Contextuality in distributed systems

Nasos Evangelou-Oost
Joint work: Callum Bannister, Ian Hayes

3 April 2023

The University of Queensland
Australia

Contextuality

1. ***Contextuality*** is a general phenomenon where bits of information agree ***locally but not globally***.
2. Arises in many contexts:
   1. quantum mechanics: quantum contextuality, nonlocality, Bell–Kochen–Specker theorem;
   2. relational databases: join-inconsistency = a database that does not admit a universal relation;
   3. constraint satisfaction problems;
   4. logic: liar paradoxes, etc.
   5. ...
3. Contextuality has a general & formal definition in the theory of ***information algebras*** (Abramsky and Carù 2019).
4. Point of this talk: ***sequential inconsistency*** in the specification of a distributed system is also a form of contextuality, and we can study it using information algebras.

Information algebras I

An information algebra Φ consists of the following data:

Running example: relational databases

- A topological space $(X, \mathcal{T})$, comprising domains of information and their connectivity/proximity.

  X is the discrete space on a set of attributes;

  $$X := \{\mathtt{id}, \mathtt{name}, \mathtt{class}\}$$

- A (functorial) assignment $U \mapsto \Phi_U$, sending each $U \in \mathcal{T}$ to a set Φ_U of ***valuations*** = pieces of information involving variables of U.

  Φ_U is the set of relations on U; e.g.

  $$r := \{(1, \mathtt{Alice}), (2, \mathtt{Bob}), (3, \mathtt{Clara})\} \in \Phi_{\{\mathtt{id}, \mathtt{name}\}}$$

Information algebras II

- A ***query*** operation
$$(\phi, U) \mapsto \phi^{\downarrow U} \in \Phi_U$$
where $U \subseteq \mathrm{d}\phi$.

Projection of relations; e.g.

$$r = \{(1, \texttt{Alice}), (2, \texttt{Bob}), (3, \texttt{Clara})\} \in \Phi_{\{\texttt{id},\texttt{name}\}}, \qquad r^{\downarrow\{\texttt{name}\}} = \{(\texttt{Alice}), (\texttt{Bob}), (\texttt{Clara})\} \in \Phi_{\{\texttt{name}\}}$$

- An associative, commutative ***combine*** operation
$$(\phi, \psi) \mapsto \phi \otimes \psi \in \Phi_{\mathrm{d}\phi \cup \mathrm{d}\psi}$$

Natural join of relations; e.g.

$$r = \{(1, \texttt{Alice}), (2, \texttt{Bob}), (3, \texttt{Clara})\} \in \Phi_{\{\texttt{id},\texttt{name}\}}$$
$$s := \{(1, \texttt{Math}), (3, \texttt{English})\} \in \Phi_{\{\texttt{id},\texttt{class}\}}$$

then

$$r \otimes s = r \bowtie s = \{(1, \texttt{Alice}, \texttt{Math}), (3, \texttt{Clara}, \texttt{English})\} \in \Phi_{\{\texttt{id},\texttt{name},\texttt{class}\}}$$

Information algebras III

These operations must satisfy some axioms; the most important is the ***combination axiom***:

$$(\phi \otimes \psi)^{\downarrow \mathrm{d}\phi} = \phi \otimes \psi^{\downarrow \mathrm{d}\phi \cap \mathrm{d}\psi}$$

We often consider ***ordered/adjoint information algebras***, where each Φ_U is a ***poset***. Then

$$\phi \leq \psi$$

means ϕ is 'less informative' than ψ.

Contextuality in information algebras I

Definition

1. A ***knowledgebase*** is a subset $K := \{\phi_i\}_{i\in I} \subseteq \Phi$ of valuations.
2. K ***locally agrees*** iff $\forall i, j \in I$
$$\phi_i^{\downarrow \mathrm{d}\phi_i \cap \mathrm{d}\phi_j} = \phi_j^{\downarrow \mathrm{d}\phi_i \cap \mathrm{d}\phi_j}$$
3. K ***globally agrees*** iff $\exists \phi.\ \mathrm{d}\phi = \bigcup_{i\in I} \mathrm{d}\phi_i, \quad \forall i \in I,$
$$\phi^{\downarrow \mathrm{d}\phi_i} = \phi_i$$

Definition

A knowledgebase K is called ***contextual*** iff it locally agrees but not globally agrees..

For relational databases, a knowledgebase is a set of relations = a database. A contextual database is said to be 'join-inconsistent' / not admitting a 'universal relation' (see Abramsky and Carù 2019 for an example)

Proposition (Abramsky and Carù 2019)

Let Φ be an adjoint information algebra.
Let $K := \{\phi_1, \ldots, \phi_n\} \subseteq \Phi$ be a knowledgebase.
Let

$$\gamma := \bigotimes K = \bigotimes_{i=1}^{n} \phi_i$$

Then K agrees globally if and only if, for each i,

$$\gamma^{\downarrow \mathrm{d}\phi_i} = \phi_i$$

Moreover, γ is the most informative valuation with this property.

'*A system is distributed if the message transmission delay is not negligible compared to the time between events in a single process.*' (Lamport 1978)

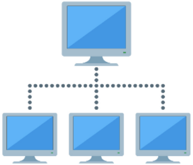

Distributed systems as information algebras I

- $(X, \mathcal{T})$ is the set of all variables/memory locations and the topology encodes their connectivity/proximity.

  E.g. three connected systems can be represented by a triangle (3 vertices, 3 edges).

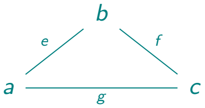

- Φ_U is the set of all possible subsets of ***sequences of states***, or ***(relative) traces*** on U.

  E.g., if each of the variables a, e, g can take on the values $0, 1$, then the set of all possible traces on $U := \{e, a, g\}$ is the set of all sequences over the alphabet $\{0, 1\}^3$.

Distributed systems as information algebras II

- The ***query*** operation

$$(\phi, U) \mapsto \phi^{\downarrow U}$$

gives the subset of all ***stuttering-reduced*** traces $\phi^{\downarrow U}$ on U, where each $t' \in \phi^{\downarrow U}$ is the restriction of some $t \in \phi$.

E.g. let

$$\phi := \left\{ \begin{array}{c} e \\ a \\ g \end{array} \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} \right\} \in \Phi_{\{e,a,g\}}$$

corresponding to the trace on $\{e, a, g\}$,

$$(0, 0, 1) \rightsquigarrow (0, 1, 1) \rightsquigarrow (1, 1, 1) \rightsquigarrow (0, 1, 1)$$

Projecting ϕ *onto* $\{a\}$ *gives*

$$\phi|_{\{a\}} = \left\{ \begin{bmatrix} 0 & 1 \end{bmatrix} \right\}$$

corresponding to the trace on $\{a\}$,

$$0 \rightsquigarrow 1$$

Distributed systems as information algebras III

- ***Combine*** is defined
$$\phi \otimes \psi := \{t \in \Phi_{\mathrm{d}\phi \cup \mathrm{d}\psi} \mid t_{\downarrow \mathrm{d}\phi} \in \phi, t_{\downarrow \mathrm{d}\psi} \in \psi\}$$
Here $t \mapsto t_{\downarrow U}$ is the projection on traces.
  - When $\mathrm{d}\phi \cap \mathrm{d}\psi = \emptyset$, $\otimes$ is a 'kind' of ***shuffle***:

    Let $\phi = \left\{ \begin{bmatrix} e_0 & e_1 \end{bmatrix} \right\} \in \Phi_{\{e\}}$ *and* $\psi = \left\{ \begin{bmatrix} f_0 & f_1 \end{bmatrix} \right\} \in \Phi_{\{f\}}$.

    Then
$$\phi \otimes \psi = \left\{ \begin{bmatrix} e_0 & e_1 \\ f_0 & f_1 \end{bmatrix}, \begin{bmatrix} e_0 & e_0 & e_1 \\ f_0 & f_1 & f_1 \end{bmatrix}, \begin{bmatrix} e_0 & e_1 & e_1 \\ f_0 & f_0 & f_1 \end{bmatrix} \right\} \in \Phi_{\{e,f\}}$$
  - When $\mathrm{d}\phi = \mathrm{d}\psi$, then $\otimes$ is set ***intersection***.
  - The general case is a mix of these effects.

Sequential consistency & contextuality in distributed systems

Sequential consistency is a property of distributed systems:

'*... the result of any execution is the same as if the operations of all the processors were executed in some sequential order, and the operations of each individual processor appear in this sequence in the order specified by its program.*' (Lamport 1979)

A specification of a distributed system is associated to a knowledgebase K.

Local agreement of K is essential for correctness of its specification: any two elements of K must agree on the common part of their domains.

But if K is ***contextual***, K locally agrees, yet there is no set of traces on the global state space that projects to the traces of each process in the system.

In line with Lamport's definition, we say the specification of a distributed system is sequentially consistent iff the associated knowledgebase both locally and globally agrees.

Example: dining philosophers I

(Example adapted from (Goguen 1992).)

3 philosophers p_0, p_1, p_2 sit at a circular table to eat, with one chopstick c_0, c_1, c_2 between each pair.

This defines a triangle with vertices p_0, p_1, p_2 and edges c_0, c_1, c_2 (***left***).

The triangle is associated to a topological space $(X, \mathscr{T})$ whose open sets are the upper sets of the face poset of the triangle (***right***).

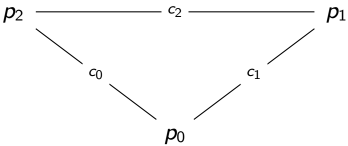

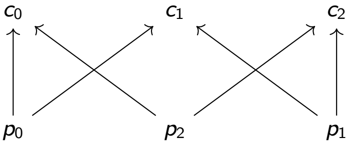

Example: dining philosophers II

Philosophers can ***think*** or ***eat***. To eat, they must hold both adjacent chopsticks.

Define states:

$$\Omega_{p_i} := \{t, e\}$$

where t = 'thinking', e = 'eating', and

$$\Omega_{c_i} := \{i-1, *, i\}$$

where $*$ = chopstick on the table, and $i-1, i$ stand for philosophers resp. p_{i-1}, p_i (who may hold chopstick c_i).

Define a cover of the space X, encoding the division of the specification into independent asynchronous processes:

$$\mathcal{U} = \{U_0, U_1, U_2\} = \{\{c_0, p_0, c_1\}, \{c_1, p_1, c_2\}, \{c_2, p_2, c_0\}\}$$

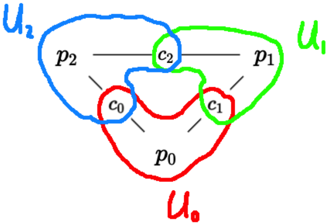

Example: dining philosophers III

Consider a specification containing all traces according to the following protocol:

Legal state transitions on $\{c_i\}$, for $x \in \Omega_{c_i}$:

$$* \mapsto x \qquad \text{(chopstick may be picked up)}$$

$$x \mapsto * \qquad \text{(chopstick may be put down)}$$

$$x \mapsto x \qquad \text{(chopstick may remain in current state)}$$

Legal state transitions on U_i:

$$\begin{array}{rlll}
 & \rightsquigarrow (*, t, *) & & \text{(1: \textit{begin thinking, without chopsticks})} \\
(l, x, r) & \rightsquigarrow (l', x, r') & l, l', r, r' \neq i & \text{(2: \textit{if no chopsticks, continue w. o. constraining nbhrs})} \\
(l, t, r) & \rightsquigarrow (l', e, r') & l, l', r, r' \neq i & \text{(3: \textit{if no chopsticks and thinking, become hungry w. o. constraining nbhrs})} \\
(l, e, *) & \rightsquigarrow (l', e, i) & l, l' \neq i & \text{(4: \textit{if no chopsticks, pick up right one w. o. constraining left nbhr})} \\
(l, e, i) & \rightsquigarrow (l', e, i) & l, l' \neq i & \text{(5: \textit{remain having right chopstick w. o. constraining left nbhr})} \\
(*, e, i) & \rightsquigarrow (i, e, i) & & \text{(6: \textit{have right chopstick, pick up left})} \\
(i, e, i) & \rightsquigarrow (*, t, *) & & \text{(7: \textit{both chopsticks and eating, put it down and think})}
\end{array}$$

Example: dining philosophers IV

Let $K := \{\phi_0, \phi_1, \phi_2\}$ where

$$\phi_i := \left\{ \begin{array}{c} c_i \\ p_i \\ c_{i+1} \end{array} \begin{bmatrix} * & * & * & (i-1) & * & i & * & * \\ t & e & e & e & e & e & t & t \\ * & * & i & i & i & i & * & (i+1) \end{bmatrix} \right\}, \qquad \mathrm{d}\phi_i = U_i = \{c_i, p_i, c_{i+1}\}$$

E.g., ϕ_1 contains a single trace encoding the following sequence of events:

1. p_1 becomes hungry (rule 3);
2. p_1 picks up the right chopstick c_2 (rule 4);
3. p_0 picks up the left chopstick c_1 (rule 5);
4. p_0 puts down the left chopstick c_1 (rule 5);
5. p_1 picks up the left chopstick c_1 (rule 6);
6. p_1 eats and puts down both chopsticks c_1 and c_2 (rule 7);
7. p_2 picks up the right chopstick c_2 (rule 2).
8. p_2 puts down the right chopstick c_2 (rule 2).

Example: dining philosophers V

$$\phi_i := \left\{ \begin{array}{c} c_i \\ p_i \\ c_{i+1} \end{array} \begin{bmatrix} * & * & * & (i - 1) & * & i & * & * & * \\ t & e & e & e & e & e & t & t & t \\ * & * & i & i & i & i & * & (i+1) & * \end{bmatrix} \right\}$$

K locally agrees:

$$\phi_i^{\downarrow\{c_i\}} = \left\{ \begin{bmatrix} * & (i-1) & * & i & * \end{bmatrix} \right\} = \phi_{i-1}^{\downarrow\{c_i\}}$$

But K does not globally agree!

1. ϕ_0 says p_0 picks up c_1 before p_2 picks up c_0,
2. ϕ_2 says p_2 picks up c_0 before p_1 picks up c_2,
3. ϕ_1 says p_1 picks up c_2 before p_0 picks up c_1,

Together these events form a causal loop, which is physically impossible & not representable as a trace on X!

$\therefore$ the specification defined by K is ***contextual/sequentially inconsistent***.

Summary

Existing theory:

1. Contextuality is a general phenomenon where information locally agrees but globally disagrees.
2. Information algebras are a convenient formalism for reasoning about contextuality.
3. There is an efficient computational theory for detecting contextuality in information algebras.

Our contribution:

1. Contextuality manifests in the specification of distributed systems as a violation of sequential consistency.
2. Specifications form a complete 'refinement' lattice whose ordering $\preceq$ means the LHS has less nondeterminism but more parallelism than the RHS. We assign a knowledgebase to each specification to detect contextuality.

Thank you!

References

Abramsky, Samson and Giovanni Carù (2019). 'Non-locality, contextuality and valuation algebras: a general theory of disagreement'. In: ***Philos. Trans. Roy. Soc. A*** 377.2157, pp. 20190036, 22. ISSN: 1364-503X. DOI: 10.1098/rsta.2019.0036.

Evangelou, Alexander (2022). ***Hand-drawn Penrose triangle***.

Evangelou-Oost, Nasos, Callum Bannister and Ian J. Hayes (2023). 'Contextuality in Distributed Systems'. In: ***Relational and Algebraic Methods in Computer Science***. Ed. by Roland Glück, Luigi Santocanale and Michael Winter. Cham: Springer International Publishing, pp. 52–68. ISBN: 978-3-031-28083-2.

Goguen, Joseph A. (1992). 'Sheaf Semantics for Concurrent Interacting Objects'. In: ***Math. Struct. Comput. Sci.*** 2.2, pp. 159–191. DOI: 10.1017/S0960129500001420.

Lamport, Leslie (1978). 'Time, Clocks, and the Ordering of Events in a Distributed System'. In: ***Commun. ACM*** 21.7, pp. 558–565. DOI: 10.1145/359545.359563.

— (1979). 'How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs'. In: ***IEEE Trans. Computers*** 28.9, pp. 690–691. DOI: 10.1109/TC.1979.1675439.

macrovector_official (n.d.). ***Free vector network icons set***. URL: https://www.freepik.com/free-vector/network-icons-set_4378386.htm.

Pouly, Marc and Jürg Kohlas (2012). ***Generic inference: a unifying theory for automated reasoning***. John Wiley & Sons.