

Asynchronously Communicating Visibly Pushdown Systems

Domagoj Babić

Google Research

joint work with Zvonimir Rakamarić
School of Computing, University of Utah

OVERVIEW

- ▶ Motivation
- ▶ Communicating Finite State Machines
- ▶ Relations
- ▶ Two Extensions: Synchronization and (Visible) Stack
- ▶ Decidability of Reachability
- ▶ Limitations and Future Work

MOTIVATION

- ▶ Debugging of large-scale distributed systems:
 - ▶ Counters and statistics
 - ▶ Unit tests
 - ▶ Progressive deployment
 - ▶ Difficult localization of root causes of failures

MOTIVATION

- ▶ Debugging of large-scale distributed systems:
 - ▶ Counters and statistics
 - ▶ Unit tests
 - ▶ Progressive deployment
 - ▶ Difficult localization of root causes of failures
- ▶ Penetration of formal methods

MOTIVATION

- ▶ Debugging of large-scale distributed systems:
 - ▶ Counters and statistics
 - ▶ Unit tests
 - ▶ Progressive deployment
 - ▶ Difficult localization of root causes of failures
- ▶ Penetration of formal methods
 - ▶ Zero?

MOTIVATION

- ▶ Debugging of large-scale distributed systems:
 - ▶ Counters and statistics
 - ▶ Unit tests
 - ▶ Progressive deployment
 - ▶ Difficult localization of root causes of failures
- ▶ Penetration of formal methods
 - ▶ Zero?
 - ▶ But there's some hope. . .

MOTIVATION

- ▶ Debugging of large-scale distributed systems:
 - ▶ Counters and statistics
 - ▶ Unit tests
 - ▶ Progressive deployment
 - ▶ Difficult localization of root causes of failures
- ▶ Penetration of formal methods
 - ▶ Zero?
 - ▶ But there's some hope. . .
- ▶ Protocol Buffers (Google), Thrift (Facebook), Singularity (MSR)
 - ▶ Specification of message formats
 - ▶ Automated (de)serialization
 - ▶ Could we expand these to describe protocols?


MOTIVATION

- ▶ Debugging of large-scale distributed systems:
 - ▶ Counters and statistics
 - ▶ Unit tests
 - ▶ Progressive deployment
 - ▶ Difficult localization of root causes of failures
- ▶ Penetration of formal methods
 - ▶ Zero?
 - ▶ But there's some hope. . .
- ▶ Protocol Buffers (Google), Thrift (Facebook), Singularity (MSR)
 - ▶ Specification of message formats
 - ▶ Automated (de)serialization
 - ▶ Could we expand these to describe protocols?
- ▶ If so, which formal model would we use?

DESIRED MODEL PROPERTIES

- ▶ Asynchronous message passing
- ▶ Distributed concurrently executing processes (with stacks)
- ▶ Lossless (unbounded) FIFO message buffers
 - ▶ Network protocols mostly guarantees no loss of messages
 - ▶ Messages consumed mostly in-order (not in Erlang!)
 - ▶ Buffers practically unbounded (bounds often large, vary a lot)
 - ▶ Some loss possible when buffers overflow (usually indicates under-design or error)
 - ▶ Processes use multiple buffers

COMMUNICATING FINITE STATE MACHINES



- ▶ Bounded number of
 - ▶ Processes (FSMs)
 - ▶ One-directional channels (unbounded FIFOs)
- ▶ Reachability undecidable
- ▶ Decidable if [Pac87]:
 - ▶ In every composite state, CFSM has a recognizable channel property
 - ▶ I.e., all queue languages are regular and form a recognizable relation

LIMITATIONS OF RECOGNIZABLE RELATIONS

- ▶ Processes can't send multiple messages to different channels in sequence
- ▶ No (unbounded) message counting
- ▶ Can't express remote procedure calls (and returns)

RELATIONS

$$\begin{array}{ccccccc} \checkmark & & \checkmark & & \emptyset & & \emptyset \\ \textit{Rec} & \subset & \textit{Sync} & \subset & \textit{DRat} & \subset & \textit{Rat} \\ [\text{Tho90}] & & [\text{EES69}] & & [\text{RS59}] & & [\text{RS59}] \end{array}$$

- ▶ Recognizable relations
 - ▶ n -tape automaton, each tape operates independently, has its own memory
 - ▶ Finite unions of cross-products of regular languages
 - ▶ Even when concatenated, still regular
- ▶ Synchronized relations
 - ▶ n -tape automaton, all heads move in lock-step
 - ▶ Can express languages like (a^n, b^n)
- ▶ Rational relations — finite n -tape automata
- ▶ (\checkmark — denotes that containment is decidable)

EXTENDING CFSMs

1. From FSMs to Visibly Pushdown Automata
2. From *Rec* to *Sync*

COMMUNICATING VISIBLY PUSHDOWN TRANSDUCCERS

Definition (CVPT)

A CVPT is a tuple $T = (\Sigma_{rcv}, \Sigma_{snd}, Q, S, I, F, \Gamma, \Delta)$ of finite sets, where Σ_{rcv} is an input alphabet, Σ_{snd} an output alphabet, Q a set of unbounded FIFO queues, S a set of states, $I \subseteq S$ a set of initial states, $F \subseteq S$ a set of final states, Γ an alphabet of stack symbols, and Δ a transition relation.

COMPOSITE CONFIGURATION

Configuration

$$C = (s, \sigma, \vec{\varrho}) = (s, \sigma, \varrho_1, \dots, \varrho_{|Q|}) \in \mathcal{S} \times \Gamma^* \times \prod q \in Q(\Sigma_{snd_q} \cup \Sigma_{rcv_q})^*$$

Composite Configuration

$$\vec{C} = (C_1, \dots, C_n)$$

TRANSITION RELATION FOR CVPTs

- ▶ Σ_{rcv} partitioning: $\Sigma_{rcv} = \Sigma_c \cup \Sigma_r \cup \Sigma_i$ (\cup is disjoint union)

- ▶ $\Delta = \delta_c \cup \delta_r \cup \delta_i$

Call $\delta_c \subseteq \mathcal{S} \times \Sigma_c \times (\Sigma_{snd} \cup \{\epsilon\}) \times \Gamma \times \mathcal{S}$

Return $\delta_r \subseteq \mathcal{S} \times \Sigma_r \times \Gamma \times (\Sigma_{snd} \cup \{\epsilon\}) \times \mathcal{S}$

Internal $\delta_i \subseteq \mathcal{S} \times (\Sigma_i \cup \{\epsilon\}) \times (\Sigma_{snd} \cup \{\epsilon\}) \times \mathcal{S}$

- ▶ If $(s_1, q_1? \overline{m}_1, q_2!m_2, \gamma, s_2) \in \delta_c$, then

$$C \xrightarrow[c]{q_1? \overline{m}_1 / q_2! m_2, \gamma} C[s_1 \leftarrow s_2, \sigma \leftarrow \gamma \cdot \sigma, \overline{m}_1 \cdot \varrho_1 \leftarrow \varrho_1, \varrho_2 \leftarrow \varrho_2 \cdot m_2]$$

- ▶ If $(s_1, q_1? \underline{m}_1, \gamma, q_2!m_2, s_2) \in \delta_r$, then


$$C \xrightarrow[r]{q_1? \underline{m}_1, \gamma / q_2! m_2} C[s_1 \leftarrow s_2, \gamma \cdot \sigma \leftarrow \sigma, \underline{m}_1 \cdot \varrho_1 \leftarrow \varrho_1, \varrho_2 \leftarrow \varrho_2 \cdot m_2]$$

- ▶ If $(s_1, q_1? m_1, q_2!m_2, s_2) \in \delta_i$, then

$$C \xrightarrow[i]{q_1? m_1 / q_2! m_2} C[s_1 \leftarrow s_2, m_1 \cdot \varrho_1 \leftarrow \varrho_1, \varrho_2 \leftarrow \varrho_2 \cdot m_2]$$


EXAMPLE OF A SYSTEM OF CVPTs

SERVER:




DATABASE:

$q_5?dcnt/q_4!dsnd$



CLIENT:





CVPT COMPOSITION

- ▶ Even for VPL inputs, CVPTs can generate context-free outputs [RS08]

Definition (Composition Property)

A CVPT T_j is said to be *composable* if a projection of its output language (i.e., a transduction of some VPL L) onto the input alphabet of any T_i is a VPL.

EXAMPLES OF CVPTs NOT SATISFYING THE COMPOSITION PROPERTY



- ▶ The left one generates a context-free language $a^n b^n$
- ▶ The right one reverses calls and returns, generating $\underline{b}^n \bar{a}^n$

SYNCHRONIZED TREE RELATIONS (TREES)

- ▶ Mapping η for a Word:

$$\underline{c}_1 \cdot \bar{c}_2 \cdot a_2 \cdot \bar{c}_3 \cdot a_3 \cdot \underline{c}_3 \cdot \underline{c}_2 \cdot a_5 \cdot \bar{c}_4 \cdot a_6$$


- ▶ First, we translate VPL words to trees

$$\eta(\epsilon) = \#;$$


$$\eta(\bar{c}w) = \bar{c}(\eta(w), \#), \text{ if no } \underline{c} \text{ matching } \bar{c}$$

$$\eta(\bar{c}w\underline{c}w') = \bar{c}(\eta(w), \eta(\underline{c}w'))$$


$$\eta(a\underline{w}) = a(\eta(w)), \text{ if } a \in \Sigma_i \cup \Sigma_r$$



SYNCHRONIZED TREE RELATIONS (OVERLAP)



SYNCHRONIZED TREE RELATIONS (OVERLAP)



Definition (Synchronized Configuration Property)

An asynchronous system of CVPTs has the *synchronized configuration property* iff in every reachable composite control state \vec{s} , the encoding $[\eta(\mathbf{L}_{qs}(\vec{s}))]$ is a synchronized tree relation, i.e.,

$\{[\eta(\sigma_1), \dots, \eta(\sigma_n), \eta(\varrho_1), \dots, \eta(\varrho_k)] \mid (\sigma_1, \dots, \sigma_n, \varrho_1, \dots, \varrho_k) \in \mathbf{L}_{qs}(\vec{s})\} \in \text{Sync}^{\mathcal{V}}$
 (where $\mathbf{L}_{qs}(\vec{s}) = \{ \bar{c} \cdot \varrho, \bar{c} \cdot \sigma \mid \bar{c}_0 \xrightarrow{*} \bar{c} \wedge \bar{c} \cdot s = \vec{s} \}$)

DECIDABILITY OF REACHABILITY

Theorem

Reachability is decidable for a system of composable CVPTs with the synchronized configuration property.

CONSISTENCY

Definition (Consistency)

Let $M = (T_1, \dots, T_n)$ be a system of CVPTs. We say that relation $\mathbf{L} \subseteq \prod_{1 \leq i \leq n} S_i \times \prod_{q \in Q} \Sigma_{rcvq}^* \times \prod_{1 \leq i \leq n} \Gamma_i^*$ is *consistent* (with respect to M) if $\vec{C}_1 \xrightarrow{*} \vec{C}_2$ and $(\vec{C}_1.\varrho, \vec{C}_1.\sigma) \in \mathbf{L}(\vec{C}_1.s)$ imply $(\vec{C}_2.\varrho, \vec{C}_2.\sigma) \in \mathbf{L}(\vec{C}_2.s)$.

PROOF SKETCH

- ▶ Existential proof
- ▶ Two semi-algorithms
- ▶ If \vec{C} reachable — can be reached by BFS
- ▶ If \vec{C} unreachable — must exist a consistent synchronized relation \mathbf{L} , such that $(\vec{C}.q, \vec{C}.\sigma) \notin \mathbf{L}(\vec{C}.s)$.
- ▶ Since synchronized tree languages can be identified by enumeration [Gol67], we can enumerate consistent relations

ALGORITHM BASED ON THE PROOF?

- ▶ Would be hopelessly inefficient, but. . .

ALGORITHM BASED ON THE PROOF?

- ▶ Would be hopelessly inefficient, but. . .
- ▶ Seems like regular (tree) model checking is the way to go

ALGORITHM BASED ON THE PROOF?

- ▶ Would be hopelessly inefficient, but. . .
- ▶ Seems like regular (tree) model checking is the way to go
- ▶ If programmers provided the invariant for at least one edge in every loop in the composite state graph — model checking can be done efficiently

RESYNCHRONIZABLE RELATIONS

- ▶ Accepted by n -tape automata with an *a-priori* bounded delay among heads
- ▶ Heads may not be perfectly synchronized
- ▶ Can be reduced to a finite union of component-wise products of synchronized relations and finite sets
- ▶ Example: $R = (b^m aab^k, c^m d^k)$
 - $k=0$ $(b^m, c^m) \cdot (a^2, \epsilon^2)$
 - $k=1$ $(b^m, c^m) \cdot (a^2 b, d \cdot \epsilon^2)$
 - $k=2$ $(b^m, c^m) \cdot (a^2 bb, d^2 \cdot \epsilon^2)$
 - $k \geq 2$ $(b^m aab^{k-2}, c^m d^k) \cdot (bb, \epsilon^2)$

SWITCHING AUTOMATA

- ▶ Generalization of resynchronizable languages
- ▶ Tapes can be switched on and off
- ▶ If the number of switches is bounded, accepted relations can be translated into a synchronized relation

LIMITATIONS AND FUTURE WORK

- ▶ Synchronized relations still somewhat limited
 - ▶ Not necessarily perfectly synchronized
 - ▶ Some relaxations possible (e.g., resynchronizable relations)

LIMITATIONS AND FUTURE WORK

- ▶ Synchronized relations still somewhat limited
 - ▶ Not necessarily perfectly synchronized
 - ▶ Some relaxations possible (e.g., resynchronizable relations)
 - ▶ Are there more expressive natural families of relations in between *Sync* and *DRat*?

LIMITATIONS AND FUTURE WORK

- ▶ Synchronized relations still somewhat limited
 - ▶ Not necessarily perfectly synchronized
 - ▶ Some relaxations possible (e.g., resynchronizable relations)
 - ▶ Are there more expressive natural families of relations in between *Sync* and *DRat*?
 - ▶ More general theory of resynchronizable relations needed

LIMITATIONS AND FUTURE WORK

- ▶ Synchronized relations still somewhat limited
 - ▶ Not necessarily perfectly synchronized
 - ▶ Some relaxations possible (e.g., resynchronizable relations)
 - ▶ Are there more expressive natural families of relations in between *Sync* and *DRat*?
 - ▶ More general theory of resynchronizable relations needed
- ▶ Checking that relation is synchronized is undecidable

LIMITATIONS AND FUTURE WORK

- ▶ Synchronized relations still somewhat limited
 - ▶ Not necessarily perfectly synchronized
 - ▶ Some relaxations possible (e.g., resynchronizable relations)
 - ▶ Are there more expressive natural families of relations in between *Sync* and *DRat*?
 - ▶ More general theory of resynchronizable relations needed
- ▶ Checking that relation is synchronized is undecidable
 - ▶ Is it possible to guarantee (say, by a type system) that all relations are synchronized by construction?

LIMITATIONS AND FUTURE WORK

- ▶ Synchronized relations still somewhat limited
 - ▶ Not necessarily perfectly synchronized
 - ▶ Some relaxations possible (e.g., resynchronizable relations)
 - ▶ Are there more expressive natural families of relations in between *Sync* and *DRat*?
 - ▶ More general theory of resynchronizable relations needed
- ▶ Checking that relation is synchronized is undecidable
 - ▶ Is it possible to guarantee (say, by a type system) that all relations are synchronized by construction?
- ▶ Couldn't find a real system that naturally fits into the model
 - ▶ Large majority of systems too simple
 - ▶ Suggestions?

REFERENCES



S. EILENBERG, C. C. ELGOT, AND J. C. SHEPHERDSON.

SETS RECOGNIZED BY n -TAPE AUTOMATA.

Journal of Algebra, 13:447–464, 1969.



E. MARK GOLD.

LANGUAGE IDENTIFICATION IN THE LIMIT.

Info. and Control, 10(5):447–474, 1967.



JAN K. PACHL.

PROTOCOL DESCRIPTION AND ANALYSIS BASED ON A STATE TRANSITION MODEL WITH CHANNEL EXPRESSIONS.

IN *Intl. Conf. on Protocol Specification, Testing and Verification (PSTV)*, PAGES 207–219, 1987.



M. O. RABIN AND D. SCOTT.

FINITE AUTOMATA AND THEIR DECISION PROBLEMS.

IBM Journal of Research and Development, 3:114–125, 1959.



JEAN-FRANÇOIS RASKIN AND FRÉDÉRIC SERVAIS.

VISIBLY PUSHDOWN TRANSDUCERS.

IN *Intl. Colloquium on Automata, Languages and Programming (ICALP), Part II*, PAGES 386–397, 2008.



WOLFGANG THOMAS.

ON LOGICAL DEFINABILITY OF TRACE LANGUAGES.

IN *ASMICS workshop, Technical University of Munich, Report TUM-I9002*, PAGES 172–182, 1990.