The analytic semantics of weakly consistent parallelism

Jade Alglave (UCL, MSR-Cambridge, UK)
Patrick Cousot (NYU, Emer. ENS)

Jayadev Misra Fest Workshop

Programming: Logics, Models, Algorithms, and Concurrency

U.T. Austin

April 29th–30th, 2016

Analytic semantics

Weak consistency models (WCM)

- Sequential consistency:
  reads read(p,x) are implicitly *coordinated* with writes w(q,x)

- WCM:
  No implicit coordination (depends on architecture, program dependencies, and explicit fences)

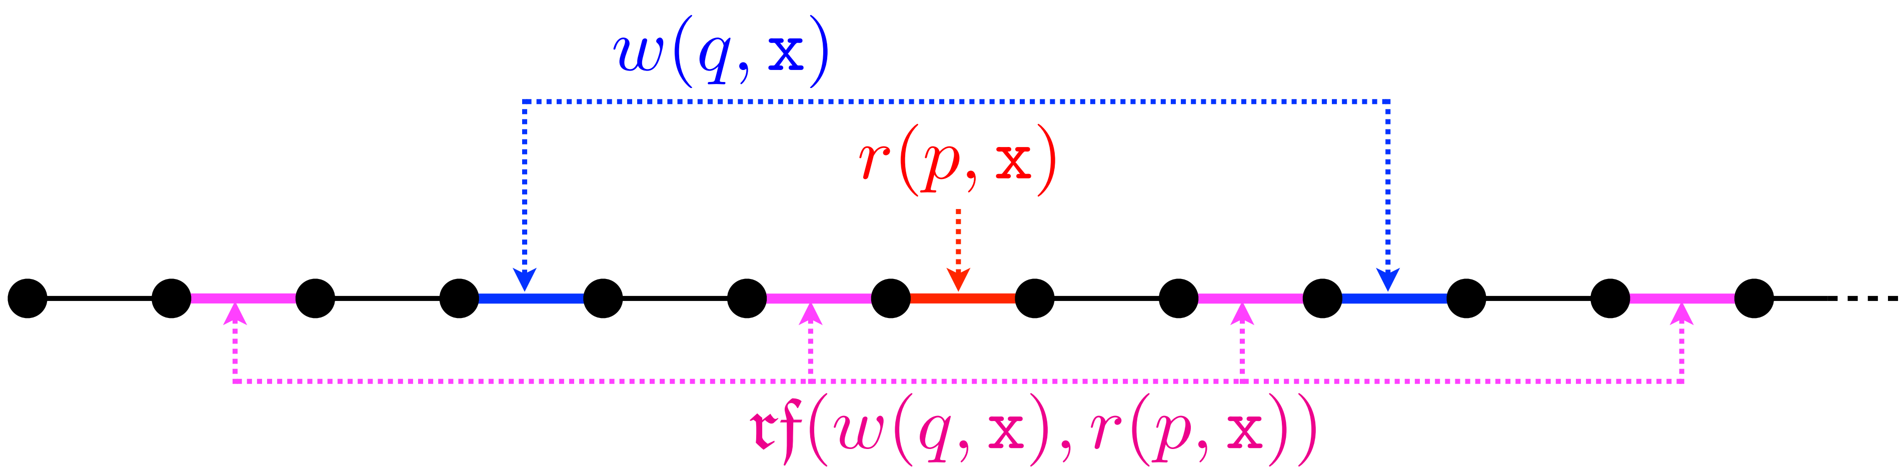

Analytic semantic specification

- Anarchic semantics:

  describes computations, no constraints on communications

- cat specification (Jade Alglave & Luc Maranget):

  imposes architecture-dependent communication constraints

- Hierarchy of anarchic semantics:

  many different styles to describe the same computations (e.g. interleaved versus true parallelism)

Example: load buffer (LB)

- Program:

```
{ x = 0; y = 0; }
P0          | P1         ;
r[] r1 x    | r[] r2 y   ;
w[] y 1     | w[] x   1  ;
exists(0:r1=1 /\ 1:r2=1)
```

- Example of execution trace $t \in S^{\perp}[\![\mathtt{P}]\!]$:

$t = w(\text{start}, \mathrm{x}, 0)\ w(\text{start}, \mathrm{y}, 0)\ r(\mathtt{P0}, \mathrm{x}, 1)\ \mathfrak{rf}[w(\mathtt{P1}, \mathrm{x}, 1), r(\mathtt{P0}, \mathrm{x}, 1)])\ w(\mathtt{P0}, \mathrm{y}, 1)\ r(\mathtt{P1}, \mathrm{y}, 1)$
$w(\mathtt{P1}, \mathrm{x}, 1)\ \mathfrak{rf}[w(\mathtt{P0}, \mathrm{y}, 1), r(\mathtt{P1}, \mathrm{y}, 1)]\ r(\text{finish}, \mathrm{x})\ \mathfrak{rf}[w(\mathtt{P1}, \mathrm{x}, 1), r(\text{finish}, \mathrm{x}, 1)]$
$r(\text{finish}, \mathrm{y}, 1)\ \mathfrak{rf}[w(\mathtt{P0}, \mathrm{y}, 1), r(\text{finish}, \mathrm{y}, 1)]$

- Abstraction to cat *candidate execution* $\alpha_{\Xi}(t)$

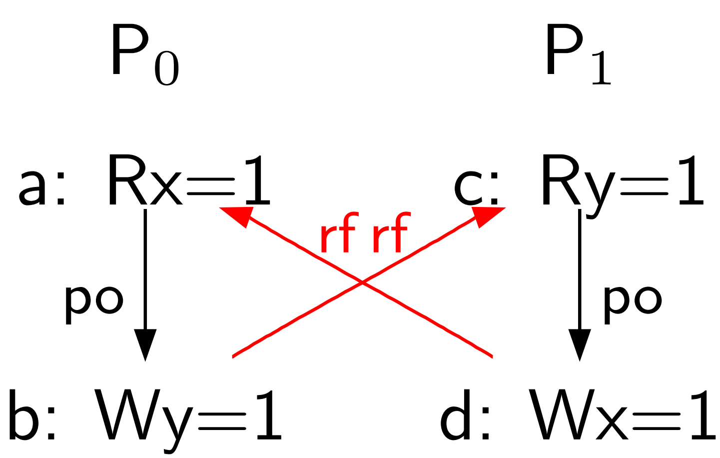

Example: load buffer (LB), cont'd

- cat specification:

```
acyclic (po | rf)+
```

The cat semantics rejects this execution $\alpha_{\Xi}(t)$:

$$[\![\texttt{cat}]\!]\,(\alpha_{\Xi}(t)) \;=\; \textsf{false}$$

The WCM semantics

- Abstraction to a candidate execution:

$$\alpha_{\Xi}(t) \triangleq \langle \alpha_e(t),\ \alpha_{po}(t),\ \alpha_{rf}(t),\ \alpha_{iw}(t),\ \alpha_{fw}(t) \rangle$$

$$\alpha_{\Xi}(\mathcal{S}) \triangleq \{\langle t,\ \alpha_{\Xi}(t)\rangle \mid t \in \mathcal{S}\}$$

- The cat semantics:

$$\alpha_{\text{🐱}}[\![\texttt{cat}]\!](\mathcal{S}) \triangleq \{t \mid \langle t,\ \Xi\rangle \in \mathcal{S} \wedge \text{🐱}[\![\texttt{cat}]\!](\Xi)\}$$

- The WCM semantics:

$$\alpha_{\text{🐱}}[\![\texttt{cat}]\!] \circ \alpha_{\Xi}(\mathcal{S}[\![\mathrm{P}]\!])$$

GC: $\langle \wp(\mathfrak{E}^{+\infty}),\ \subseteq\rangle \underset{\alpha_{\Xi}}{\overset{\gamma_{\Xi}}{\leftrightarrows}} \langle \wp(\mathfrak{E}^{+\infty} \times \Xi),\ \subseteq\rangle \underset{\alpha_{\text{🐱}}[\![\texttt{cat}]\!]}{\overset{\gamma_{\text{🐱}}[\![\texttt{cat}]\!]}{\leftrightarrows}} \langle \wp(\mathfrak{E}^{+\infty}),\ \subseteq\rangle$

Definition of the anarchic semantics

Axiomatic parameterized definition of the anarchic semantics

- The semantics $\mathcal{S}^{\perp}[\![\mathtt{P}]\!]$ is a finite/infinite sequence of *interleaved* events of processes satisfying well-formedness conditions.
- Example: computation (local variable assignment)

register assignment event by process p in trace τ

unique event stamp θ

$$\forall p \in \mathbb{P}\mathbb{i} \;.\; \forall k \in]1, 1+|\tau|[\;.\; \forall \ell \in \mathbb{L}(p) \;.\; \forall v \in \mathcal{D} \;. \qquad (\mathsf{Wf}_{22}(\tau))$$
$$(\exists \theta \in \mathfrak{P}(p) \;.\; \overline{\tau}_k = \mathfrak{a}(\langle p,\ \ell,\ \mathtt{r}\ \mathtt{:=}\ e,\ \theta\rangle, v))$$
$$\Longrightarrow (\ell \in \mathsf{N}^p(\tau, k) \wedge \mathsf{action}(p, \ell) = \mathtt{r}\ \mathtt{:=}\ e \wedge v = \mathsf{E}^p[\![e]\!](\tau, k-1)) \;.$$

control of process p is at label ℓ

action of process p is at label ℓ is a register assignment

value v of e is evaluated by past-travel

Axiomatic parameterized definition of the anarchic semantics

- Example: communication
  - a read event is initiated by a read action:

read event by process p in trace τ

unique media variable (L-value)

$$\forall p \in \mathbb{P}\mathrm{i} \,.\, \forall k \in]1, 1+|\tau|[\,.\, \forall \ell \in \mathbb{L}(p) \,. \qquad (\mathsf{Wf}_{23}(\tau))$$
$$(\exists \theta \in \mathfrak{P}(p) \,.\, (\overline{\tau}_k = \mathfrak{r}(\langle p,\ \ell,\ \mathtt{r} := \mathtt{x},\ \theta\rangle, \mathtt{x}_\theta)))$$
$$\Longrightarrow (\ell \in N^p(\tau, k) \wedge \mathsf{action}(p, \ell) = \mathtt{r} := \mathtt{x}) \,.$$

  - a read must read-from ($\mathfrak{rf}$) a write (fairness):

$$\forall p \in \mathbb{P}\mathrm{i} \,.\, \forall i \in]1, 1+|\tau|[\,.\, \forall r \in \mathfrak{Rf}(p) \,. \qquad (\mathsf{Wf}_{26}(\tau))$$
$$(\overline{\tau}_i = r) \Longrightarrow (\exists j \in]1, 1+|\tau|[\,.\, \exists w \in \mathfrak{Wi} \,.\, \overline{\tau}_j = \mathfrak{rf}[w, r]) \,.$$

Axiomatic parameterized definition of the anarchic semantics

- Predictive evaluation of media variables:

$$\begin{aligned}\mathcal{V}^p_{(32)}[\![\mathtt{x}_\theta]\!](\tau, k) \triangleq v \ \textbf{where} \ & \exists! i \in [1, 1+|\tau|[\ .\ (\overline{\tau}_i = \mathfrak{r}(\langle p,\ \ell,\ \mathtt{r} := \mathtt{x},\ \theta\rangle, \mathtt{x}_\theta)) \wedge \\ & \exists! j \in [1, 1+|\tau|[\ .\ (\overline{\tau}_j = \mathfrak{rf}[\mathfrak{w}(\langle p',\ \ell',\ \mathtt{x} := e',\ \theta'\rangle, v), \overline{\tau}_i])\end{aligned}$$

- Local path-based evaluation of an expression:

$$\begin{aligned}\mathcal{E}^p_{(30)}[\![\mathtt{r}]\!](\tau, k) &\triangleq v \quad \text{if } k > 1 \wedge \big((\overline{\tau}_k = \mathfrak{a}(\langle p,\ \ell,\ \mathtt{r} := e,\ \theta\rangle, v)) \vee \\ &\qquad (\overline{\tau}_k = \mathfrak{r}(\langle p,\ \ell,\ \mathtt{r} := \mathtt{x},\ \theta\rangle, \mathtt{x}_\theta) \wedge \mathcal{V}^p[\![\mathtt{x}_\theta]\!](\tau, k) = v)\big) \\ \mathcal{E}^p_{(30)}[\![\mathtt{r}]\!](\tau, 1) &\triangleq \mathcal{I}[\![0]\!] \qquad \textit{i.e. } \overline{\tau}_1 = \epsilon_{\mathsf{start}} \text{ by } \mathsf{Wf}_{15}(\tau) \\ \mathcal{E}^p_{(30)}[\![\mathtt{r}]\!](\tau, k) &\triangleq \mathcal{E}^p_{(30)}[\![\mathtt{r}]\!](\tau, k-1) \qquad \text{otherwise.}\end{aligned}$$

Abstractions of the anarchic semantics

Abstractions

- Semantics:

$$S^{\perp}[\![\mathrm{P}]\!] \triangleq \boldsymbol{\lambda}\langle \mathcal{B},\ \mathsf{sat},\ \mathcal{D},\ I,\ \mathfrak{S},\ V,\ E,\ N\rangle \bullet \{\tau \in \mathfrak{T}[\![\mathrm{P}]\!]|_{\cong} \mid \mathsf{Wf}_{15}(\tau) \wedge \ldots \wedge \mathsf{Wf}_{29}(\tau)\}$$

parameters of the semantics — trace well-formedness conditions

- Examples of abstractions:
  - Choose data (e.g. ground values, uninterpreted symbolic expressions, interpreted symbolic expressions i.e. "symbolic guess")
  - Bind parameters (e.g. how expressions are evaluated)
  - …

The hierarchy of interleaved semantics

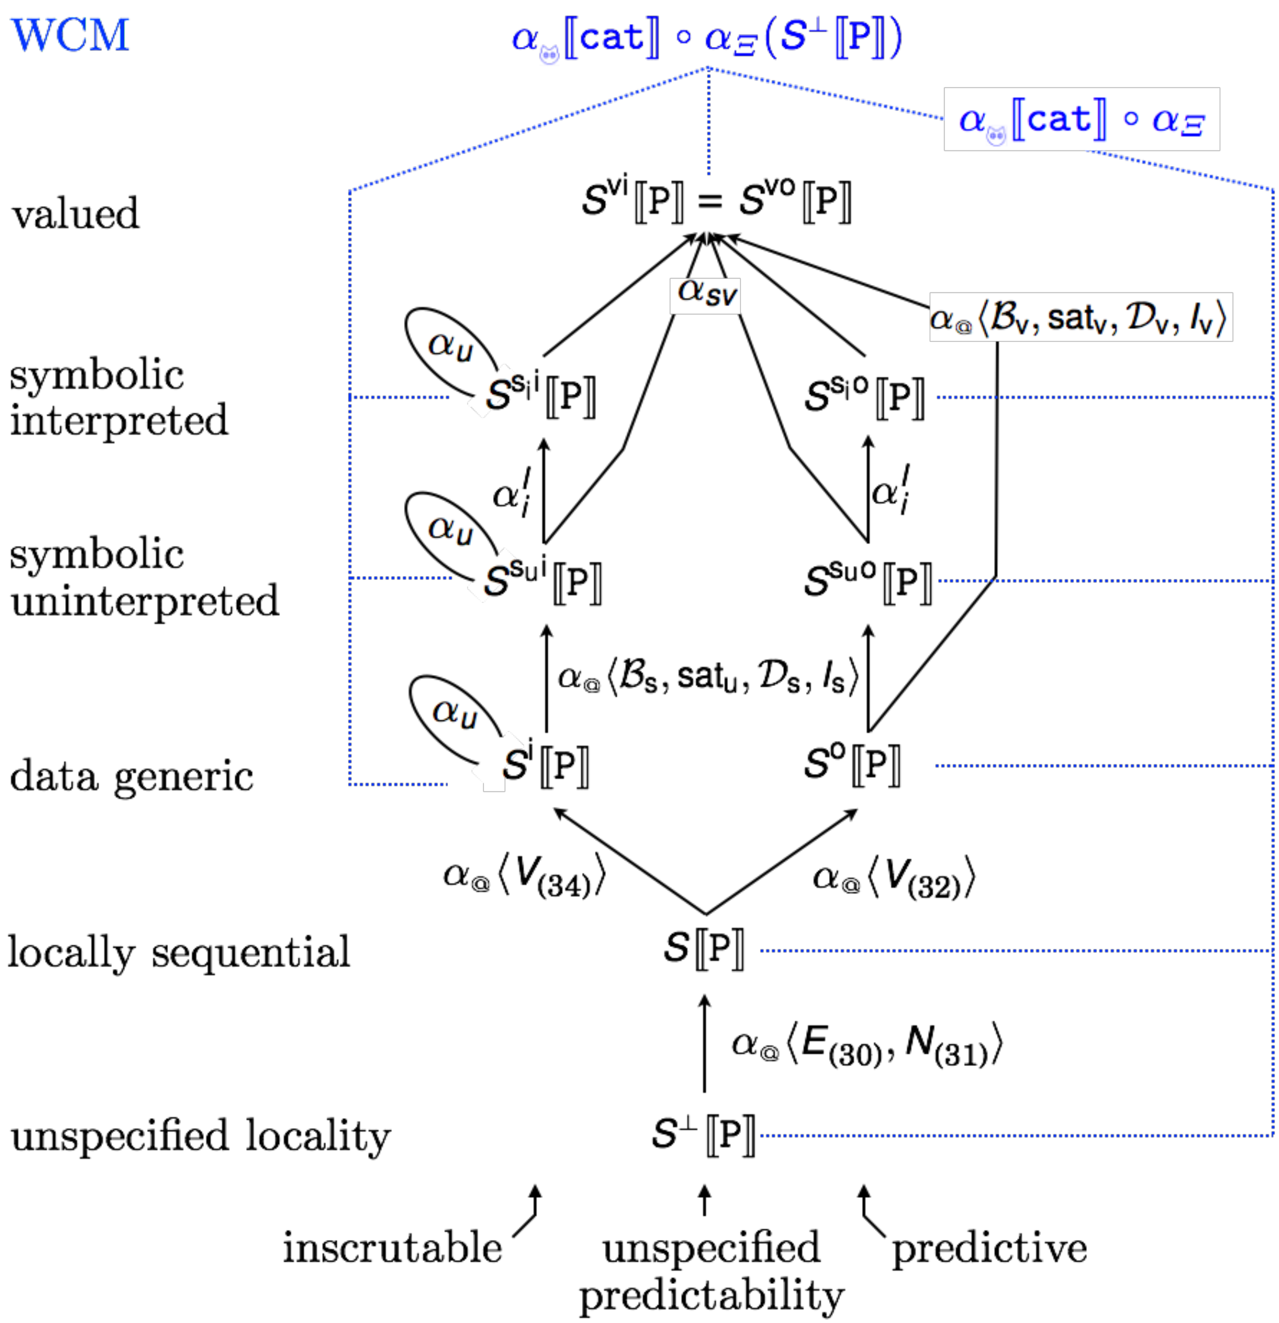

True parallelism

- Extract from interleaved executions:
  - The subtrace of each process (sequential execution)
  - The rf communication relation (interactions between processes)

$\Longrightarrow$ no more global time !

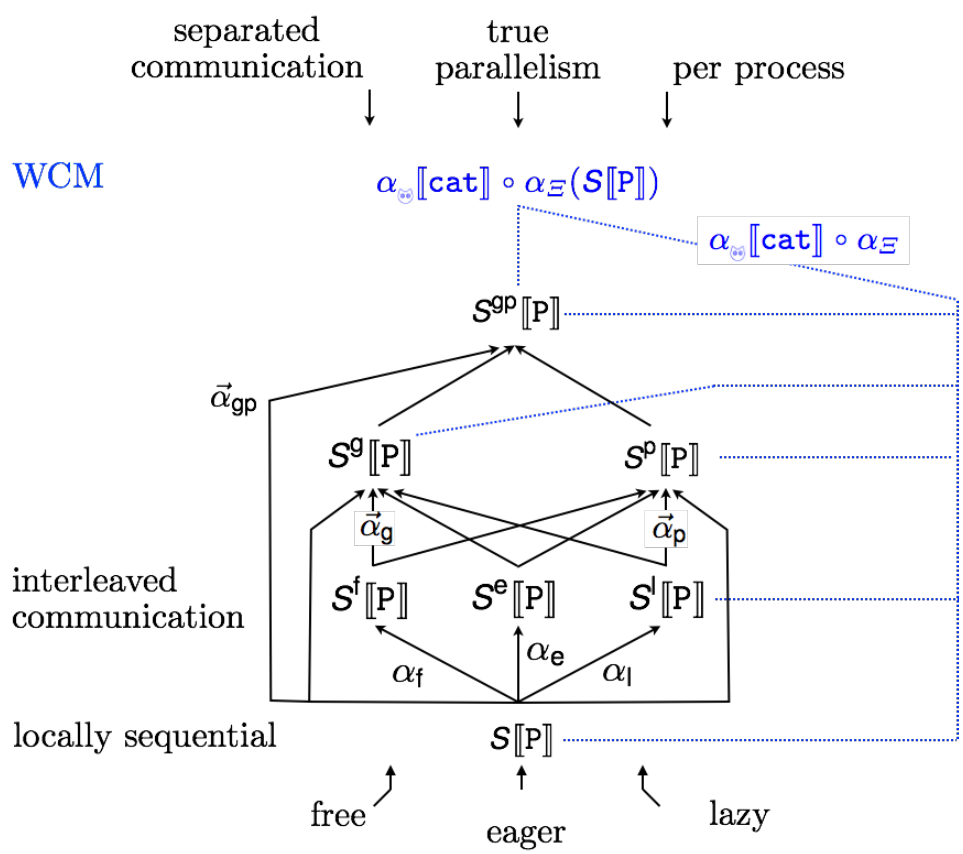

States

- At each point in a trace, the state abstracts the past computation history up to that point
- Example: classical environment (assigning values to register at each point k of the trace):

$$\rho^p(\tau, k) \triangleq \boldsymbol{\lambda}\, \mathtt{r} \in \mathbb{R}(p) \bullet \mathit{E}^p[\![\mathtt{r}]\!](\tau, k)$$

$$\nu^p(\tau, k) \triangleq \boldsymbol{\lambda}\, \mathtt{x}_\theta \bullet \mathit{V}^p_{(32)}[\![\mathtt{x}_\theta]\!](\tau, k)$$

Prefixes, transitions, …

- Abstract traces by their prefixes:

$$\overleftarrow{\alpha}(\mathcal{S}) \triangleq \bigcup\{\overleftarrow{\alpha}(\tau) \mid \tau \in \mathcal{S}\}$$
$$\overleftarrow{\alpha}(\tau) \triangleq \{\tau\langle\!\![j\langle\!\![\mid j \in [1, 1+|\tau|[\}$$
$$\tau\langle\!\![j\langle\!\![\triangleq \langle \xrightarrow{\overline{\tau}_i} \underline{\tau}_i \mid i \in [1, 1+j[\rangle$$

- and transitions: extract transitions from traces

$\Longrightarrow$ communication fairness is lost, inexact abstraction,
$\Longrightarrow$ add fairness condition

Effect of the cat specification on the hierarchy

Exactness and cat preservation

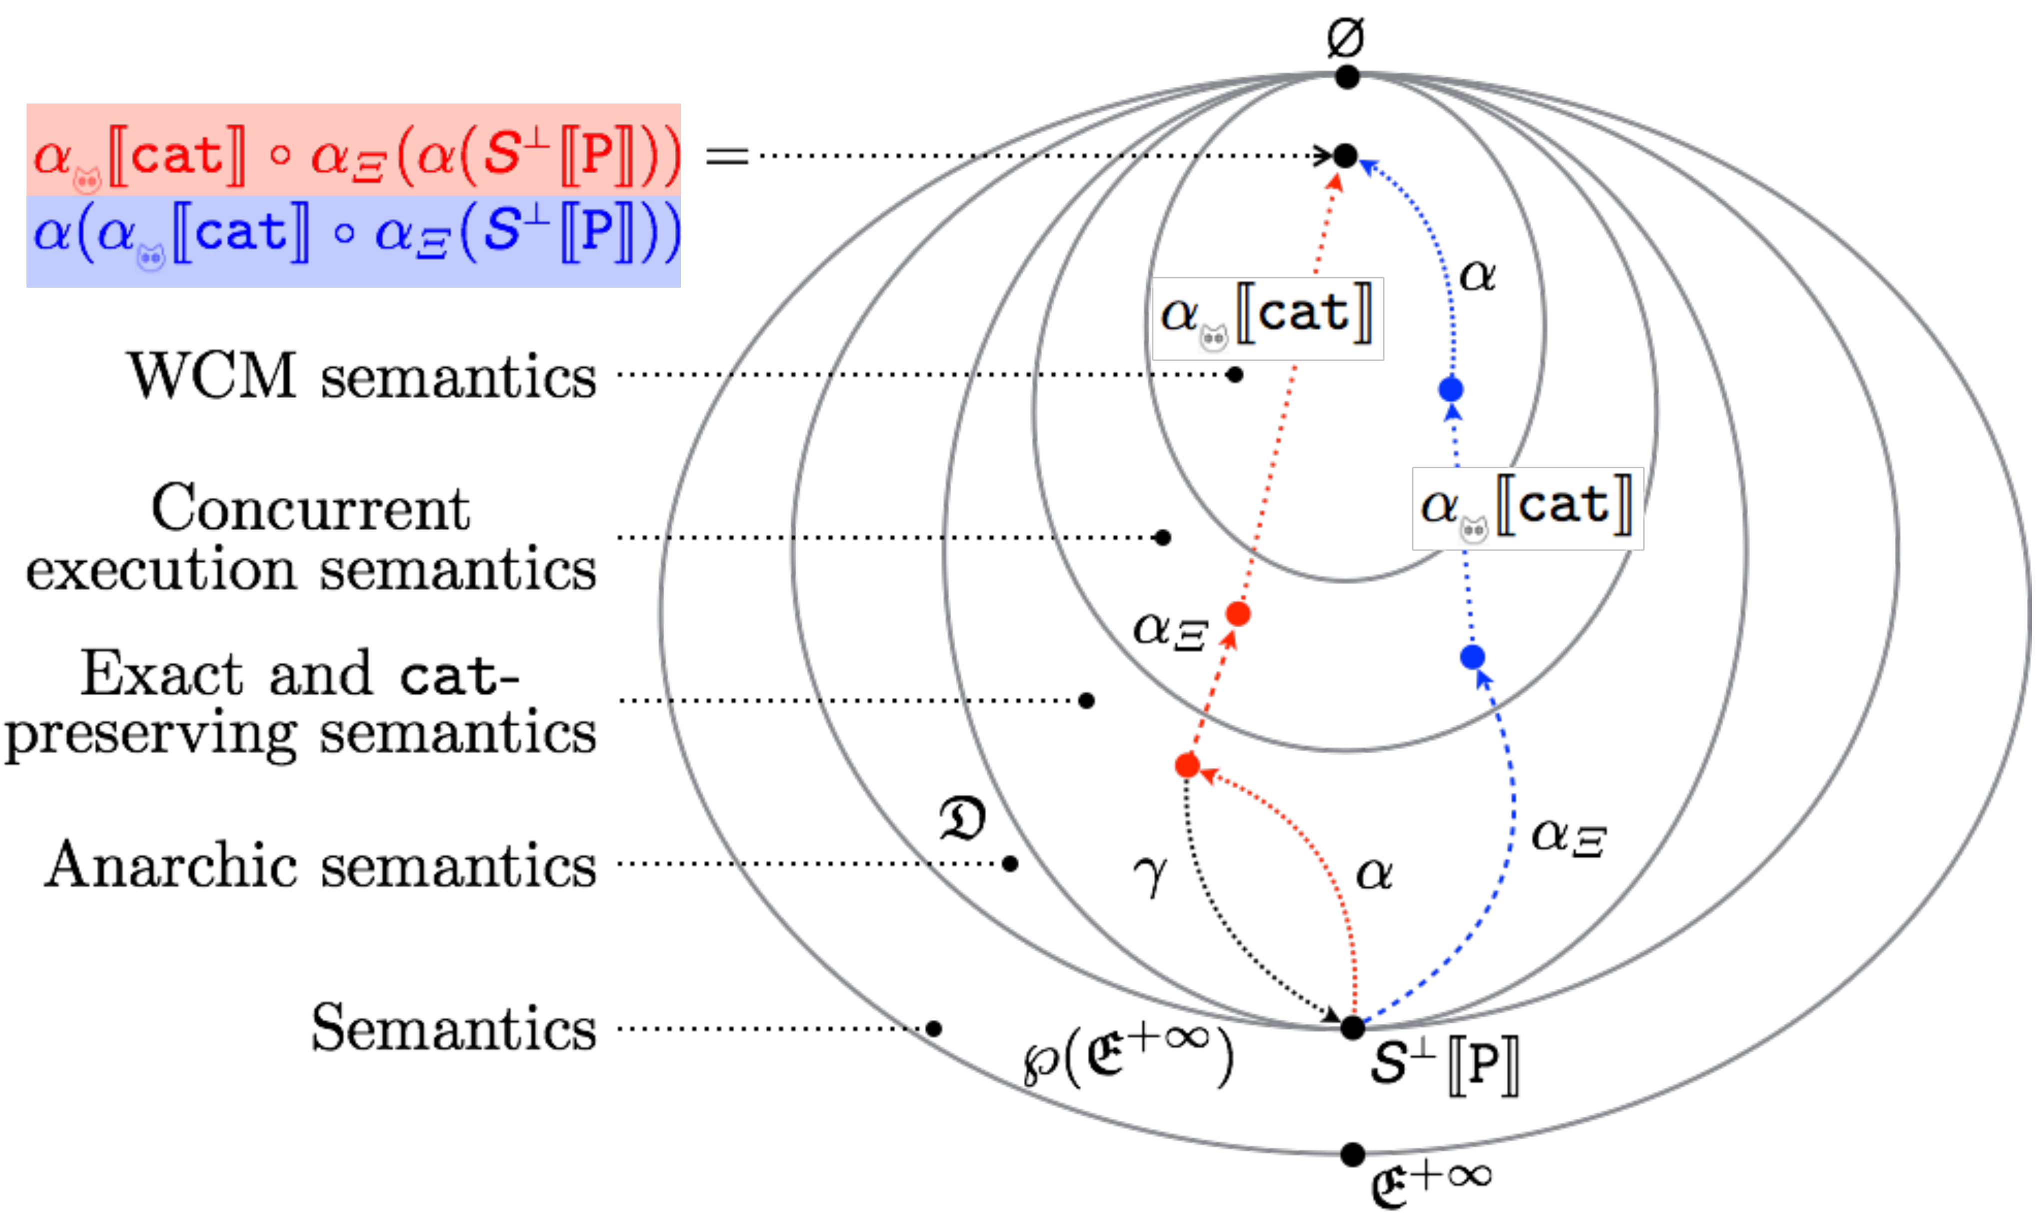

The cat abstraction

- The same cat specification $\alpha_{🐱}[\![\mathtt{cat}]\!]$ applies equally to any concurrent execution abstraction α_{Ξ} of any interleaved/truly parallel semantics in the hierarchy

- The appropriate level of abstraction to specify WCM:
  - No states, only marker (e.g. fence), r, w, rf(w,r) events
  - No values in events
  - No global time (only po order of events per process)
  - Time of communications forgotten (only rf of who communicates with whom)

Conclusion

Conclusion

- The hierarchy of anarchic semantics describe the same computations and potential communications in very different styles
- The cat semantics restricts communications to a machine/network architecture in the same way for all semantics in the hierarchy
- This idea of parameterized semantics at various levels of abstraction is useful for
  - Verification
  - Static analysis

The End