

From 20 Years of Beowulf MPI & Cluster Computing

Aiichiro Nakano

*Collaboratory for Advanced Computing & Simulations
Department of Computer Science
Department of Physics & Astronomy
Department of Chemical Engineering & Materials Science
Department of Biological Sciences
University of Southern California*

Email: anakano@usc.edu

<http://dl.acm.org/citation.cfm?id=2737911&CFID=717321656&CFTOKEN=85686577>

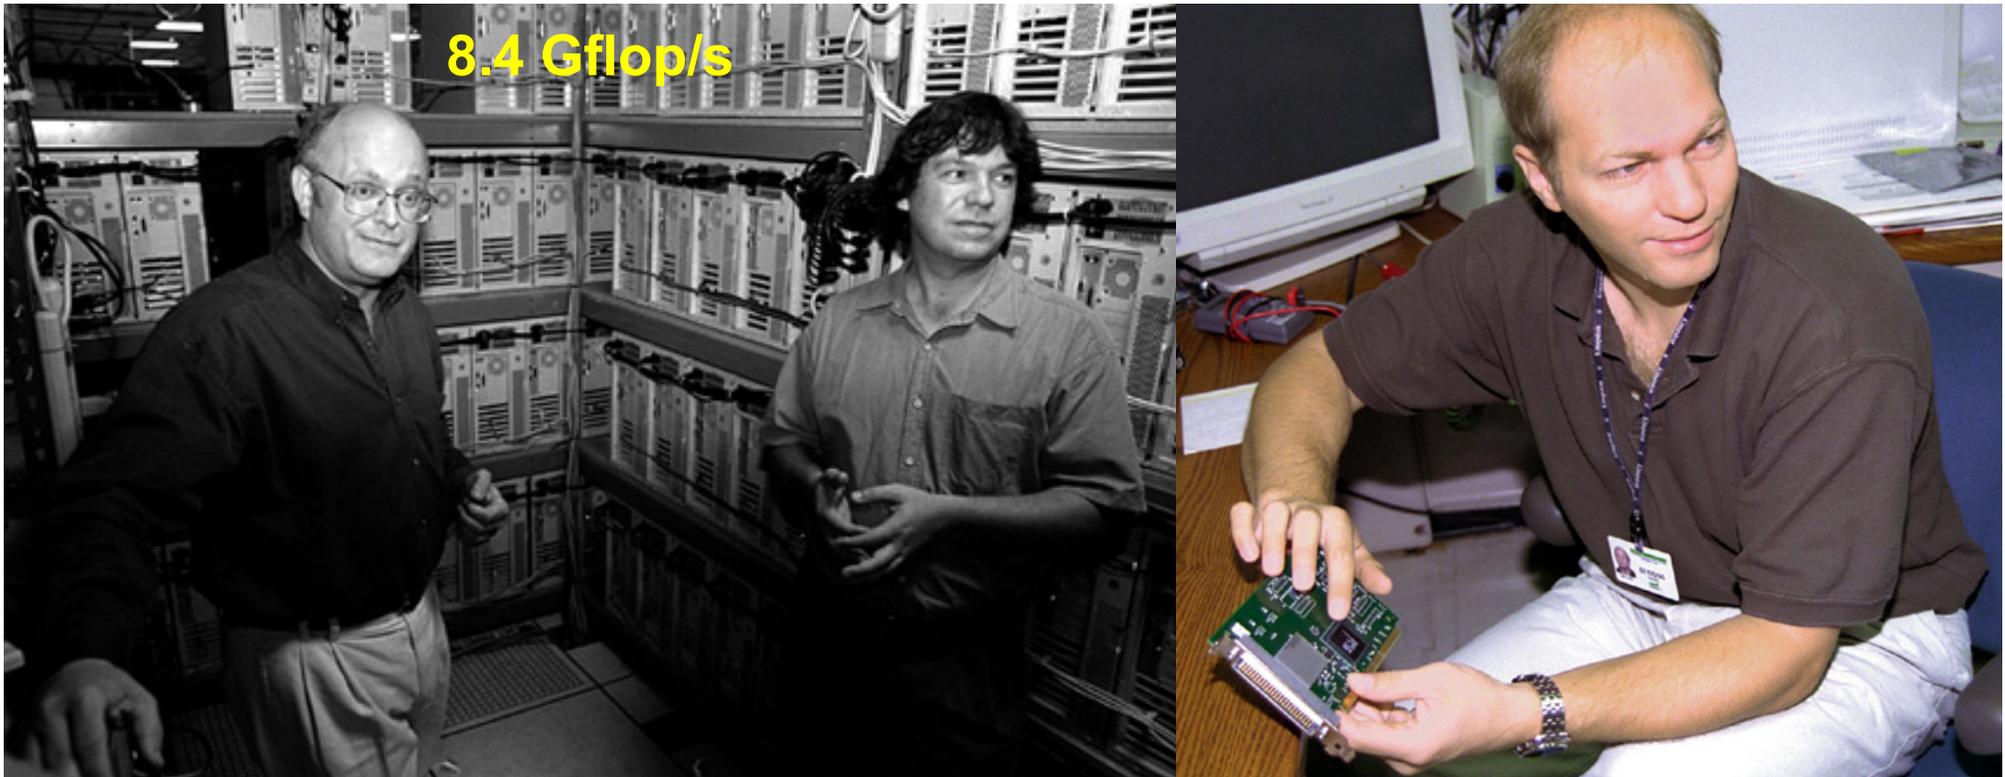


<http://www.crest.iu.edu/beowulf14>



Beowulf Cluster

- Parallel computer made of commodity components
- Open software (Linux; network driver developed)
- Started by Thomas Sterling & Don Becker in '94



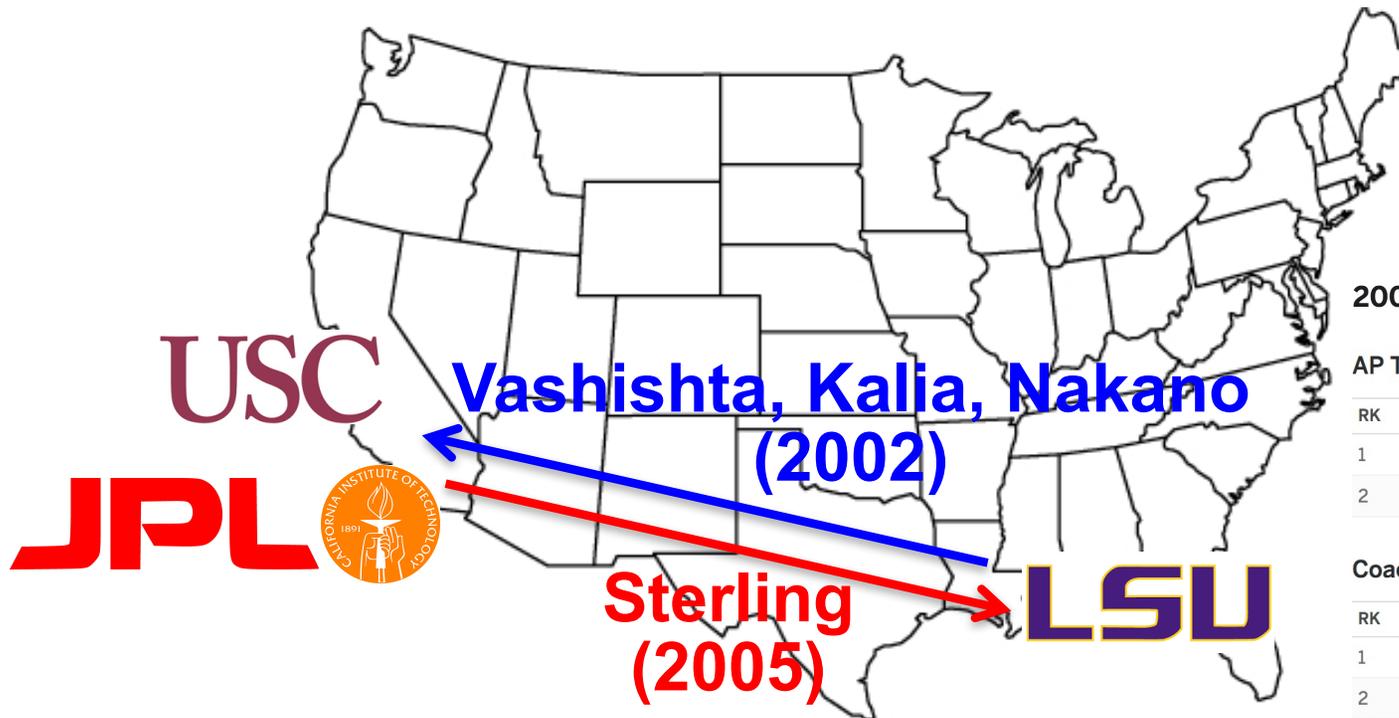
(**Left**) Caltech's Thomas Sterling (left) and John Salmon discuss the building of *Naegling*, one of the largest Beowulf computers. Each of its 120 processors performs 70 million floating-point operations per second (flops) across a range of applications. (**Right**) Don Becker, Goddard Space Flight Center, holds a personal computer network adapter. He has developed, and is constantly updating, software to drive nearly all adapters for use with the Linux operating system.

<http://www.hq.nasa.gov/hpcc/insights/vol7/beowulf.htm>

Thomas Sterling and Me

Top500 List - November 2002

Rank	Site	System	Cores	Rmax (GFlop/s)	Rpeak (GFlop/s)	Power (kW)
1	Japan Agency for Marine -Earth Science and Technology Japan	Earth-Simulator NEC	5120	35860.0	40960.0	3200
17	Louisiana State University United States	SuperMike - P4 Xeon 1.8 GHz - Myrinet Atipa Technology	1024	2207.0	3686.4	



2003 College Football Rankings - Postseason

AP Top 25

RK	TEAM	REC	PTS	TREND
1	USC (48)	12-1	0	—
2	LSU (17)	13-1	0	—

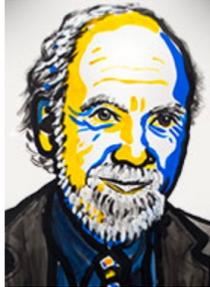
Coaches Poll

RK	TEAM	REC	PTS	TREND
1	LSU (60)	13-1	1572	—
2	USC (3)	12-1	1514	—

Digress: LIGO



© Nobel Media. Ill. N. Elmehed
Rainer Weiss
Prize share: 1/2



© Nobel Media. Ill. N. Elmehed
Barry C. Barish
Prize share: 1/4



© Nobel Media. Ill. N. Elmehed
Kip S. Thorne
Prize share: 1/4

The Nobel Prize in Physics 2017 was divided, one half awarded to Rainer Weiss, the other half jointly to Barry C. Barish and Kip S. Thorne *"for decisive contributions to the LIGO detector and the observation of gravitational waves"*.



Gravitational-wave research in Louisiana



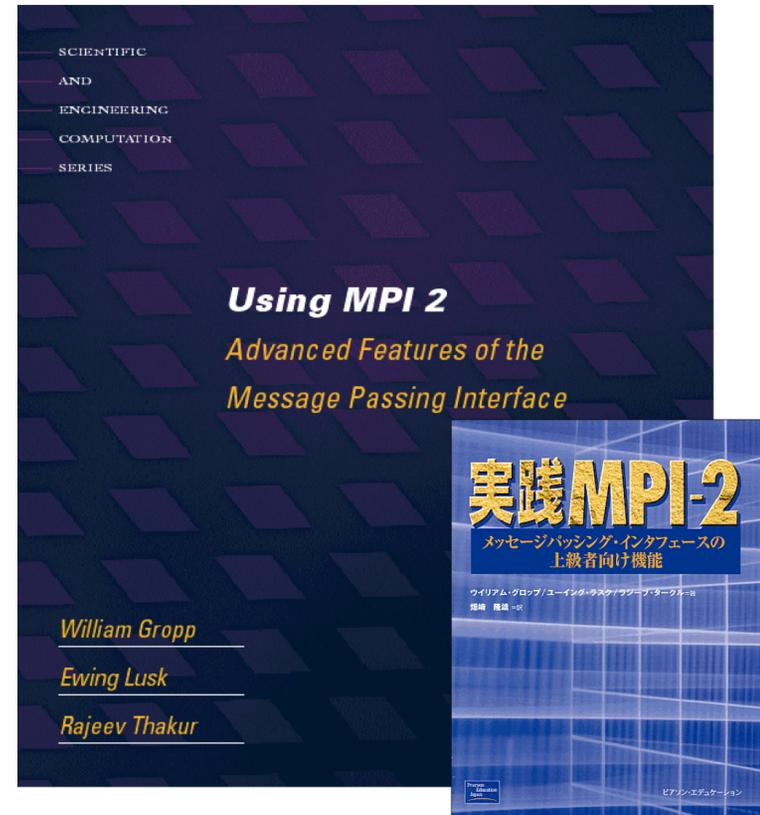
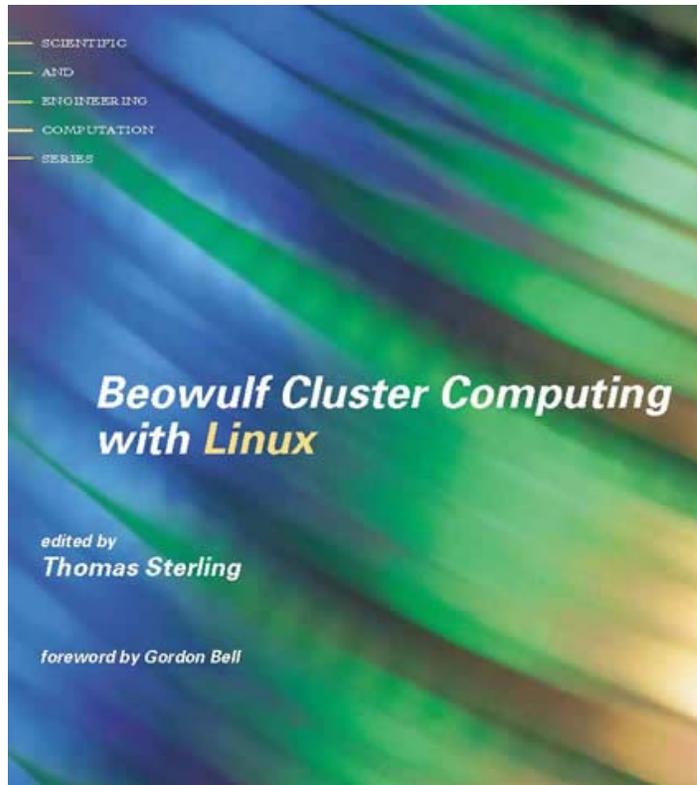
LIGO (Laser Interferometer Gravitational-Wave Observatory) in Livingston, LA

- Early 70's: Bill Hamilton started building a cryogenic gravitational-wave detector at LSU
- Mid 90's: LIGO construction started



20 Years of Beowulf and MPI

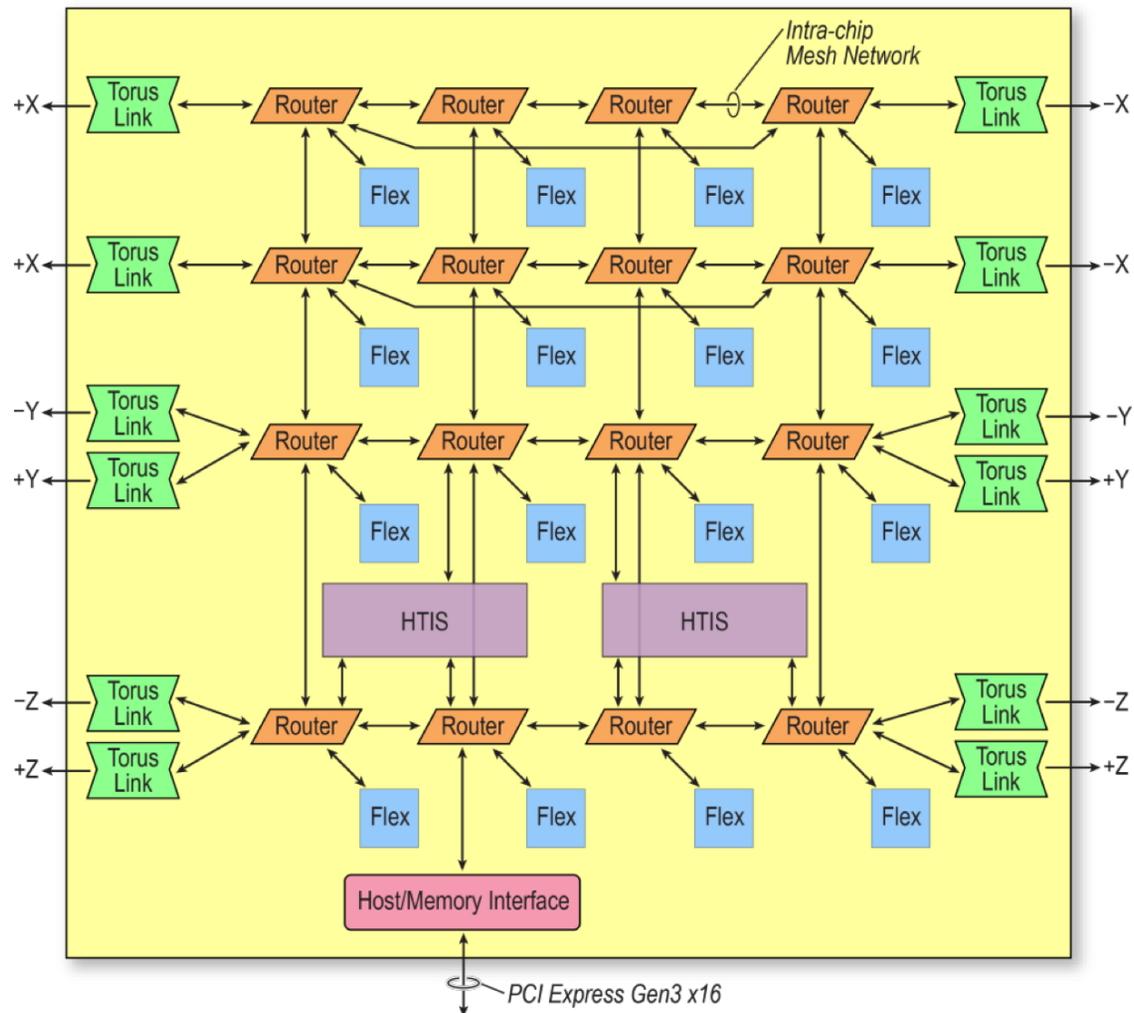
- Dominant parallel-computing paradigm for the past 20 years:
Distributed processes communicating via message passing



- First Beowulf ('94)
- Evolving by embracing multicore & accelerators per computing node
- MPI 1 ('94)
- Evolving: *Using Advanced MPI*

Counter-Approach: Anton 2

- Unified on-chip & inter-node networks
- New algorithm: Gaussian series expansion of the Coulombic interaction (no FFT required)



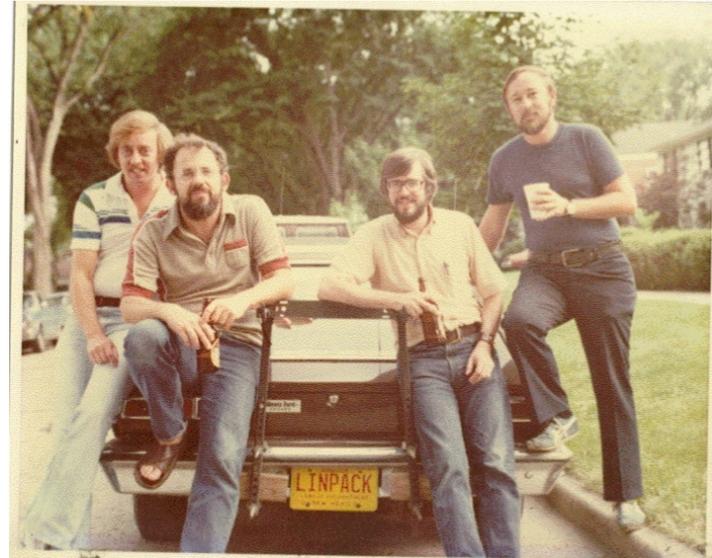
Massive Data Analytics

- Scalable data-analytics/machine-learning algorithms are critically needed, e.g., $O(N^2) \rightarrow O(N)$ pair statistics
- **Seven computational giants**
 1. Basic statistics
 2. Generalized N -body problem
 3. Graph-theoretic computations
 4. Linear algebraic computations
 5. Optimization
 6. Integration
 7. Alignment problems

National Research Council, *Frontiers of Massive Data Analytics* ('13)
<https://www.nap.edu/catalog/18374/frontiers-in-massive-data-analysis>

$O(N^3)$ Linpack to $O(N)$ HPCG

- High performance conjugate gradient (HPCG) proposed toward exaflop/s, but ...

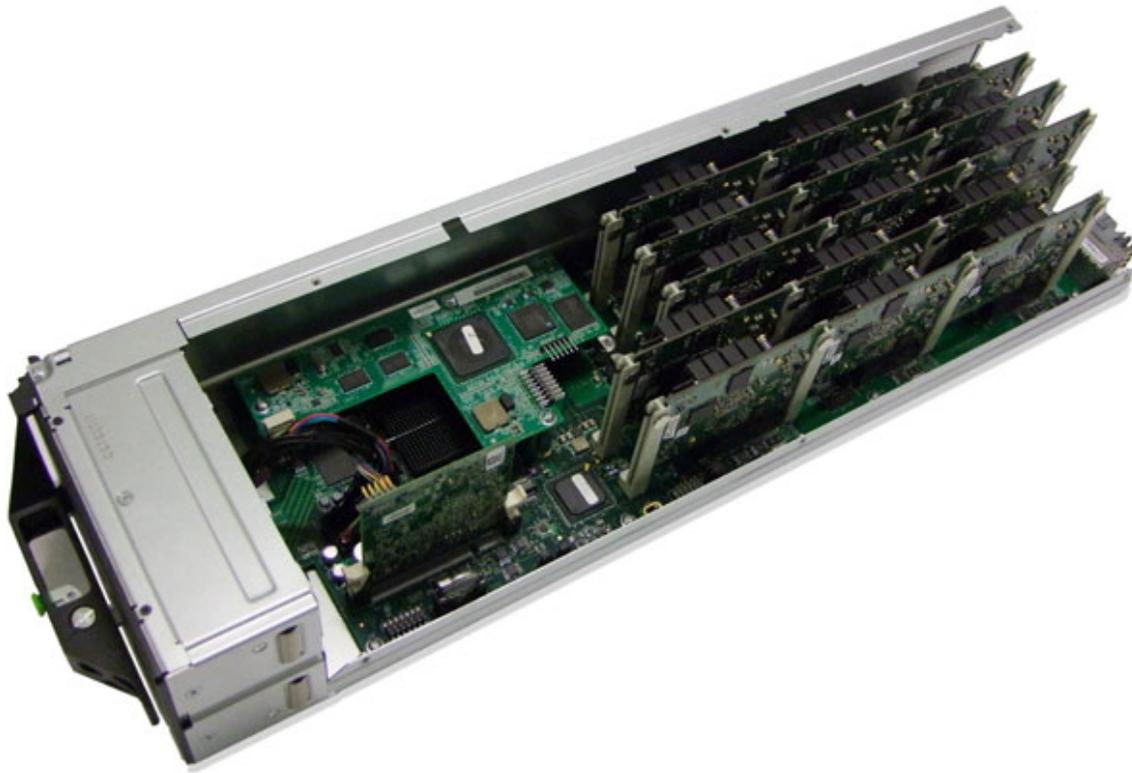


Site	Computer	Cores	HPL Rmax (Pflops)	HPL Rank	HPCG (Pflops)	HPCG/HPL
NSCC / Guangzhou	Tianhe-2 NUDT, Xeon 12C 2.2GHz + Intel Xeon Phi 57C + Custom	3,120,000	33.9	1	.580	1.7%
RIKEN Advanced Inst for Comp Sci	K computer Fujitsu SPARC64 VIIIfx 8C + Custom	705,024	10.5	4	.427	4.1%
DOE/OS Oak Ridge Nat Lab	Titan, Cray XK7 AMD 16C + Nvidia Kepler GPU 14C + Custom	560,640	17.6	2	.322	1.8%
DOE/OS Argonne Nat Lab	Mira BlueGene/Q, Power BQC 16C 1.60GHz + Custom	786,432	8.59	5	.101#	1.2%
Swiss CSCS	Piz Daint, Cray XC30, Xeon 8C + Nvidia Kepler 14C + Custom	115,984	6.27	6	.099	1.6%

HPL
HPCG

Smartphones as Exascale Nodes

- Building an exaflop/s computer from commodity components (again, but with smartphones this time?)



Blades of Glory: Mont-Blanc's prototype contains 15 nodes made up of ARM-core processors.

IEEE Spectrum (May '14)

Or Raspberry Pi?

