

The Topology of Asynchronous Byzantine Colorless Tasks

Joint with
Hammurabi Mendes
Christine Tasson



BROWN

Overview

Basic concepts of
combinatorial
topology

Duality between
combinatorial & continuous
mechanisms

First *characterization* of tasks
that can be solved in the
asynchronous Byzantine model

Road Map

Colorless Tasks

Operational Model

Combinatorial Model

Building Blocks

Crash Failure Solvability

Byzantine Failure Solvability

Road Map

Colorless Tasks

Operational Model


Combinatorial Model

Building Blocks


Crash Failure Solvability

Byzantine Failure Solvability

Tasks

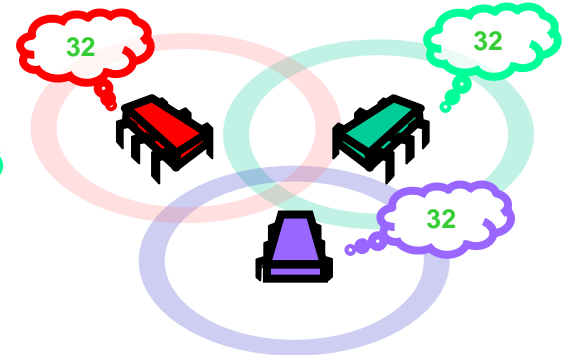




Possible set of *input* values




Finite computation


Possible set of *output* values



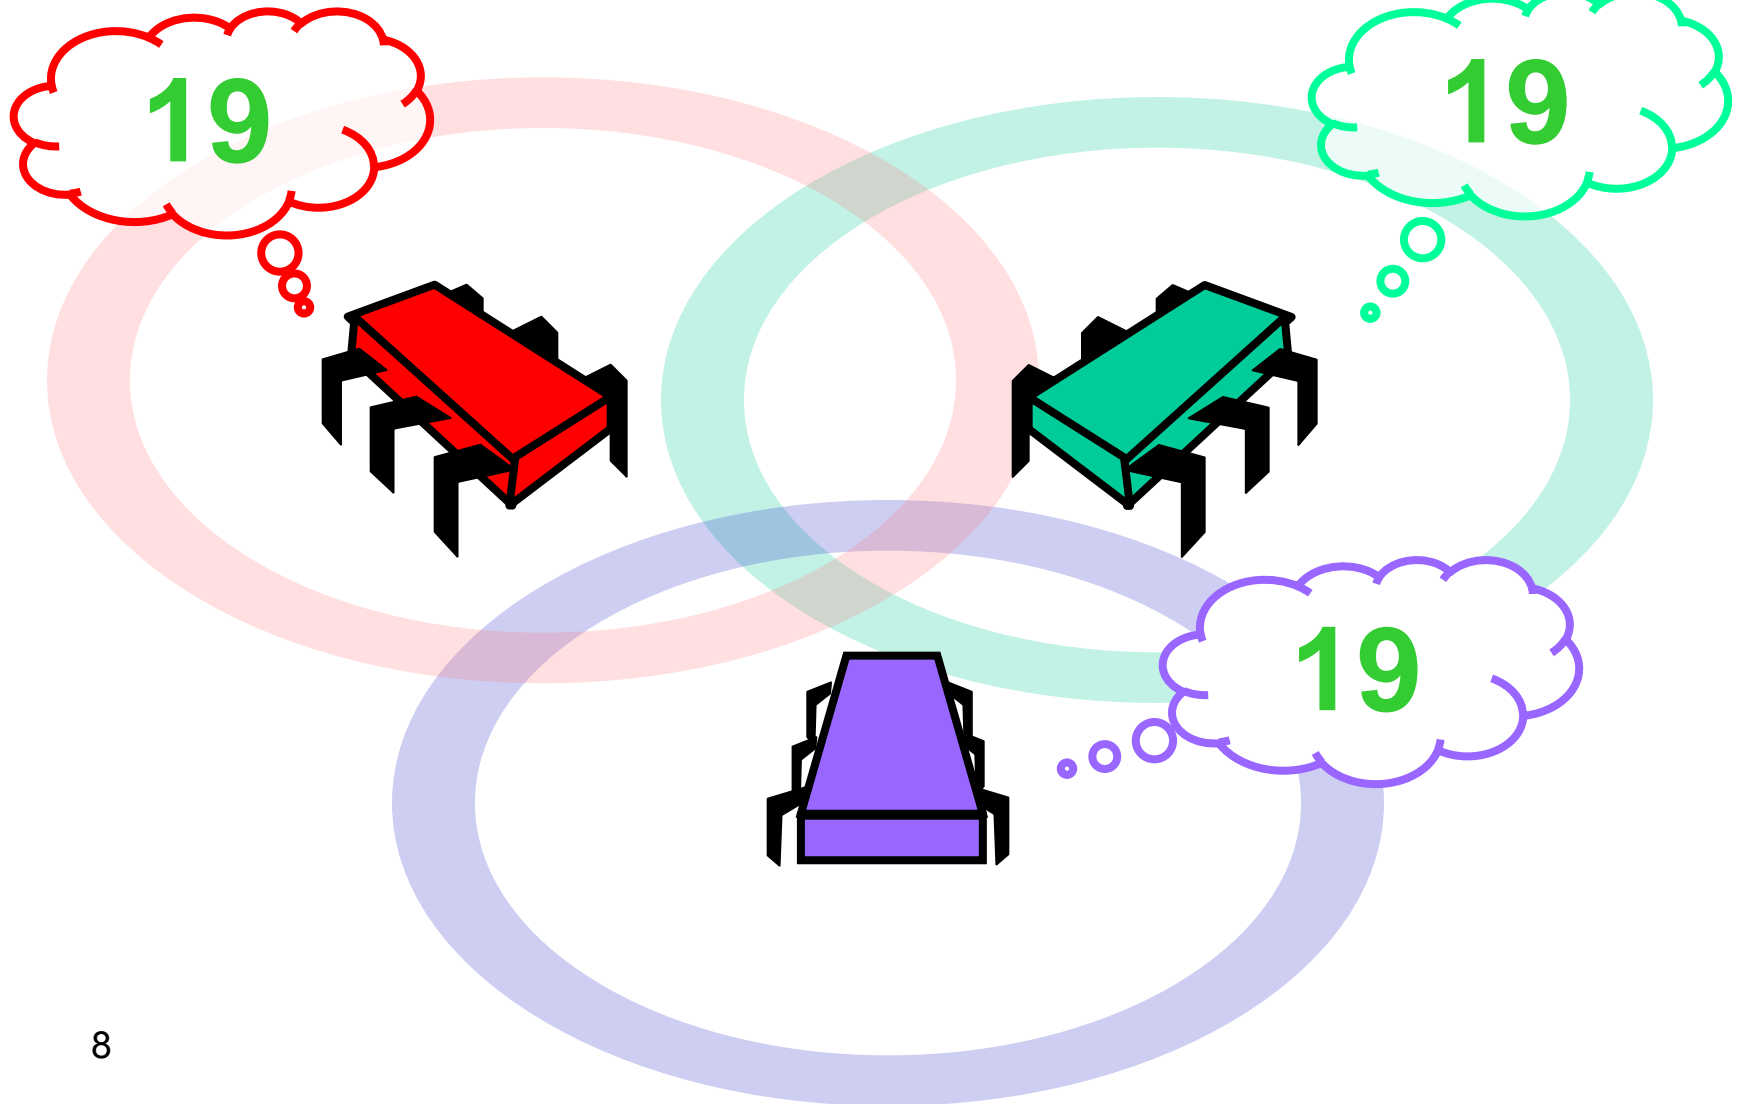
Consensus Start




Communication




Consensus Finish




k-Set Agreement Start




Communication




k -Set Agreement Finish



Colorless Tasks



Colorless Tasks



Colorless Tasks

The *set* of input values ...

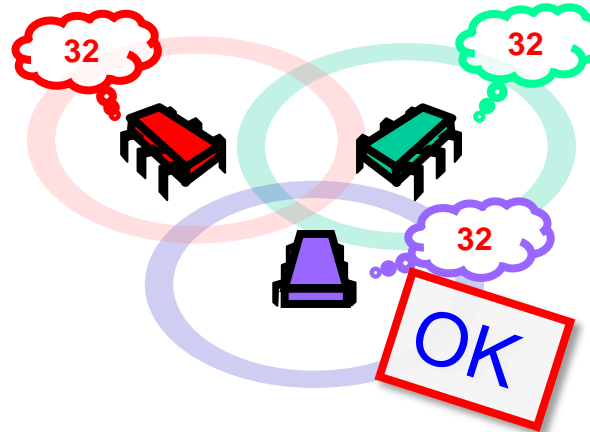
determines the *set* of output values.

Number and identities irrelevant...

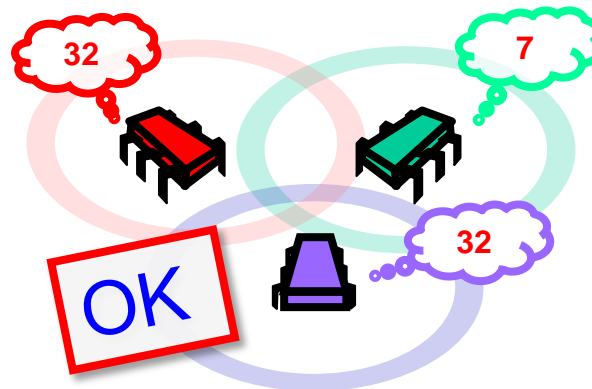
for both input and output values



Examples

Consensus




k -set agreement





Majority
Not OK!



Road Map

Colorless Tasks

Operational Model

Combinatorial Model

Building Blocks

Crash Failure Solvability

Byzantine Failure Solvability

Failures



Crash failures: processes halt

How many?

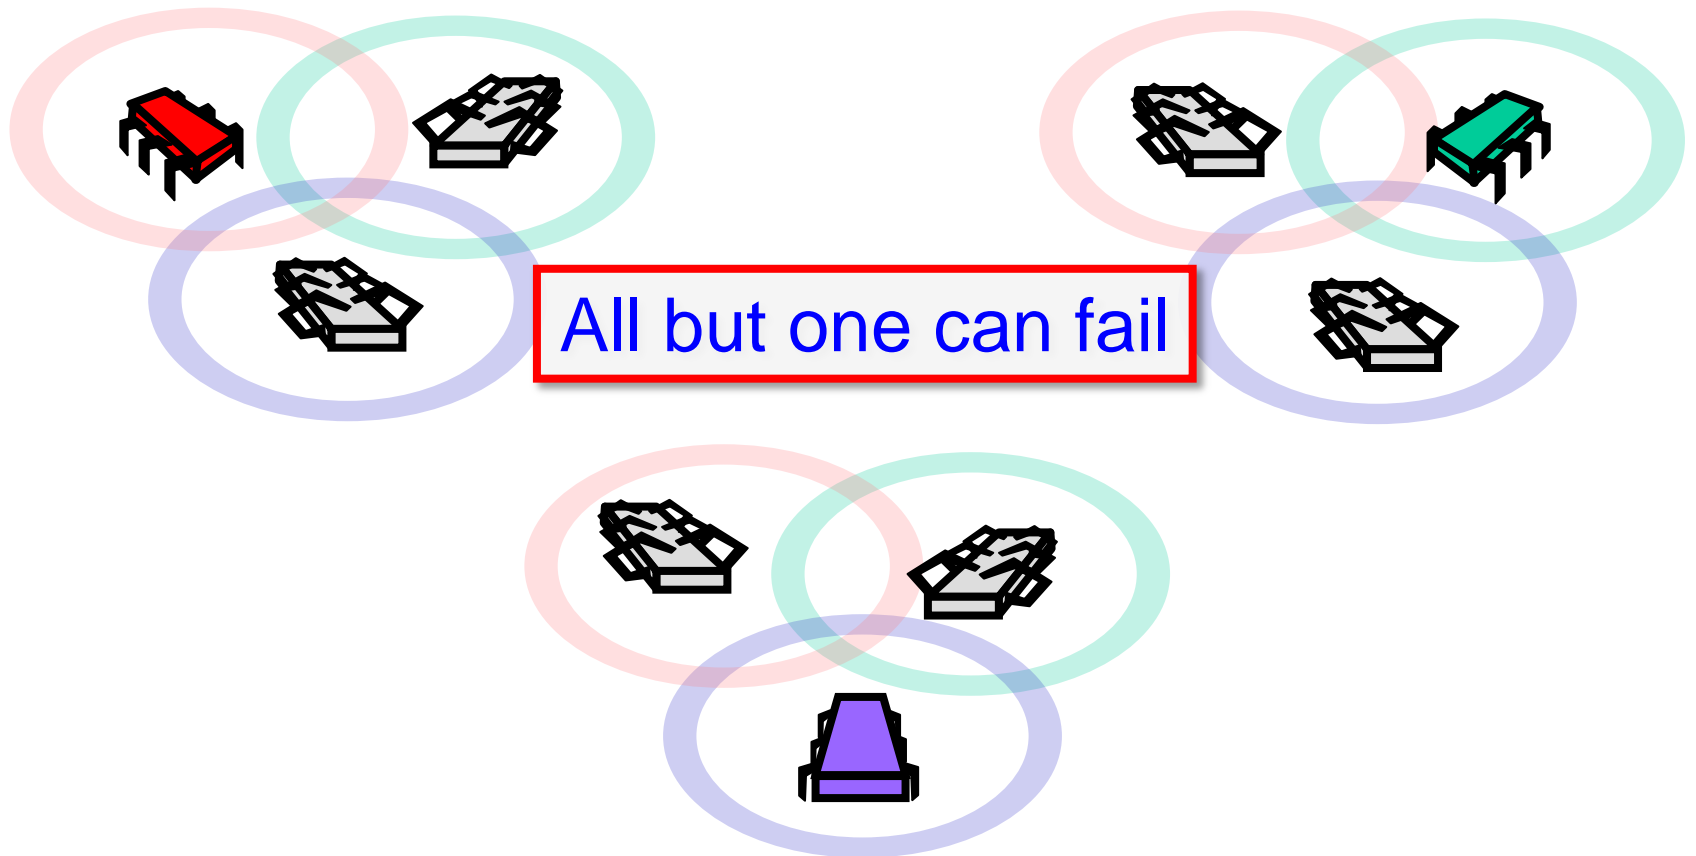
Failures



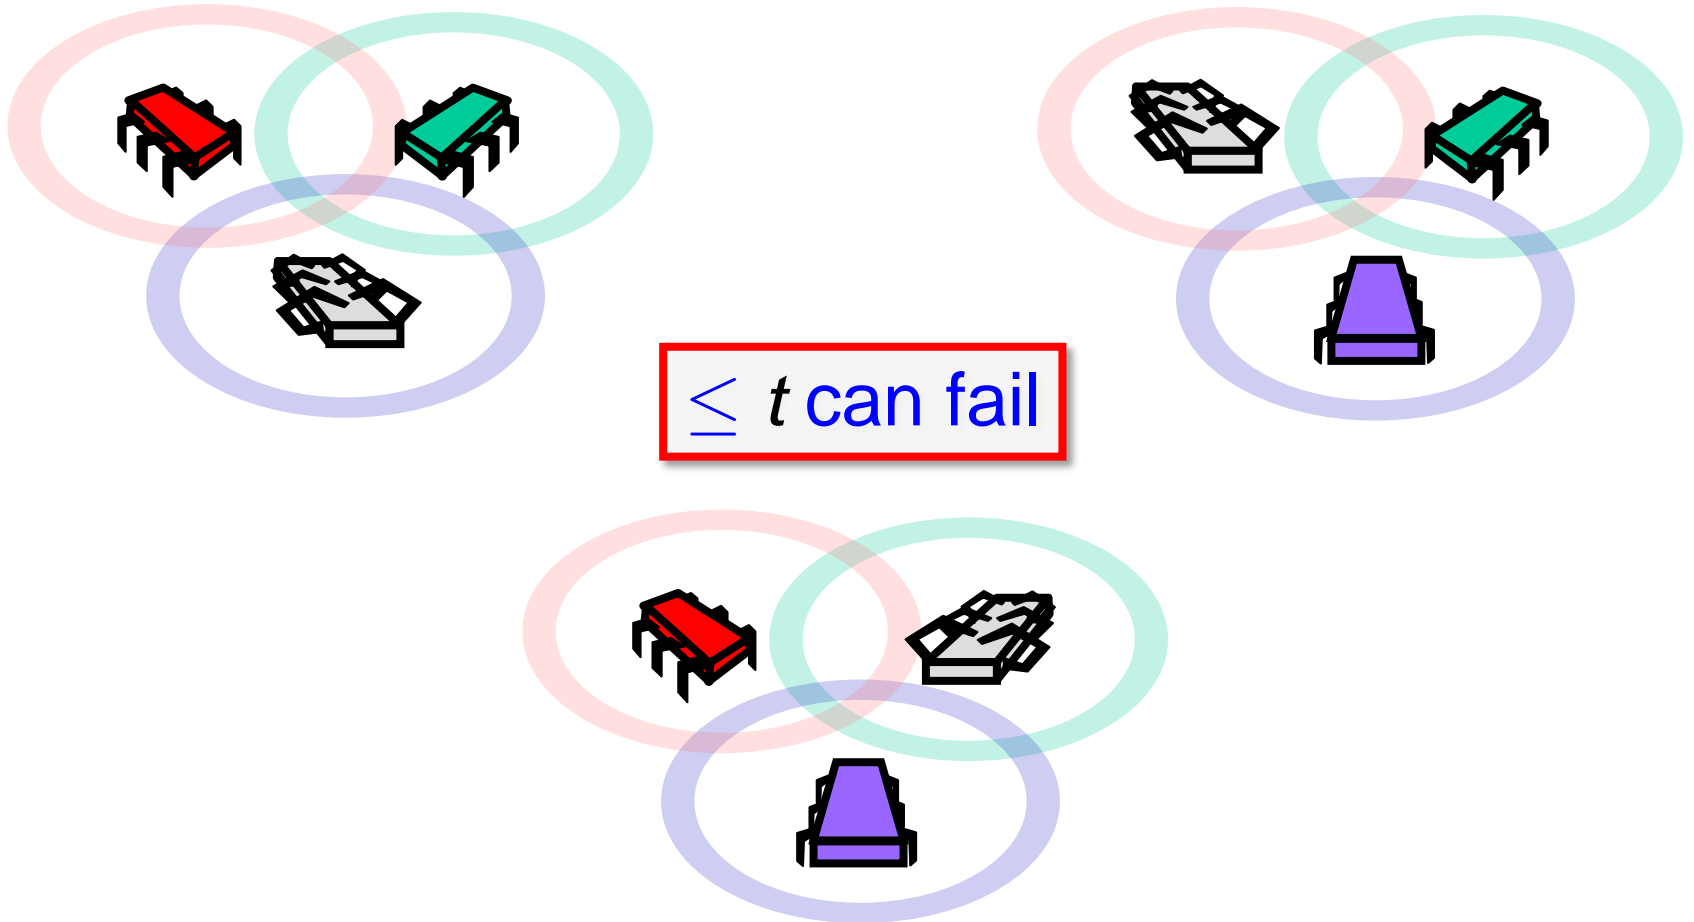
Crash failures: processes misbehave

How many?

Resilience: Wait-Free



Resilience: t -resilient



Road Map

Colorless Tasks

Operational Model

Combinatorial Model

Building Blocks

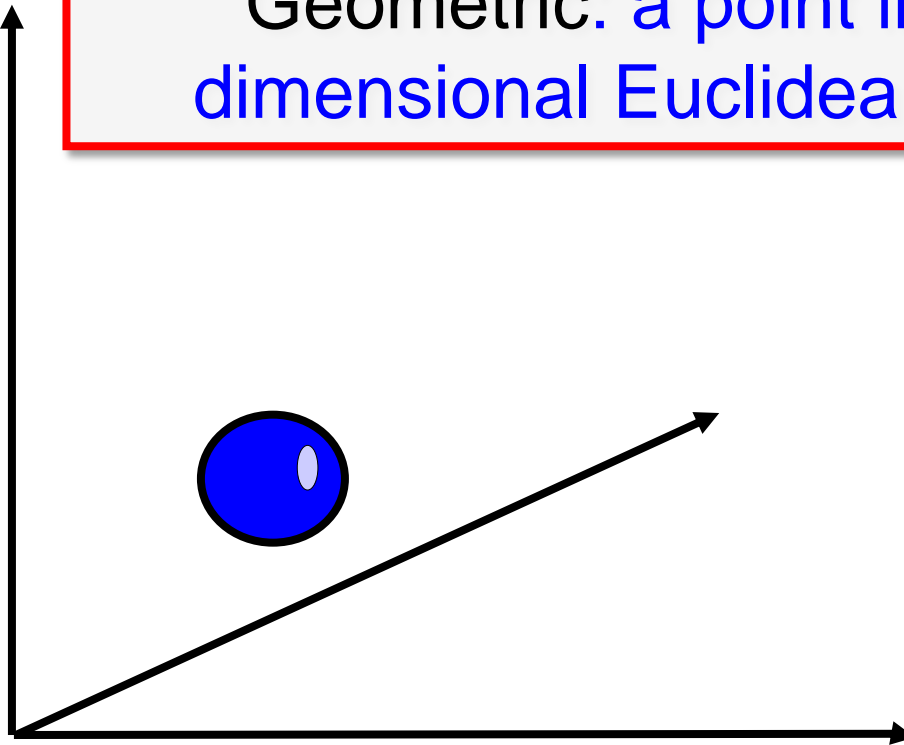
Crash Failure Solvability

Byzantine Failure Solvability

A Vertex

Combinatorial: an element of a set

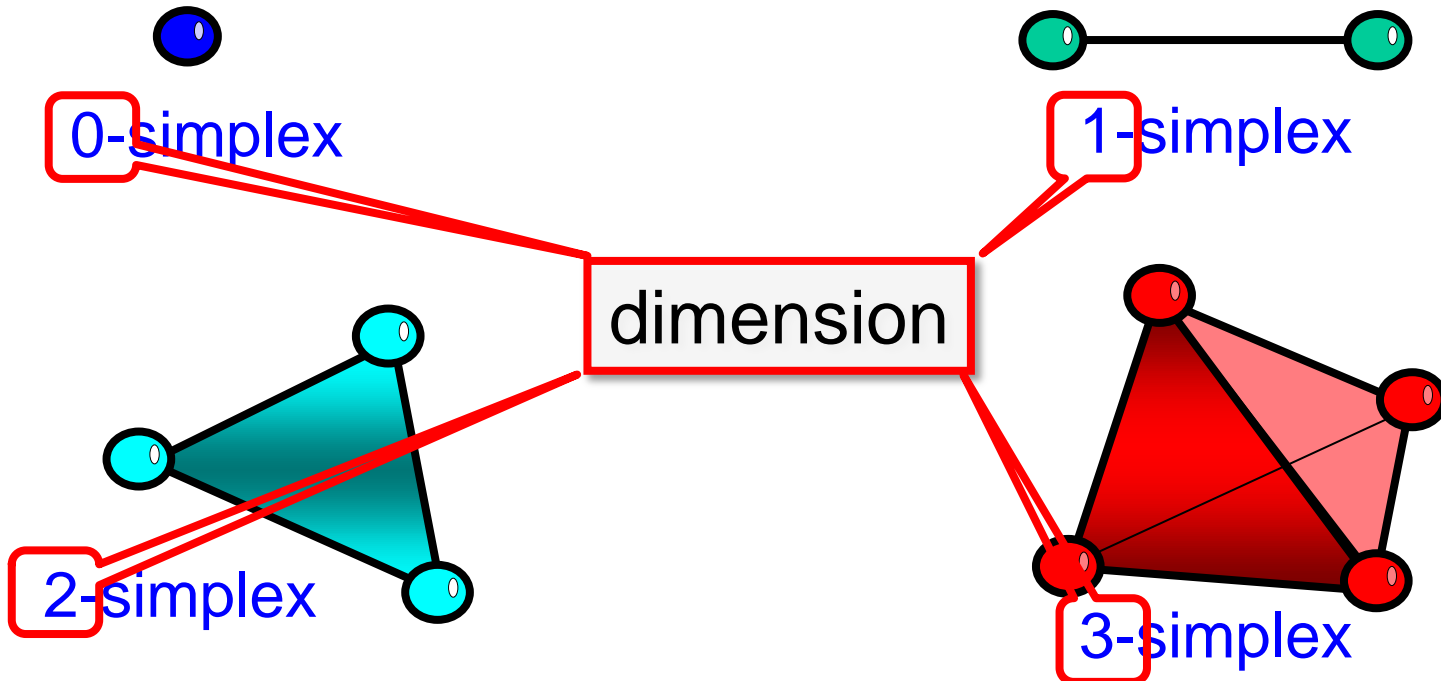
Geometric: a point in high-dimensional Euclidean Space



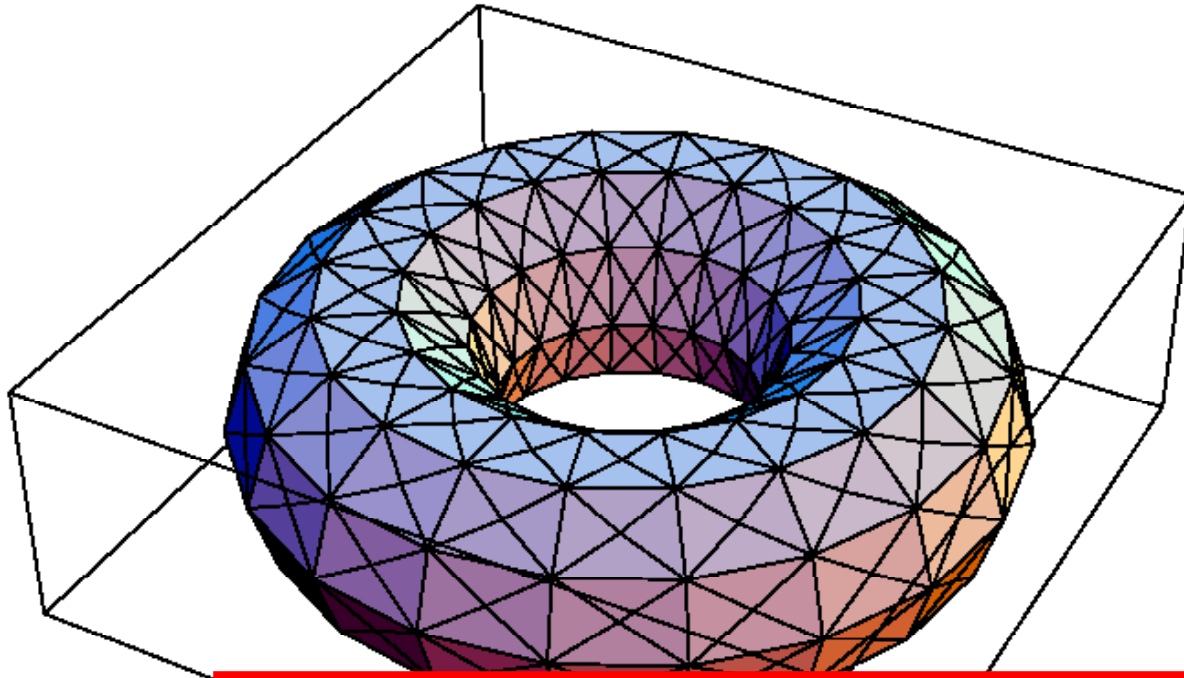
Simplexes

Combinatorial: a set of vertexes

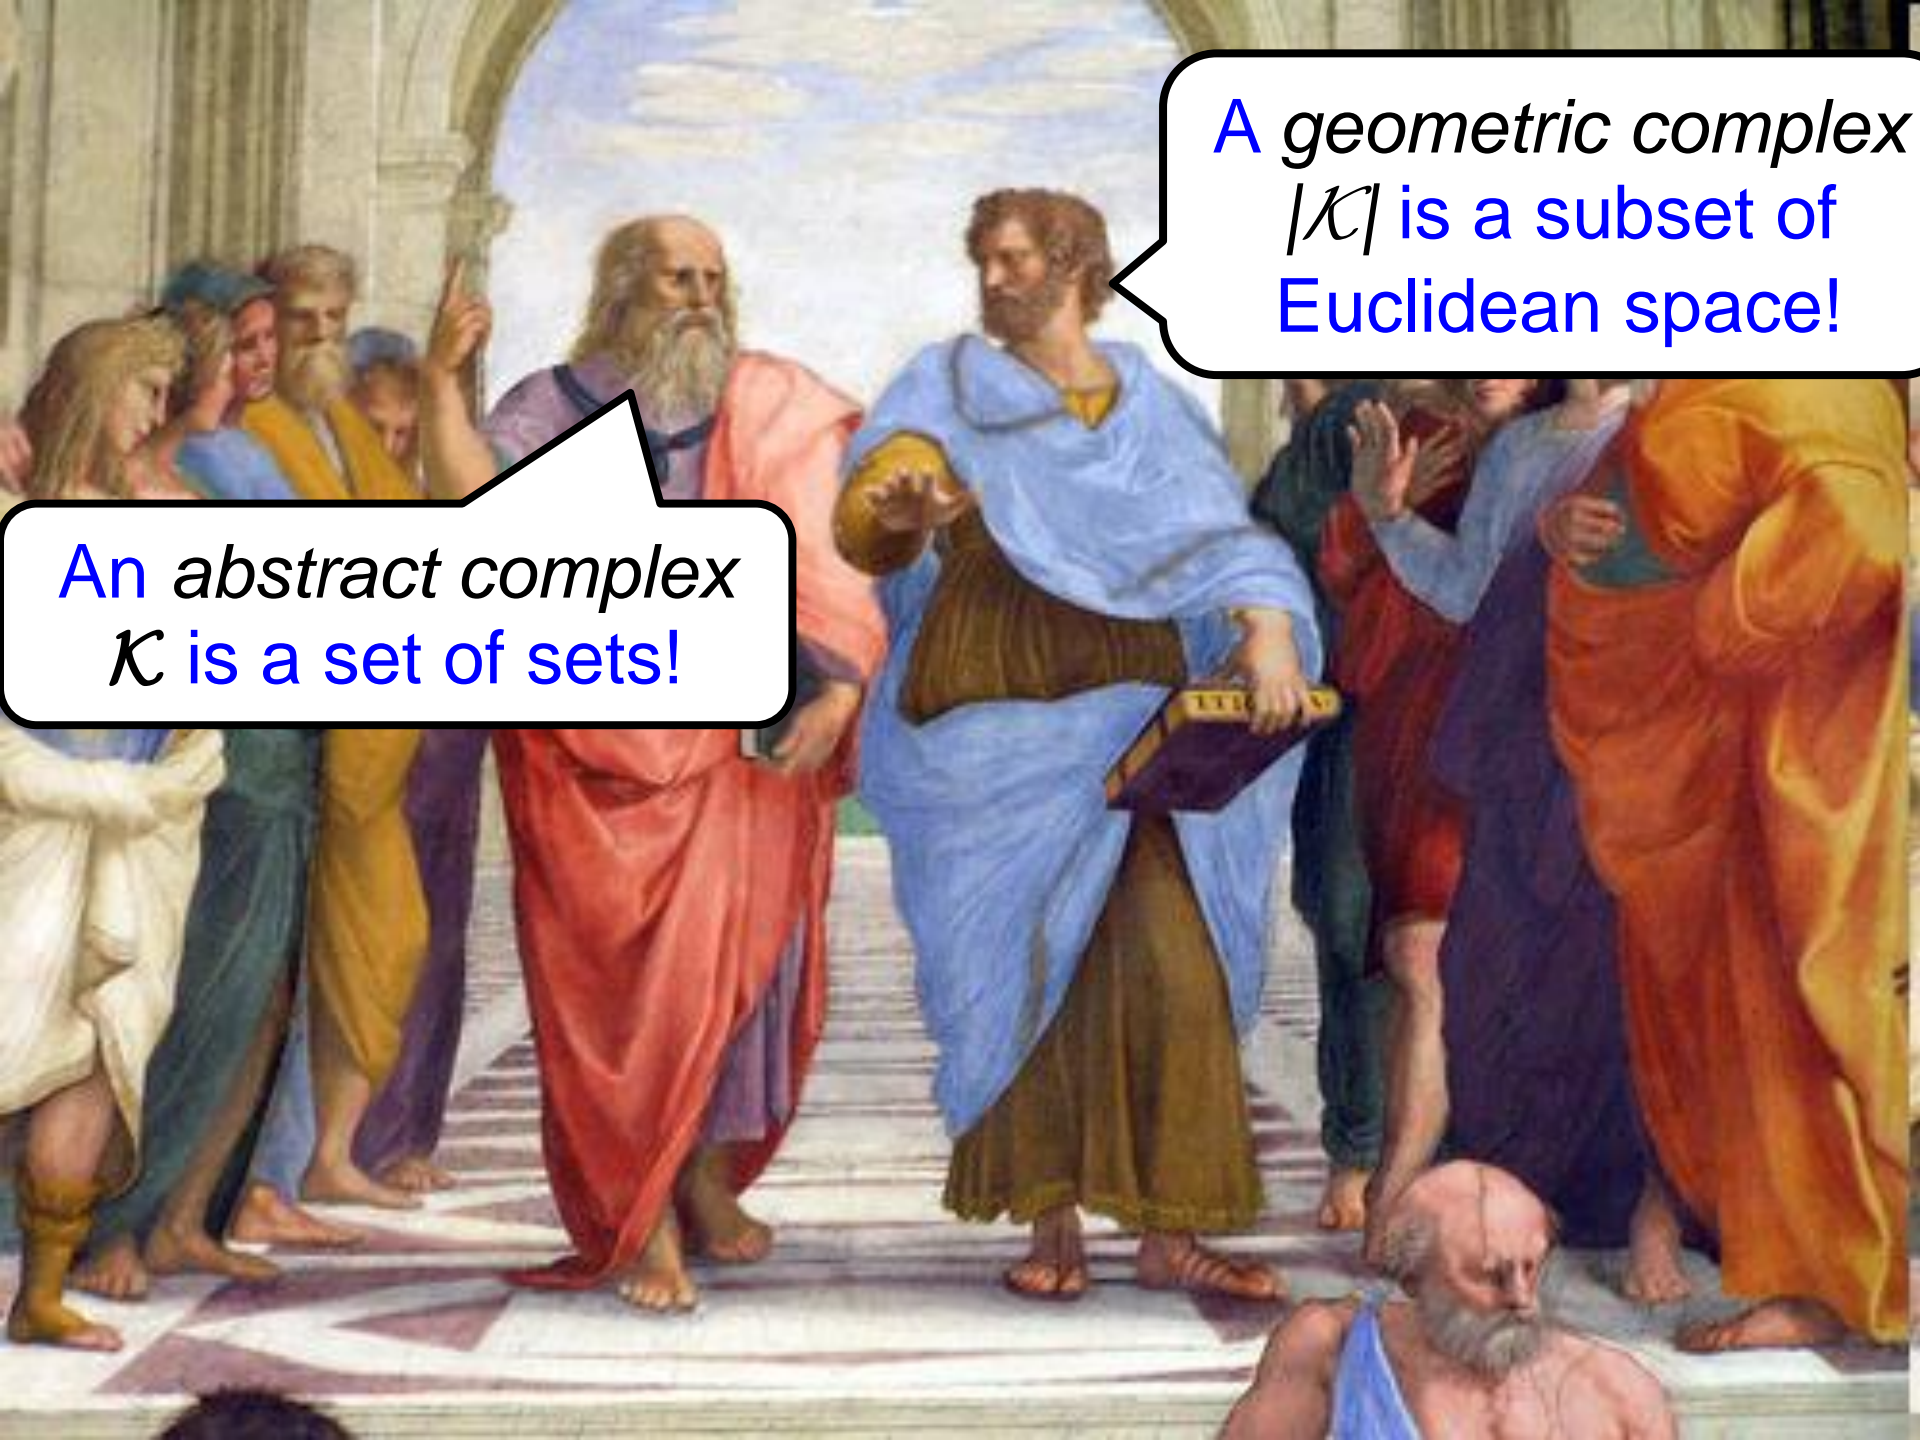
Geometric: convex hull of points in general position



Simplicial Complex



A set of simplexes closed under inclusion.




A geometric complex $|\mathcal{K}|$ is a subset of Euclidean space!

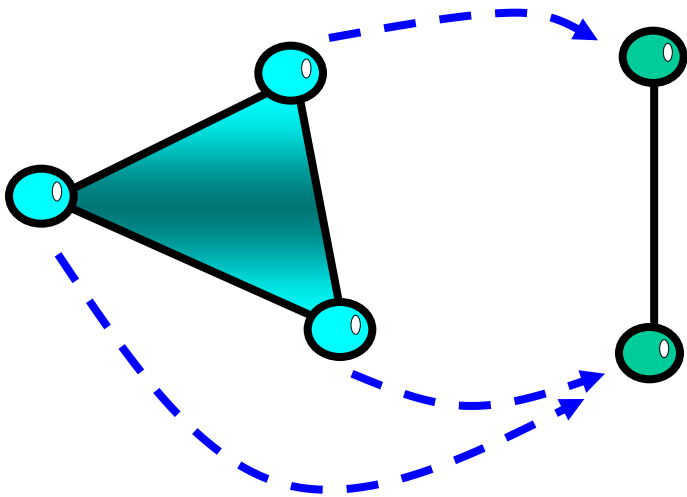
An abstract complex \mathcal{K} is a set of sets!

Simplicial Maps

Vertex-to-vertex map ...



Simplicial Map




Vertex-to-vertex map ...

that sends simplexes to
simplexes

piece-wise linear map
on geometric simplexes

Carrier Map



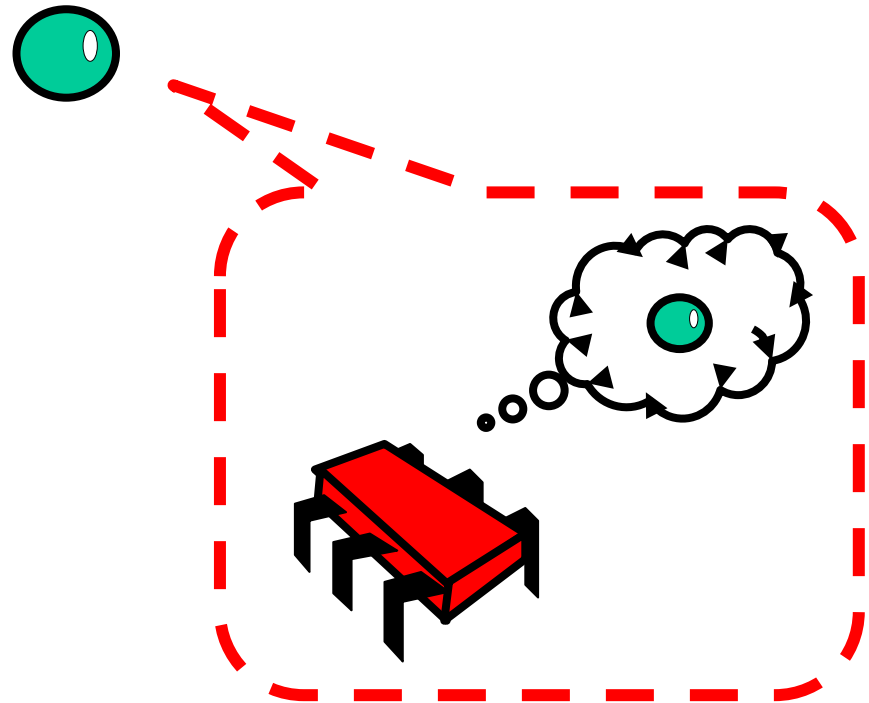
Maps simplex ...

to subcomplex.


Monotonic: if $\sigma \subset \tau$ then $\Phi(\sigma) \subset \Phi(\tau)$

Always OK to discard inputs


Vertex = Input or Output Value



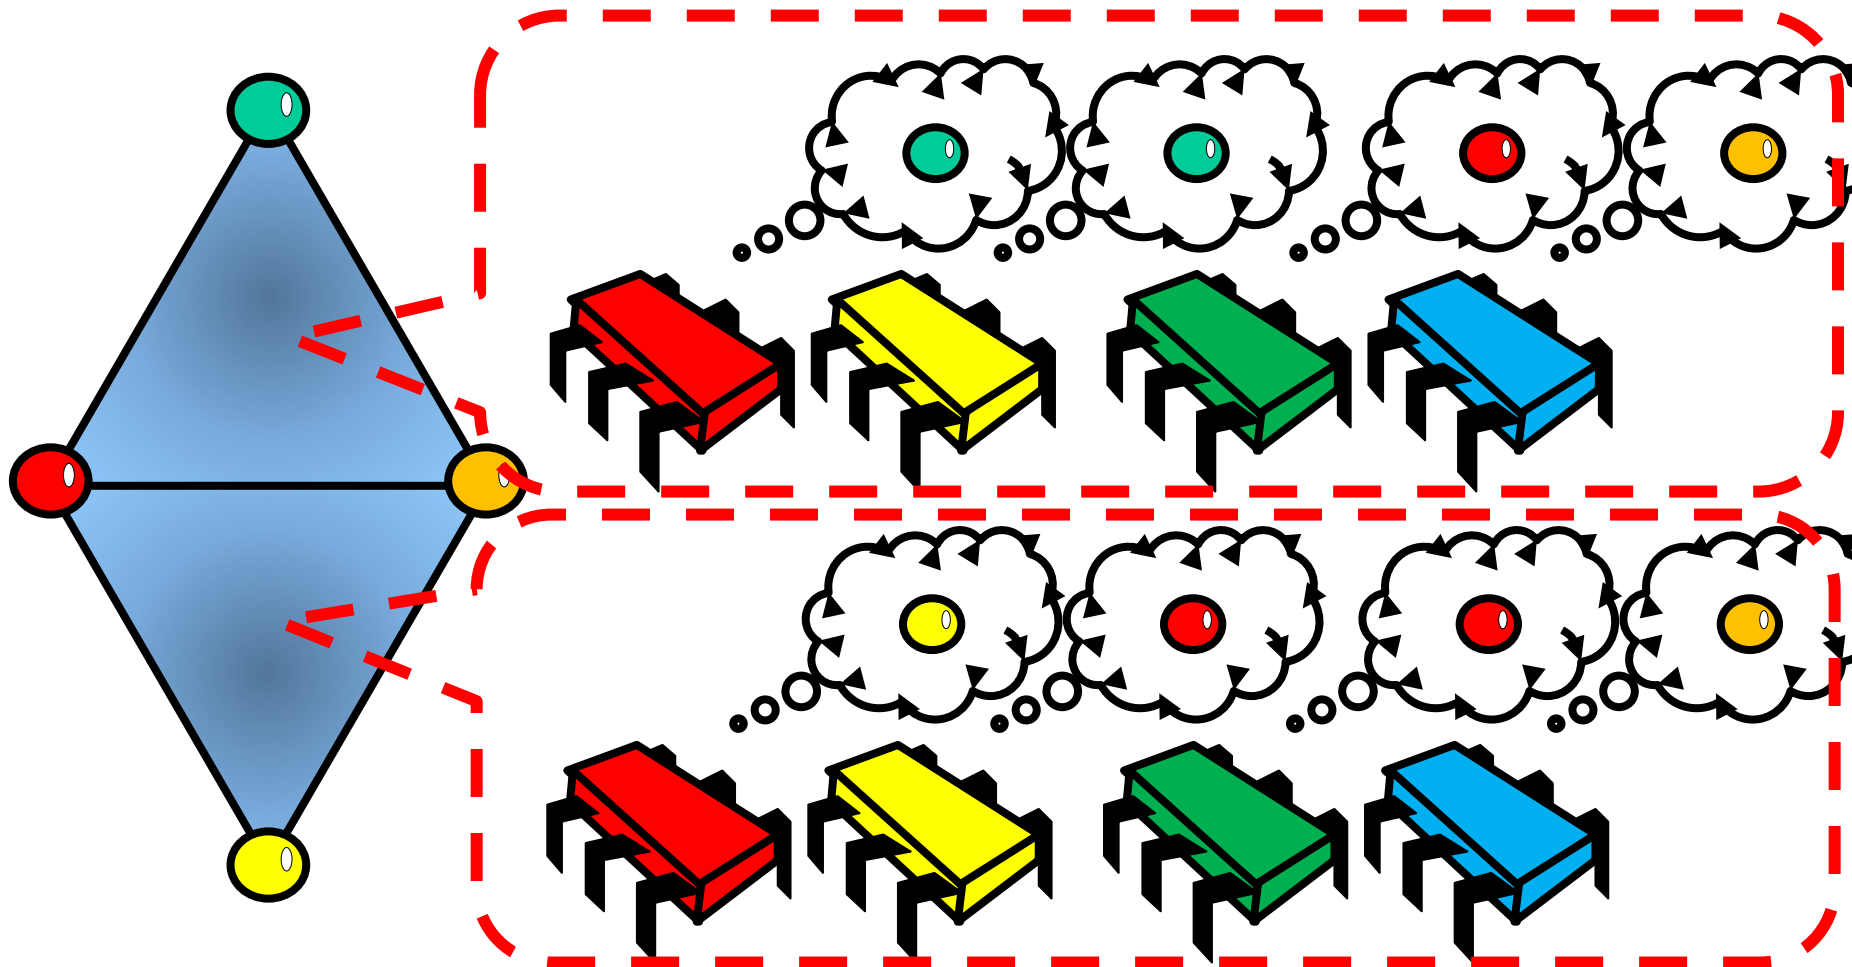
Vertex = Input or Output Value



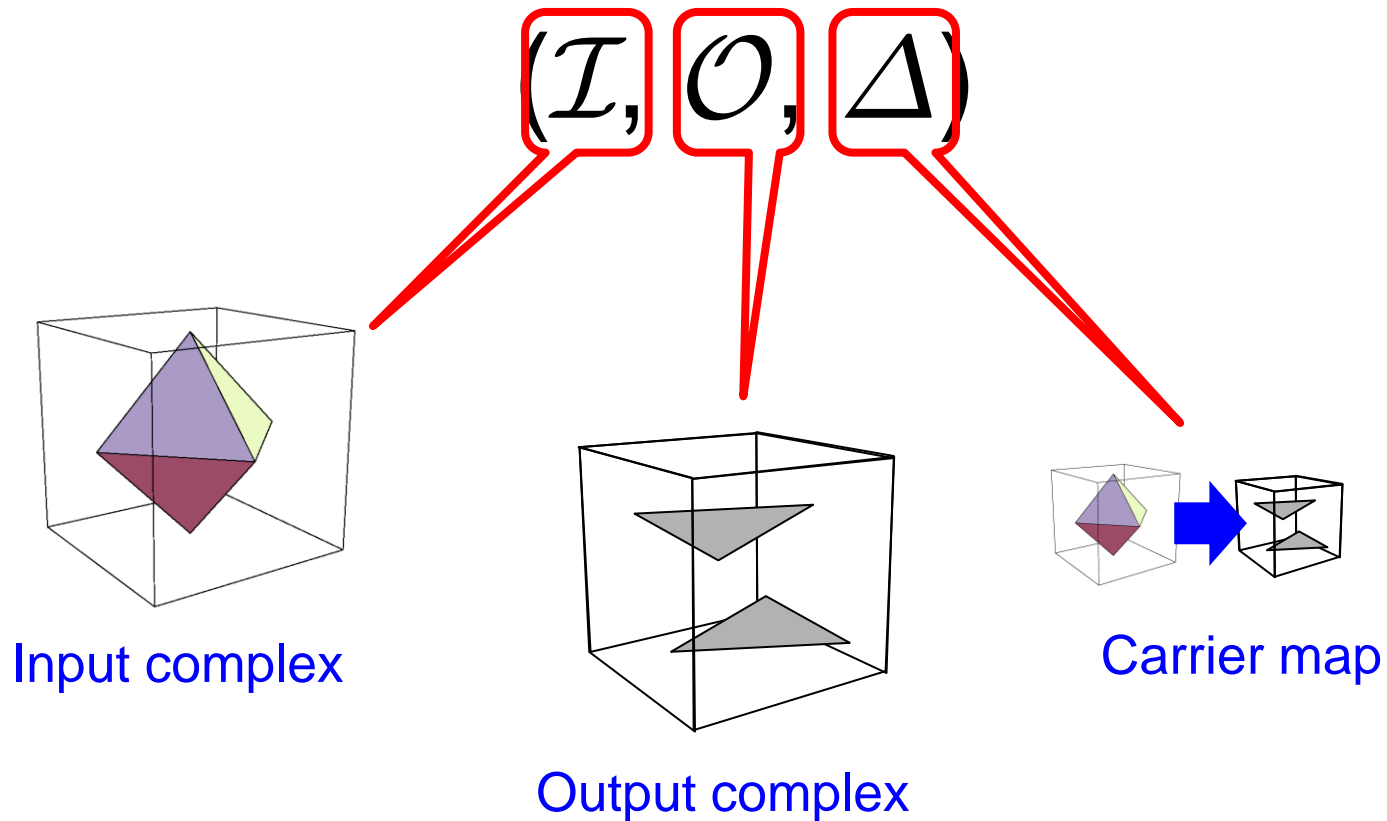
Simplex = Compatible Values




Simplex = Compatible Values




Task Specification



Consensus

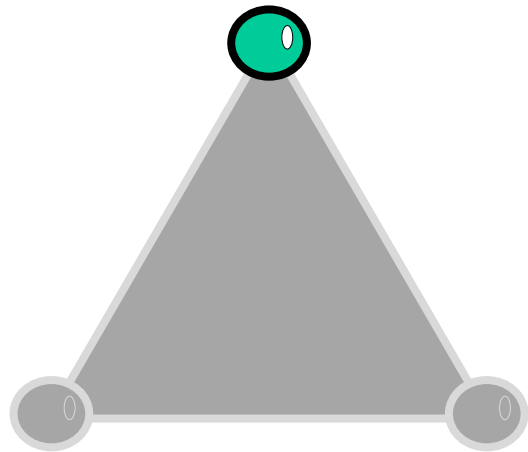


Input Complex





Output Complex

Carrier Map




Input Complex

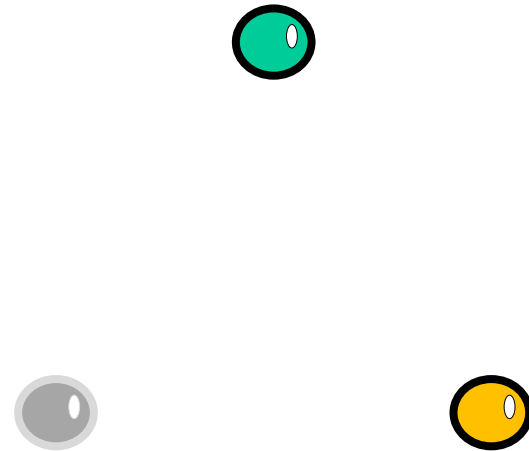



Output Complex

Consensus




Input Complex





Output Complex

Consensus



Input Complex



Output Complex

Protocol

```
view = my input value;  
for (i = 0; i < r; i++) {  
    broadcast view;  
    view += messages received;  
}  
return  $\delta$ (view)
```

Finite program

Protocol

```
view = my input value;  
for (i = 0; i < n; i++)  
    broadcast(view, i);  
    view += messages received;  
}  
return  $\delta$ (view)
```

Start with input value

Protocol

```
view = my input value;  
for (i = 0; i < r; i++) {  
    broadcast  
    view = messages received;  
}  
return  $\delta$ (view)
```

Run for fixed number of rounds

Protocol

```
view = my input value;  
for (i = 0; i < r; i++) {  
    broadcast view;  
    view += Send current view to others  
}  
return  $\delta$ (view)
```

Protocol

```
view = my input value;  
for (i = 0; i < r; i++) {  
    broadcast view;  
    view += messages received;  
}  
return  $\delta$ (view)
```

**Concatenate messages
received to view
(full-information protocol)**

Protocol

```
view = my input value;  
for (i = 0; i < r; i++) {  
    broadcast view;  
    view += messages received;  
}  
return  $\delta$ (view)
```

finally, apply task-specific
decision map to view

Protocol Complex

Vertex: possible view

Full information: messages
sent & received

Simplex: compatible set of views

Each execution defines a simplex

Road Map

Colorless Tasks

Operational Model

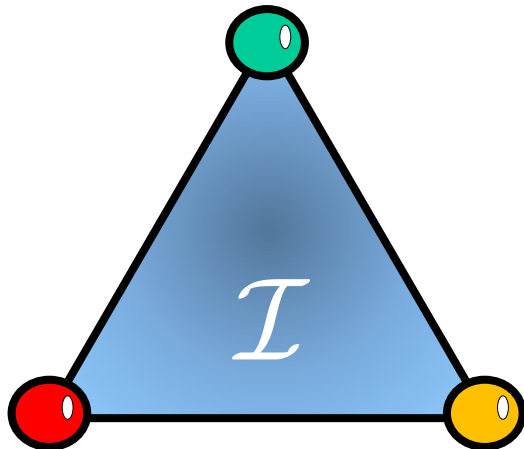
Combinatorial Model

Building Blocks

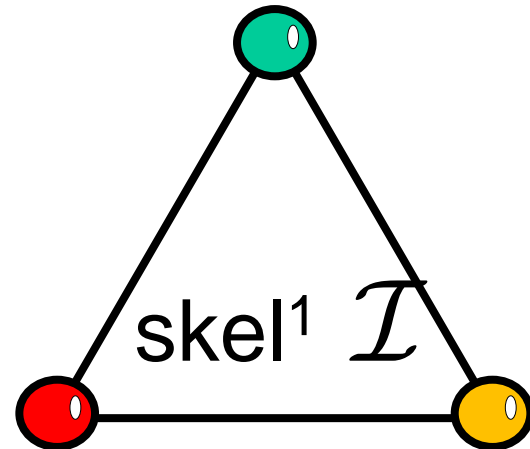
Crash Failure Solvability

Byzantine Failure Solvability

k -Set Agreement



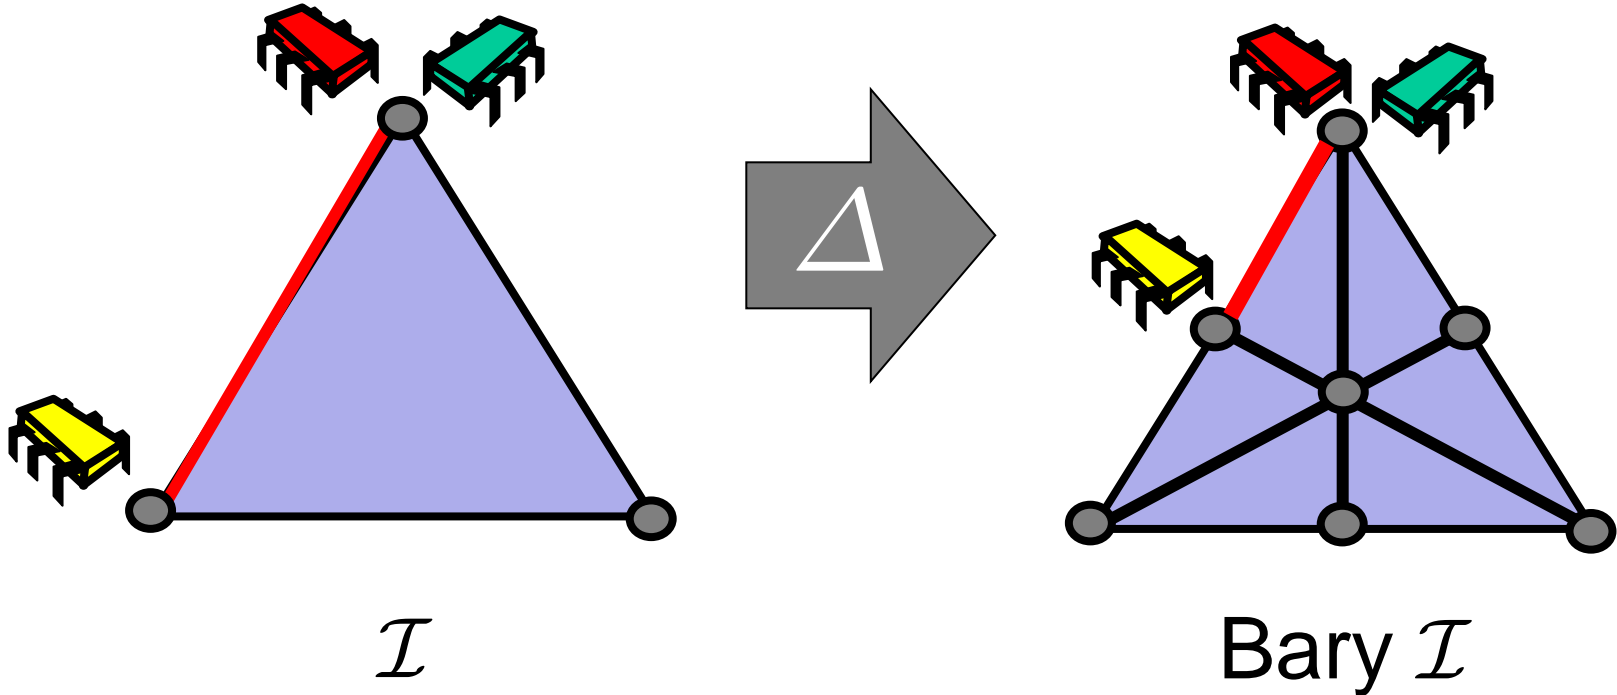
Input Complex



Output Complex

Δ is k -skeleton operator

Barycentric Agreement



Δ is barycentric subdivision operator

Closely related to snapshot

k -Set Agreement in Crash Failure Model

Theorem

There exists a t -resilient $(t+1)$ -set agreement protocol
 $(\mathcal{I}, \text{skel}^t \mathcal{I}, \text{skel}^t(\cdot))$

Proof

Broadcast value, wait for all but t
values, decide least one.

k -Set Agreement in Byzantine Failure Model

Theorem

There exists a t -resilient $(t+1)$ -set agreement protocol
($\mathcal{I}, \text{skel}^t \mathcal{I}, \text{skel}^t(\cdot)$)
Iff $n+1 > t \cdot (\dim \mathcal{I} + 2)$

new!

Cost of Byzantine failures!

The Necessary Part

Byzantine processes cannot influence decisions!

Non-Faulty processes cannot “believe” value with $< t+1$ witnesses

If $n+1 > t \cdot (\dim \mathcal{I} + 2)$ then some value has at least $t+1$ witnesses

The Sufficient Part

Variation of *reliable broadcast* protocol of [Bracha 87] and [Shrikanth & Toueg 87]

Non-Faulty processes agree on values sent by others, even faulty processes.

If one non-faulty process receives a message, so do the others (liveness)

Barycentric Agreement in Crash Failure Model

Theorem

There is a t -resilient barycentric agreement protocol
 $(\mathcal{I}, \text{Bary}^N \mathcal{I}, \text{Bary}^N(\cdot))$

Proof

Variation of *stable vectors* algorithm of
[Attiya et al. 90]

Barycentric Agreement in Byzantine Failure Model

Theorem

There is a t -resilient barycentric agreement protocol

$(\mathcal{I}, \text{bary}^t \mathcal{I}, \text{bary}^t(\cdot))$

Iff $n+1 > t \cdot (\dim \mathcal{I} + 2)$

new!

Cost of Byzantine failures!

The Necessary Part

Byzantine processes cannot influence decisions!

Non-Faulty processes cannot “believe” value with $< t+1$ witnesses

If $n+1 > t \cdot (\dim \mathcal{I} + 2)$ then some value has at least $t+1$ witnesses

The Sufficient Part

Byzantine variation of *stable vectors*
algorithm of [Attiya et al. 90]

Use reliable broadcast to spread values

Ignore values with fewer than $t+1$ witnesses ...

Road Map

Colorless Tasks

Operational Model

Combinatorial Model

Building Blocks

Crash Failure Solvability

Byzantine Failure Solvability

Solvability for Crash Failures

Theorem

There is a t -resilient protocol for task
 $(\mathcal{I}, \mathcal{O}, \Delta)$

Iff there is a continuous map

$$f: |\text{skel}^t \mathcal{I}| \rightarrow |\mathcal{O}|$$

carried by Δ

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

The Necessary Part

We can model the set of process views after a full-information protocol as a *protocol complex* \mathcal{P}

There is a simplicial decision map

$$\delta: \mathcal{P} \rightarrow \mathcal{O} \text{ carried by } \Delta$$

WLOG, \mathcal{P} is isomorphic to $\text{bary}^N \text{skel}^t \mathcal{I}$

$$\delta: \text{bary}^N \text{skel}^t \mathcal{I} \rightarrow \mathcal{O}$$

induces a piece-wise linear map

$$|\delta|: |\text{bary}^N \text{skel}^t \mathcal{I}| \rightarrow |\mathcal{O}|$$

$$|\delta|: |\text{skel}^t \mathcal{I}| \rightarrow |\mathcal{O}|$$

carried by Δ

The Sufficient Part

$$f: |\text{skel}^t \mathcal{I}| \rightarrow |\mathcal{O}|$$

has a simplicial approximation for some $N > 0$

$$\phi: \text{bary}^N \text{skel}^t \mathcal{I} \rightarrow \mathcal{O}$$

Step 1: use t -set agreement protocol to go from vertex of \mathcal{I} to vertex of $\text{skel}^t \mathcal{I}$

Step 2: use repeated barycentric agreement to go from vertex of $\text{skel}^t \mathcal{I}$ to vertex of $\text{bary}^N \text{skel}^t \mathcal{I}$

Step 3: from vertex $v \in \text{bary}^N \text{skel}^t \mathcal{I}$, decide $\phi(v)$

Road Map

Colorless Tasks

Operational Model

Combinatorial Model

Building Blocks

Crash Failure Solvability

Byzantine Failure Solvability

Solvability for Byzantine Failures

Theorem

There is a t -resilient protocol for task
 $(\mathcal{I}, \mathcal{O}, \Delta)$

Iff there is a continuous map

$$f: |\text{skel}^t \mathcal{I}| \rightarrow |\mathcal{O}|$$

carried by Δ

$$\text{Iff } n+1 > t \cdot (\dim \mathcal{I} + 2)$$

new!

Cost of Byzantine failures!

The Necessary Part

We can model the set of process views after a full-information protocol as a *protocol complex* \mathcal{P}

There is a simplicial decision map

$$\delta: \mathcal{P} \rightarrow \mathcal{O} \text{ carried by } \Delta$$

WLOG, \mathcal{P} is isomorphic to $\text{bary}^N \text{skel}^t \mathcal{I}$

$$\delta: \text{bary}^N \text{skel}^t \mathcal{I} \rightarrow \mathcal{O}$$

induces a piece-wise linear map

$$|\delta|: |\text{bary}^N \text{skel}^t \mathcal{I}| \rightarrow |\mathcal{O}|$$

$$|\delta|: |\text{skel}^t \mathcal{I}| \rightarrow |\mathcal{O}|$$

carried by Δ

The Necessary Part

We can model the set of process views after a full-information protocol as a *protocol complex* \mathcal{P}

There is a simplicial decision map $\delta: \mathcal{P} \rightarrow \mathcal{O}$ carried by Δ

Under the assumption that \mathcal{P} is isomorphic to $\text{skel}^t \text{bary}^N \mathcal{I}$

Technical details more involved

δ induces a piece-wise linear map $|\delta|: |\mathcal{P}| \rightarrow |\mathcal{O}|$

$|\delta|: |\text{skel}^t \text{bary}^N \mathcal{I}| \rightarrow |\mathcal{O}|$

$|\delta|: |\text{skel}^t \mathcal{I}| \rightarrow |\mathcal{O}|$

carried by Δ

The Sufficient Part

$$f: |\text{skel}^t \mathcal{I}| \rightarrow |\mathcal{O}|$$

has a simplicial approximation for some $N > 0$

$$\phi: \text{bary}^N \text{skel}^t \mathcal{I} \rightarrow \mathcal{O}$$

Step 1: use t -set agreement protocol to go from vertex of \mathcal{I} to vertex of $\text{skel}^t \mathcal{I}$

Step 2: use repeated barycentric agreement to go from vertex of $\text{skel}^t \mathcal{I}$ to vertex of $\text{bary}^N \text{skel}^t \mathcal{I}$

Step 3: from vertex $v \in \text{bary}^N \text{skel}^t \mathcal{I}$, decide $\phi(v)$

The Sufficient Part

$$f: |\text{skel}^t \mathcal{I}| \rightarrow |\mathcal{O}|$$

has a *simplicial approximation* for some $N > 0$

$$\text{bary}^N \mathcal{I} \rightarrow \mathcal{O}$$

Basic Idea is the same

Step 1: use t -set agreement protocol to
go from vertex of \mathcal{I} to vertex of $\text{skel}^t \mathcal{I}$

Step 2: use repeated t -set agreement
to go from vertex of $\text{skel}^t \mathcal{I}$
to vertex of $\text{skel}^t \text{bary}^N \mathcal{I}$

Technical details more involved

Step 3: from vertex $v \in \text{skel}^t \text{bary}^N \mathcal{I}$,
decide $\phi(v)$

Conclusions

Many have looked at specific tasks ...

consensus

approximate agreement

k -set agreement

Often with much weaker validity!

First to look at general (colorless) tasks ...

First to characterize what can and can't be solved

Conclusions

The *language* of combinatorial topology (*vertex, simplex, skeleton, simplicial map ...*) allows us to *state and prove* such results succinctly

Important to exploit the duality of combinatorial and continuous model (such as *simplicial approximation*)

Here, we did not need “advanced” concepts like connectivity, but they are needed elsewhere, such as the synchronous model ...

Open Problems

Colored tasks?

Complexity?

“Rational” adversaries?

Mechanism design?

Long-lived computations?

Randomized?

COMBINATORIAL
TOPOLOGY
&
DISTRIBUTED
COMPUTING



MK
MOSKVA UNIVERSITY

*Maurice Herlihy,
Dmitry Feichtner-Kozlov, Sergio Rajsbaum*