

First-Order Linear Temporal Logic for Testing Distributed Protocols

José João Ferreira, Nuno Policarpo, José Fragoso Santos,
Alcino Cunha and Alessandro Gianola

INESC-ID / Instituto Superior Técnico,
University of Lisbon, Portugal

OVERLAY 2025 - Bologna, Italy, October 26
Automata, formal languages and logic

The importance of distributed systems and protocols

The importance of distributed systems and protocols

- large-scale social networks

The importance of distributed systems and protocols

- large-scale social networks
- financial infrastructures

The importance of distributed systems and protocols

- large-scale social networks
- financial infrastructures
- cloud computing and storage platforms

The importance of distributed systems and protocols

- large-scale social networks
- financial infrastructures
- cloud computing and storage platforms
- blockchain networks

The importance of distributed systems and protocols

- large-scale social networks
- financial infrastructures
- cloud computing and storage platforms
- blockchain networks

...

The need for correctness: from algorithms to implementations

Examples of bugs found in implementations

- FaB Paxos [1]
- Raft [2]
- Tendermint [3]
- Gasper [4]

- [1] I. Abraham, G. Gueta, D. Malkhi, L. Alvisi, R. Kotla, and J. Martin, “Revisiting fast practical byzantine fault tolerance,” *CoRR*, vol. abs/1712.01367, 2017.
- [2] C. Jensen, H. Howard, and R. Mortier, “Examining Raft’s behaviour during partial network failures,” in *Proceedings of the 1st Workshop on High Availability and Observability of Cloud Systems*, ACM, 2021, pp. 11–17.
- [3] C. Cachin and M. Vukolic, “Blockchain consensus protocols in the wild,” *CoRR*, vol. abs/1707.01873, 2017.
- [4] J. Neu, E. N. Tas, and D. Tse, “Ebb-and-flow protocols: A resolution of the availability-finality dilemma,” in *42nd IEEE Symposium on Security and Privacy*, IEEE, 2021, pp. 446–465.

How can we validate distributed protocols?

Approaches for validating distributed protocols

- **deductive verification**
 - TLAPS for Paxos or Raft [1]
 - Coq for Raft [2]
 - EPR for Paxos [3]

- [1] K. Chaudhuri, D. Doligez, L. Lamport, and S. Merz, “Verifying safety properties with the TLA+ proof system,” in *Proceedings of the 5th International Joint Conference on Automated Reasoning*, Springer, 2010, pp. 142–148.
- [2] J. R. Wilcox et al., “Verdi: A framework for implementing and formally verifying distributed systems,” in *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ACM, 2015, pp. 357–368.
- [3] O. Padon, G. Losa, M. Sagiv, and S. Shoham, “Paxos made EPR: decidable reasoning about distributed protocols,” *Proceedings of the ACM on Programming Languages*, vol. 1, no. OOPSLA, 108:1–108:31, 2017.

Approaches for validating distributed protocols

- **model checking**
 - Alloy for Chord [1]
 - TLC for Pastry [2]

[1] P. Zave, "Using lightweight modeling to understand Chord," *Comput. Commun. Rev.*, vol. 42, no. 2, pp. 49–57, 2012.

[2] N. Azmy, S. Merz, and C. Weidenbach, "A machine-checked correctness proof for Pastry," *Science of Computer Programming*, vol. 158, pp. 64–80, 2018.

Approaches for validating distributed protocols

- **fuzzing**

- Mallory for Raft [1]
- Chronos for Zab and GHOST [2]

[1] R. Meng, G. Pîrlea, A. Roychoudhury, and I. Sergey, "Greybox fuzzing of distributed systems," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, ACM, 2023, pp. 1615–1629.

[2] Y. Chen, F. Ma, Y. Zhou, M. Gu, Q. Liao, and Y. Jiang, "Chronos: Finding timeout bugs in practical distributed systems by deep-priority fuzzing with transient delay," in *IEEE Symposium on Security and Privacy*, 2024, pp. 1939–1955.

Approaches for validating distributed protocols

- **monitoring**
 - TLA+ for 2PC [1]
 - TLA+ for Raft [2]
 - ATL for Chord [3]

- [1] H. Cirstea, M. A. Kuppe, B. Loillier, and S. Merz, "Validating traces of distributed programs against TLA+ specifications," in *Proceedings of the 22nd Int. Conf. on Software Eng. and Formal Methods*, Springer, 2024, pp. 126–143.
- [2] H. Howard, M. A. Kuppe, E. Ashton, A. Chamayou, and N. Crooks, "Smart casual verification of the confidential consortium framework," in *22nd USENIX Symposium on Networked Systems Design and Impl.*, 2025, pp. 259–276.
- [3] N. Policarpo, J. F. Santos, A. Cunha, J. Leitão, and P. Á. Costa, "Specifying distributed hash tables with allen temporal logic," in *13th IEEE/ACM International Conference on Formal Methods in Software Engineering*, IEEE, 2025, pp. 63–73.

Online vs offline monitoring

First-Order Linear Temporal Logic

Why is FOLTL suited for...

Why is FOLTL suited for...

- ...distributed protocols?

Why is FOLTL suited for...

- ...distributed protocols?
- ...online monitoring?

Why is FOLTL suited for...

- ...distributed protocols?
- ...online monitoring?
- ...offline monitoring?

Why is FOLTL suited for distributed protocols?

- properties are concerned with time and data
- LTL is not expressive enough
- with past operators, specifications become more natural

Why is FOLTL suited for online monitoring?

- (semi-)decidable fragments of FOLTL and tools
 - TDL using `jUnitRV` [1]
 - ALTL_f using `ADATool` [2]
 - LTL_f^{MT} using `BLACK` [3]
- SMT solvers, tableaux procedures, and automata-based methods

- [1] N. Decker, M. Leucker, and D. Thoma, "Monitoring modulo theories," *Int. J. Softw. Tools Technol. Transf.*, vol. 18, no. 2, pp. 205–225, 2016.
- [2] P. Felli, M. Montali, F. Patrizi, and S. Winkler, "Monitoring arithmetic temporal properties on finite traces," in *Proceedings of the 37th AAAI Conference on Artificial Intelligence*, AAAI Press, 2023, pp. 6346–6354.
- [3] M. Faella and G. Parlato, "A unified automata-theoretic approach to ltl_f modulo theories," in *Proceedings of the 27th European Conference on Artificial Intelligence*, IOS Press, 2024, pp. 1254–1261.

Why is FOLTL suited for offline monitoring?

- applicable, though largely unexplored
- no need for SAT tools
- temporal reasoning, quantification, and predicates


Syntax

$$t := c \mid x \mid f(t_1, \dots, t_n)$$

$$\begin{aligned} \phi := & p(t_1, \dots, t_n) \mid t_1 = t_2 \mid \neg\phi \mid \phi \vee \phi \mid \exists x. \phi \\ & \mid X\phi \mid Y\phi \mid \phi U \phi \mid \phi S \phi \end{aligned}$$

First-Order Linear Temporal Logic

Semantics



Two-Phase Commit (2PC)

2PC

- **prepare**(n, t):
the coordinator requests node n to prepare for transaction t
- **v-commit**(n, t) / **v-abort**(n, t):
node n votes to commit/abort transaction t
- **i-commit**(t) / **i-abort**(t):
the coordinator instructs all nodes to commit/abort transaction t

2PC

LIVENESS **Eventual decision**: if the coordinator requests any node n to prepare for any transaction t , all nodes will eventually be instructed to commit or abort that transaction

$$\forall n, t. G\left(\text{prepare}(n, t) \Rightarrow F(\mathbf{i}\text{-commit}(t) \vee \mathbf{i}\text{-abort}(t))\right)$$

2PC

SAFETY Atomic commitment: if the coordinator instructs all nodes to commit any transaction t , then all nodes n must have voted to commit that transaction; if it instructs to abort t , then at least one node n must have voted to abort it

$$\forall t. G \left(\left(\mathbf{i}\text{-commit}(t) \Rightarrow \forall n. O \mathbf{v}\text{-commit}(n, t) \right) \wedge \left(\mathbf{i}\text{-abort}(t) \Rightarrow \exists n. O \mathbf{v}\text{-abort}(n, t) \right) \right)$$

Distributed Hash Table (DHT)

DHT

- **b-store_z**(n, k, v):
the client requests node n to store the key-value pair (k, v)
- **e-store_z**(n'):
the operation completes at node n'
- **b-lookup_z**(n, k):
the client requests the value for key k from node n
- **e-lookup_z**(n', v):
the operation completes at the node n' storing (k, v) and returns v

DHT

LIVENESS **Operation termination:** every store or lookup operation eventually completes, that is, for any operation z , once it starts, it will eventually end

$$\forall z. G \left(\left(\mathbf{b-store}_z(\dots) \Rightarrow F \mathbf{e-store}_z(-) \right) \wedge \left(\mathbf{b-lookup}_z(\dots) \Rightarrow F \mathbf{e-lookup}_z(\dots) \right) \right)$$

DHT

SAFETY **Lookup consistency:** if a lookup z obtains a value v for a given key k , then the (k, v) pair was previously written by a store operation

$$\forall k, v. G \left(\forall z. (\mathbf{b}\text{-lookup}_z(-, k) \wedge F \mathbf{e}\text{-lookup}_z(-, v)) \right. \\ \left. \Rightarrow O \mathbf{b}\text{-store}(-, k, v) \right)$$

FOLTL modeling style for DHT

FOLTL modeling style for DHT

- operations modeled with paired begin-end events rather than single predicates spanning multiple time points

FOLTL modeling style for DHT

- operations modeled with paired begin-end events rather than single predicates spanning multiple time points
- contrasts with 2PC example, showing FOLTL's flexibility

FOLTL modeling style for DHT

- operations modeled with paired begin-end events rather than single predicates spanning multiple time points
- contrasts with 2PC example, showing FOLTL's flexibility
- fragments with syntactic variations may suit protocols with continuous operations and clear input-output distinction

FOLTL for online monitoring

FOLTL for online monitoring

- general FOLTL is undecidable for satisfiability checking, but some fragments are (semi-)decidable

FOLTL for online monitoring

- general FOLTL is undecidable for satisfiability checking, but some fragments are (semi-)decidable
- satisfiability tools (e.g., Alloy, BLACK) could be adapted for runtime trace validation

FOLTL for online monitoring

- general FOLTL is undecidable for satisfiability checking, but some fragments are (semi-)decidable
- satisfiability tools (e.g., Alloy, BLACK) could be adapted for runtime trace validation
- **challenge:** adapt solvers for online monitoring

FOLTL for offline monitoring

FOLTL for offline monitoring

- may not handle checking liveness properties

FOLTL for offline monitoring

- may not handle checking liveness properties
- domains are trace-bounded, so it avoids satisfiability checking

FOLTL for offline monitoring

- may not handle checking liveness properties
- domains are trace-bounded, so it avoids satisfiability checking
- monitoring complete traces is decidable and efficient

FOLTL for offline monitoring

- may not handle checking liveness properties
- domains are trace-bounded, so it avoids satisfiability checking
- monitoring complete traces is decidable and efficient
- **challenge:** build scalable monitors for quantifier-heavy specifications on large traces

FOLTL for testing distributed protocols

challenges:

FOLTL for testing distributed protocols

challenges:

- identify useful FOLTL fragments for online/offline monitoring

FOLTL for testing distributed protocols

challenges:

- identify useful FOLTL fragments for online/offline monitoring
- instrument distributed systems to collect traces

FOLTL for testing distributed protocols

challenges:

- identify useful FOLTL fragments for online/offline monitoring
- instrument distributed systems to collect traces
- design a monitor to check trace validation against FOLTL formulas

Thank you

First-Order Linear Temporal Logic for Testing Distributed Protocols

José João Ferreira, Nuno Policarpo, José Frago Santos,
Alcino Cunha and Alessandro Gianola

INESC-ID / Instituto Superior Técnico,
University of Lisbon, Portugal

OVERLAY 2025 - Bologna, Italy, October 26
Automata, formal languages and logic