Connectivity

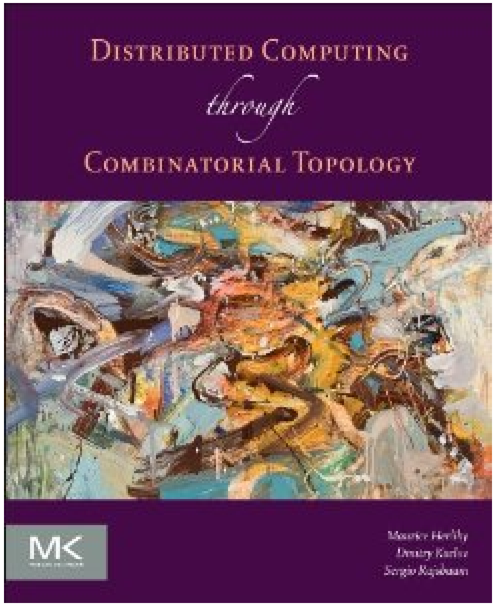

Companion slides for

Distributed Computing

Through Combinatorial Topology

Maurice Herlihy & Dmitry Kozlov & Sergio Rajsbaum

Previously

Used Sperner's Lemma to show *k*-set agreement impossible when protocol complex is a manifold.

But in many models, protocol complexes are not manifolds …

Road Map

Consensus Impossibility

- General theorem
- Application to read-write models

k-set agreement Impossibility

- General theorem
- Application to read-write models

Road Map

Consensus Impossibility

General theorem

Application to read-write models

k-set agreement Impossibility

General theorem

Application to read-write models

A Path

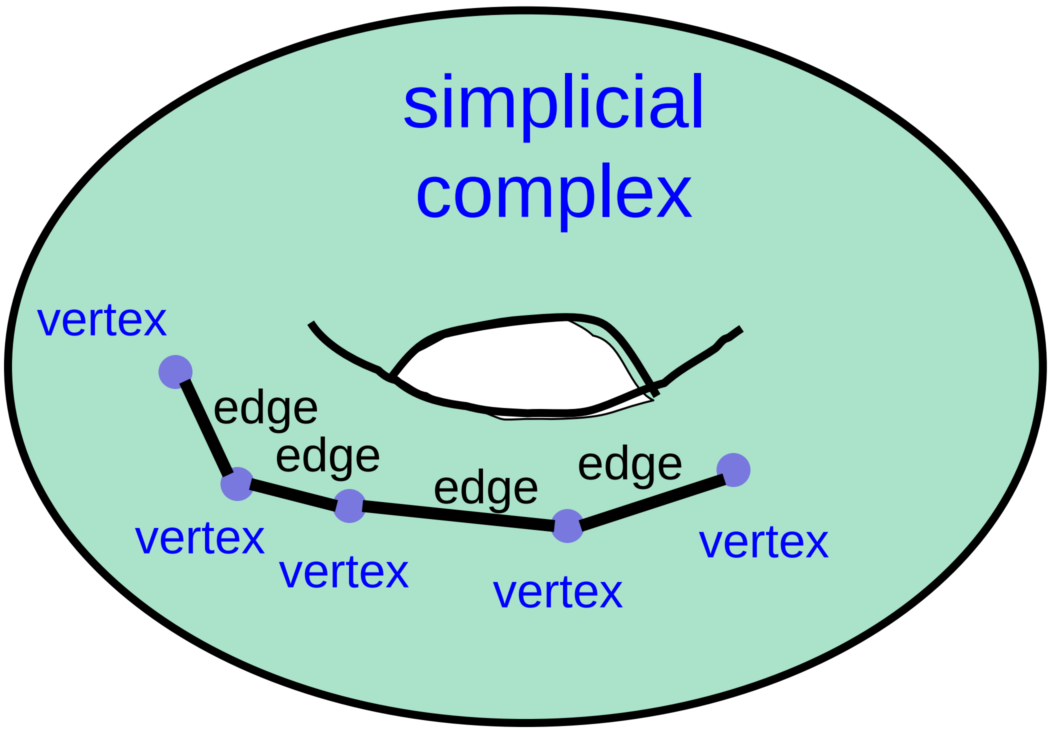
simplicial
complex
vertex
edge
edge
edge
edge
vertex
vertex
vertex
vertex

Path Connected

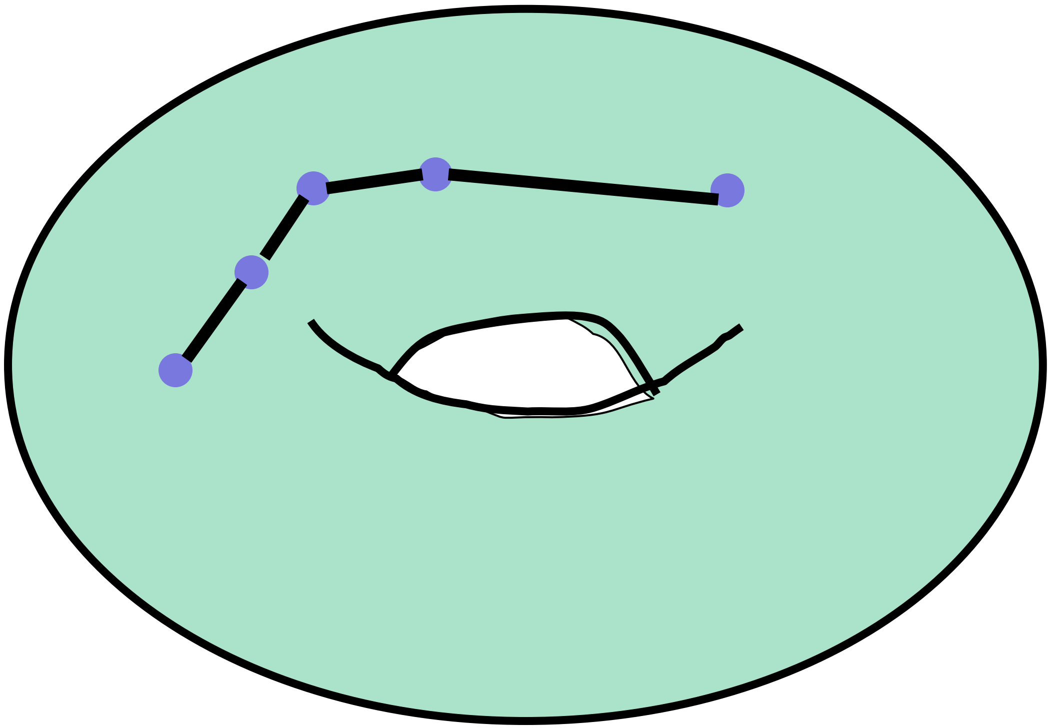

Any two vertexes can be linked by a path

Theorem

If, for protocol $(\mathcal{I},\mathcal{P},\mathit{\Xi})$ …

Theorem

If, for protocol $(\mathcal{I},\mathcal{P},\Xi)$ ...

For every *n*-simplex σ^n,

$\Xi(\sigma^n)$ is path-connected …

Theorem

If, for protocol $(\mathcal{I},\mathcal{P},\Xi)$ ...

For every *n*-simplex σ^n,

$\Xi(\sigma^n)$ is path-connected ...

(If no process fails)

Theorem

If, for protocol $(\mathcal{I},\mathcal{P},\Xi)$ …

For every *n*-simplex σ^n,
$\Xi(\sigma^n)$ is path-connected …

For every *(n-1)*-simplex $\sigma^{n\text{-}1}$,
$\Xi\ (\sigma^{n\text{-}1})$ is non-empty …

Then …

(If one process fails)

Theorem

If, for protocol $(\mathcal{I},\mathcal{P},\Xi)$ …

For every *n*-simplex σ^n,
$\Xi(\sigma^n)$ is path-connected …

For every *(n-1)*-simplex $\sigma^{n\text{-}1}$,
$\Xi(\sigma^{n\text{-}1})$ is non-empty …

Then …

$\Xi(\cdot)$ cannot solve consensus.

Model Independence

Holds for message-passing or shared memory ….

Synchronous, asynchronous, or in-between …

Any adversarial scheduler …

As long as one failure is possible.

Protocol Complex Notation

All 3 participate

$\Xi(PQR)$

2 participate

$\Xi(PQ)$, $\Xi(QR)$, $\Xi(PR)$,

1 participates

$\Xi(P)$, $\Xi(Q)$, $\Xi(R)$,

Proof

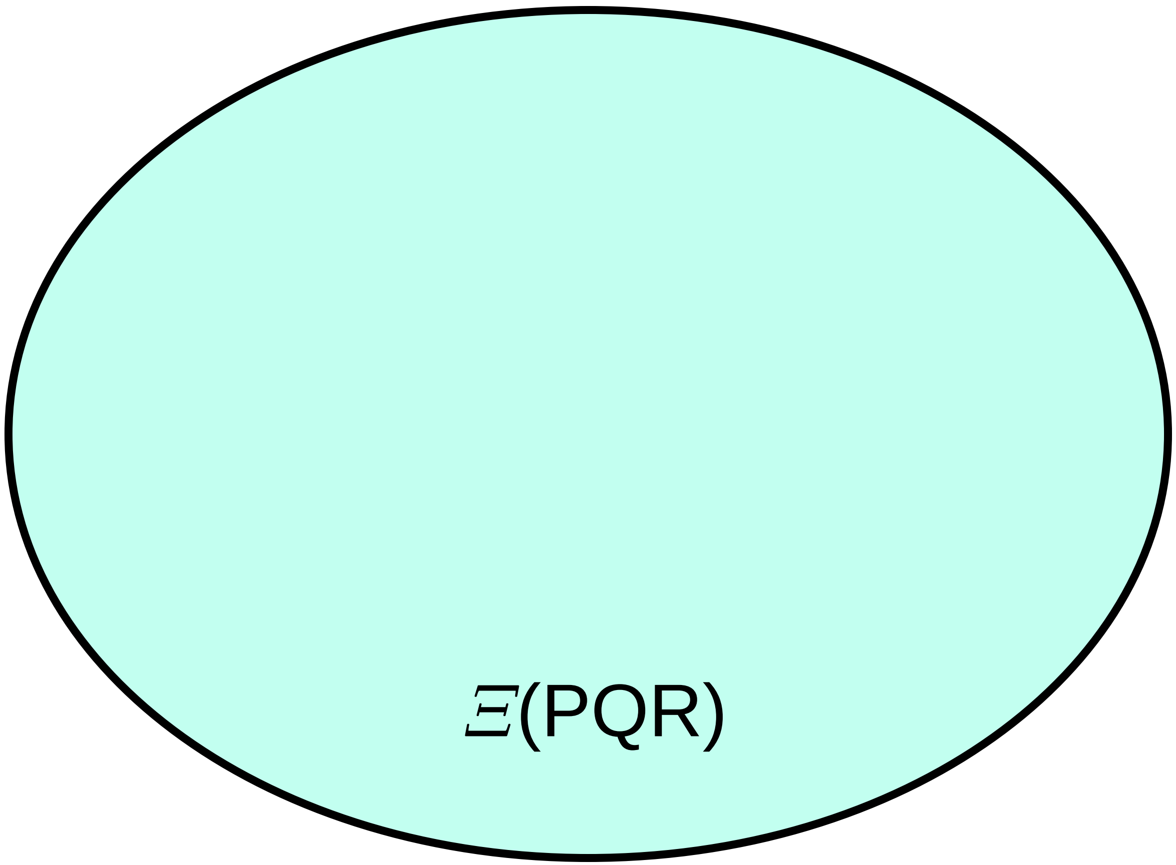

Proof

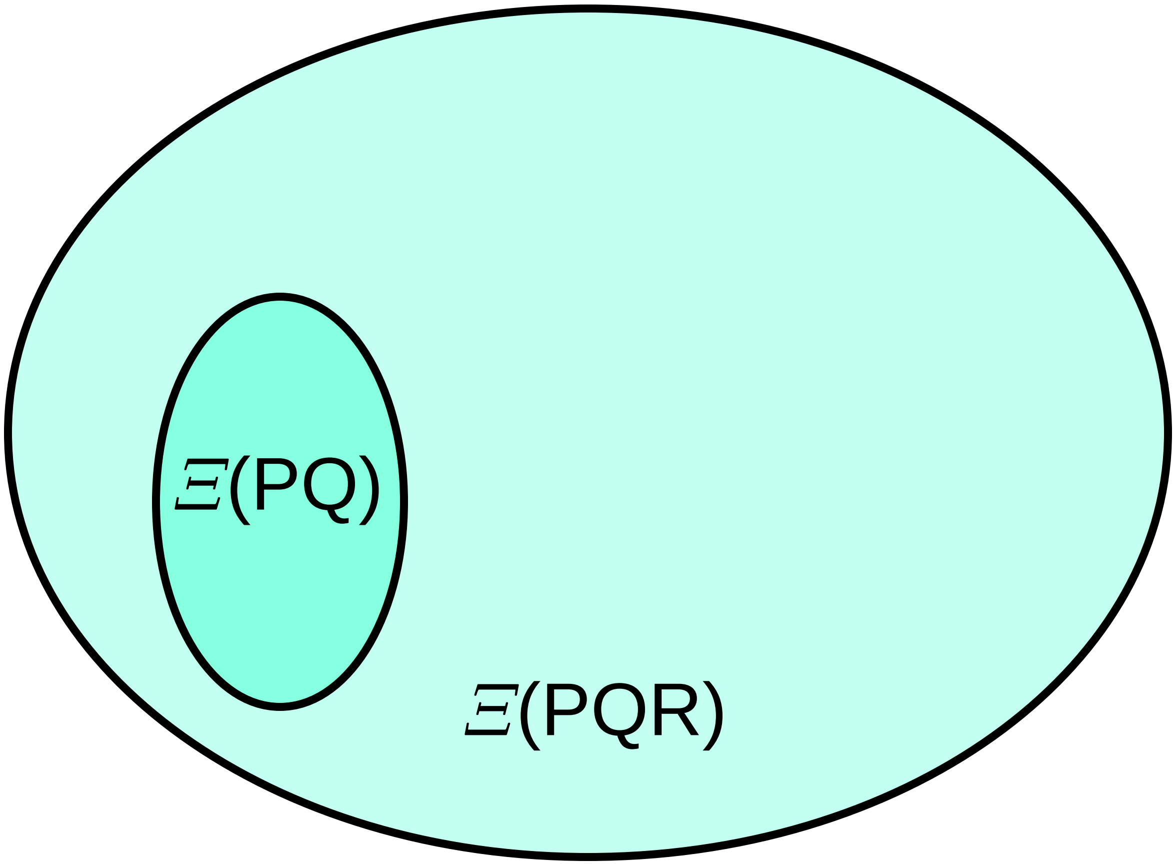

Proof

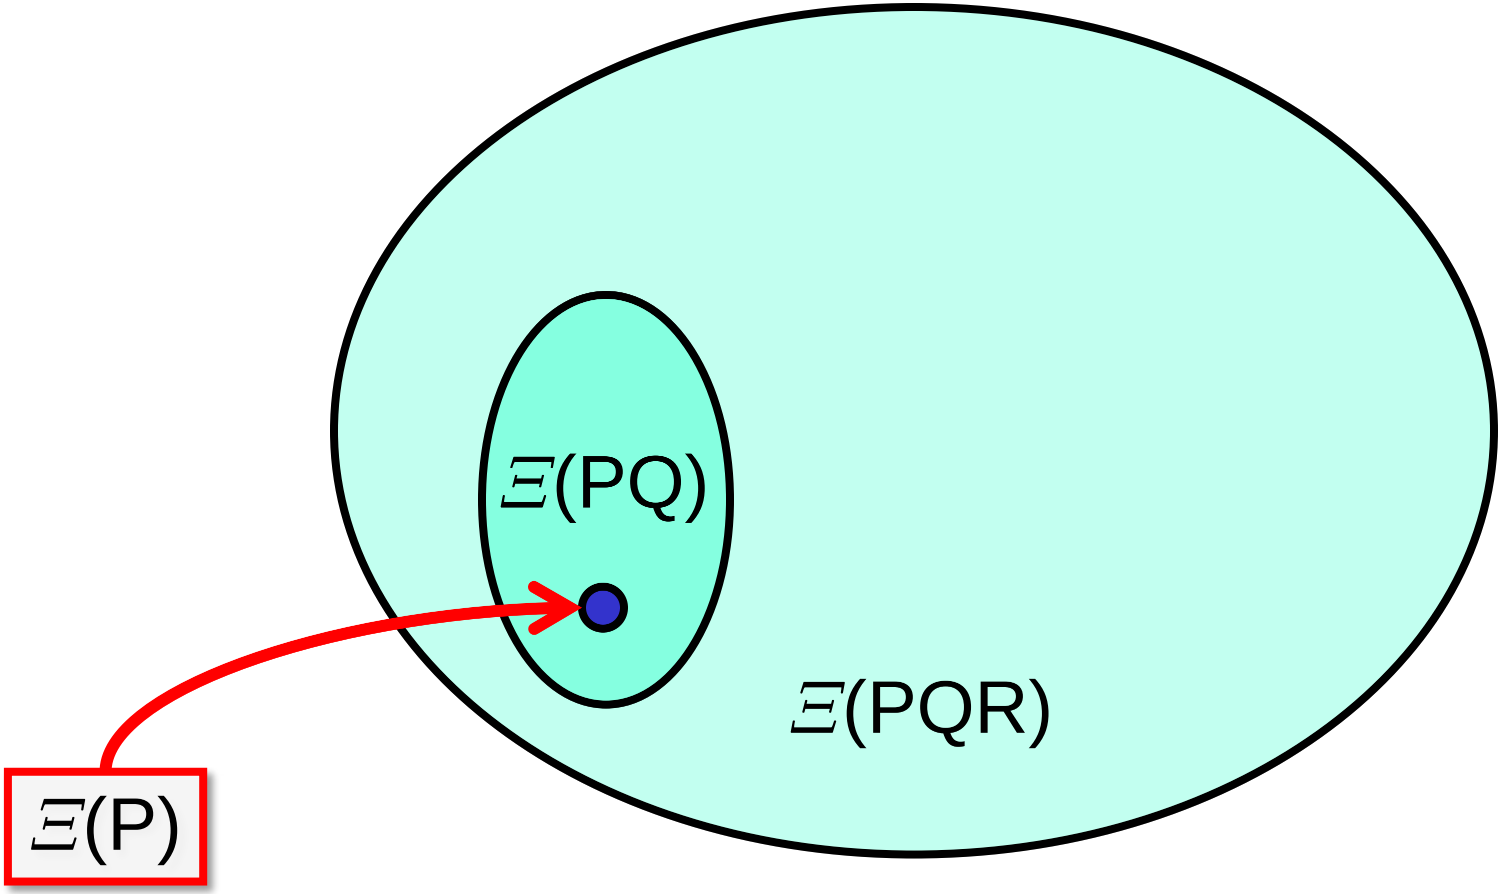

Proof

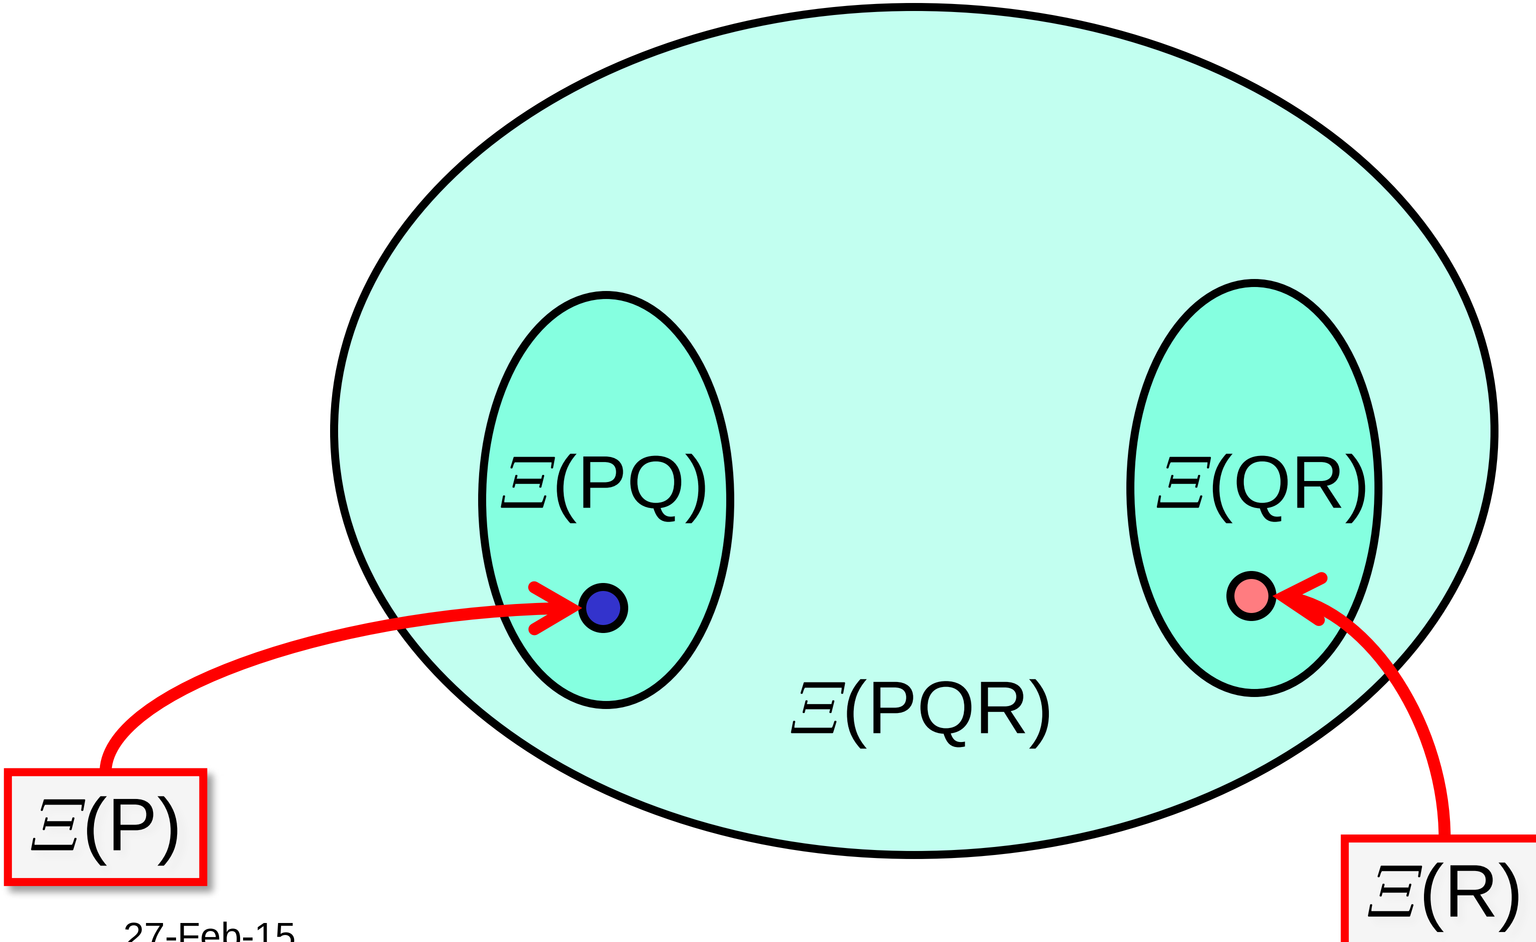

Complex is Path-Connected

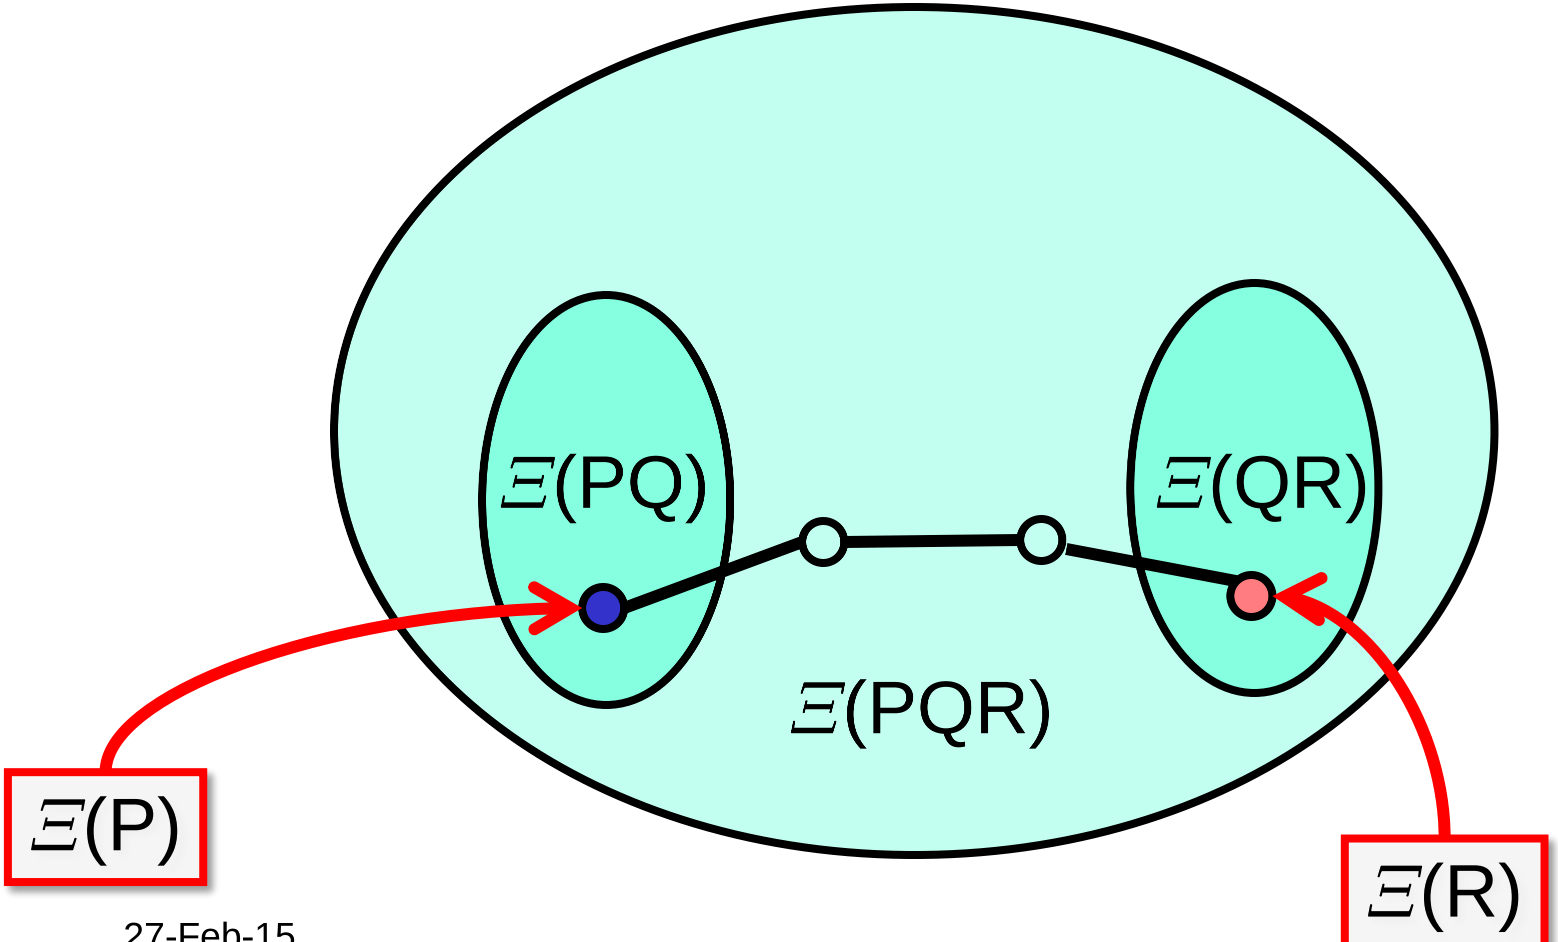

One-dimensional Sperner's Lemma

Some edge has two "colors"
An execution that decides two distinct values

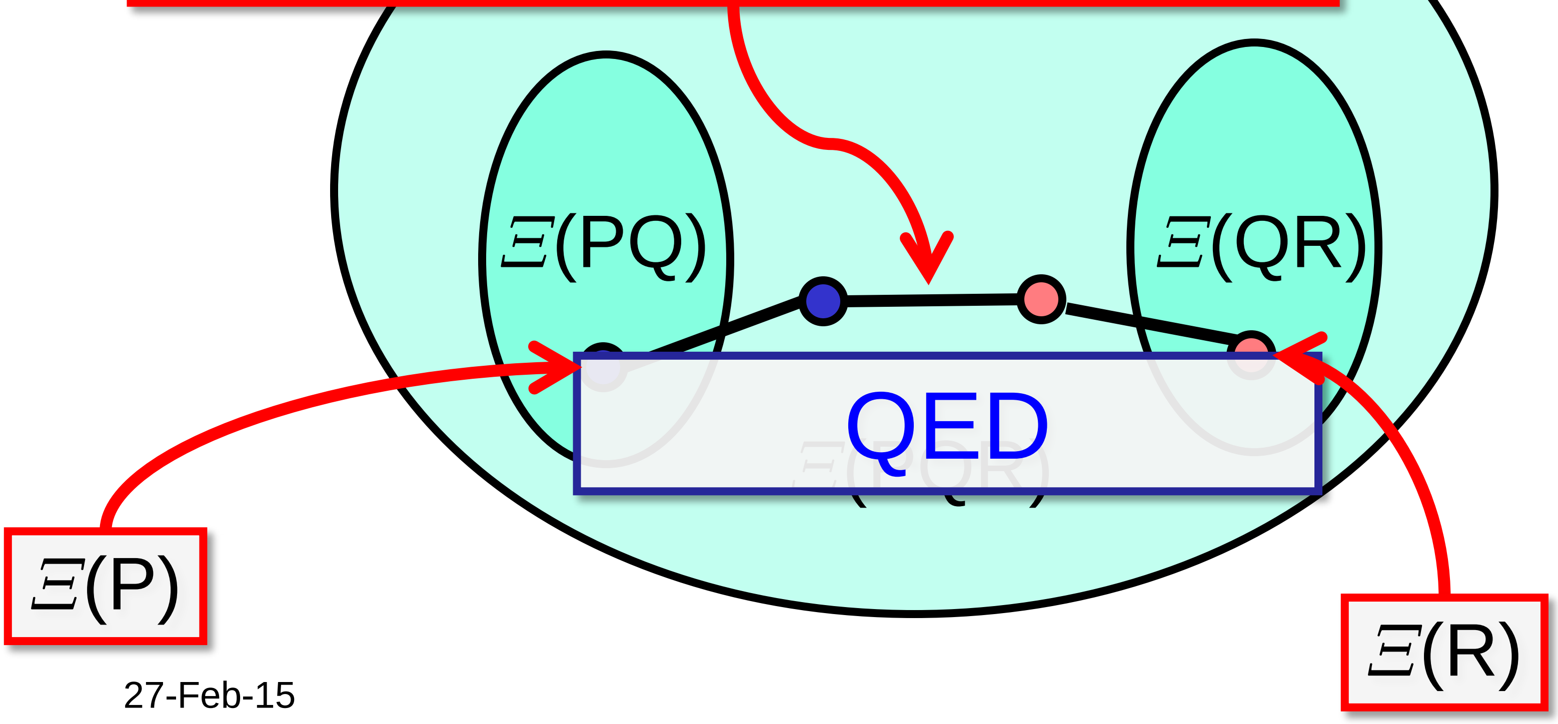

Road Map

Consensus Impossibility

General theorem

Application to read-write models

k-set agreement Impossibility

General theorem

Application to read-write models

Application

We now show that consensus is impossible in wait-free read-write memory

For every *n*-simplex σ^n,
$\Xi(\sigma^n)$ is path-connected …

For every *(n-1)*-simplex $\sigma^{n\text{-}1}$,
$\Xi(\sigma^{n\text{-}1})$ is non-empty.

Reachable Complexes

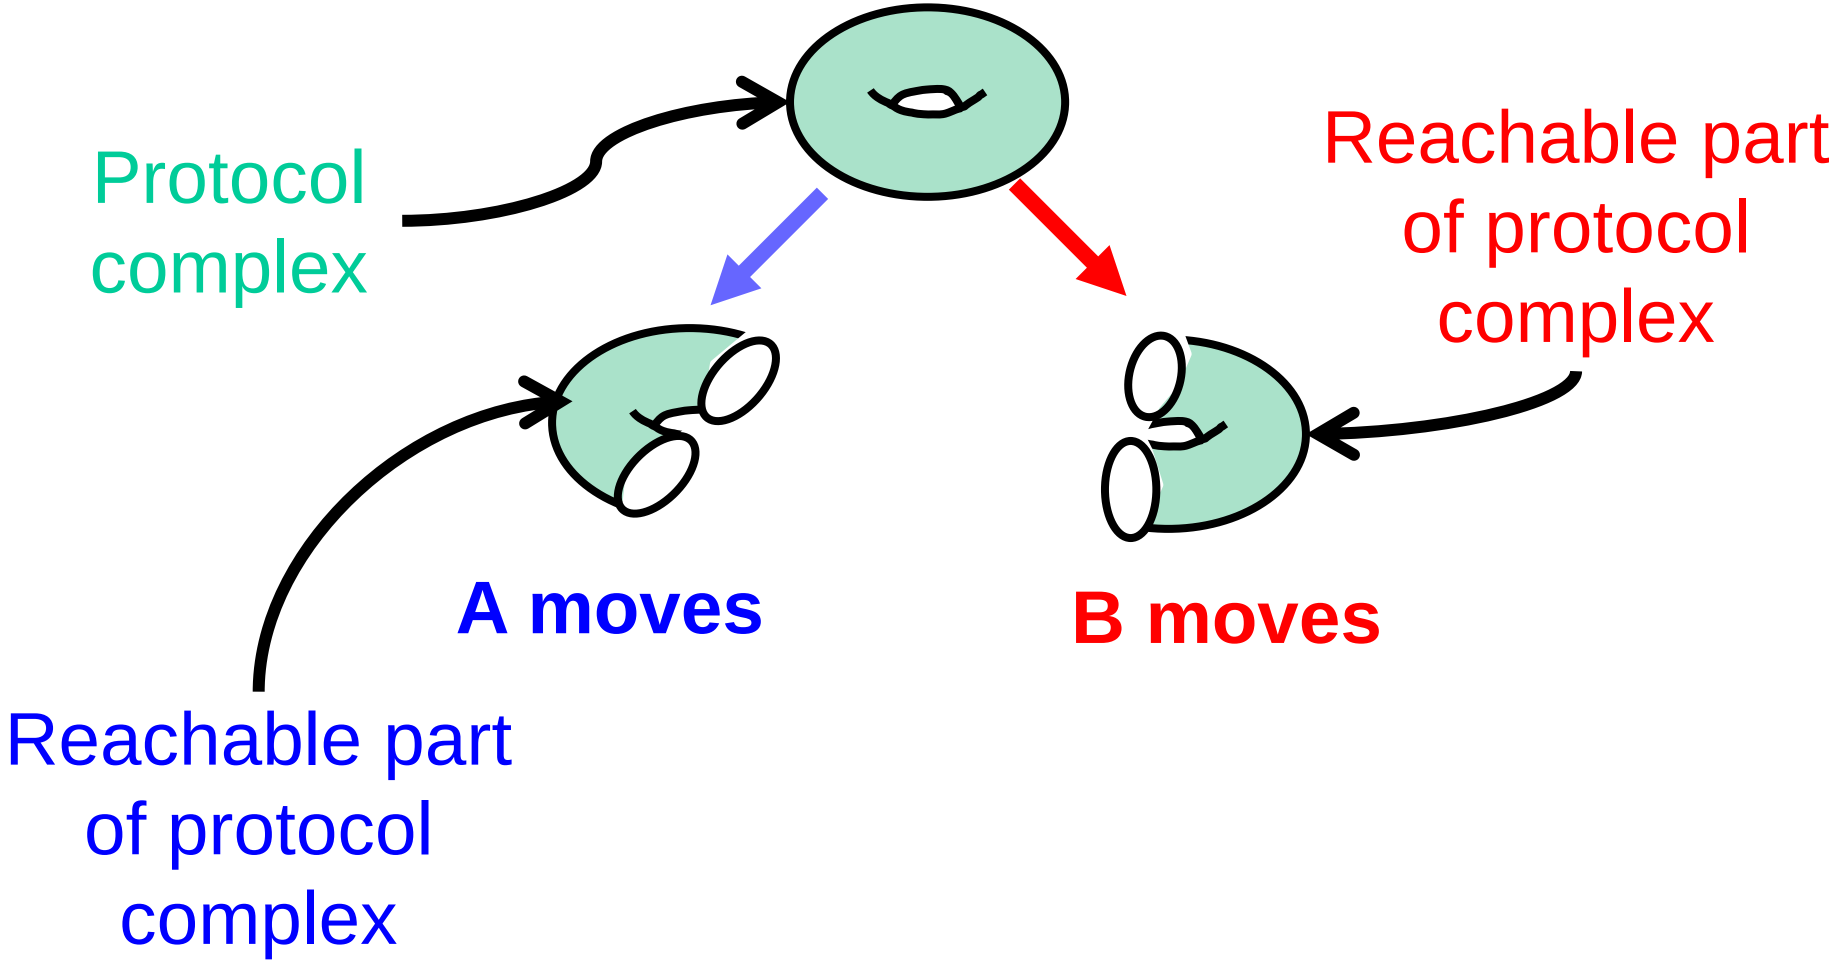

Reachable Complexes

At the end, single simplex reachable

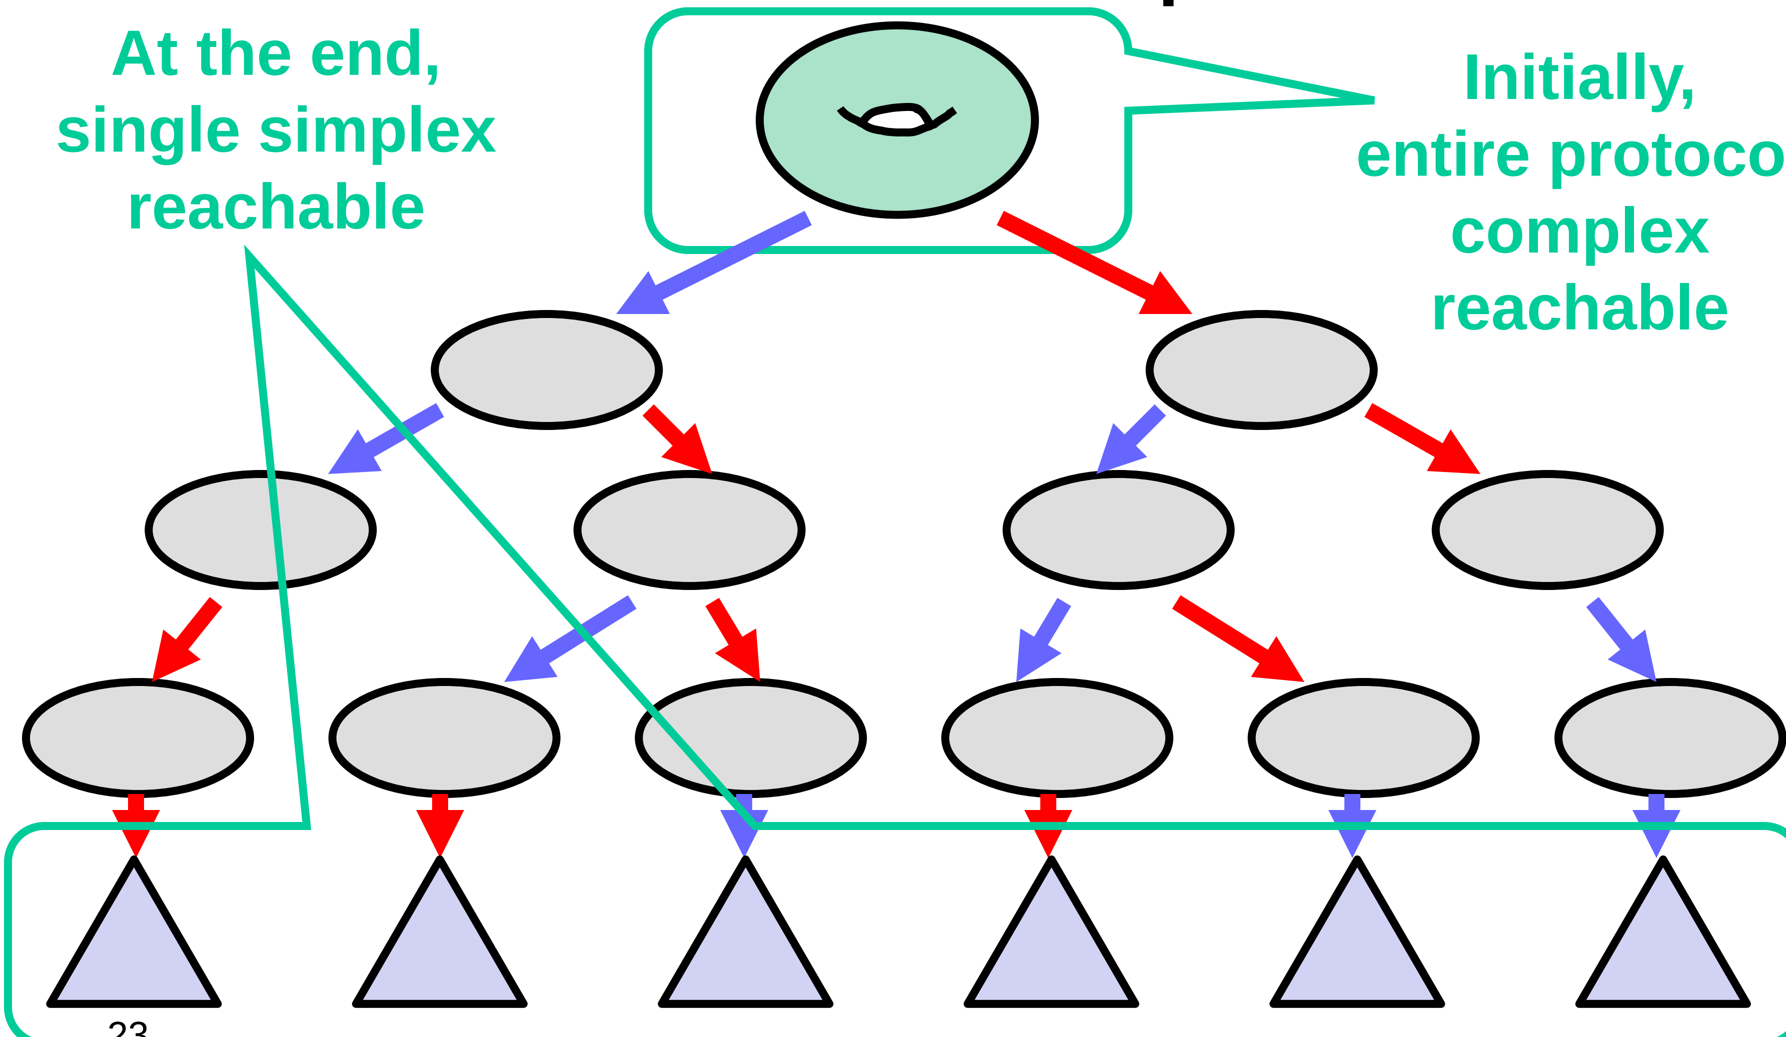

Initially, entire protocol complex reachable

Eventual Property

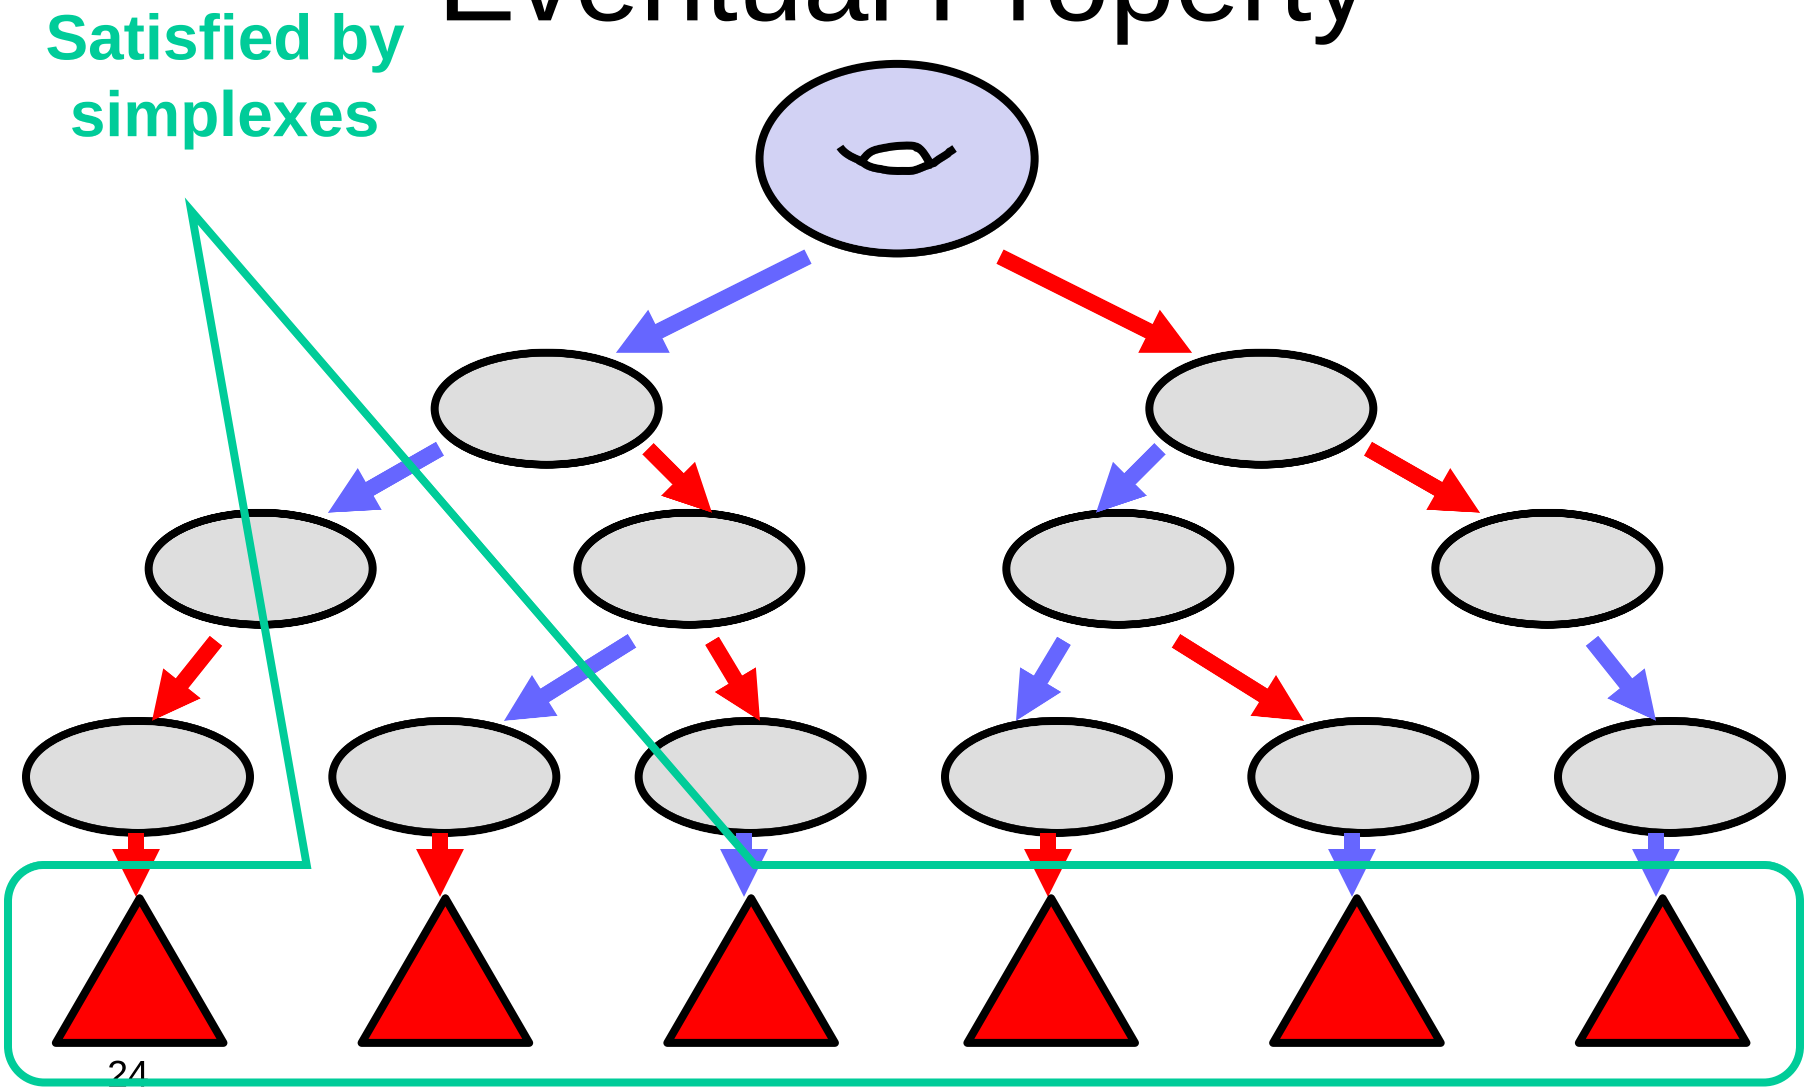

Eventual Property

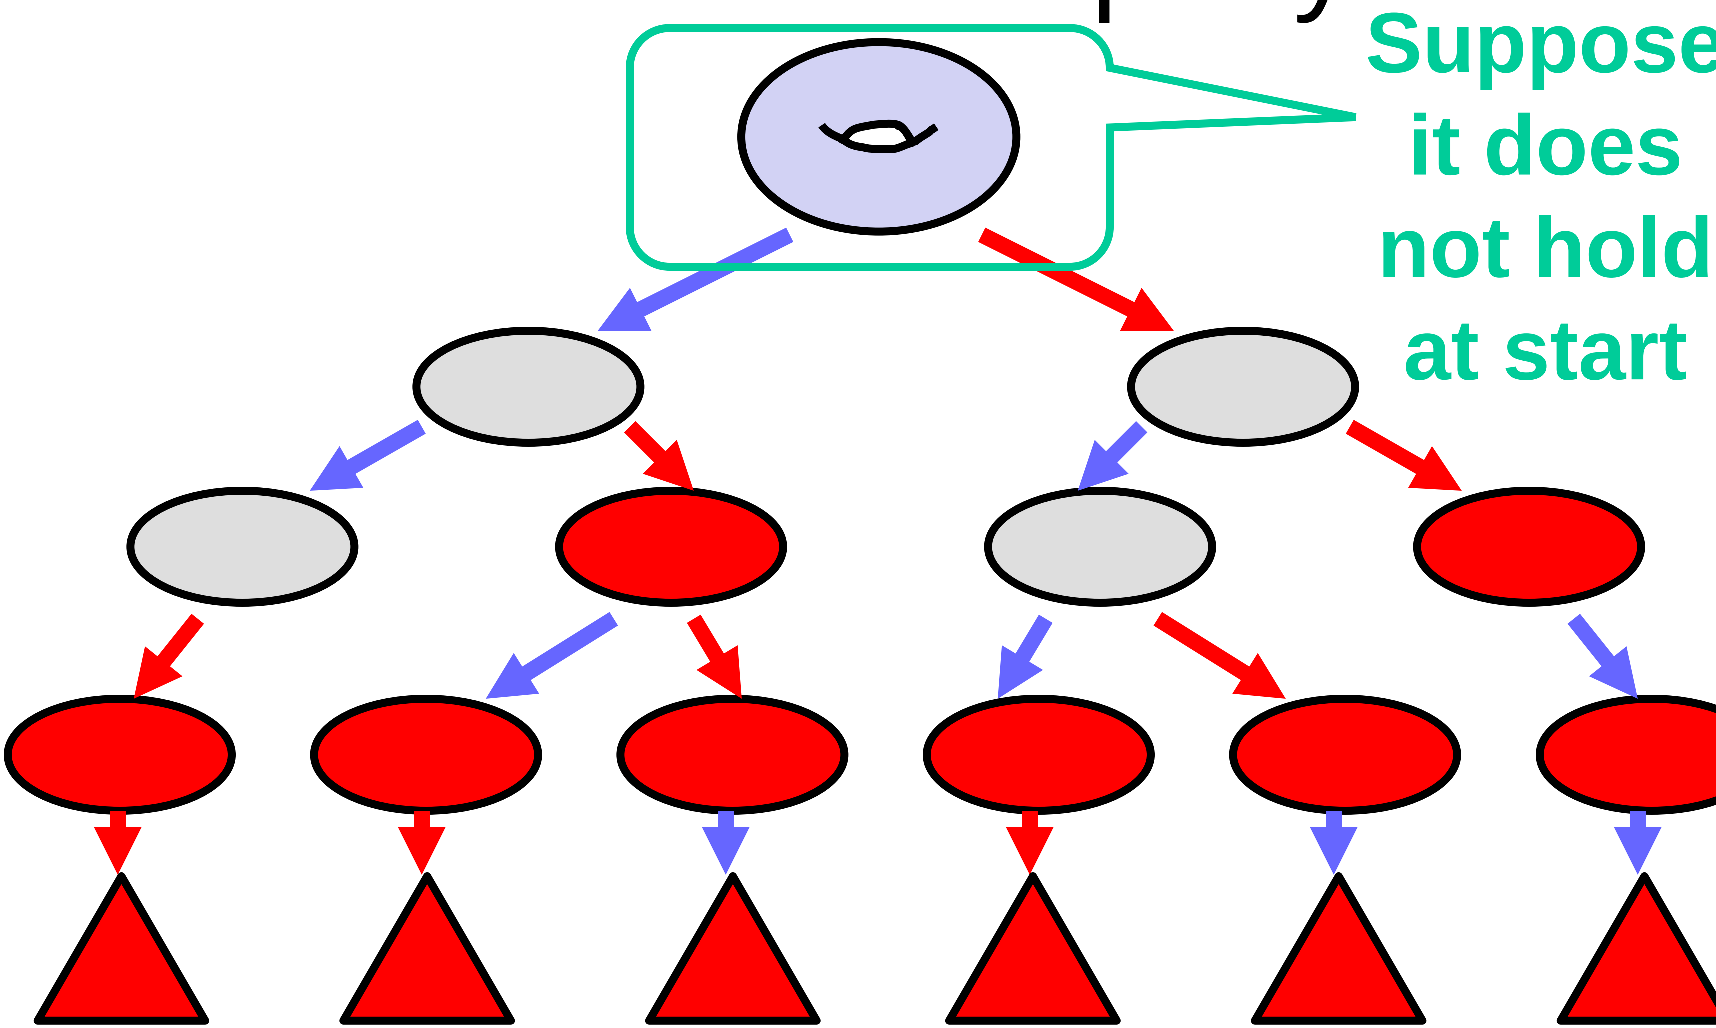

Critical States

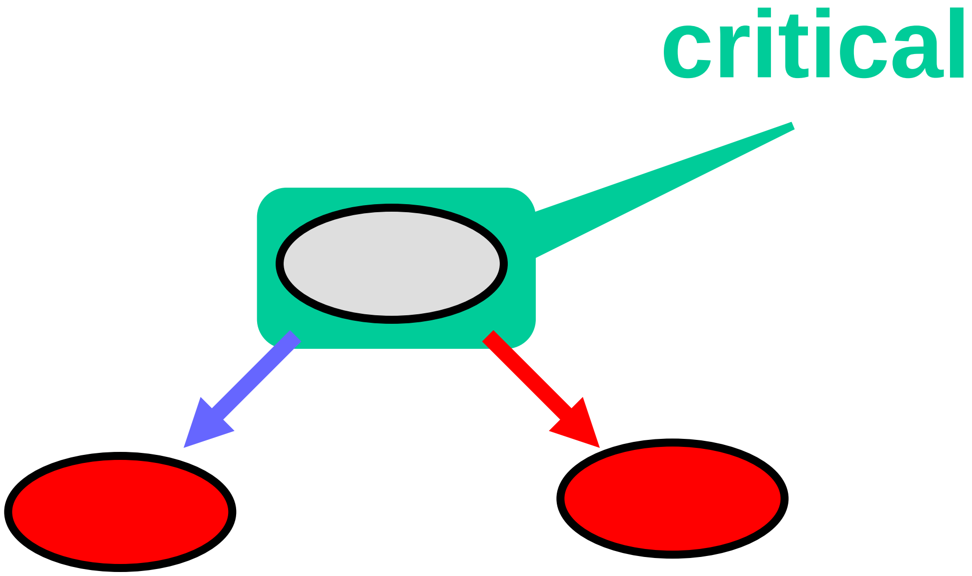

Critical States Exist

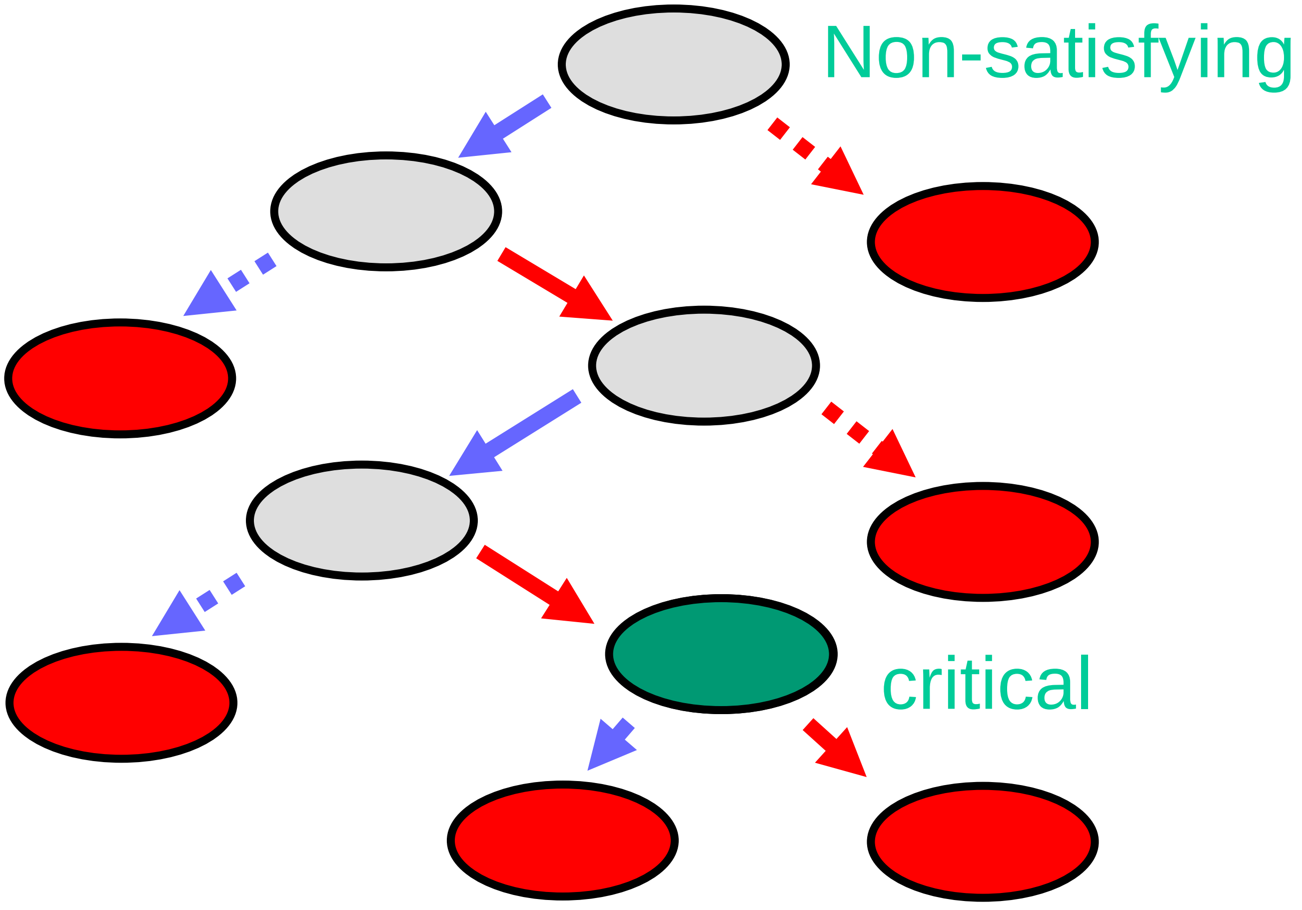

Eventual Property

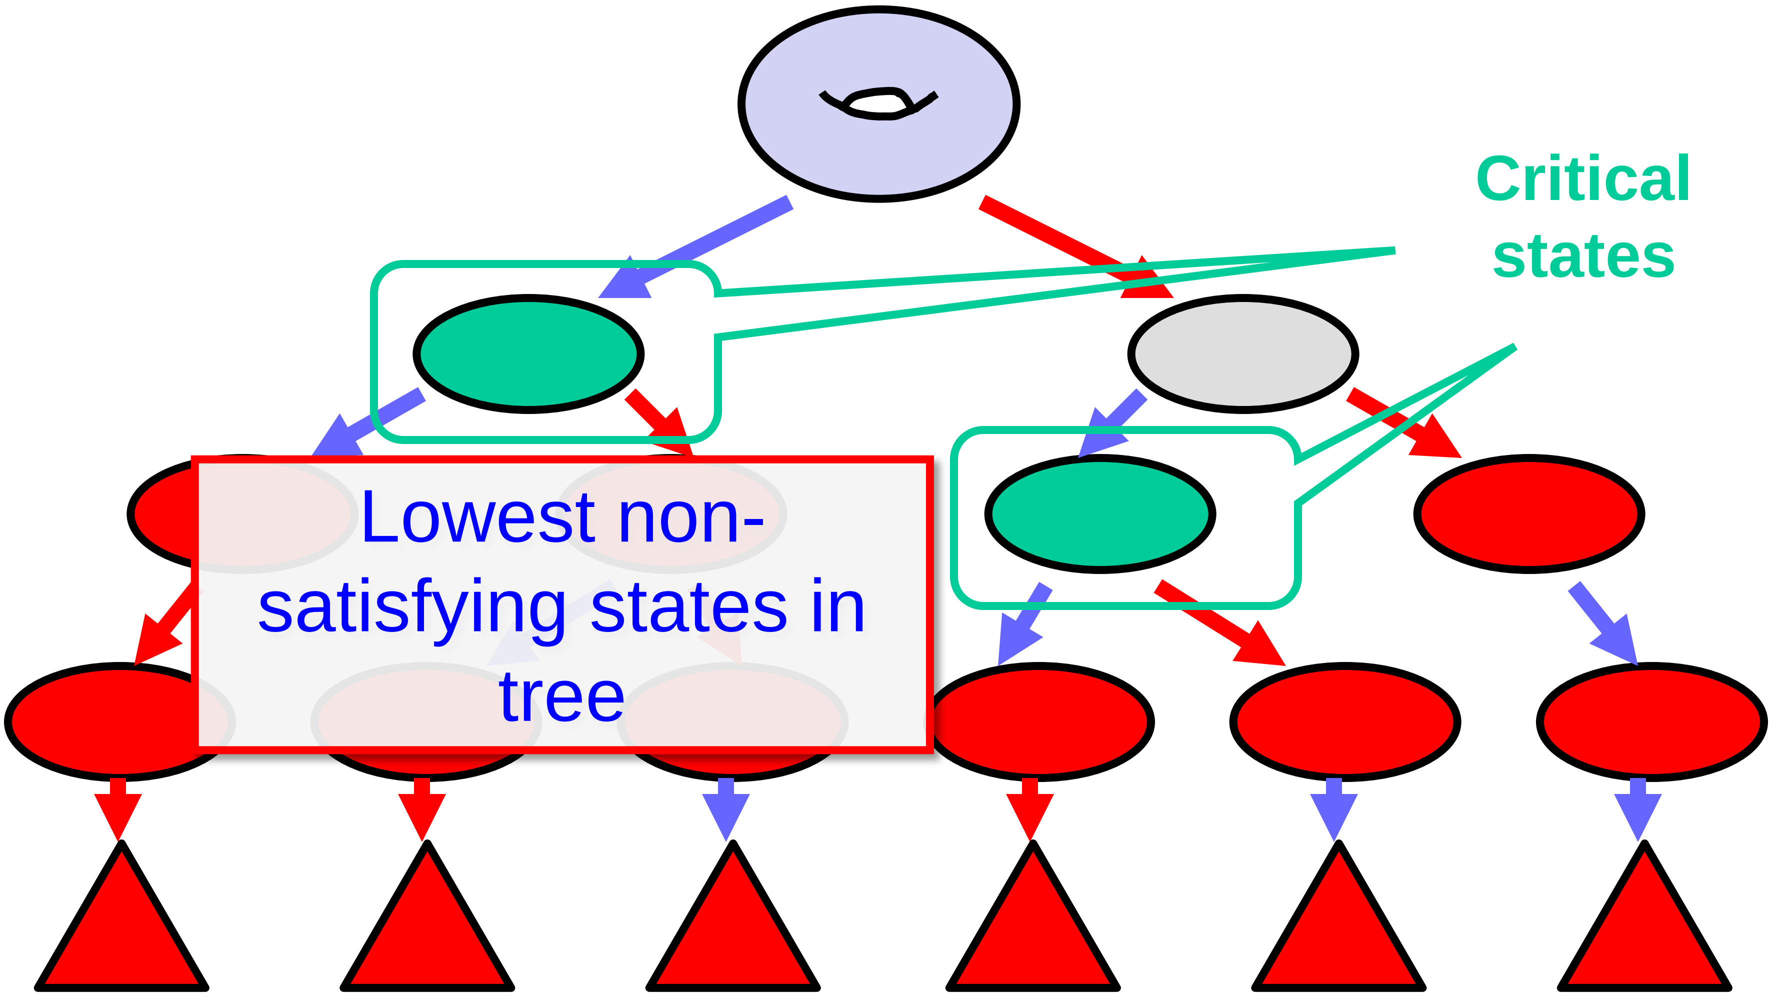

Path-Connectivity is an Eventual Property

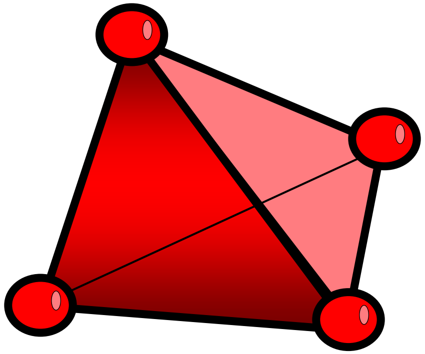

Individual simplex is path-connected

Path-Connectivity has a Critical State

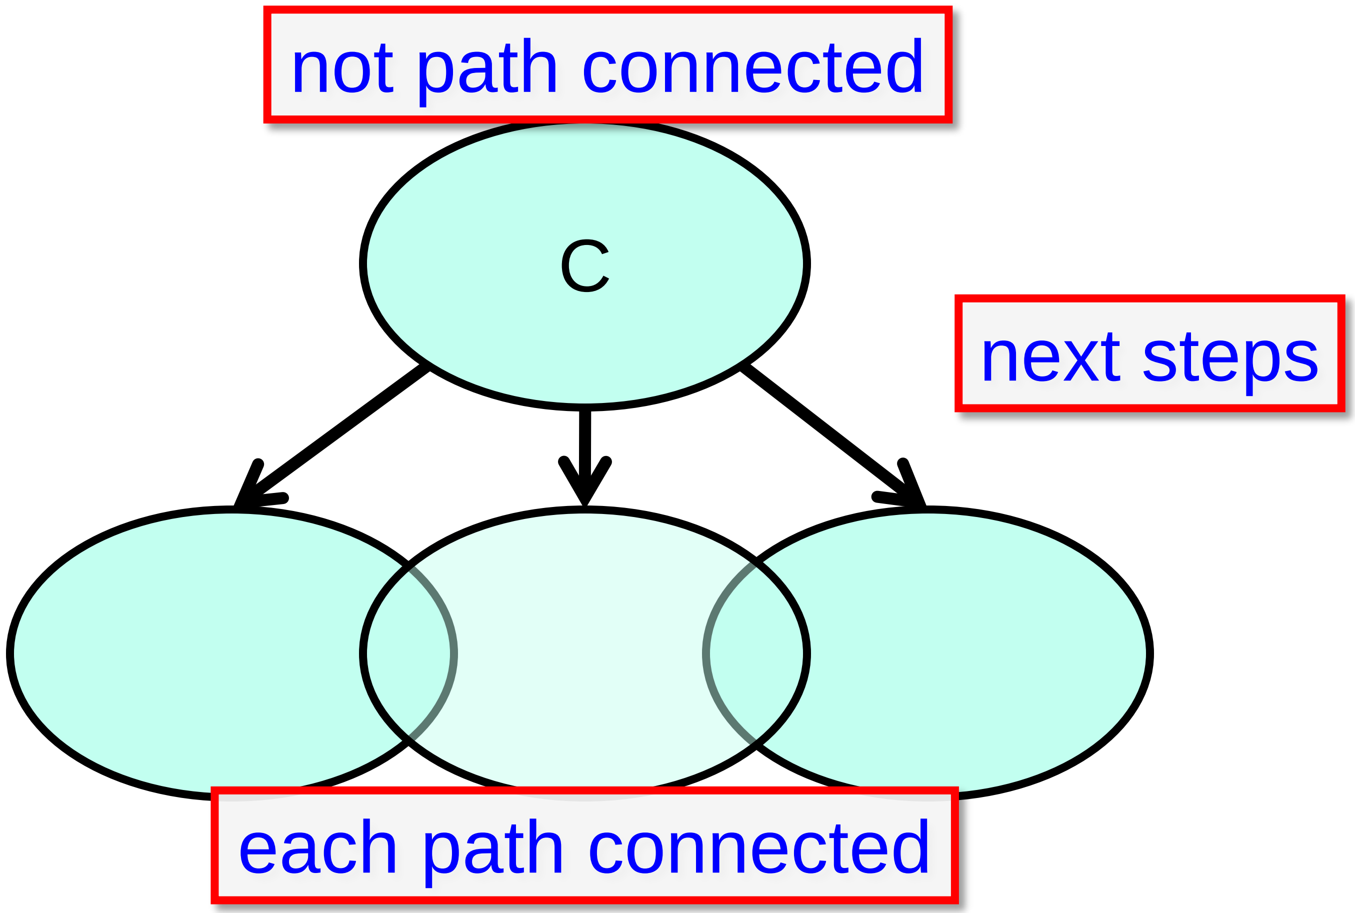

Critical State in Layered IS

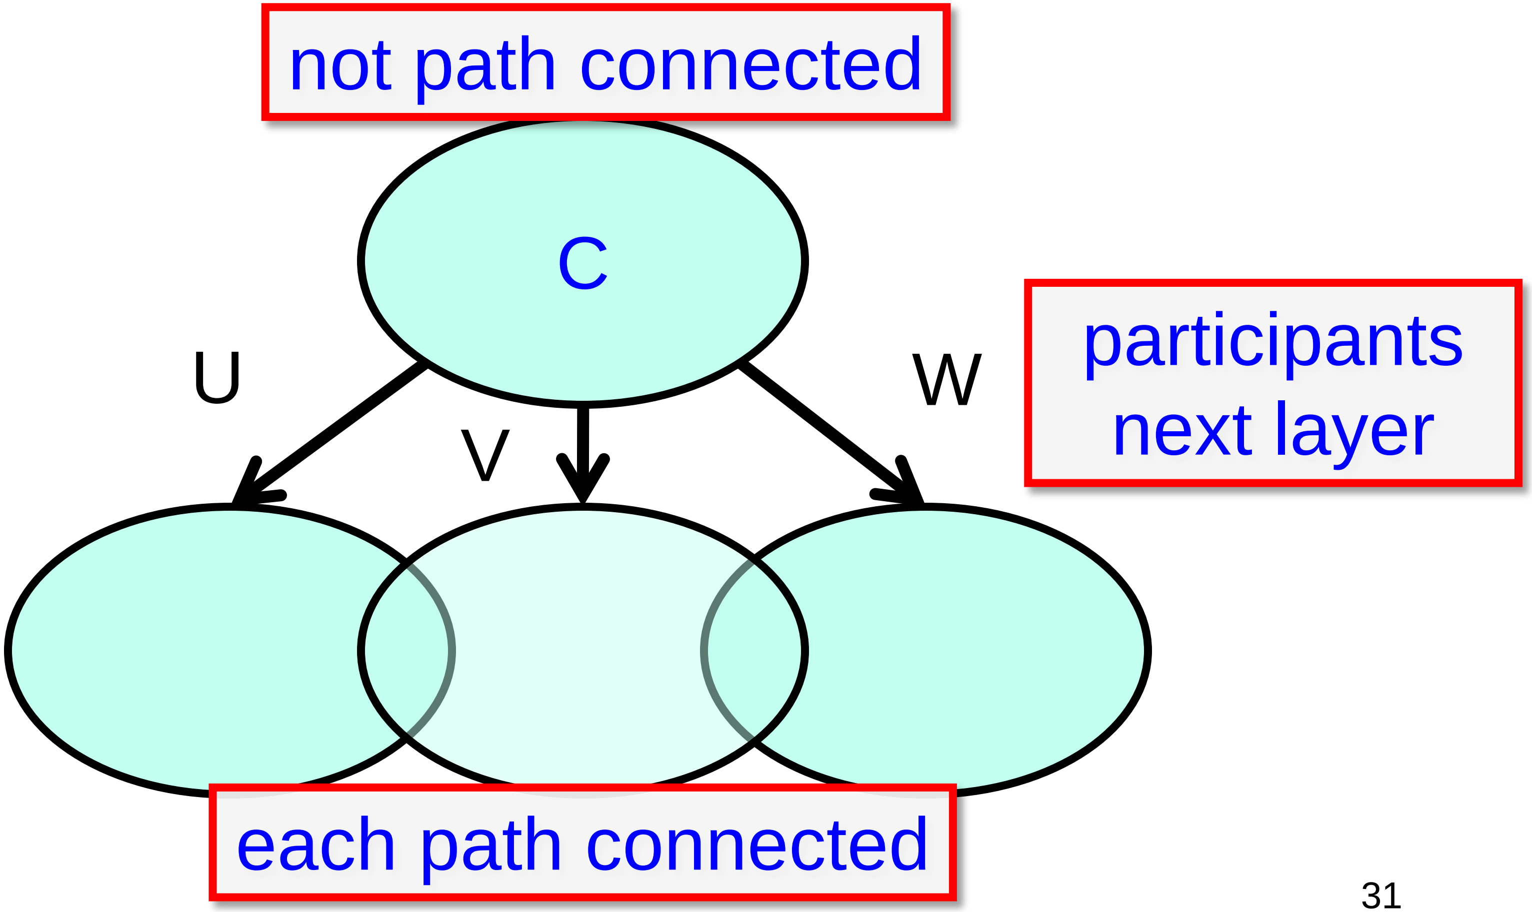

Notation

$\boxed{C \uparrow U}$

Configuration reached by running processes in U in next layer

Critical State in Layered IS

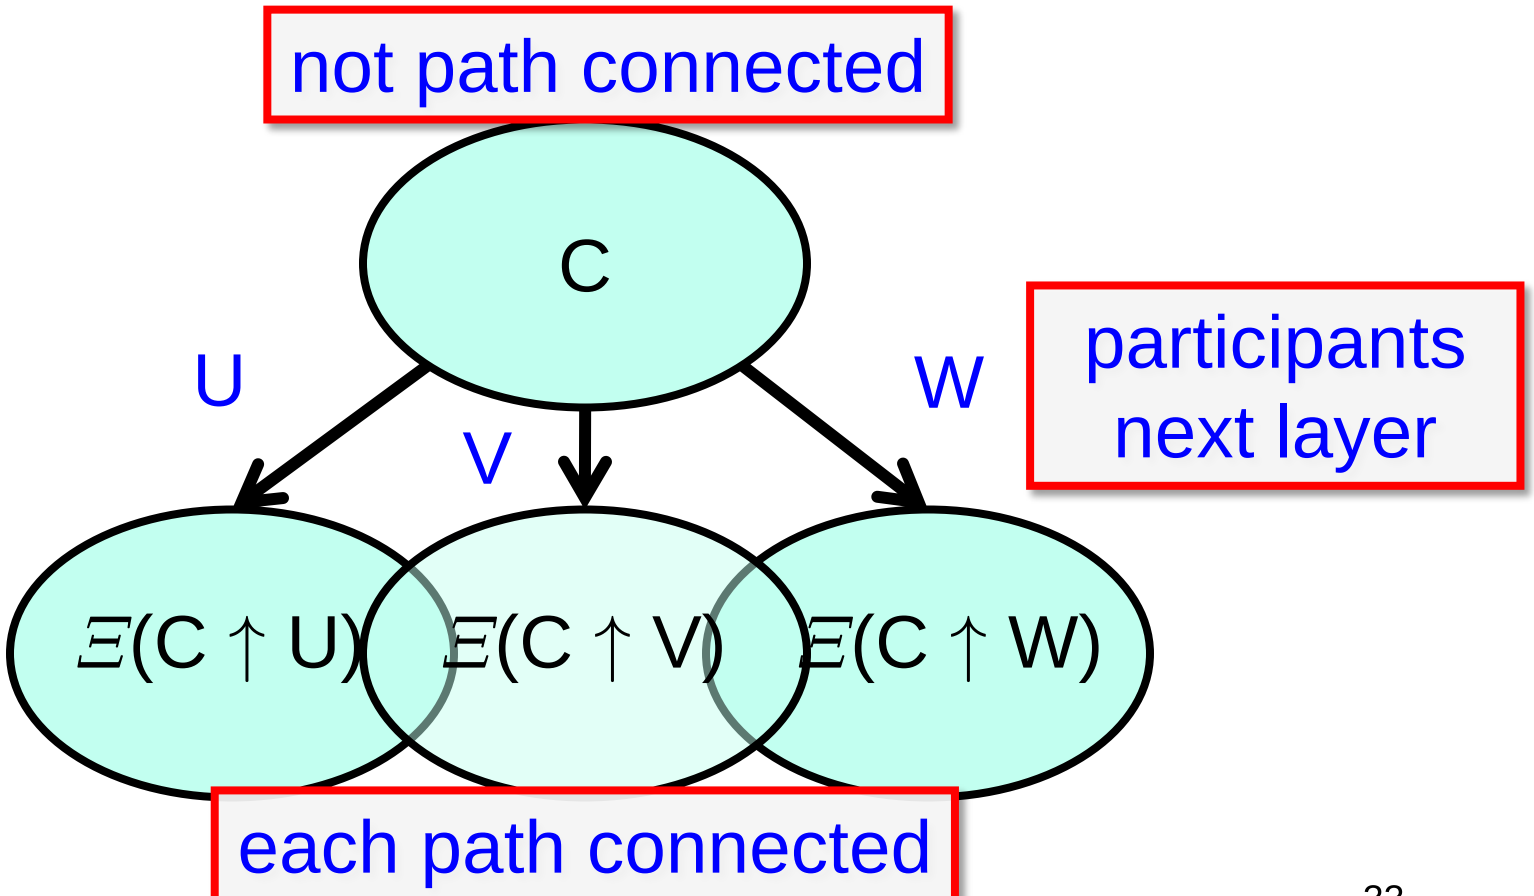

Proof Strategy

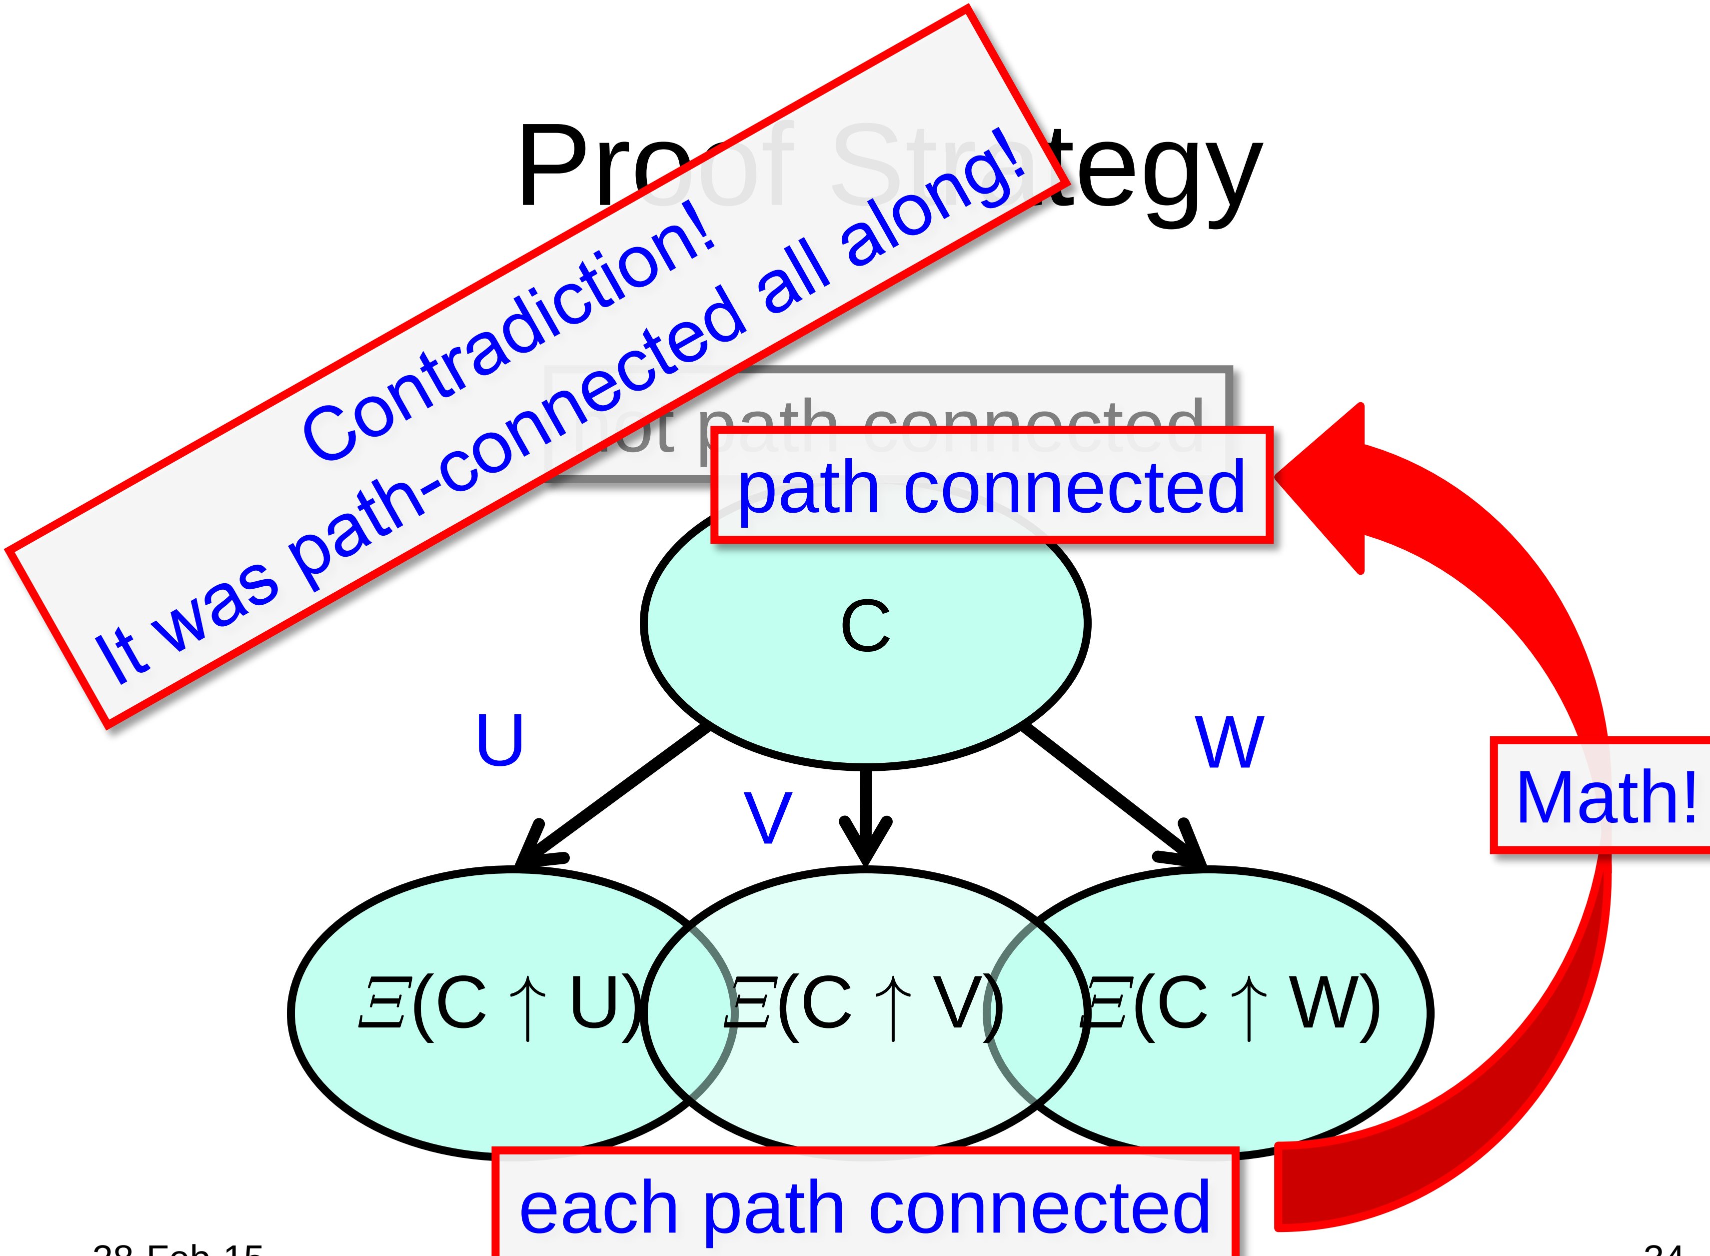

One-Dimensional Nerve Lemma

Reason about path-connectivity of a graph …

From path-connectivity of components …

And how they fit together.

Graph $\mathcal{K}$

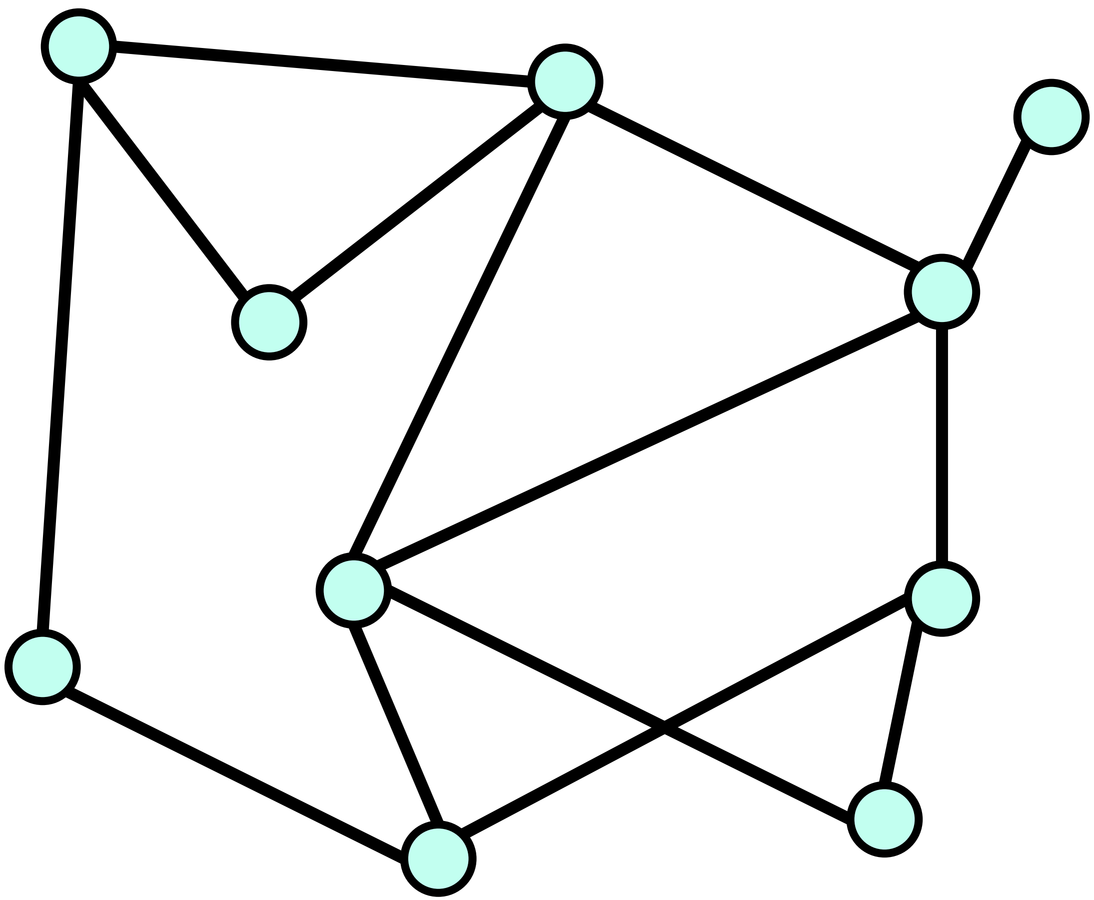

Covering $\mathcal{K}_i$

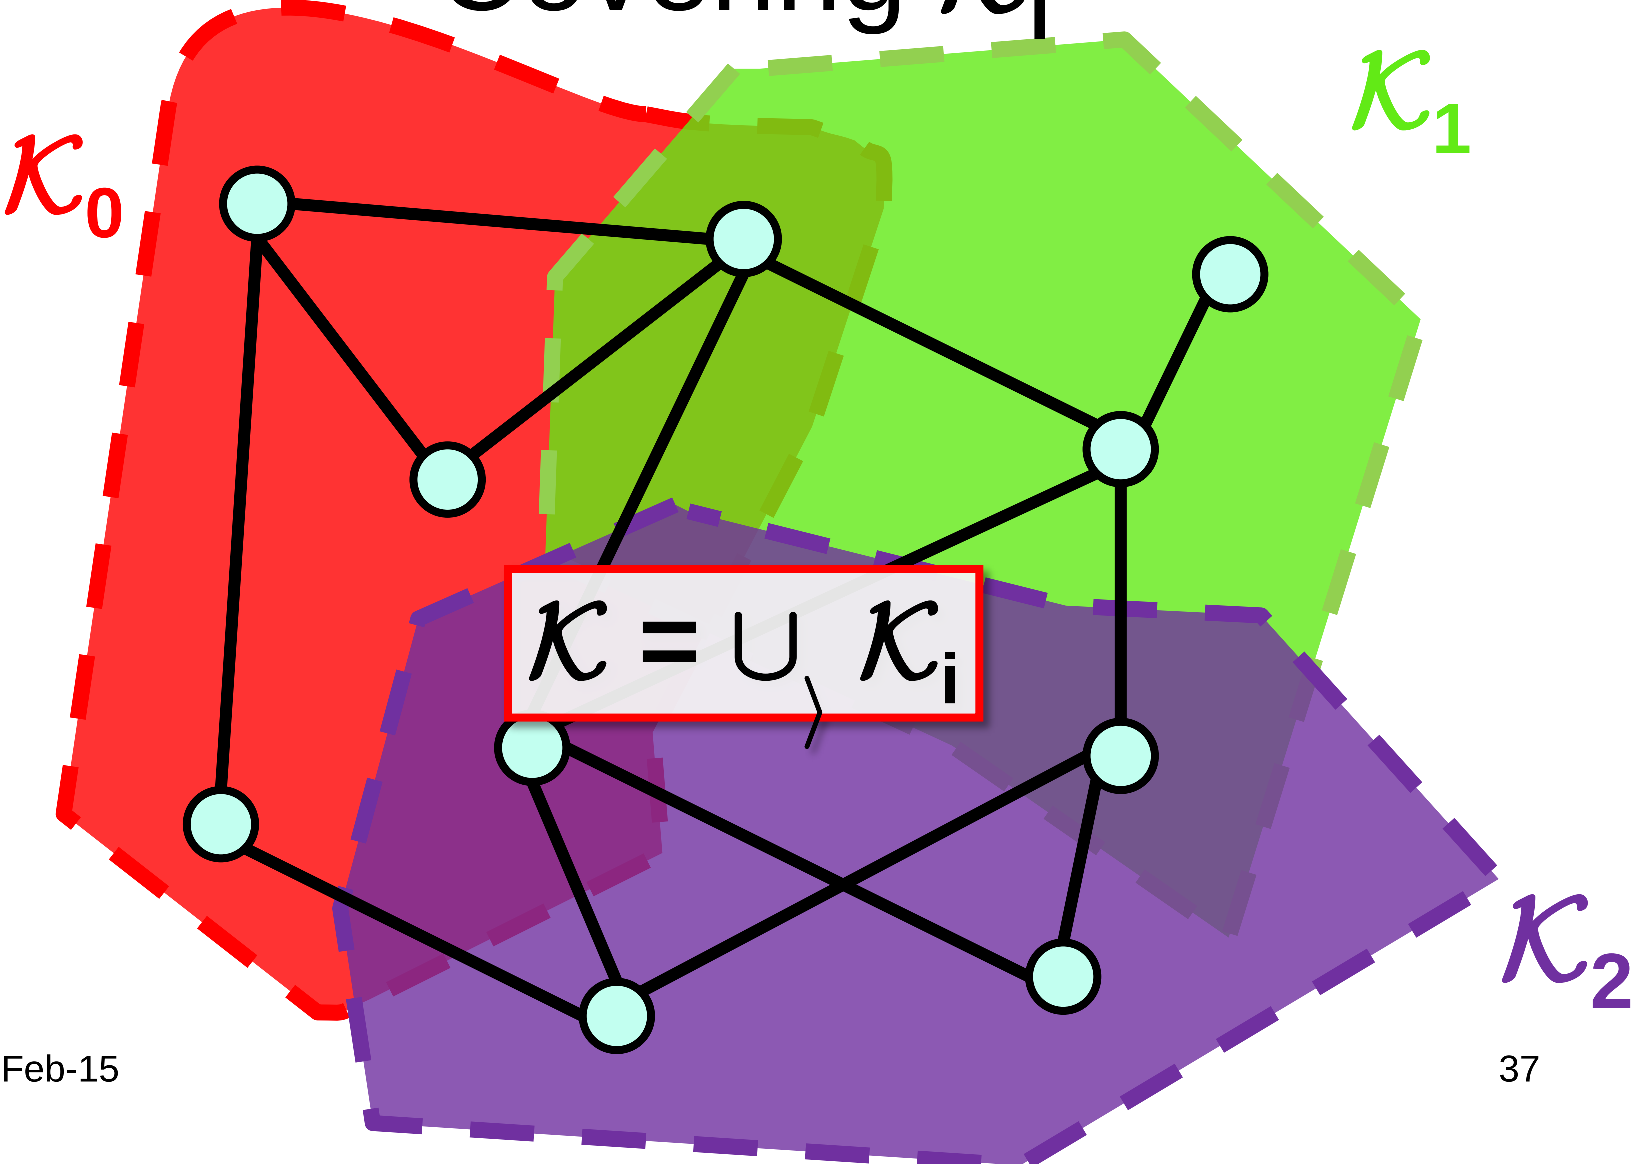

Nerve Graph Vertices

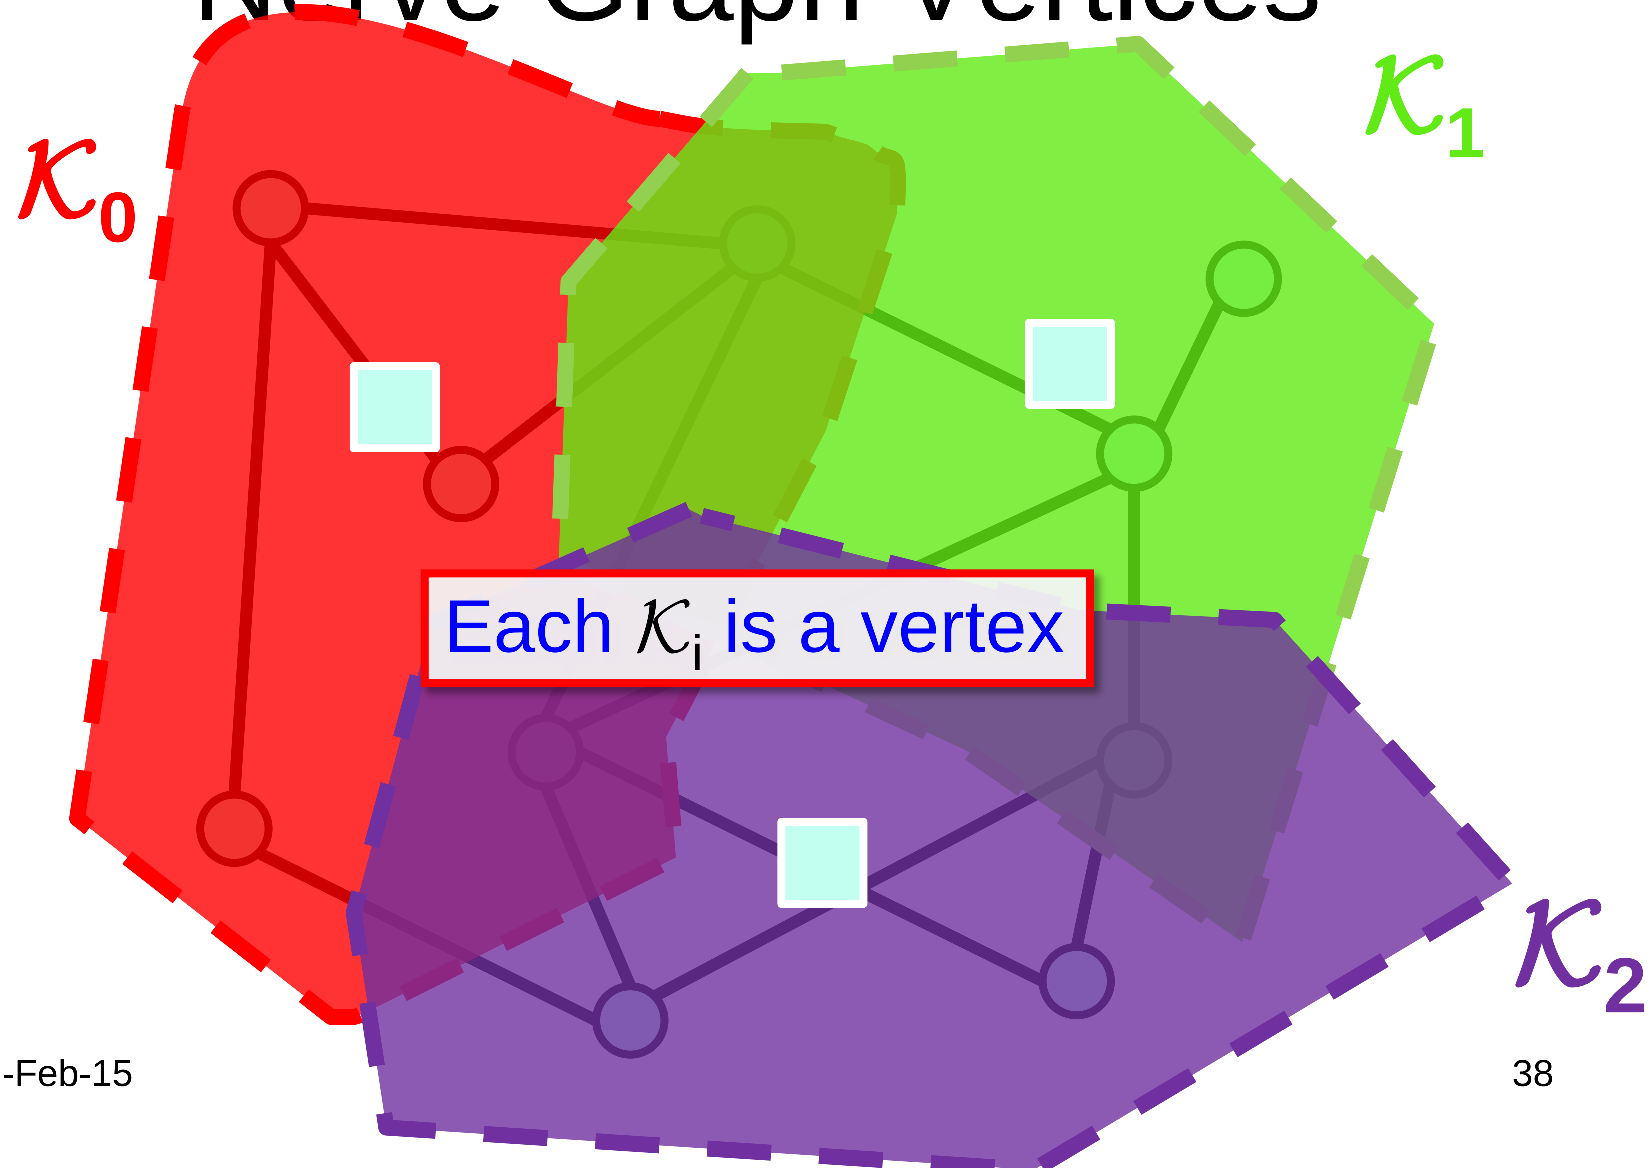

Nerve Graph Edges

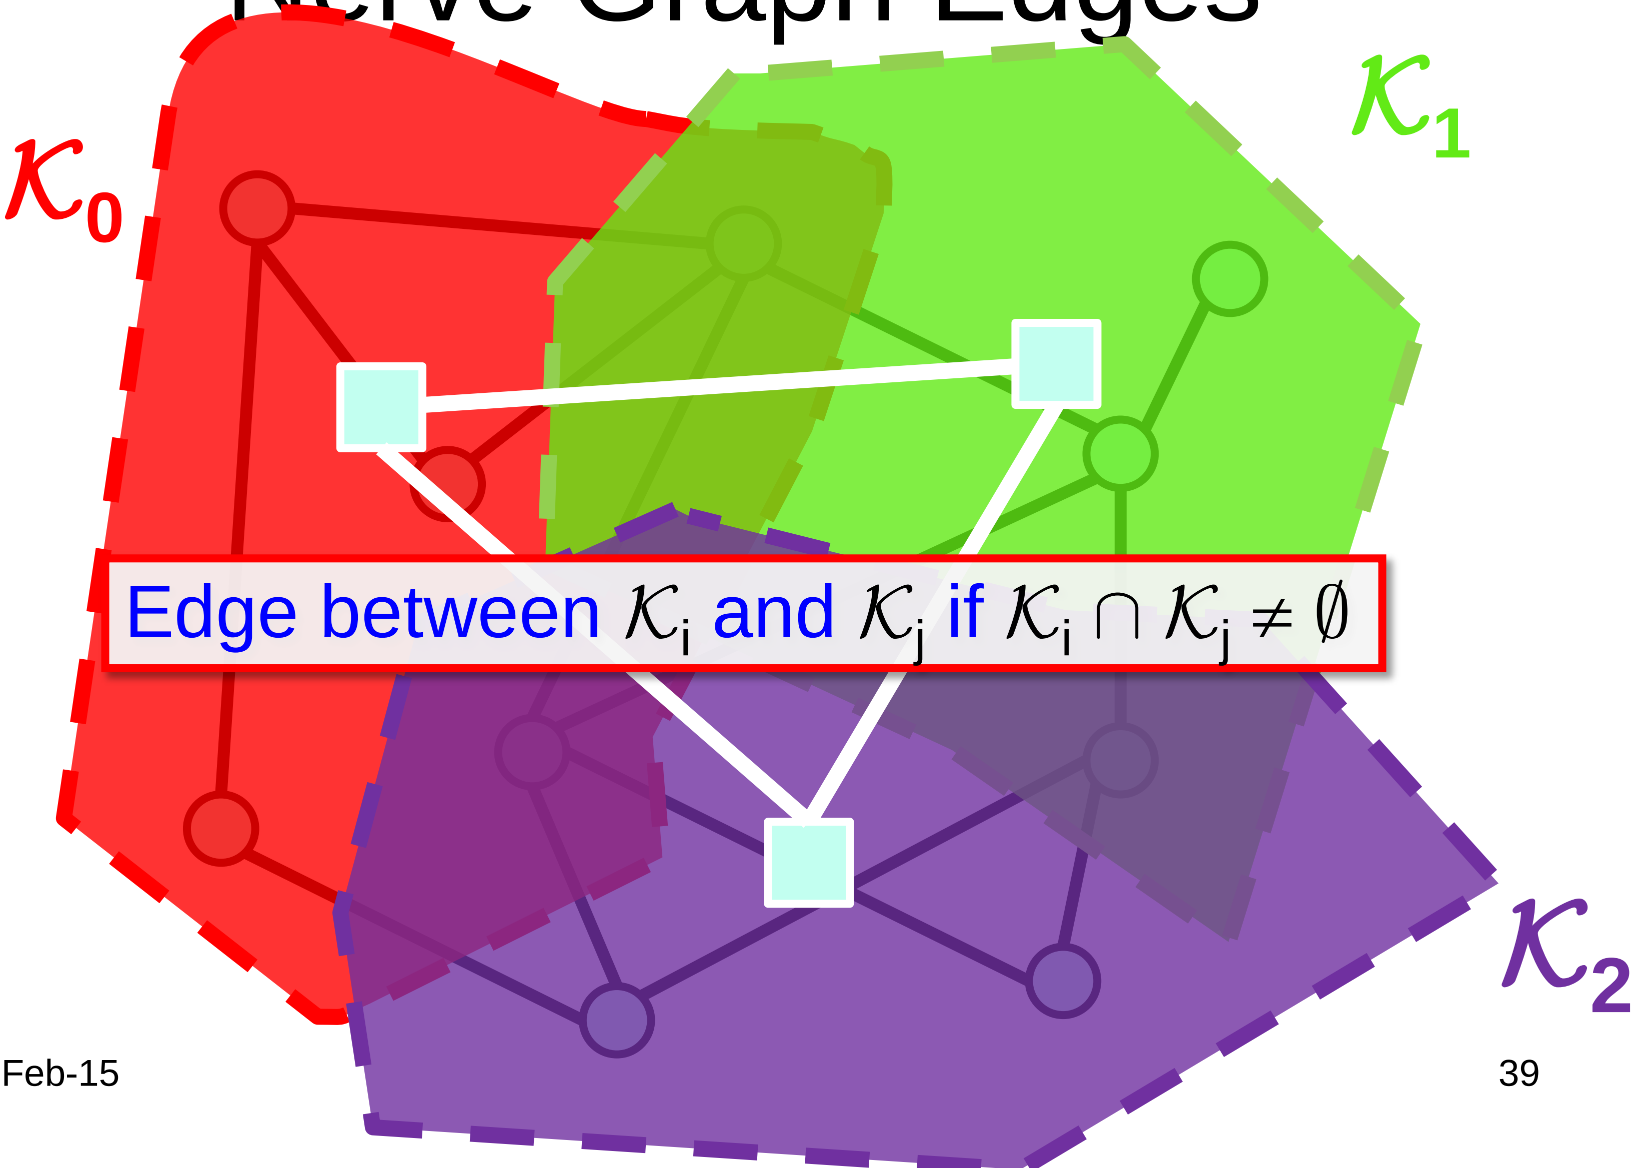

One-Dimensional Nerve Lemma

If each $\mathcal{K}_i$ is path-connected …

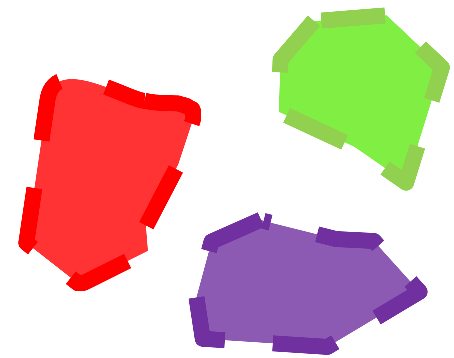

and the nerve graph $\mathcal{N}_i(\mathcal{K}_i)$ is path-connected …

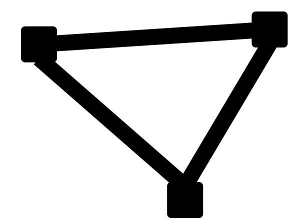

then $\mathcal{K}_i$ is path-connected …

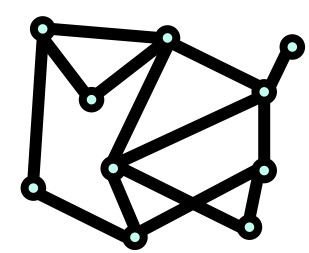

Critical State in Layered IS

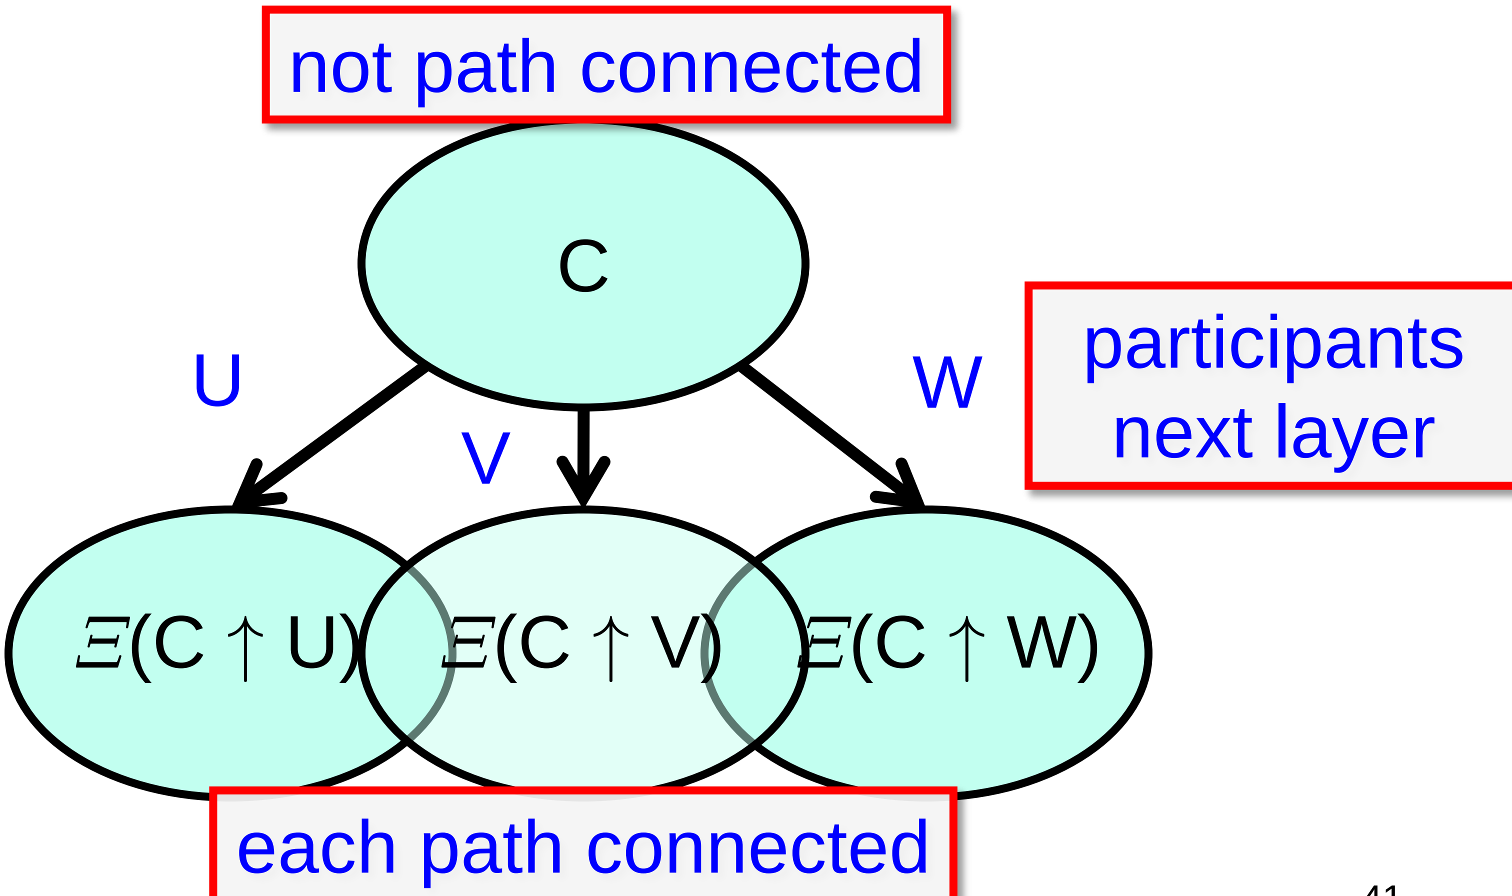

Configurations

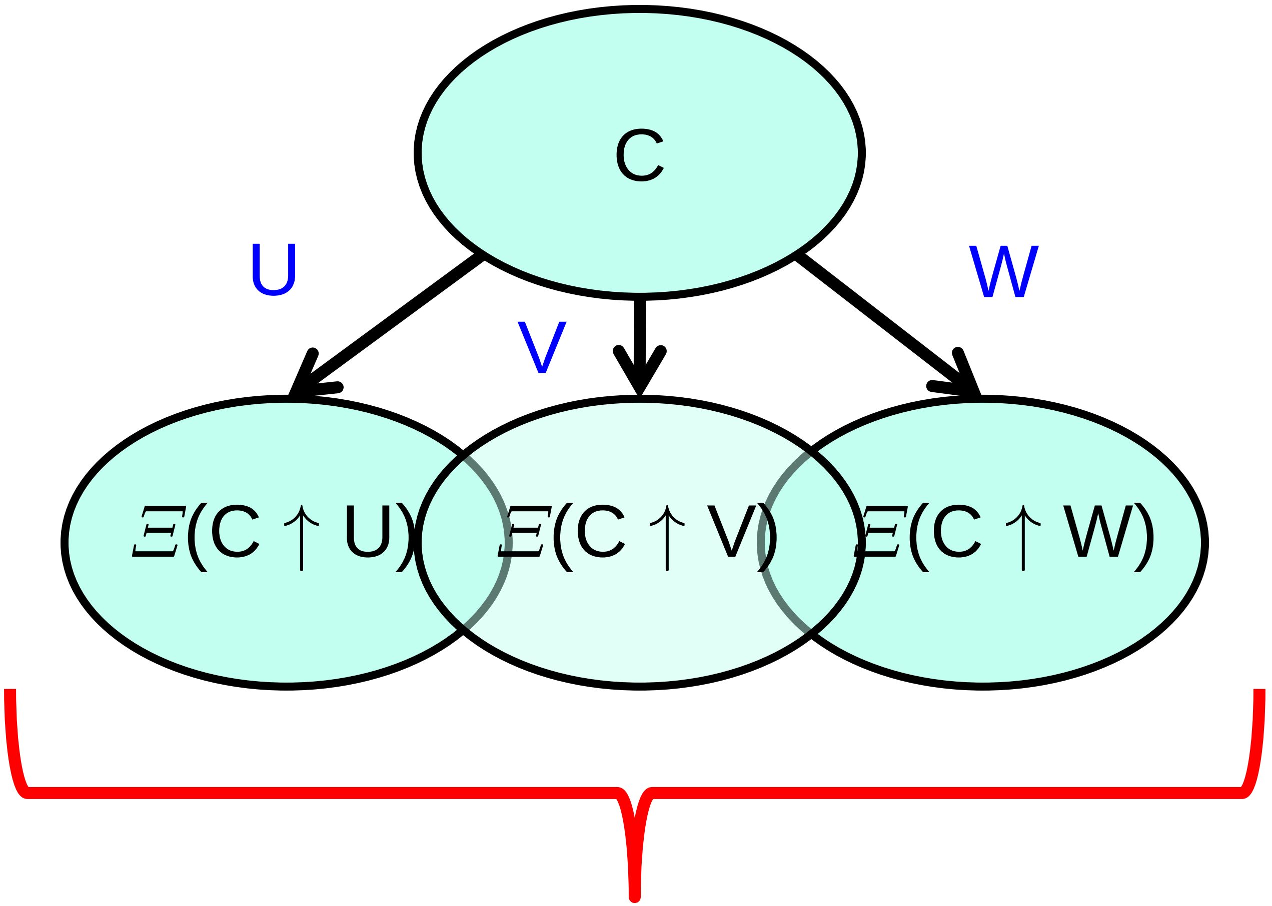

The $\Xi(C \uparrow U)$ form a *cover* for reachable complex

Configurations

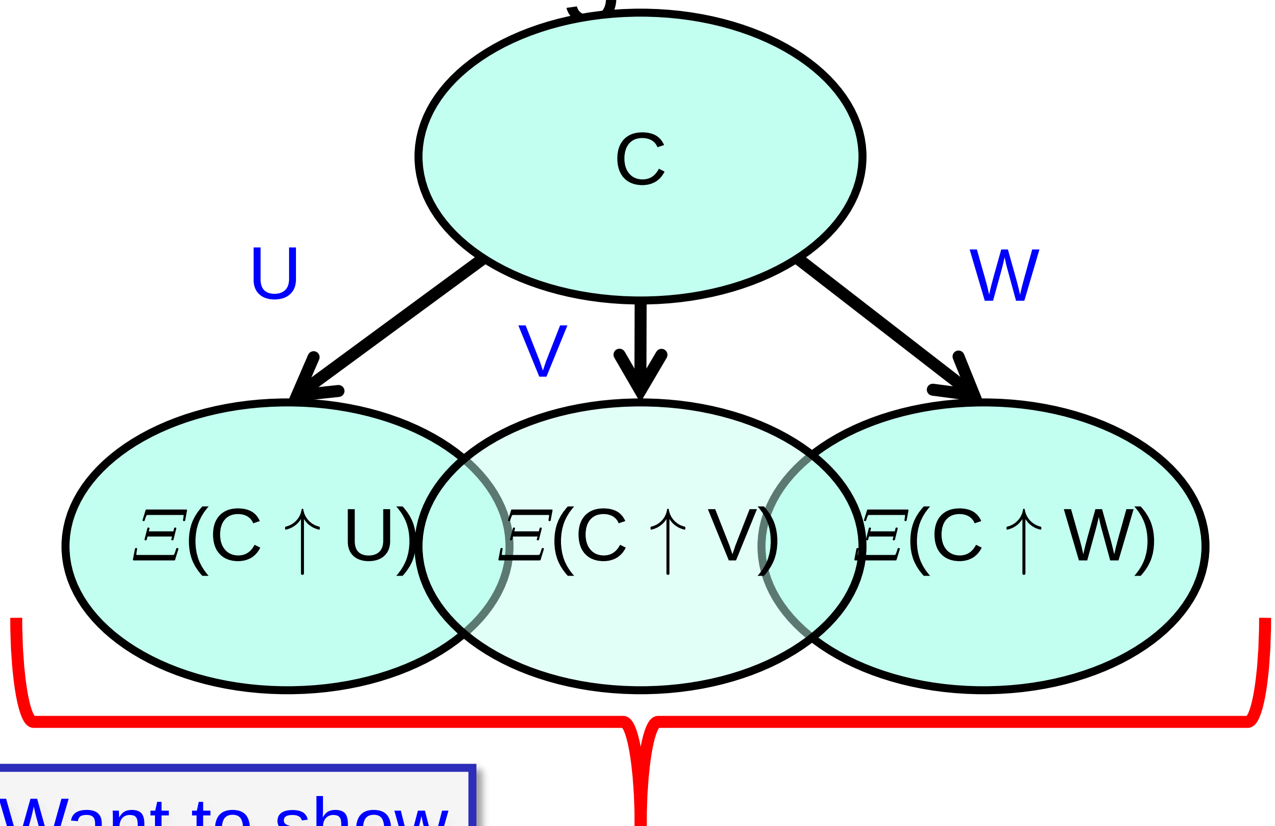

Want to show

each $\Xi(C \uparrow U)$ is path-connected

each $\Xi(C \uparrow U) \cap \Xi(C \uparrow V)$ is non-empty

Computing Intersections

U = {P,Q,R}

P	Q	R
write	write	write
snap	snap	snap

C ↑ U

V = {P,Q}

P	Q	R
write	write	
snap	snap	
		write
		snap

C ↑ V

(C ↑ V) ↑ U \ V

Memory state same

Only V can distinguish

Intersections (Case 1)

If $V \subseteq U$, then

$\Xi(C \uparrow U) \cap \Xi(C \uparrow V)$ is the complex reachable from $C \uparrow U$ in executions where no process in V takes another step

Notation

$(\Xi \downarrow U)(C)$

Complex reachable from C in executions where processes in U halt and the rest finish.

Notation

$(\Xi \downarrow U)(C)$

Complex reachable from C in executions where processes in U halt and the rest finish.

If $V \subseteq U$, then

$\Xi(C \uparrow U) \cap \Xi(C \uparrow V) = (\Xi \downarrow V)(C \uparrow U)$

Notation

$(\Xi \downarrow U)(C)$

Complex reachable from C in executions where processes in U halt and the rest finish.

If $V \subseteq U$, then

$\Xi(C \uparrow U) \cap \Xi(C \uparrow V) = (\Xi \downarrow V)(C \uparrow U)$

Computing Intersections

$U = \{P\}$

P	Q	R
write		
snap		
$C \uparrow U$	write	
	snap	

$(C \uparrow U) \uparrow V \setminus U$

$V = \{Q\}$

P	Q	R
	write	
	snap	
write	$C \uparrow V$	
snap		

$(C \uparrow V) \uparrow U \setminus V$

Memory state same

Only $U \cup V$ can distinguish

Intersections (Case 2)

Then crash processes in U ∪ V

If U, V incomparable, then

$\Xi(C \uparrow U) \cap \Xi(C \uparrow V) = (\Xi \downarrow U \cup V)(C \uparrow U \cup V)$

Run processes in U ∪ V

Lemma

$$\Xi(C \uparrow U) \cap \Xi(C \uparrow V) = (\Xi \downarrow W)(C \uparrow U \cup V)$$

Where

$$W = \begin{cases} U \text{ if } U \subseteq V \\ V \text{ if } V \subseteq U \\ U \cup V \text{ otherwise} \end{cases}$$

Lemma

The nerve graph $\mathcal{N}(\mathit{\Xi}(\mathrm{C} \uparrow \mathrm{U}))$ is path-connected

Proof

Consider vertex v = $\mathit{\Xi}(\mathrm{C} \uparrow \mathit{\Pi})$

Show every vertex has an edge to v

Proof (con't)

$\boxed{\text{Consider vertex } v = \mathit{\Xi}(C \uparrow \mathit{\Pi})}$

Proof (con't)

Consider vertex v = Ξ(C $\uparrow$ Π)

for every U $\subset$ Π consider possible edge …

54

Proof (con't)

Consider vertex $v = \Xi(C \uparrow \Pi)$

for every $U \subset \Pi$ consider possible edge …

$\Xi(C \uparrow \Pi) \cap \Xi(C \uparrow U) = (\Xi \downarrow U)(C \uparrow \Pi)$

Proof (con't)

Consider vertex v = $\Xi(C \uparrow \Pi)$

for every $U \subset \Pi$ consider possible edge …

$\Xi(C \uparrow \Pi) \cap \Xi(C \uparrow U) = (\Xi \downarrow U)(C \uparrow \Pi)$

Run everyone in next layer

56

Proof (con't)

Consider vertex v = $\Xi(C \uparrow \Pi)$

for every $U \subset \Pi$ consider possible edge …

$\Xi(C \uparrow \Pi) \cap \Xi(C \uparrow U) = (\Xi \downarrow U)(C \uparrow \Pi)$

Run everyone in next layer

crash everyone in U

Proof (con't)

Consider vertex v = $\Xi(C \uparrow \Pi)$

for every $U \subset \Pi$ consider possible edge …

$\Xi(C \uparrow \Pi) \cap \Xi(C \uparrow U) = (\Xi \downarrow U)(C \uparrow \Pi)$

Because $U \subset \Pi$,
complex non-empty,
hence edge exists

Run everyone in next layer

crash everyone in U

58

Theorem

For every input simplex σ, the layered IS protocol complex $\Xi(\sigma)$ is path-connected

Proof

Induction on n

Case n=0: $\Xi(\sigma)$ is a single vertex

Case induction step …

Proof

assume $\Xi(C)$ not path-connected

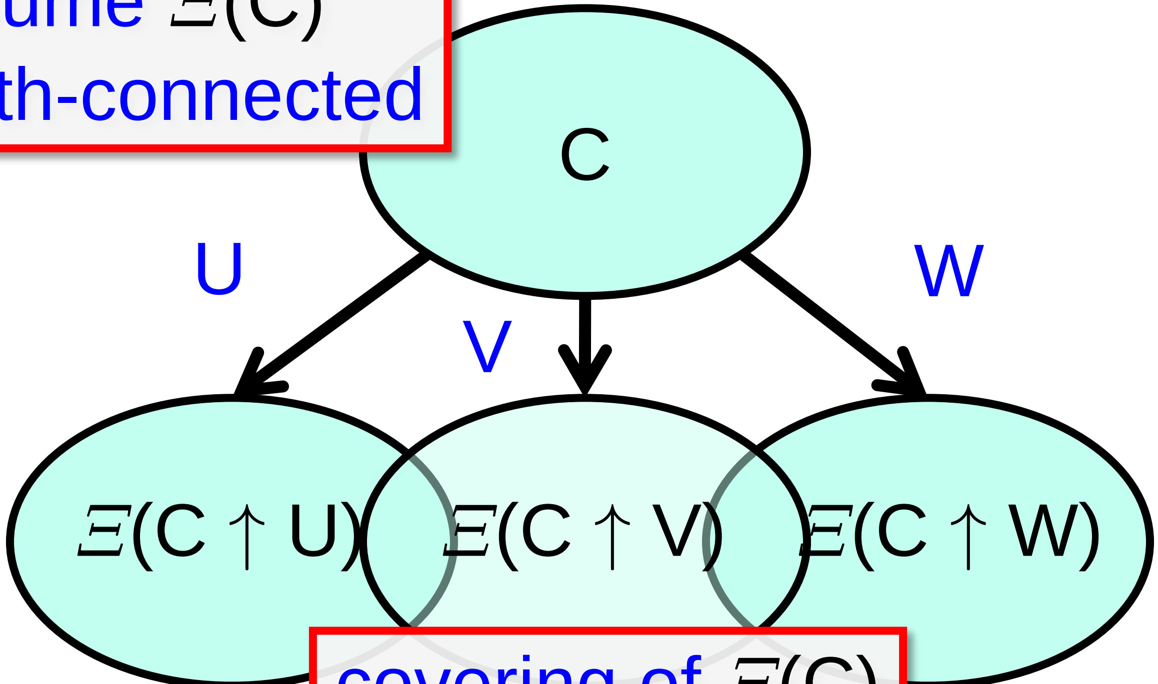

covering of $\Xi(C)$

each $\Xi(C \uparrow U)$ is path connected

their nerve graph path-connected

Proof

covering of $\Xi(C)$

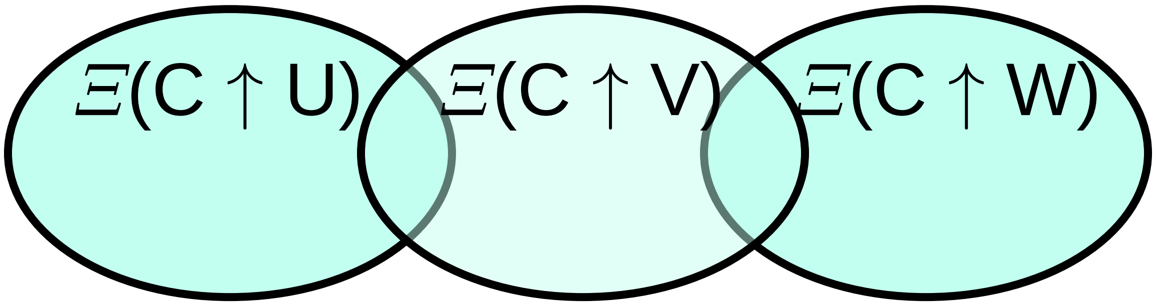

each $\Xi(C \uparrow U)$ is path connected

their nerve graph path-connected

$\Xi(C)$ is path connected by Nerve Lemma

contradiction!

Road Map

Consensus Impossibility

General theorem

Application to read-write models

k-set agreement Impossibility

General theorem

Application to read-write models

So Far …

Expressed solvability of *consensus* as a topological property of protocol complex …

And applied the result to wait-free read-write memory.

Next: do the same for *k*-set agreement!

Rethinking Path Connectivity

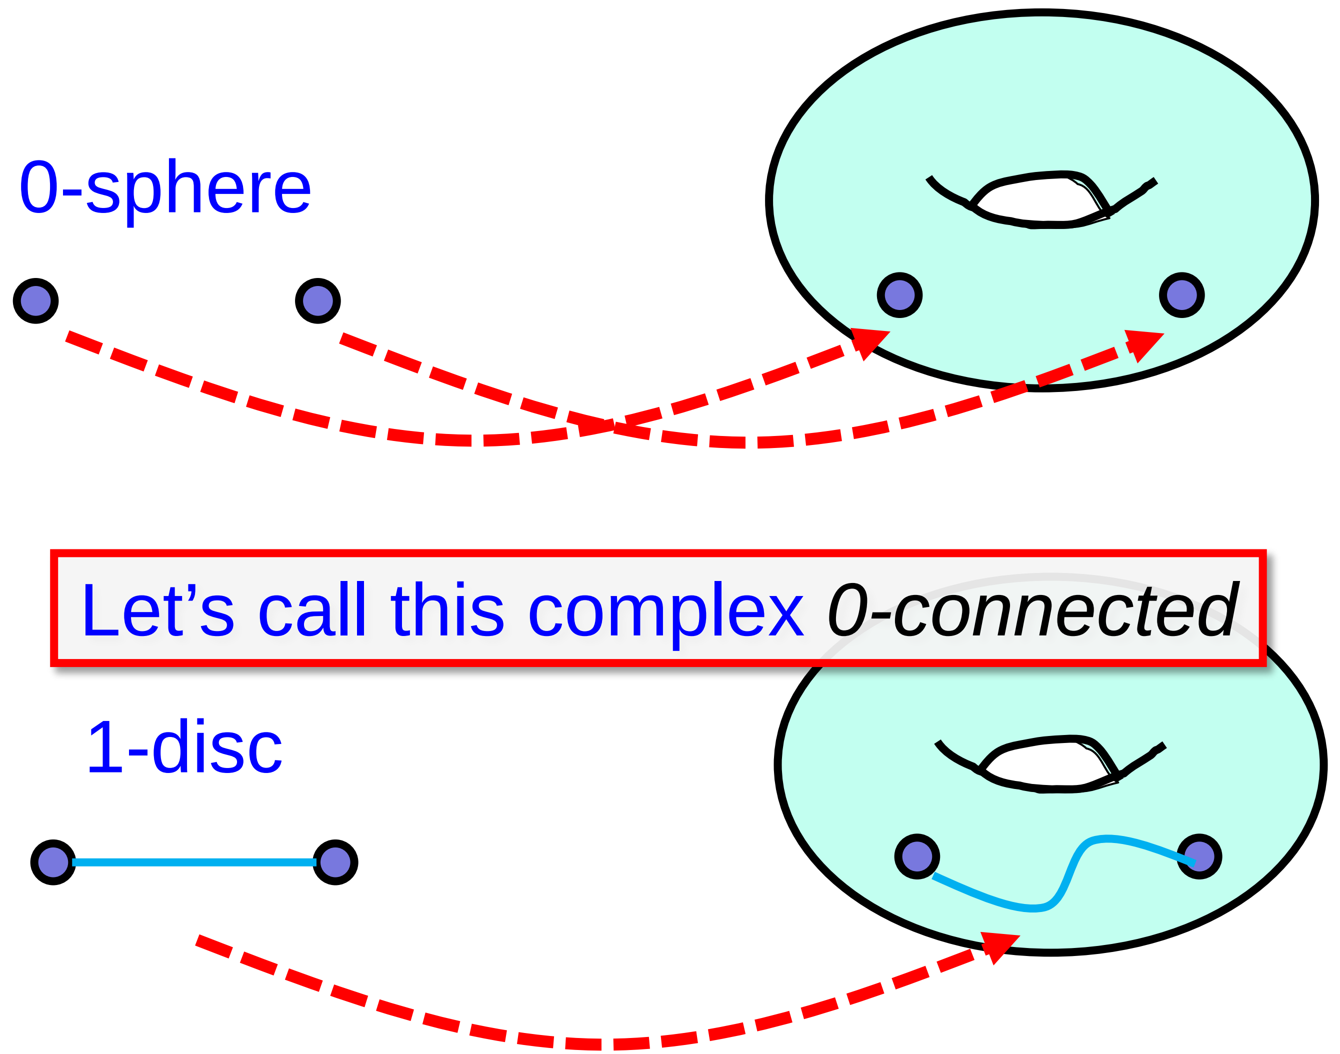

Let’s call this complex *0-connected*

1-Connectivity

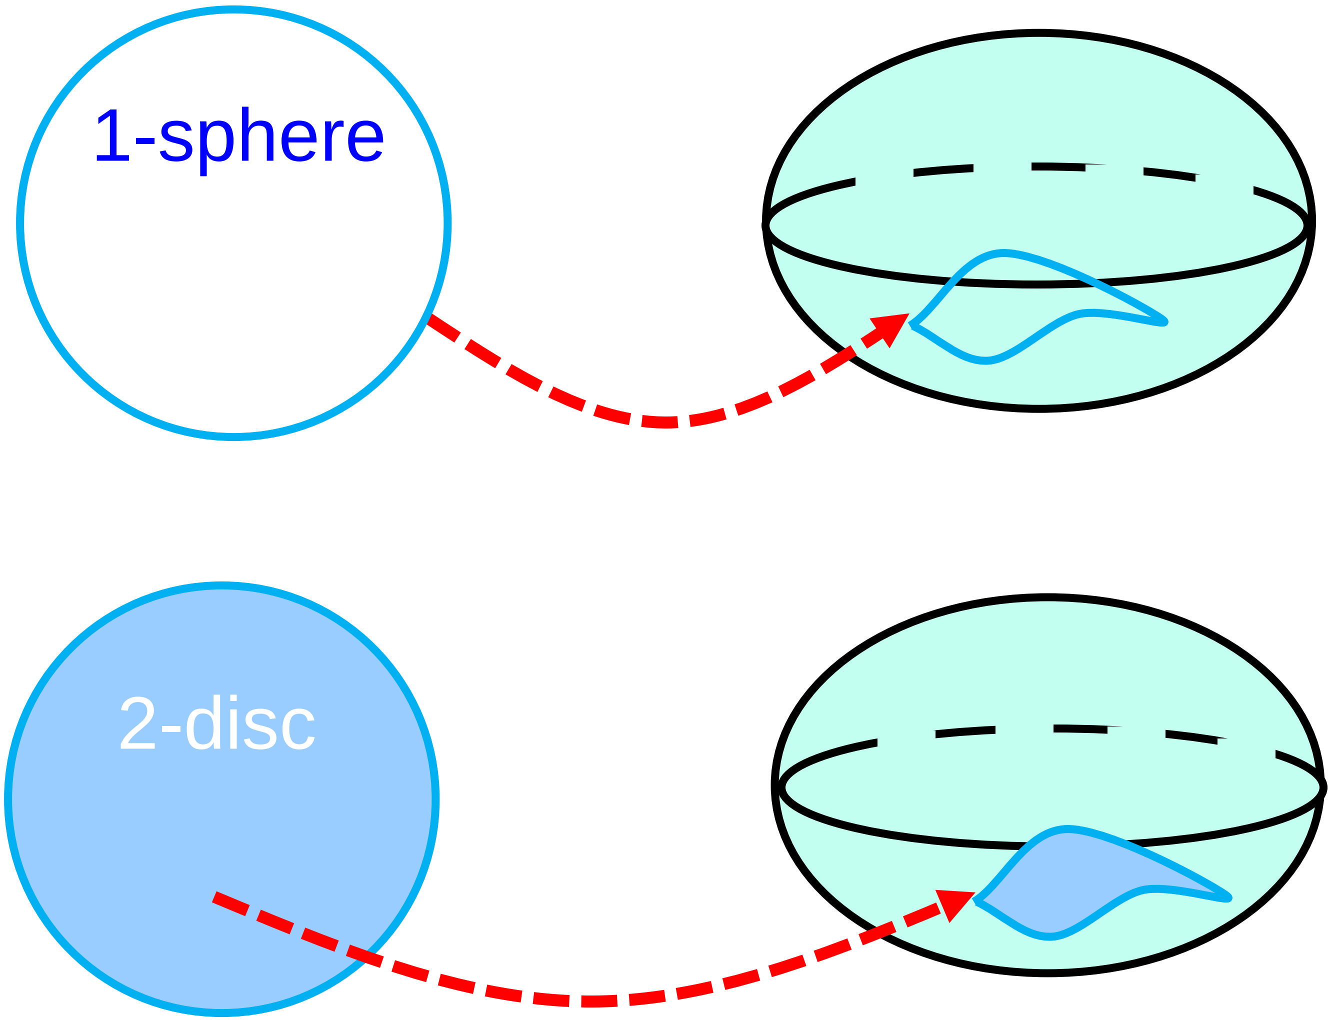

This Complex is not 1-Connected

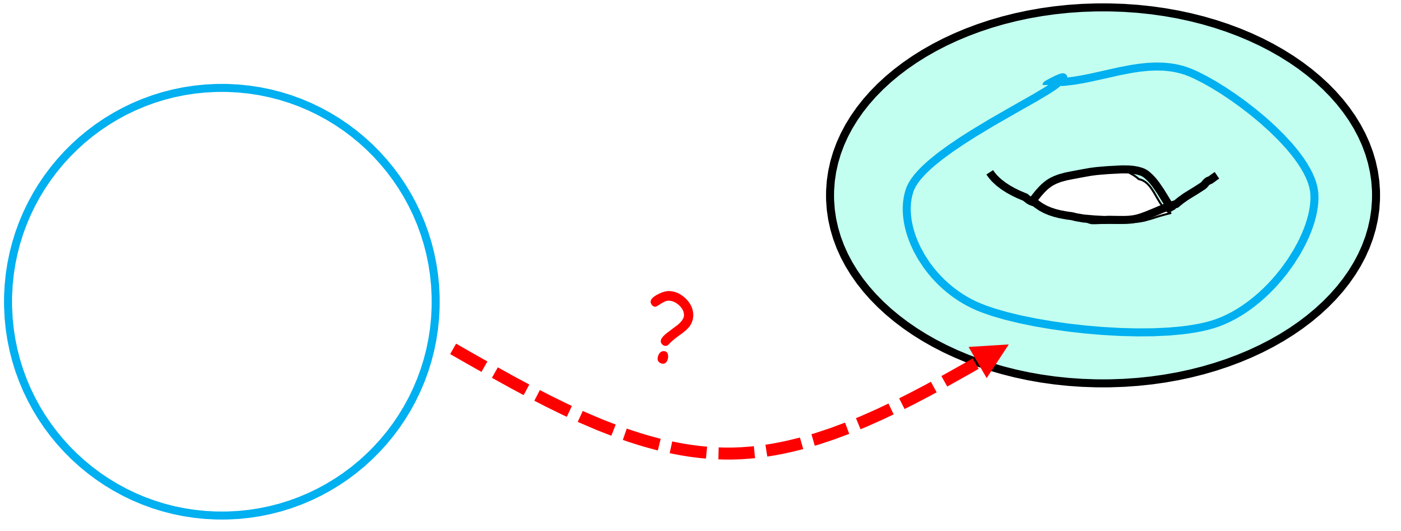

2-Connectivity

2-sphere

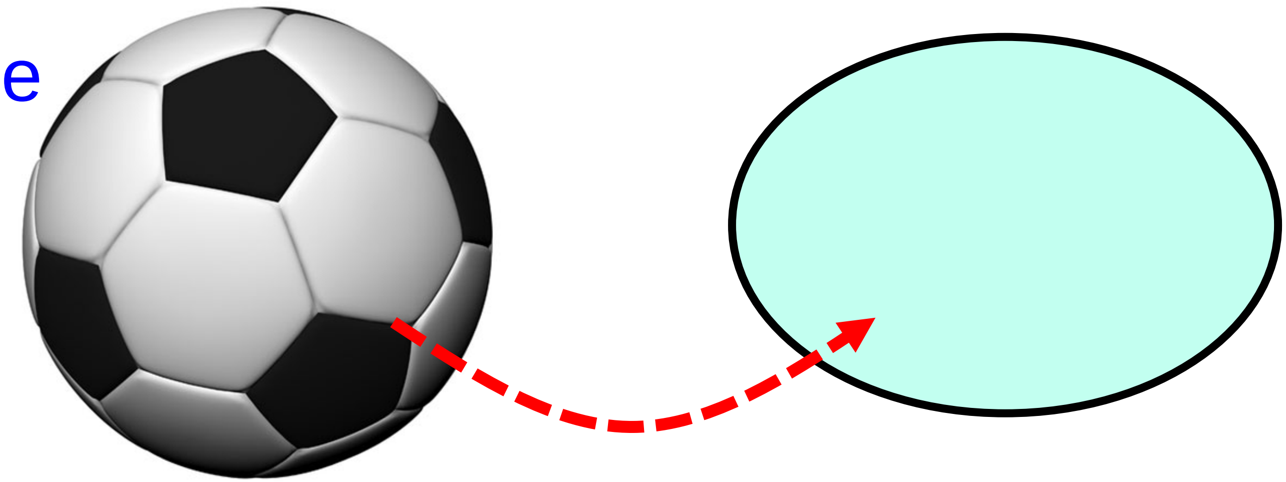

3-disk

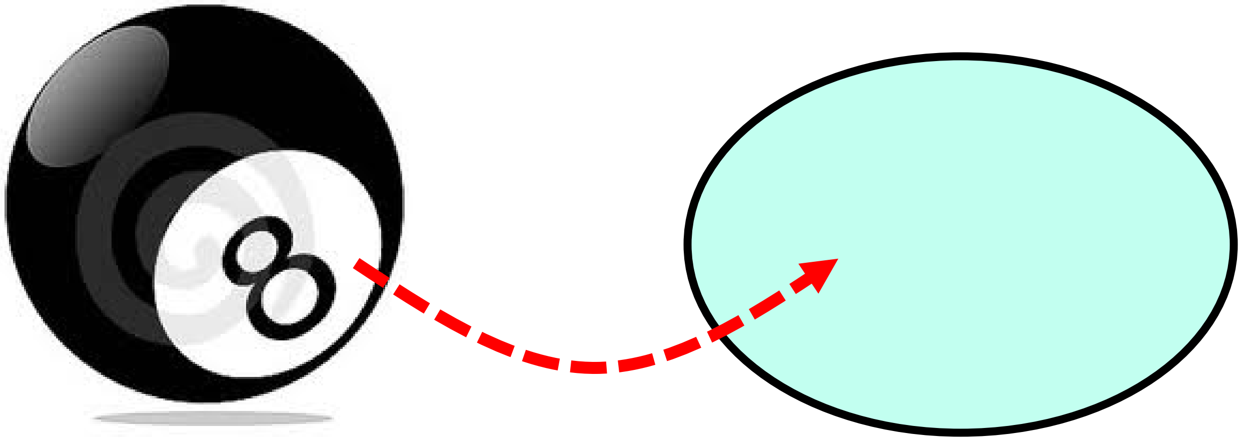

n-connectivity

$\mathcal{C}$ is *n*-connected, if, for *m* $\leq$ *n*, every continuous map of the *m*-sphere

$$f : S^m \to \mathcal{C}$$

can be extended to a continuous map of the (*m+1*)-disk

$$f : D^{m+1} \to \mathcal{C}$$

(*-1*)-connected is non-empty

Road Map

Consensus Impossibility

General theorem

Application to read-write models

k-set agreement Impossibility

General theorem

Application to read-write models

Connectivity and *k*-Set Agreement

Theorem

$(\mathcal{I},\mathcal{O},\Delta)$ an $(n+1)$-process k-set agreement task…

$(\mathcal{I},\mathcal{P},\Xi)$ a protocol …

such that $\Xi(\sigma)$ is $(k-1)$-connected for all σ in $\mathcal{I}$ …

then $(\mathcal{I},\mathcal{P},\Xi)$ cannot solve k-set agreement.

Lemma

carrier map $\Phi: \mathcal{A} \mapsto 2^{\mathcal{B}}$

such that for all $\alpha \in \mathcal{A}$,

$\Phi(\alpha)$ is ((dim α) – 1)-connected.

Then Φ has a simplicial approximation ϕ: $\mathrm{Div}^N \mathcal{A} \to \mathcal{B}$.

Lemma Proof Sketch

carrier map Φ: $\mathcal{A} \to 2^{\mathcal{B}}$

has continuous approximation f: $|\mathcal{A}| \to |\mathcal{B}|$

$f(|\sigma|) \subseteq |\Phi(\sigma)|$

Inductive construction …

Lemma Proof Sketch

continuous approximation f: $|\mathcal{A}| \to |\mathcal{B}|$

Base

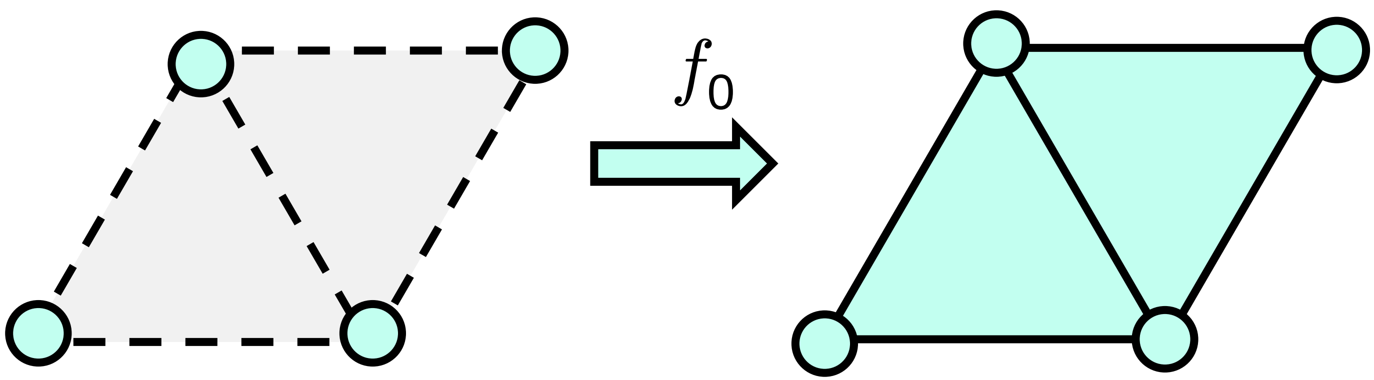

define on vertices ...

Lemma Proof Sketch

Step

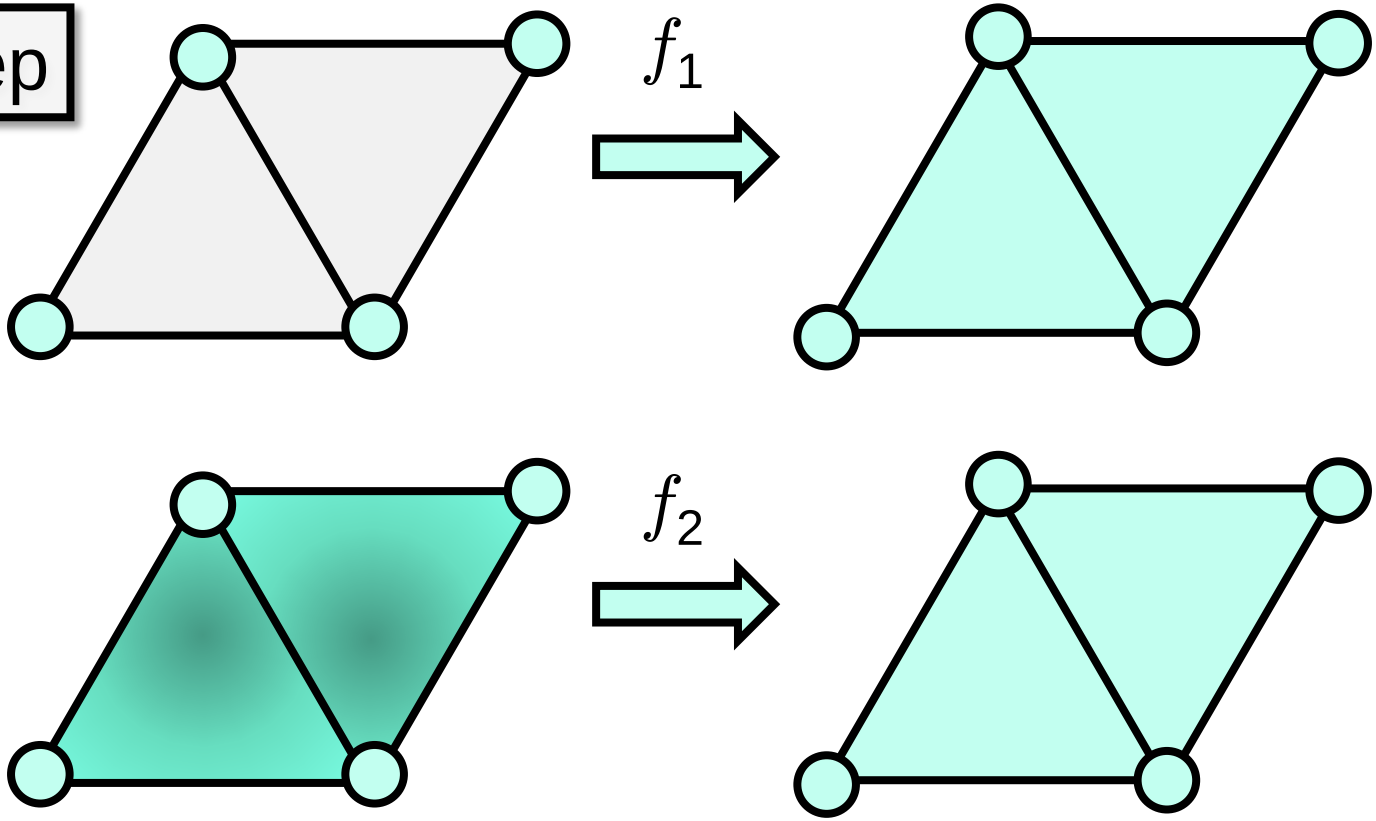

connectivity allows “filling in”

Lemma Proof Sketch

$\boxed{\text{continuous approximation } f\text{: } |\mathcal{A}| \to |\mathcal{B}|}$

take simplicial approximation
ϕ: Div $\mathcal{A} \to \mathcal{B}$ of f: $|\mathcal{A}| \to |\mathcal{B}|$

Theorem Proof Sketch

let $\sigma \in \mathcal{I}$ have k+1 distinct input values

let Δ^k be simplex labeled with k+1 values

$\partial\Delta^k$ its (k-1)-skeleton

c: $\Xi(\sigma) \to \partial\Delta^k$ well-defined simplicial map

By lemma, Ξ has simplicial approximation

ϕ: Div $\sigma \to \Xi(\sigma)$ of f: $|\mathcal{A}| \to |\mathcal{B}|$

Theorem Proof Sketch

c: $\Xi(\sigma) \to \partial\Delta^k$ well-defined simplicial map

By lemma, Ξ has simplicial approximation

ϕ: Div $\sigma \to \Xi(\sigma)$ of f: $|\mathcal{A}| \to |\mathcal{B}|$

composition Div $\sigma \to \Xi(\sigma) \to \partial\Delta^k$

defines a Sperner coloring of Div σ

some τ in Div σ maps to all of Δ^k

Contradiction!

k-Connectivity is an Eventual Property

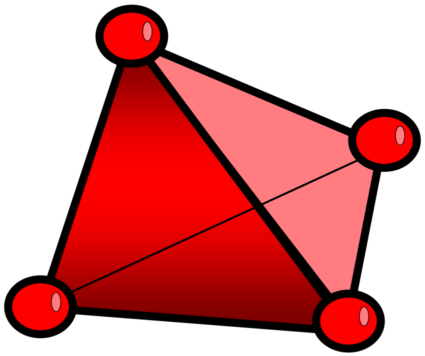

Individual simplex is
k-connected

k-Connectivity has a Critical State

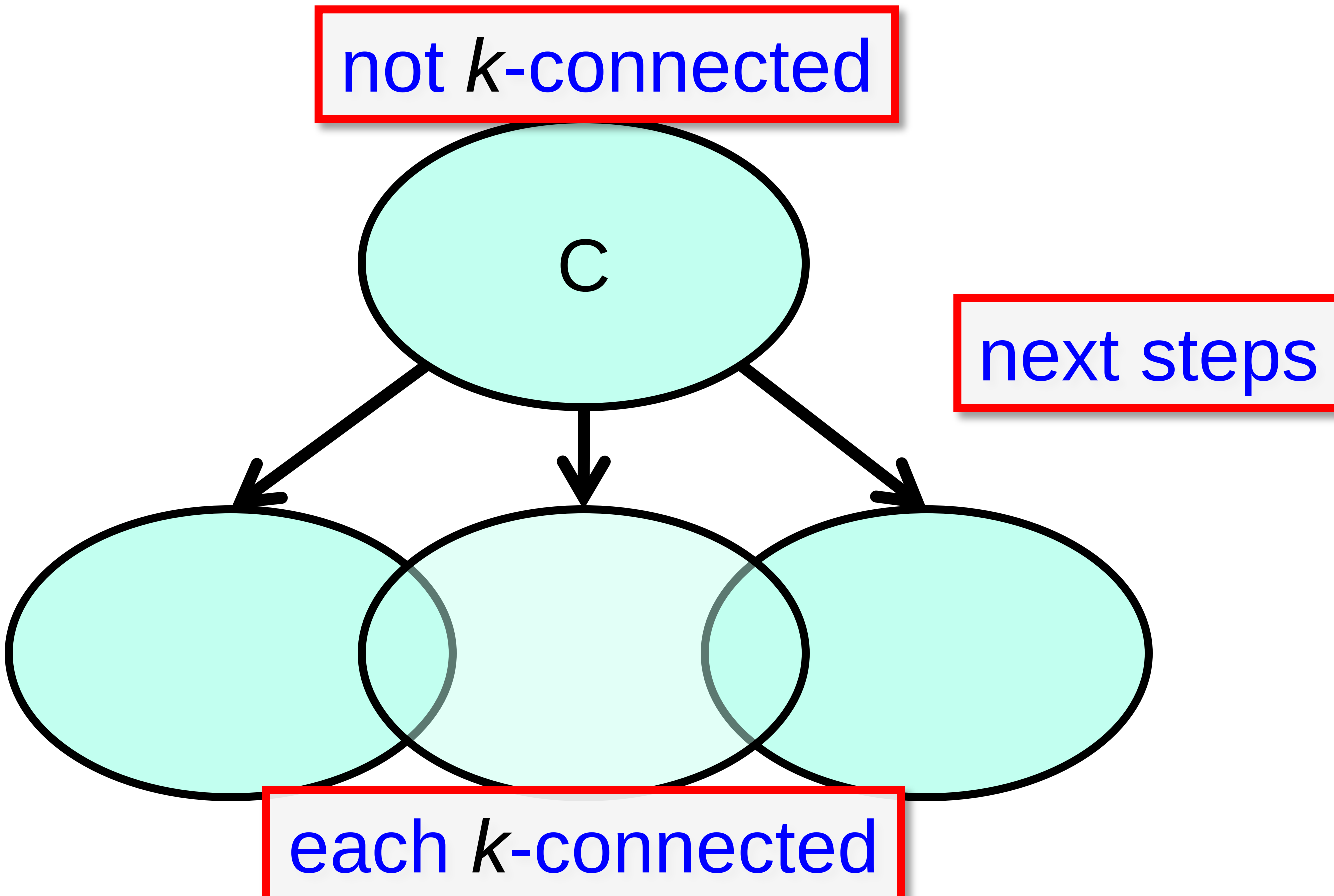

Critical State in Layered IS

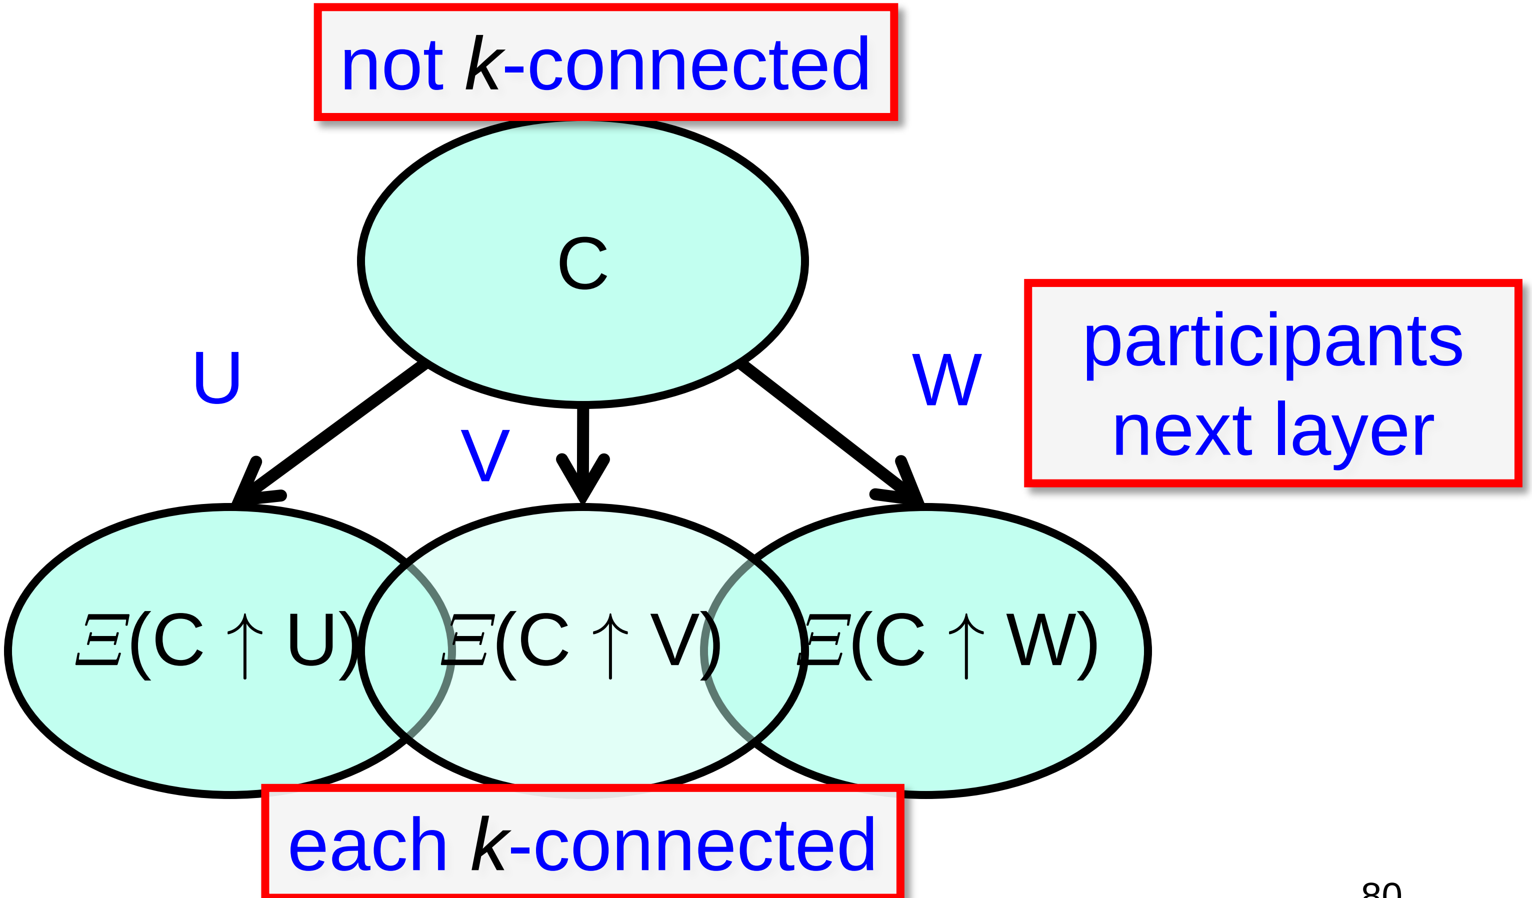

Proof Strategy

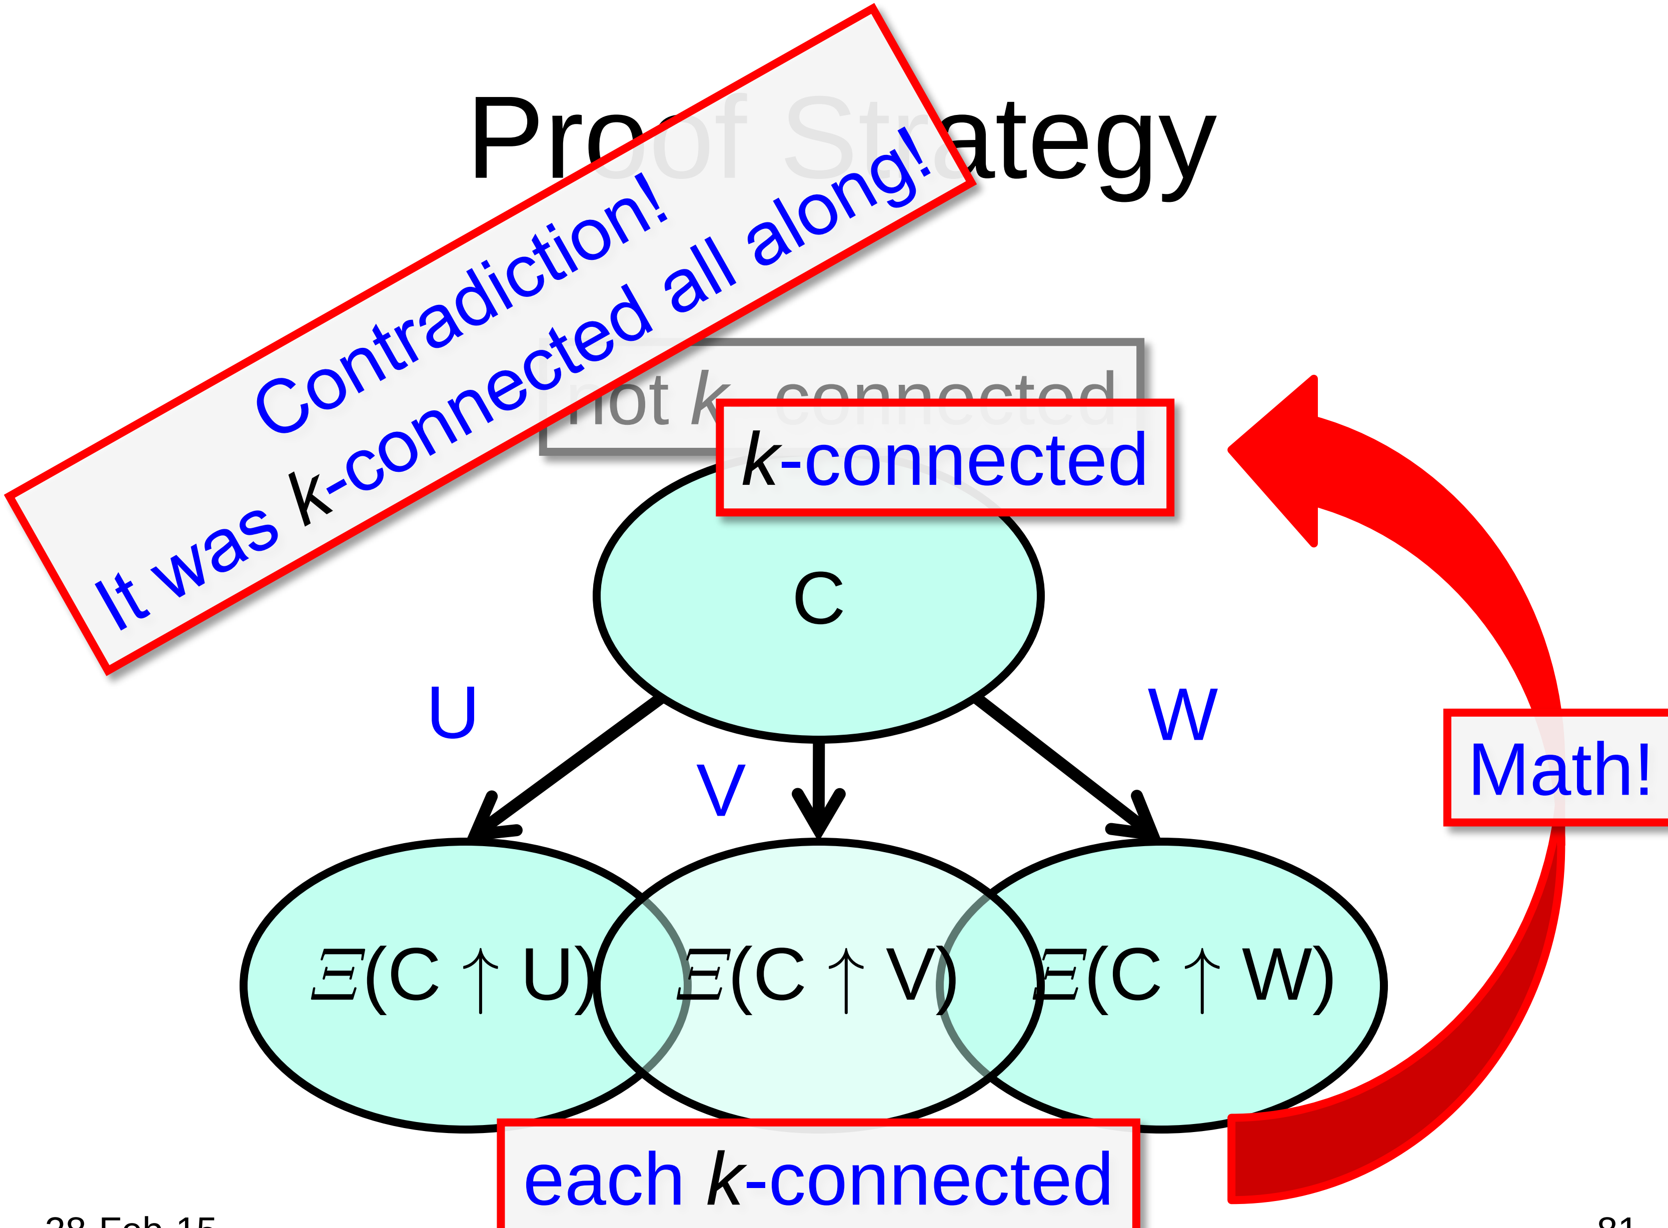

Nerve Lemma

Reason about connectivity of a complex…

From connectivity of components …

And how they fit together.

Covering

$$\mathcal{C} = \mathcal{C}_0 \cup \cdots \cup \mathcal{C}_\ell$$

Complex

Covering

$$\mathcal{C} = \mathcal{C}_0 \cup \cdots \cup \mathcal{C}_\ell$$

Complex

Covering

$$\mathcal{N}(C_0, \ldots, \mathcal{C}_\ell)$$

$\mathcal{C}_i$

Vertex

$$\mathcal{C}_U = \{\mathcal{C}_i | i \in U\}$$

Simplex, if $\cap\mathcal{C}_i$ non-empty

Nerve Example: Sphere

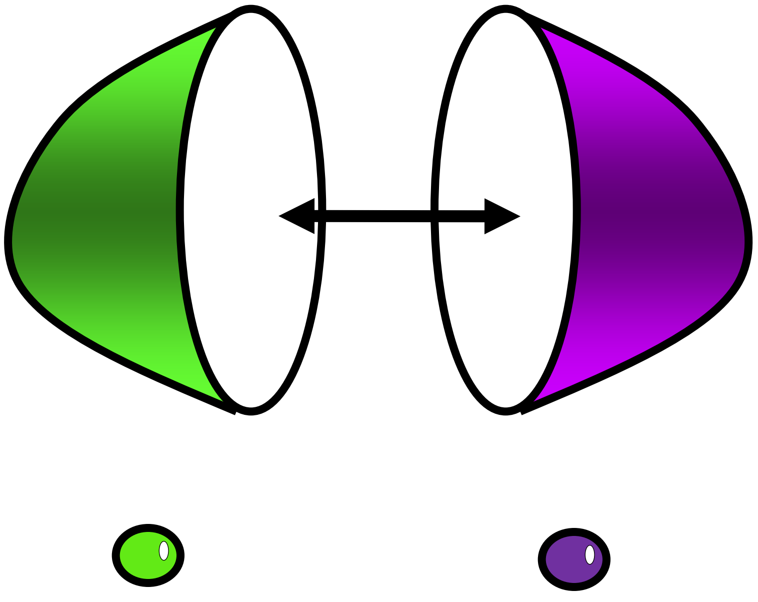

Reasoning About Connectivity

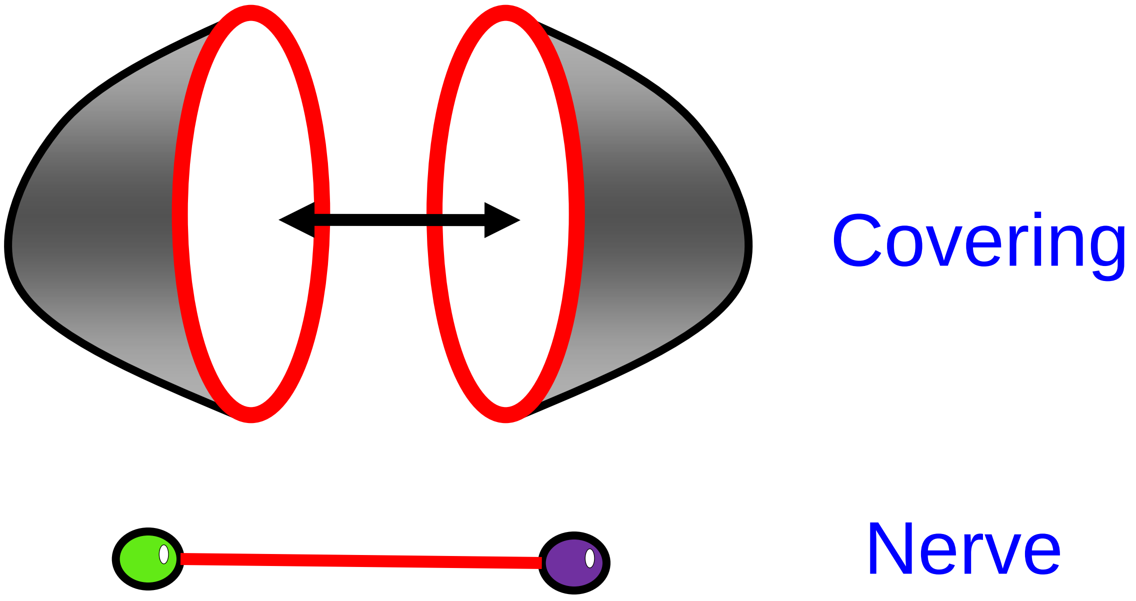

Nerve Lemma

$$\boxed{\mathcal{C}_i}$$

k-connected

If …

(*k*-|*U*|+1)-connected

$$\boxed{\mathcal{C}_U} = \bigcap_{i \in U} \mathcal{C}_i$$

…Then

$\mathcal{C}$ is k-connected …

if and only if ..

$\mathcal{N}(\mathcal{C}_0, \ldots \mathcal{C}_m)$ is k-connected.

Nerve Example: Sphere

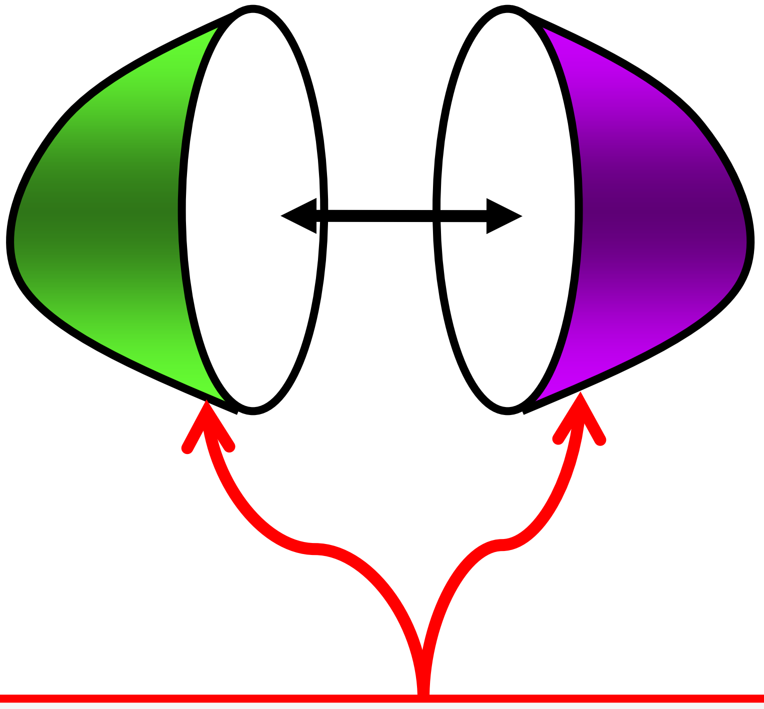

1-connected …

Reasoning About Connectivity

Can apply Nerve Lemma!

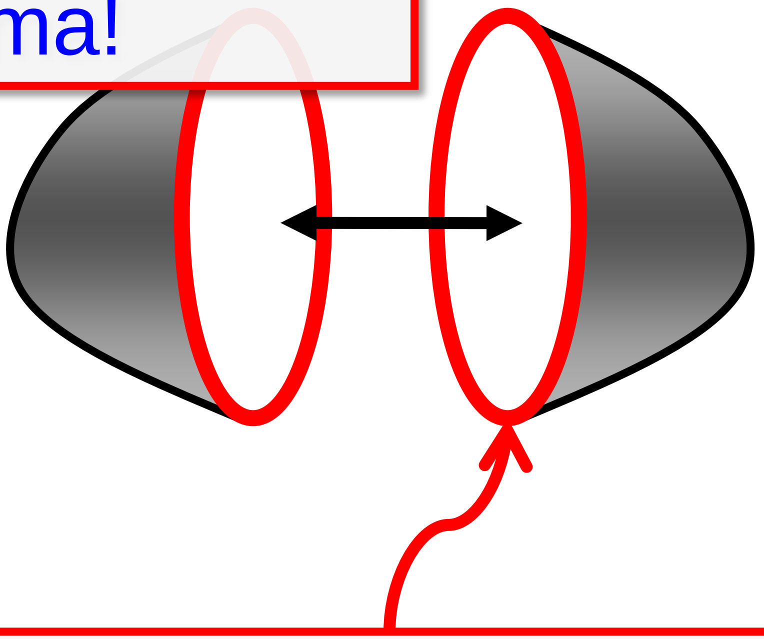

0-connected

Reasoning About Connectivity

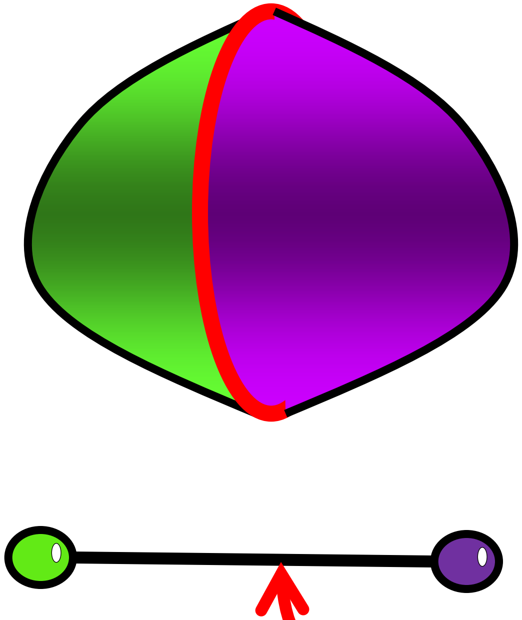

1-connected …

Reasoning About Connectivity

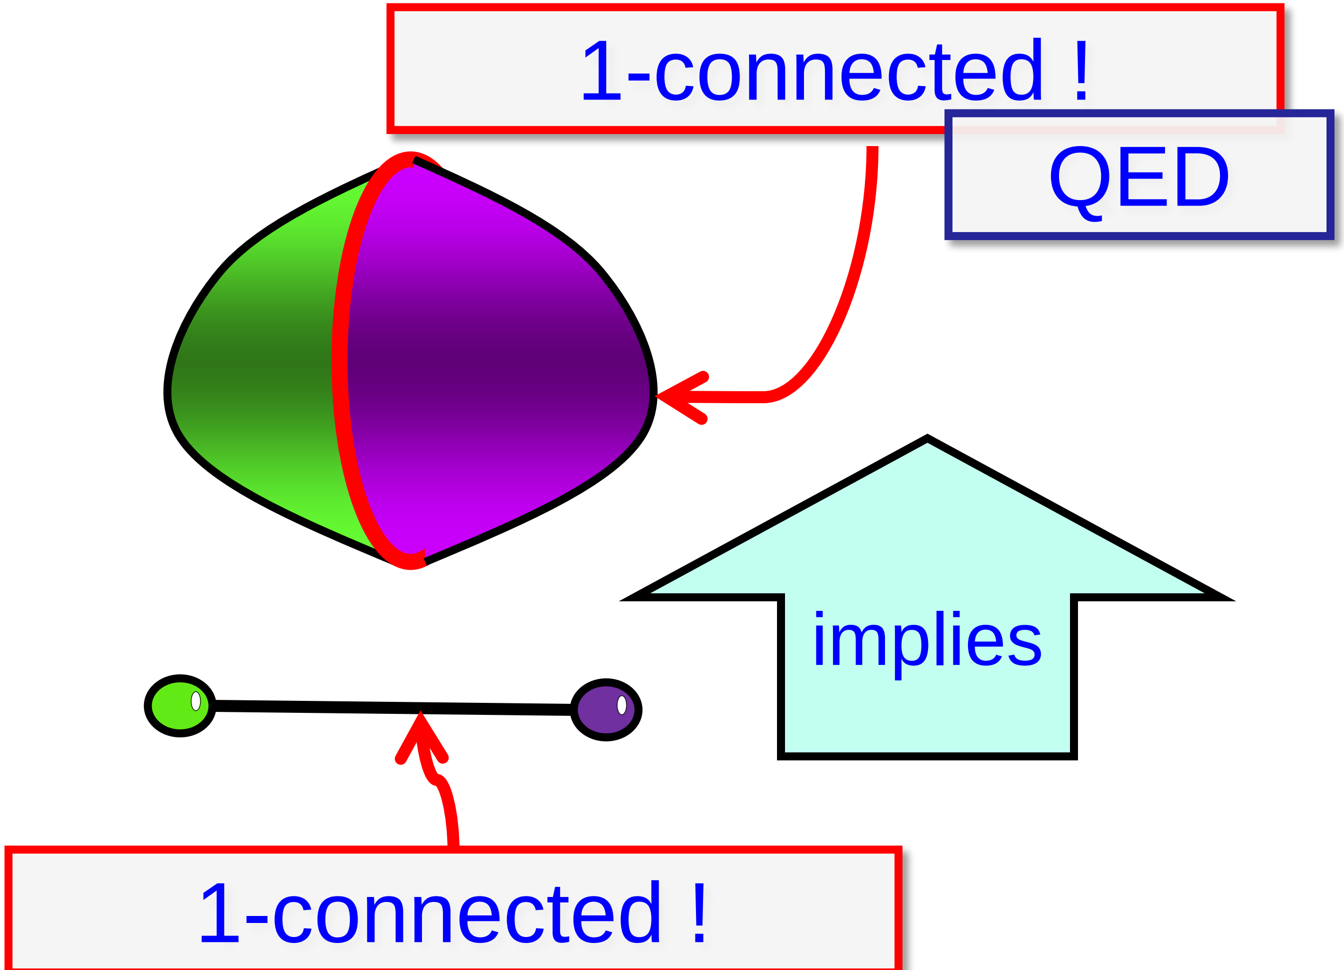

Nerve Complex Lemma

The nerve complex
$\mathcal{N}(\Xi(\mathrm{C} \uparrow \mathrm{U}) \mid \emptyset \subseteq \mathrm{U} \subseteq \Pi)$
is *n*-connected

Proof

Consider vertex v = $\Xi(\mathrm{C} \uparrow \Pi)$

Show the nerve complex is a *cone* with apex v

Nerve Complex Lemma

each $\Xi(C \uparrow U_i)$ is a vertex

each set $\{\Xi(C \uparrow U_i) \mid i = 0, \ldots, m\}$

Is a simplex if and only if

$\bigcap_i \Xi(C \uparrow U_i) \neq \emptyset$

Reasoning About Intersections

Lemma

Let $U_0, \ldots, U_m$ sets of process names …

Indexed so $|U_0| \geq \ldots \geq |U_m|$

Intersection Lemma

$\bigcap \mathit{\Xi}(C \uparrow U_i) = (\mathit{\Xi} \downarrow W)\ (C \uparrow \cup U_i)$

Where

$$W = \begin{cases} \bigcup_{i=1}^{m} U_i \text{ if } \bigcup_{i=1}^{m} U_i \subseteq U_0 \\ \bigcup_{i=0}^{m} U_i \text{ otherwise} \end{cases}$$

Proof is inductive version of earlier lemma

Corollary

If $\bigcup_i U_i = \mathit{\Pi}$ but each $U_i \neq \mathit{\Pi}$,

then $\bigcap_i \mathit{\Xi}(C \uparrow U_i) = \emptyset$.

Nerve Complex

Let vertex $v = \Xi(C \uparrow \Pi)$

Let $\sigma = \{\Xi(C \uparrow U_i)\}$ be a simplex

So $\bigcap_i \Xi(C \uparrow U_i) \neq \emptyset$

and $\bigcup_i U_i \neq \Pi$

must show that $\sigma \cup \{v\}$ is a simplex …

Intersection Lemma Proof

to show that $\sigma \cup \{v\}$ is a simplex, show that ...

$\Xi(C \uparrow \Pi) \bigcap \bigcap_i \Xi(C \uparrow U_i)$

is non-empty.

Intersection Lemma Proof

$\Xi(C \uparrow \Pi) \bigcap \left(\bigcap_I \Xi(C \uparrow U_i)\right) = \left(\Xi \boxed{\downarrow \bigcup_i U_i}\right) (C \uparrow \Pi)$

by corollary, $\bigcup_i \neq \Pi$

don't halt everyone

complex non-empty

it's a simplex!

QED

Lemma

$\bigcap_{\square \in I} \Xi(C \uparrow U_i)$ is $(n\text{-}|I|+1)$-connected

argue by induction on *n*

trivial for $n = 0$ …

Proof

$\bigcap_{\square \in I} \mathit{\Xi}(C \uparrow U_i) = (\mathit{\Xi} \downarrow W)(C \uparrow X)$

for $|W| > 0$, $W \subseteq X \subseteq \cup_i U_i$.

a protocol complex for n-$|W|$+1 processes …

either empty, or n-connected by induction hypothesis.

therefore (n-$|I|$+1)-connected

Theorem

For every input simplex σ, the layered IS protocol complex $\Xi(\sigma)$ is k-connected

Proof

Induction on n

Case n=0: $\Xi(\sigma)$ is a single vertex

Case induction step …

Proof

assume $\Xi(C)$
not k-connected

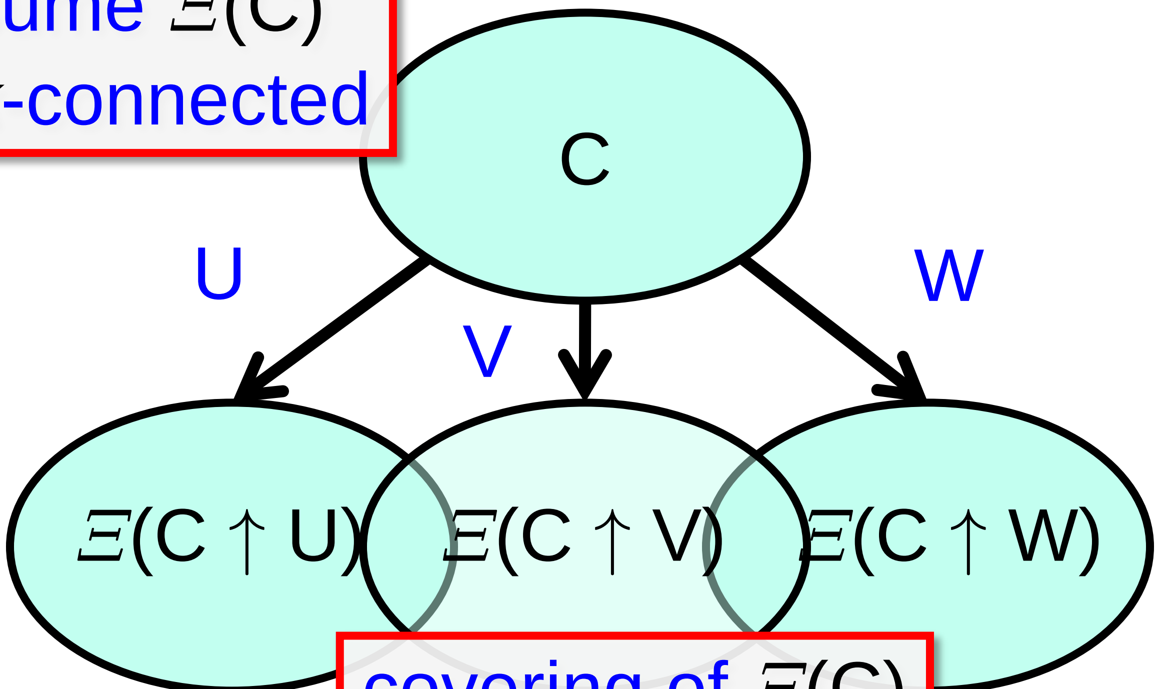

covering of $\Xi(C)$

$\bigcap_{i \in I} \Xi(C \uparrow U_i)$ is $(n-|I|+1)$-connected

their nerve complex is k-connected

Proof

covering of Ξ(C)

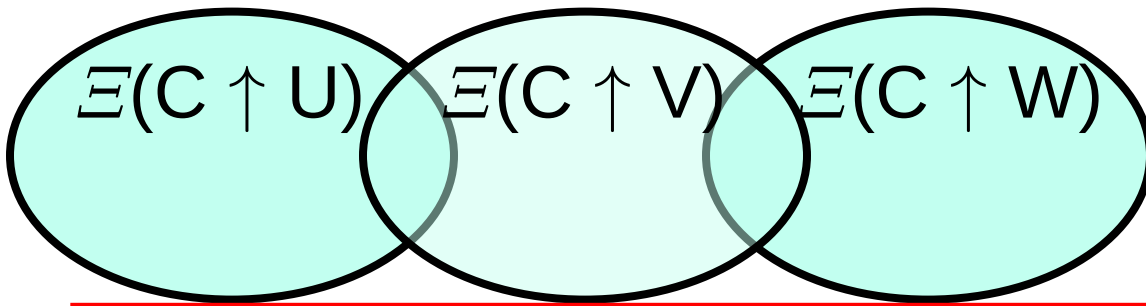

each $\bigcap_{i \in I} \Xi(C \uparrow U_i)$ is $(n-|I|+1)$-connected

their nerve graph k-connected

Ξ(C) is k-connected by Nerve Lemma

contradiction!

Conclusions

Model-independent topological properties that prevent …

consensus

k-set agreement

path-connectivity

k-connectivity

Model-specific application to wait-free read-write memory