



Enabling Efficient Group Communication in a Collaborative Environment

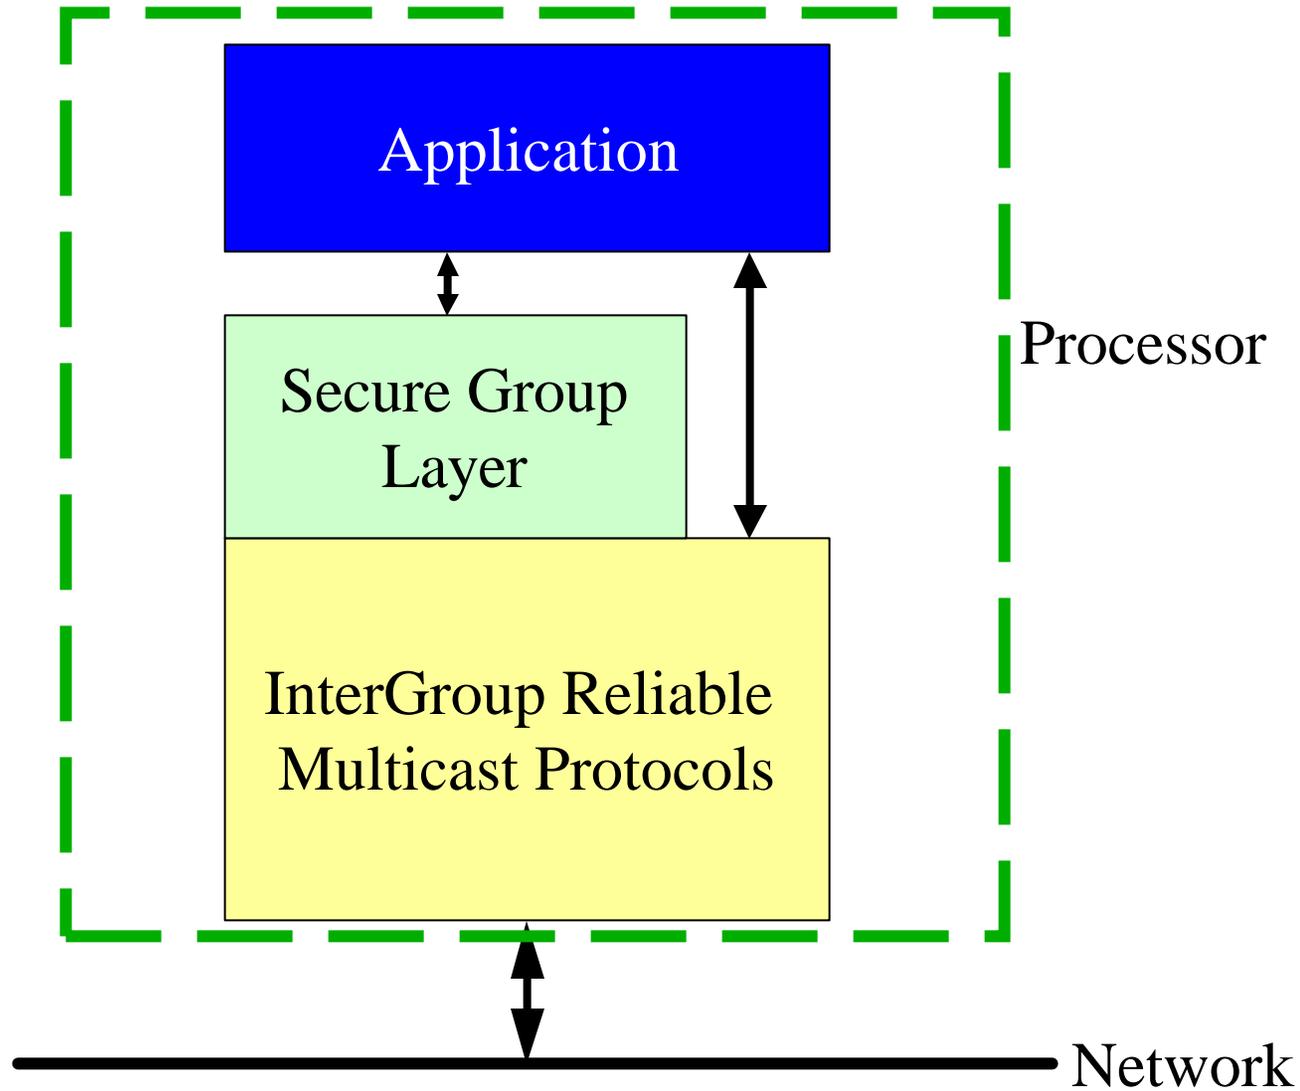
Deb Agarwal (DAAgarwal@lbl.gov),
Karlo Berket (KBerket@lbl.gov), and
Olivier Chevassut (Ochevassut@lbl.gov)

Secure Reliable Multicast Communication Goals



- Provide efficient group communication mechanism for collaborating groups
- Simplify communication between components in a distributed application
- Simplify distributed application development
- Flexible delivery capabilities to support a broad range of application needs
 - ordered and unordered communication
 - reliable and unreliable communication
- Support for collaborations spread across the Internet
- Scalable to collaborations with many members
- Supports access control, confidentiality, authenticity, and integrity among group members

Architecture



InterGroup Goals



- Any member of the group can send messages to the group
- Membership tracked with notification of membership changes
- Deliver messages at each member of the group in a consistent order
 - FIFO order, causal order, or timestamp order
 - membership changes delivered in order
- Scale to the Internet
 - groups with many members
 - heterogeneous latency between members

Why is this hard?



- Reliable and ordered delivery requires knowledge of all the messages that have been sent
- Buffer recovery requires knowing when all of the group has received messages
- Membership (for message delivery) requires consensus among the members
- Dynamic network connectivity and heterogeneous latencies on the Internet
- Partitioning of the network
- Providing guaranteed consistency of delivery order regardless of what happens

The InterGroup Protocol



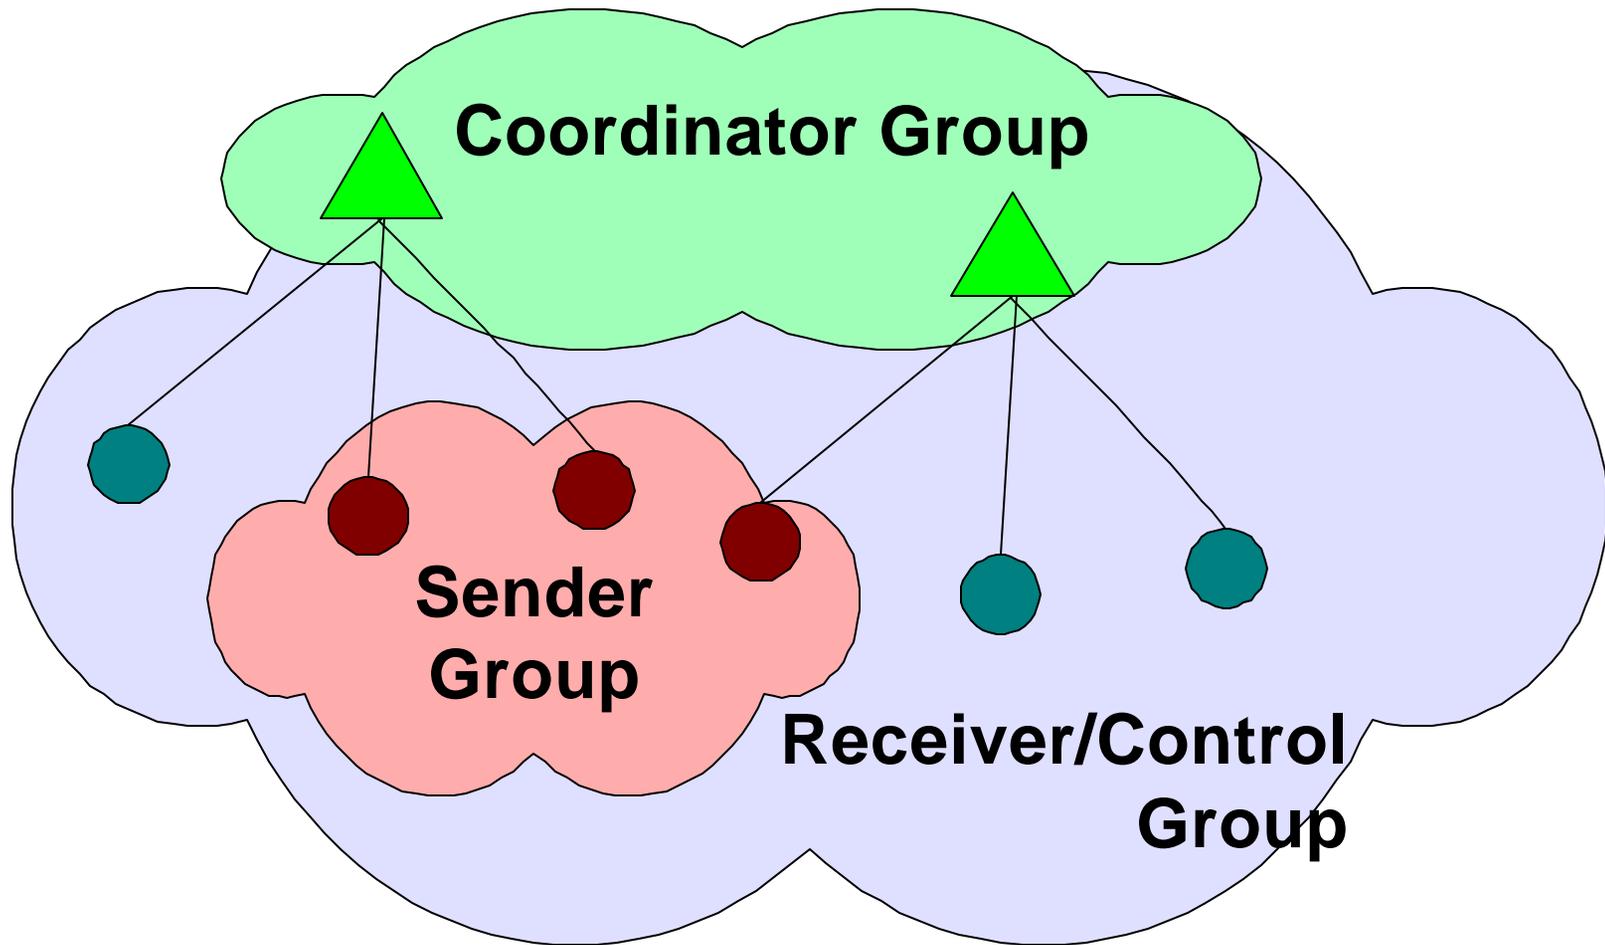
- Based on experience building the Totem protocol
- Split group into a sender and receiver group
- Sender group membership
 - processes are in the sender group only while transmitting messages
 - strictly maintained
 - very dynamic
- Receiver group membership
 - not strictly maintained
 - hierarchically organized to scale to large groups
 - used for retransmissions and garbage collection

InterGroup Implementation



- Sender group membership
 - based on Transis and Totem membership
 - 3 phases: consensus, recovery of messages, and ordering of the membership change message
 - provides virtual synchrony and extended virtual synchrony
 - membership messages ordered with respect to data messages
- Receiver group
 - 2 level hierarchy with coordinators as the branching points and receivers as the leaves
 - use SRM style self-organizing algorithm

InterGroup Schematic

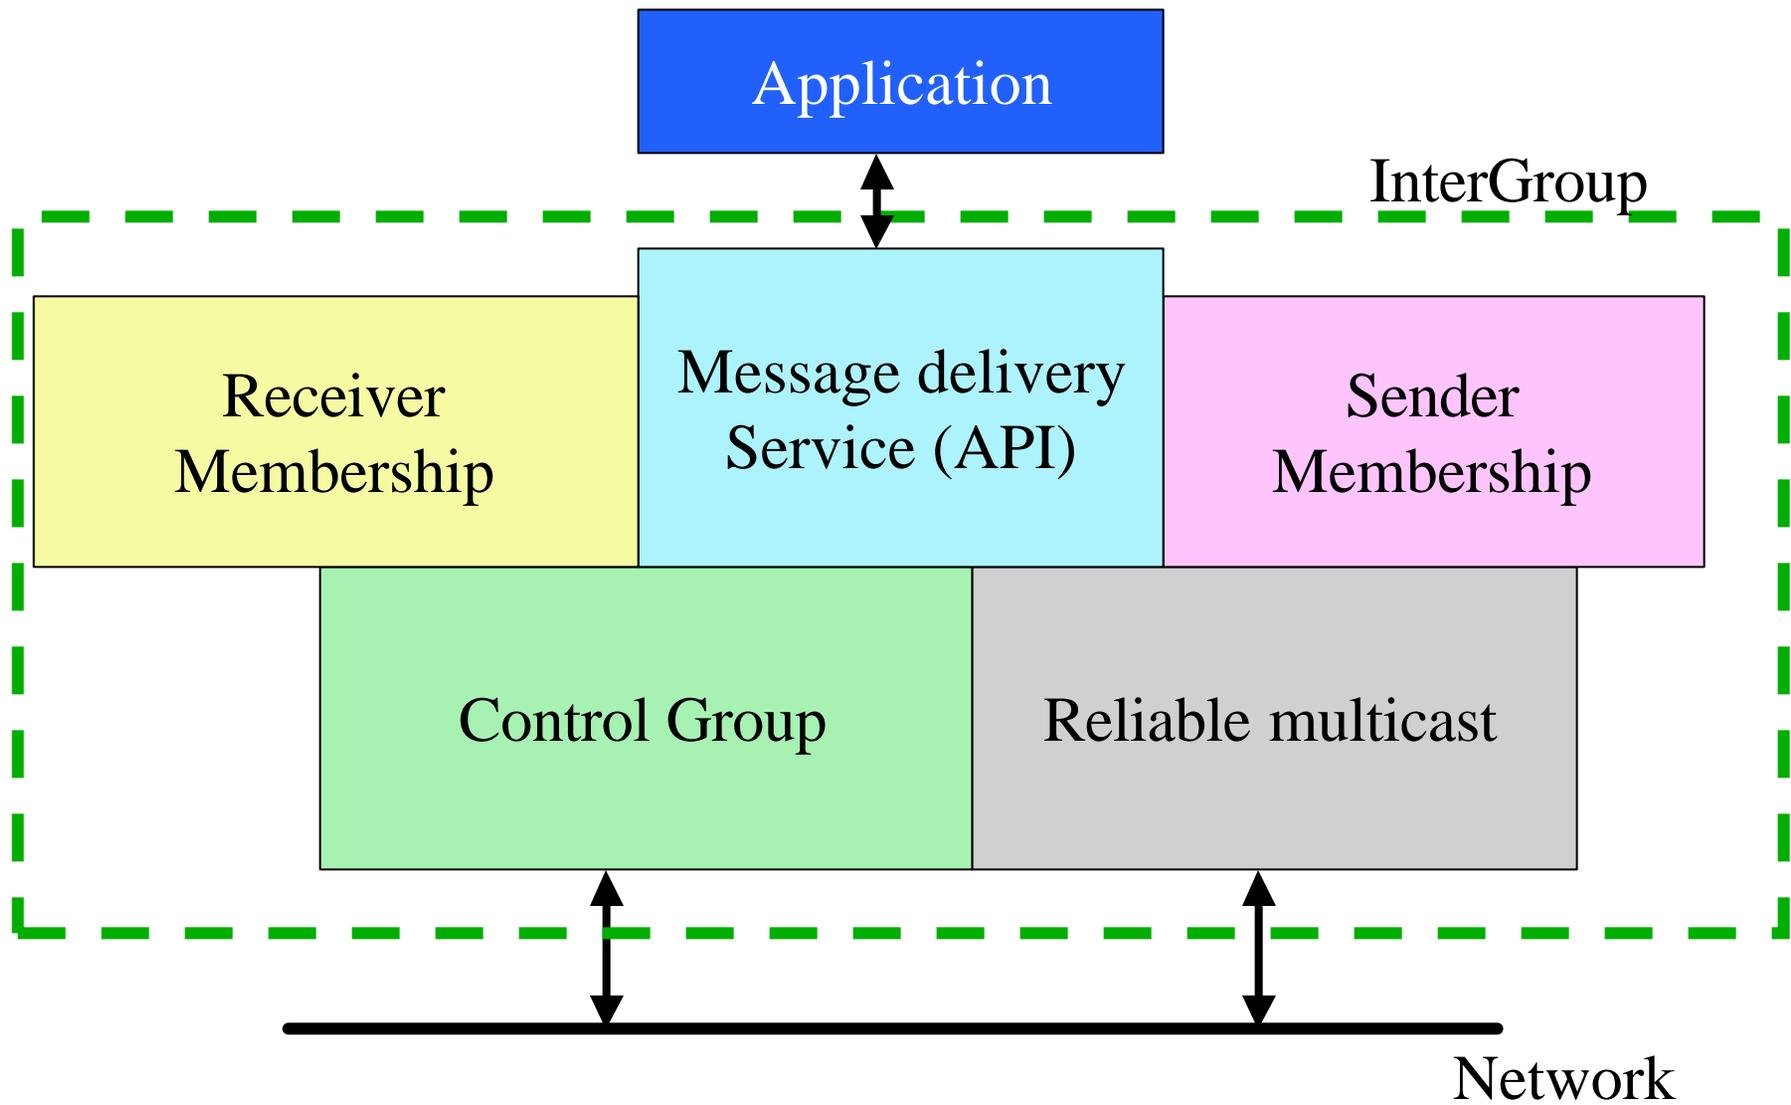


InterGroup Features



- Flexible receiver-based delivery guarantee decision
 - reliability
 - ordering
- Receivers only pay the delay of the level of ordering and reliability they require
- Sender group membership only tracks the processes currently sending messages
- Proxy sending of messages available to receivers
- Strong message delivery guarantees provided (when requested by receiver)
- Receivers use a single coordinator hierarchy for all groups

InterGroup Architecture



Secure Group Communication



- Goals
 - Provide a secure channel for the group with properties similar to Secure Socket Layer (SSL)
 - Authorization of group members (individually enforced)
 - Group key management (not centralized)
 - Group security optional
 - Portable implementation

Secure Group Implementation



- Totem system
 - Group key management uses Cliques toolkit (ISI/UCI)
 - Handles network partitioning
 - Authorization/access control using Akenti (LBNL)
- InterGroup protocols
 - Will be porting to work with the InterGroup sender group
 - Need to develop key distribution mechanisms for the receiver group

Implementation



- Alpha release of the InterGroup implementation expected in April 2001
 - written in Java
 - includes built in performance testing features
 - flow/congestion control very crude
- Secure Group Layer
 - written in C
 - similar security model to SSL
 - currently working on rigorous security analysis
 - release not yet scheduled
- URL - <http://www-itg.lbl.gov/CIF/GroupComm/>