

CORDS

—Communication Objects for Real-time Dependable Systems

Franco Travostino

The Open Group
Research Institute

www.opengroup.org/RI/PubProjPgs/CORDS.htm

Mission

First, build real-time, dependable distributed services for mission-critical applications

E.g.: control of factory floors, combat systems, medical tools

- Predictable behaviors, even in the presence of failures
- Real-time, fault tolerant, security properties in a single system
- Configurable and adaptive services throughout the various layers
- Unforgiving mechanical apparatus is often the consumer

Then, leverage services' high configurability and re-use parts to build services for less stringent environments

E.g.: white-board, video-conferencing, Web clusters

- Soft real-time guarantees
- Typically, a human is the consumer, via sophisticated CHIs
- Scalability along various dimensions

CORDS

It builds upon x-kernel
(University of Arizona)

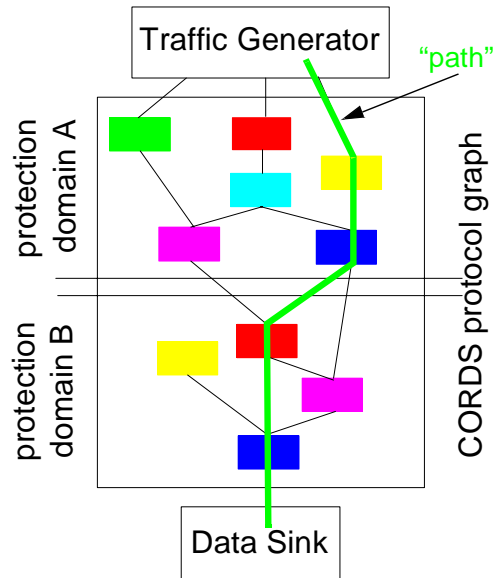
Innovation: Paths for
advanced resource control

Real-time framework

Library of protocols for
clusters and fault tolerance

Few OS dependencies in the
framework only

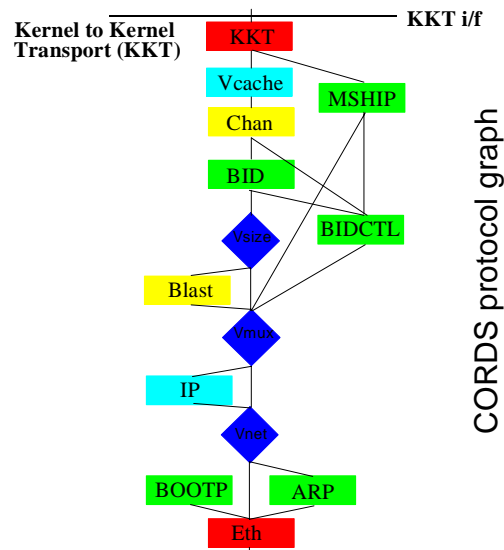
Middleware (MK, NT, UNIX)



Using CORDS for ...

KKT, a **communication fabric**
for clusters

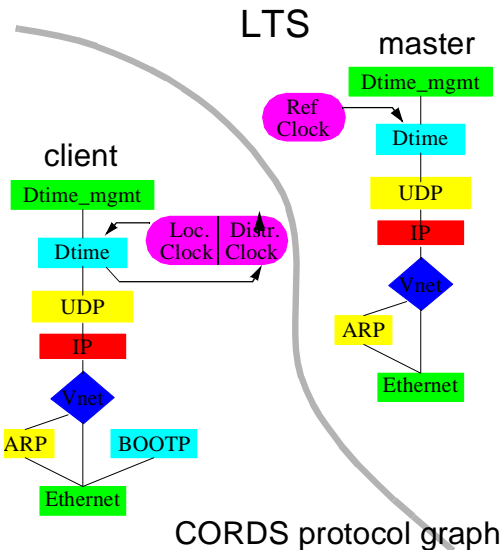
- Designed for unreliable interconnects (i.e., Ethernet)
- It supports and leverages protocol standards (i.e., IP, BOOTP)
- Optimized for very small and large messages (tens of KB)
- Used over Ethernet, FDDI, Myrinet
- Served as transport for MK/AD
- Recently being used to sustain trans-node RPCs



Using CORDS for ...(cont.)

LTS, a **distributed clock service** for external clock synchronization

- Probabilistic clock synchronization algorithm (F. Cristian)
- Upper bound on master/client deviation (unlike NTP, DTS)
- Bounded network traffic
- Failure semantics
- Example: client-master max-dev = +/- 1.25 ms; request interval = 4 secs avg., 400 ms worst case



Using CORDS for ...(cont.)

GIPC, a **group service** that provides ISIS-like functionality

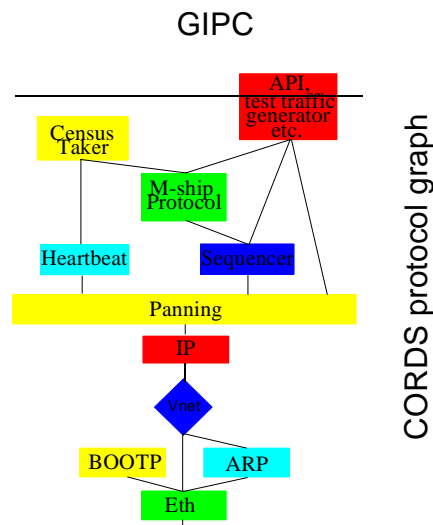
- Process groups
- Several delivery semantics (including atomic broadcasts).

with **predictable** behaviors even in adverse circumstances

- Failures
- Competing traffic

Paths enable fine-grain resource management

Demonstrated sub-second failure detection and recovery



Using CORDS for ...(cont.)

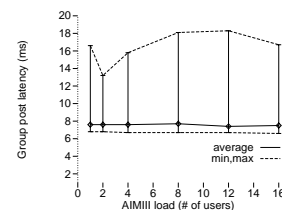

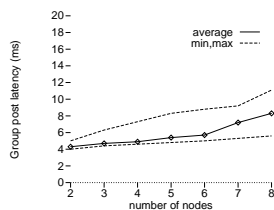

a **real-time characterization framework** that comprises

- Utilities to reconcile log trails from different systems
- CORDS built-in instrumentation and configuration tools
- ORCHESTRA—fault injection facilities contributed and ported to CORDS by University of Michigan

We used it extensively to characterize GIPC

Demos & Measurements

With CORDS, we have built and characterized real systems



CORDS' Impact

Since '94, CORDS has been funded by DARPA and commercial sponsorship

DASCOM has announced use of CORDS for secure routers

Alacron/Honeywell AVIS system uses CORDS w/ Myrinet for node interconnect

Results of CORDS/Myrinet use returned to Myricom

CORDS is base of research at UofA, UMich

CORDS is used in SHAWS, our clustered, HA Web Server

Ball-sorter demonstration featured at several public events

Summary and Conclusions

CORDS results from a LEGO™ approach to Comm Services

It exploits real-time facilities of the host OS

CORDS path programming enables real-time Communication

We have challenged CORDS in different Comm scenarios

All benefit from an open architecture & substantial code reuse

In the CORDS-based GIPC, we have forged the integration of real-time and fault management disciplines

The CORDS framework is the catalyst of this integration

There is now a CORDS community out there

References

- Travostino F., Feeney L., Bernadat P. and Reynolds F., "Building Middleware for Real-Time Dependable Distributed Services," To appear in the *Proceedings of the First IEEE International Symposium on Object-Oriented Real-Time distributed Computing (ISORC)*, Kyoto, Japan, April 20-22, 1998.
- Feeney L., Bernadat P., Travostino F., "Characterizing Group Communication Middleware for Real-time Distributed Systems," WIP Report available from the 18th IEEE Real-Time Systems Symposium, San Francisco, CA, Dec. 1997.
- Reynolds, F., and Travostino, F., "A New WWW Distributed Programming Framework", *Proceedings of etaCOM '96*, Portland, OR, May 1996.
- Travostino, F., and Menze, Edwin F. III, "The CORDS Book," Version 2.0 (formerly known as: "The x-kernel based Communication Subsystem for OSF MK: changes to the x-kernel 3.2 API"), The Open Group Research Institute, Cambridge, MA, September 6, 1996.
(<http://www.opengroup.org/RI/ri/biblio/papers/manual.ps>)
- Travostino, F., Reynolds, F., "An O-O Communication Subsystem for Real-Time Distributed Mach", *Post-Proceedings of the 1994 IEEE Workshop on Object-Oriented Real-Time Dependable Systems*, Irvine, CA, October 1994.
(<http://www.opengroup.org/os/os.coll.papers/Vol4/elementary.frame.ps>)

References (cont'd)

- Travostino, F., Menze, E.F. III, and Reynolds, F., "Paths: Programming with System Resources in Support of Real-time Distributed Applications," *Proceedings of the 2nd IEEE Workshop on Object-Oriented Real-Time Dependable Systems*, Feb 1-2 1996, Laguna Beach, CA.
(<http://www.opengroup.org/RI/ri/biblio/papers/words96.frame.ps>)
- Wells, D. M., "Using Object Frameworks to Enable Real-Time and Dependability in a Modular Operating System," presented at the IEEE Computer Society Second International Workshop on Object-oriented Real-Time Dependable Systems, February 1-2, 1996, Laguna Beach, California.
(<http://www.opengroup.org/RI/ri/biblio/papers/words96.paper.fr3.ps>)