

Simulation & Deployment of Large-scale Distributed Systems



TECHNISCHE
UNIVERSITÄT
DARMSTADT




Christof Leng

Workload Generation

- Configured in the experiment database
- Sophisticated session model
 - Background churn determined by lifetime distribution
 - Workload events to mark nodes as active or inactive
- Support for Unix signals to trigger application-level events

Example Applications

- Supports peer-to-peer and client/server applications
- Language support for Standard ML, C/C++, and Java
- Event-driven, single-threaded applications can be ported easily



BubbleStorm [2]
Search Overlay



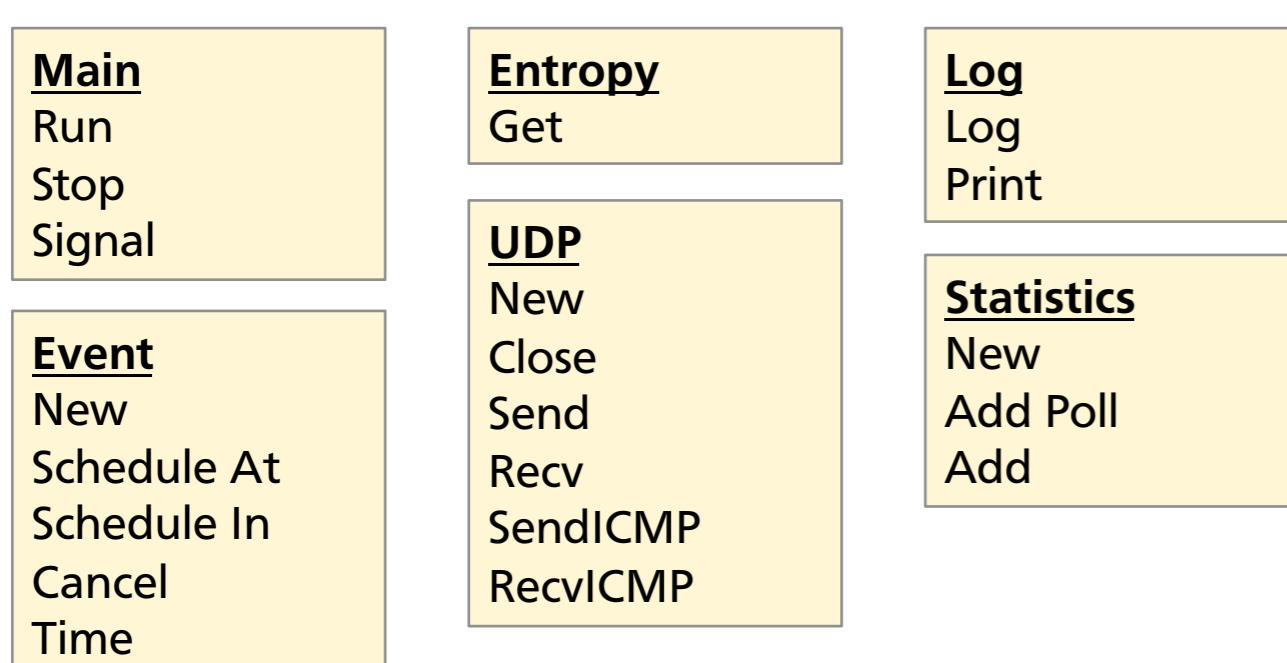
Planet PI4 [3]
Online Shooter

CUSP Transport Protocol

- Complete user-land transport protocol implementation [4]
- Containing flow control, congestion control, reliability, priorities/QoS, encryption, authentication
- Encapsulated over UDP
- Simplifies simulator implementation (simulator only needs UDP support)
- Powerful tool for the implementation of novel network applications

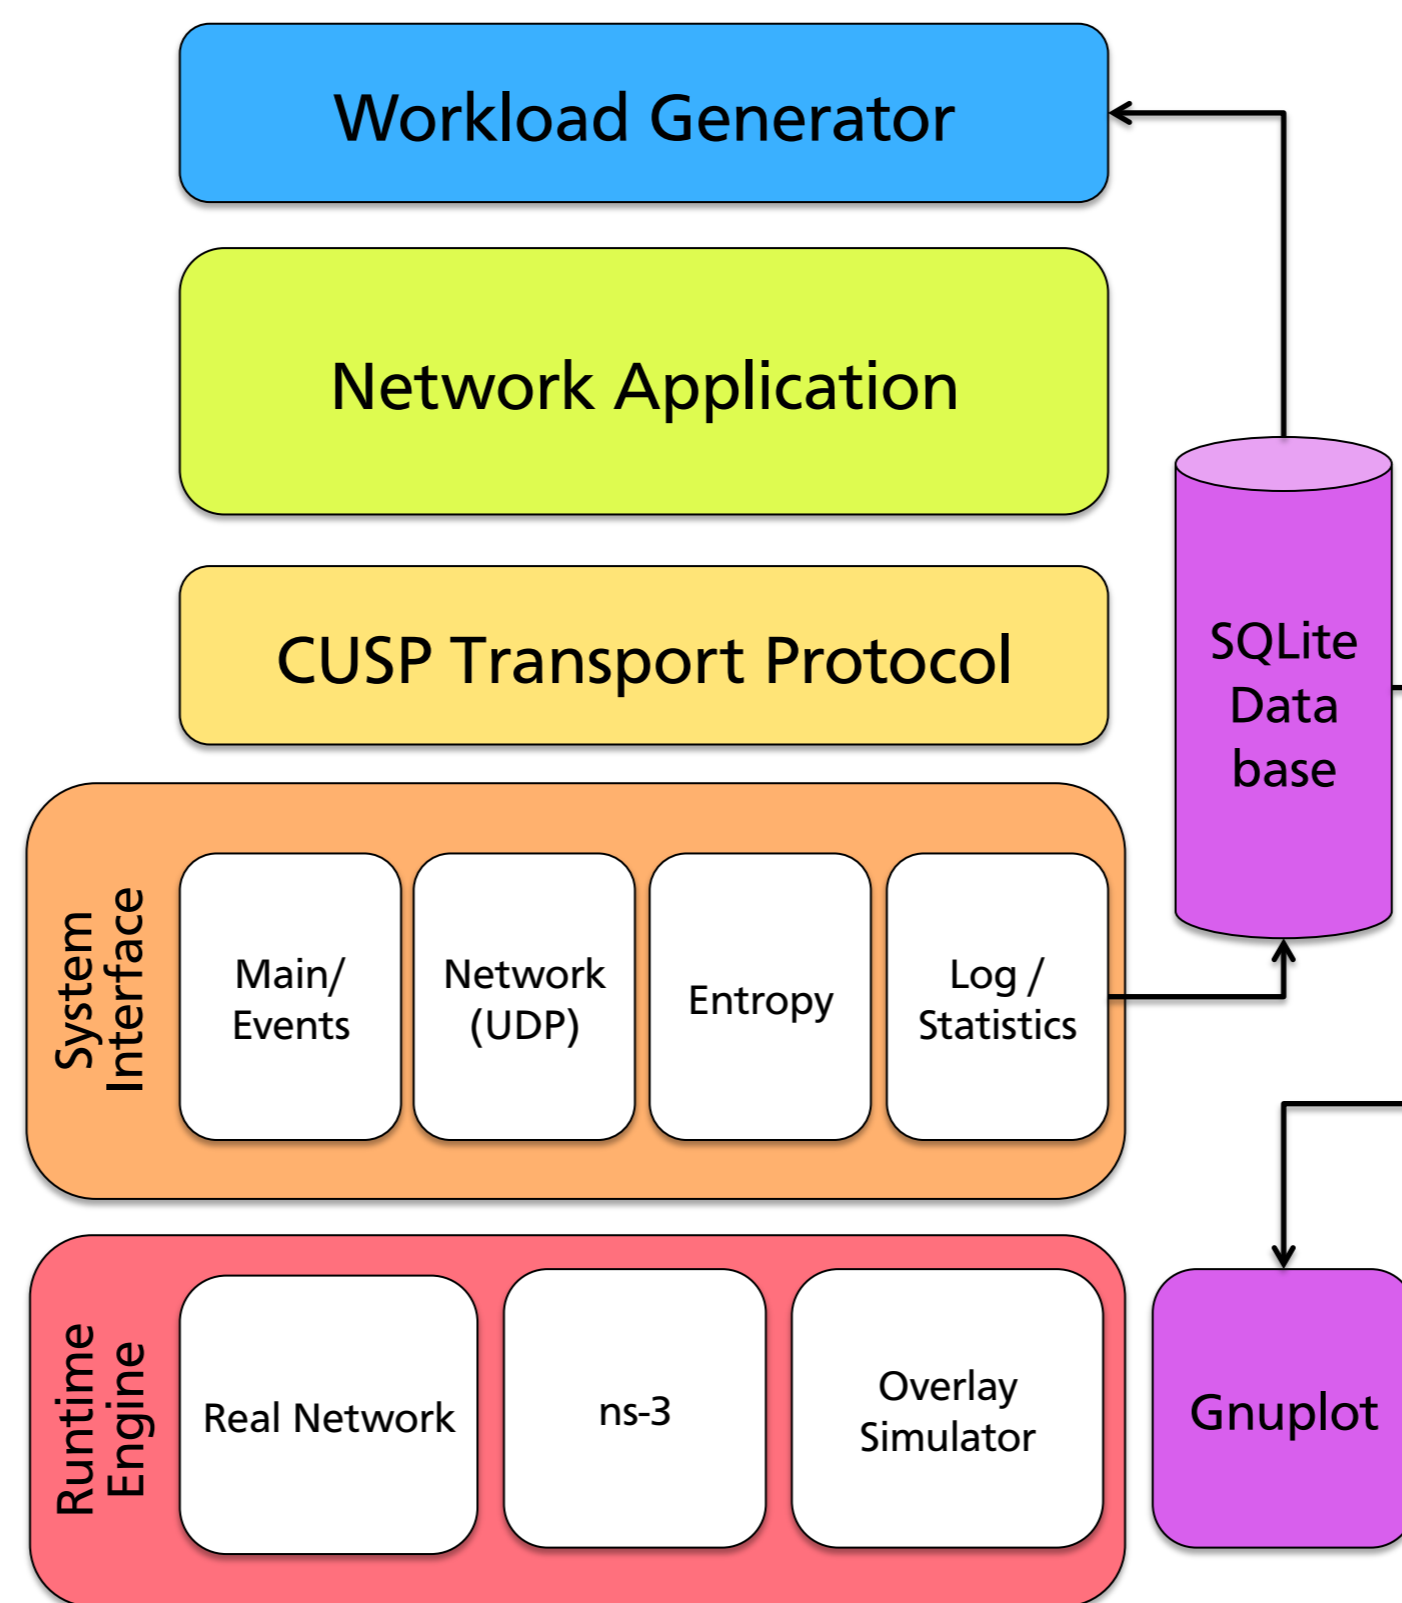
System Interface

- Narrow system interface that abstracts from runtime
 - Easy to learn for application developers
 - Simple to implement new or extend existing runtimes
- Main interfaces:



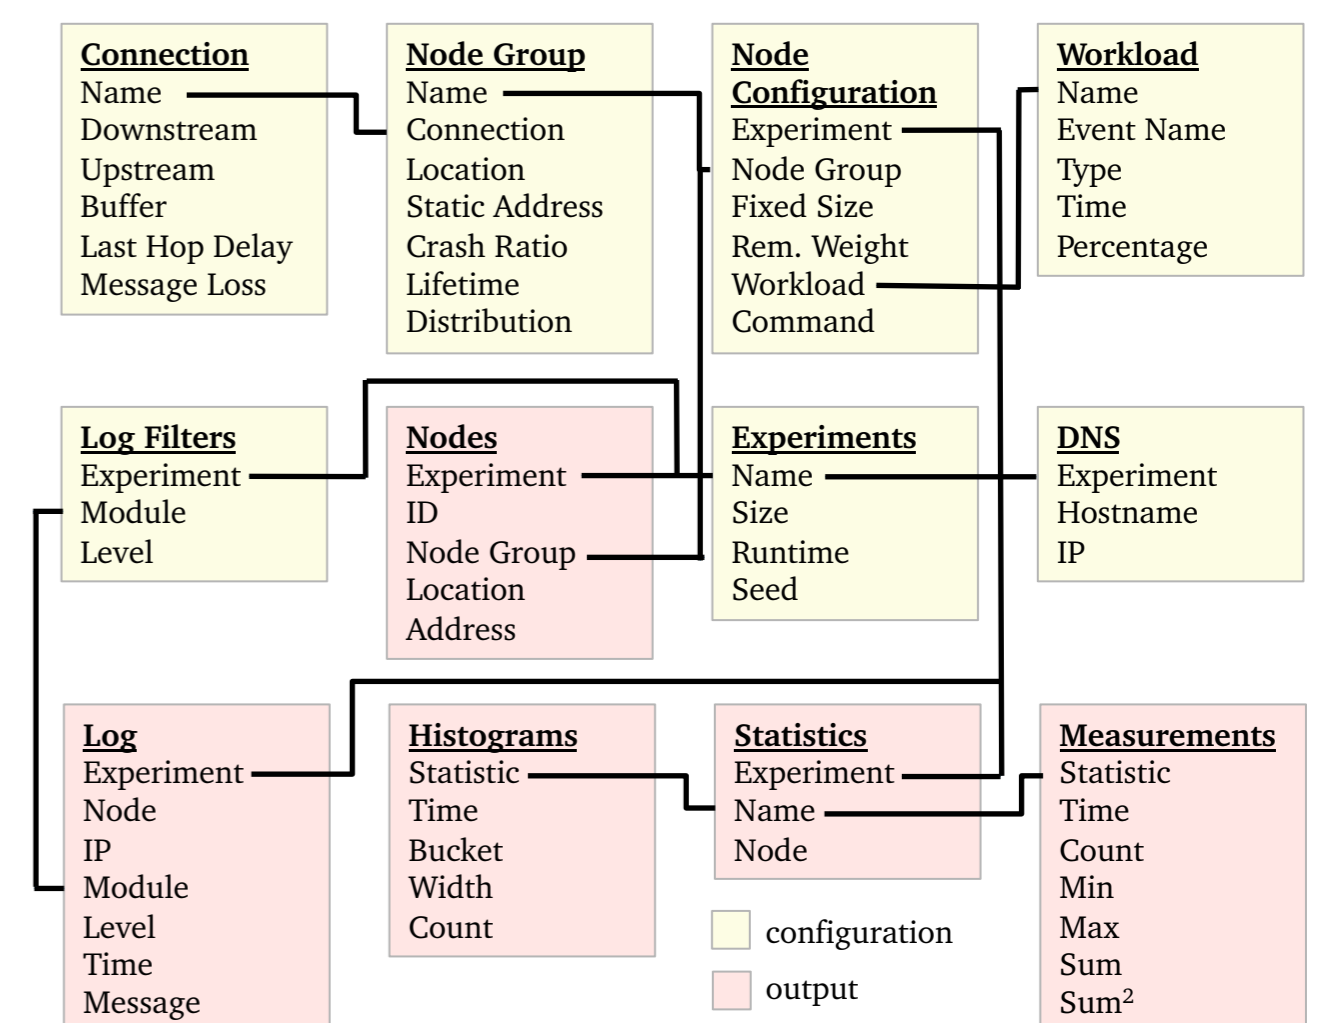
Approach

- Flexible framework** for prototype development and scientific evaluation of distributed applications
- Interchangeable runtime engines** for simulation and real-world deployment
- Narrow system interface** separates application from runtime
- Versatile experiment database** for experiment configuration and output analysis



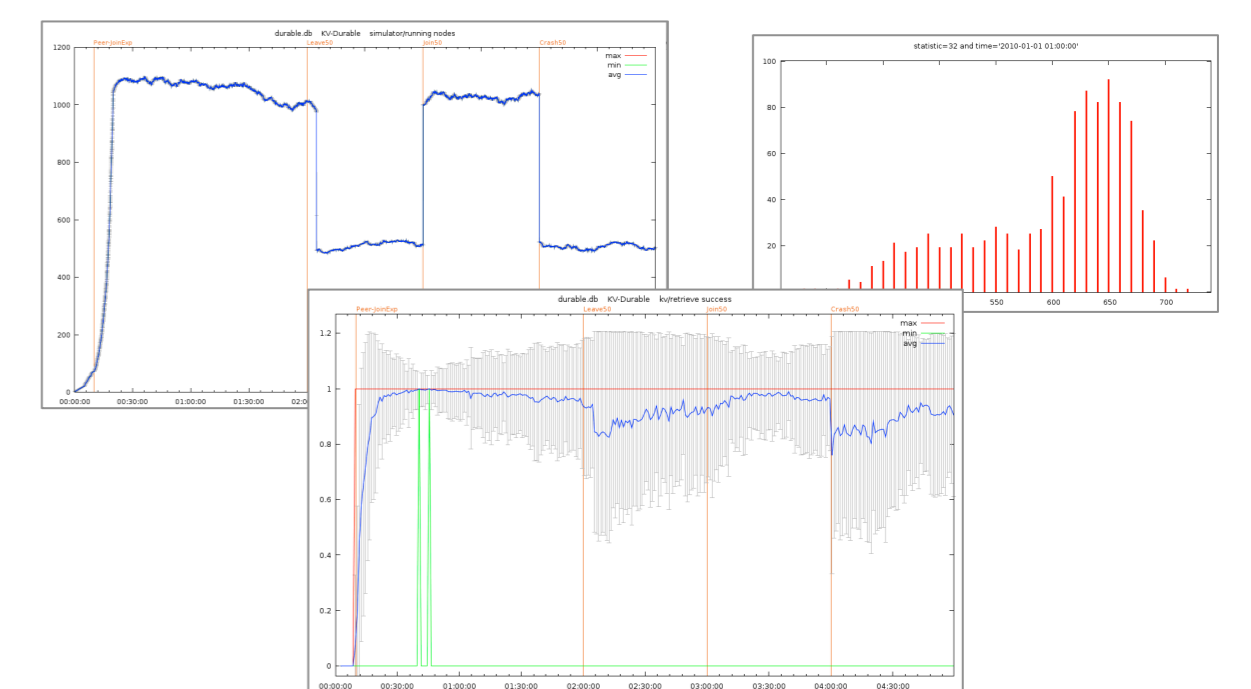
Experiment Database

- Experiment configuration and output in the same database file (using SQLite)
- Flexible experiment configuration
- Statistics and histogram data for plotting
- Rich log output for post-mortem debugging



Integrated Plotting Tools

- Plot directly from experiment database using Gnuplot scripts
- Live plotting during simulation run



Exchangeable Runtime

- Custom overlay simulator**
 - Scalable to 10k+ nodes
 - Coordinate-based delay model [5]
 - Send/receive queues with limited bandwidth
 - SQLite database for log/statistics
- ns-3 simulator [6]**
 - Full network stack simulation
 - Community standard network simulator
- Real network**
 - Real-time main loop
 - Operating system's network stack
 - /dev/urandom for entropy
 - Logging to stdout/stderr
 - Validate simulation results

References:
 [1] Konstantin Pussep, Christof Leng, Sebastian Kaune. Modeling User Behavior in P2P Systems. In Klaus Wehrle, Mesut Günes, James Groß: Modeling and Tools for Network Simulation, Springer, 2010
 [2] Wesley W. Terpstra, Jussi Kangasharju, Christof Leng, Alejandro P. Buchmann. BubbleStorm: Resilient, Probabilistic, and Exhaustive Peer-to-Peer Search. Proceedings of ACM SIGCOMM 2007
 [3] Max Lehn, Christof Leng, Robert Rehner, Tonio Triebel, Alejandro Buchmann. An Online Gaming Testbed for Peer-to-Peer Architectures. Proceedings of ACM SIGCOMM 2011
 [4] Wesley W. Terpstra, Christof Leng, Max Lehn, Alejandro P. Buchmann. Channel-based Unidirectional Stream Protocol (CUSP). Proceedings of the IEEE INFOCOM Mini Conference, 2010
 [5] Sebastian Kaune, Konstantin Pussep, Aleksandra Kovacevic, Christof Leng, Gareth Tyson, Ralf Steinmetz. Modelling the Internet Delay Space Based on Geographic Locations. Proceedings of PDP 2009
 [6] Thomas R. Henderson, Sumit Roy, Sally Floyd, George F. Riley. ns-3 project goals. Proceeding of the 2006 workshop on ns-2: the IP network simulator (WNS2 2006)

