

TORSTEN HOEFLER

DEPARTMENT OF COMPUTER SCIENCE, ETH ZÜRICH

LogGOPSim – Simple and Fast Large-Scale Simulations

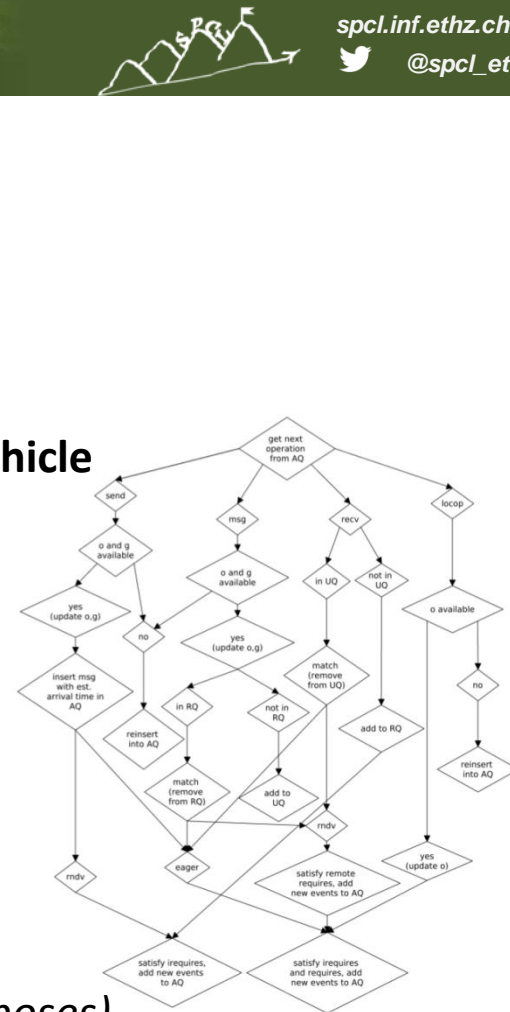


LogGOPSim: Type of analysis

- **Purpose: Performance Model for Extrapolation**
 - Large-scale algorithm development (millions of nodes) [1]
 - Application verification, noise injection, **simple research vehicle**


- **Features:**

- Input languages:
 - GOAL, MPI, Portals IV (beta)*
- Network Models:
 - LogGOPS, LogGOPS+offload, Packet-level (slow)*
- Special features (what's different from simulator XYZ?):
 - Simplicity*** (essentially a single C file, modular, research purposes)
 - Scalability*** (Millions of processes, >100k events per second)
 - Noise injection*** (from real traces, can simulate resilience overheads)
 - Accuracy*** (verified on several systems --- not special, I know)





LogGOPSim: Typical Workload

- **Typical target code:**
 - MPI (or Portals IV) applications
Trace-driven (incl. automatic trace extrapolation features)
- **Typical machine:**
 - LogGOPSim model [2]
 - Portals IV triggered operations
 - Packet-level simple congestion
- **Typical resource of interest:**
 - Network, nothing else 😊
- **Typical scale:**
 - 10's (classroom use)
 - 100's (small design studies)
 - 1000000's (research work)




LogGOPSim: Ecosystem and Integration

- **Our hope:**
 - It's simple enough that we don't need an ecosystem
- **We have:**
 - MPI tracing library (adaptors to others are simple but work-intensive)
 - Schedule generator
 - Single-threaded simulator
 - Very simple visualization (eps output)
- **Unix Philosophy: If you need a tool, add it**
 - Highly modular design
 - New network module: 2 API calls
~ 1 hour for boiler plate
 - New interface: change parser + single C file
~ 1 week of work



LogGOPSim: Self assessment

- **Strengths:**
 - Simple, modular, easy to get student started
 - Very fast
 - Scalability (in-memory trace expansion) [3]
- **Weaknesses:**
 - Doesn't simulate computation at all
 - It's a trace-driven simulation (big traces can be horrible)
- **Features planned:**
 - Full Portals IV support
 - Offload architectures
 - Better topology simulations

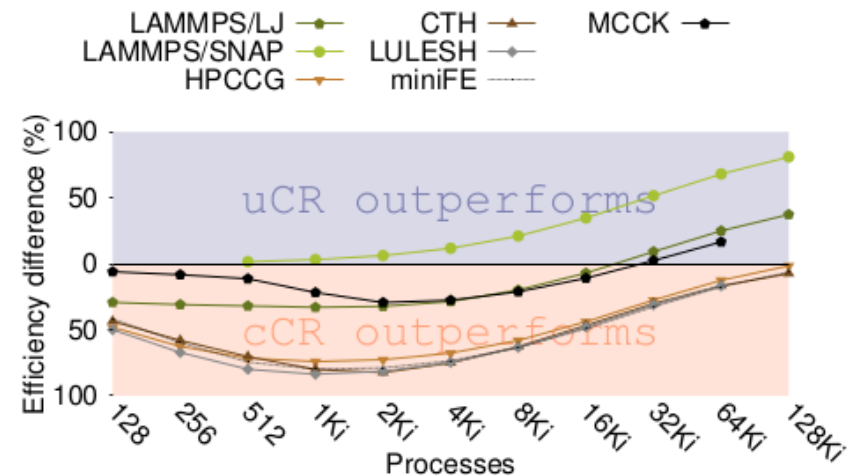


LogGOPSim: Success Stories

- **Collaboration with Sandia on Resilience Overheads at Scale**
 - Last year's PMBS paper on scalable checkpoint/restart simulation [3]
 - This year's SC paper on C/R overheads in practice [4]
 - ... much more to come in the area of FT.

- **Used in several research groups (without us knowing)**
 - E.g., KTH in the EU FP7 project "Exascale ProGRAMming Models" (EPiGRAM)
Very cool parameter sensitivity studies

- **We even used it in class**
 - Students started to extend it immediately



[3]: Levy, Topp, Ferreira, Arnold, TH, Widener: *Using Simulation to Evaluate the Performance of Resilience Strategies at Scale*, PMBS'13

[4]: Ferreira, Widener, Levy, Arnold, TH: *Understanding the Effects of Communication and Coordination on Checkpointing at Scale*, SC14