



















MPI FUNCTIONS / DATATYPES				
MPI_ALLGATHERV				
MPI_ATTR_DELETE MPI_BCAST ACH MPI_BUFFER_DETACH IS MPI_CAT_CREATE MPI_CAT_SHIFT TE MPI_COMM_DUP				
MPI_COMM_REMOTE_GRO MPI_COMM_TEST_INTER FREE MPI_ERRHANDLER_GET				
NG MPI_FINALIZE MPI_GET_ELEMENTS ATE MPI_GRAPH_GET				
HBURS_COUNT MPI_GROUP_COMPARE				
E INCL MPI GROUP RANK				
N MPI_IBSEND				
CREATE MPI_INTERCOMM_MERGE				
E MPI OP CREATE				
MPI_PCONTROL				
MPI_REDUCE				
MPI_RSEND_INIT MPT_SEND				
MPI_SSEND				
MPI_TEST				
MPI_TEST_CANCELLED				
COUDDS MPI_TYPE_EXTENT				
T MPI_TYPE_UB				
MPI_WAITALL				
TICTUC				

COMMON MPI FUNCTIONS			
MPI - no re	covery for process crashes, network partitions		
Communica	ation among grouped processes:(groupID, proces	ssID)	
IDs used to	route messages in place of IP addresses		
Operation	Description		
MPI_bsend	Append outgoing message to a local send buffer		
MPI_send	Send message, wait until copied to local/remote buffer		
MPI_ssend	Send message, wat until transmission starts		
MPI_sendrecv	Send message, wait for reply		
MPI_isend	Pass reference to outgoing message and continue		
MPI_issend	Pass reference to outgoing messages, wait until receipt star	rt	
MPI_recv	Receive a message, block if there is none		
MPI_irecv	Check for incoming message, do not block!		
March 4, 2019	TCSS558: Applied Distributed Computing [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma	L13.12	

























- RabbitMQ, Apache QPid
 - Implement Advanced Message Queueing Protocol (AMQP)
- Apache Kafka
 - Dumb broker (message store), similar to a distributed log file
 - Smart consumers intelligence pushed off to the clients
 - Stores stream of records in categories called topics
 - Supports voluminous data, many consumers, with minimal O/H
 - Kafka does not track which messages were read by each consumer
 - Messages are removed after timeout
 - Clients must track their own consumption (Kafka doesn't help)
 - Messages have key, value, timestamp
 - Supports high volume pub/sub messaging and streams

March 4, 2019

TCSS558: Applied Distributed Computing [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma

L13.23











































	COORDINATION - 2	
 Synchronization challenges begin with <u>time</u>: How can we synchronize computers, so they all agree on the time? How do we measure and coordinate when things happen? 		
 Fortunately, for synchronization in distributed systems, it is often sufficient to only agree on a relative ordering of events E.g. not actual time 		
March 4, 2019	TCSS558: Applied Distributed Computing [Winter 2019] L13.45 School of Engineering and Technology, University of Washington - Tacoma L13.45	







































































VECTOR CLOCKS - 4		
 Disclaimer: Without know is not possible relationship or 	ing actual information contained in messages, it e to state with certainty that there is a causal r perhaps a conflict	
 Vector clocks We never know 	can help us suggest possible causality w for sure	
March 4, 2019	TCSS558: Applied Distributed Computing [Winter 2019] School of Engineering and Technology, University of Washington - Tacoma	



























