1 Distributed systems
Examples: client/server, the web, the internet, DNS, Gnutela, BitTorrent, "cloud", NCSA, AWS datacenter
Definitions: "A collection of independent computers that appear thing together", "A distributed system is a collection of entities each of which is autonomous, programmable, asynchronous and munication medium."
Goals: heterogeneity, robustness, availability, transparency (hides internal workings from users), concurrency, efficiency, scalability, security, openness
Chewda
Chewda</

$$x_0 = n, y_0 = 1, \frac{dx}{dt} = -\beta xy$$

$$x = \frac{n(n+1)}{n+e^{\beta}(n+1)t}, y = \frac{(n+1)}{1+ne^{-\beta}(n+1)t}$$

$$y = (n+1) - \frac{1}{n^{cv-2}}$$



Topology-aware: for random selection, core routers face O(N) matructured data, lots of reads N regions O(A) in the probability $1-1/r_{\rm eff}$. (Norm workford) such systems 1) Run on a lot of with O(A) and O(A) workford O(A) and O

The specifier, "A matrix despection, of entropy of the specifier of the specif

Time



FIFO Ordering: Example Update [2,0,0,0] **Causal ordering**: send events that are causally-related must be re-ceived in same causality-ordering order at all receivers; send entire



13 Leader election 1^{13} Leader 1^{13} Leader election 1^{13} Leader 1^{13}



Maekawa: process associated with voting set; intersection of any two voting sets must be non-empty (quorums!), K sets, M sets, process, $K = M = \sqrt{N}$ is best; $2\sqrt{N}$ per enter, \sqrt{N} per exit, intersection of each of the set of

Retransmit request	Filter duplicate requests	Re-execute function or retransmit reply	RPC Semantics
Yes	No	Re-execute	At least once
Yes	Yes	Retransmit	At most once
No	NA	NA	Maybe

Idempotent: applied multiple times without any side effects; can
be used with at-least-once semantics20 Sensor networksMarshalling: convert req/res into common, platform independent
representations: either completes and commits all operations, of
aborts with no effectNode: sensors, microprocessor, comms link, power
Transmission: RF (store and forward, bidirectional), noptical (di-
rectoral altennas, costly broadcast, line of sight needed, passive
routing ("wormhole"), unidirectional links)ACID: stomic (all or nothing), consistency (if starts in consistent
state, transaction ends in consistent state), isolation (transactions), durability
(all effects saved in permanent storage after transaction completed
successfully)Definization: transmit is expensive, so compute instead of trans-
mit; build trees among nodes, calculate summaries in trees, trans-
altent iff some ordering of transactions is serially coving21 Structure of potential
21 Structure of potential

16 Replication control

Chubby: Paxos-like consensus; advisory locks only (no muter guaranteed unless every clicht checks locks before access)
15 RPCs & concurrency control
Local Procedure Call: exactly-once semantics RPC Semantics: at most once (Java RMI), at least once
RPC, maybe/best-effort (COBRA)

do write Invalidate downsides: two processes writing concurrently (flip-flopping, *false sharing*) Update: multiple can have in W state; on write, multicast up-dated part of page; preferred when lots of sharing, writes are to small variables, or large page sizes

23 Spark

