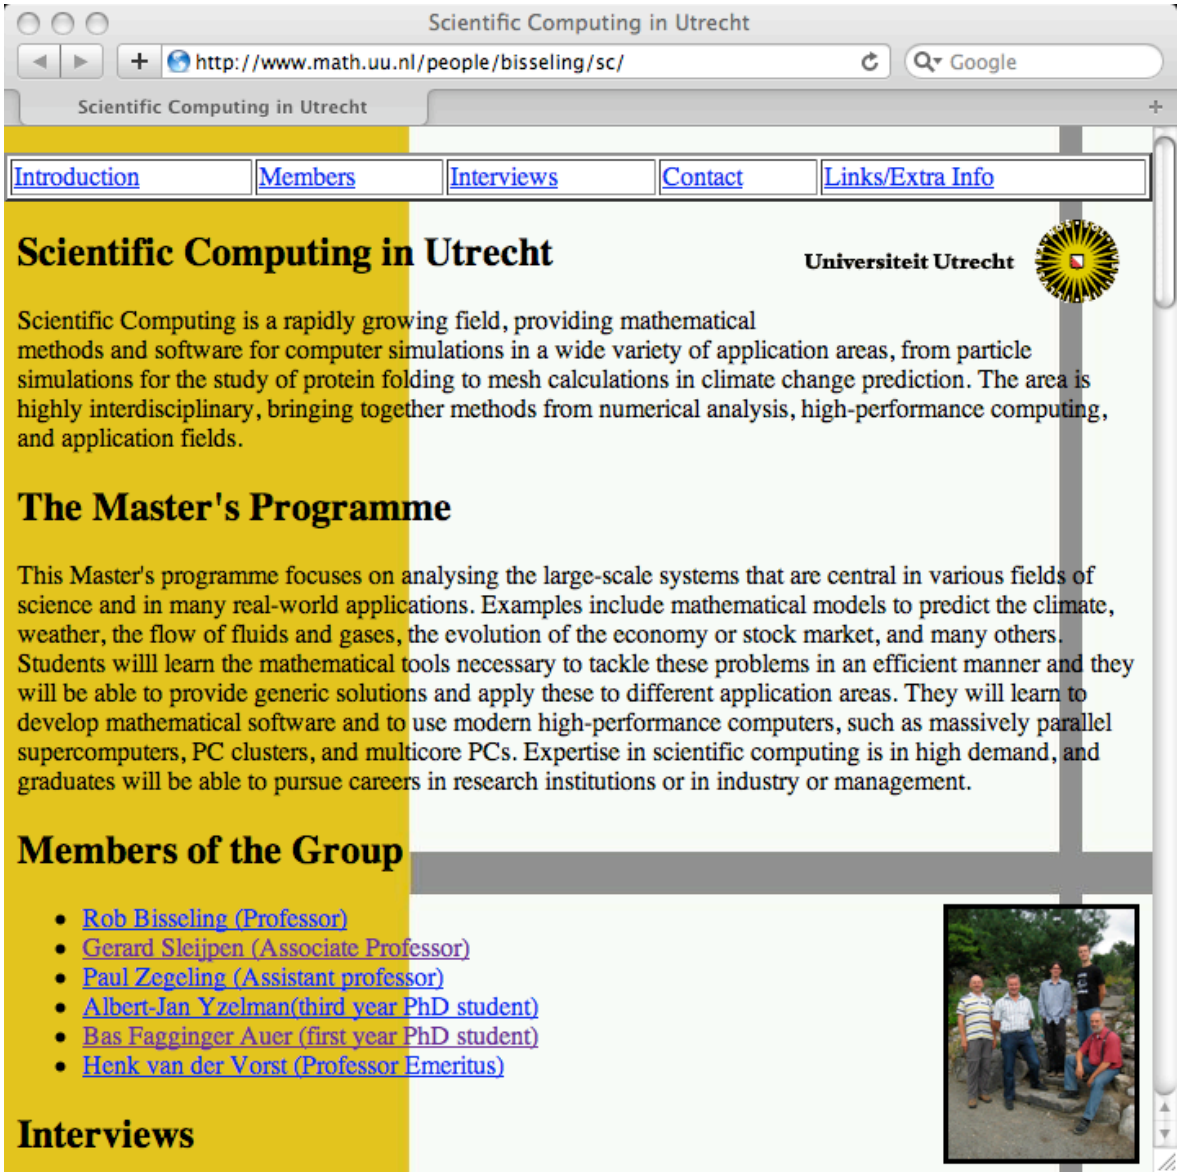

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Universiteit Utrecht

Programme consists of 10 MSc courses + thesis

- Numerical linear algebra (Gerard Sleijpen & Martin van Gijzen)
- Scientific computing laboratory (Albert-Jan Yzelman)
- Parallel algorithms (Rob Bisseling)
- Numerical PDEs (Paul Zegeling/Rob Stevenson)
- Modelling and simulation (Gerard Barkema)
- choice of: Computational biology (Paulien Hogeweg) or Climate modelling (Henk Dijkstra)
- 4 elective courses
- thesis (9 months research)

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Interdisciplinary programme

- Home base: Mathematics Institute
- Joint thesis supervision:
  - oceanographic modelling using GPUs
  - X-ray spectroscopy computations for space research using OpenMP
  - computer simulation of colloids on a PC
- Courses in other disciplines as electives
- Students from other disciplines take a specific course from the SC offerings
- Strong students from the Utrecht BSc/MSc programme with double major mathematics/physics are often interested, discovering the joys of computation later in life

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Universiteit Utrecht

Parallel Algorithms course: book

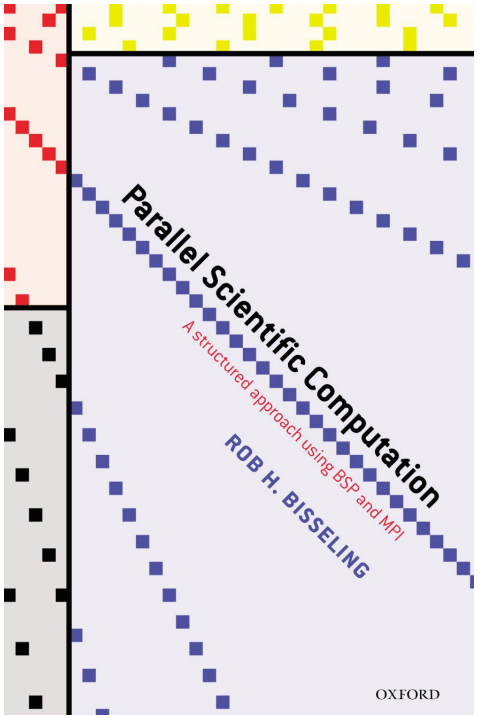

- Developed lecture material from 1993–2003. Appeared as a book in 2004.
- The book took a long time to write, but then it decays also slowly, if at all. Machines come and go, but LU decomposition stays.

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Universiteit Utrecht

Parallel Algorithms course: additional material

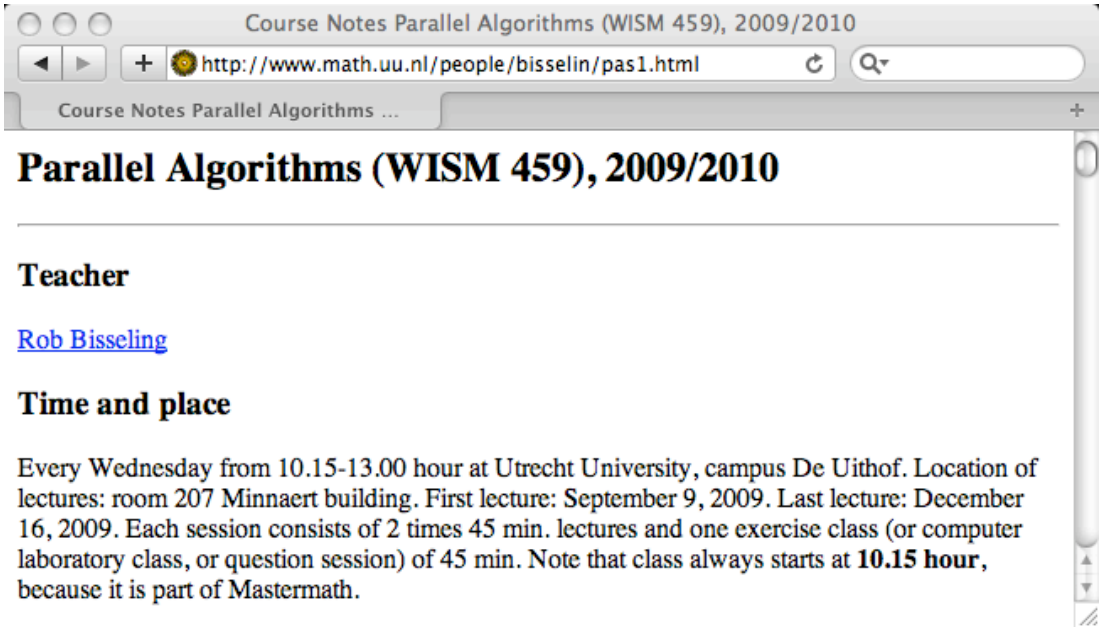

- Software package: BSPedupack, contains BSPlib programs for benchmarking, inner-product computation, dense LU decomposition, Fast Fourier Transform, sparse matrix–vector multiplication.
- Software package: MPIedupack (MPI in BSP style).
- Slides, including LaTeX sources (in Prosper, not yet Beamer). 26 lectures of 45 min. Good for other teachers, and to get some students started in LaTeX.

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Universiteit Utrecht

Parallel Algorithms course: computer laboratory

- Introductory lab class: run benchmarks to obtain BSP parameters for computation, communication, synchronisation
- First assignment: create list of primes by parallel Eratosthenes sieve
- Second assignment: choice of exercises, varies every year. Write a parallel program for:
  - Dense Cholesky
  - Compression by the Ziv-Lempel algorithm (LZ77), find repeats in "yabbadabbadoo"
  - Decimals of π, using FFT to speed up large integer arithmetic (a lot of work, only for the ambitious). You get a good grade if you reach a million correct decimals
  - Sparse Conjugate Gradient
  - Wavelet transform
- Students may work in pairs on a program, but must hand in individual reports, to deter free riders while fostering collaboration

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Visualising partitioning for parallelism

Splitting the sparse matrix `lns3937` into 5 parts. Film made using MondriaanMovie by Bas Fagginger Auer, part of Mondriaan v3.0, to be released Spring 2010. See the poster! Film can be found at http://www.math.uu.nl/people/bisseling/oratiefilmpje.avi

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Pictures of a revolution: the guillotine

King Louis XVI of France executed at the Place de la Concorde in Paris, January 23, 1793. Source: http://www.solarnavigator.net/history/french_revolution.htm

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Universiteit Utrecht

The parallel computing revolution

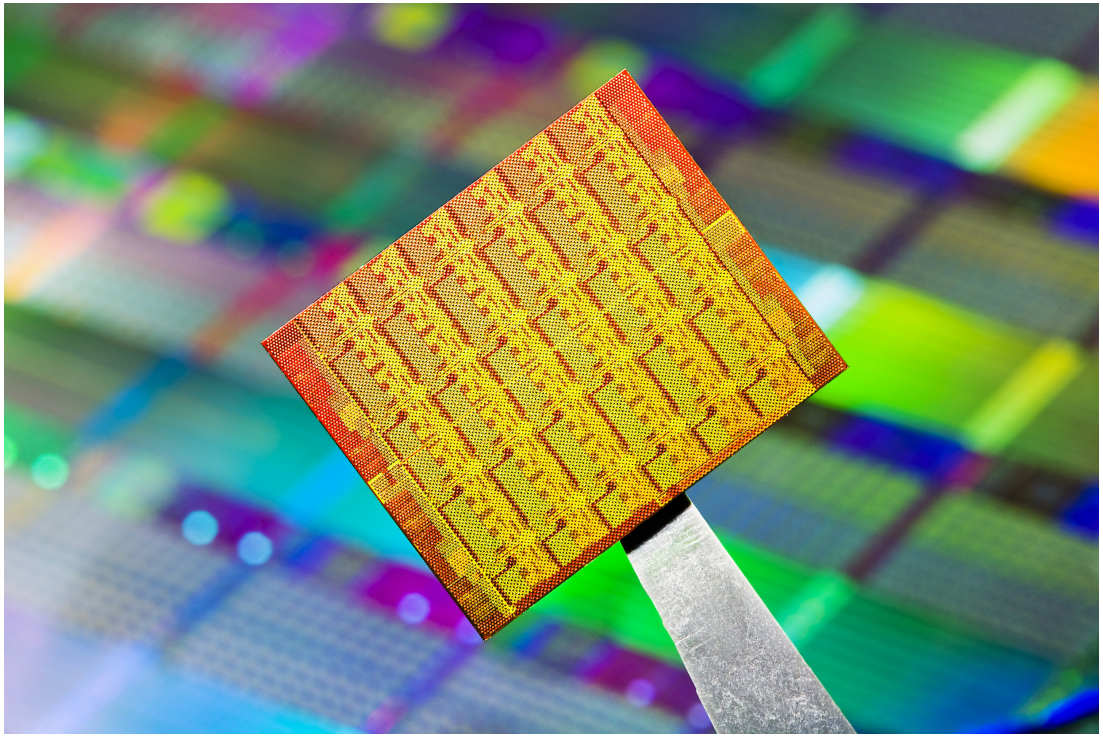

Intel Single-Chip Cloud computer with 48 cores, announced December 2, 2009. Energy consumption from 25 to 125 Watt, depending on use. Each pair of cores has a variable clock frequency. Source: http://techresearch.intel.com

Outline

Past

History

Hardware

Software

Present

Master in SC

Parallel SC

Future

Since 2008: what a typical student takes to class

- Parallel computing on every student's lap: e.g. a MacBook Pro with an Intel Core 2 Duo processor.
- Next year quadcore, octacore?
- Should this be our target architecture, or alternatively, our development tool for running on supercomputers? My view: both.

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Difficulties students encounter

- Parallel software has been less tested than sequential software.
- Students are haunted by bugs, usually their own, but you never know ...
- They need to install Linux on a PC or buy a Mac with Unix under the hood, and they have to install OpenMPI and BSPonMPI. They become mature this way.
- They need to learn the interactive and batch system on a supercomputer.
- Things change all the time. The teacher does not know all either.

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future

Food for thought

- Why are computer science curricula not radically changed, by reinstating parallel computing courses? These had been removed over the past decade in many universities.
- Do we have to wait with teaching parallel algorithms until a clear picture has emerged about a consensus parallel programming model? This may never happen.
- Simpler models are needed, even simpler than BSP.
- Hardware-oblivious approaches are the way to go, like cache-oblivious reordering for sparse matrix–vector multiplication (Yzelman & Bisseling, *SISC* 2009).
- Visualising algorithms becomes more and more important, in teaching the YouTube generation, but also in research.

Outline
Past
History
Hardware
Software
Present
Master in SC
Parallel SC
Future