

3rd Workshop on Personal Computer based Networks Of Workstations (PC-NOW 2000)

Clusters composed of fast personal computers are now well established as cheap and efficient platforms for distributed and parallel applications. The main drawback of a standard NOWs is the poor performance of the standard inter-process communication mechanisms based on RPC, sockets, TCP/IP, Ethernet. Such standard communication mechanisms perform poorly both in terms of throughput as well as message latency.

Several prototypes developed around the world have proved that re-visiting the implementation of the communication layer of a standard Operating System kernel, a low cost hardware platform composed of only commodity components can scale up to several tens of processing nodes and deliver communication and computation performance exceeding the one delivered by the conventional high-cost parallel platforms.

This workshop provides a forum to discuss issues related to the design of efficient NOW/Clusters based on commodity hardware and public domain operating systems as compared to custom hardware devices and/or proprietary operating systems.

Workshop Organizers

G. Chiola (DISI, U. Genoa, I)
G. Conte (CE, U. Parma, I)
L.V. Mancini (DSI, U. Rome, I)

Sponsors

IEEE TFCC
(Task Force on Cluster Computing)

Program Committee

Program Chair: C. Anglano (U. Piemonte Or., I)
M. Baker (CSM, U. Portsmouth, UK)
L. Bougè (ENS Lyon, F)
G. Chiola (DISI, U. Genoa, I)
G. Ciaccio (DISI, U. Genoa, I)
G. Conte (CE, U. Parma, I)
H.G. Dietz (ECE, Purdue U., USA)
W. Gentsch (GENIAS Software GmbH, D)
G. Iannello (DIS, U. Napoli, I)
Y. Ishikawa (RWCP, J)
K. Li (Princeton U., USA)
L.V. Mancini (DSI, U. Roma 1, I)
T.G. Mattson (Intel Corp., USA)
W. Rehm (Informatik, T.U. Chemnitz, D)
P. Rossi (ENEA HPCN, Bologna, I)
P. Roe (Queensland U. of Tech., AUS)
D.B. Skillikorn (Queens U., CAN)
D. Tavangarian (Informatik, U. Rostock, D)
B. Tourancheau (LHPC, U. Lyon, F)

Referees

C. Anglano	W. Gentsch	W. Rehm
O. Aumage	G. Iannello	P. Roe
M. Baker	Y. Ishikawa	P. Rossi
G. Chiola	L.V. Mancini	D. Tavangarian
G. Ciaccio	T.G. Mattson	B. Tourancheau
G. Conte	J.-F. Mehaut	R. Westrelin
M. Fischer	R. Namyst	

Accepted Papers

Session 1: Cluster Interconnect Design and Implementation

- M. Trams, W. Rehm, D. Balkanski, and S. Simeonov “Memory Management in a combined VIA/SCI Hardware”
- M. Fischer, et al. “ATOLL, a new switched, high speed Interconnect in comparison to Myrinet and SCI”
- R.R. Hoare “ClusterNet: An Object-Oriented Cluster Network”

Session 2: Off-the-shelf Clusters Communication

- M. Baker, S. Scott, A. Geist, and L. Browne “GigaBit Performance under NT”
- H.A. Chen, Y.O. Carrasco, and A.W. Apon “MPI Collective Operations over IP Multicast”

Session 3: Multiple Clusters and Grid Computing

- S. Lalis, and A. Karipidis ‘An Open Market-Based Architecture for Distributed Computing”
- M. Barreto, R. Avila, and Ph. Navaux “The MultiCluster Model to the Integrated Use of Multiple Workstation Clusters”

Session 4: Data Intensive Applications

- S.H. Chung, et al. “Parallel Information Retrieval on an SCI-Based PC-NOW”
- M. Exbrayat, and L. Brunie ‘A PC-MOW Based Parallel Extension for a Sequential DBMS”

Other Activities

In addition to the presentation of contributed papers an invited talk will be scheduled at the workshop.