

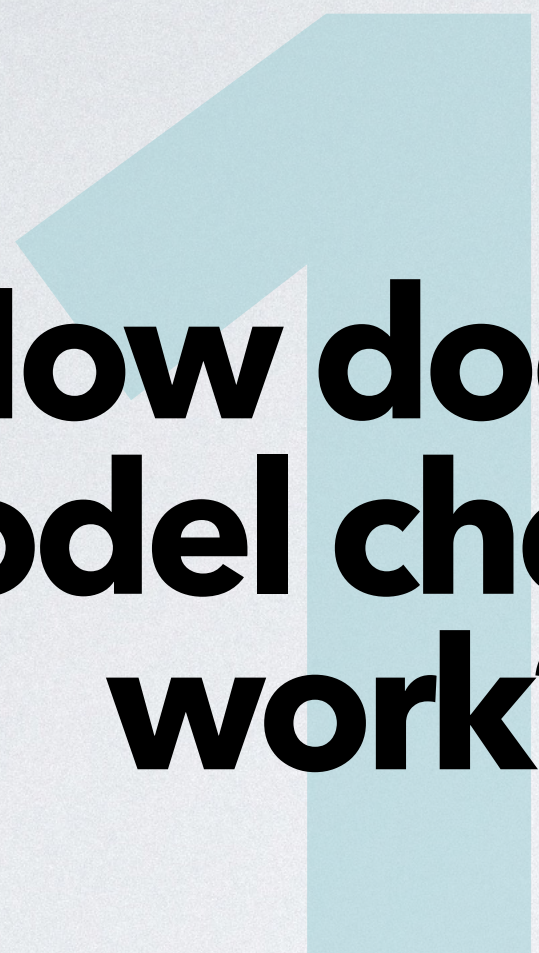
Pragmatic model checking: from theory to implementations

Niels Lohmann


Universität
Rostock



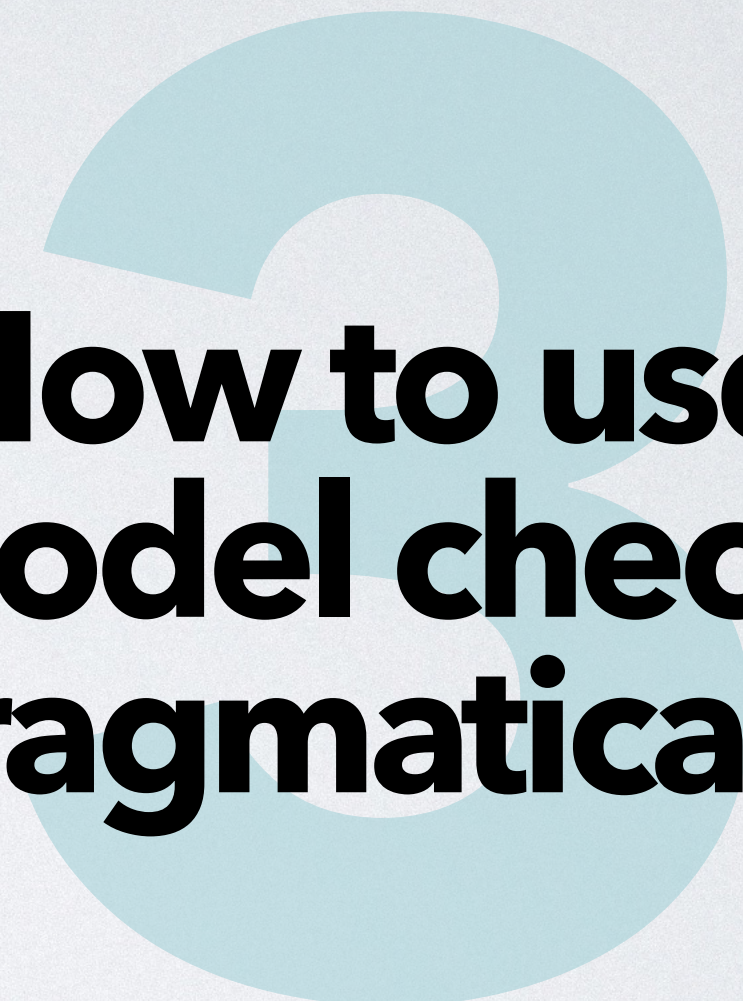
Traditio et Innovatio



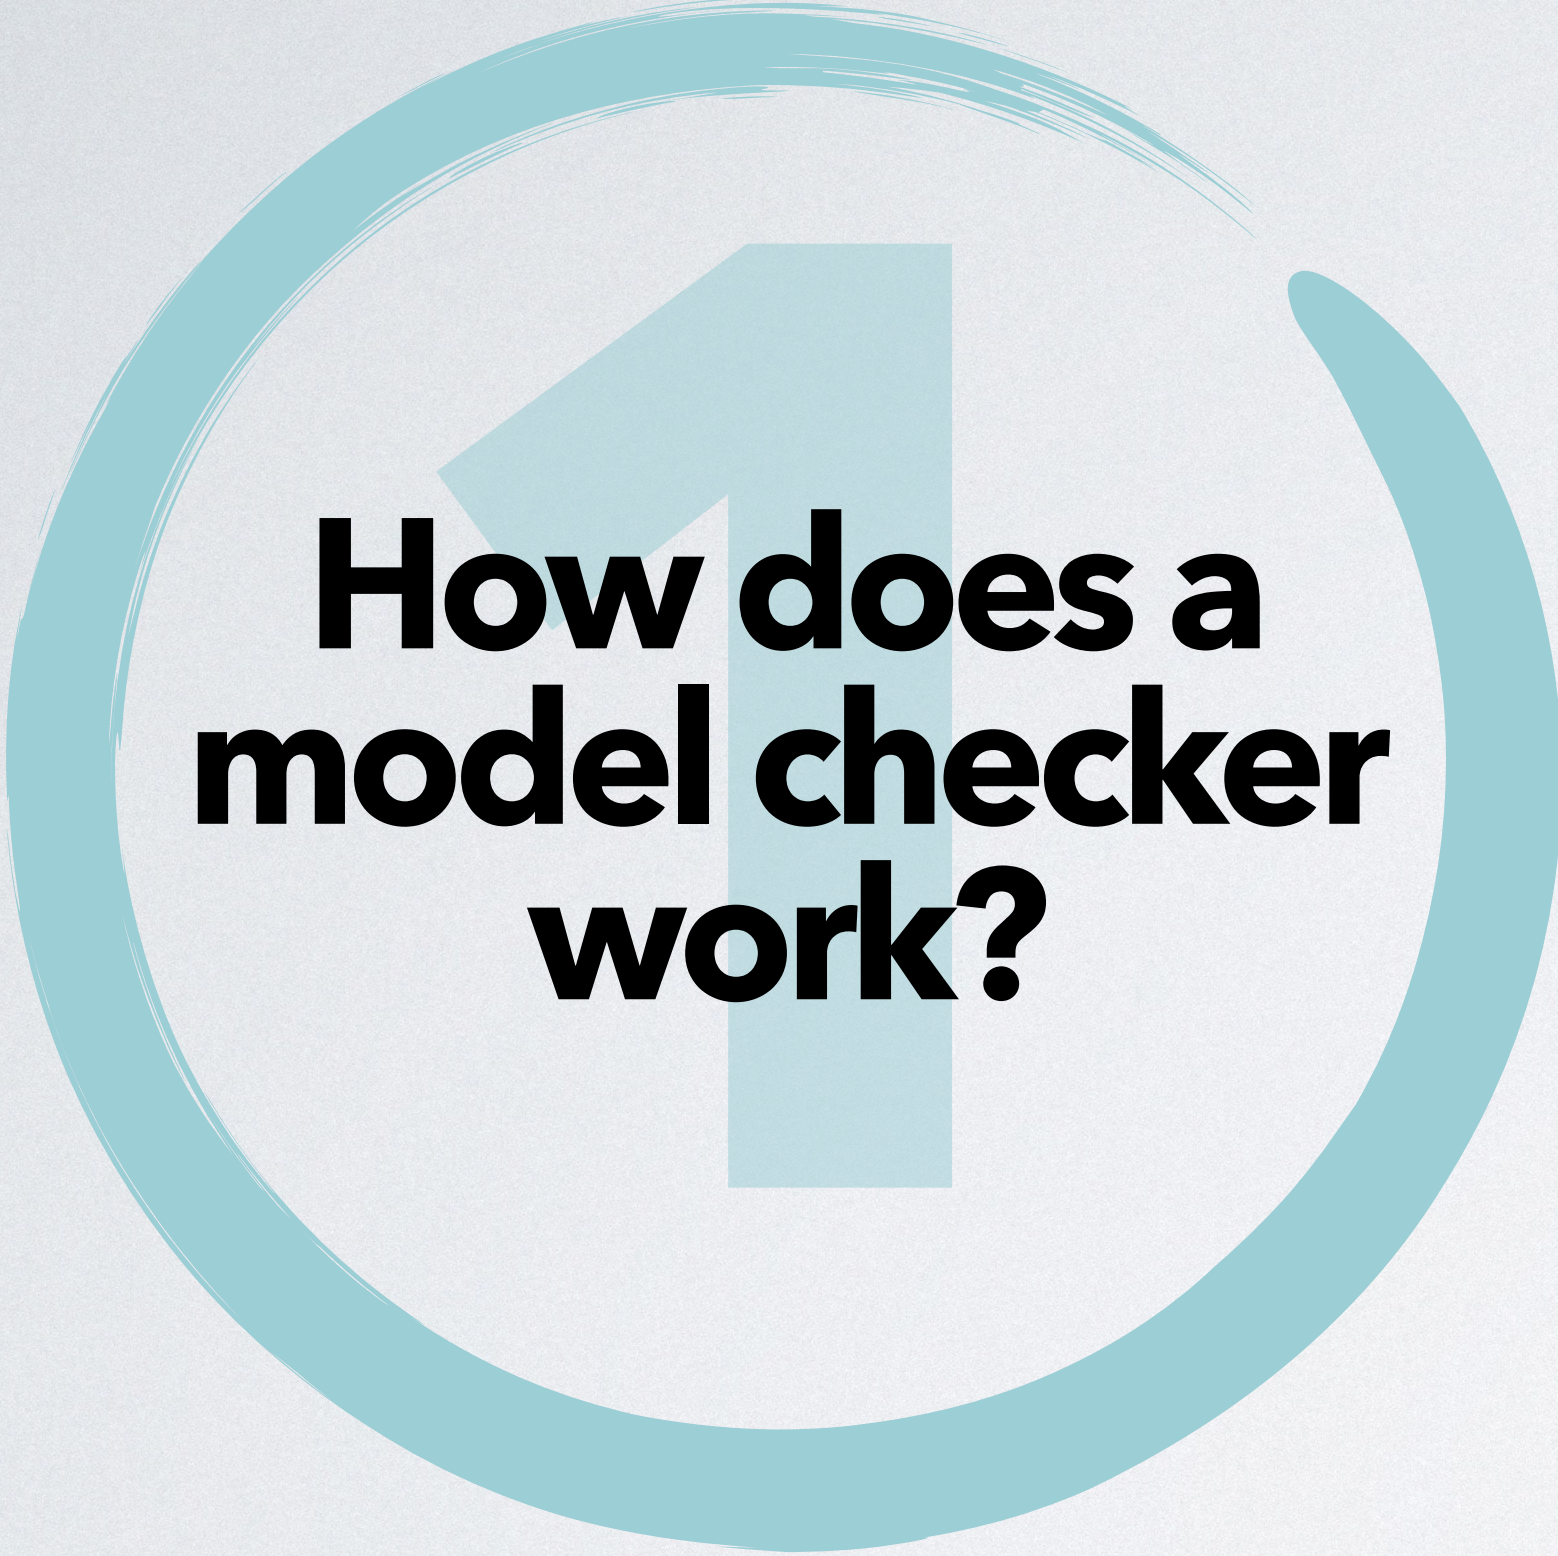
**How does a
model checker
work?**



**How to build an
effective model
checker?**



**How to use a
model checker
pragmatically?**

A large teal circle with a thick, hand-drawn style border. Inside the circle is a large, light teal number '1'.

**How does a
model checker
work?**

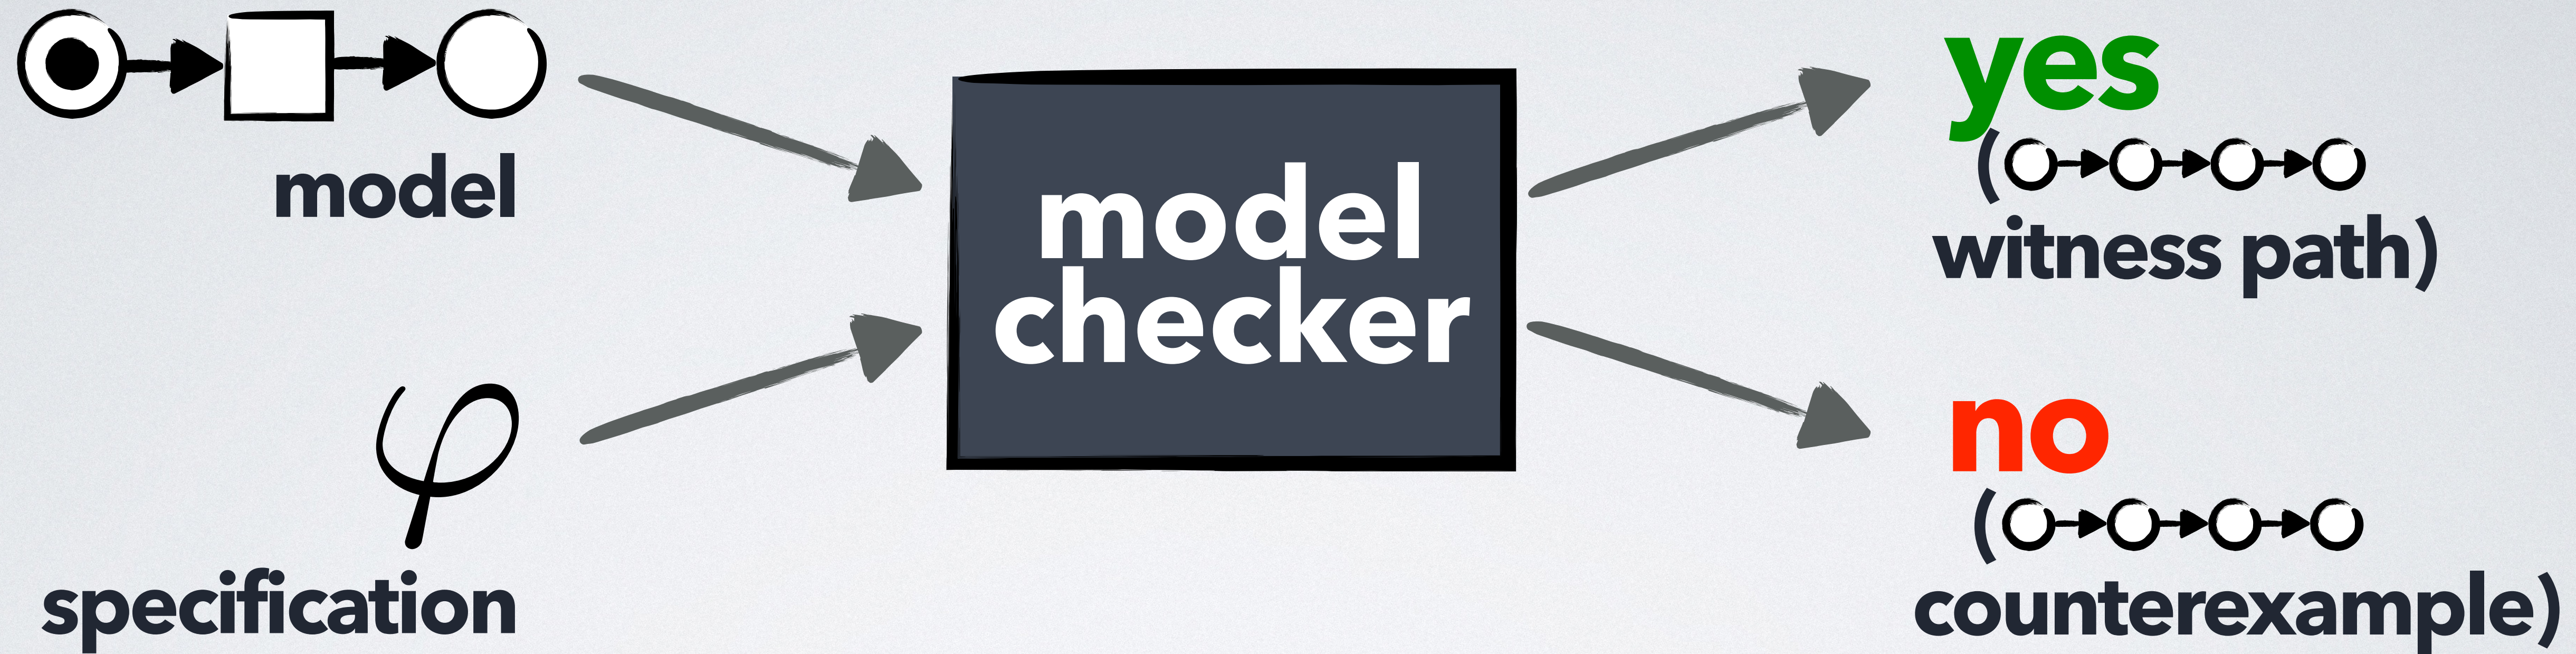
A large, light teal number '2' with a thick, hand-drawn style border.

**How to build an
effective model
checker?**

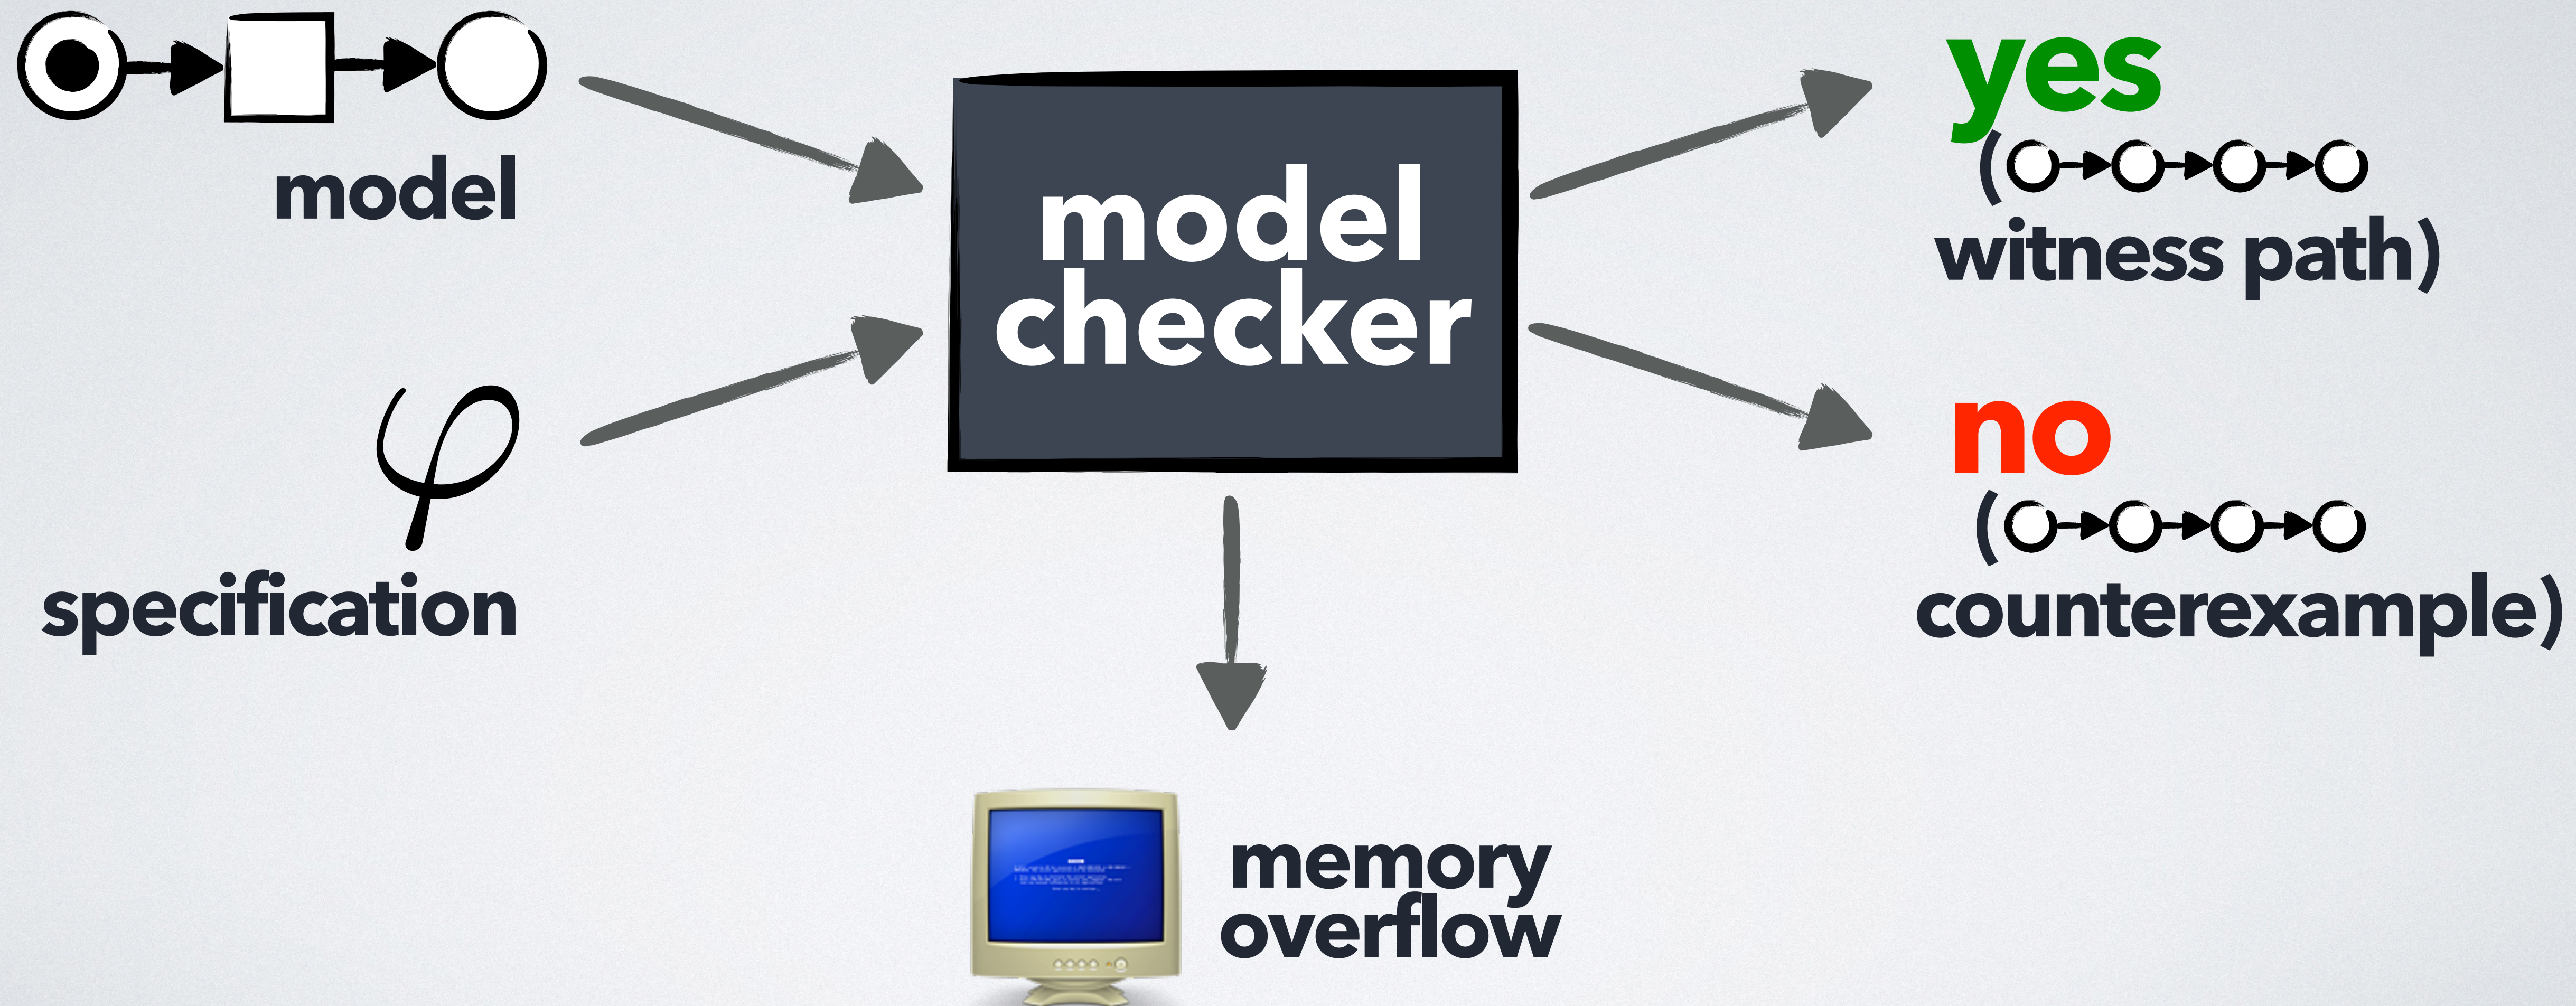
A large, light teal number '3' with a thick, hand-drawn style border.

**How to use a
model checker
pragmatically?**

Model checking in a nutshell



Model checking in a nutshell



State explosion

business process* with 66 parallel branches * **modeled by IBM customers using the IBM Websphere business modeler**

State explosion

business process* with 66 parallel branches

*** modeled by IBM
customers using the
IBM Websphere
business modeler**

state space:

$2^{66} \approx 2.37 \cdot 10^{19}$ markings

State explosion

business process* with 66 parallel branches

*** modeled by IBM
customers using the
IBM Websphere
business modeler**

state space:

$2^{66} \approx 2.37 \cdot 10^{19}$ **markings**

memory for state space:

608 exabytes** (66 bit per marking)

**** mega, giga, peta,
tera, exa...**

State explosion

business process* with 66 parallel branches * **modeled by IBM customers using the IBM Websphere business modeler**

state space:

$2^{66} \approx 2.37 \cdot 10^{19}$ **markings**

memory for state space:

608 exabytes** (66 bit per marking)

** **mega, giga, peta, tera, exa...**

time for state space generation

7799 years (10 CPU cycles/marking at 3 GHz)

State explosion

business process* with 66 parallel branches

*** modeled by IBM
customers using the
IBM Websphere
business modeler**

state space:

$2^{66} \approx 2.37 \cdot 10^{19}$ **markings**

memory for state space:

608 exabytes** (66 bit per marking)

**** mega, giga, peta,
tera, exa...**

time for state space generation

7799 years (10 CPU cycles/marking at 3 GHz)

energy consumption for state space generation

2,9 megatons of TNT (at 50 watts)

State explosion

business process* with 66 parallel branches

*** modeled by IBM
customers using the
IBM Websphere
business modeler**

state space:

$2^{66} \approx 2.37 \cdot 10^{19}$ **markings**

memory for state space:

608 exabytes** (66 bit per marking)

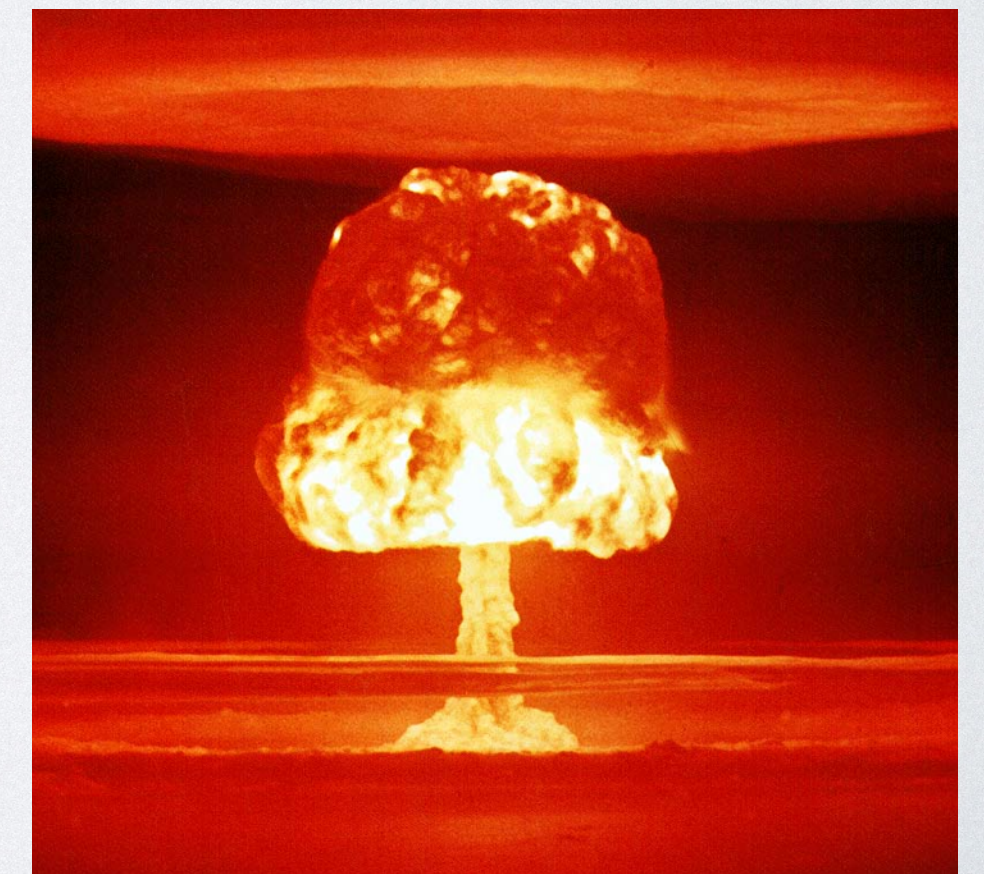
**** mega, giga, peta,
tera, exa...**

time for state space generation

7799 years (10 CPU cycles/marking at 3 GHz)

energy consumption for state space generation

2,9 megatons of TNT (at 50 watts)



The core algorithm

```
markings = []  
search( $m_0, \varphi$ )
```

```
1 def search( $m, \varphi$ ):  
2     check( $m, \varphi$ )  
3     markings.add( $m$ )  
4     for  $t$  in enabled( $m$ ):  
5          $m' = \text{fire}(m, t)$   
6         if not  $m'$  in markings:  
            search( $m', \varphi$ )
```

- **search** is a simple depth first search (+ SCC detection)
- **check** depends on the nature of the property and may terminate search
- **enabled** and **fire** implement the Petri net firing rule

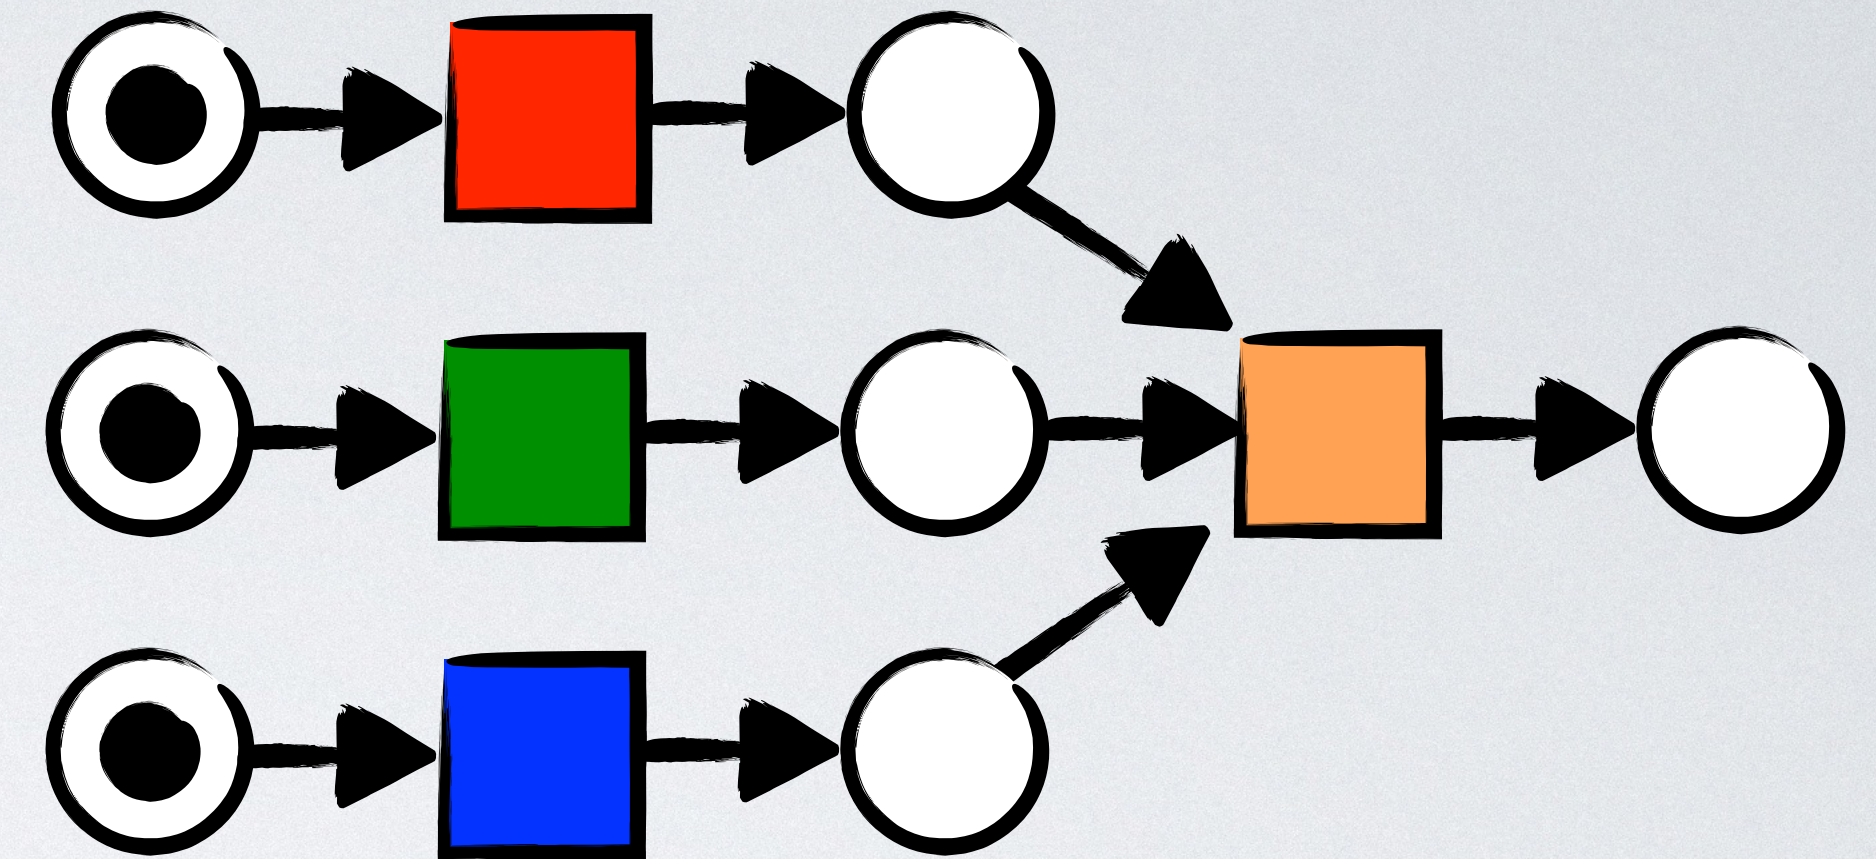
Reduction techniques

apply **theoretic** results to

- 1. store fewer markings
 - 2. fire fewer transitions /
generate fewer markings
 - 3. store markings more efficiently
- } while **preserving the property**

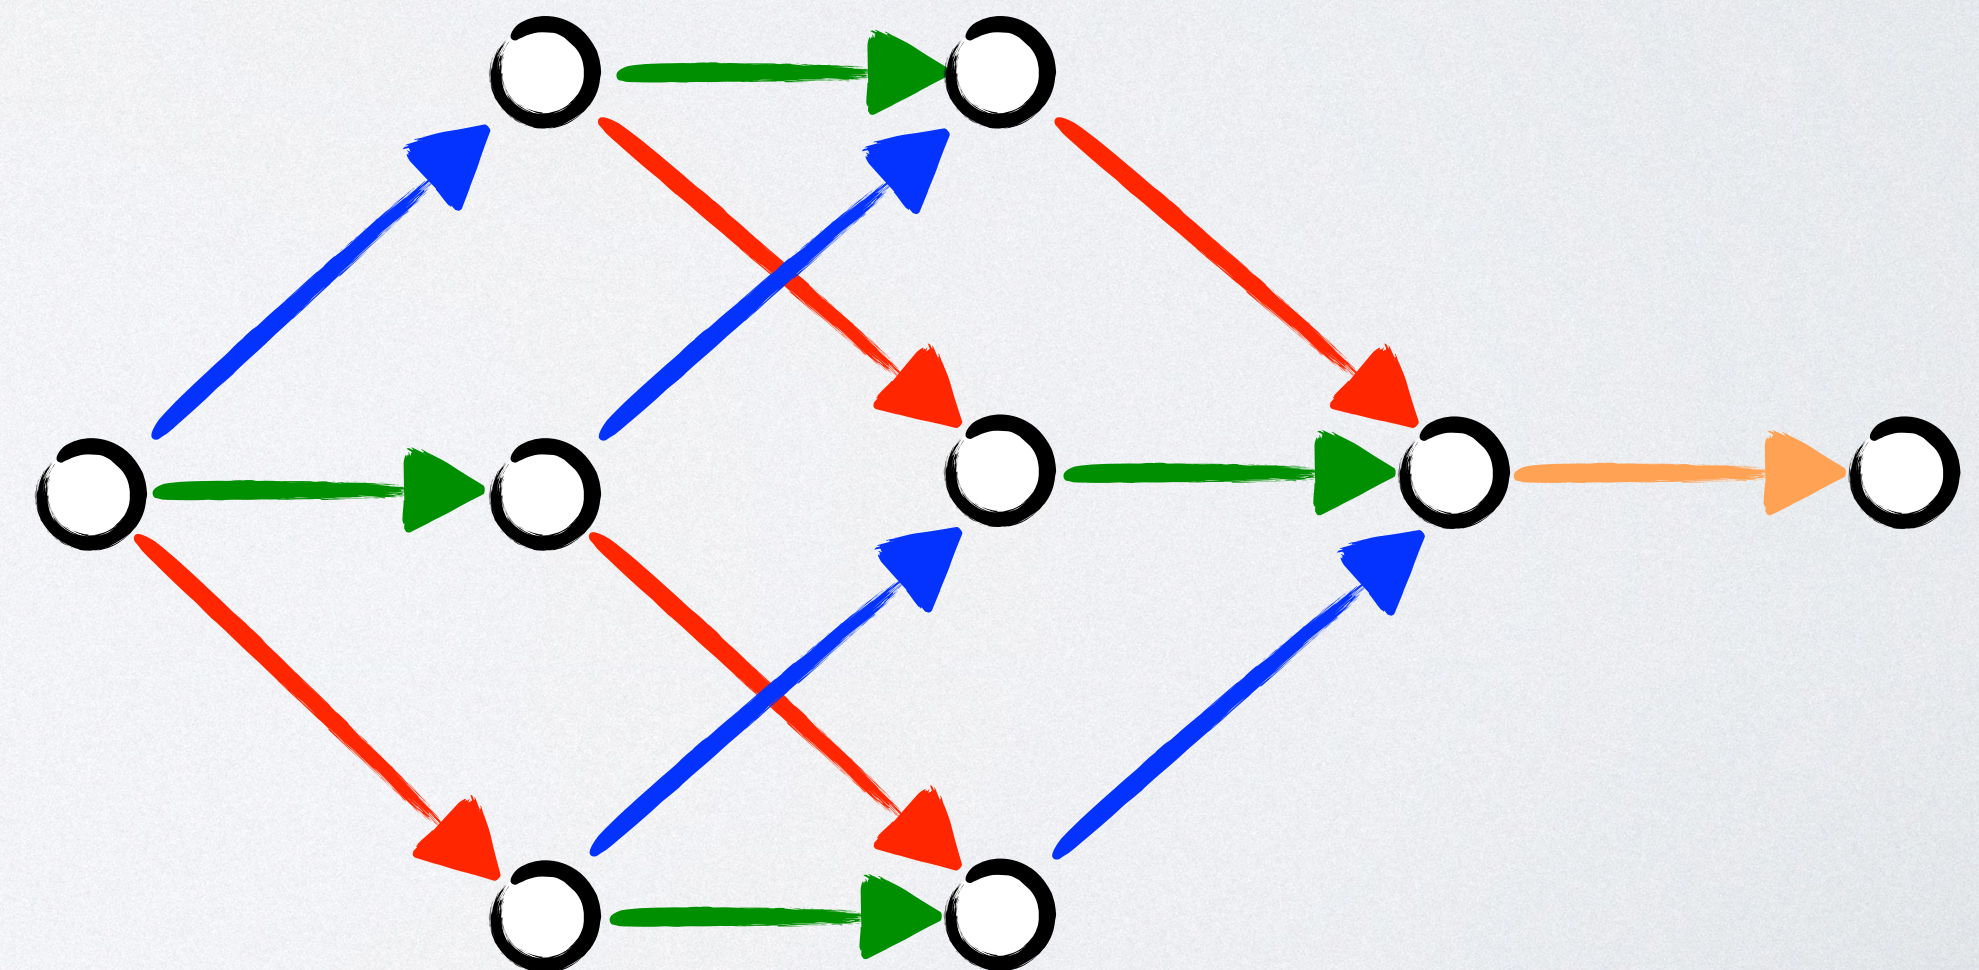
1. Partial order reduction

observation: concurrent transitions can be fired in **any order**, yielding the same final state



idea: fix one ordering by **postponing** the firing of some transitions

implementation: search on Petri net structure

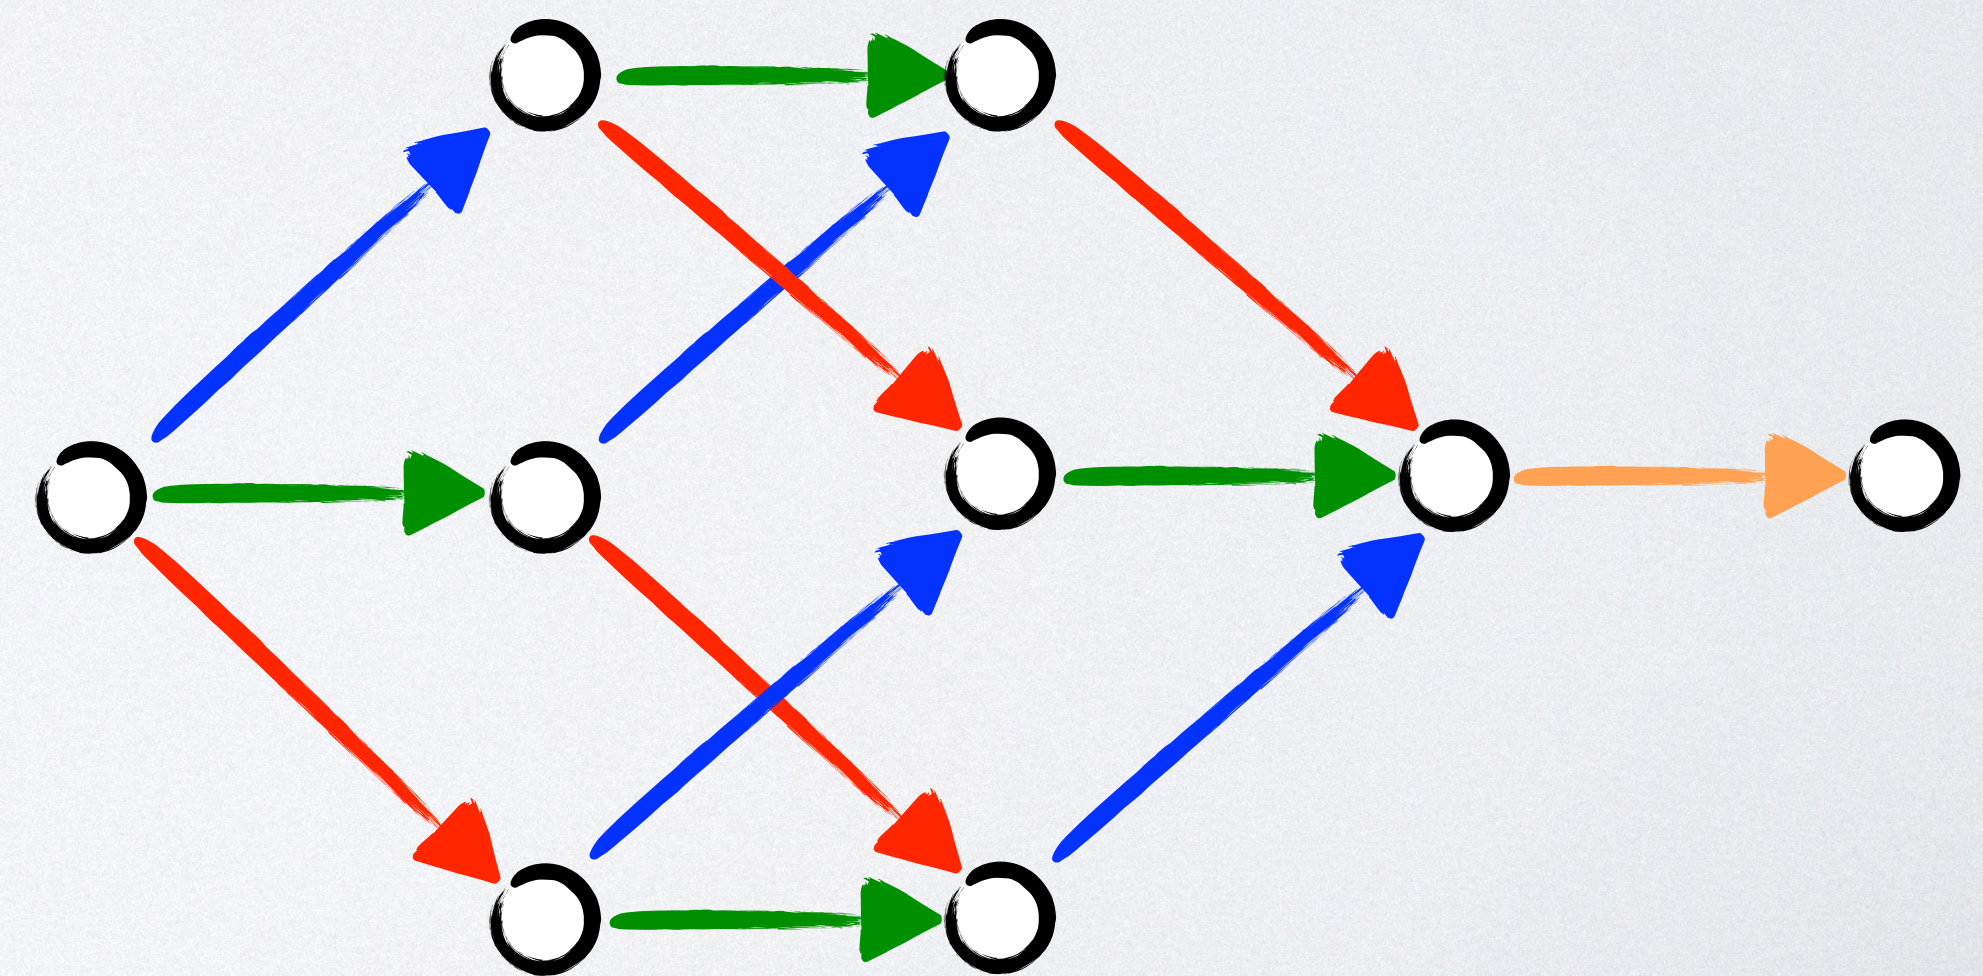
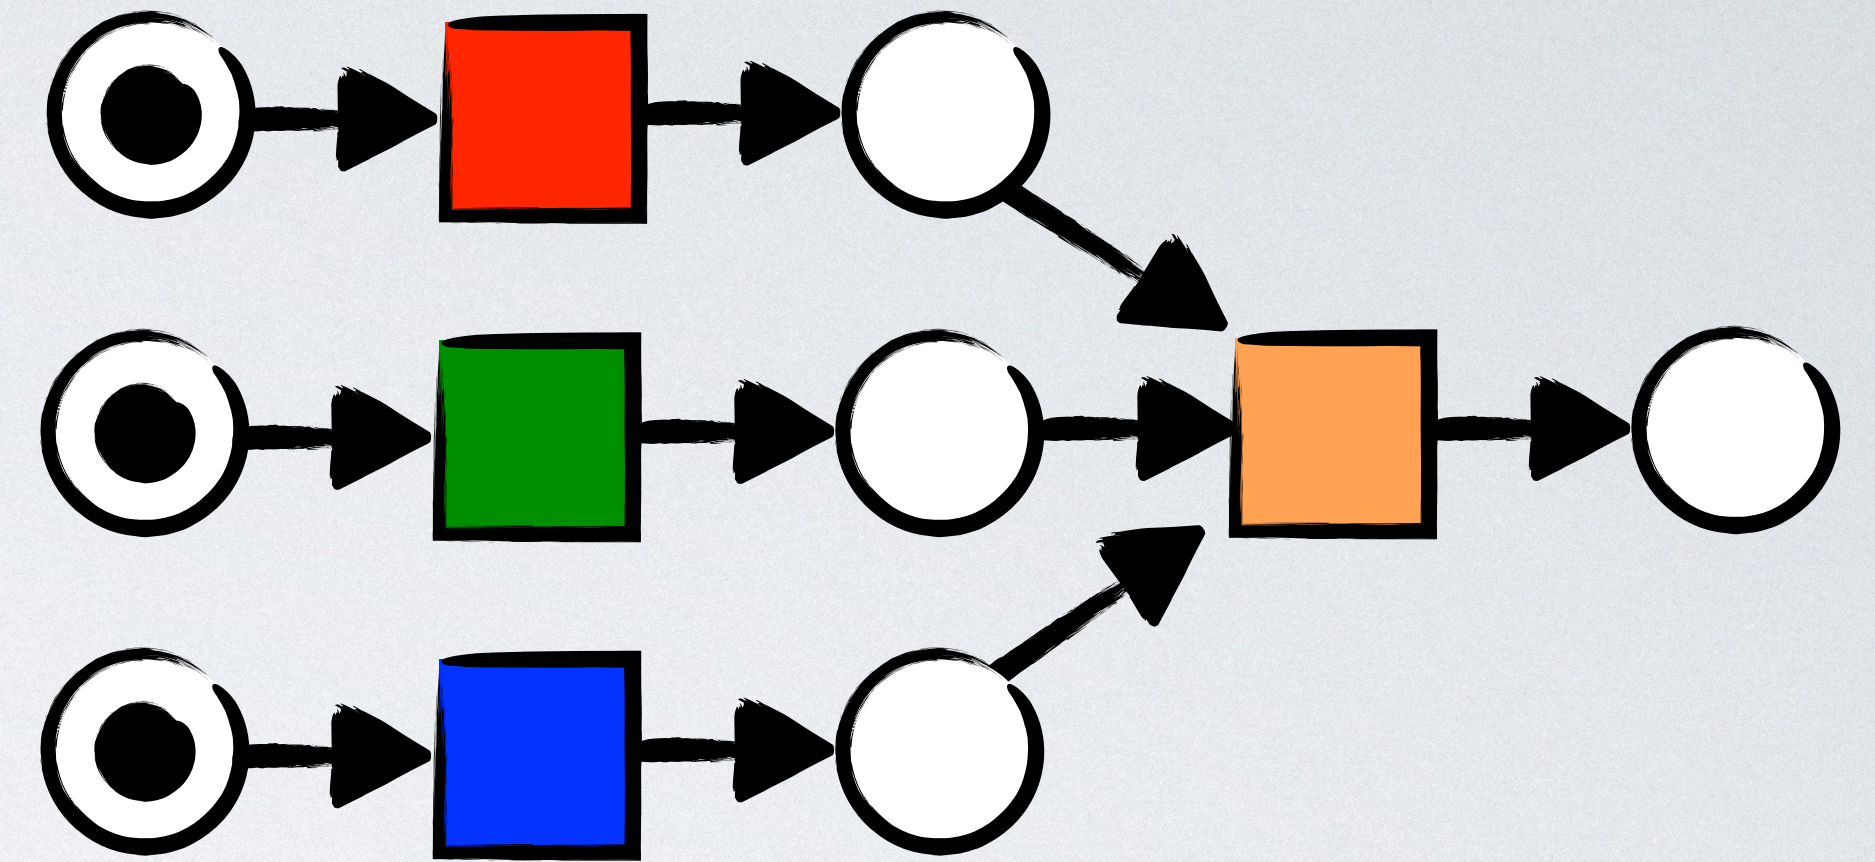


1. Partial order reduction

algorithm (sketch):

- choose one activated transition
- until fixed point is reached:
add all conflicting transitions

only fire these transitions

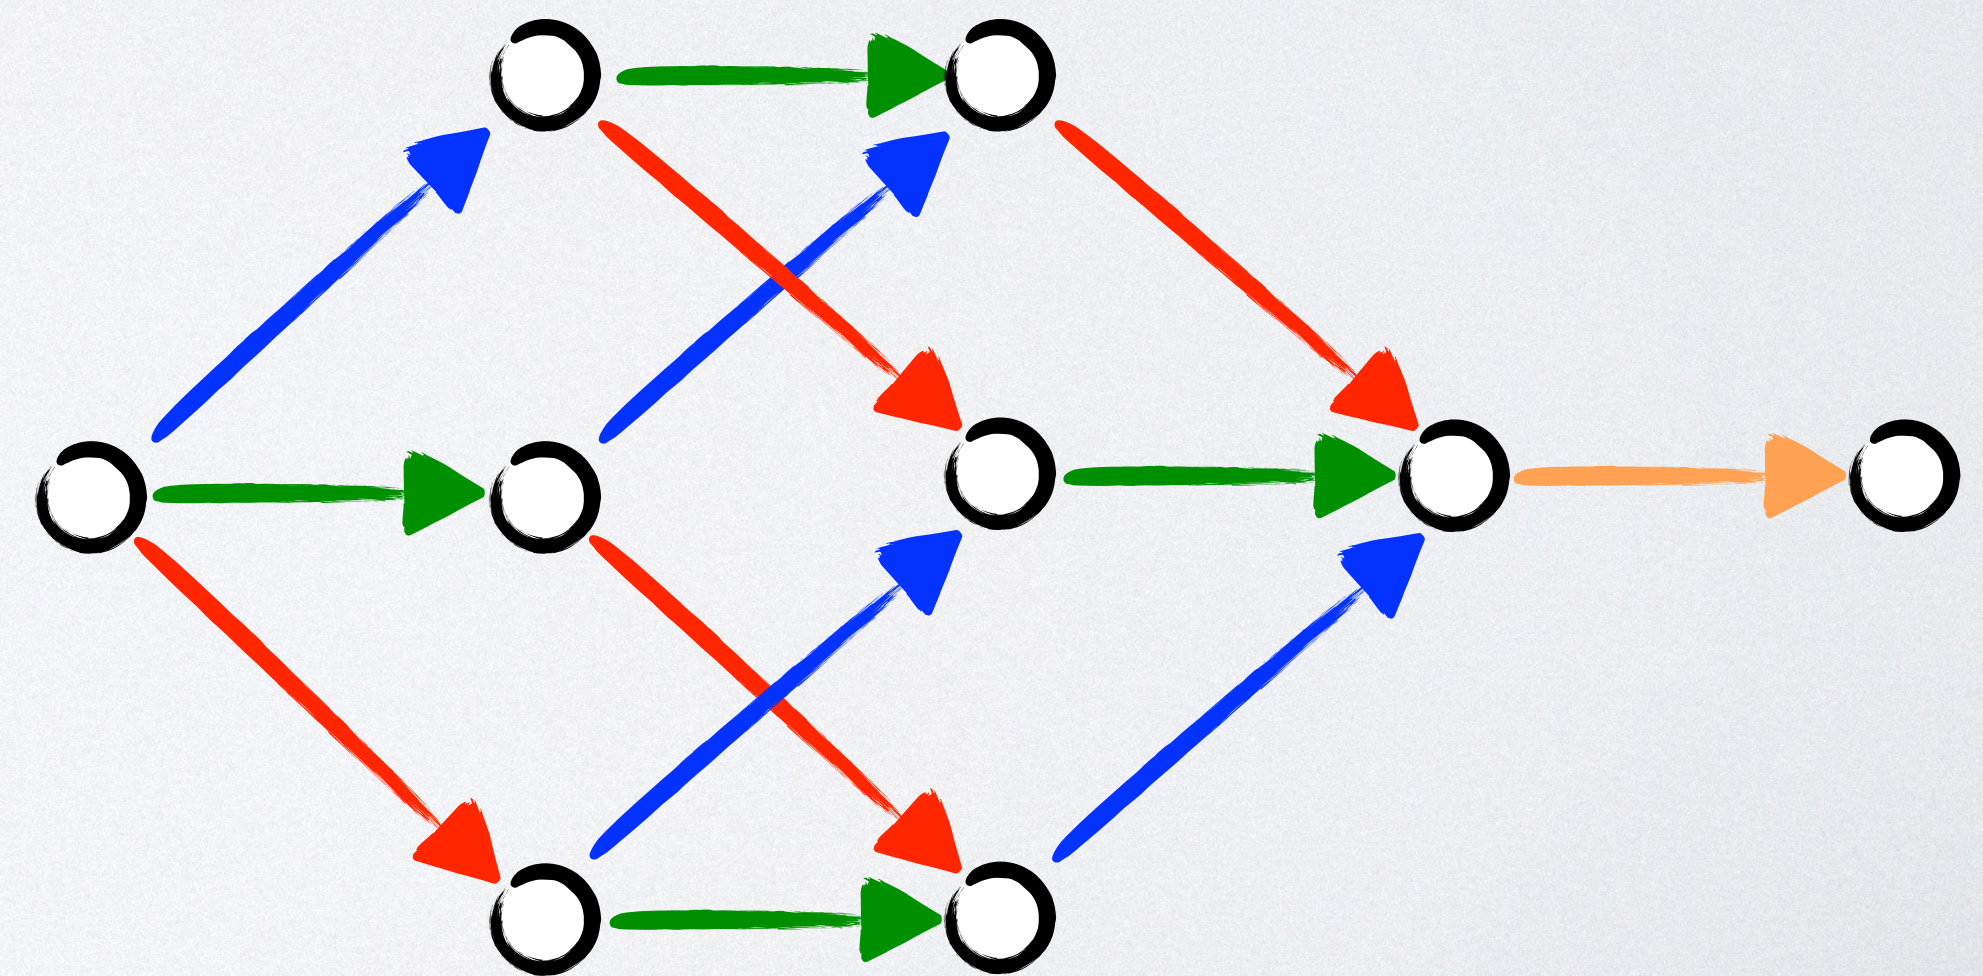
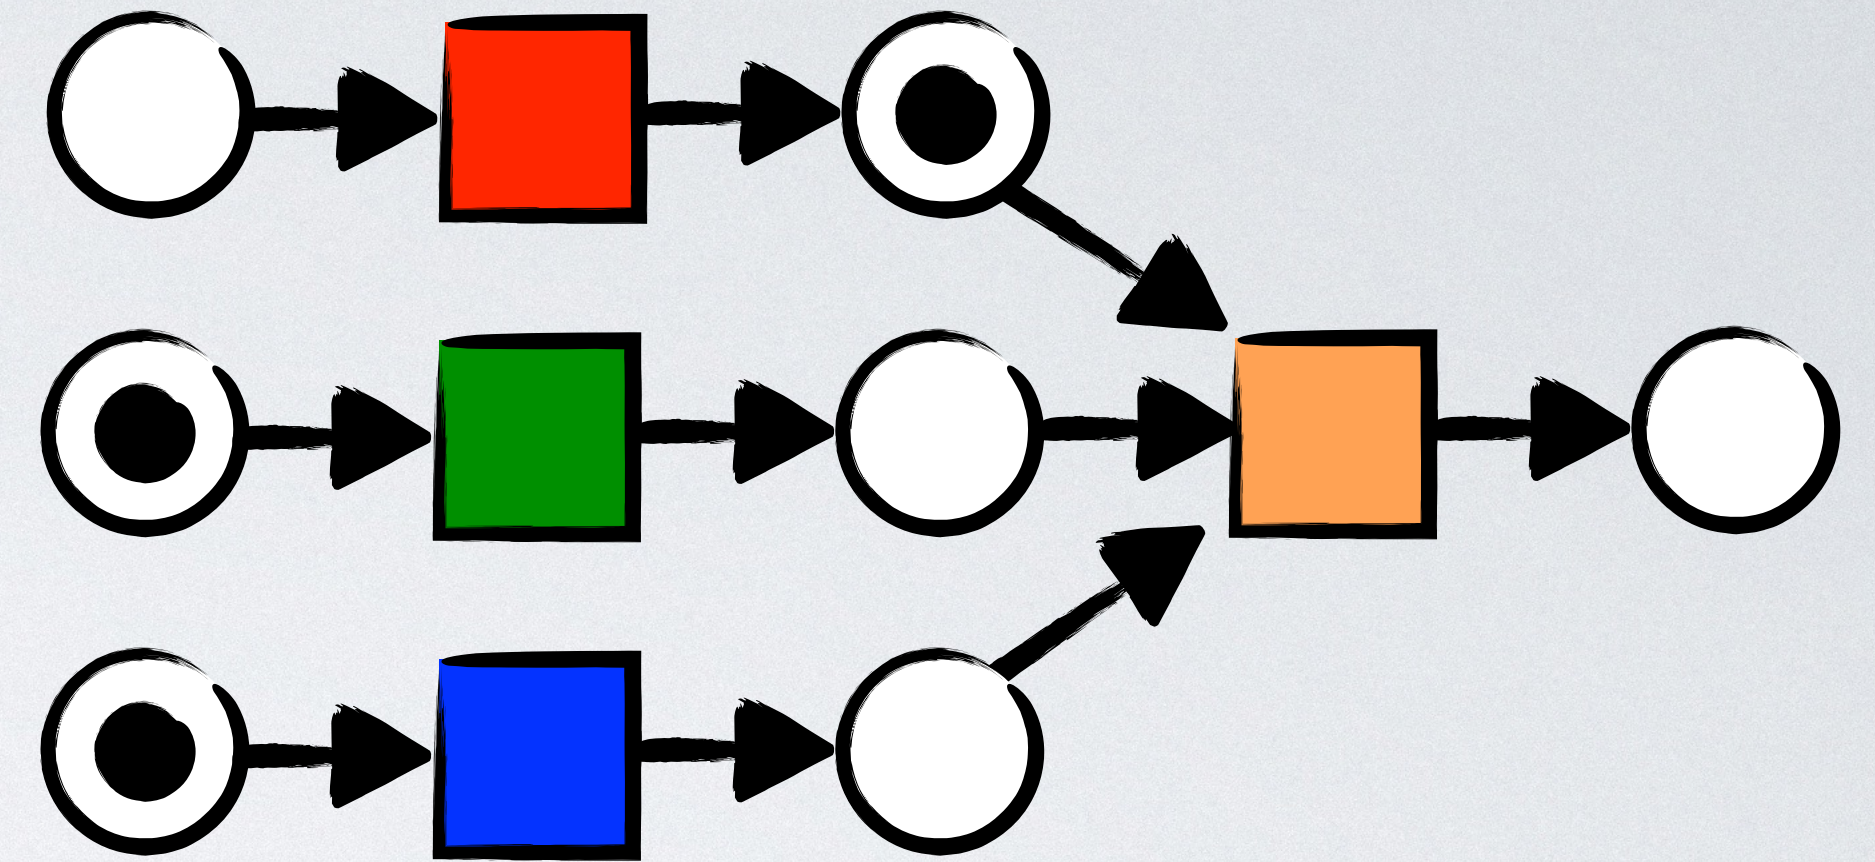


1. Partial order reduction

algorithm (sketch):

- choose one activated transition
- until fixed point is reached:
add all conflicting transitions

only fire these transitions

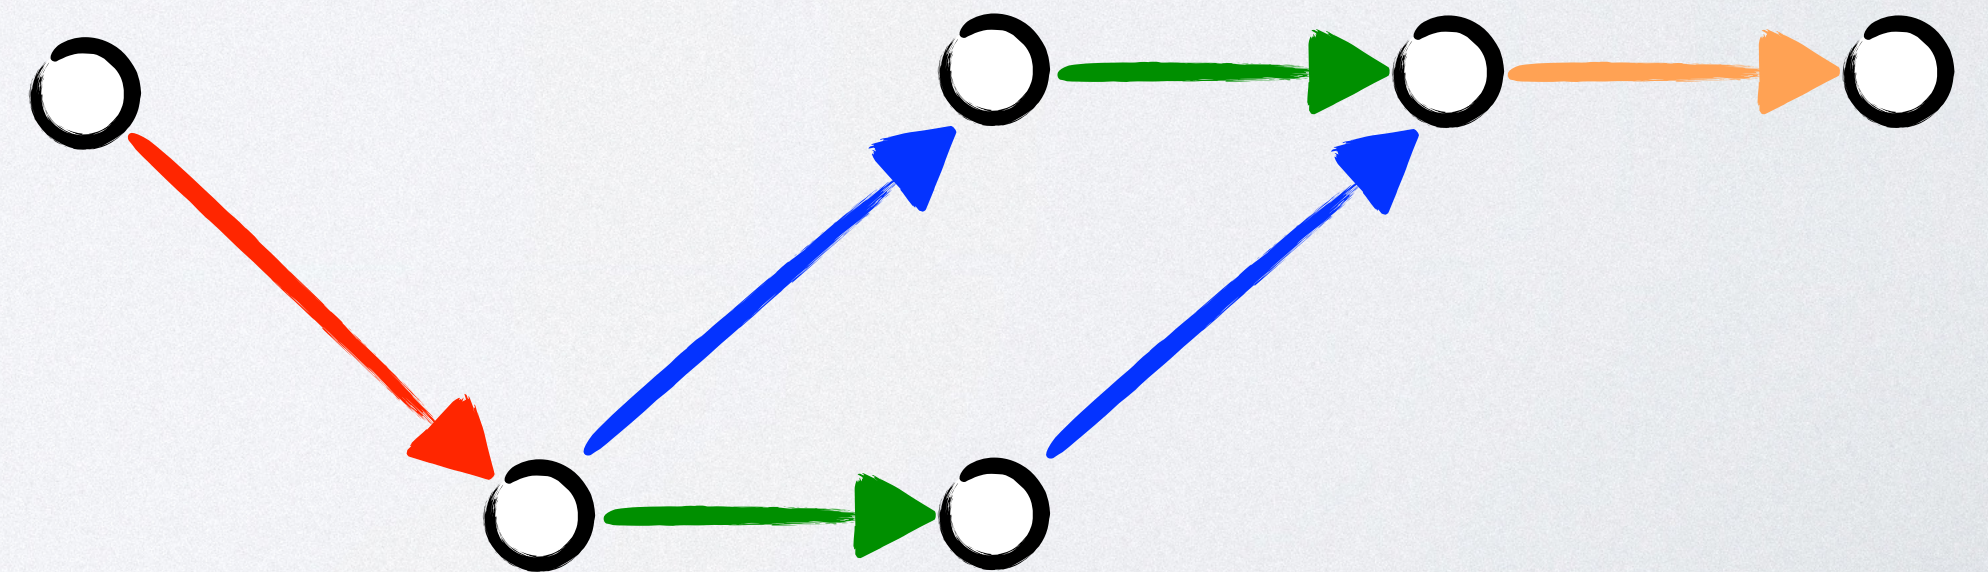
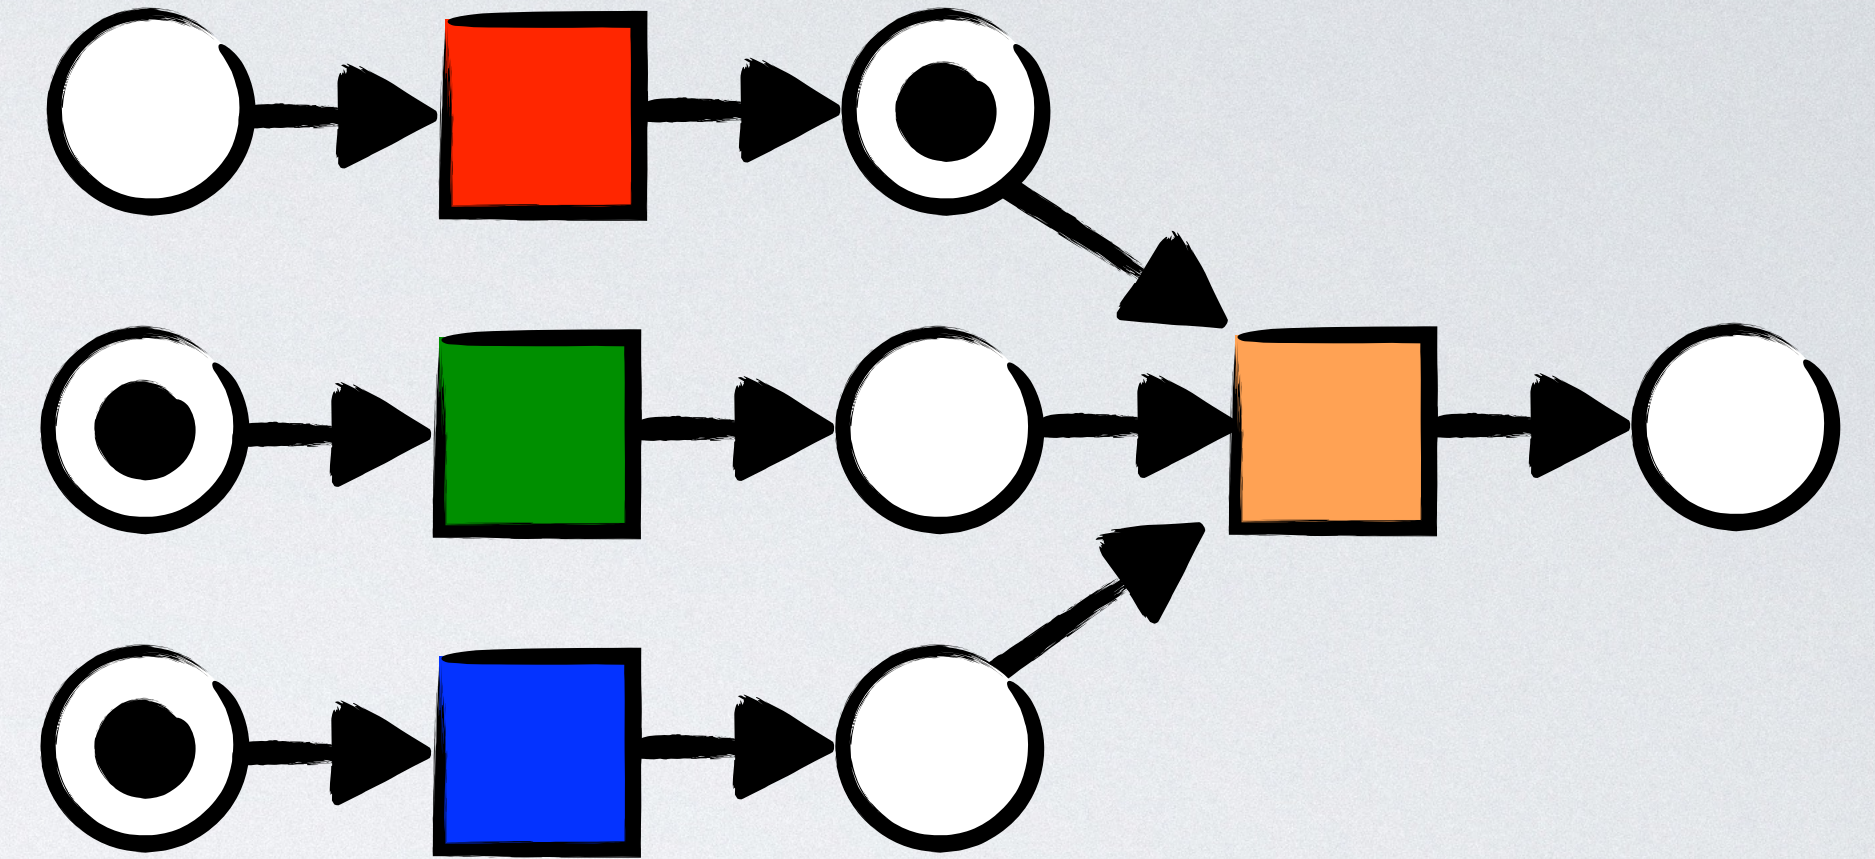


1. Partial order reduction

algorithm (sketch):

- choose one activated transition
- until fixed point is reached:
add all conflicting transitions

only fire these transitions

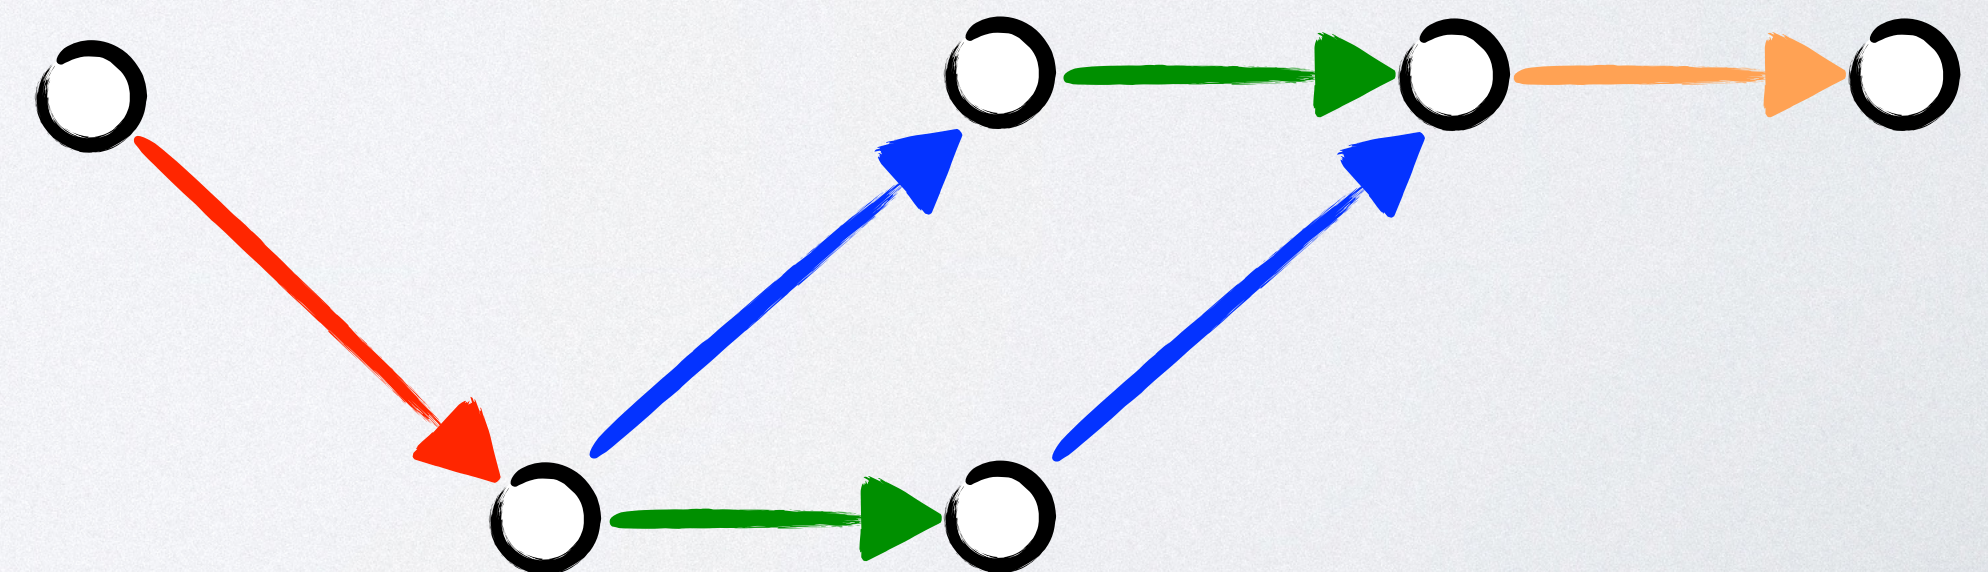
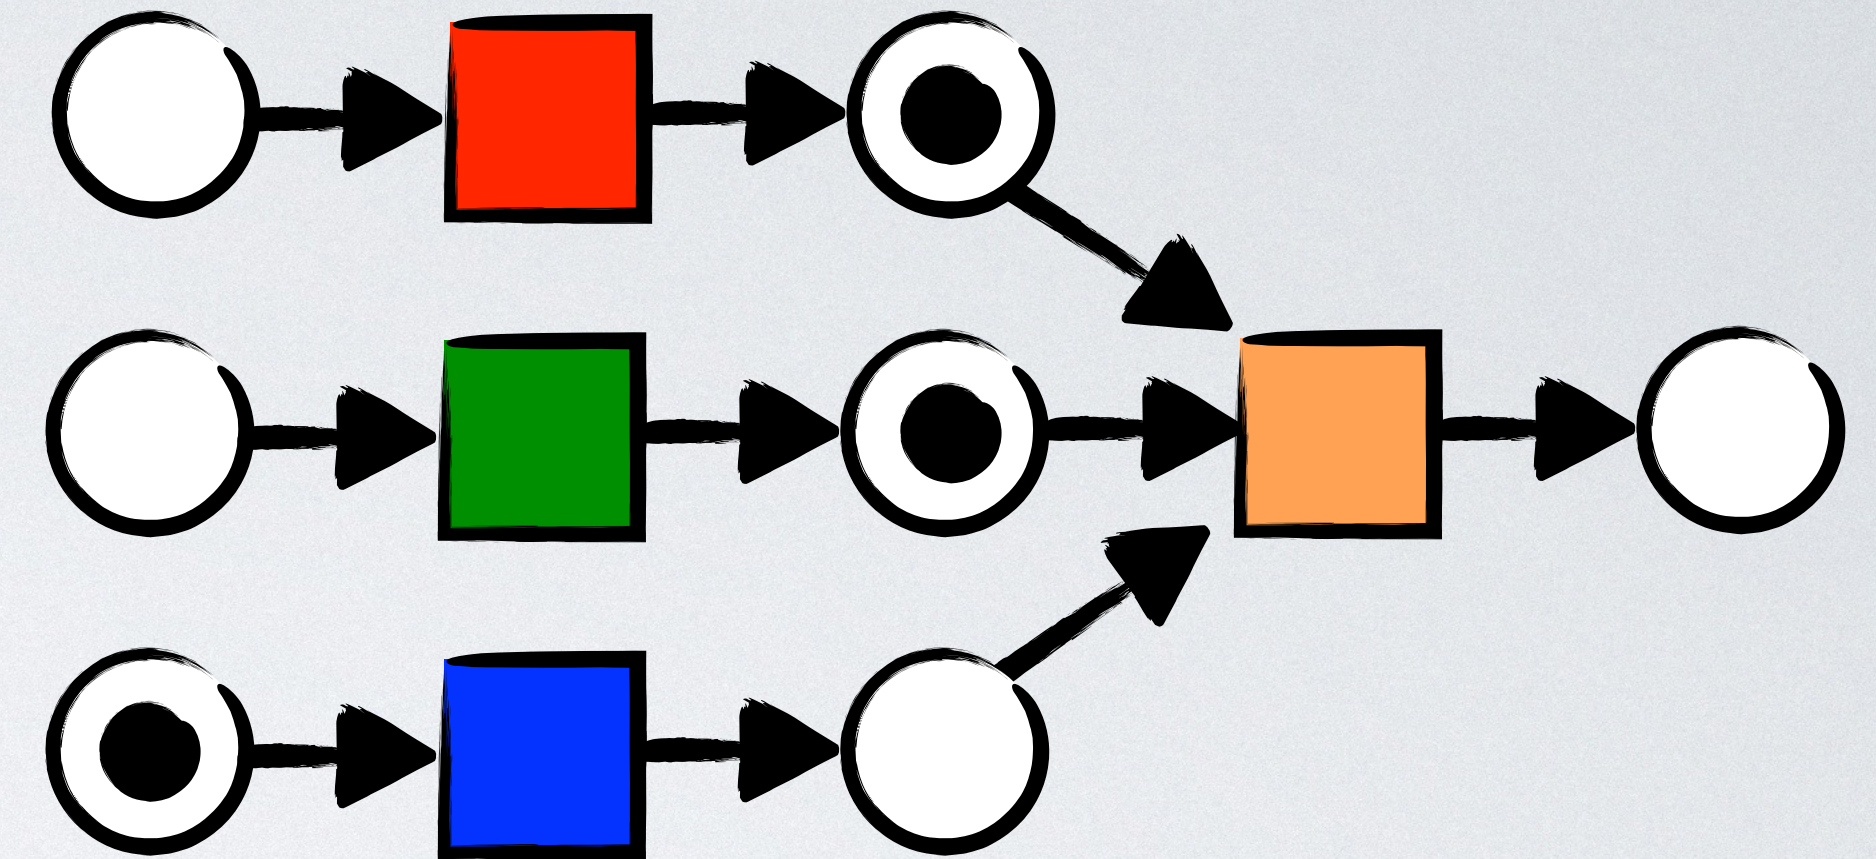


1. Partial order reduction

algorithm (sketch):

- choose one activated transition
- until fixed point is reached:
add all conflicting transitions

only fire these transitions

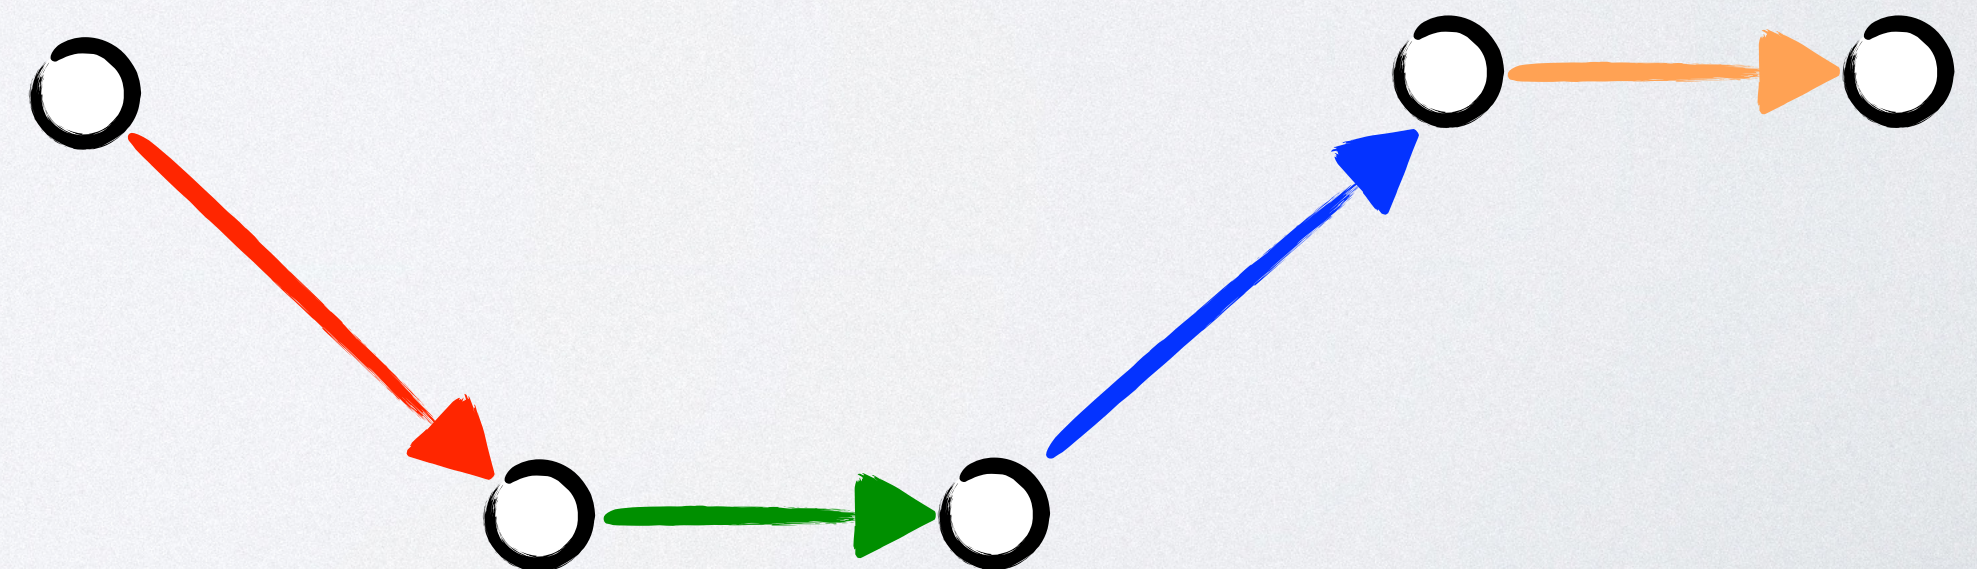
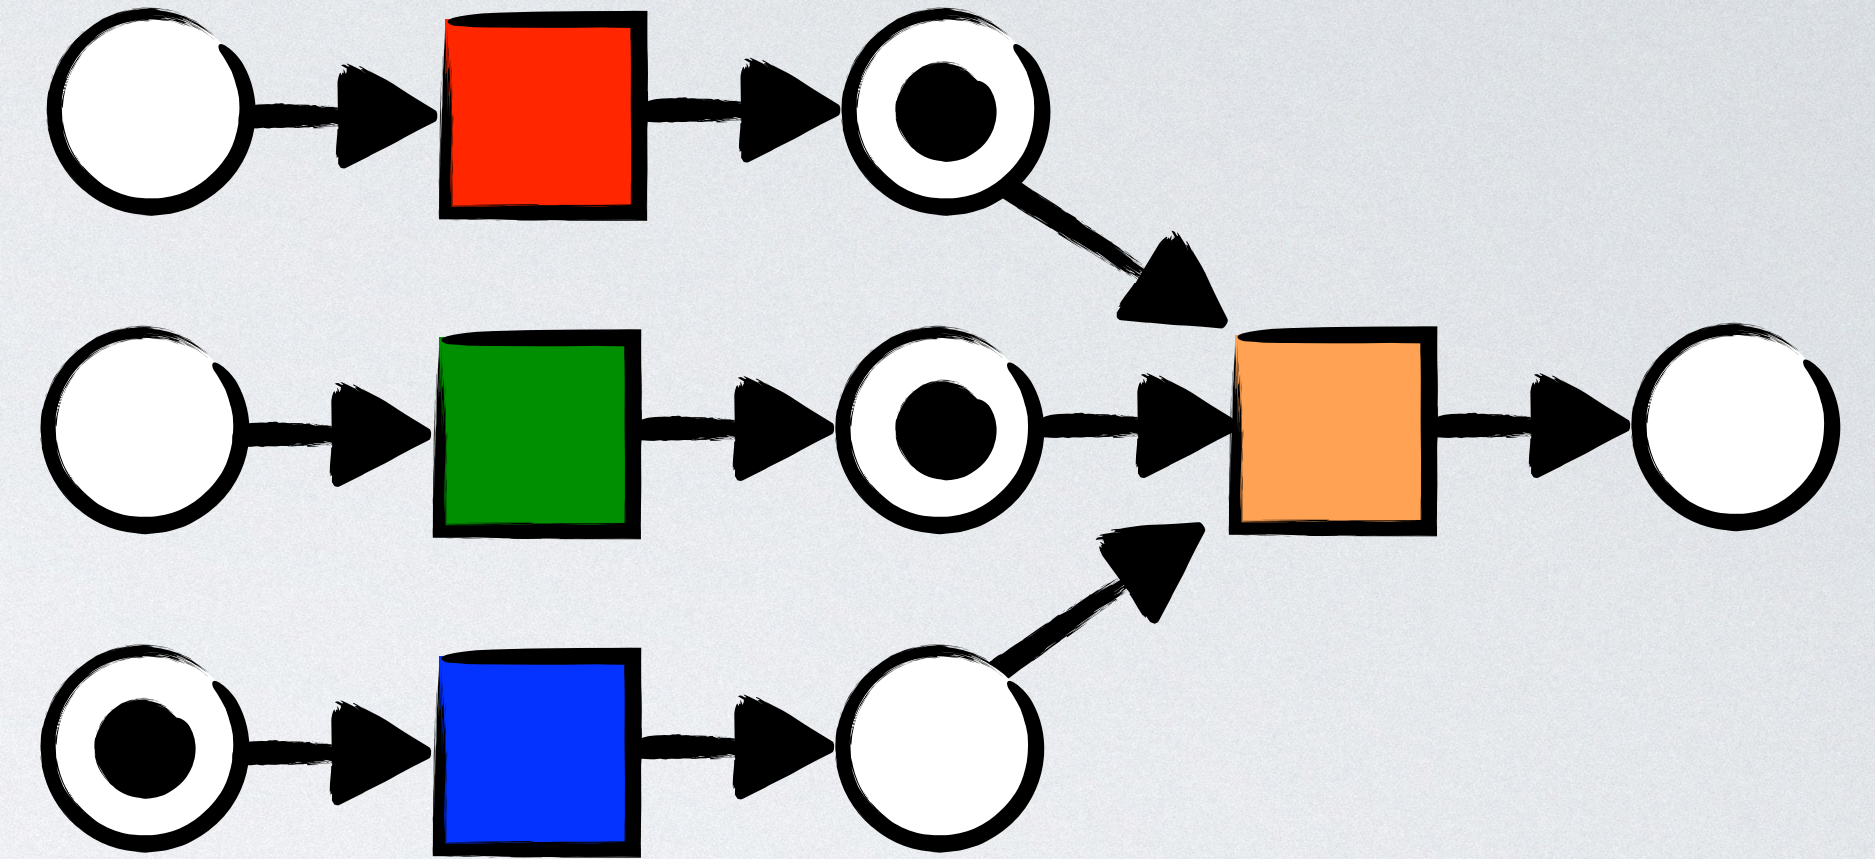


1. Partial order reduction

algorithm (sketch):

- choose one activated transition
- until fixed point is reached:
add all conflicting transitions

only fire these transitions

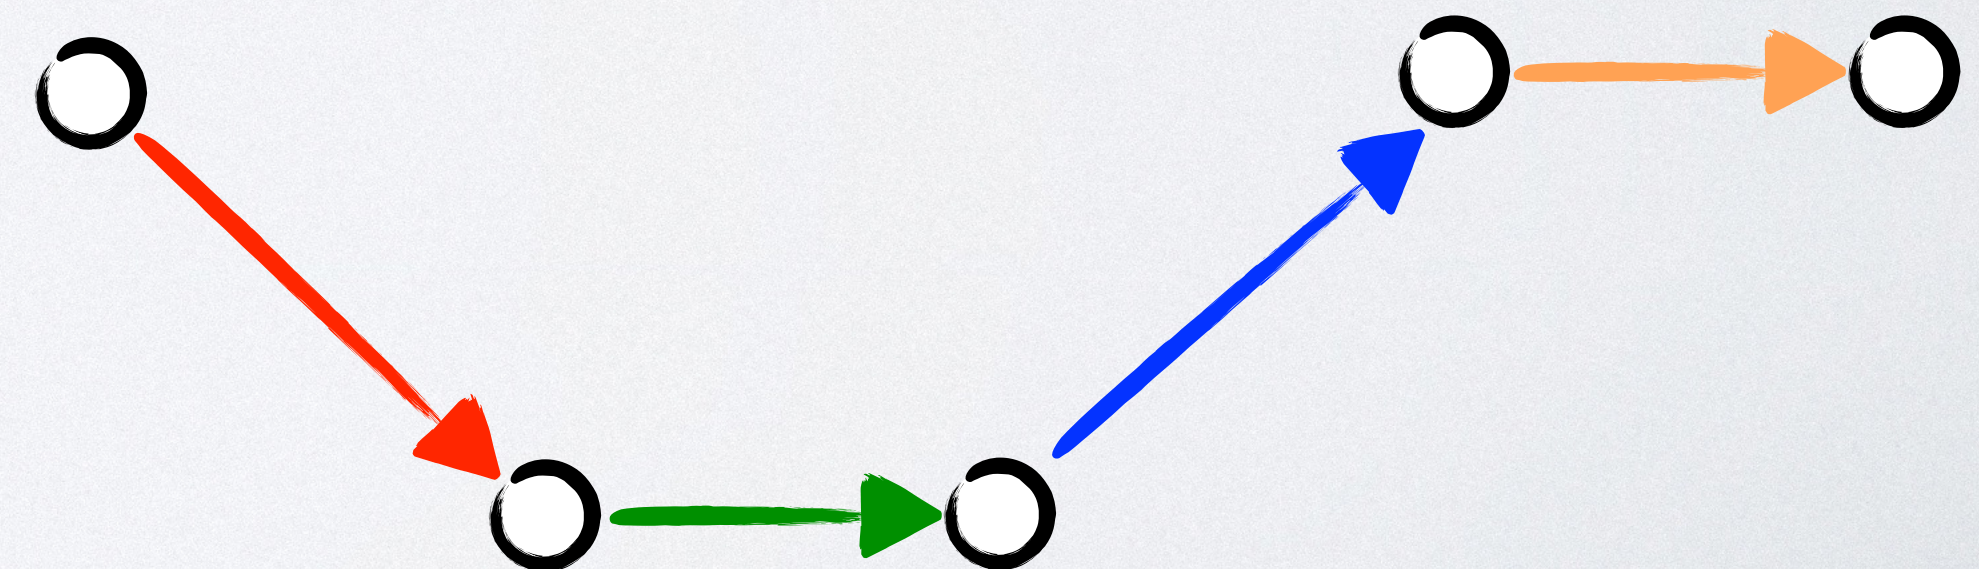
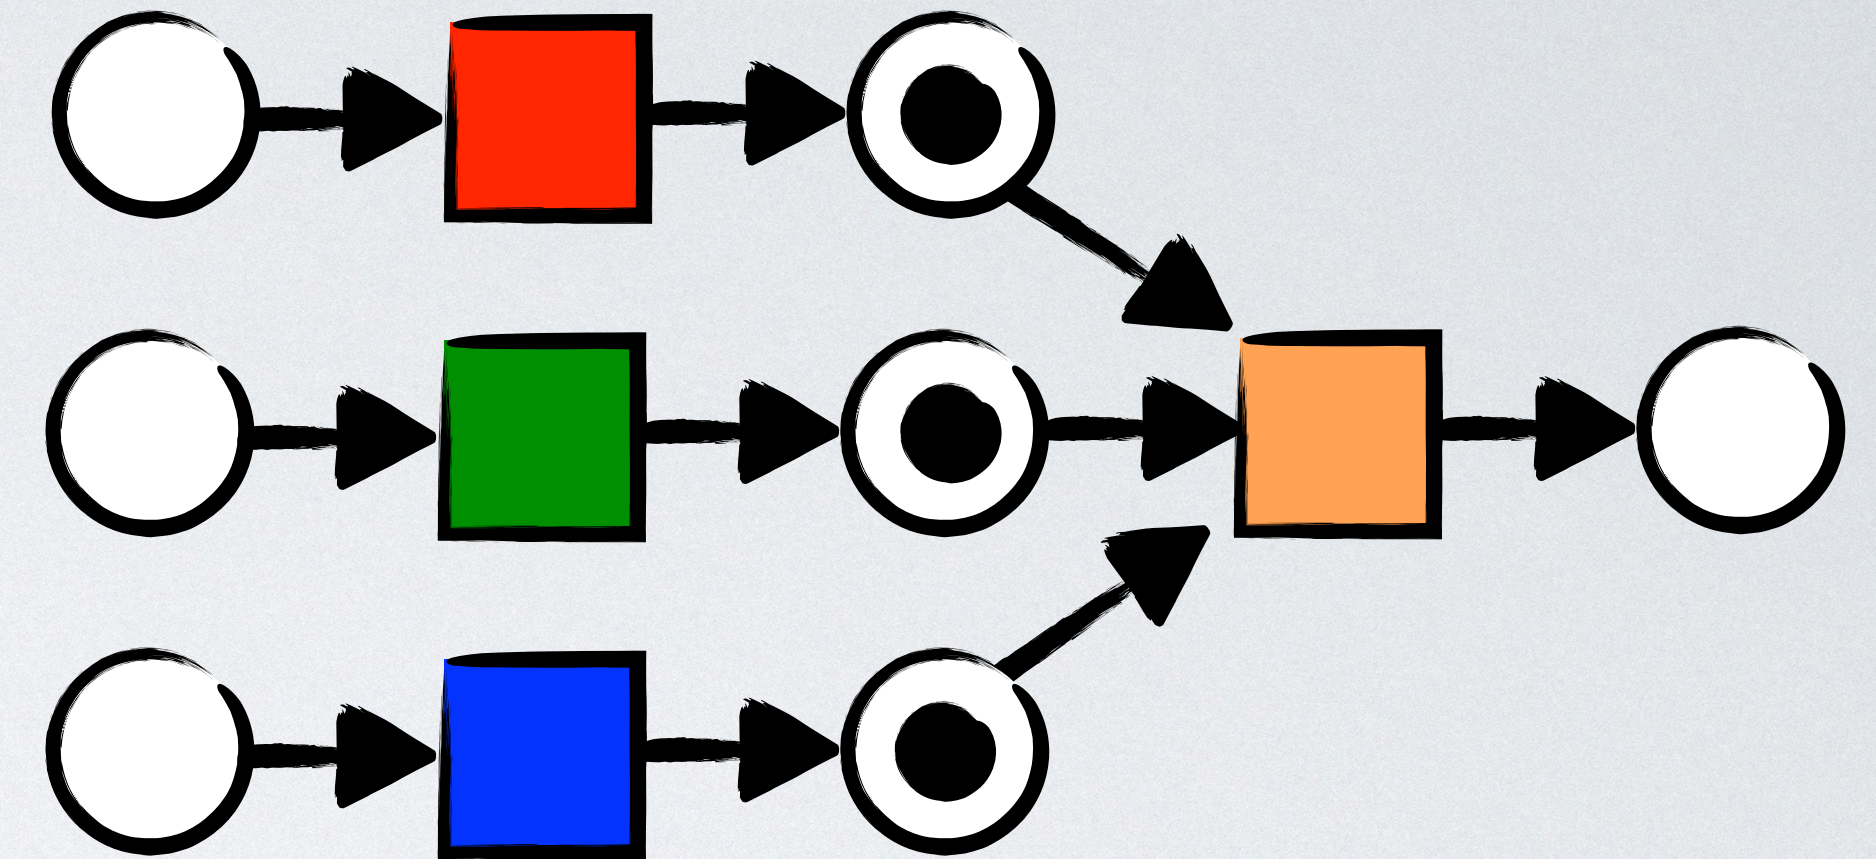


1. Partial order reduction

algorithm (sketch):

- choose one activated transition
- until fixed point is reached:
add all conflicting transitions

only fire these transitions

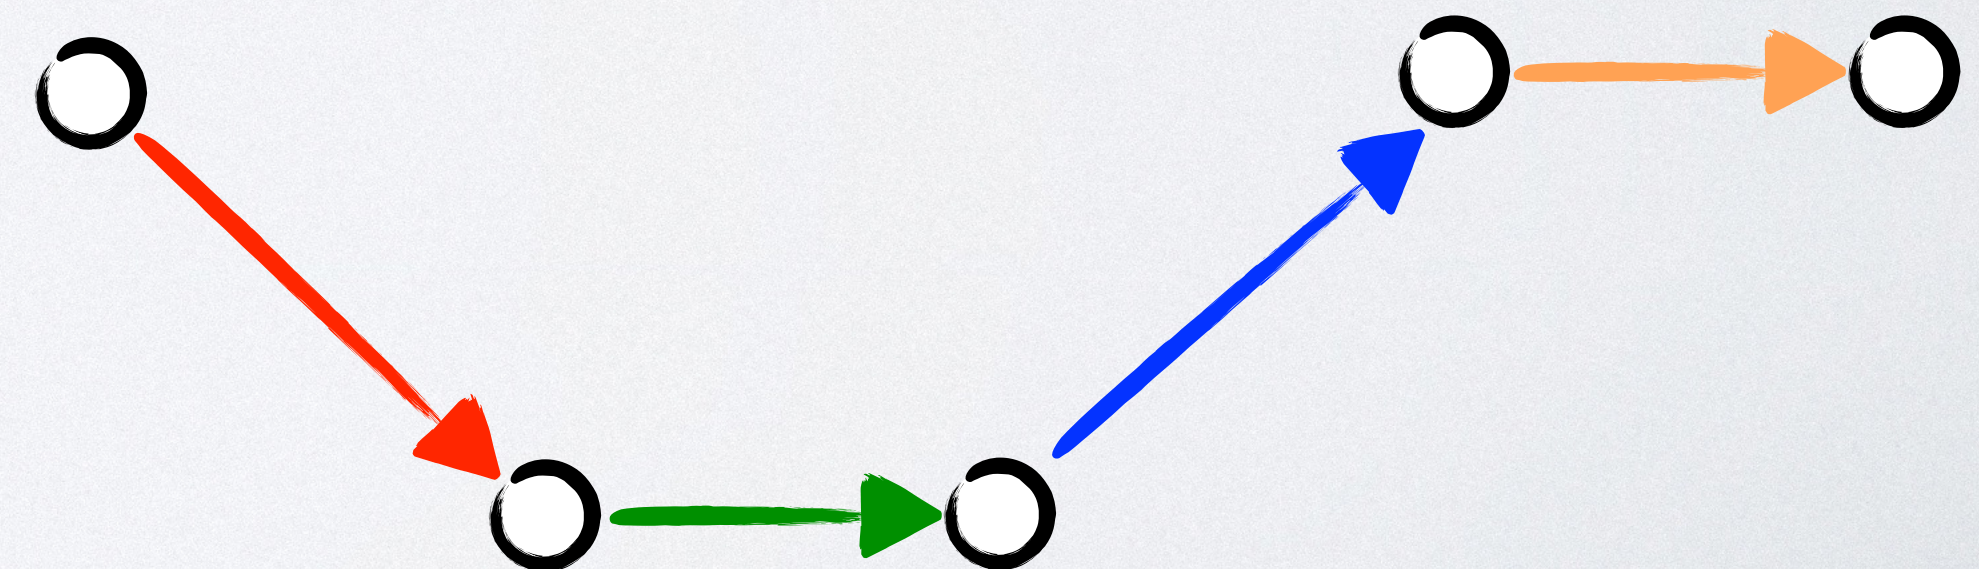
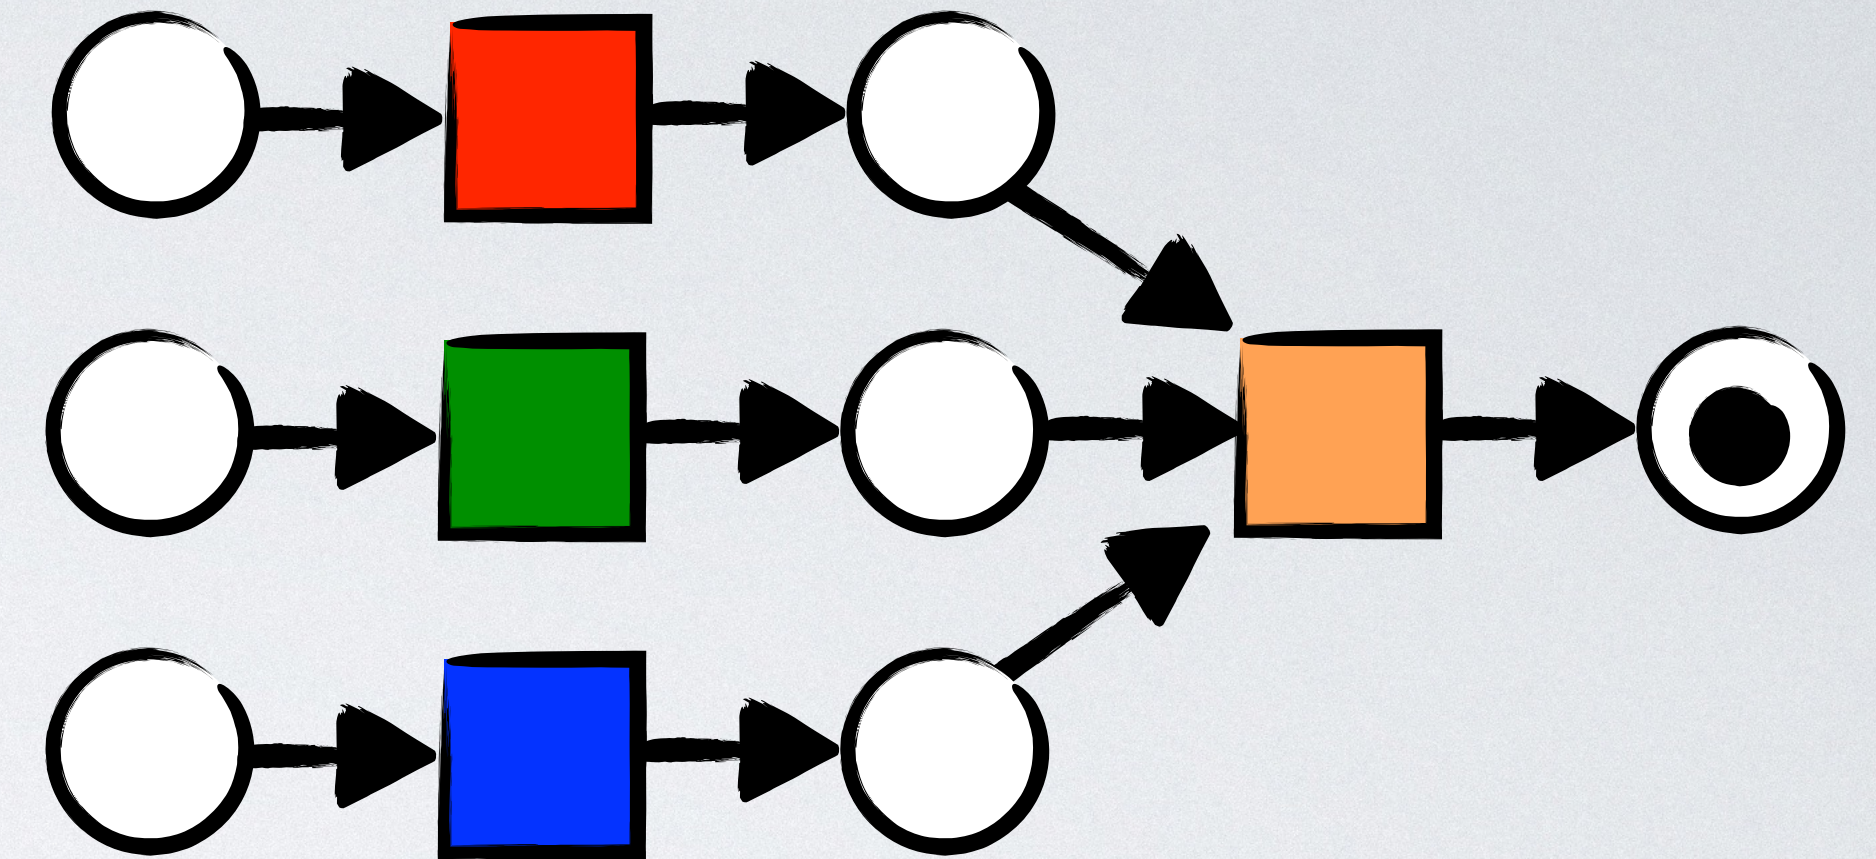


1. Partial order reduction

algorithm (sketch):

- choose one activated transition
- until fixed point is reached:
add all conflicting transitions

only fire these transitions

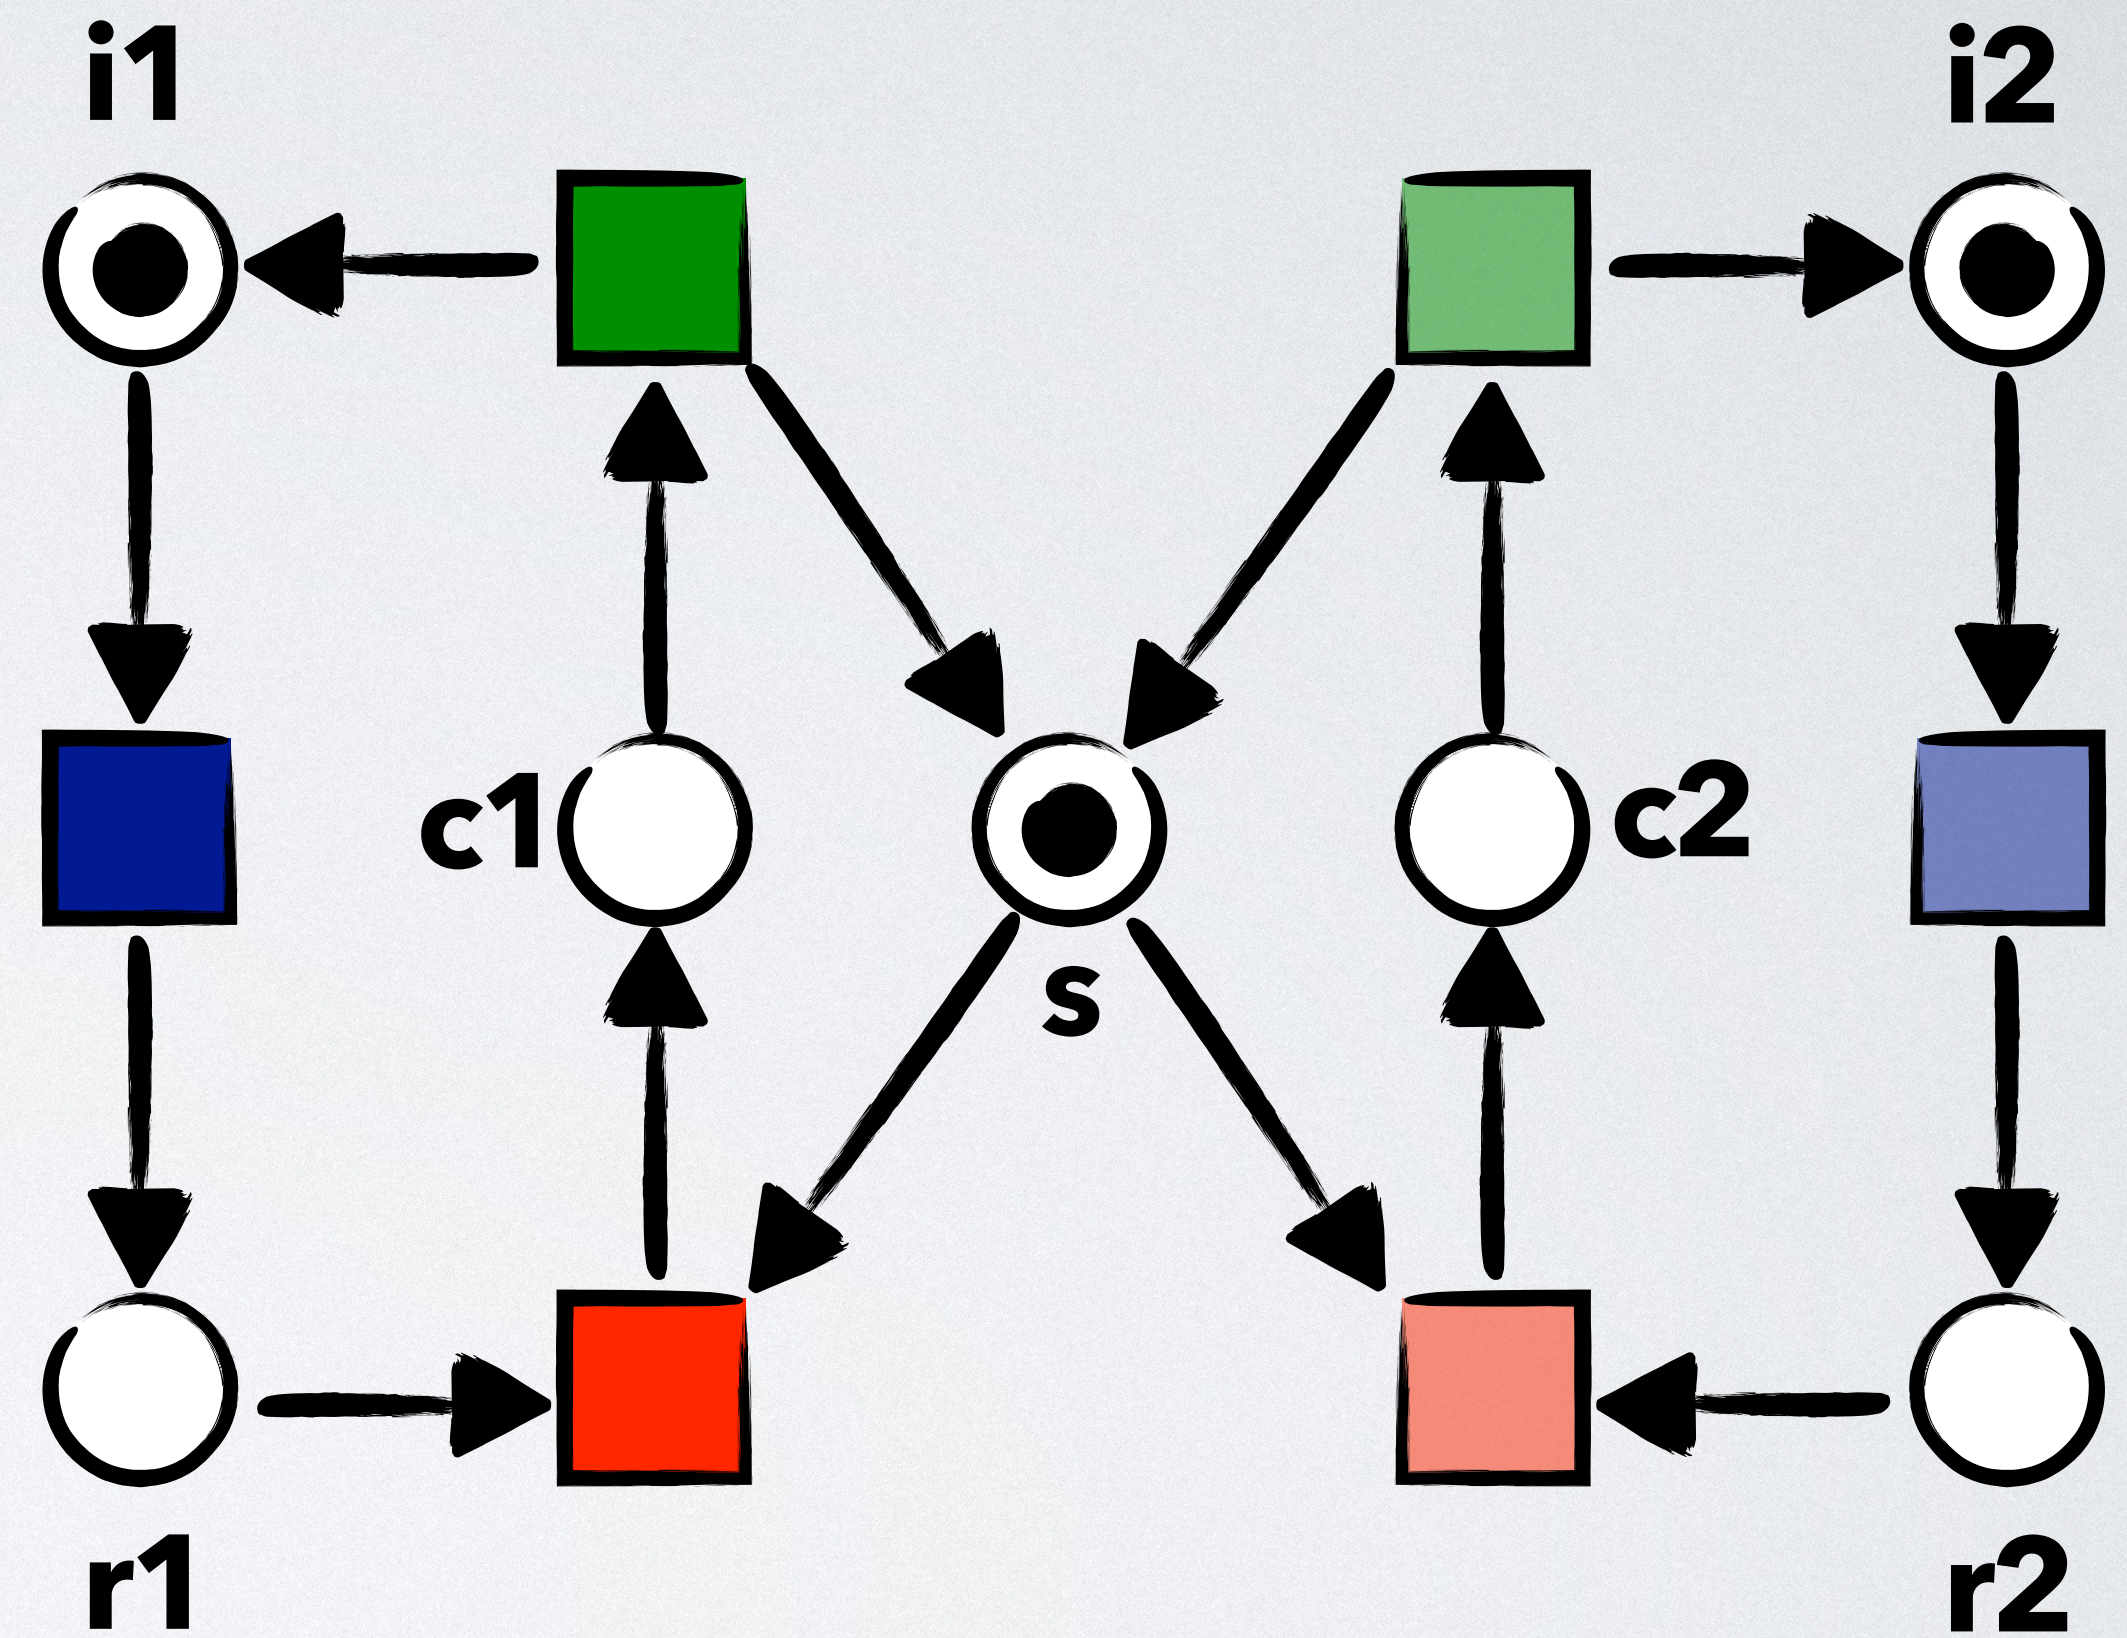


2. Symmetry reduction

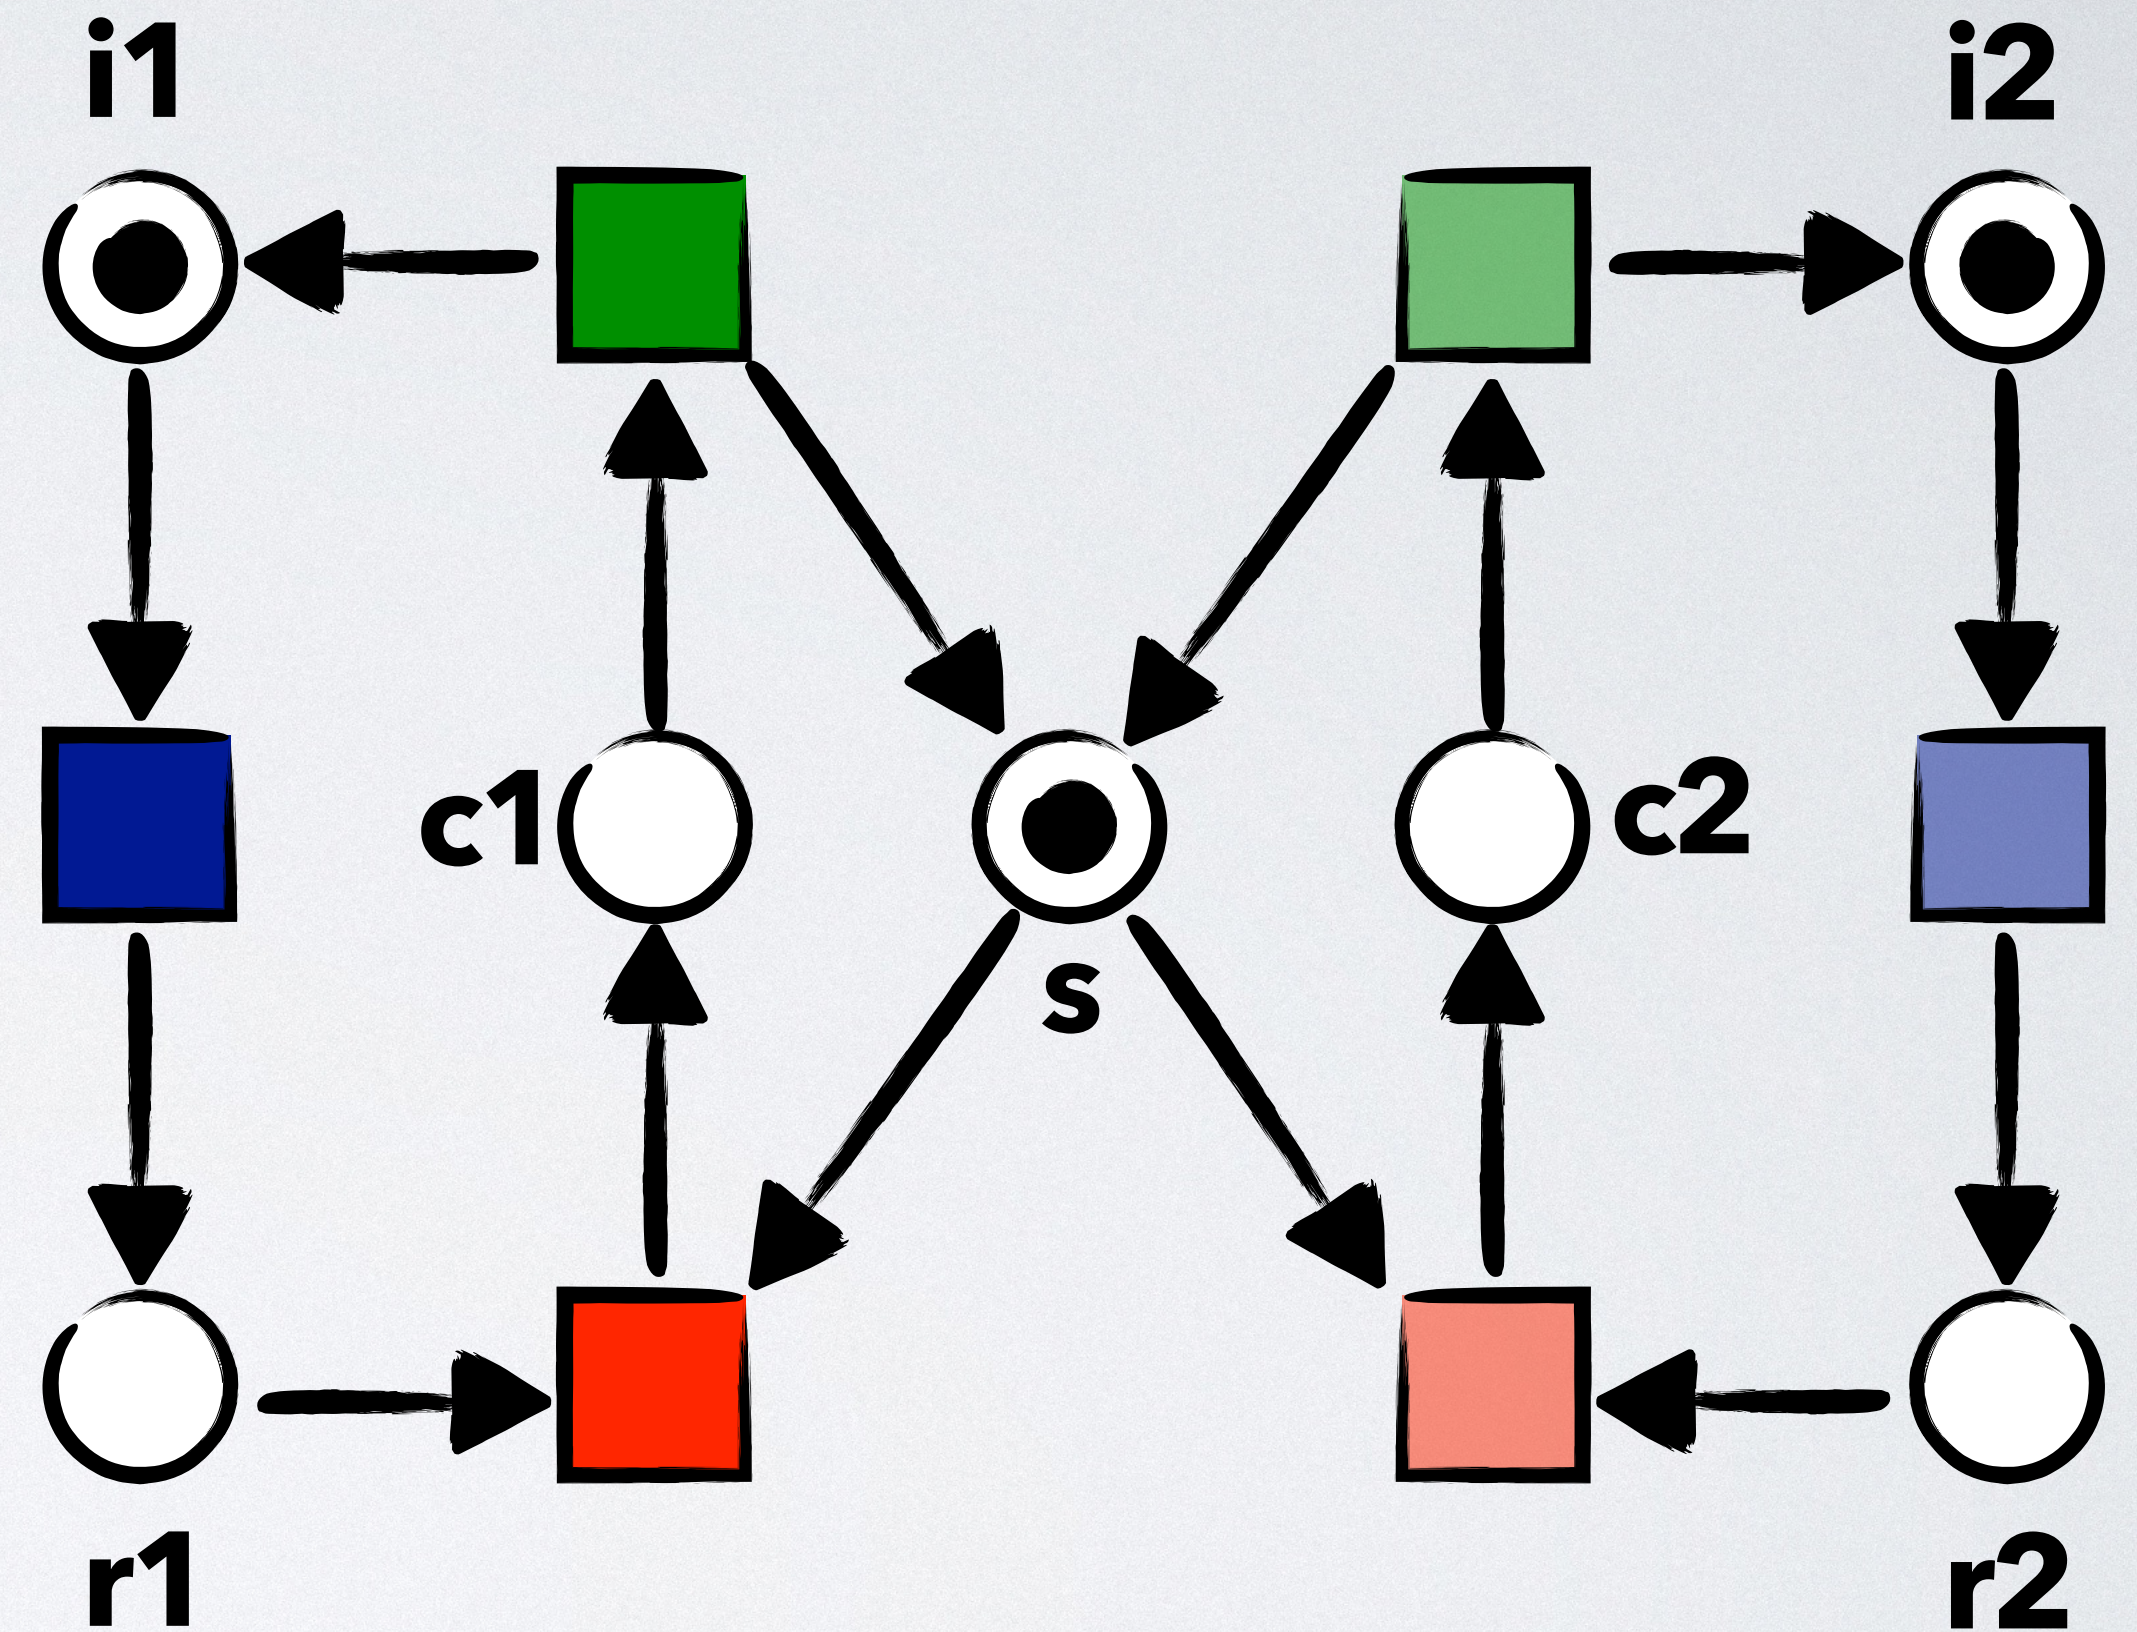
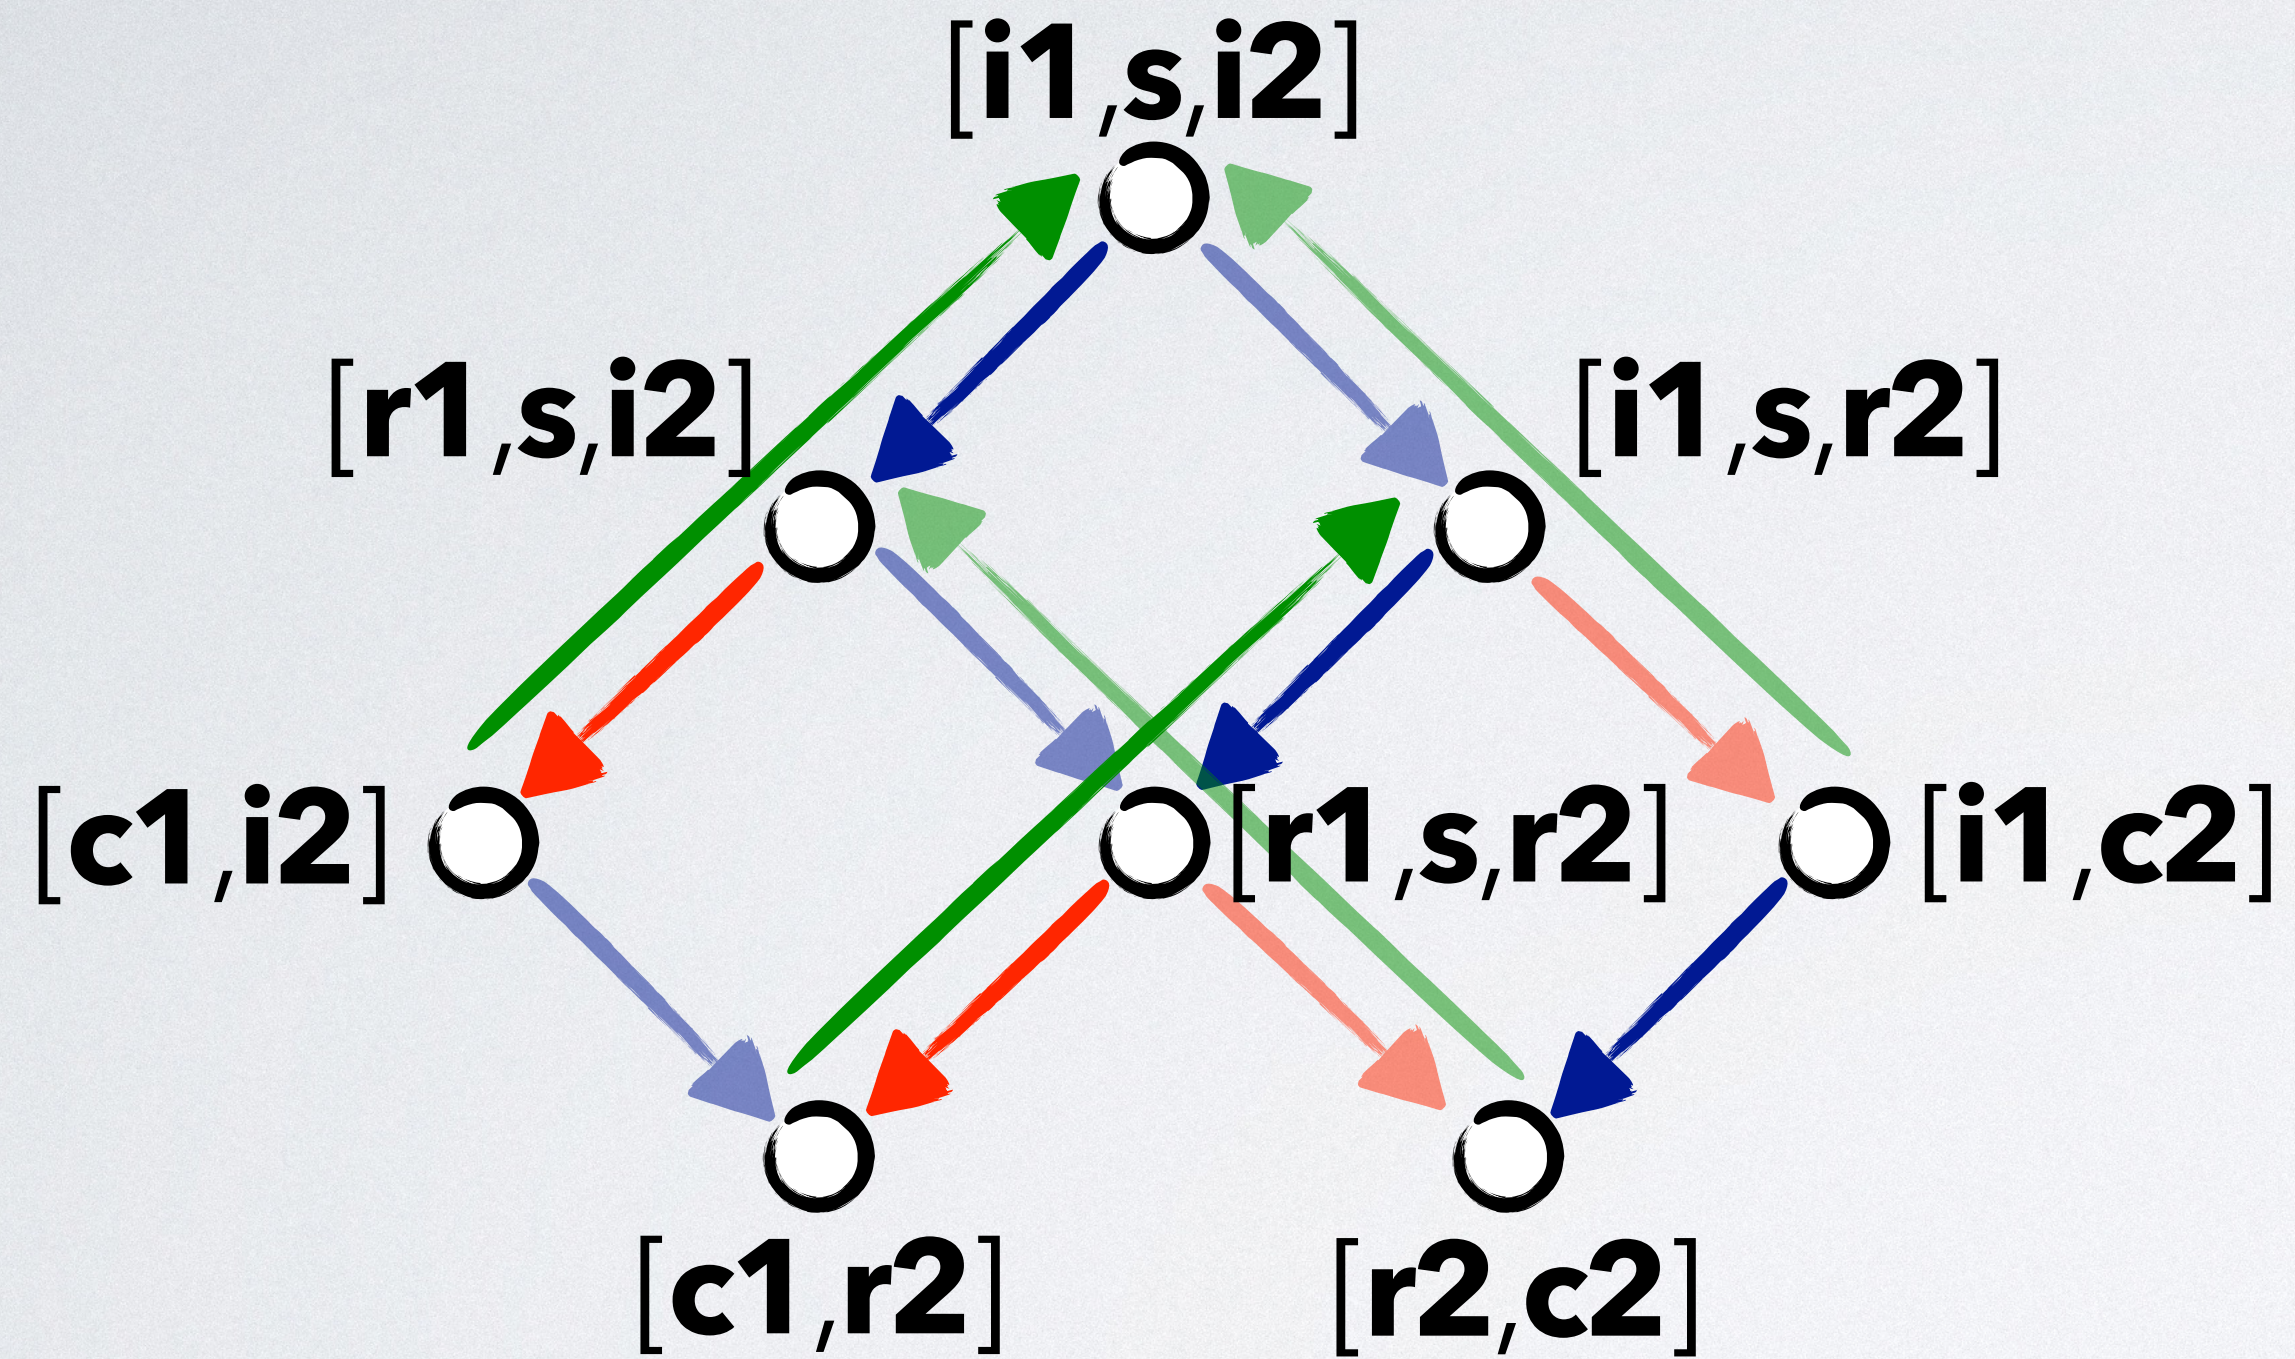
observation: symmetric models (e.g. due to copies of components) have **symmetric behavior**

idea: do not store markings if a symmetric copy is already stored

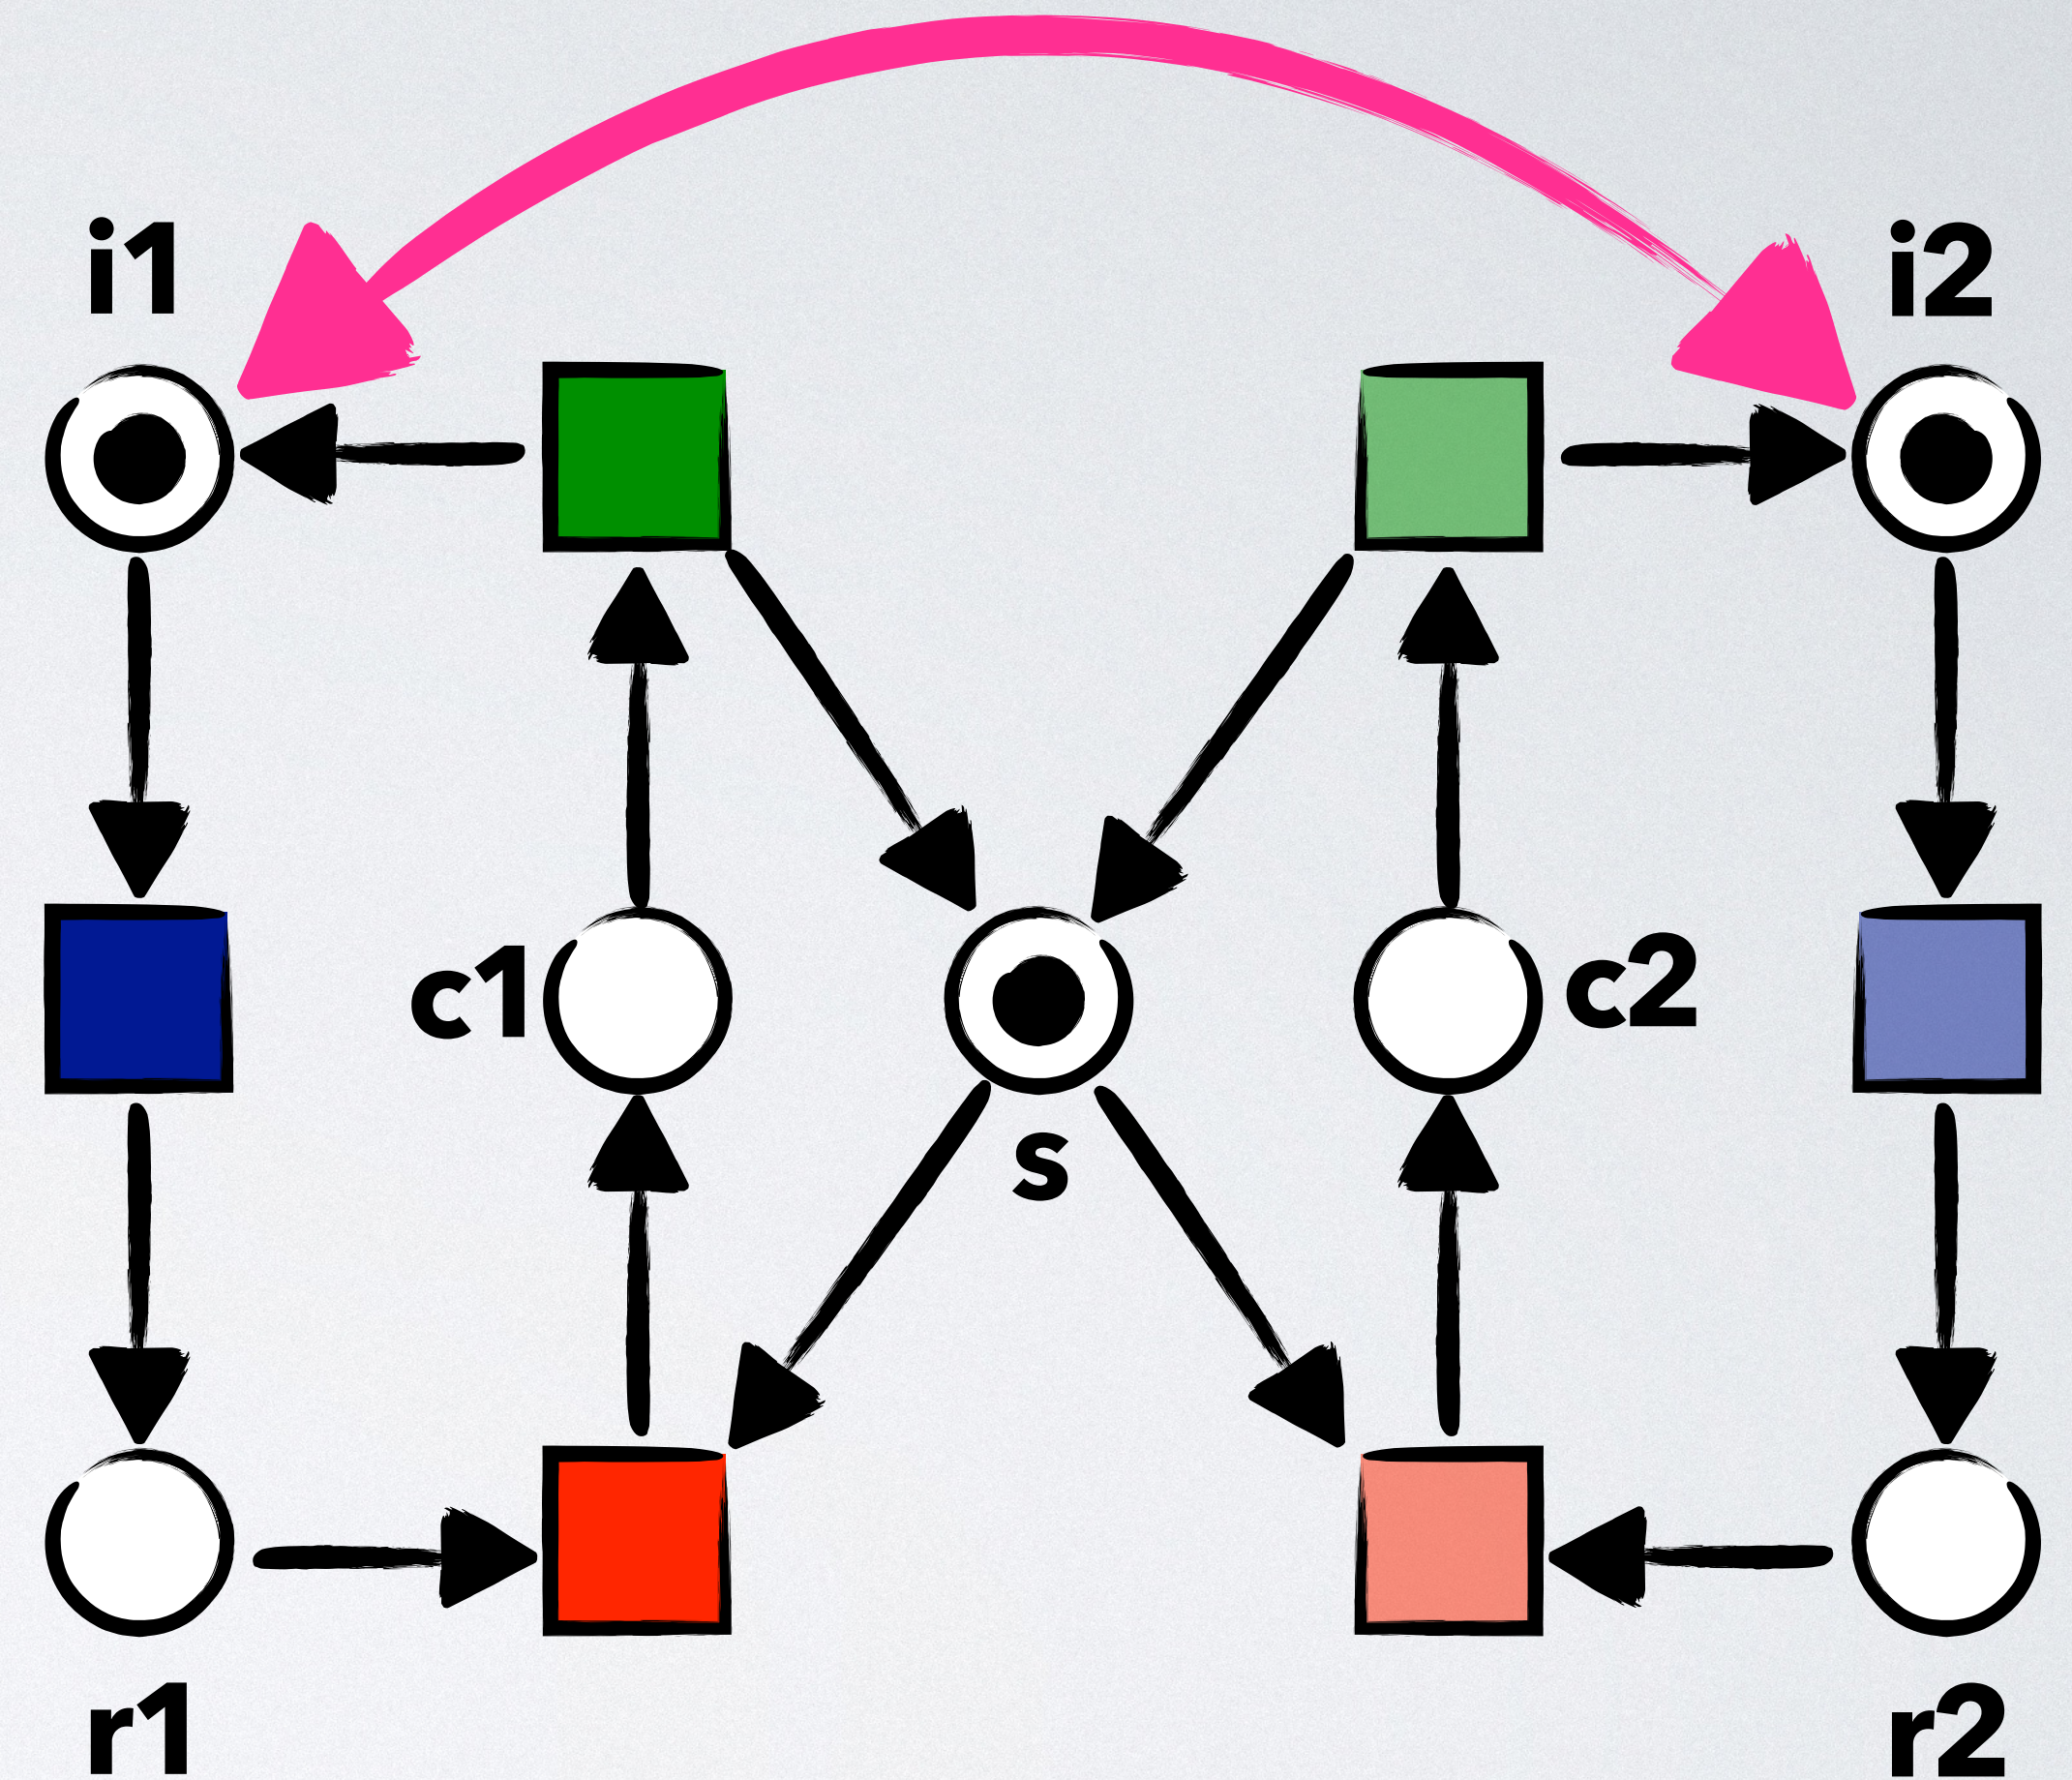
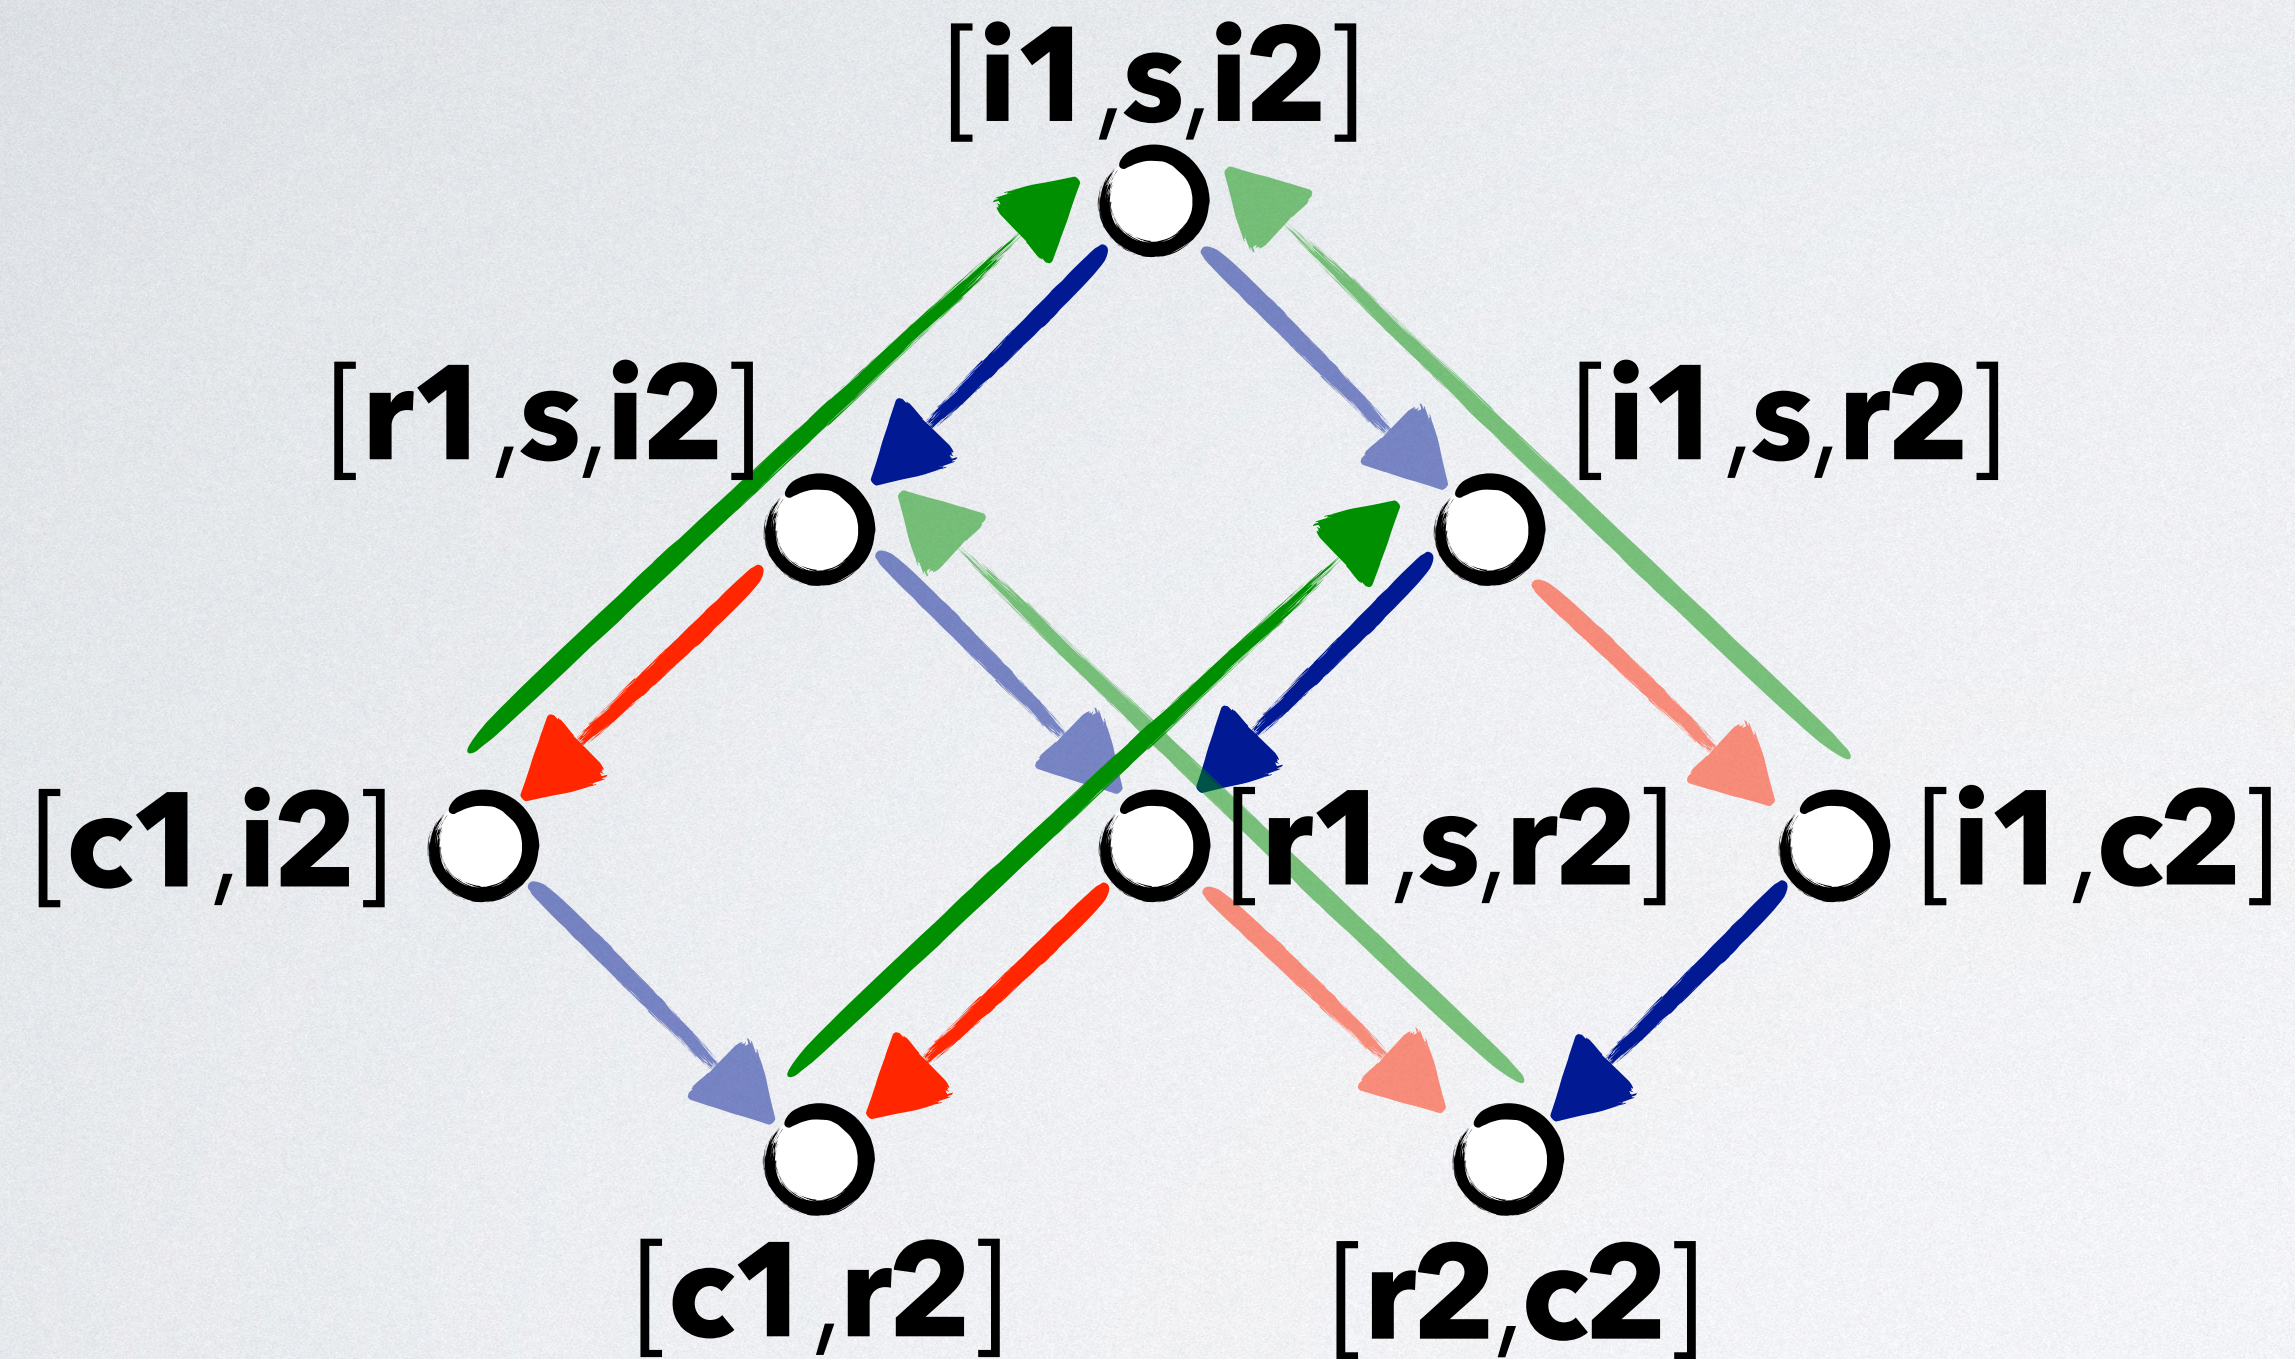
implementation: Petri net graph automorphisms



2. Symmetry reduction

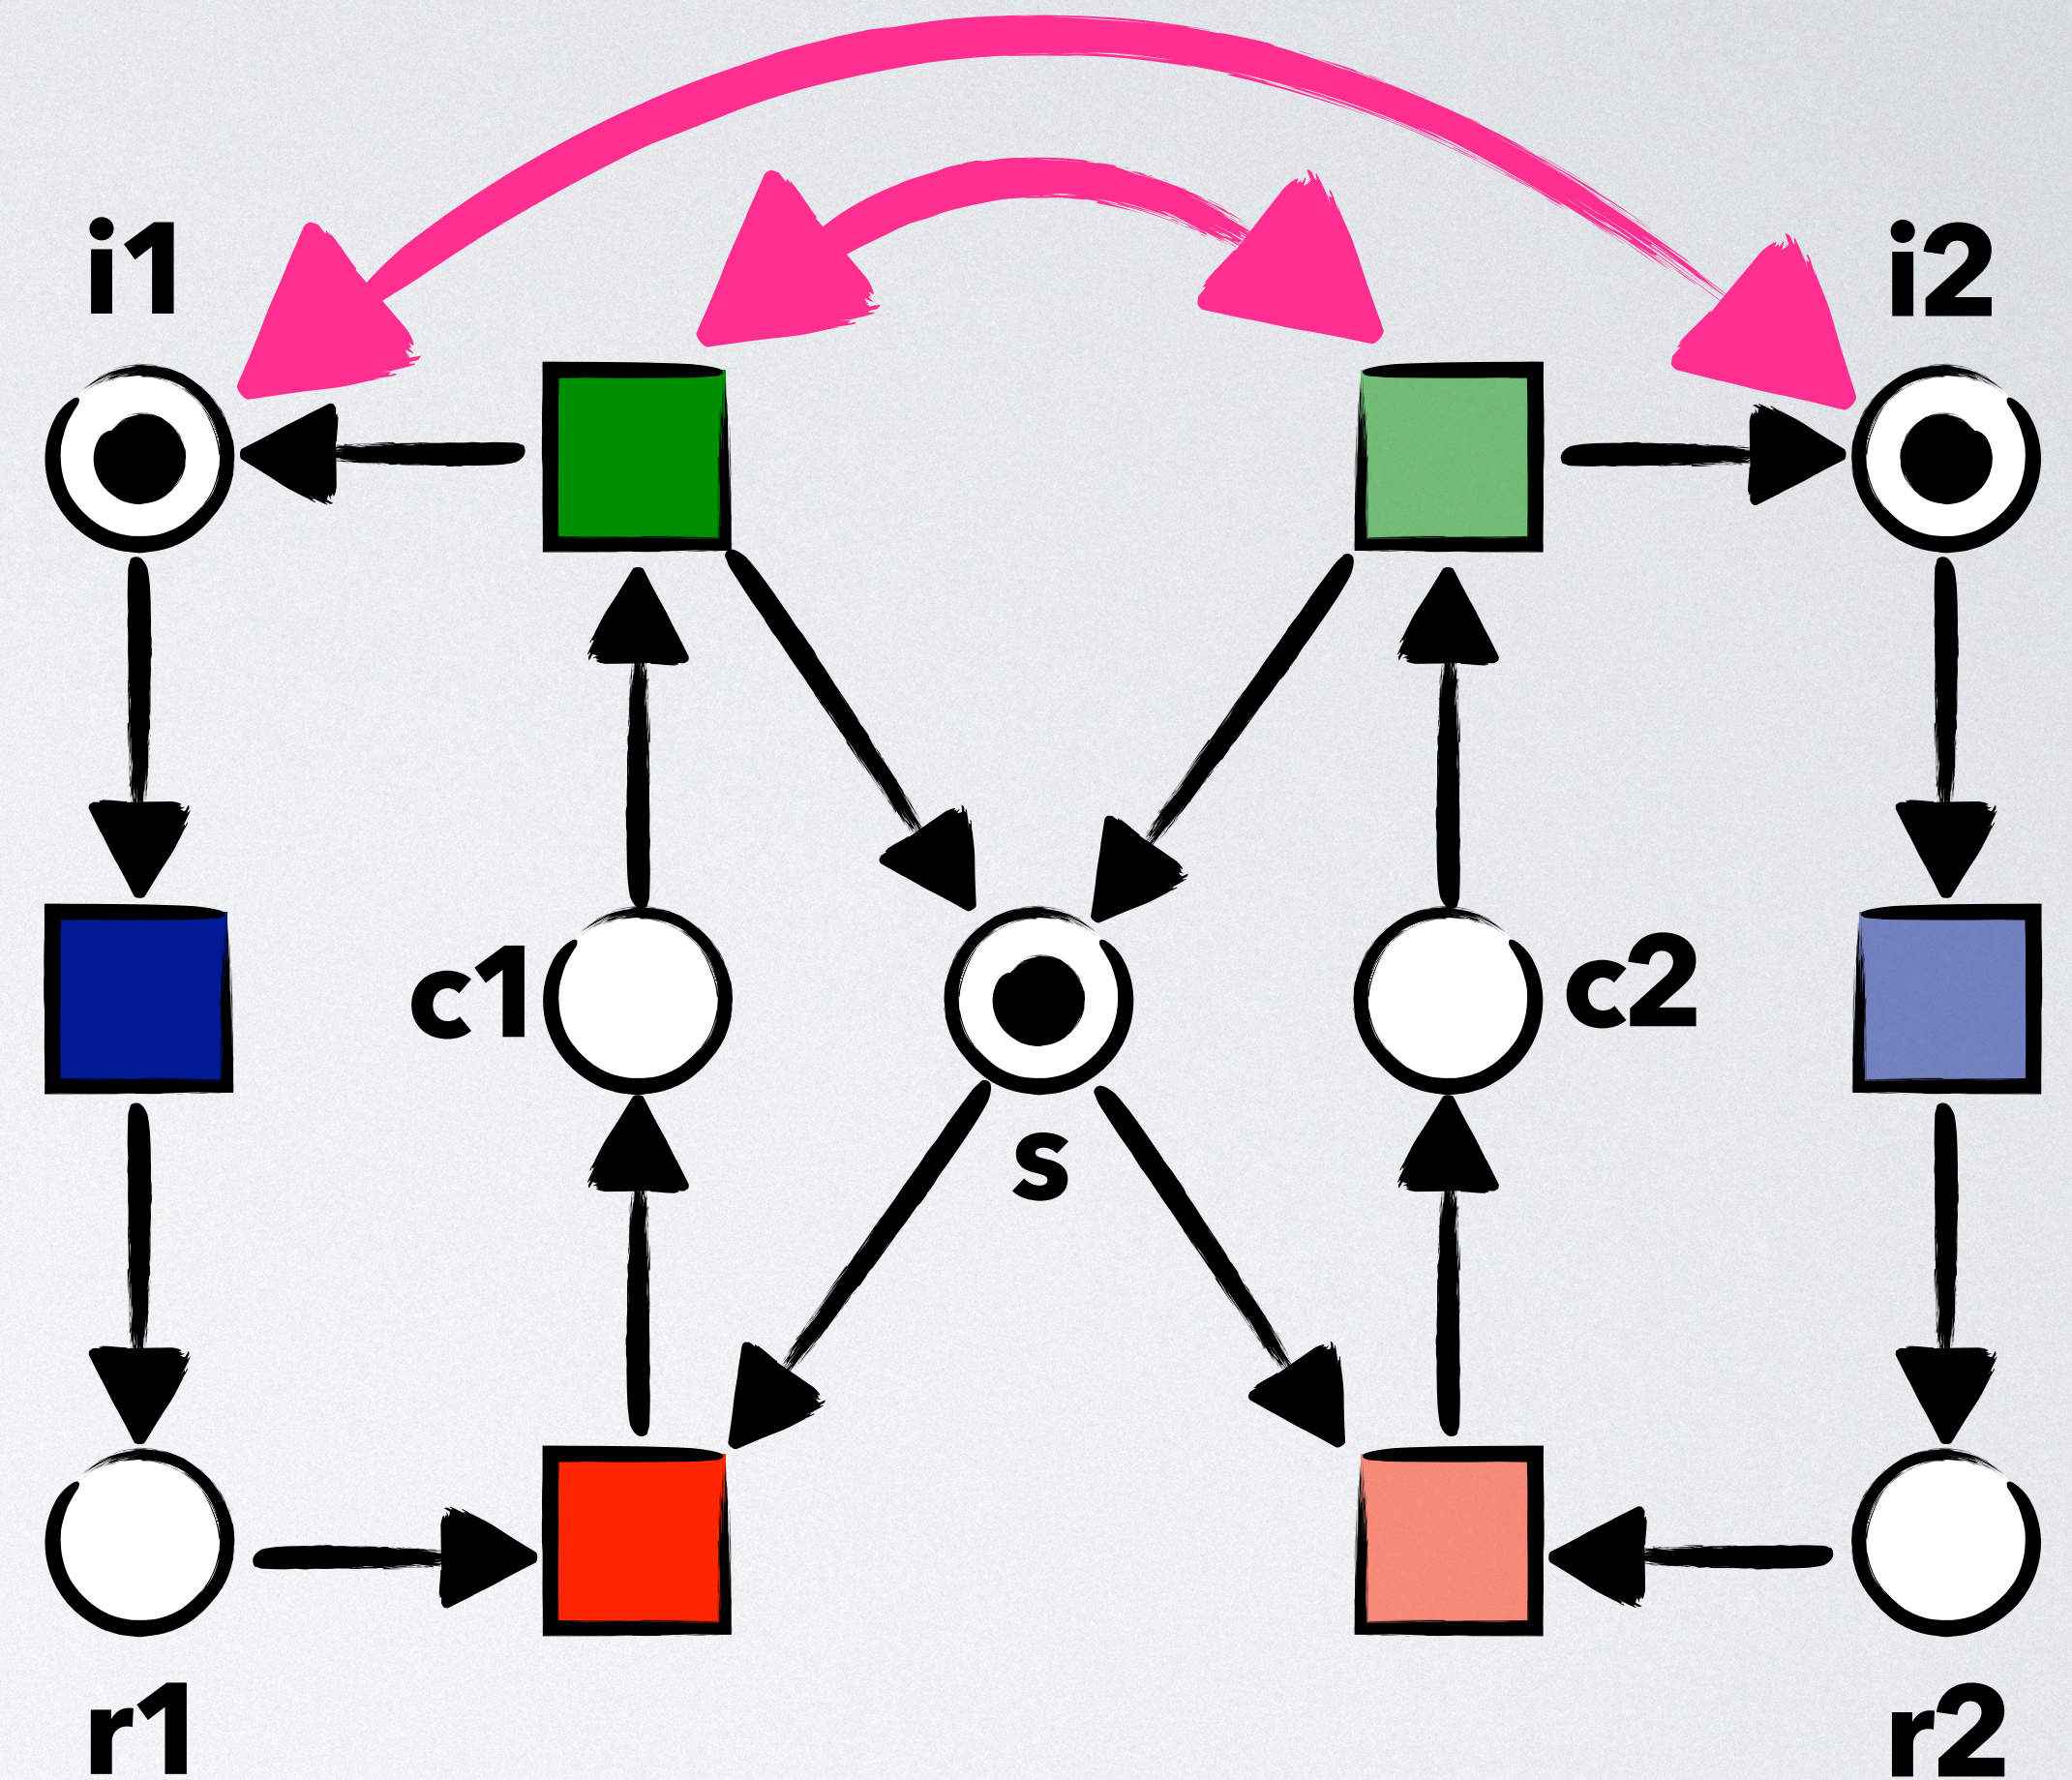
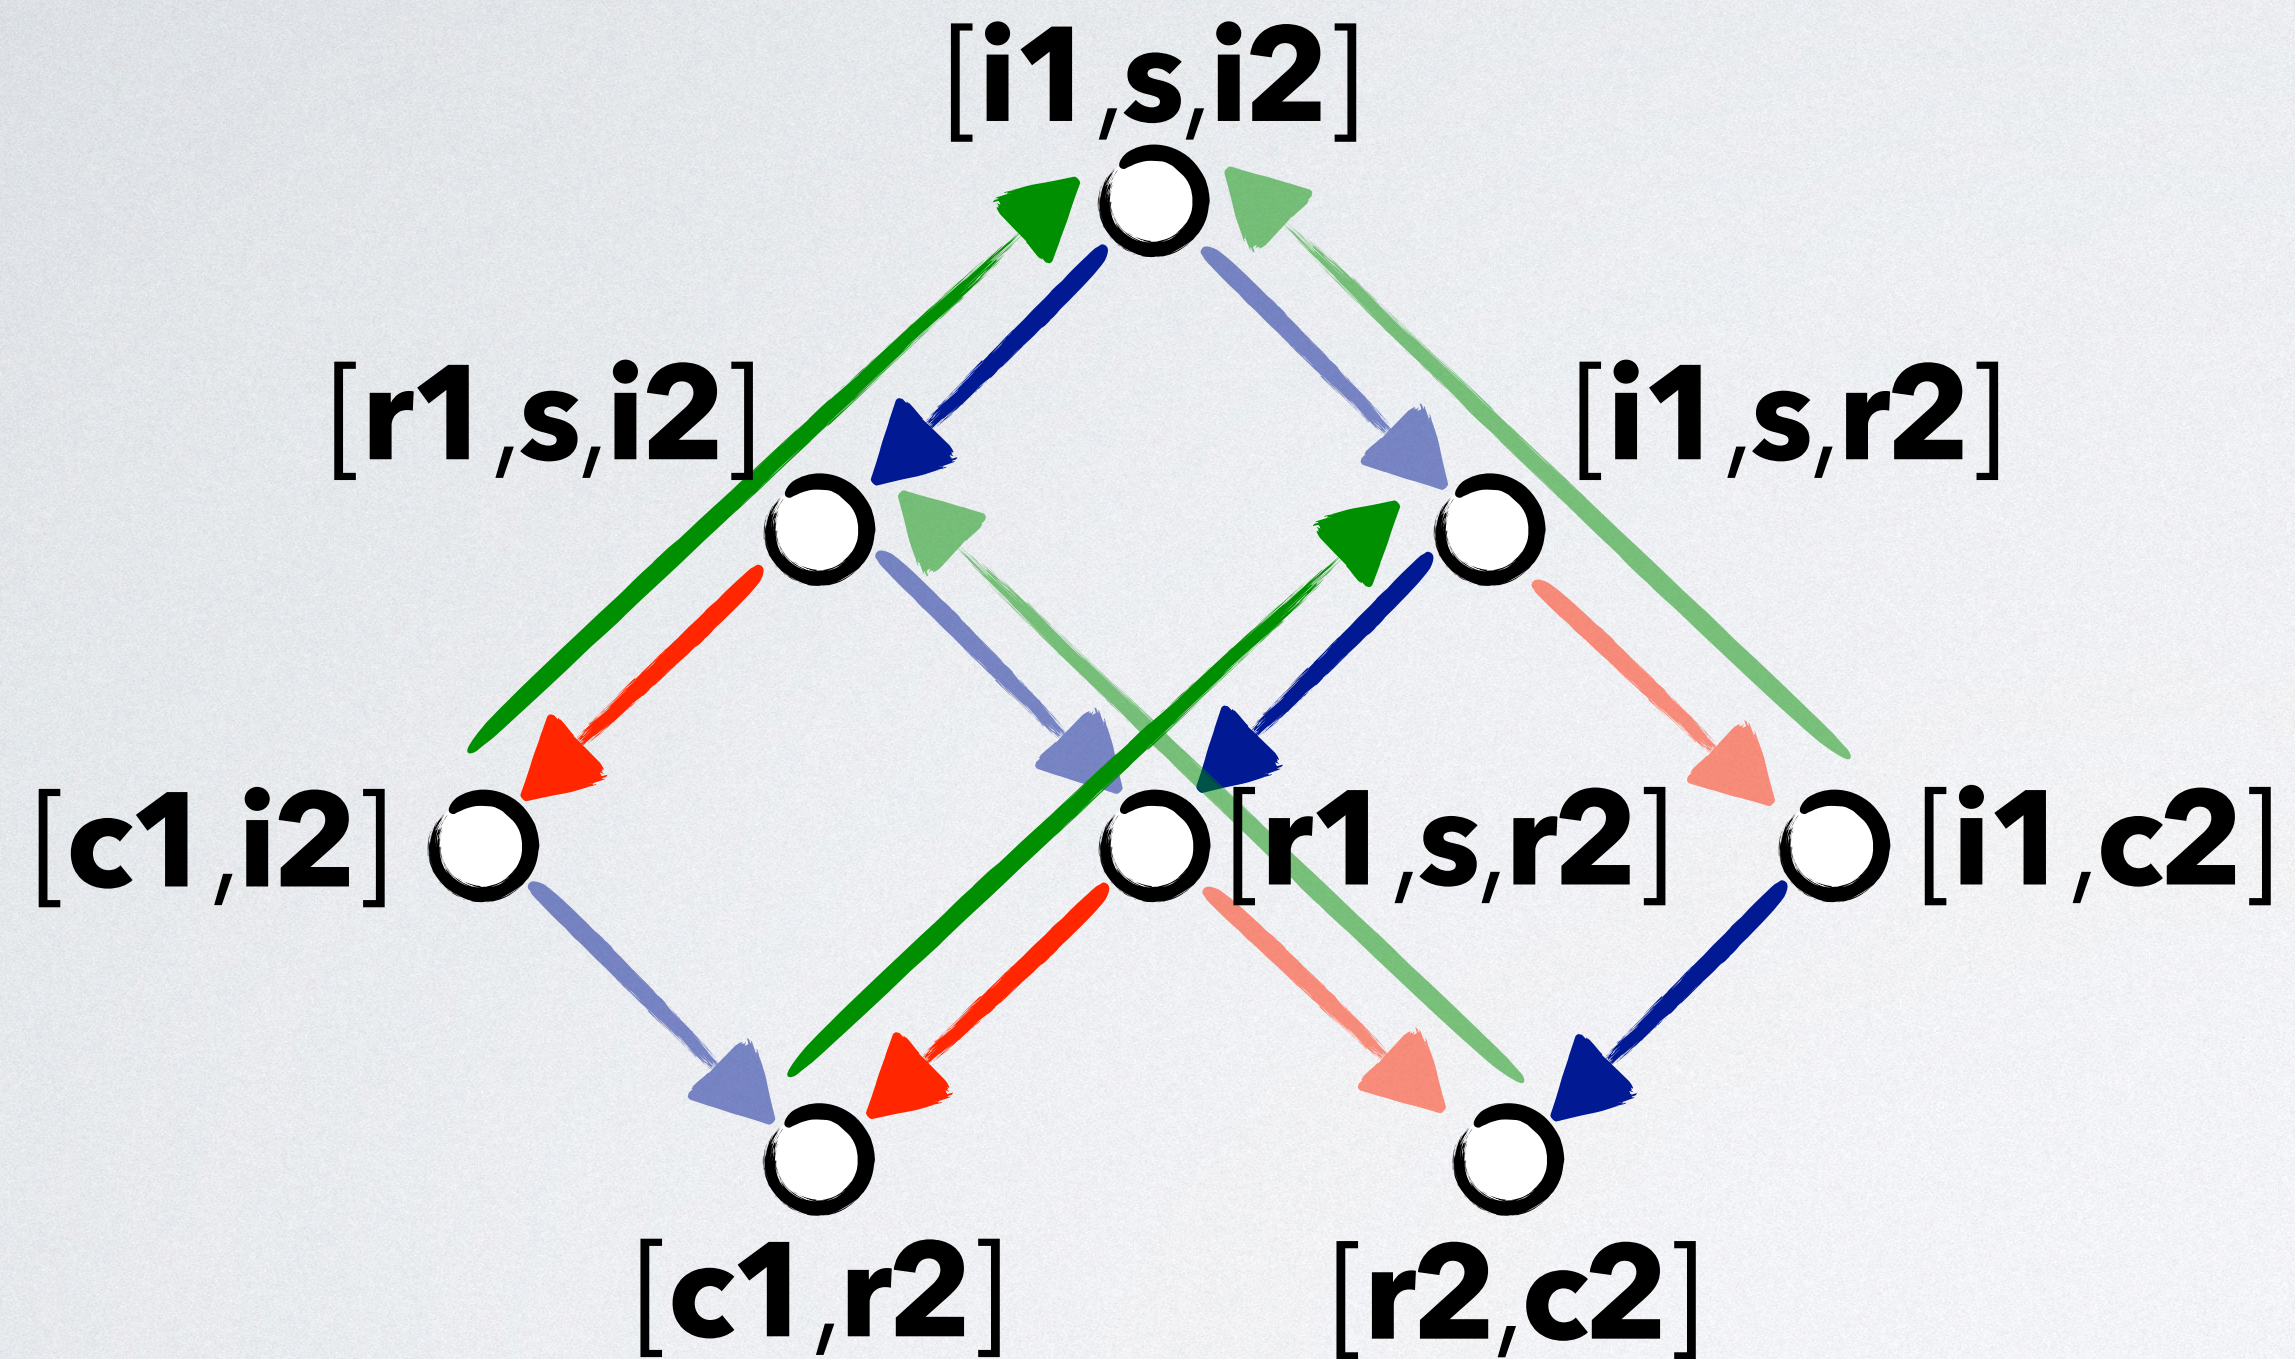


2. Symmetry reduction



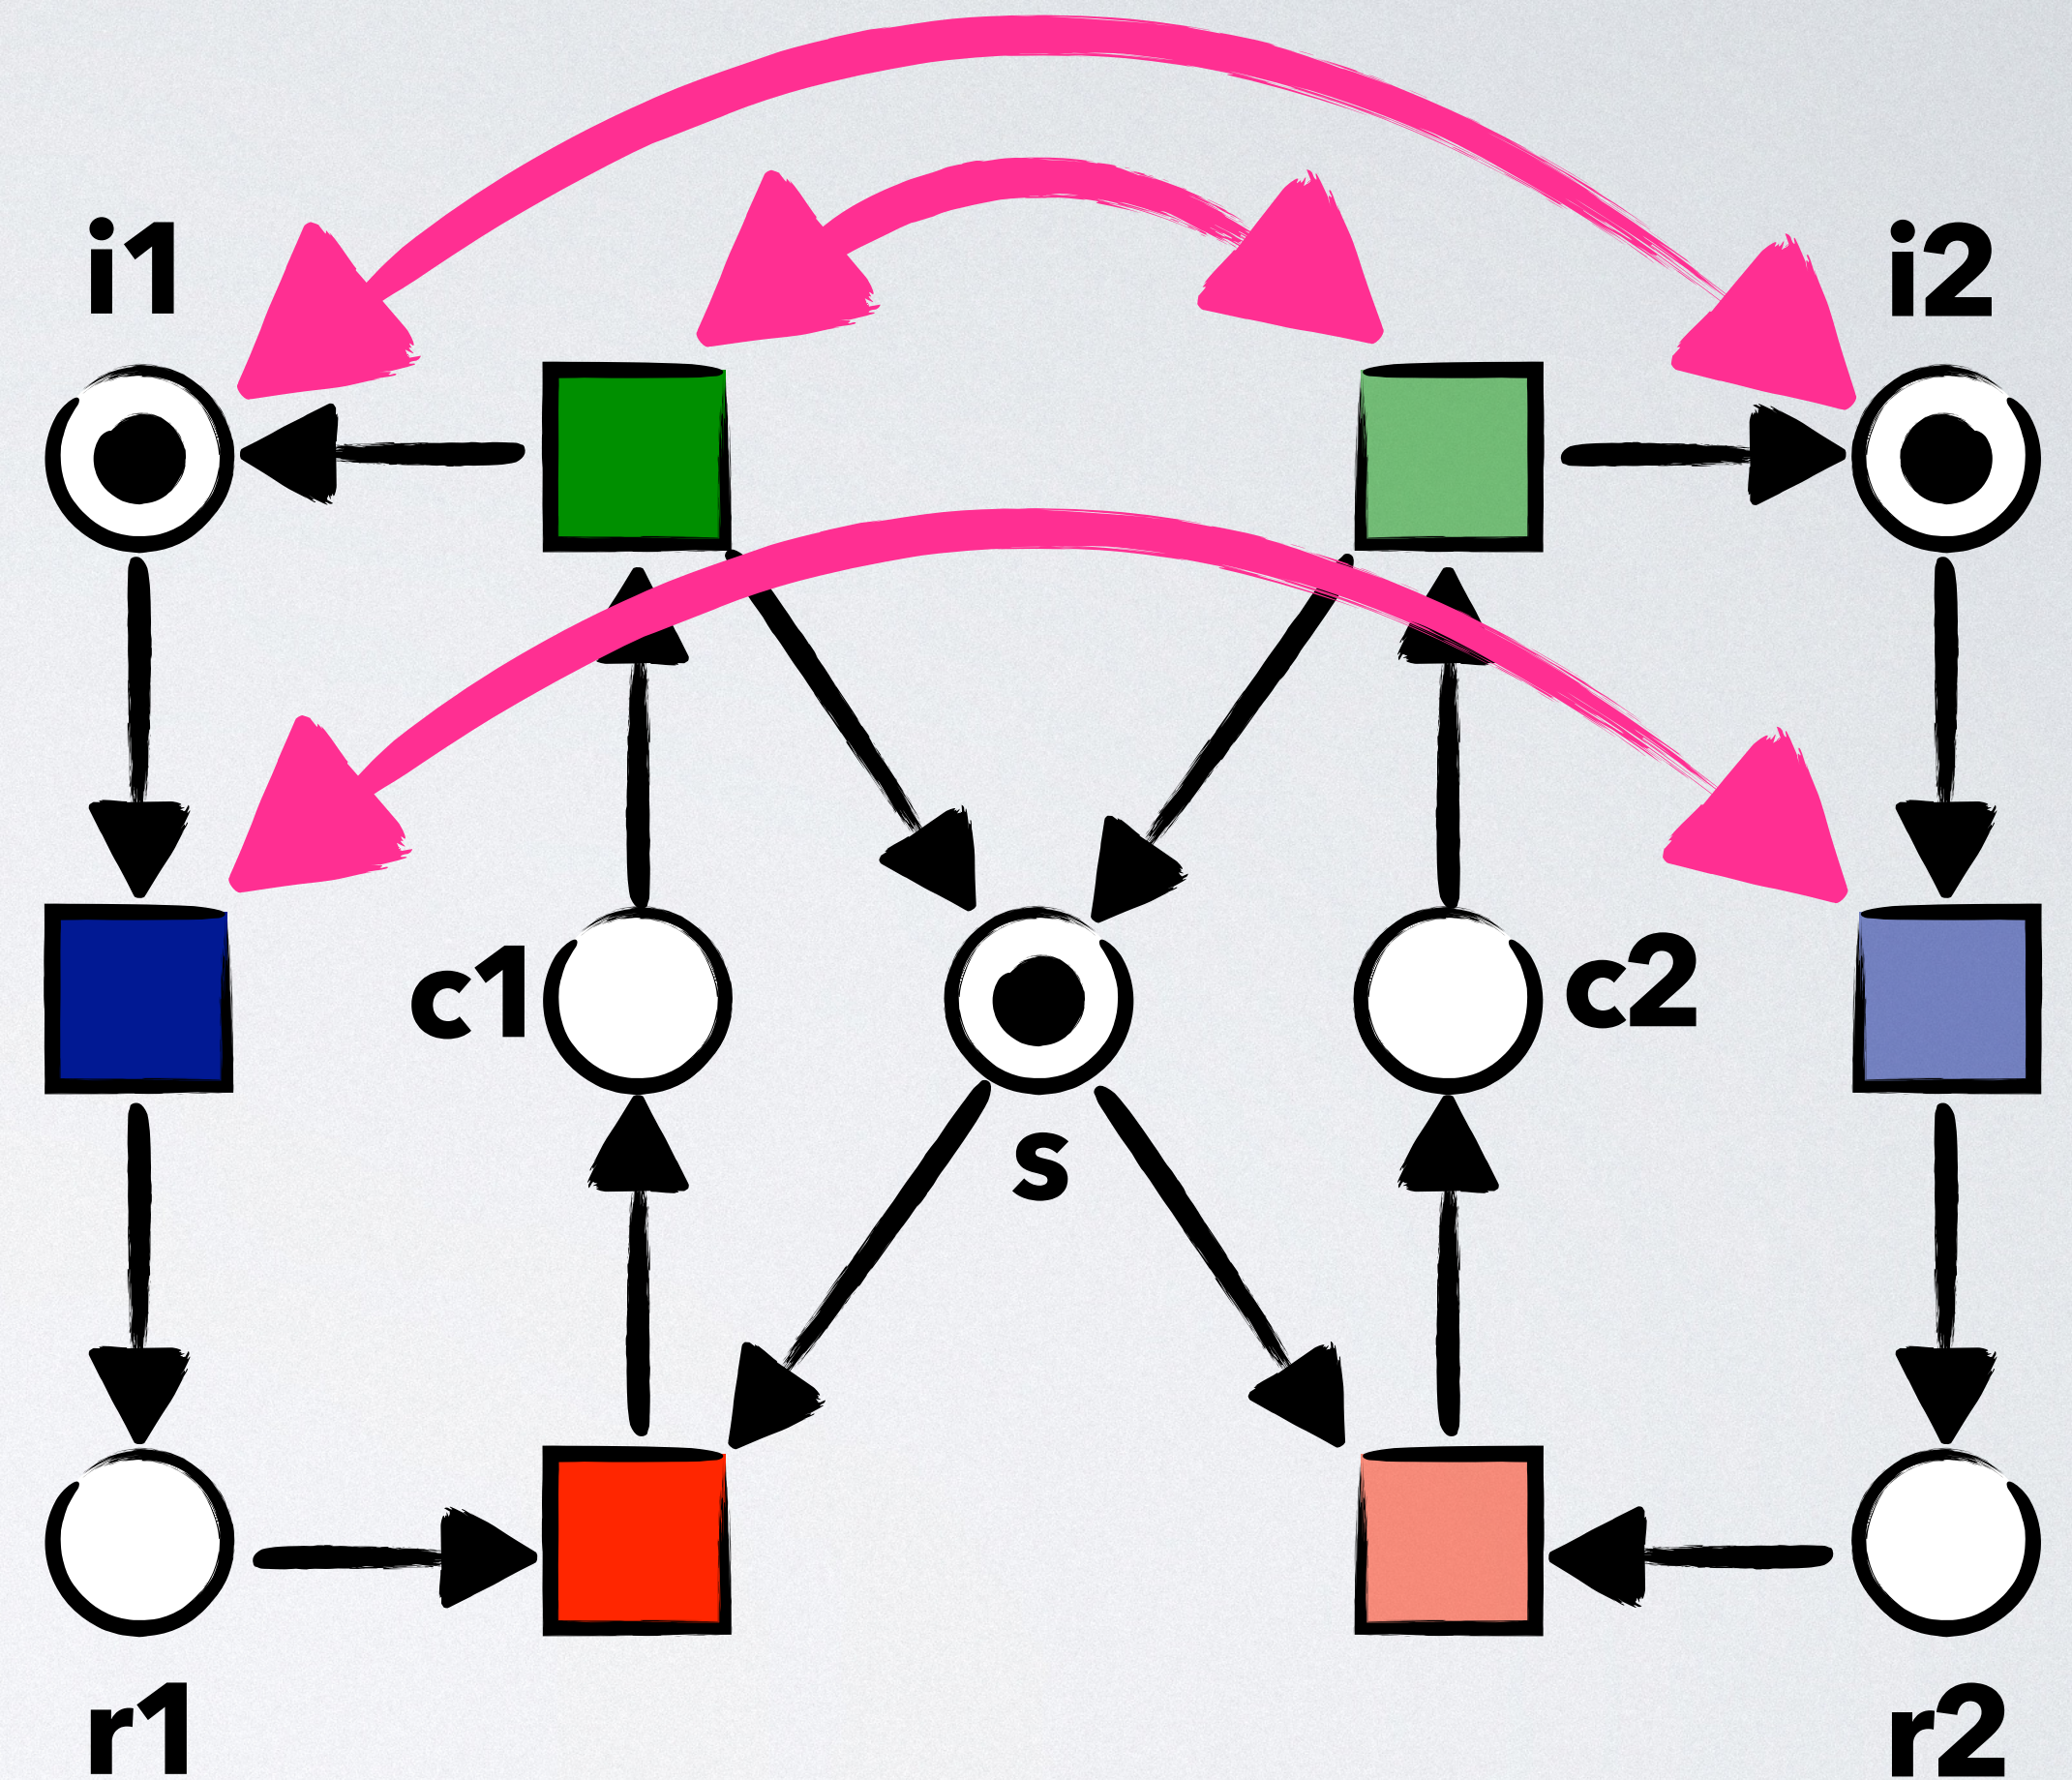
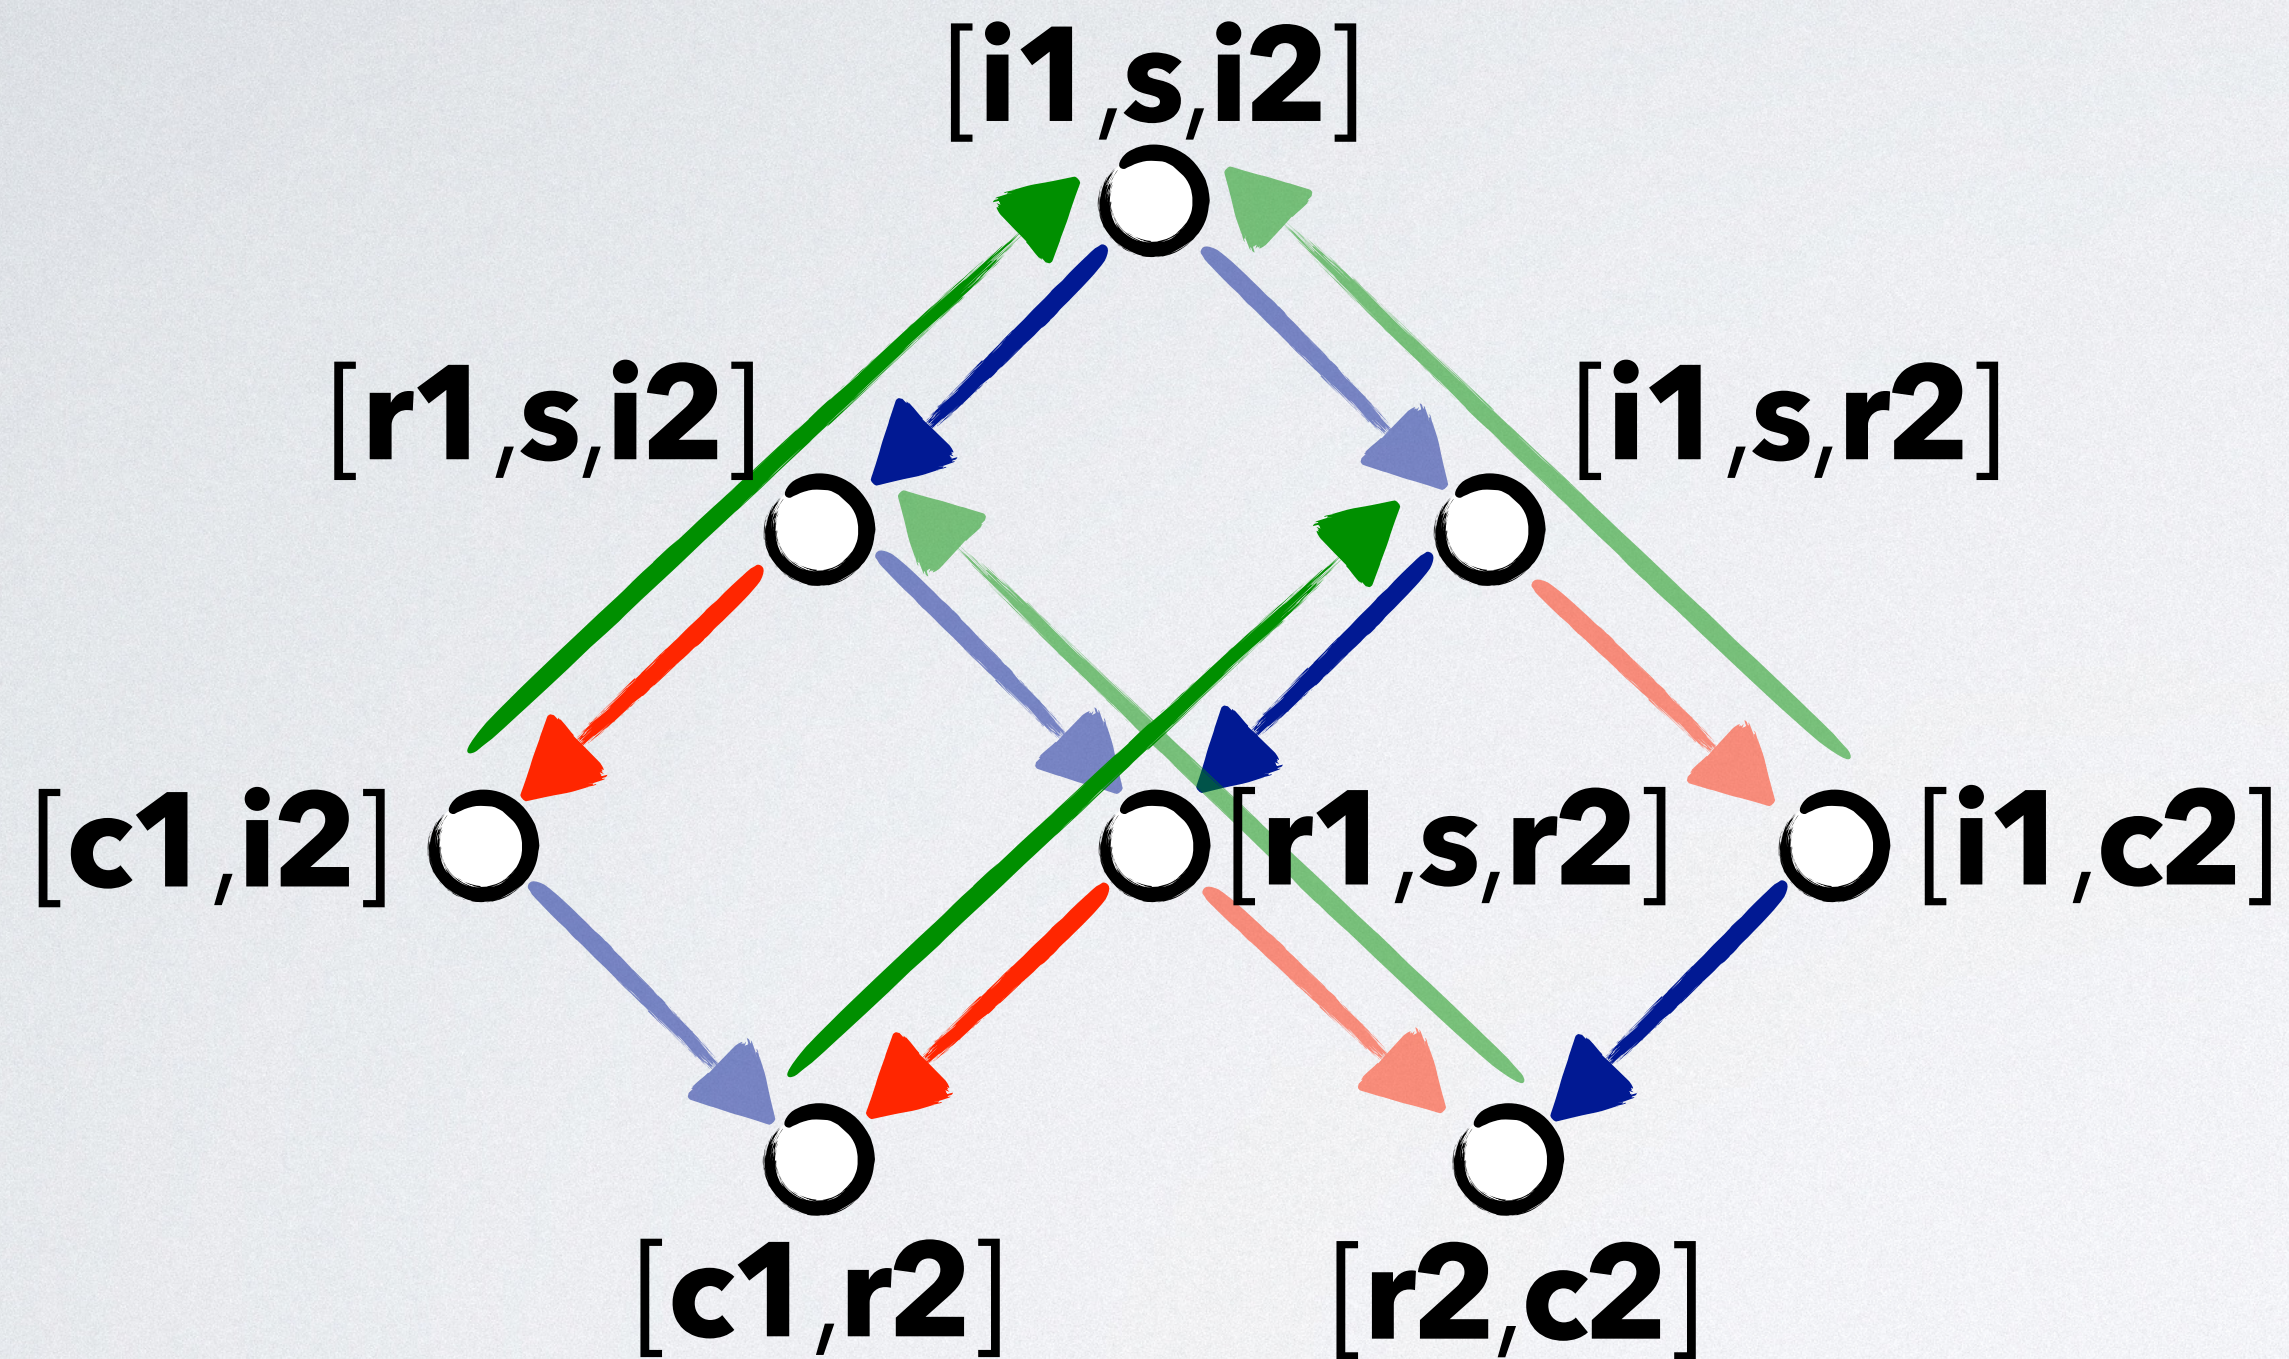
graph
automorphisms

2. Symmetry reduction



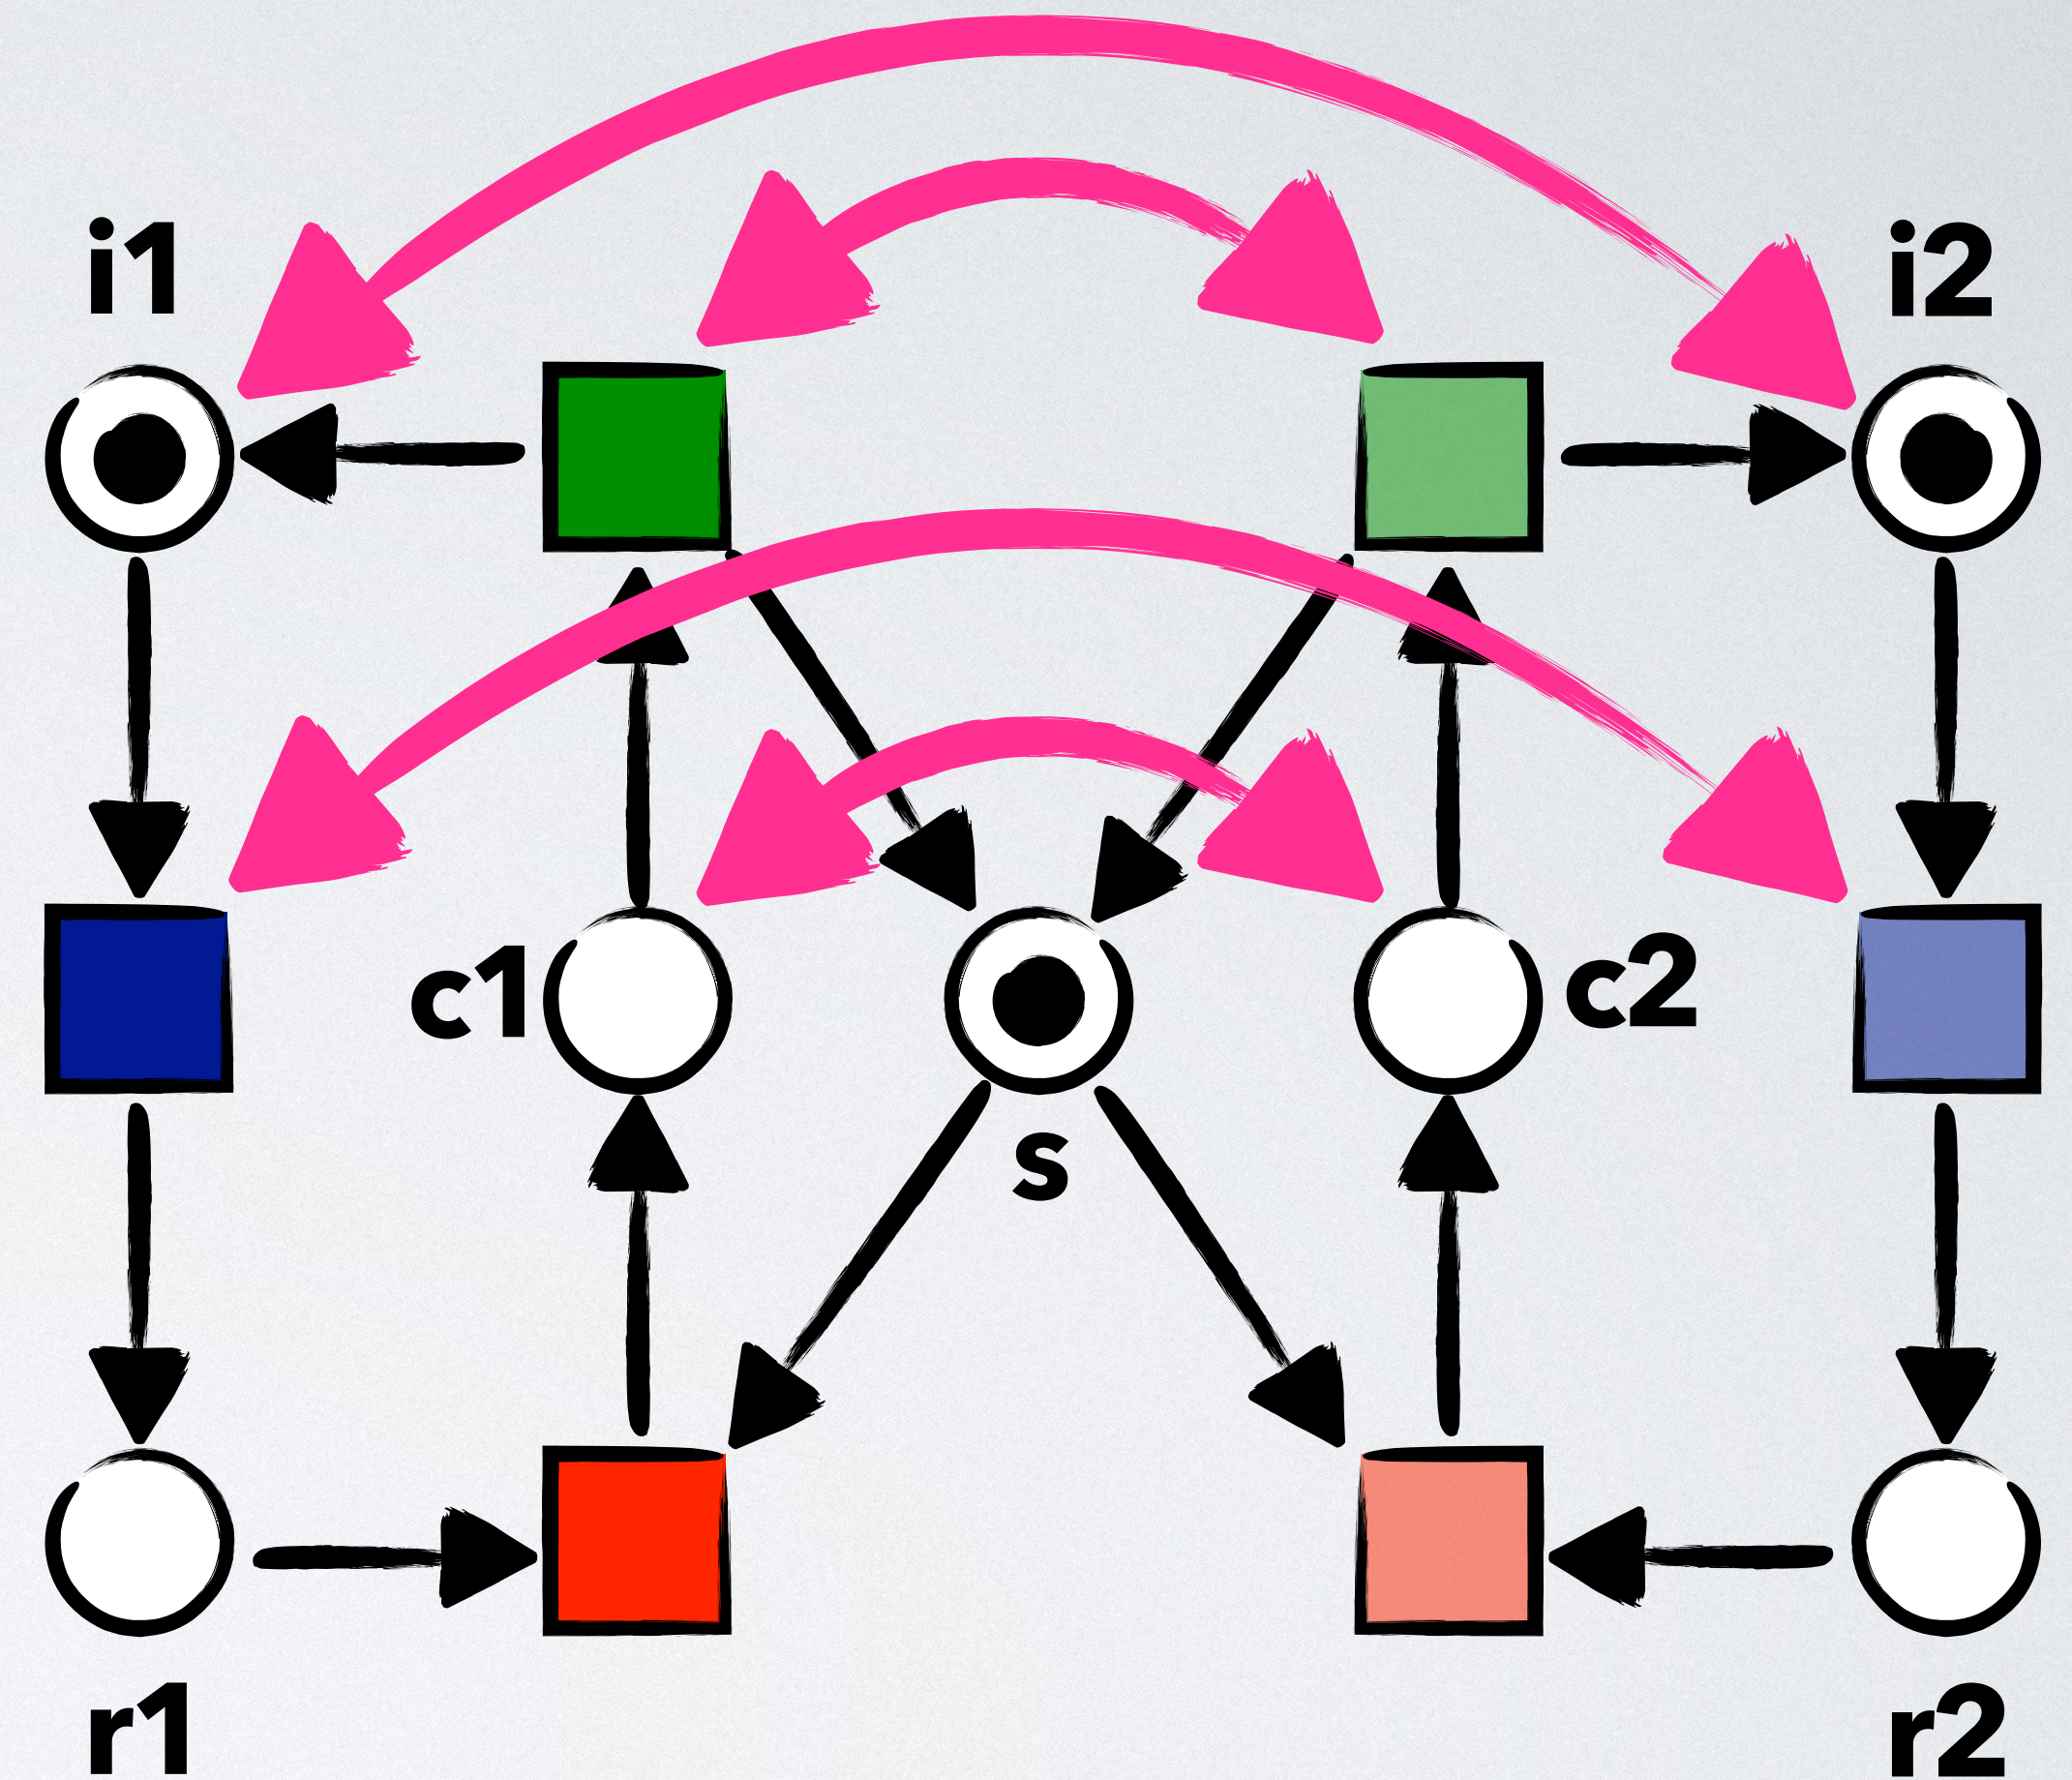
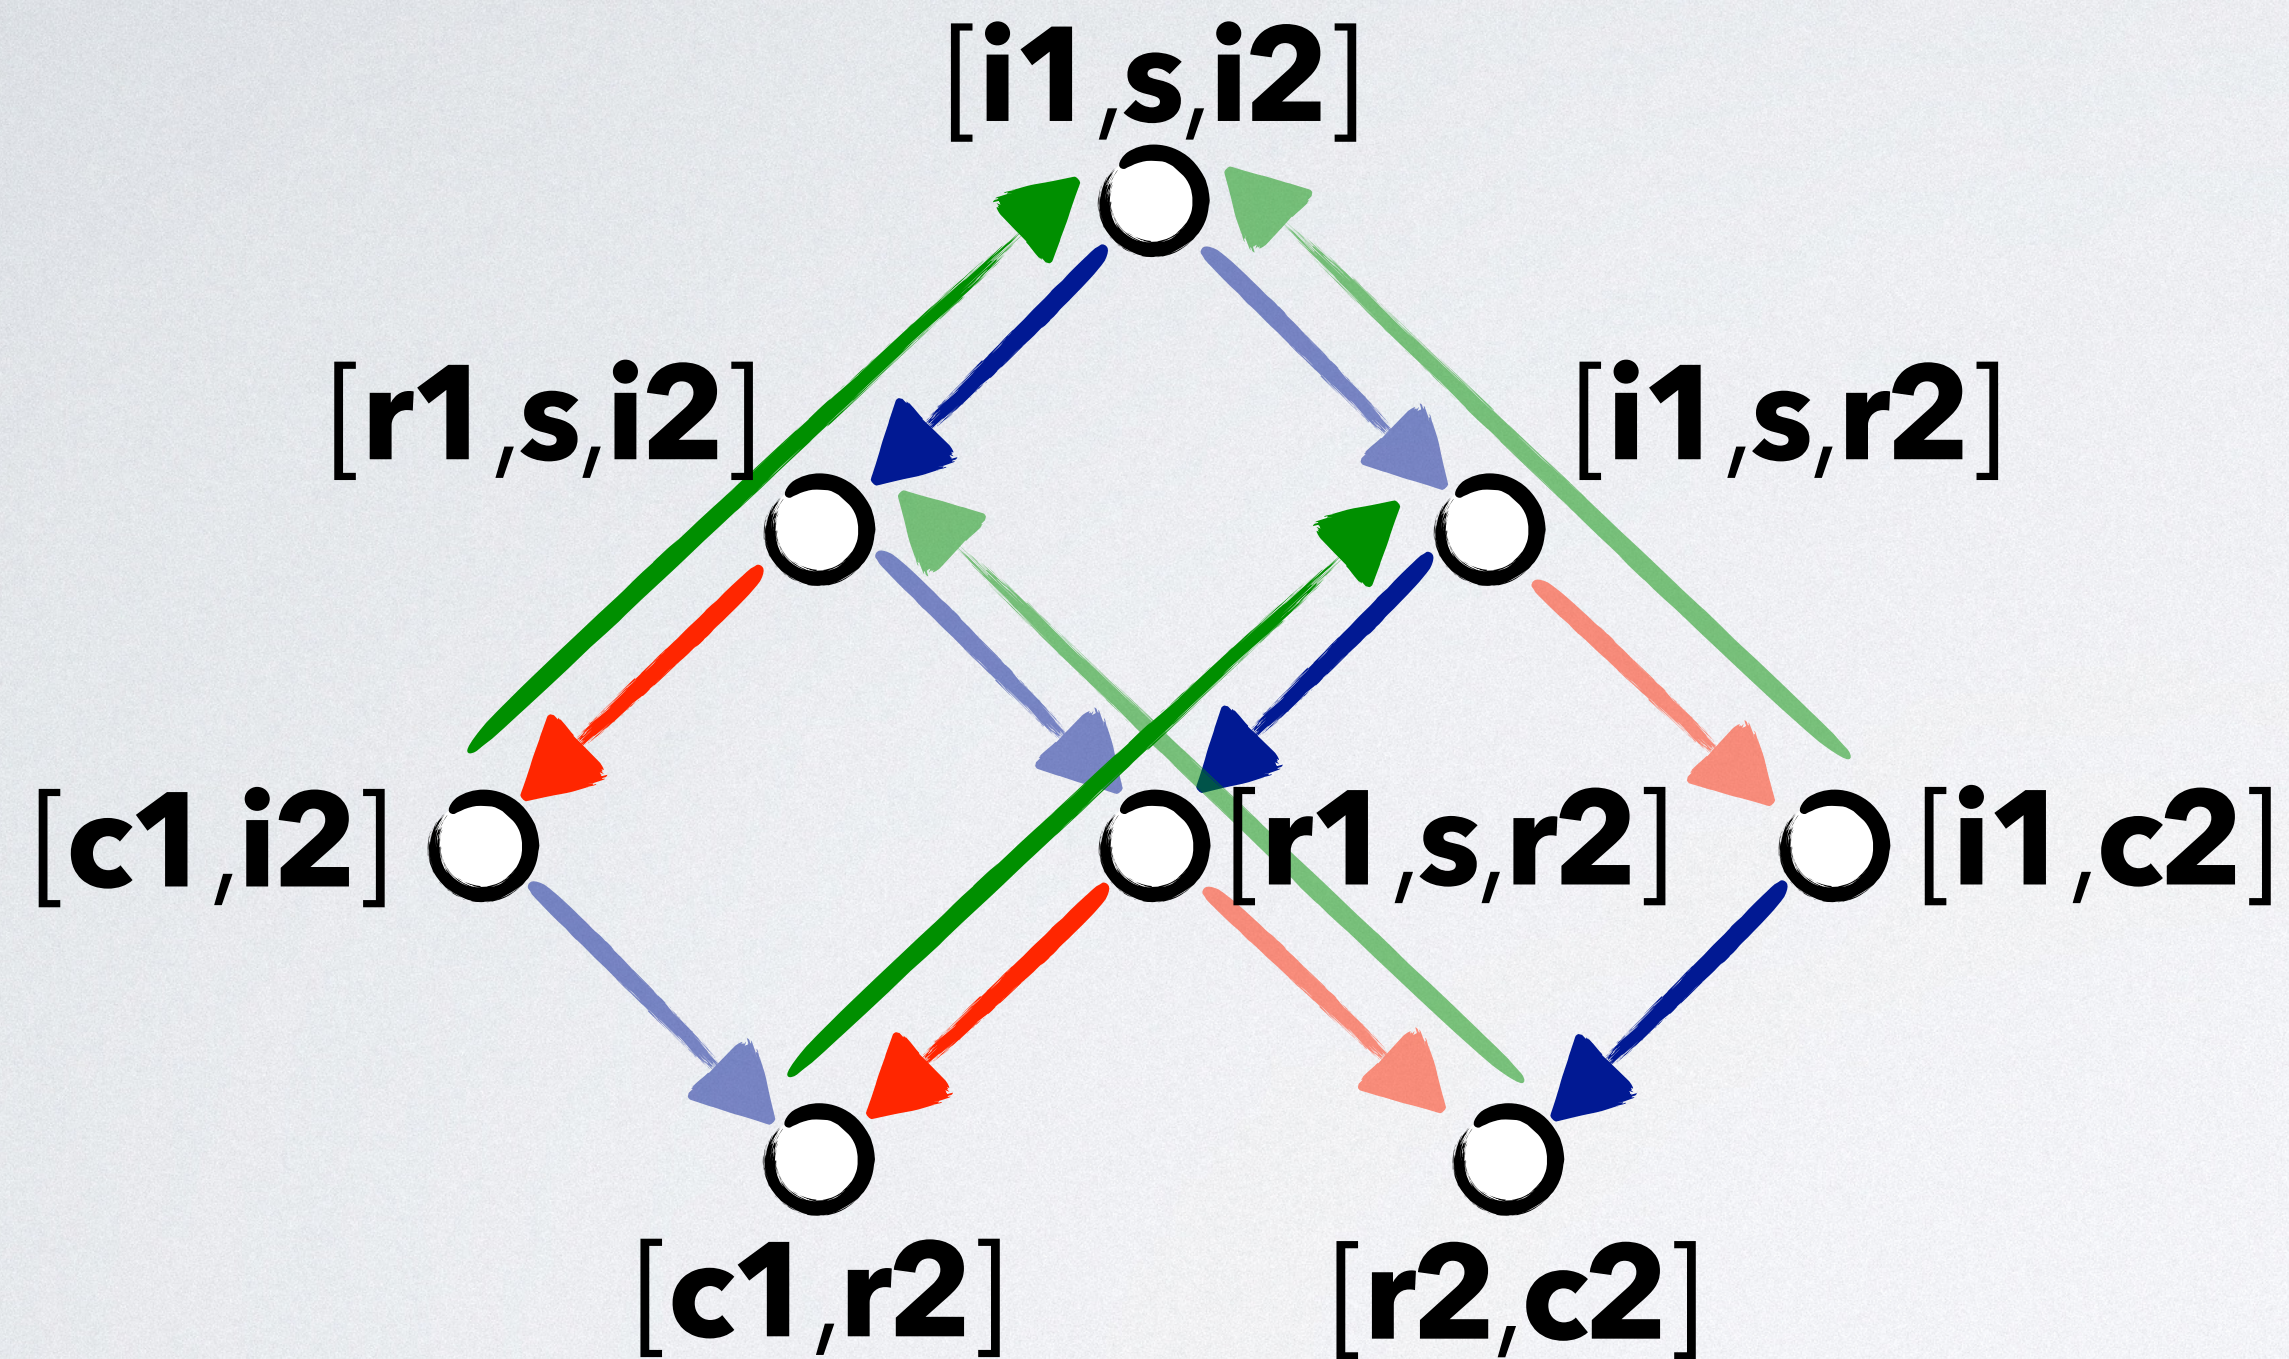
graph
automorphisms

2. Symmetry reduction



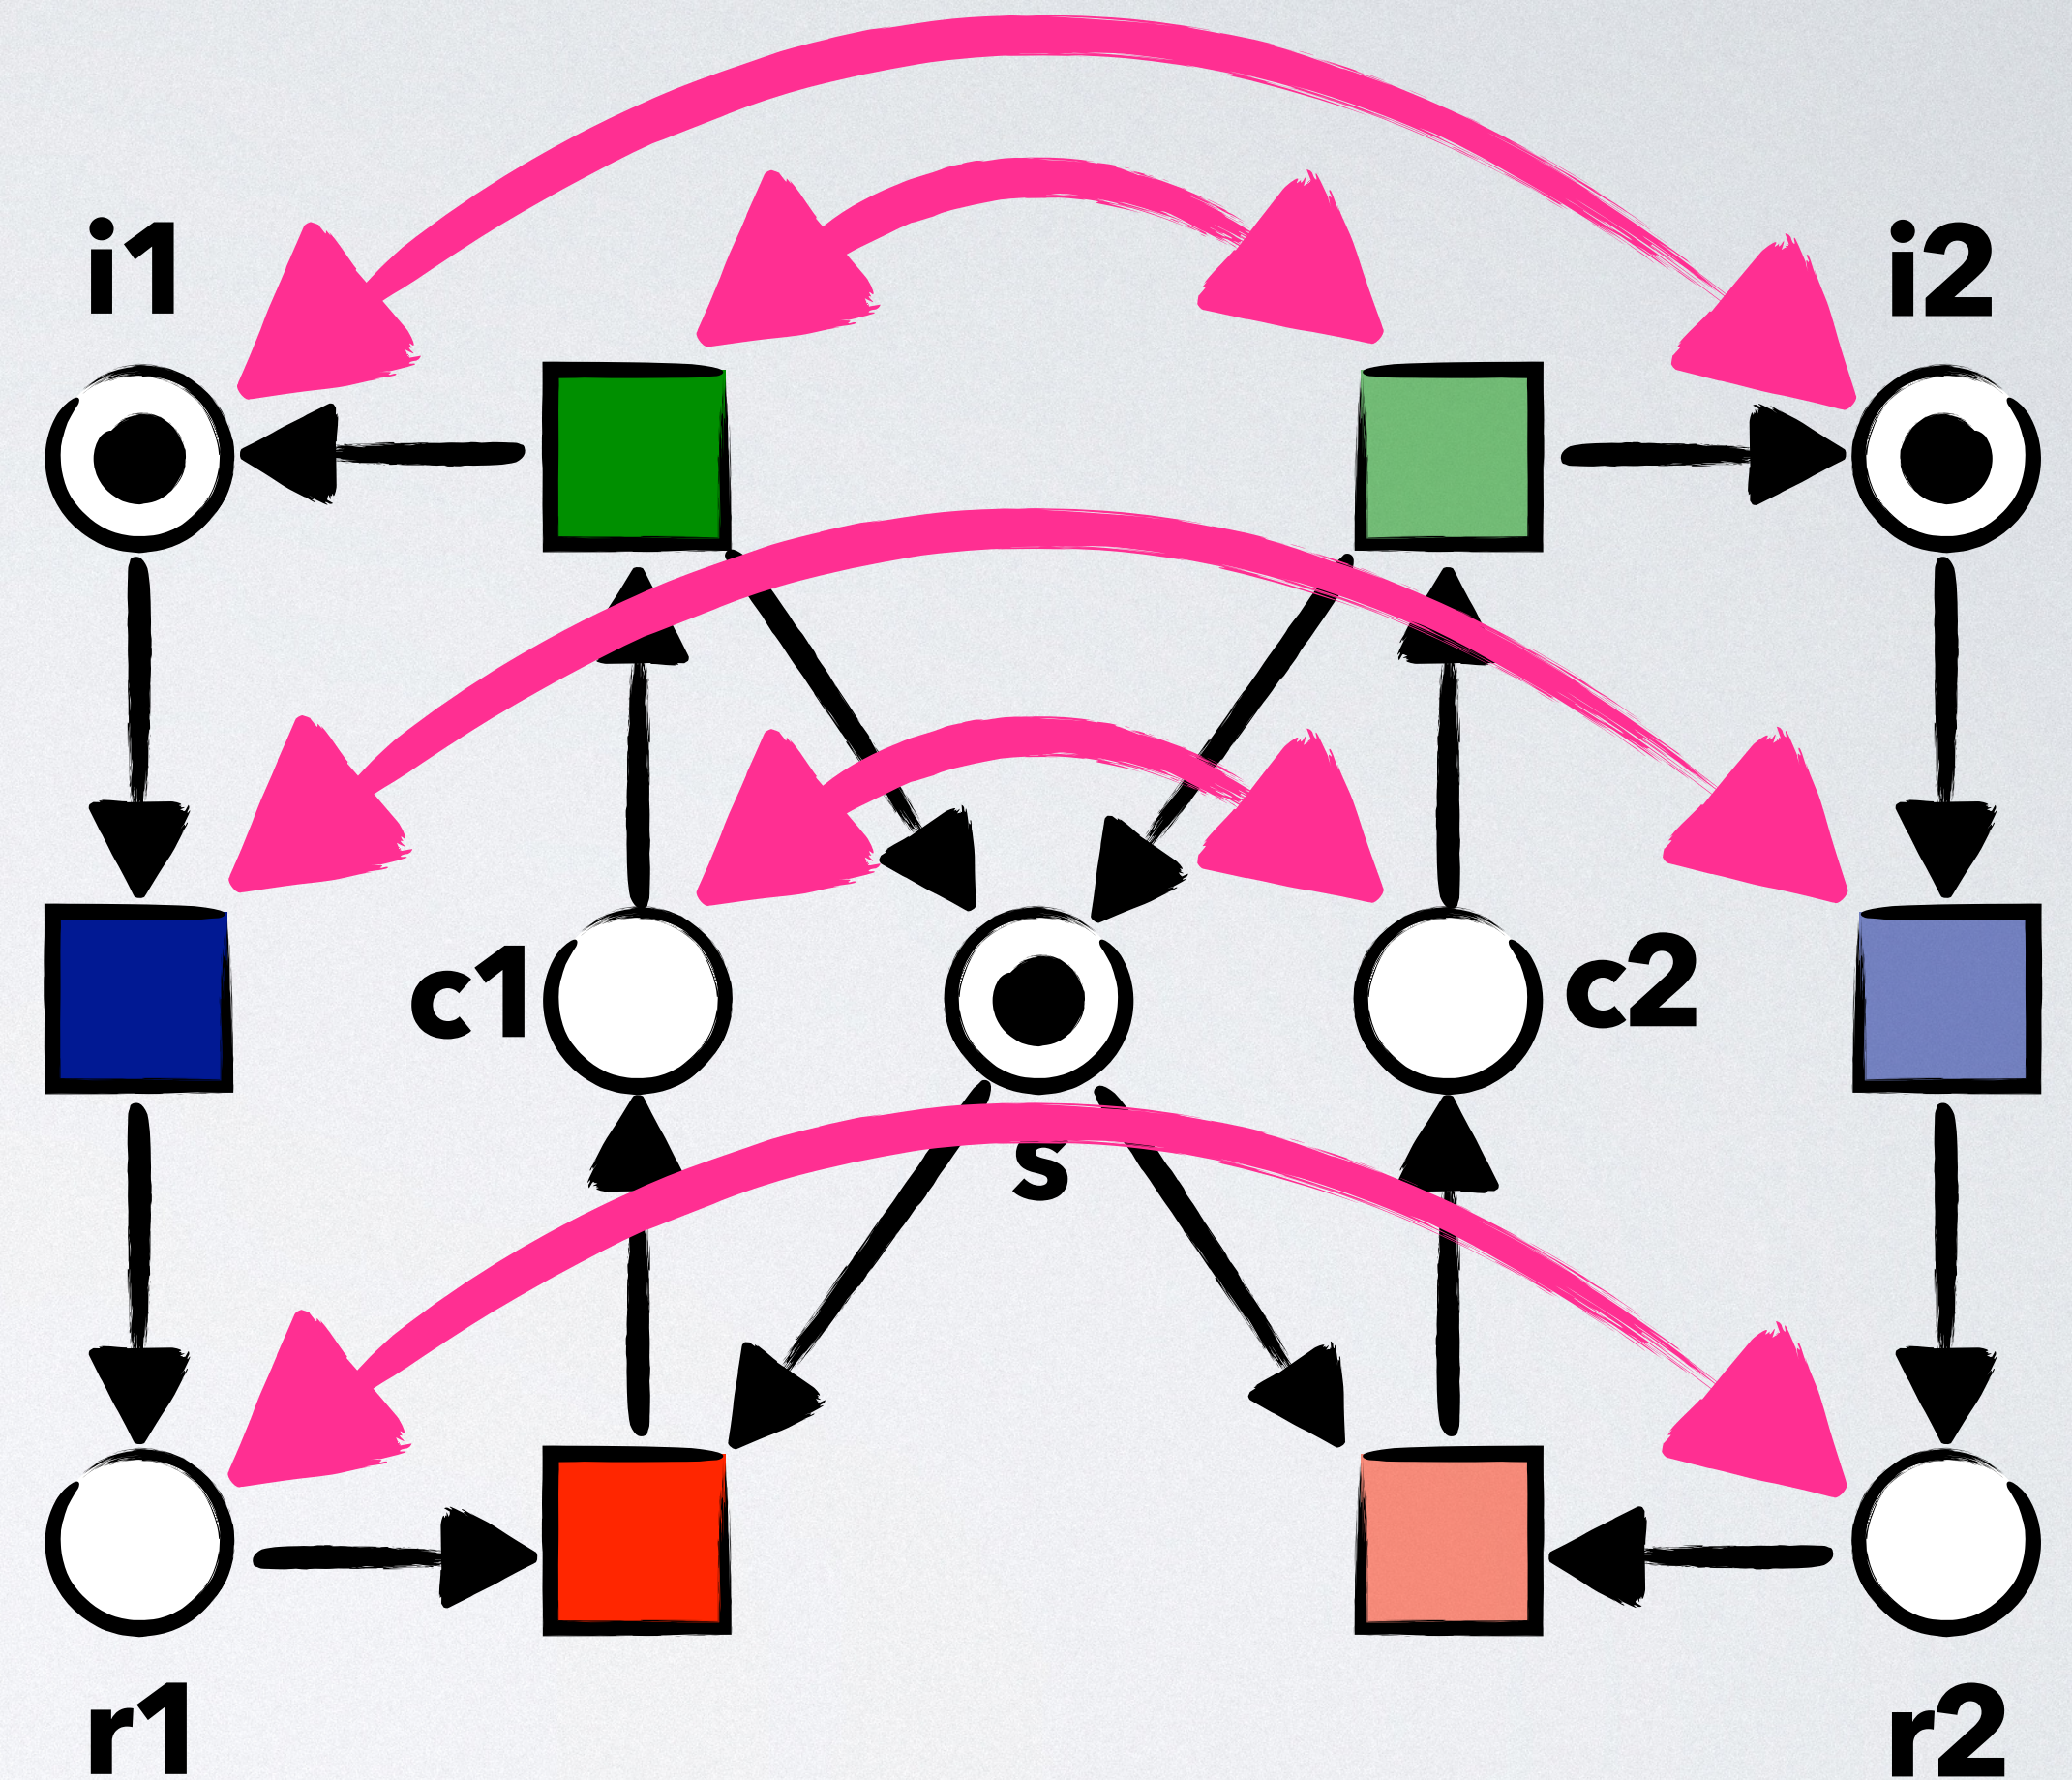
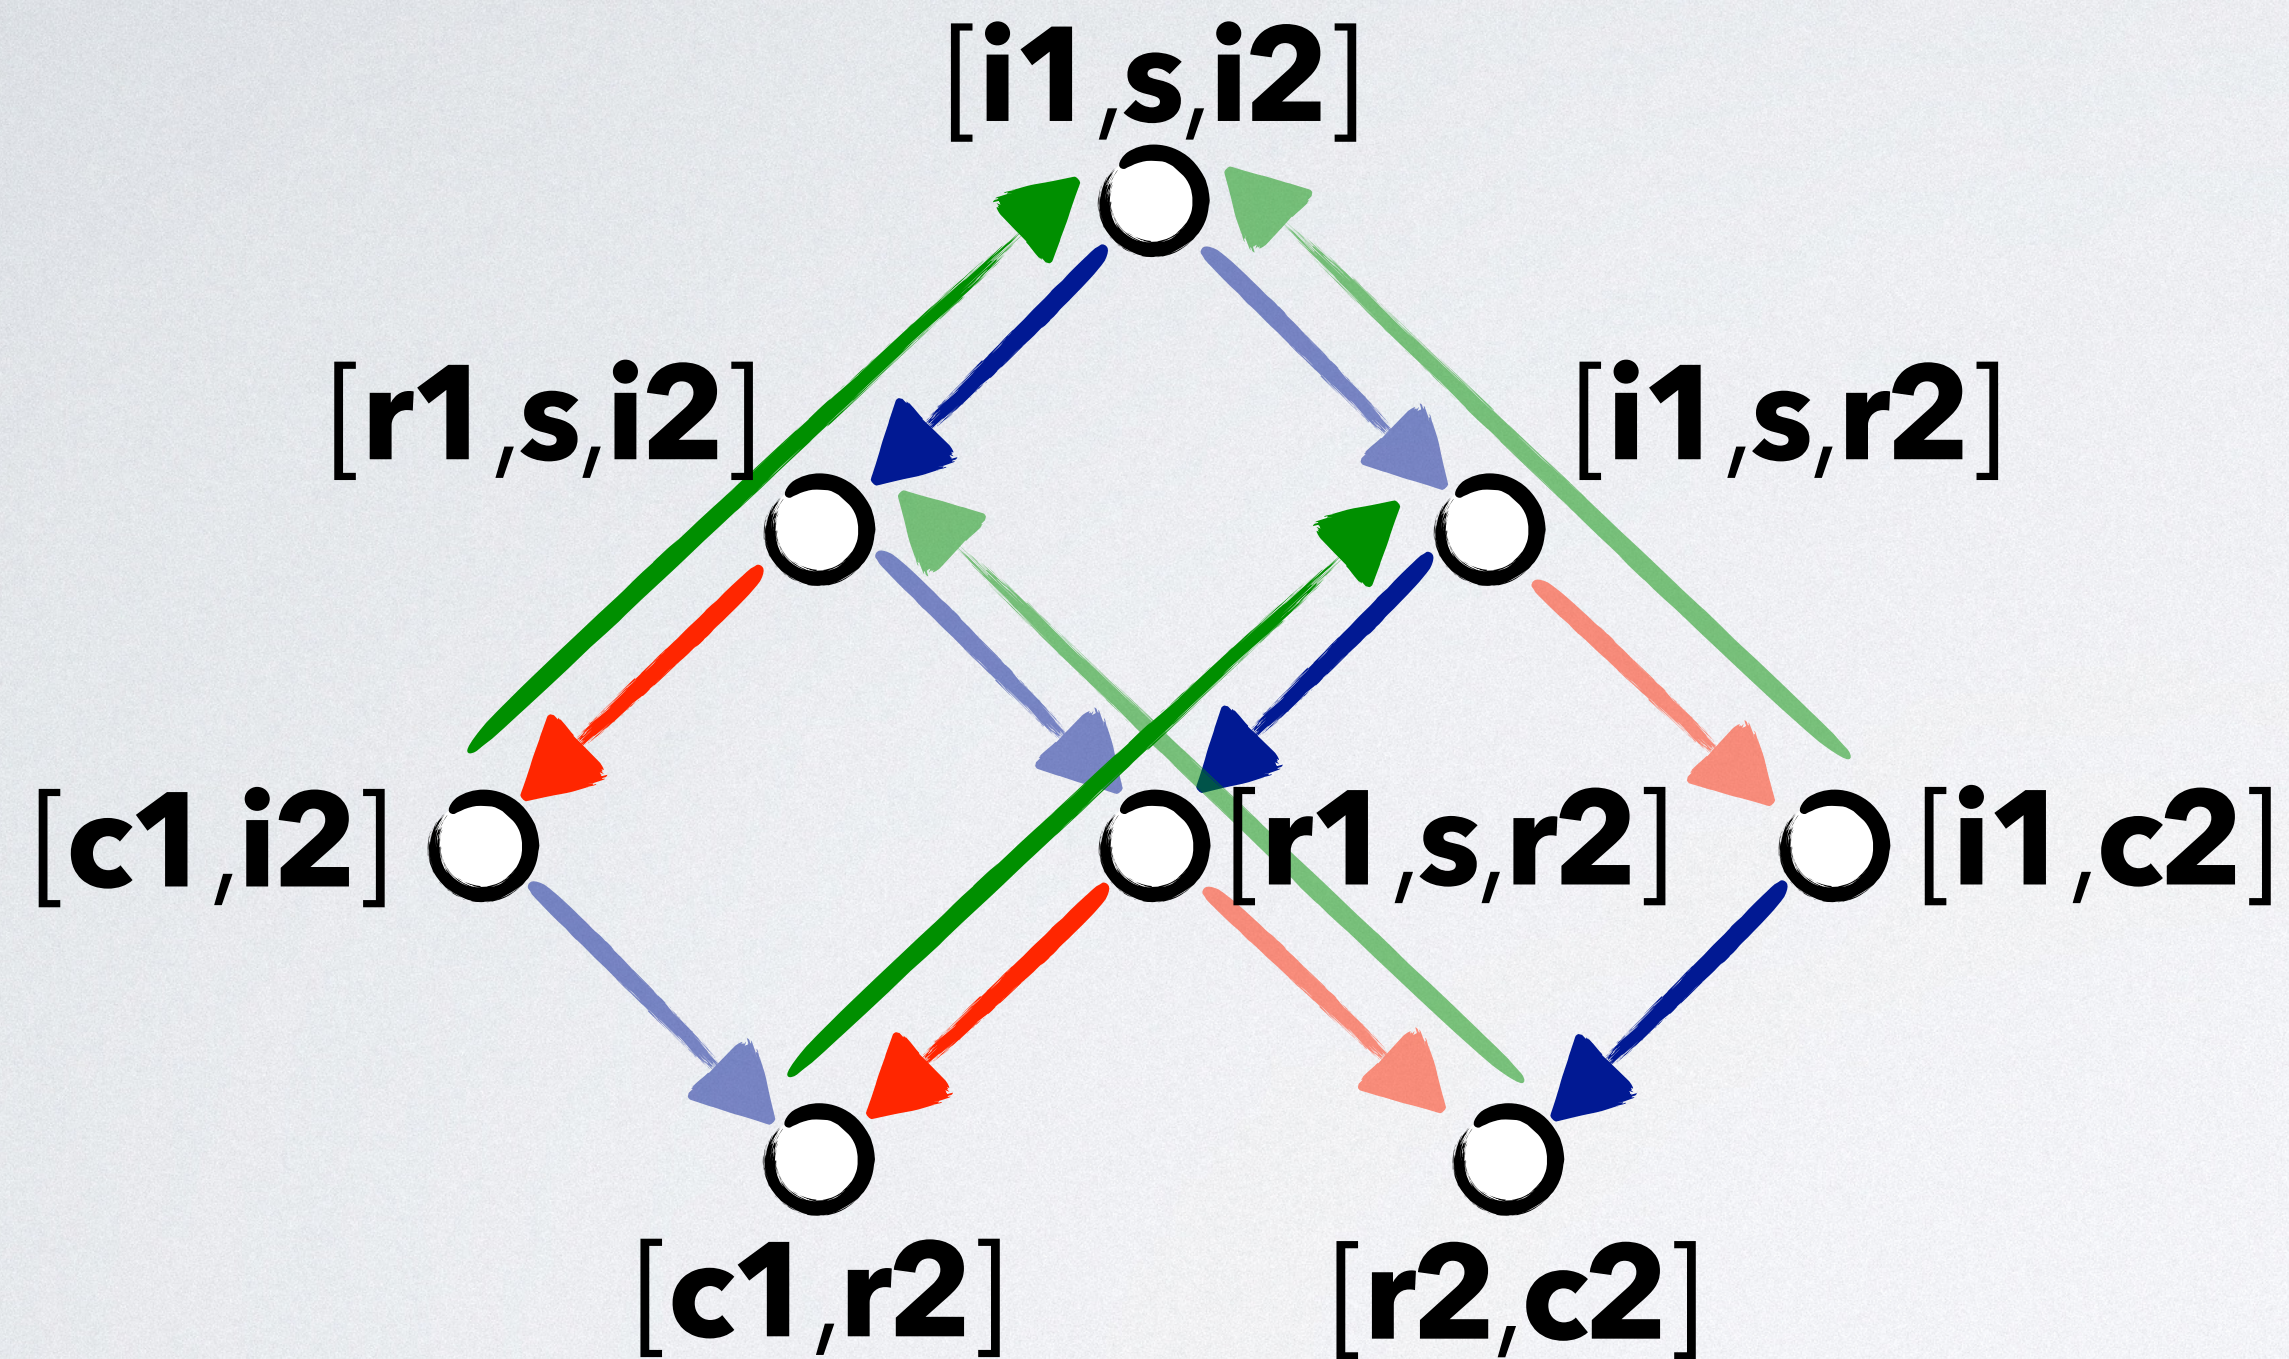
graph
automorphisms

2. Symmetry reduction



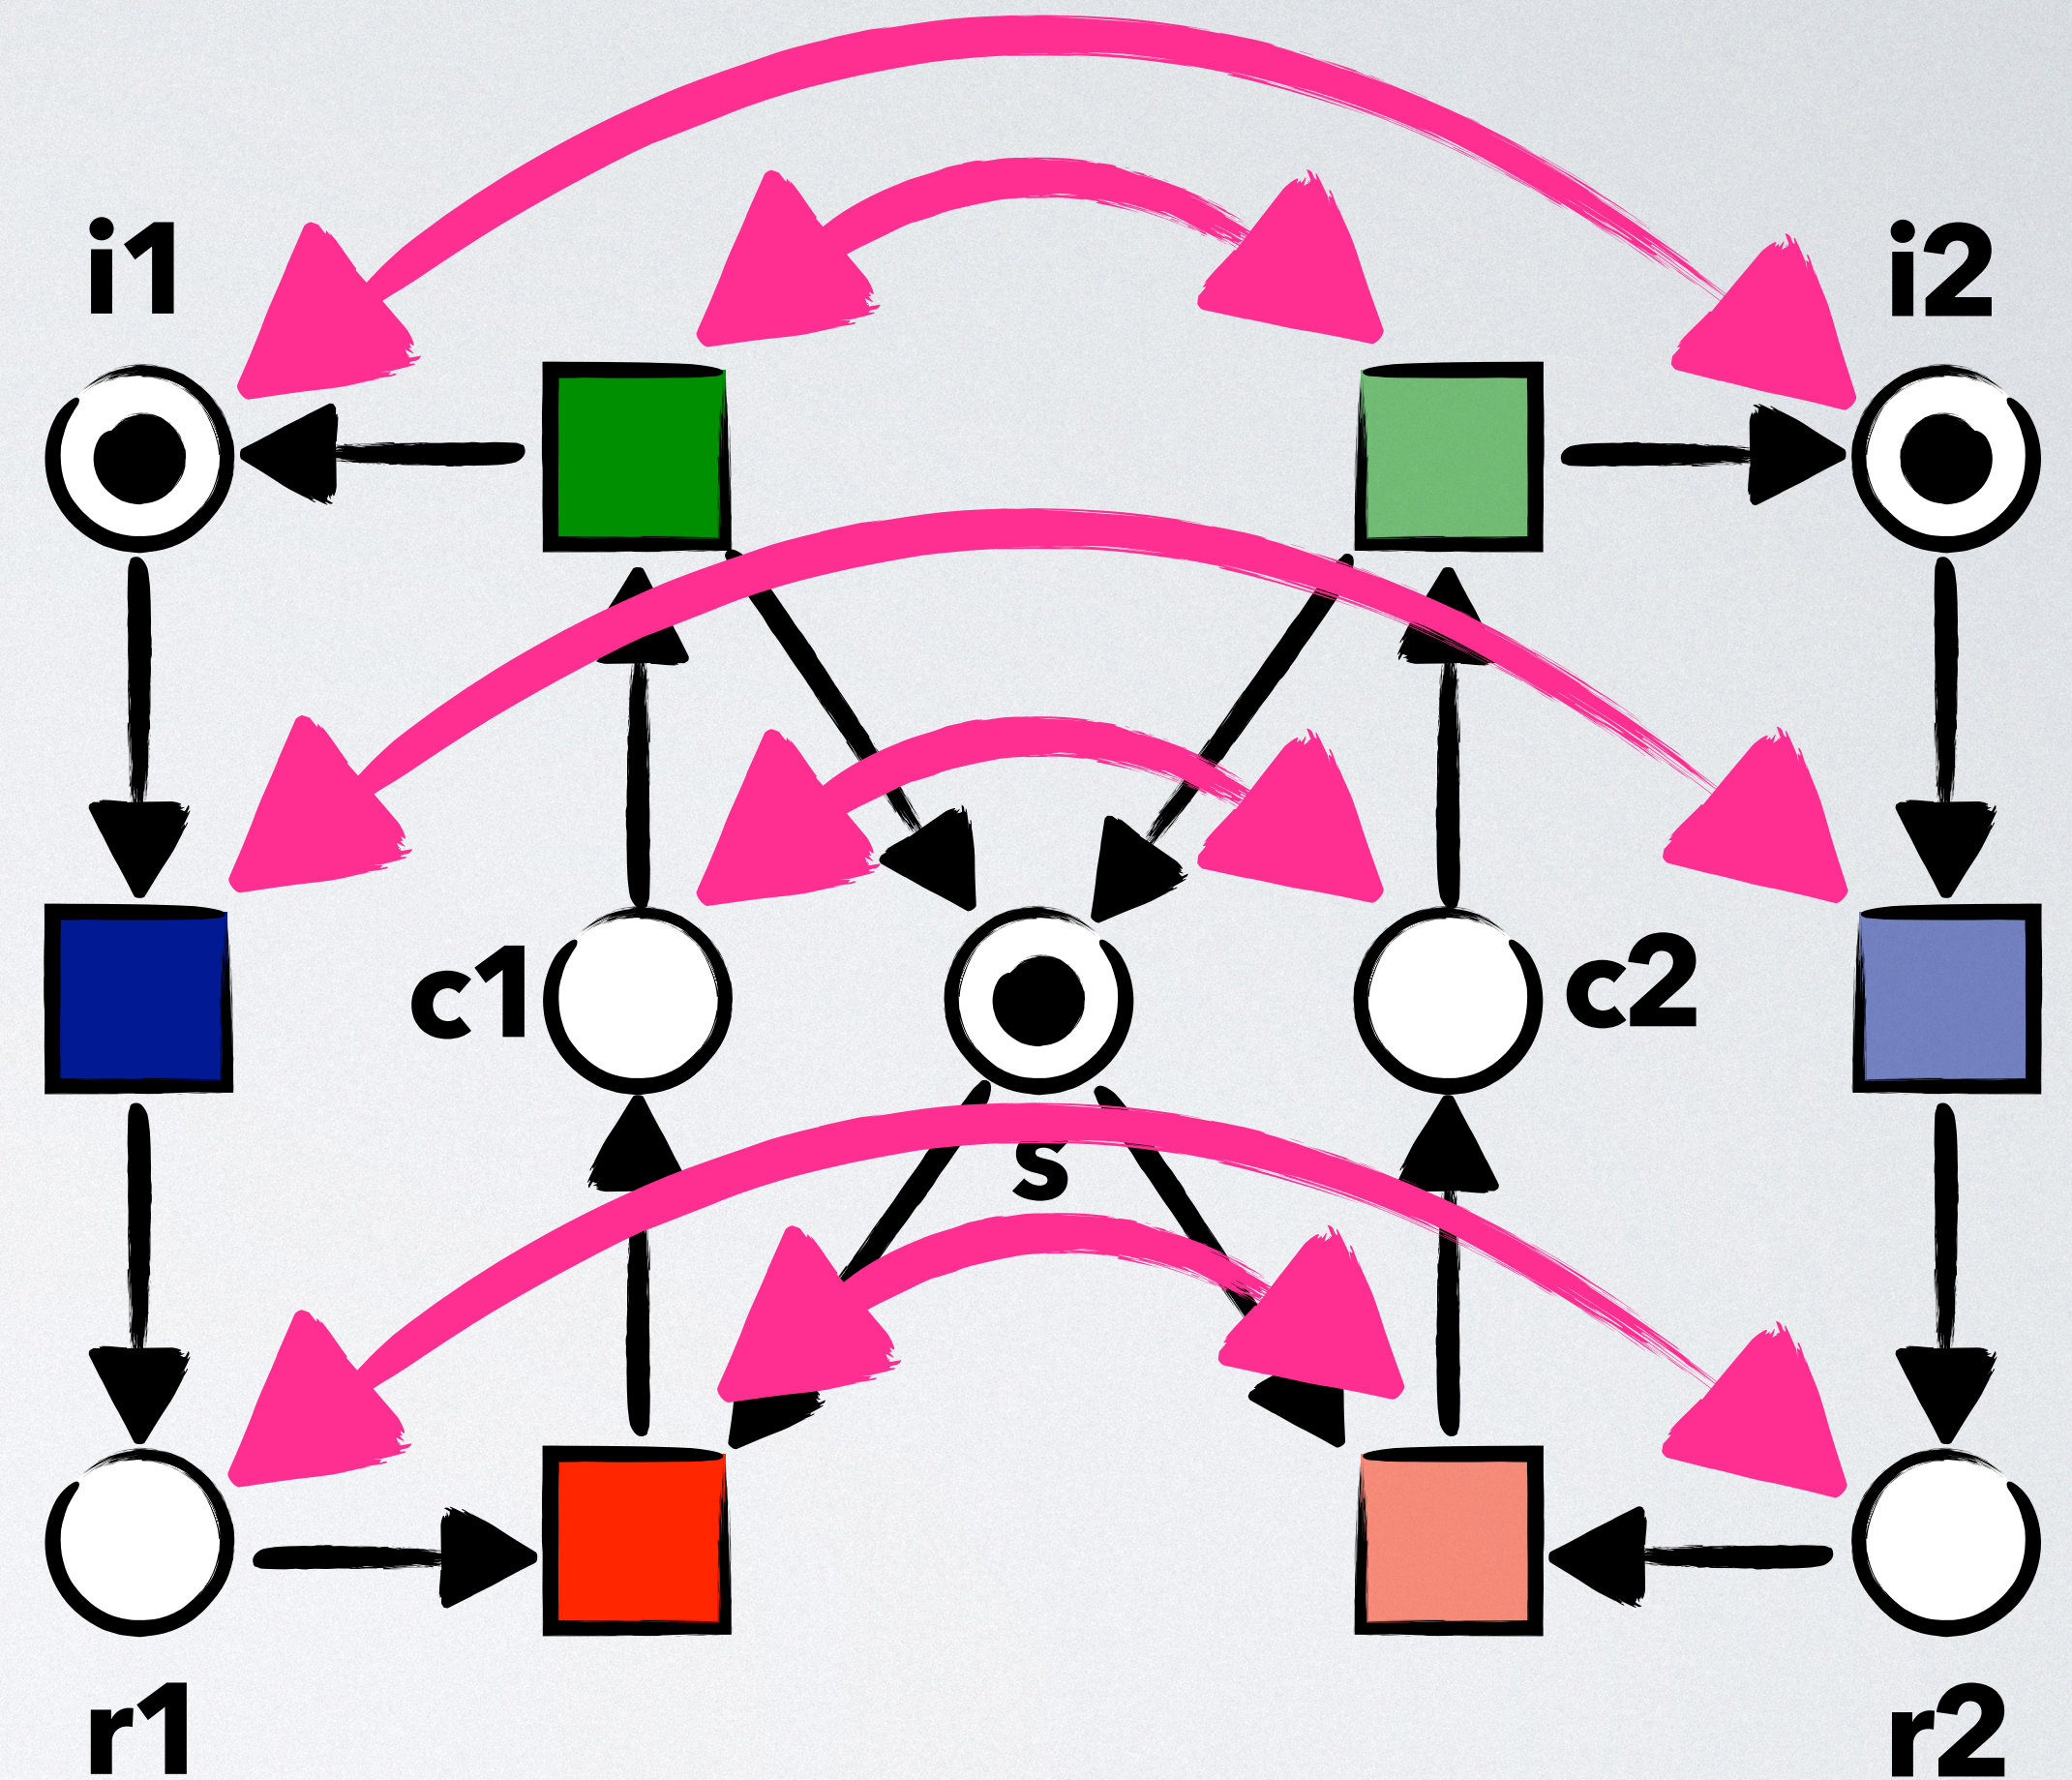
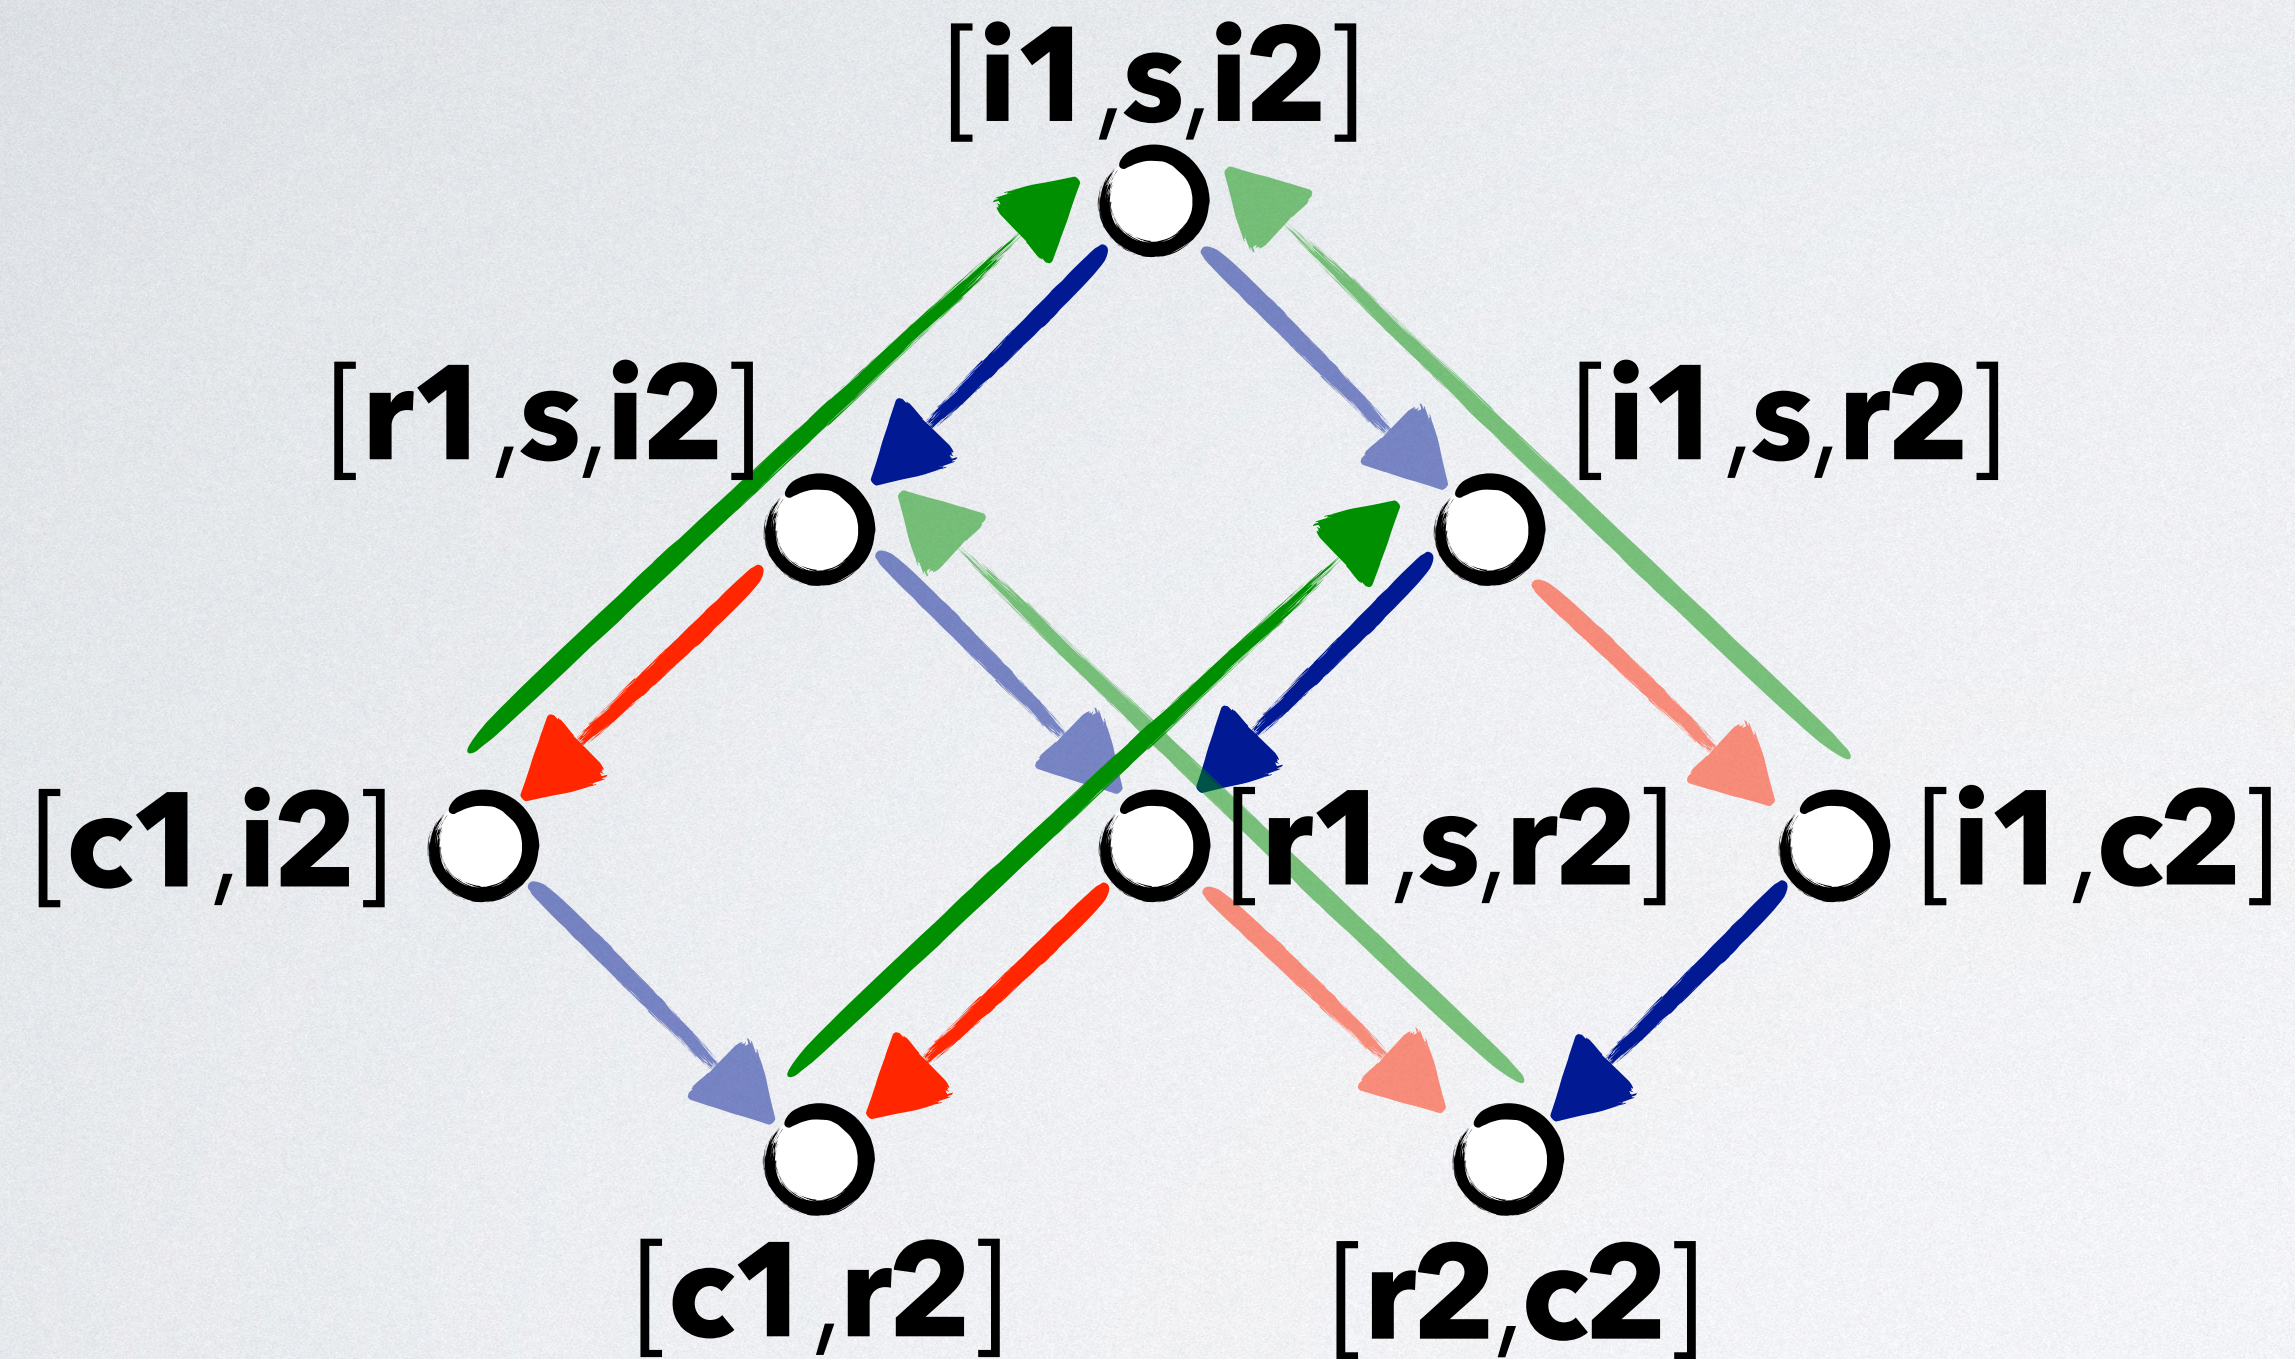
graph
automorphisms

2. Symmetry reduction



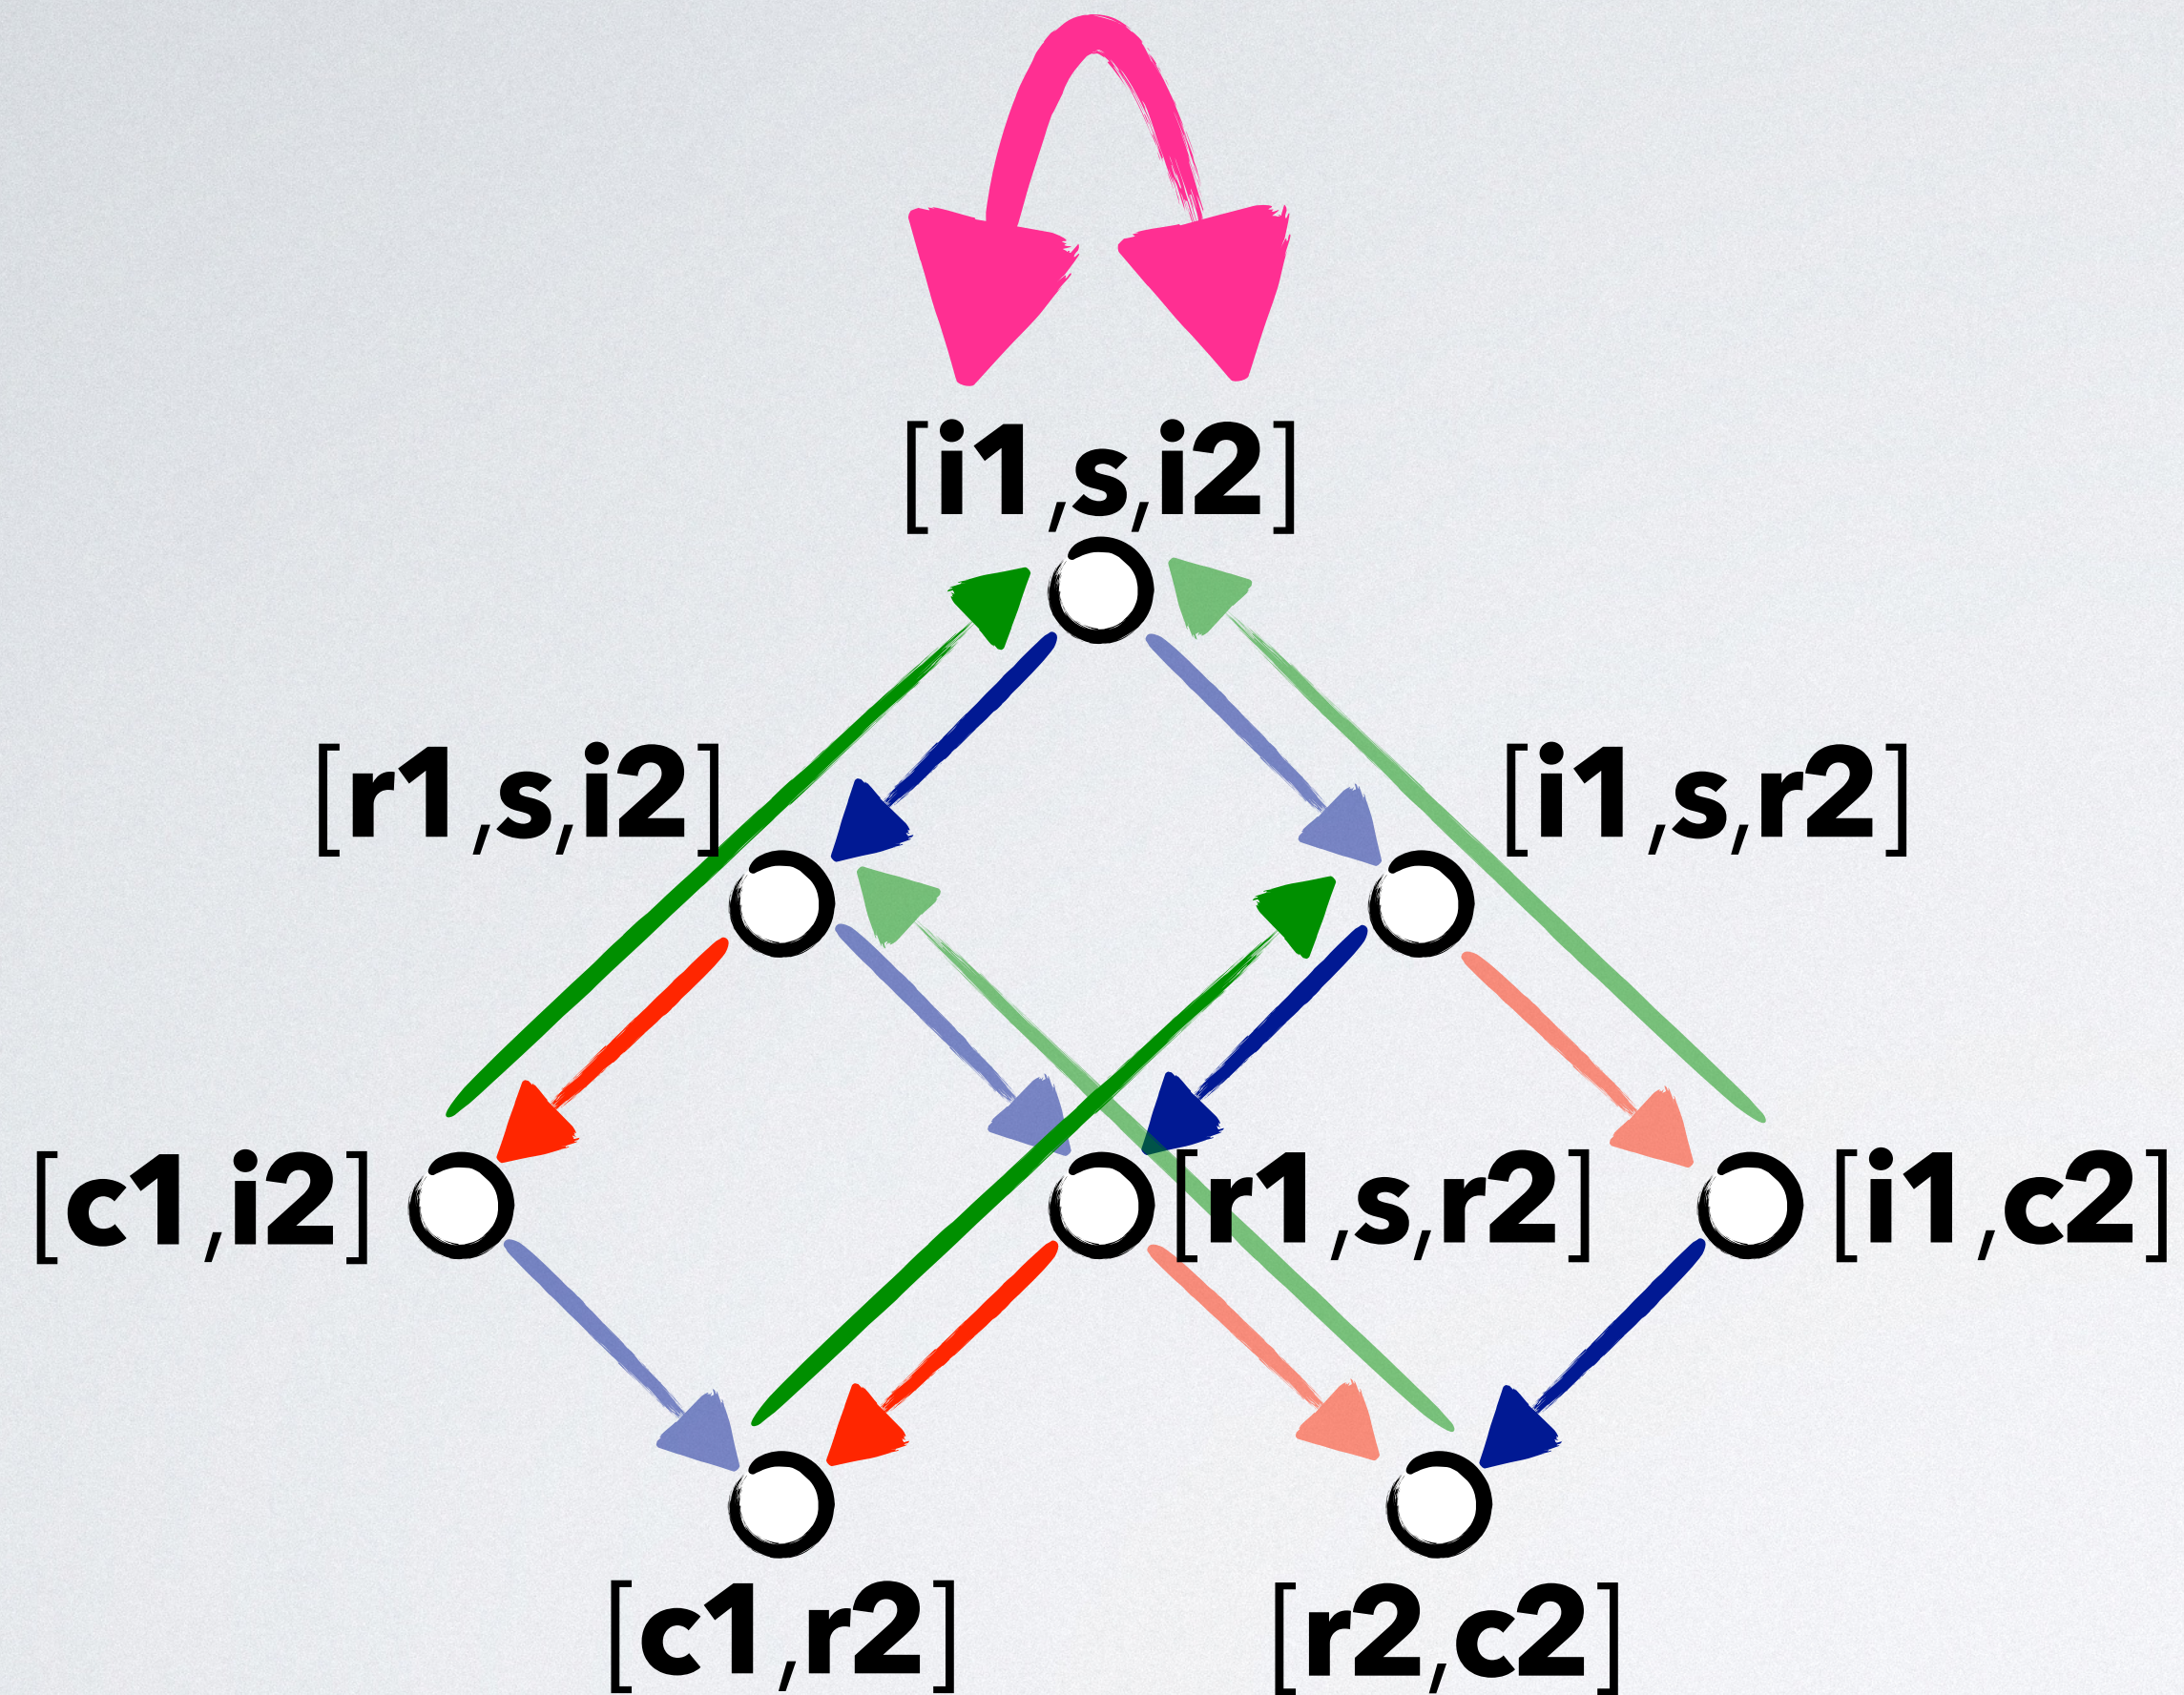
graph
automorphisms

2. Symmetry reduction

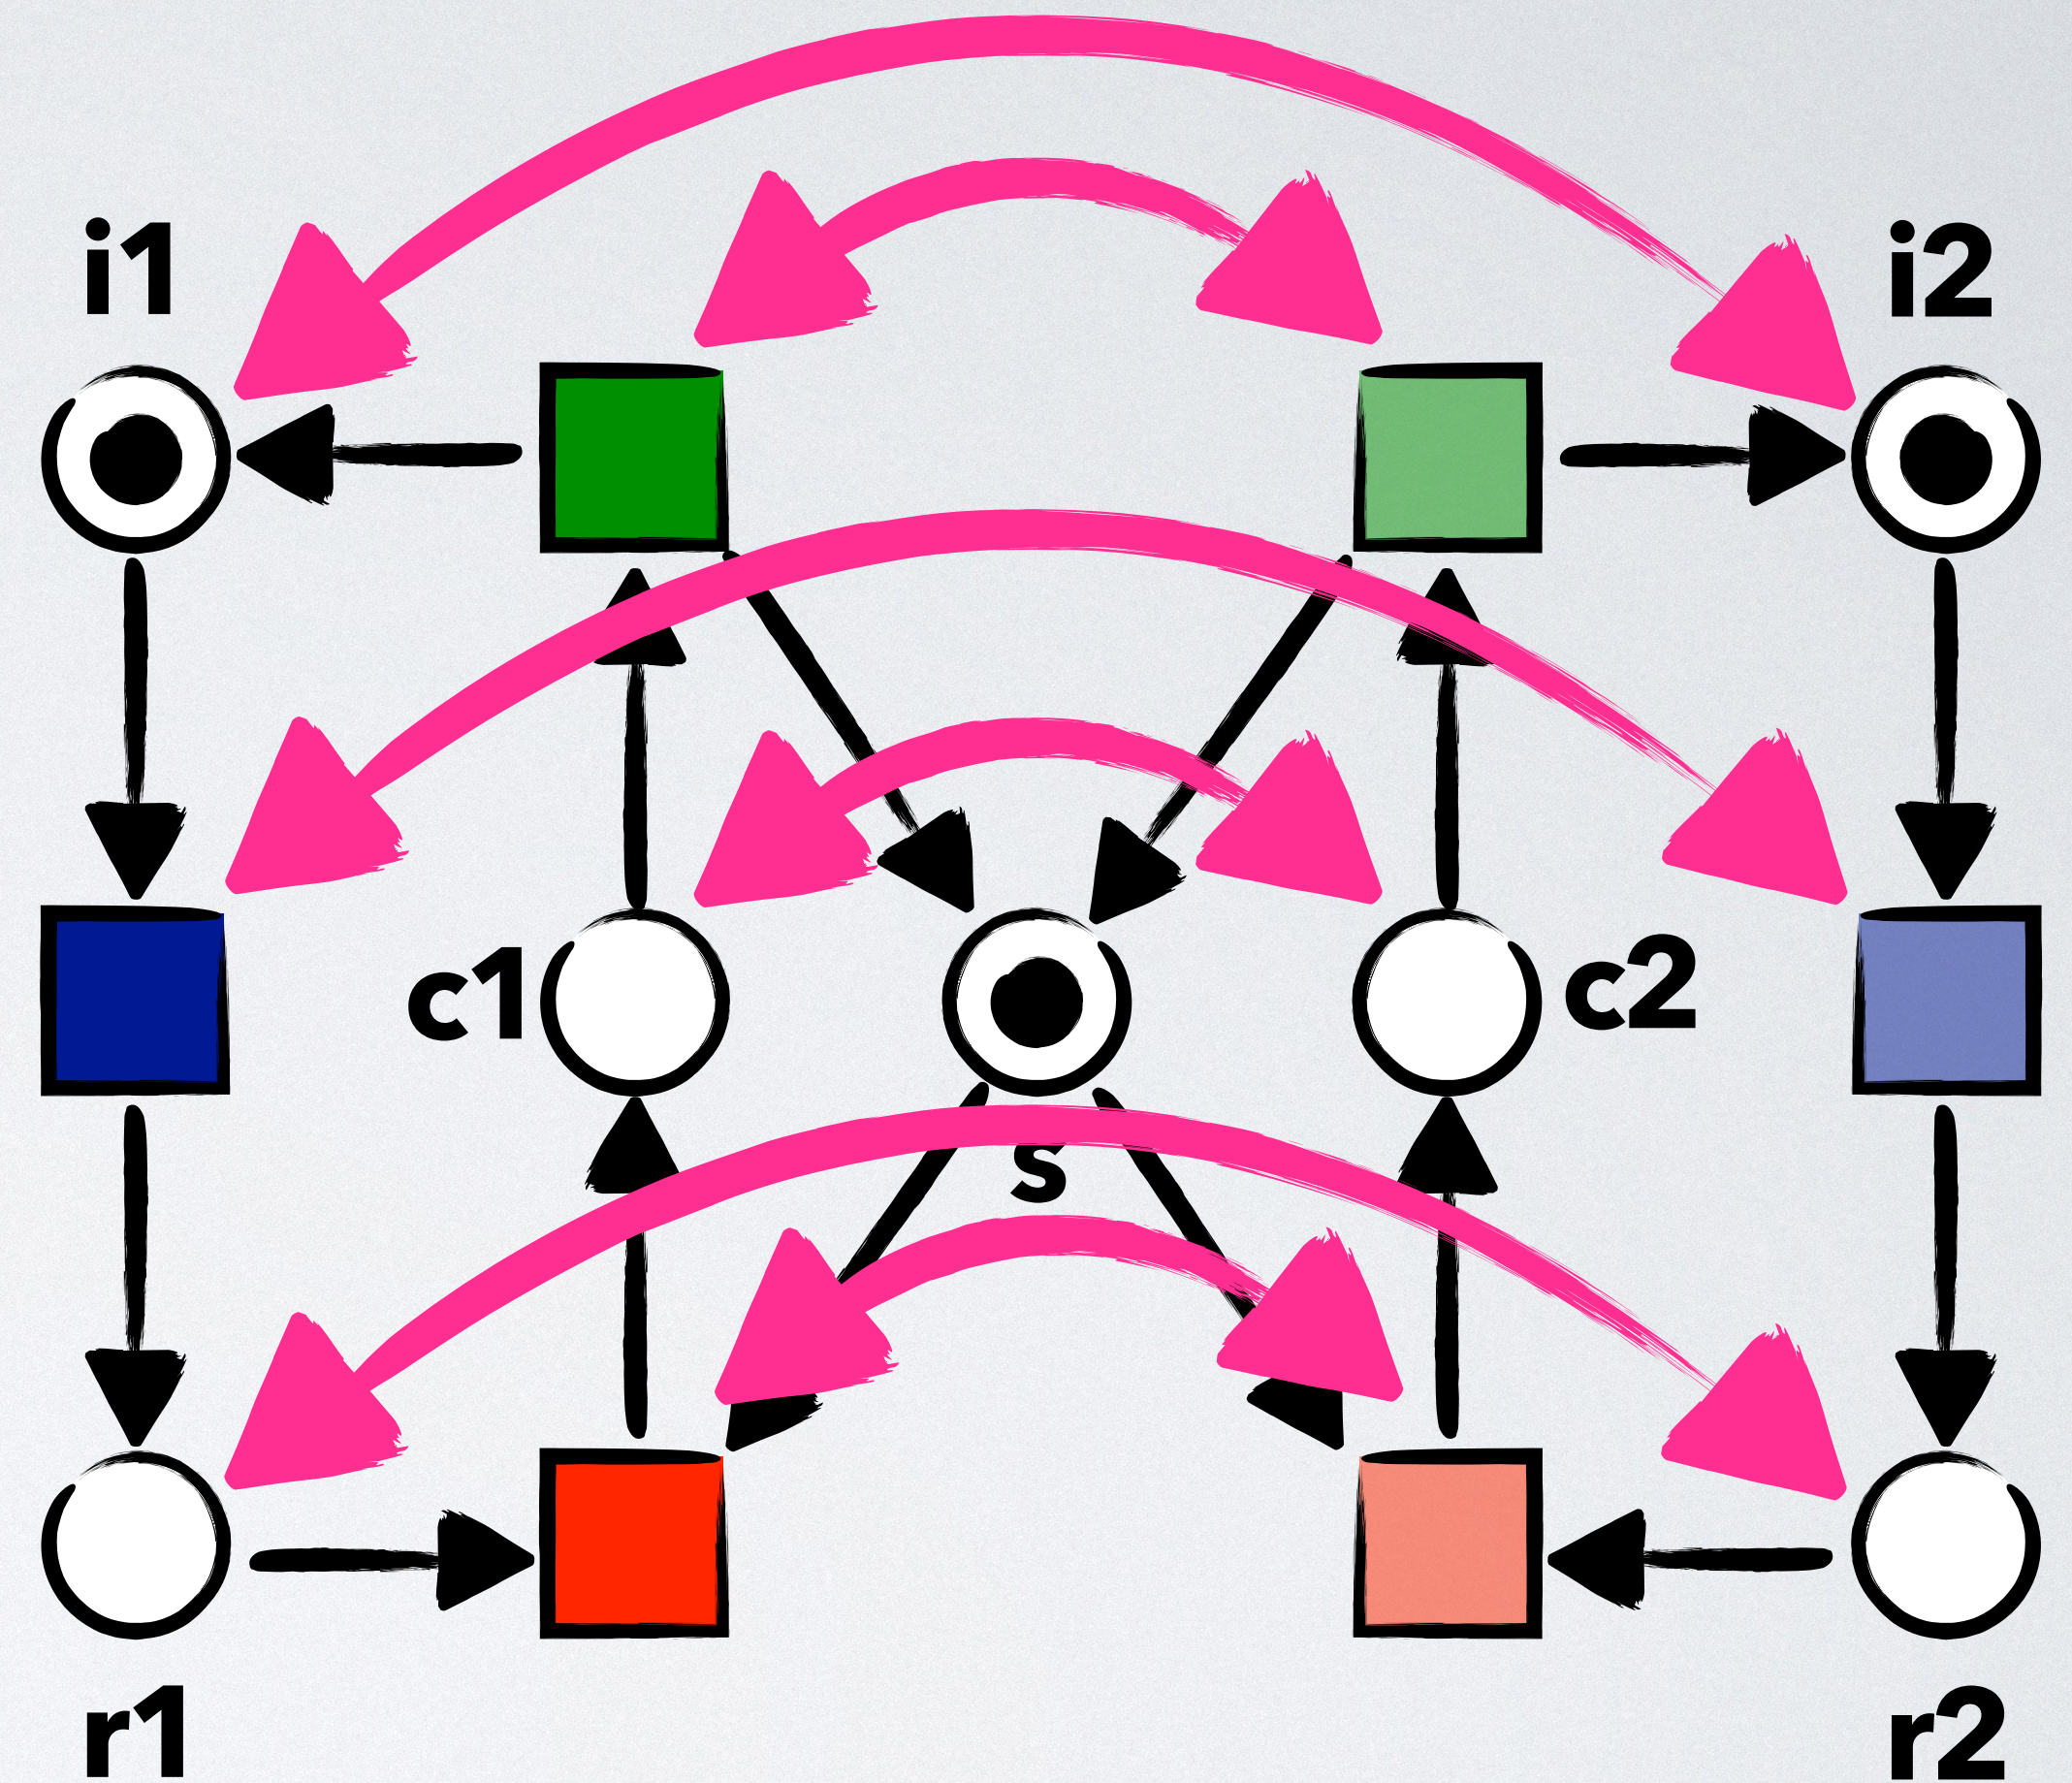


graph
automorphisms

2. Symmetry reduction

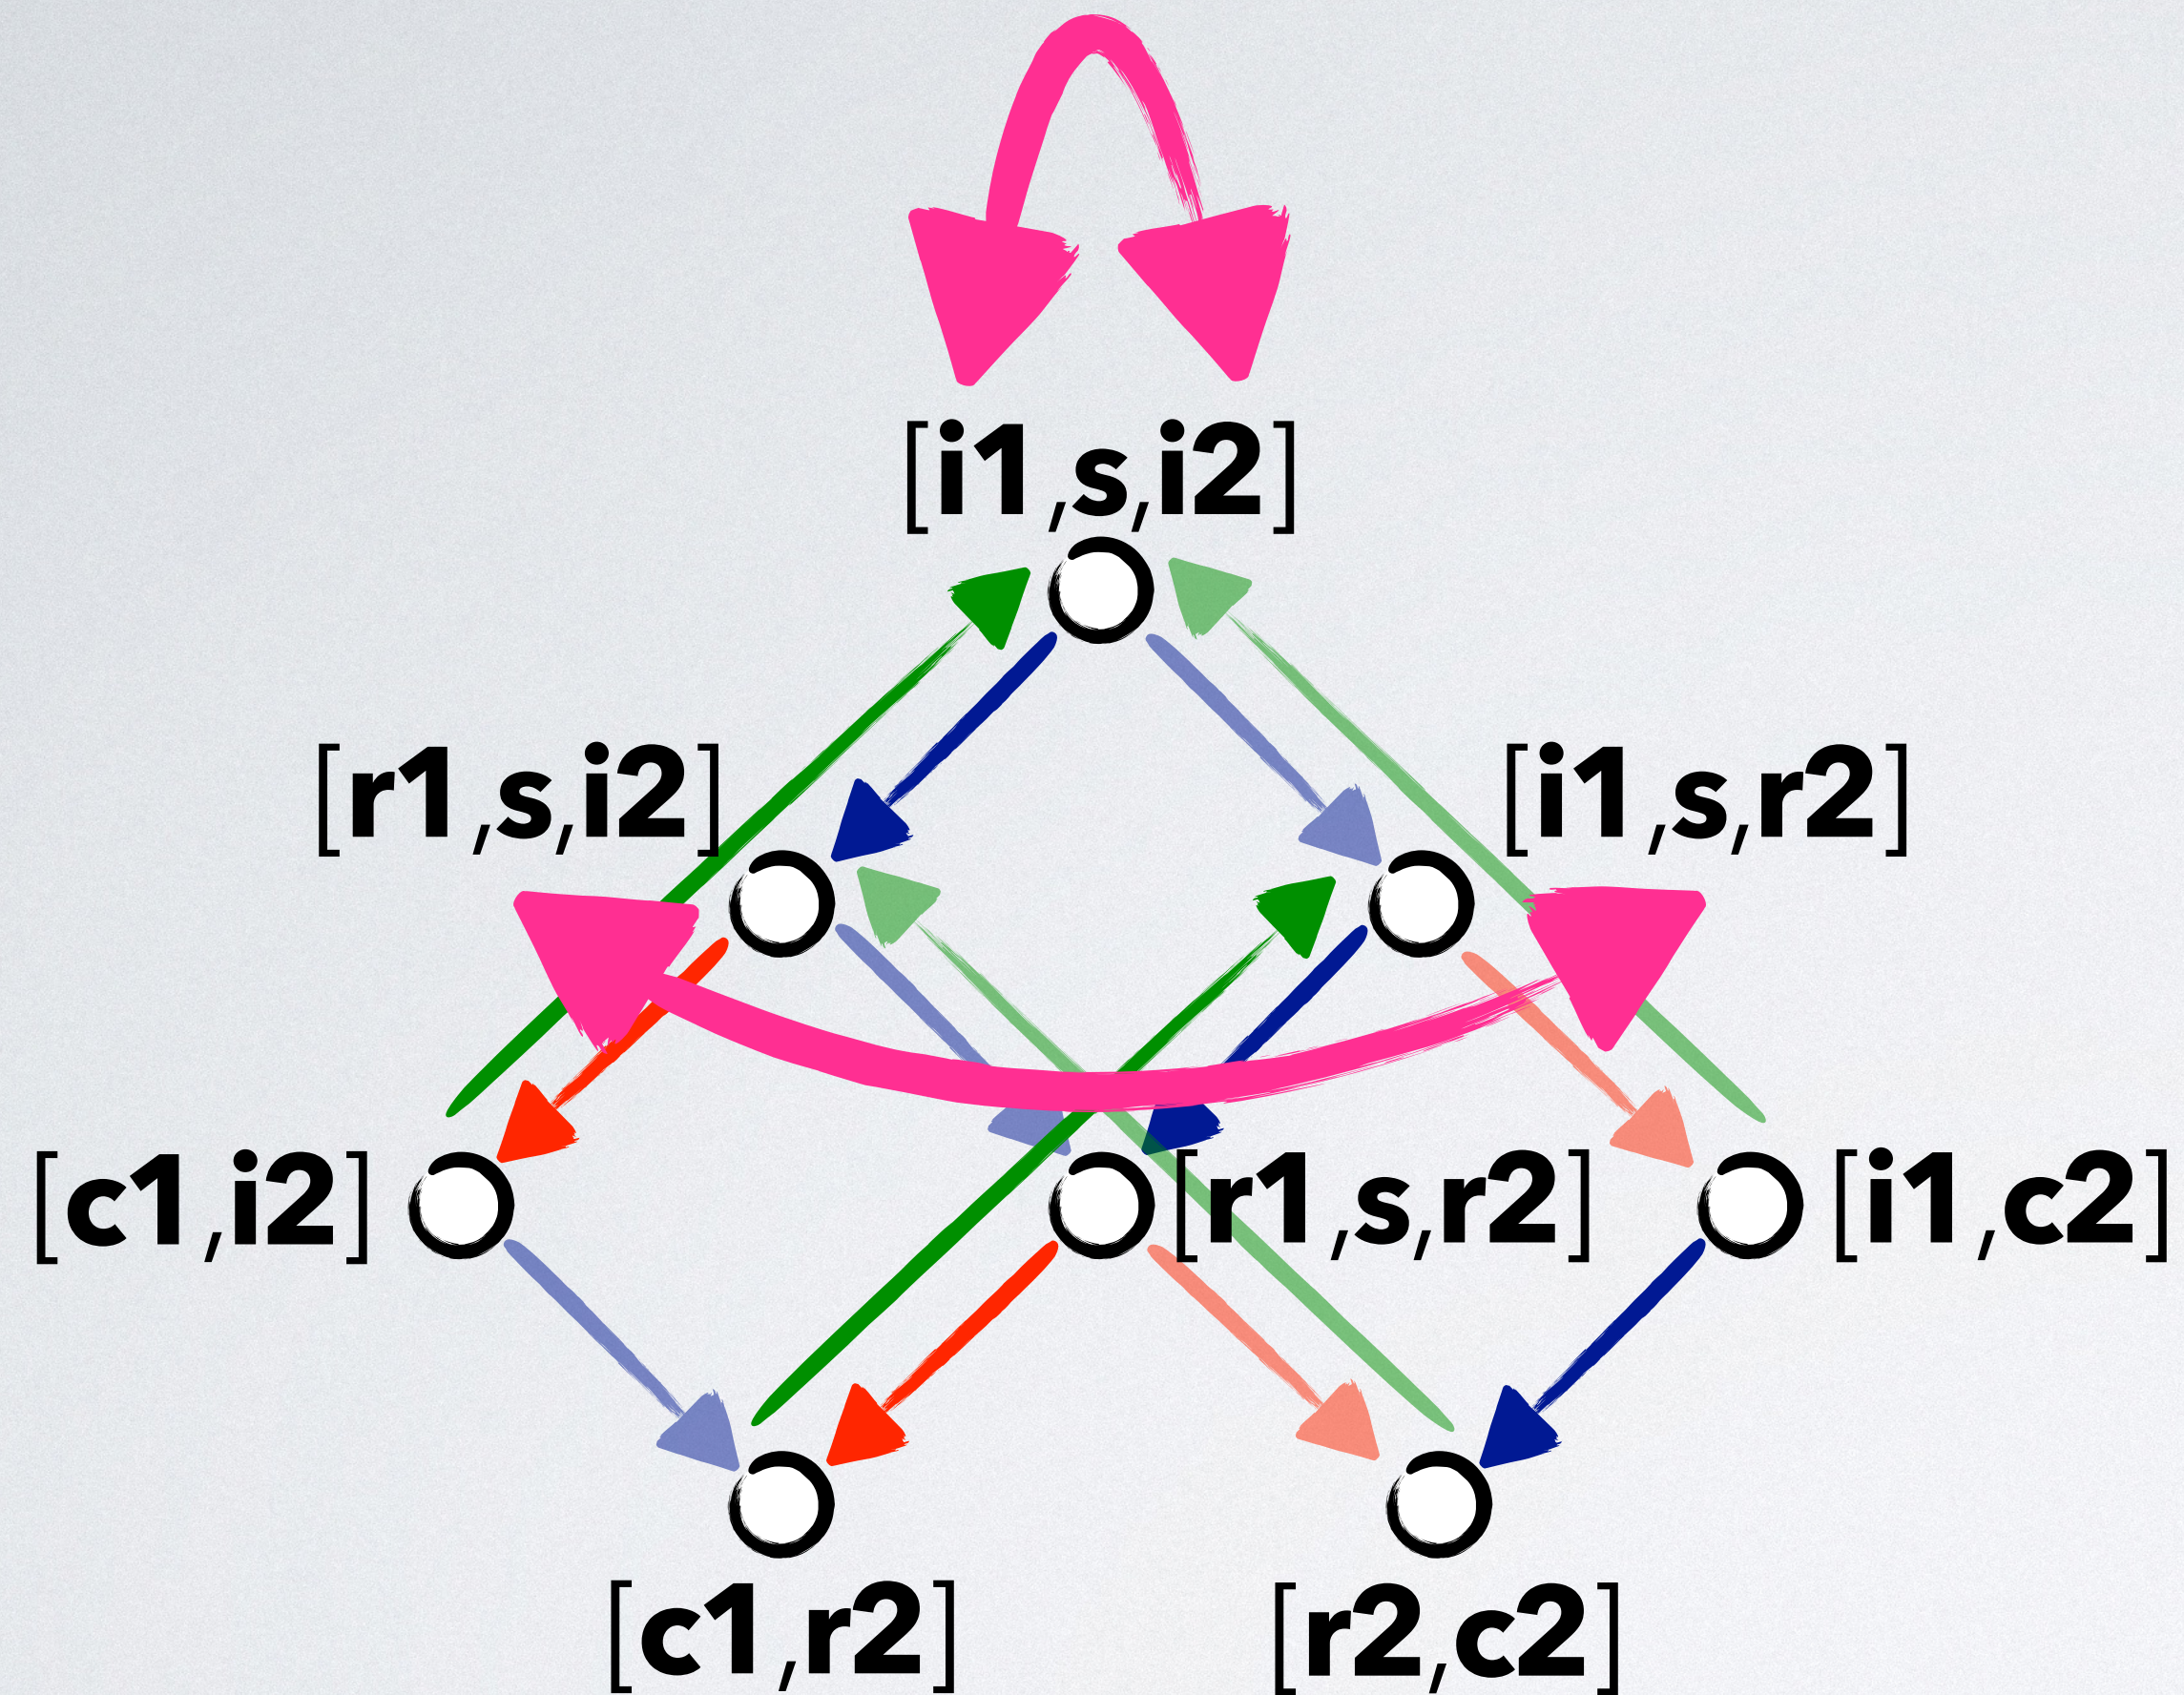


symmetric markings

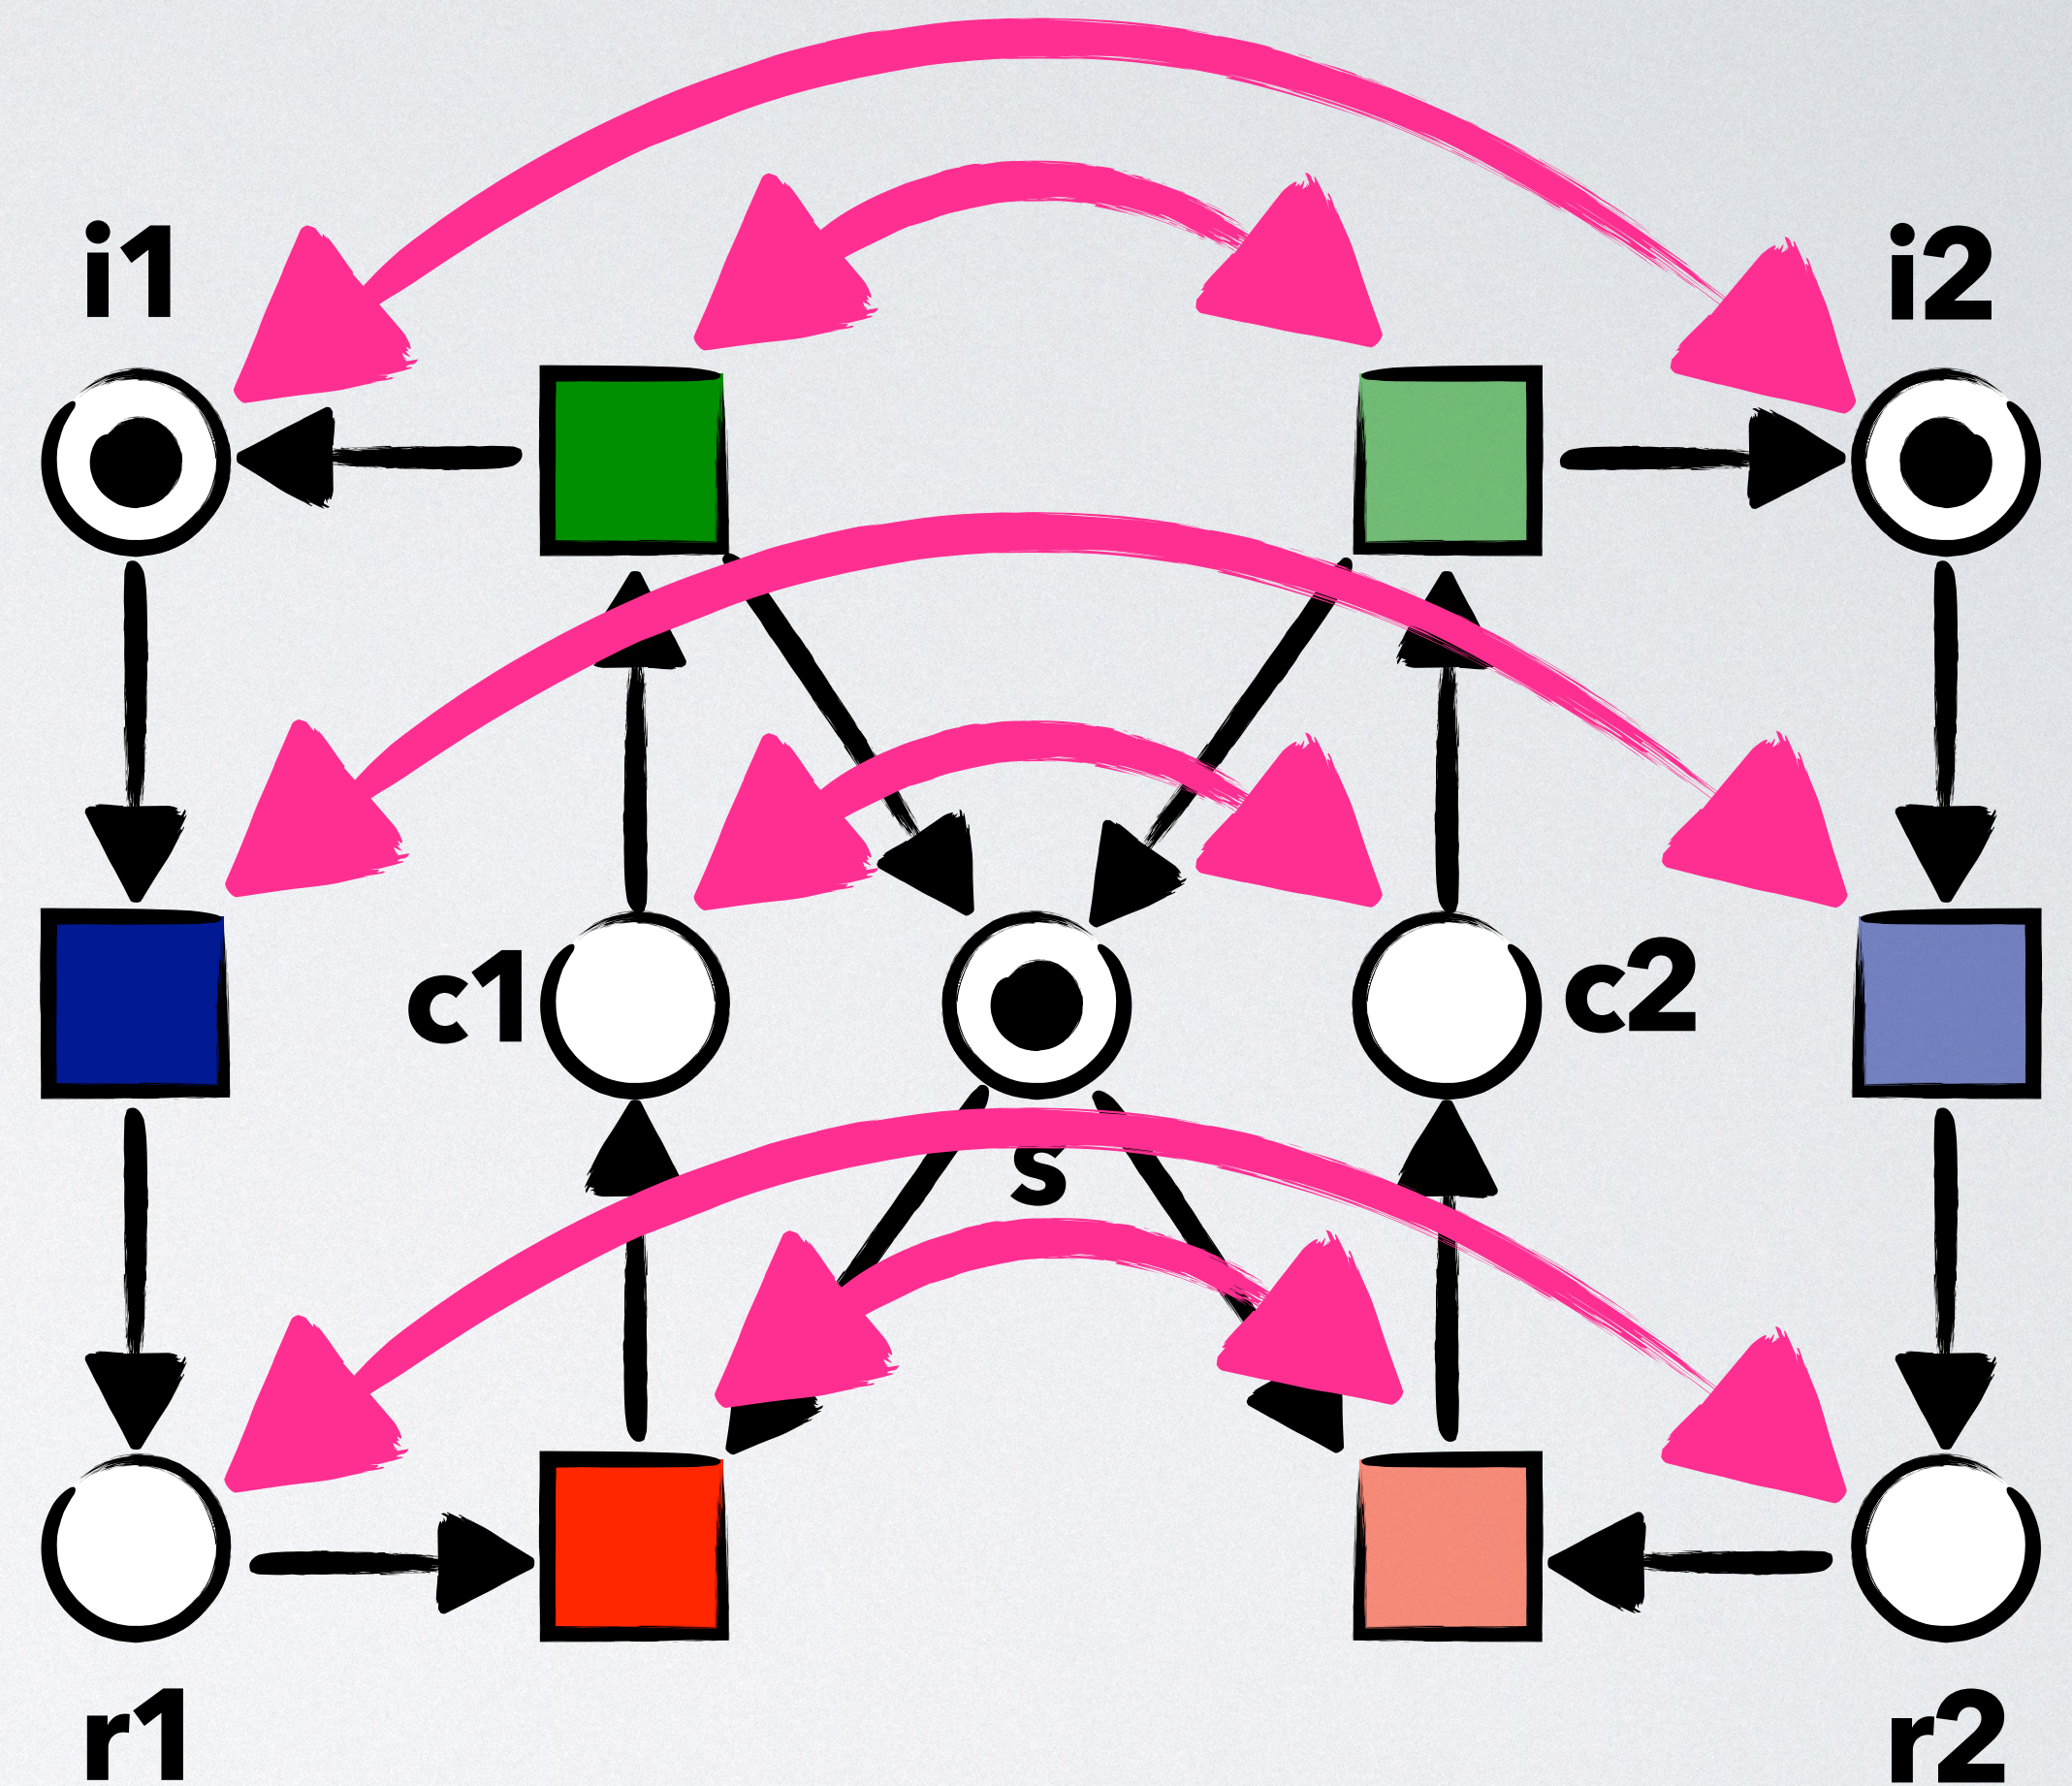


graph
automorphisms

2. Symmetry reduction

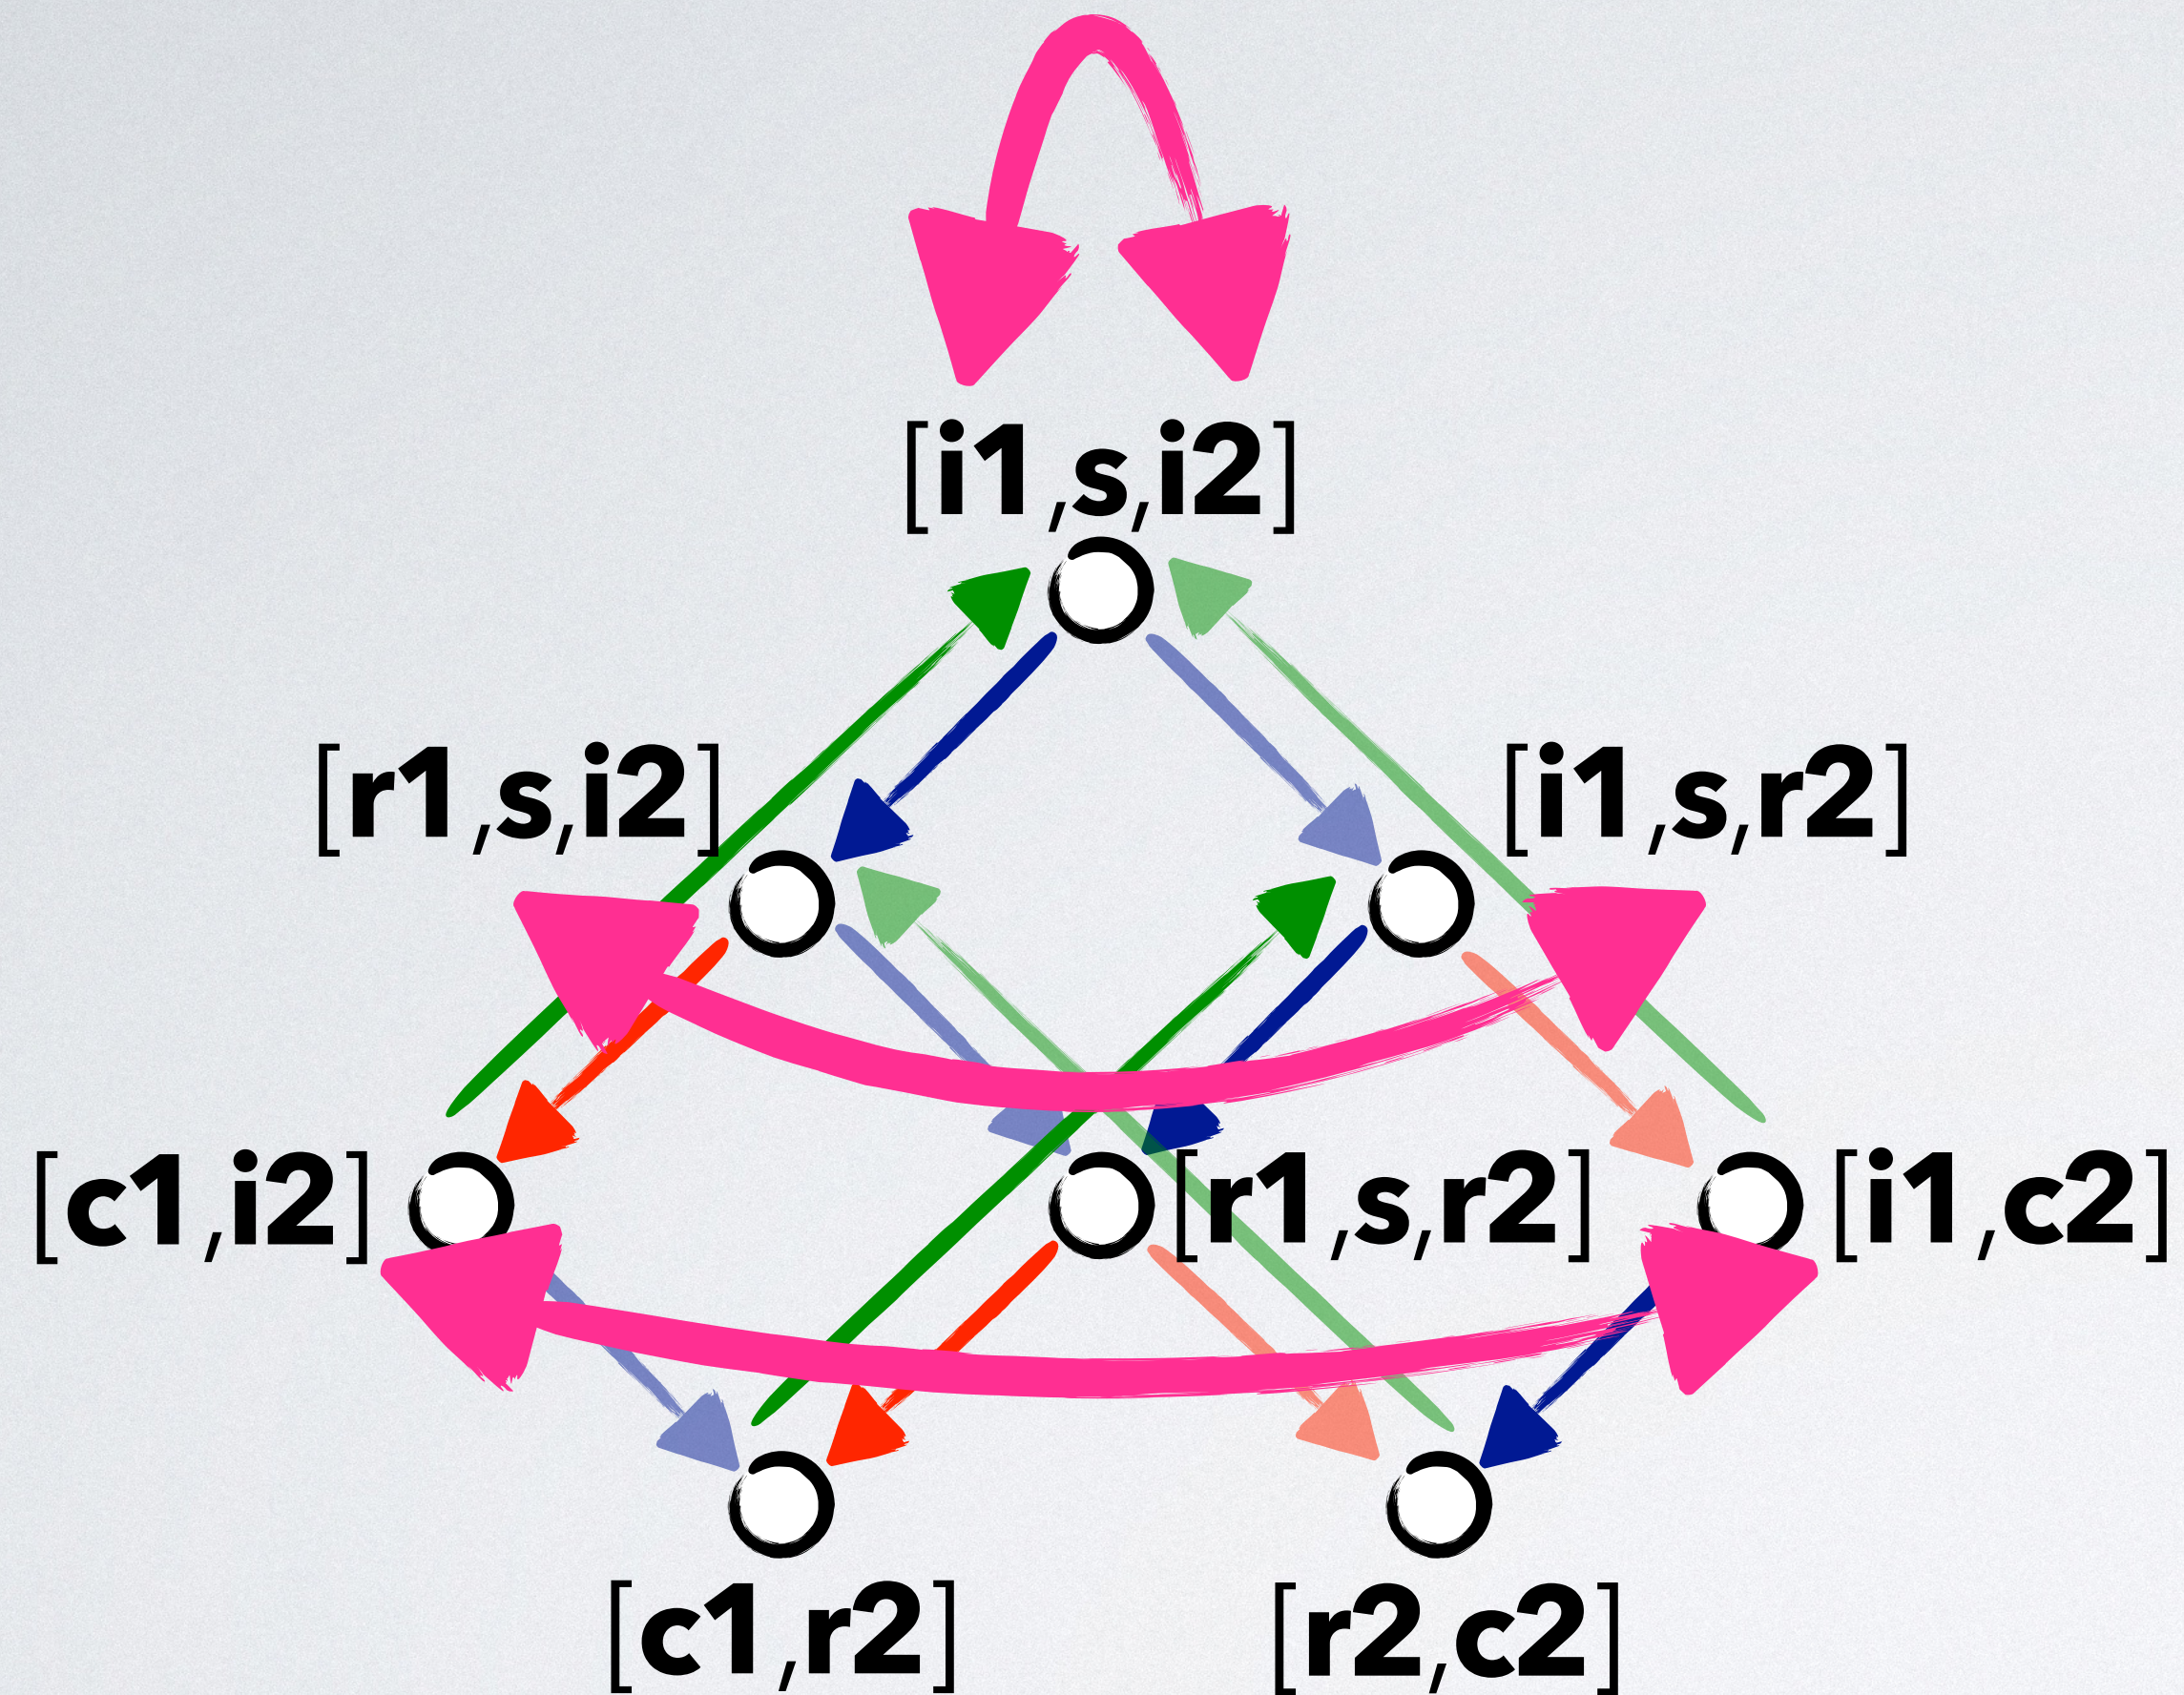


symmetric markings

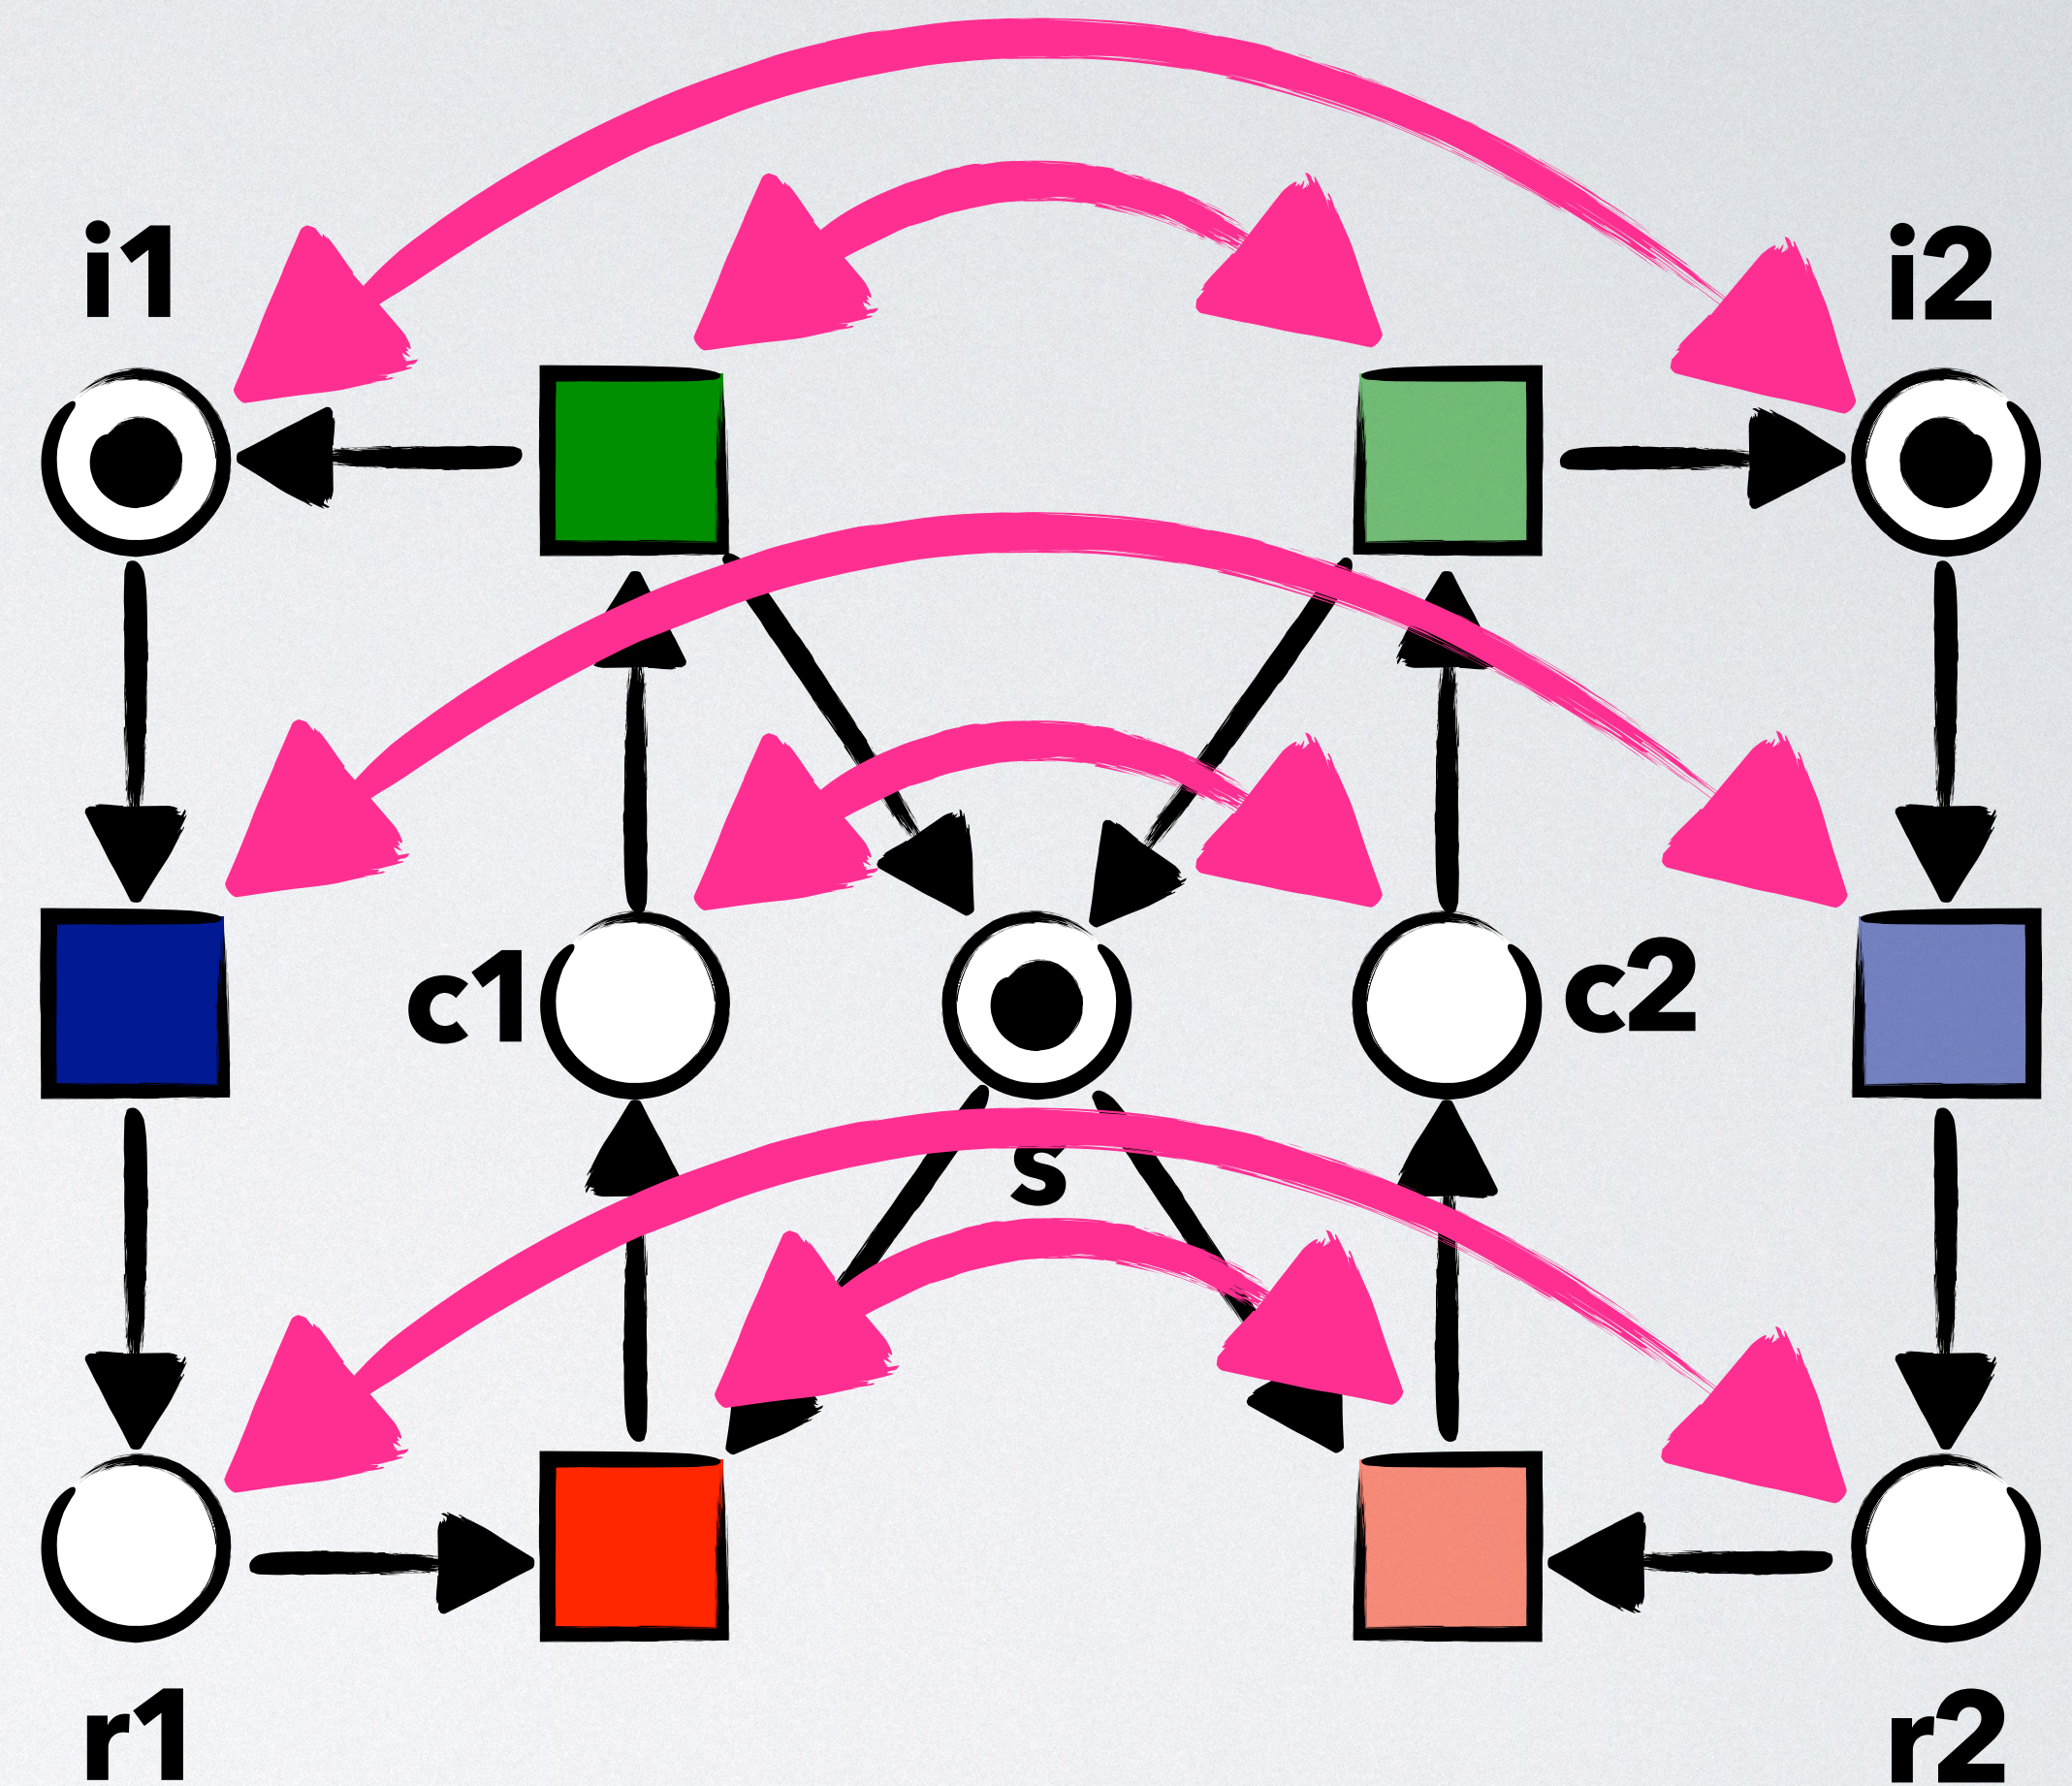


graph automorphisms

2. Symmetry reduction

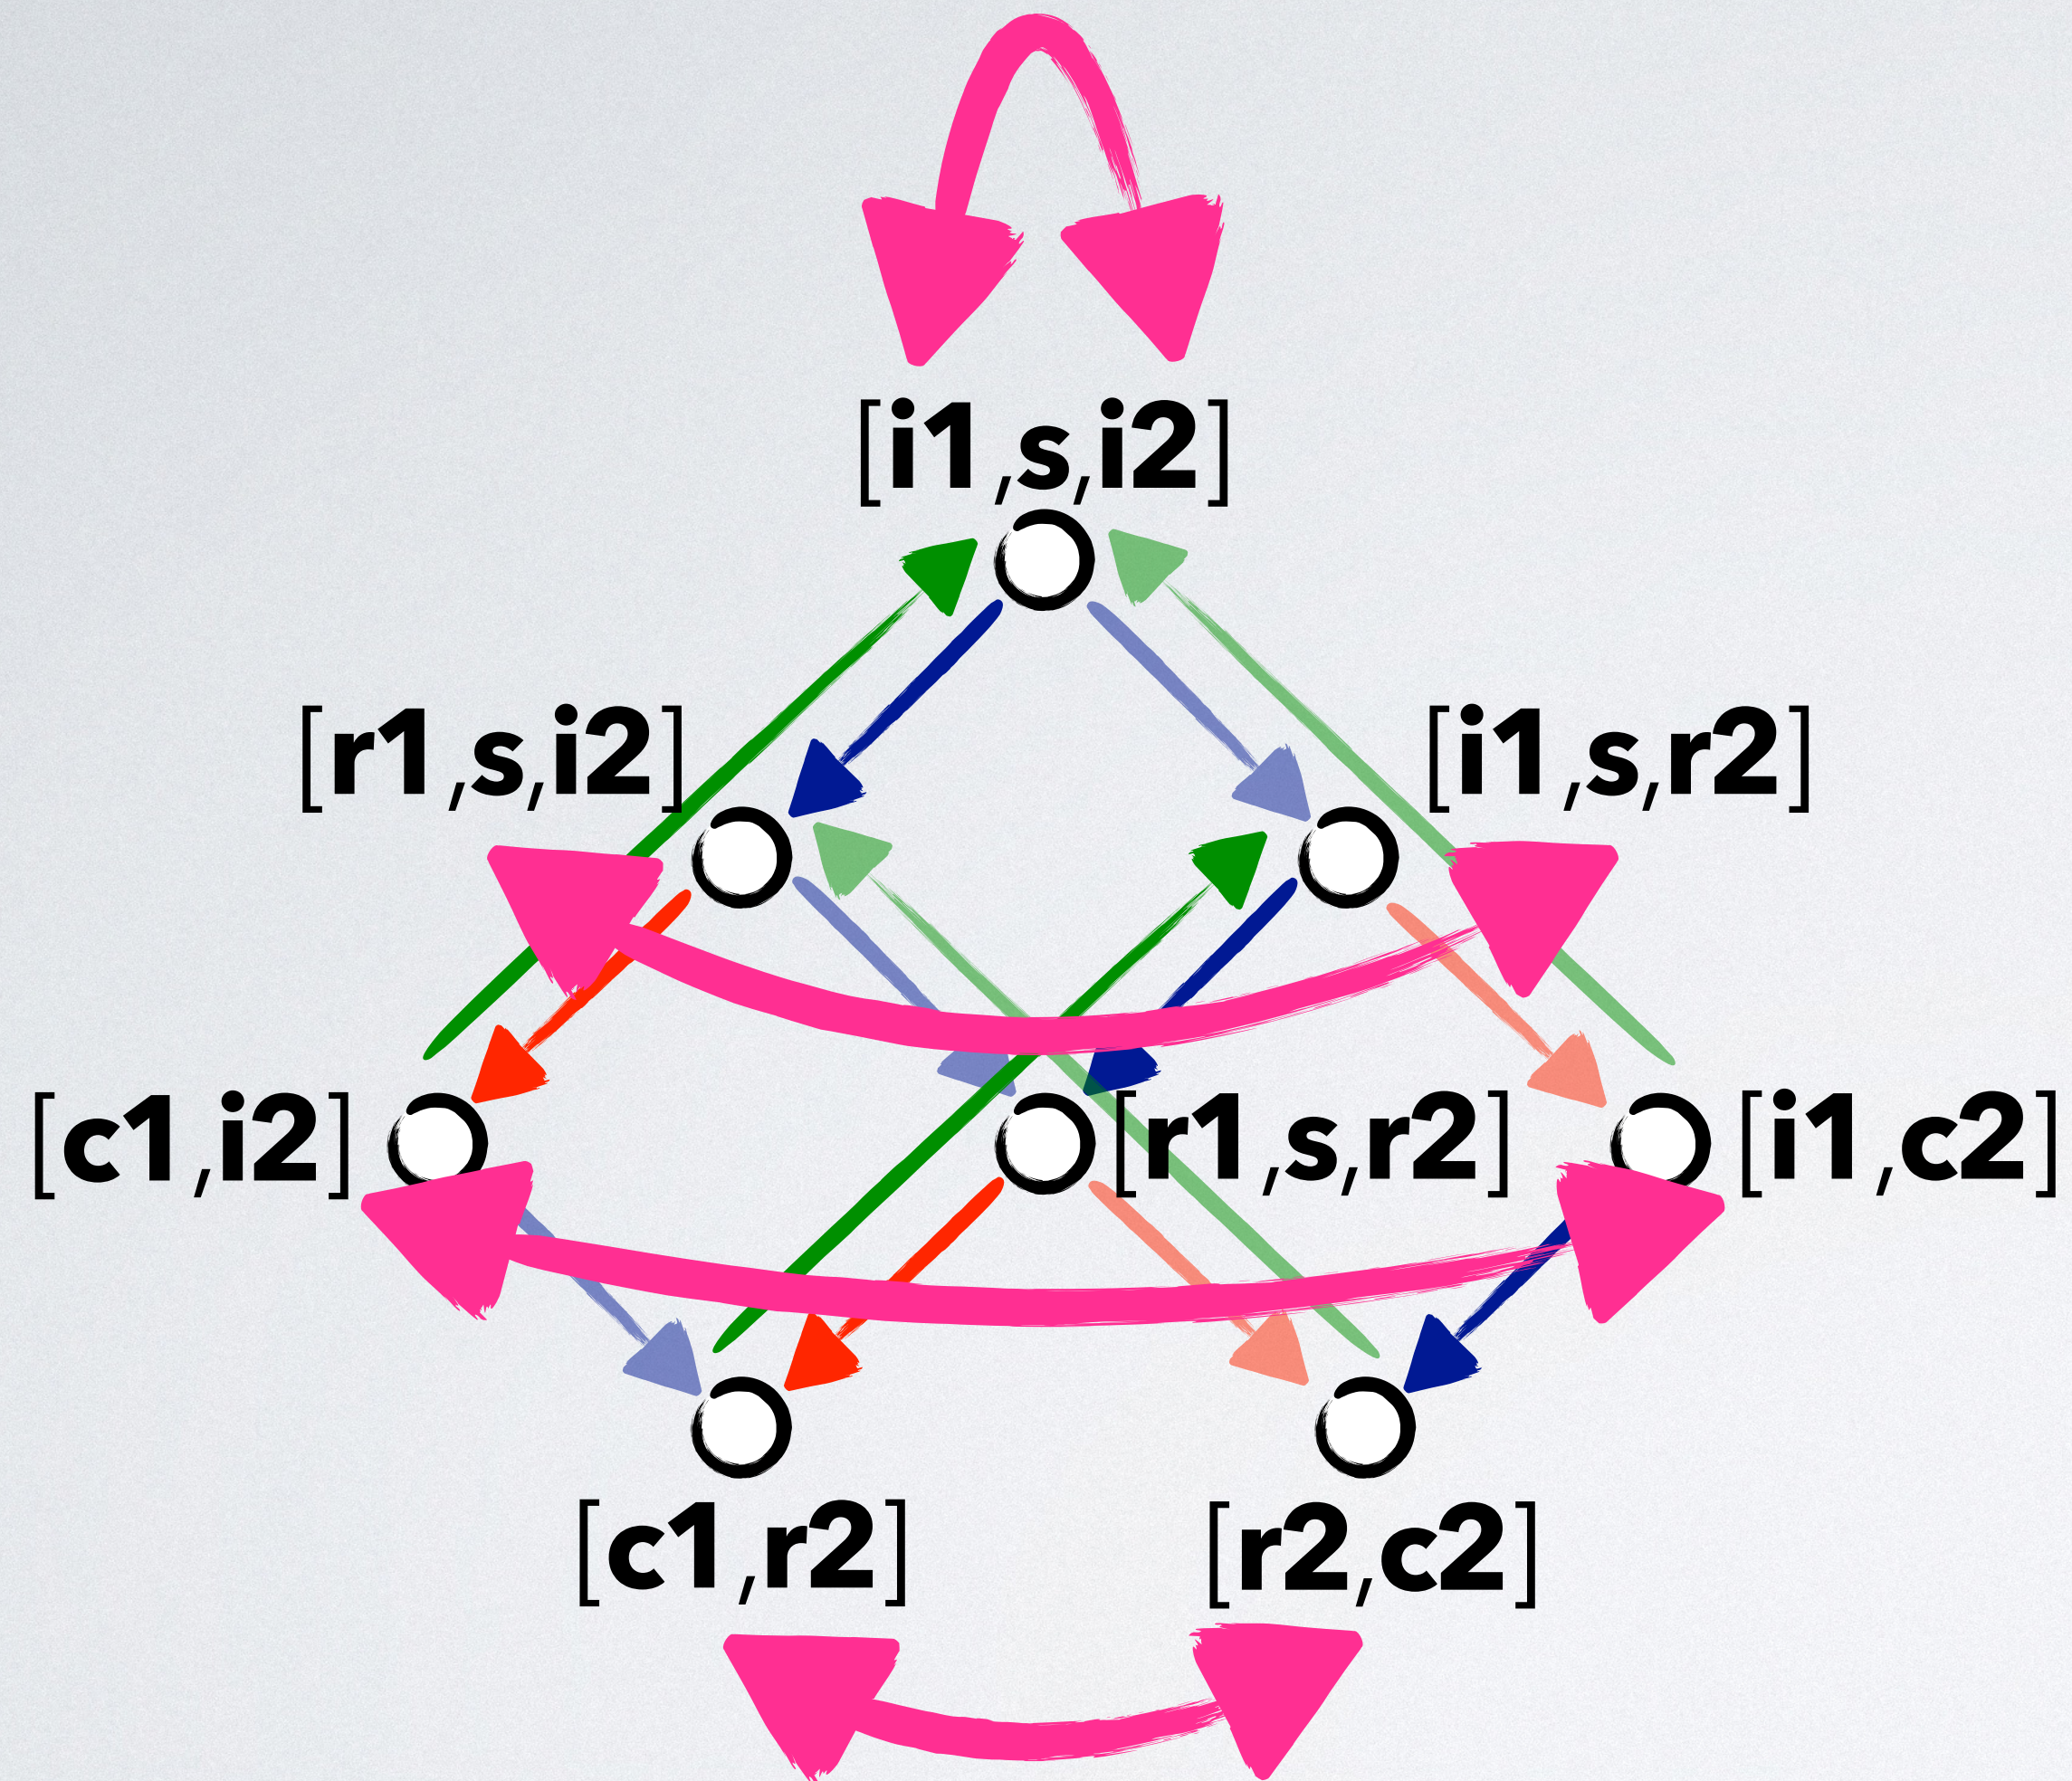


symmetric markings

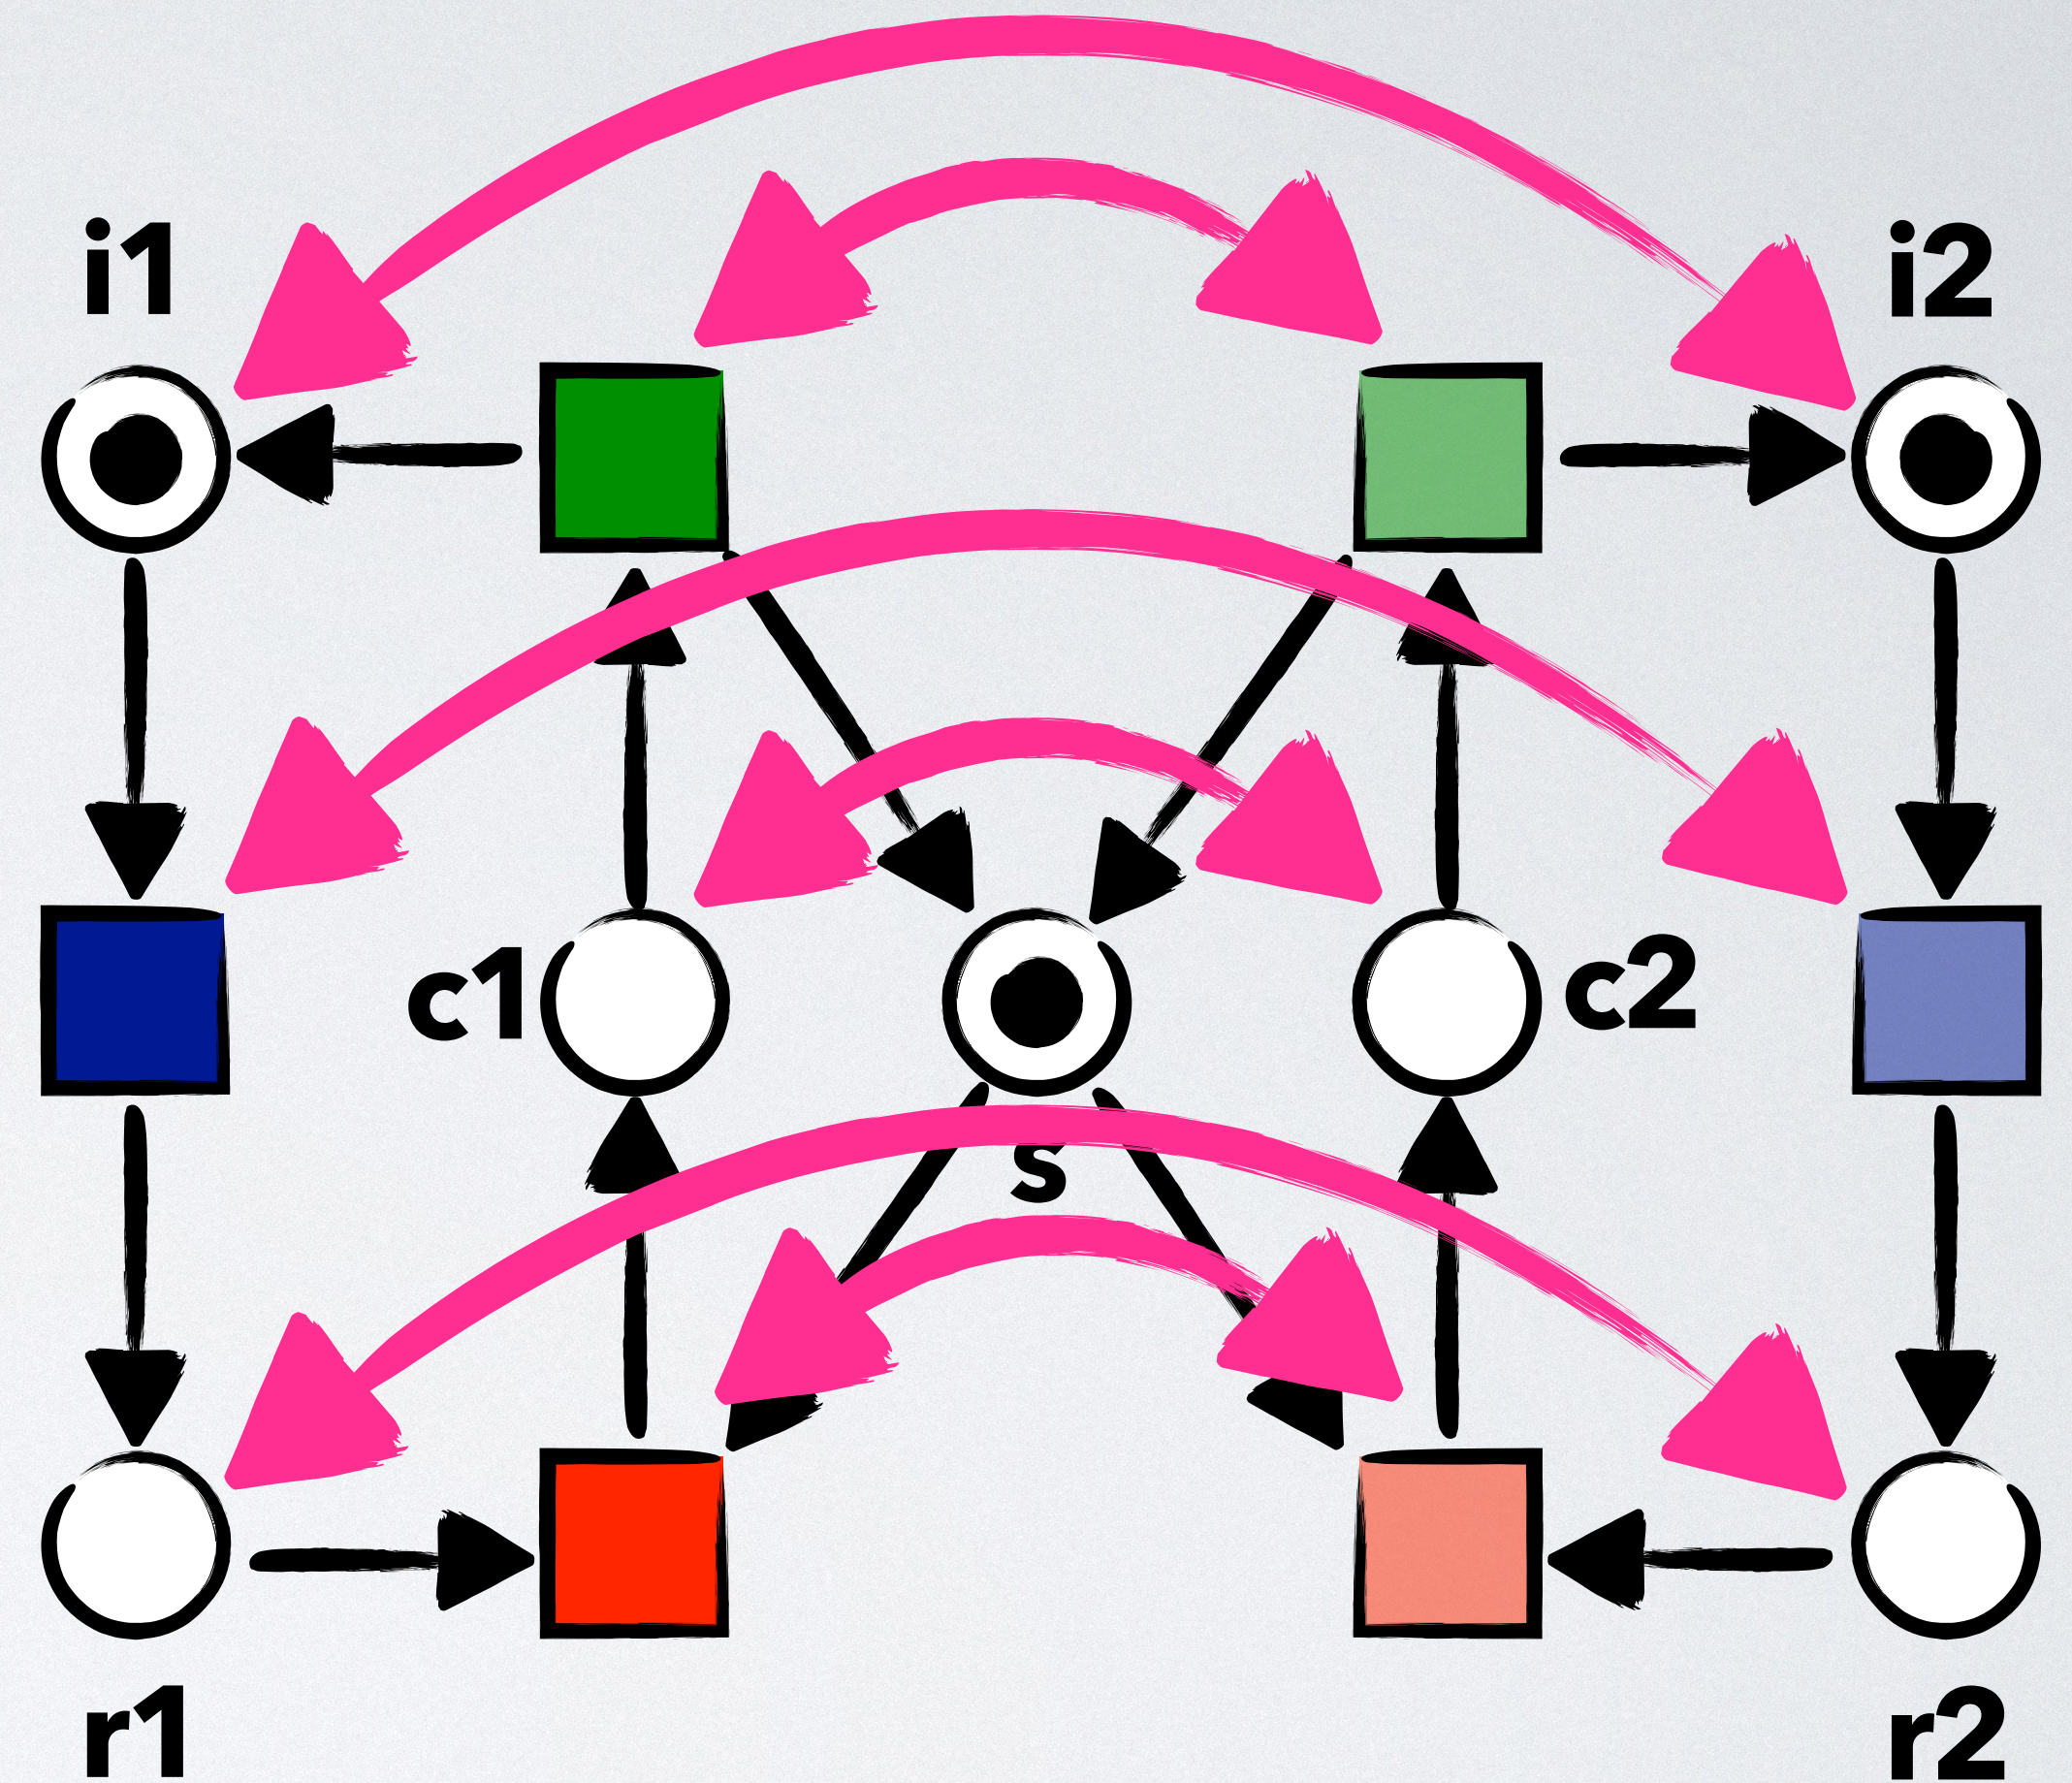


graph automorphisms

2. Symmetry reduction

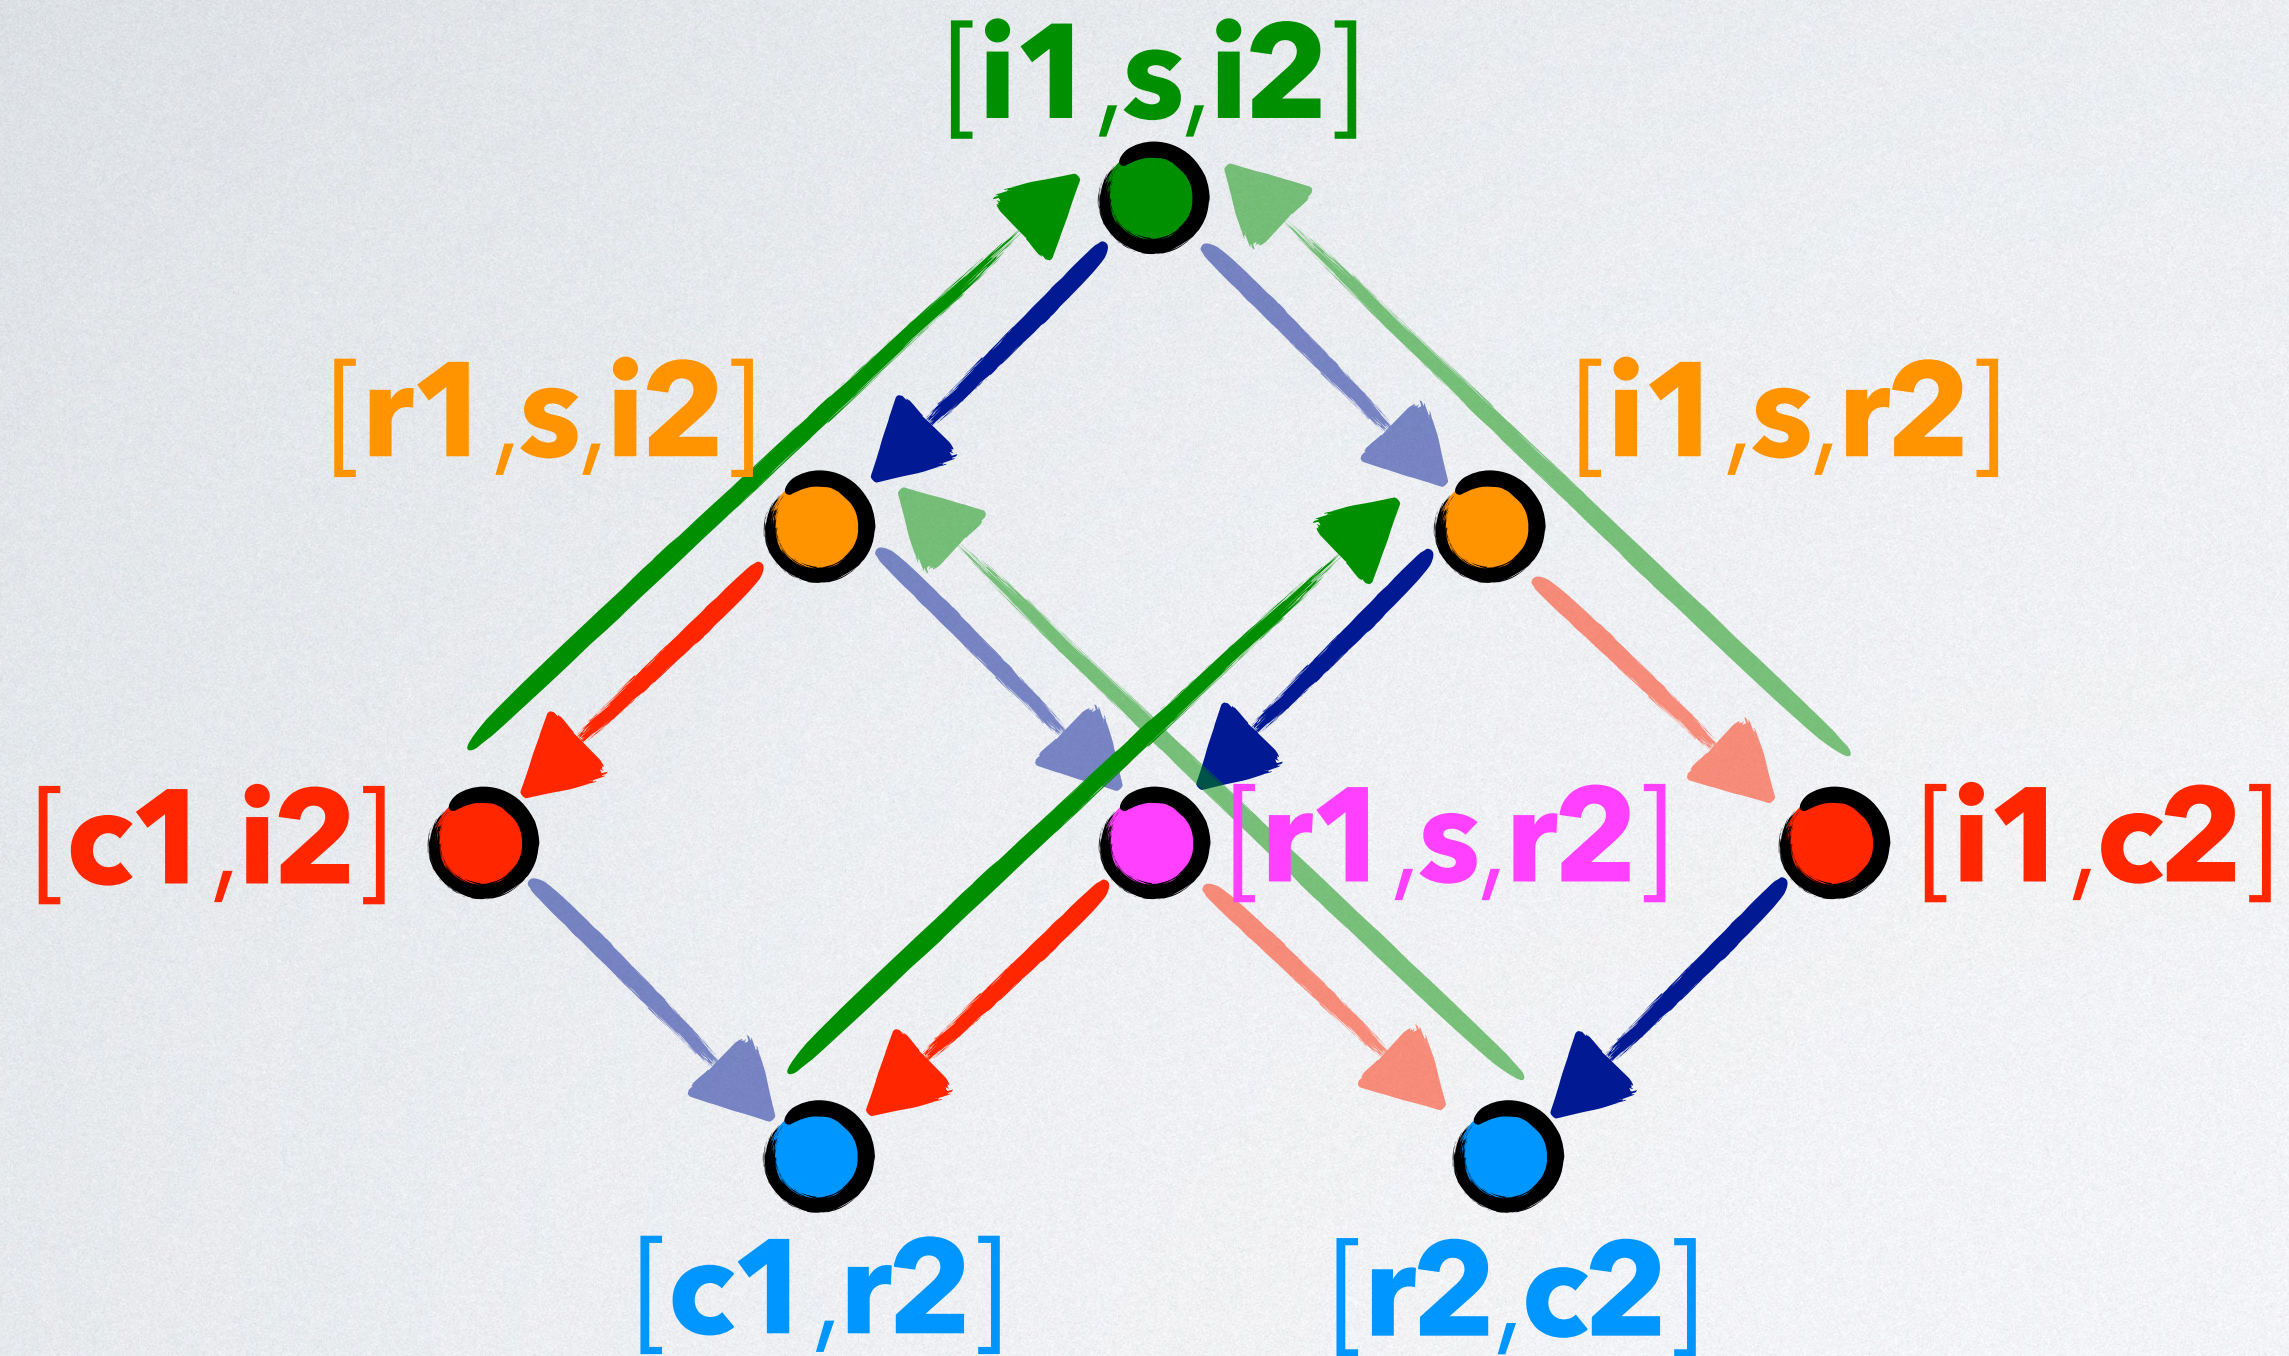


symmetric markings

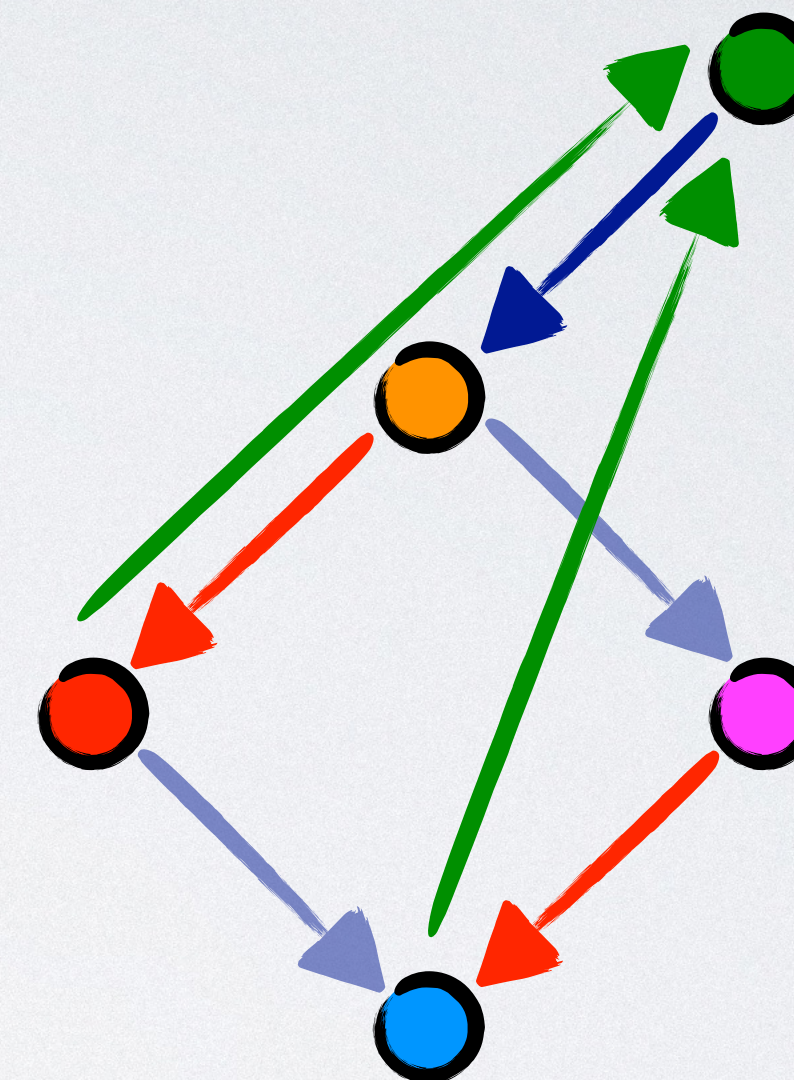


graph automorphisms

2. Symmetry reduction



complete state space



reduced state space

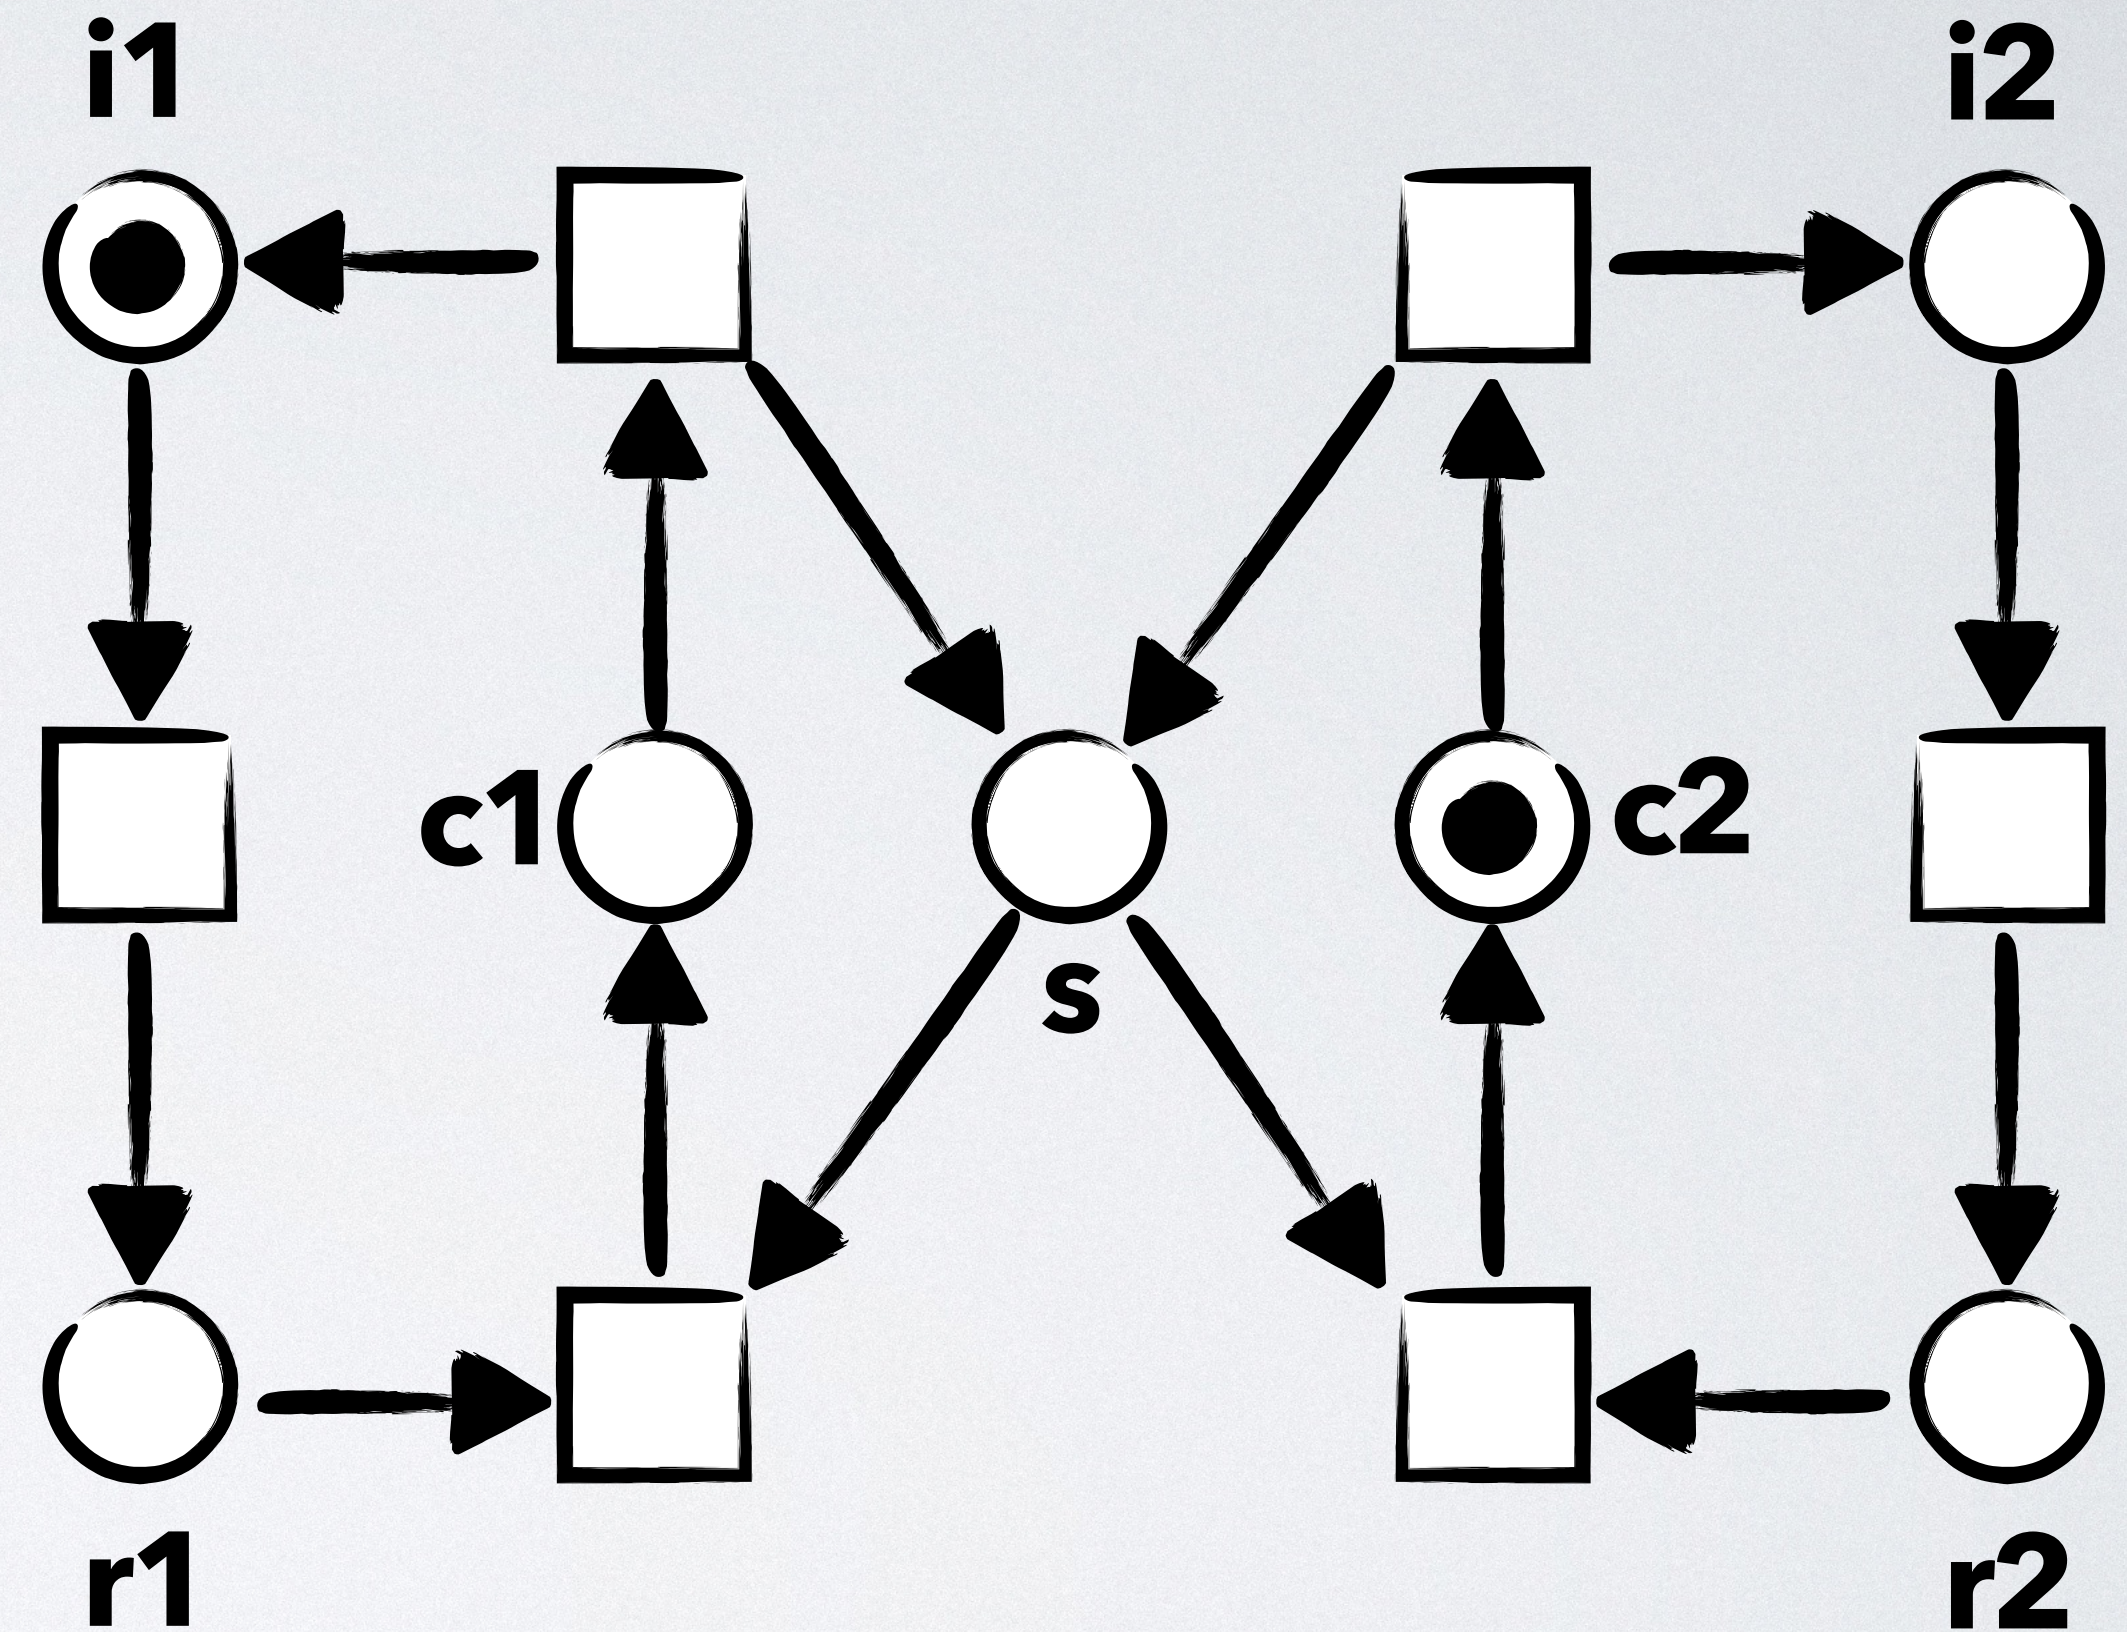
3. State compression

observation: the value of some places **depend** on the value of others

idea: do not store markings of "**implicit places**"

implementation: Petri net place invariants

$$m = \begin{matrix} & i1 & r1 & c1 & s & i2 & r2 & c2 \\ [1, & 0, & 0, & 0, & 0, & 0, & 1 \end{matrix}$$

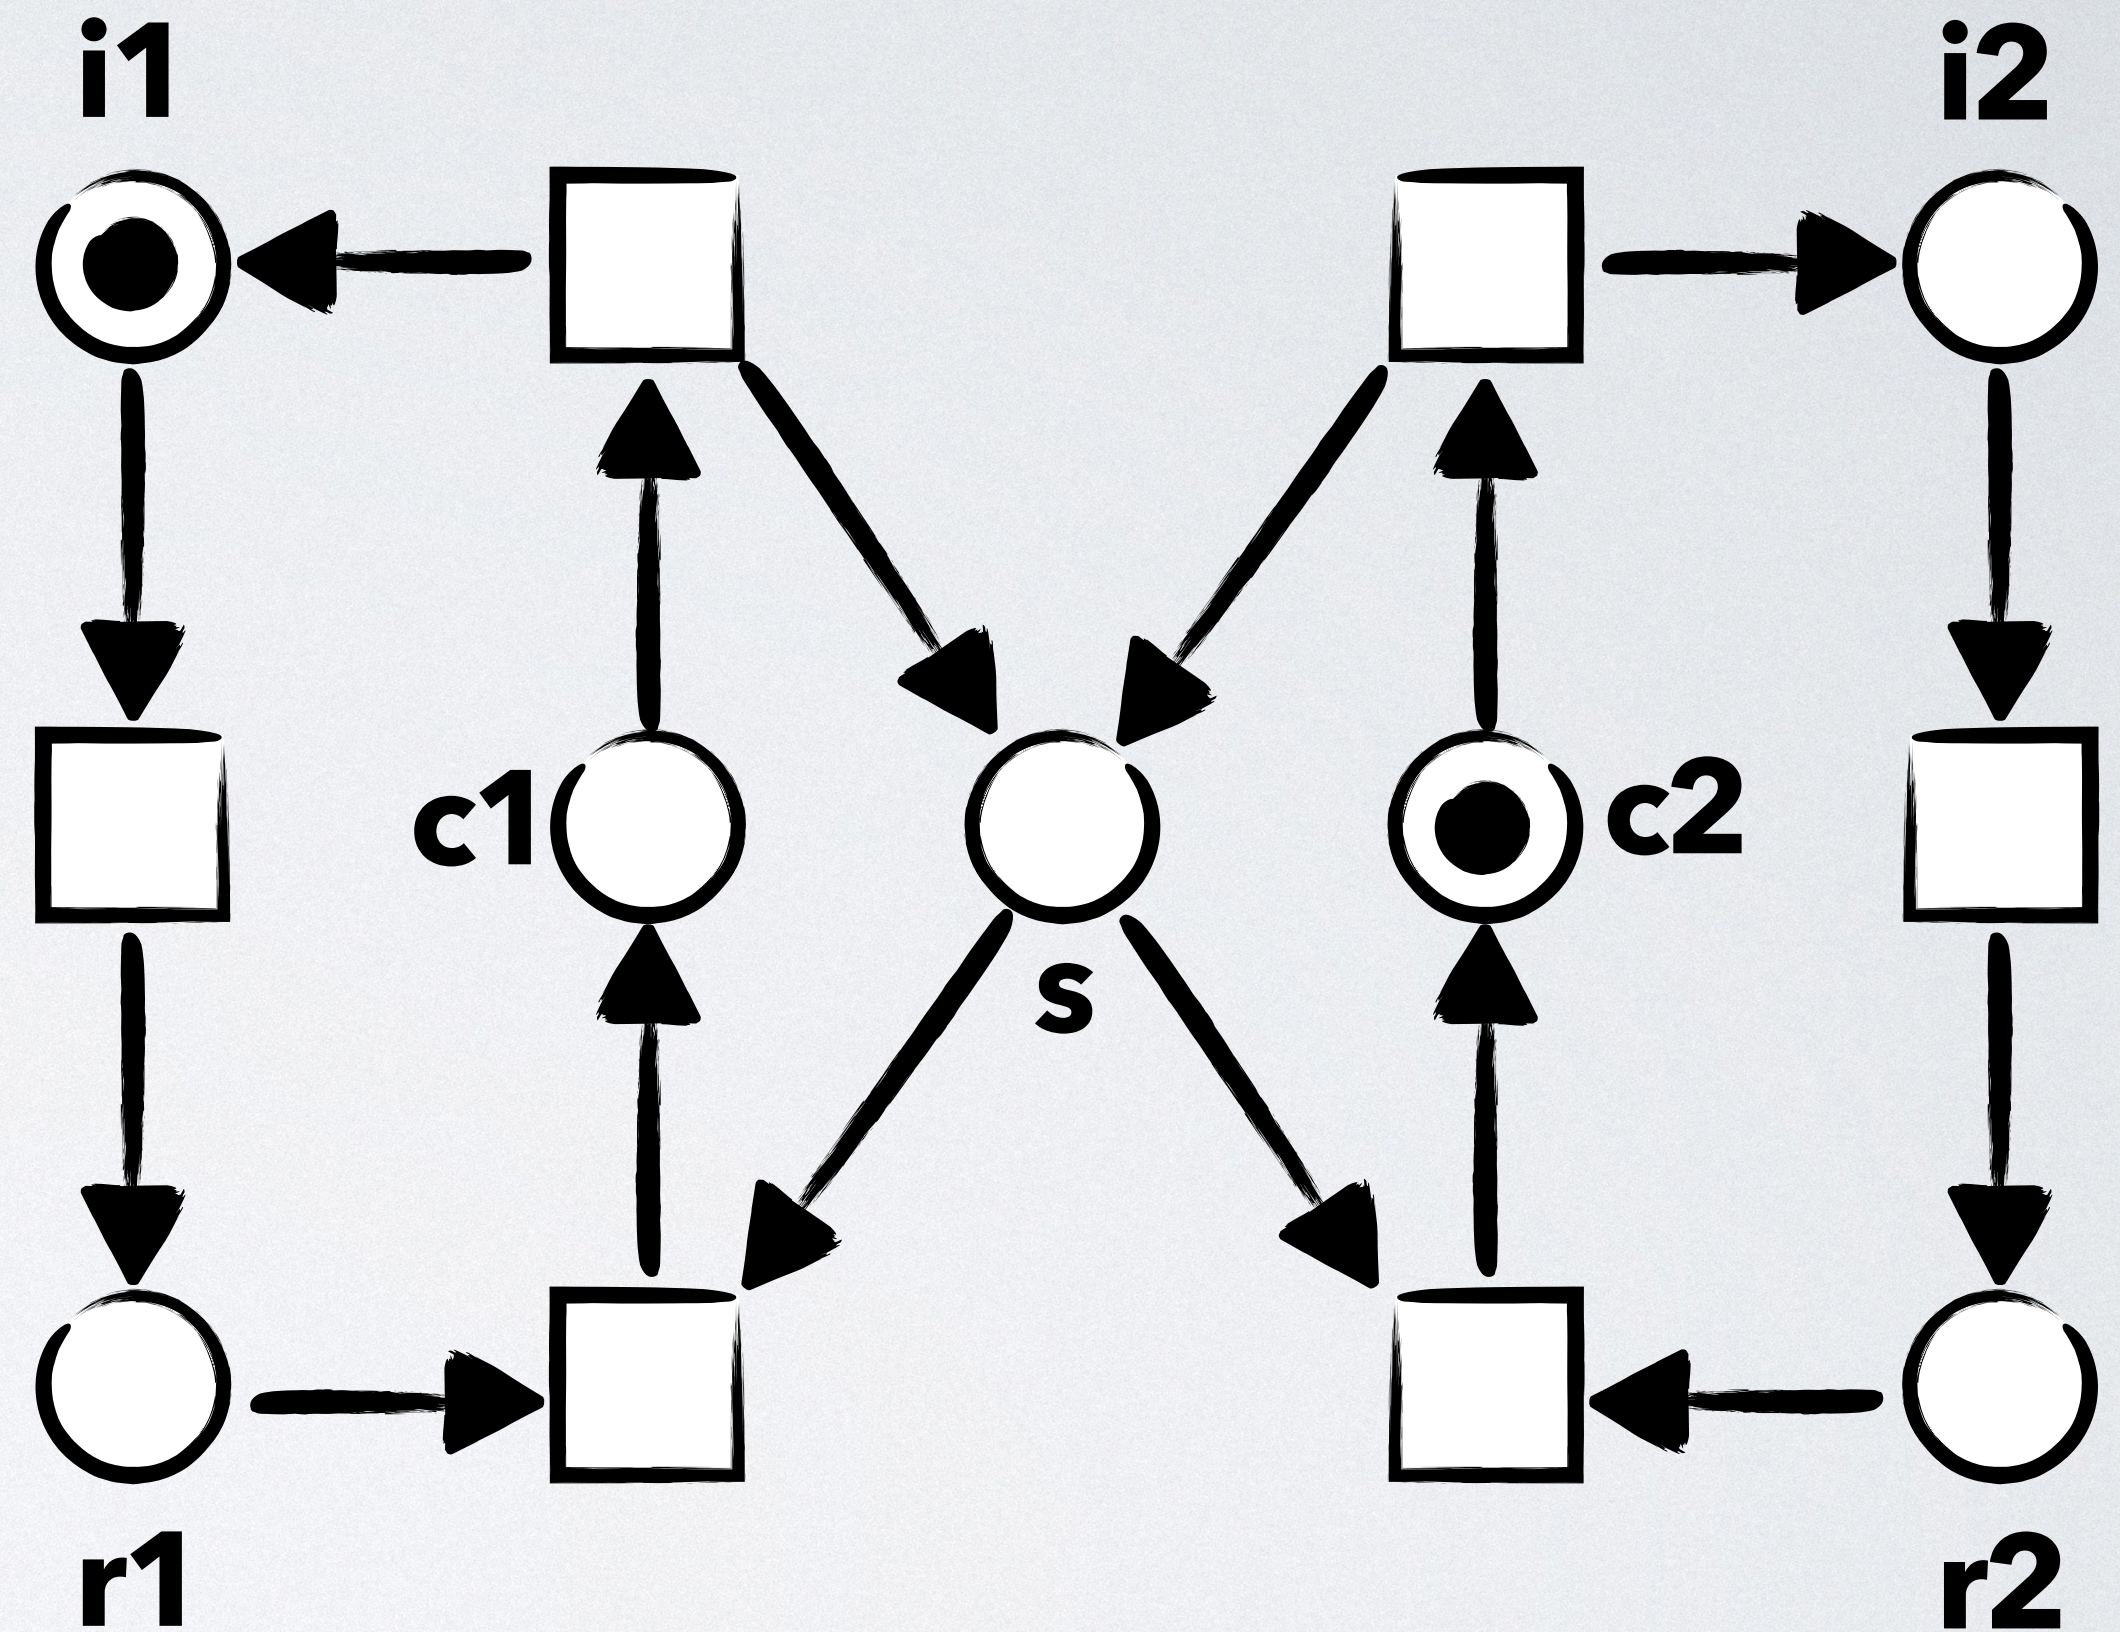


3. State compression

place invariants: for all
reachable markings m :

$$\begin{aligned}m(\mathbf{c1}) + m(\mathbf{c2}) + m(\mathbf{s}) &= 1 \\m(\mathbf{i1}) + m(\mathbf{c1}) + m(\mathbf{r1}) &= 1 \\m(\mathbf{i2}) + m(\mathbf{c2}) + m(\mathbf{r2}) &= 1\end{aligned}$$

$$m = \begin{matrix} & \mathbf{i1} & \mathbf{r1} & \mathbf{c1} & \mathbf{s} & \mathbf{i2} & \mathbf{r2} & \mathbf{c2} \\ m = & [1, 0, 0, 0, 0, 0, 1] \end{matrix}$$



3. State compression

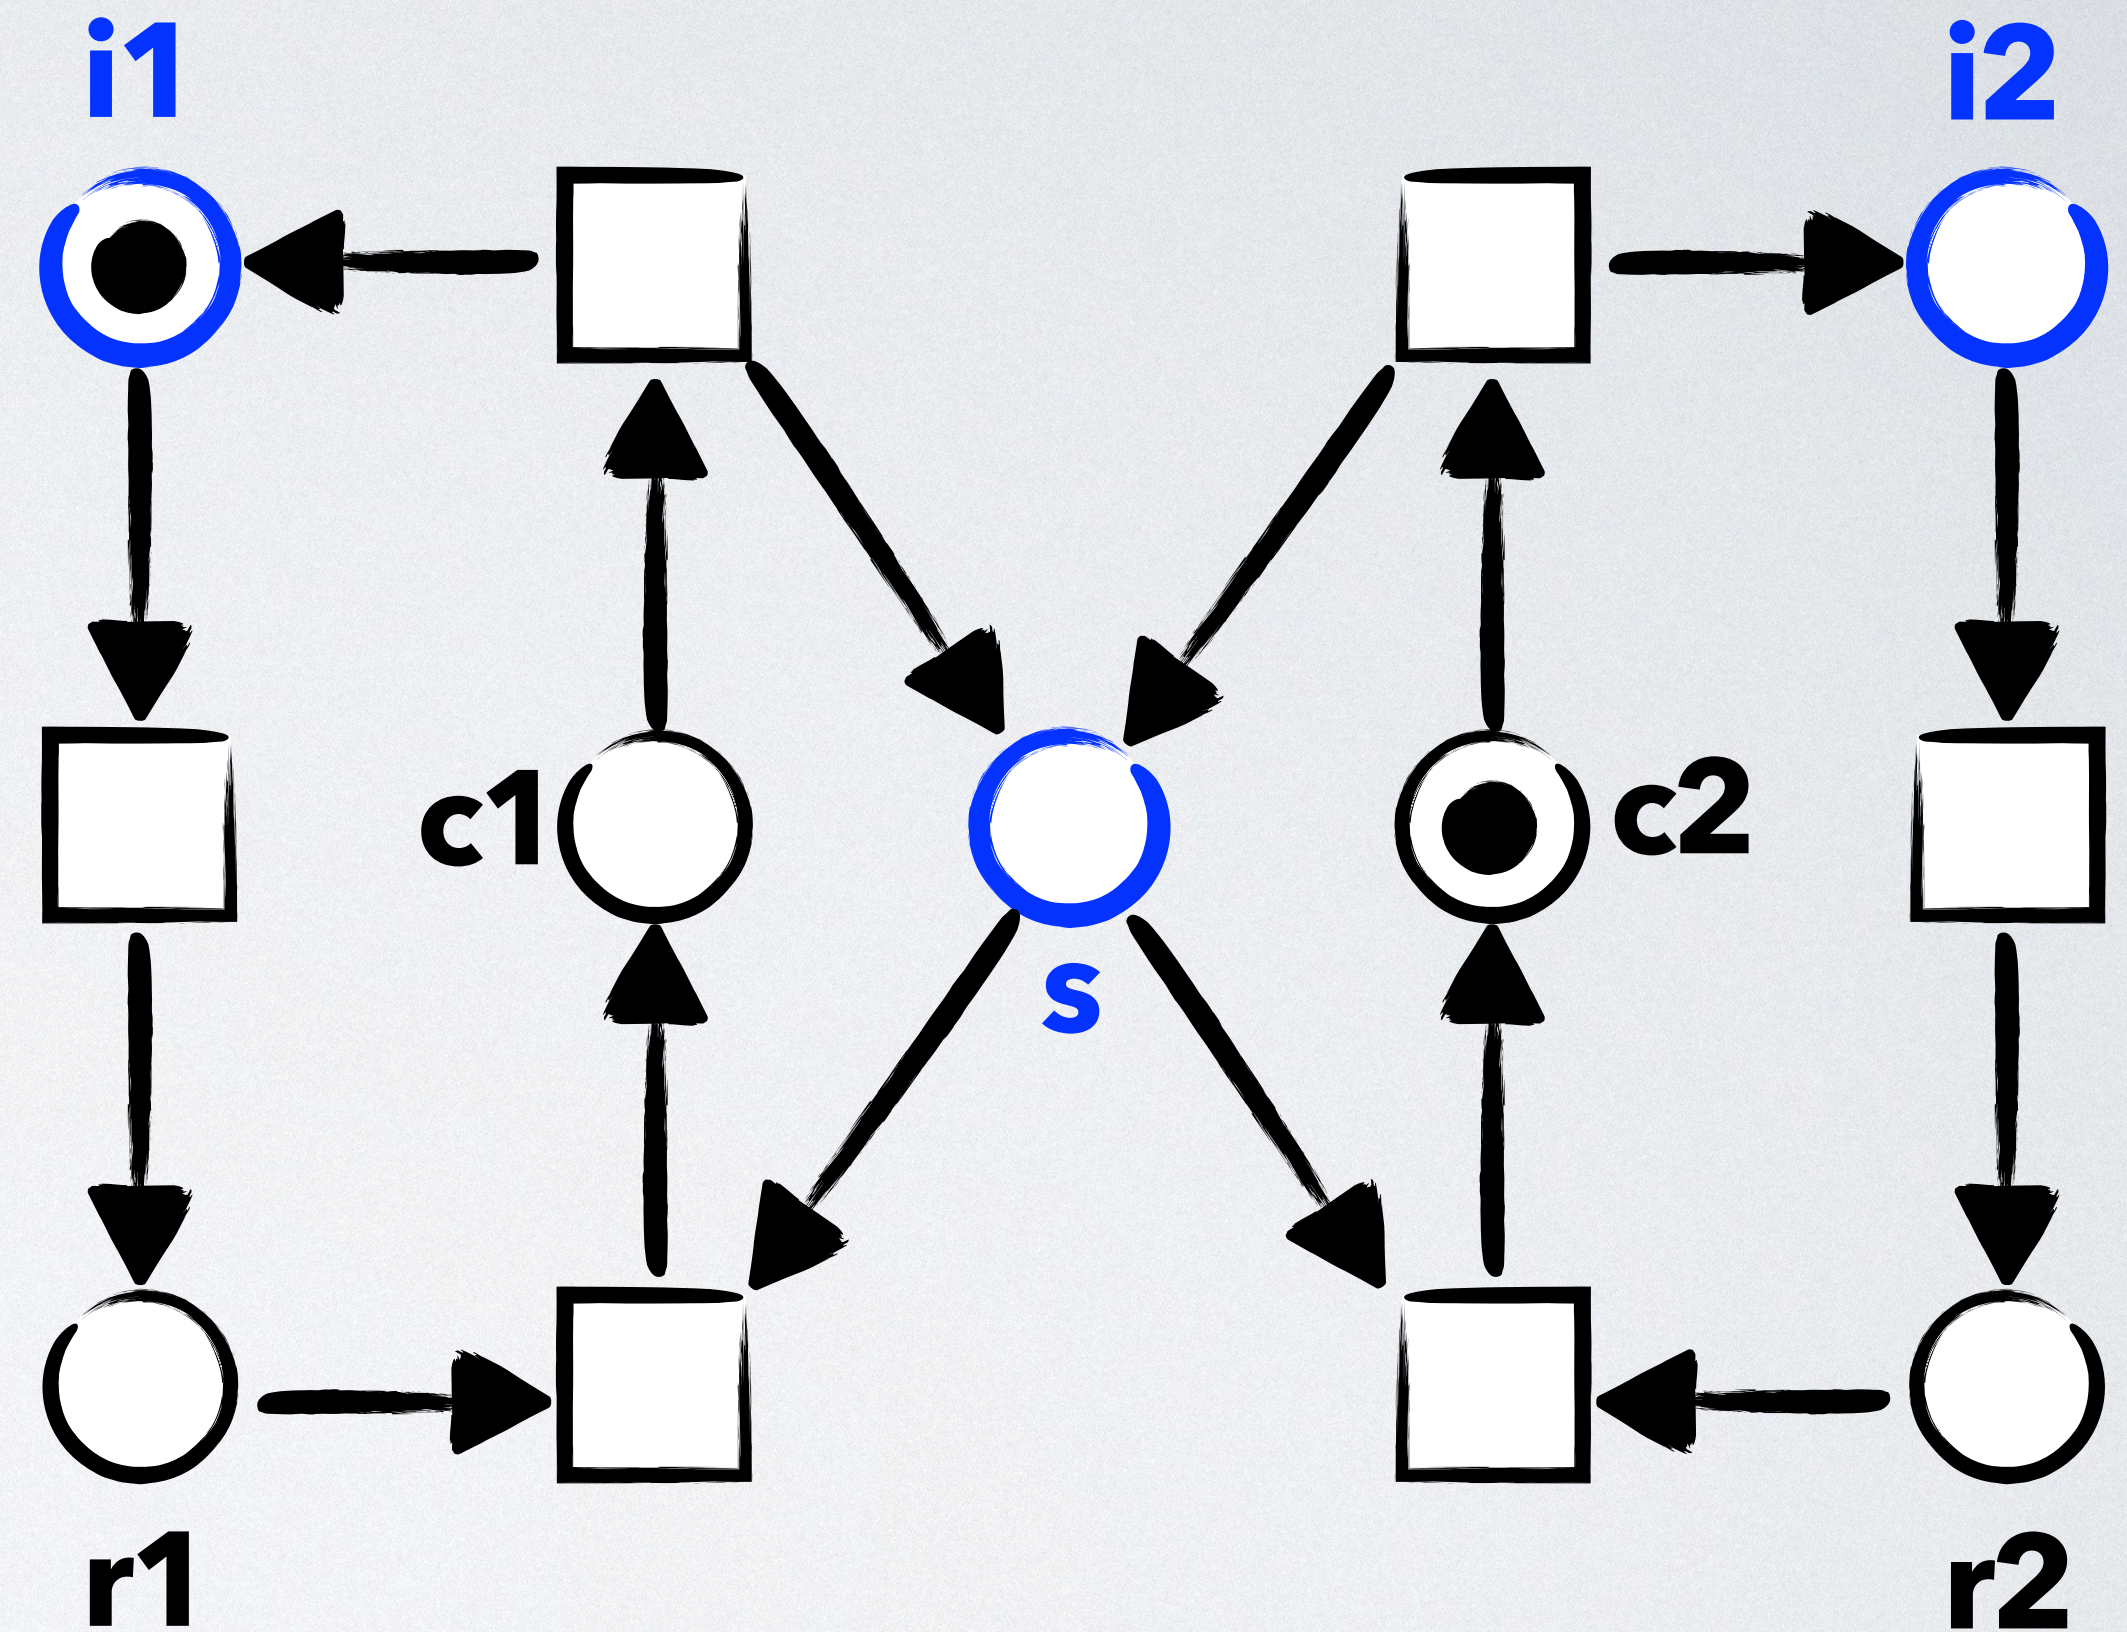
place invariants: for all
reachable markings m :

$$\begin{aligned}m(\mathbf{c1}) + m(\mathbf{c2}) + m(\mathbf{s}) &= 1 \\m(\mathbf{i1}) + m(\mathbf{c1}) + m(\mathbf{r1}) &= 1 \\m(\mathbf{i2}) + m(\mathbf{c2}) + m(\mathbf{r2}) &= 1\end{aligned}$$

therefore:

$$\begin{aligned}m(\mathbf{s}) &= 1 - m(\mathbf{c1}) - m(\mathbf{c2}) \\m(\mathbf{i1}) &= 1 - m(\mathbf{c1}) - m(\mathbf{r1}) \\m(\mathbf{i2}) &= 1 - m(\mathbf{c2}) - m(\mathbf{r2})\end{aligned}$$

$$m = \begin{matrix} \mathbf{i1} & \mathbf{r1} & \mathbf{c1} & \mathbf{s} & \mathbf{i2} & \mathbf{r2} & \mathbf{c2} \\ \hline 1 & 0 & 0 & 0 & 0 & 0 & 1 \end{matrix}$$



i1, **i2**, and **s** are **implicit**

Reduction techniques: implementation

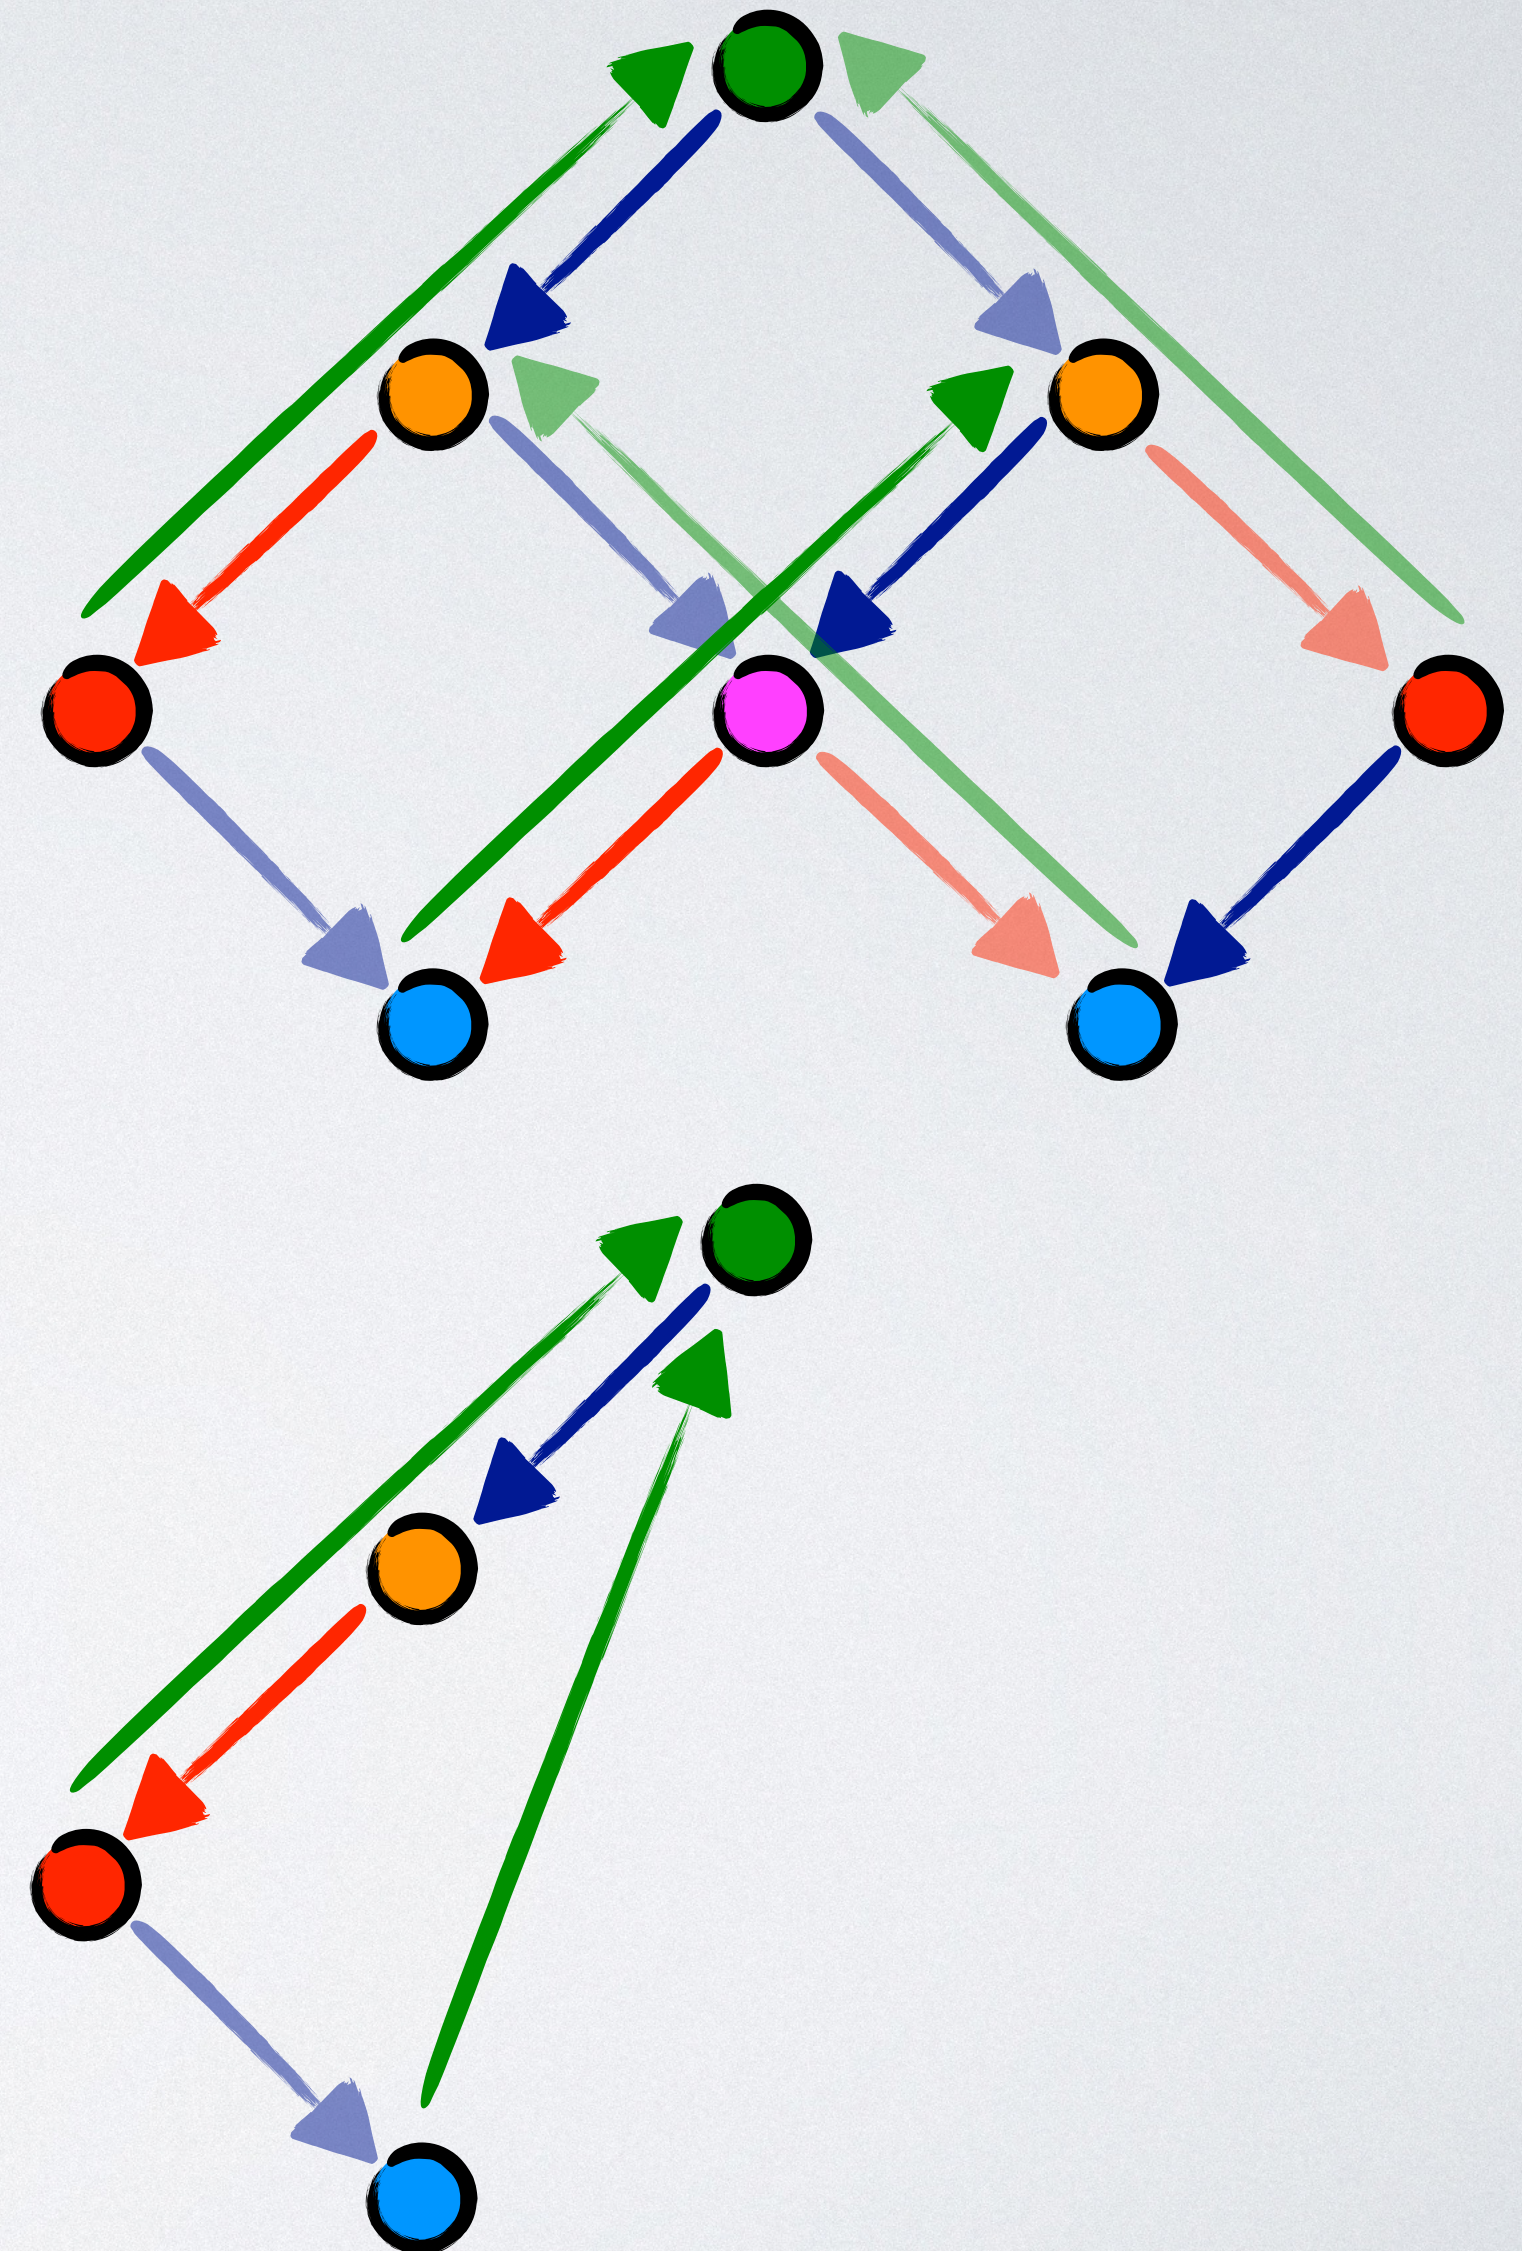
```
markings = []  
c = compressor()  
s = symmetries()  
search( $m_0, \varphi$ )
```

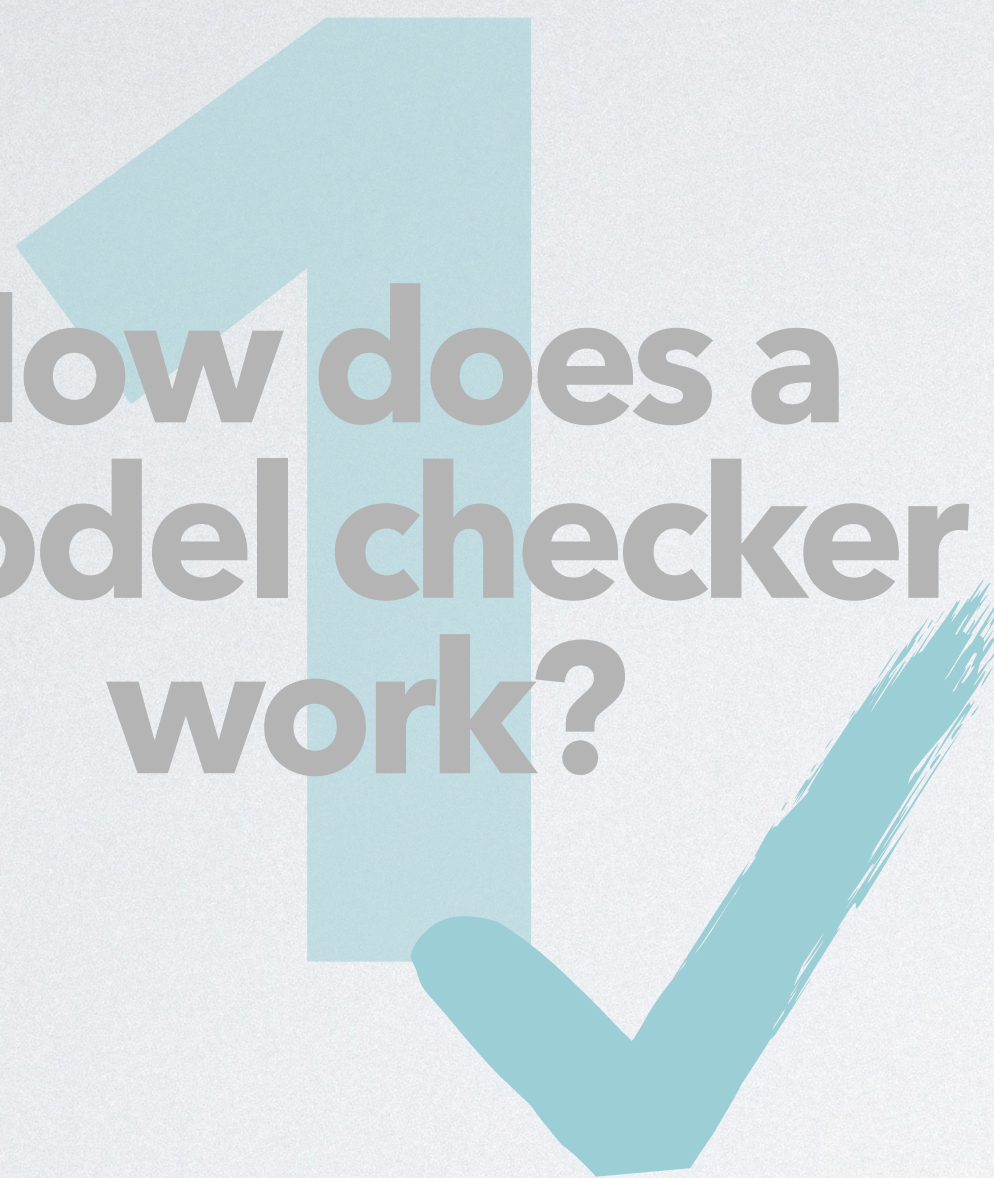
```
1  def search( $m, \varphi$ ):  
2      check( $m, \varphi$ )  
3      markings.add(c.compress( $m$ ))  
4      for t in selection(enabled( $m$ )):  
5           $m' = \text{fire}(m, t)$   
6          if not s.symm( $m'$ ) in markings:  
            search( $m', \varphi$ )
```

- **compressor** and **symmetries** preprocess the net
- **compress** creates shorter marking vectors
- **selection** chooses some enabled transitions
- **symm** checks if symmetric marking is already stored

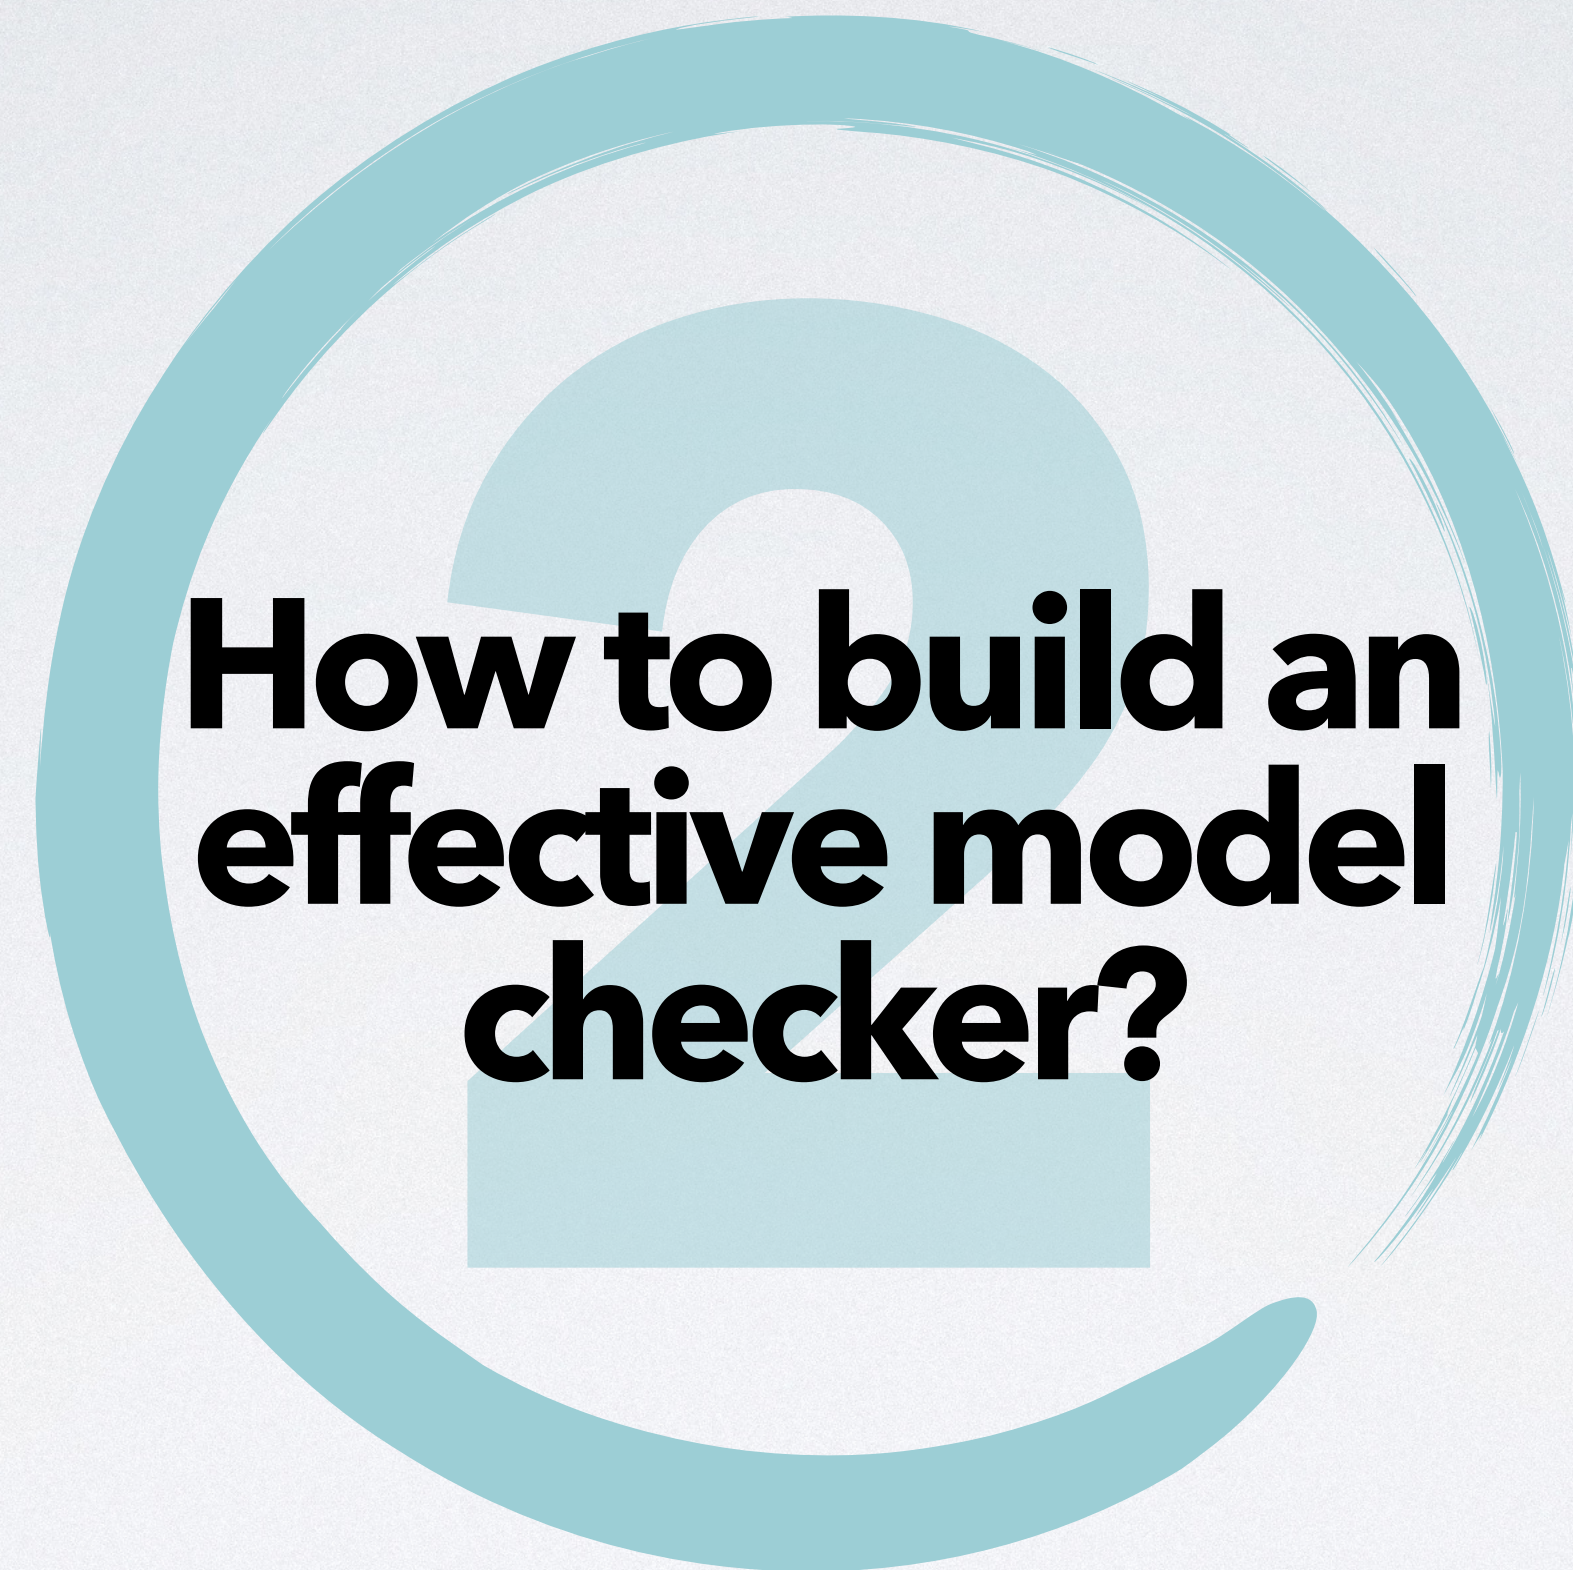
Reduction techniques: summary

- can be combined
- a lot of them work on the Petri net structure
- preprocessing pays off
- some are Petri net exclusive
- decades of research

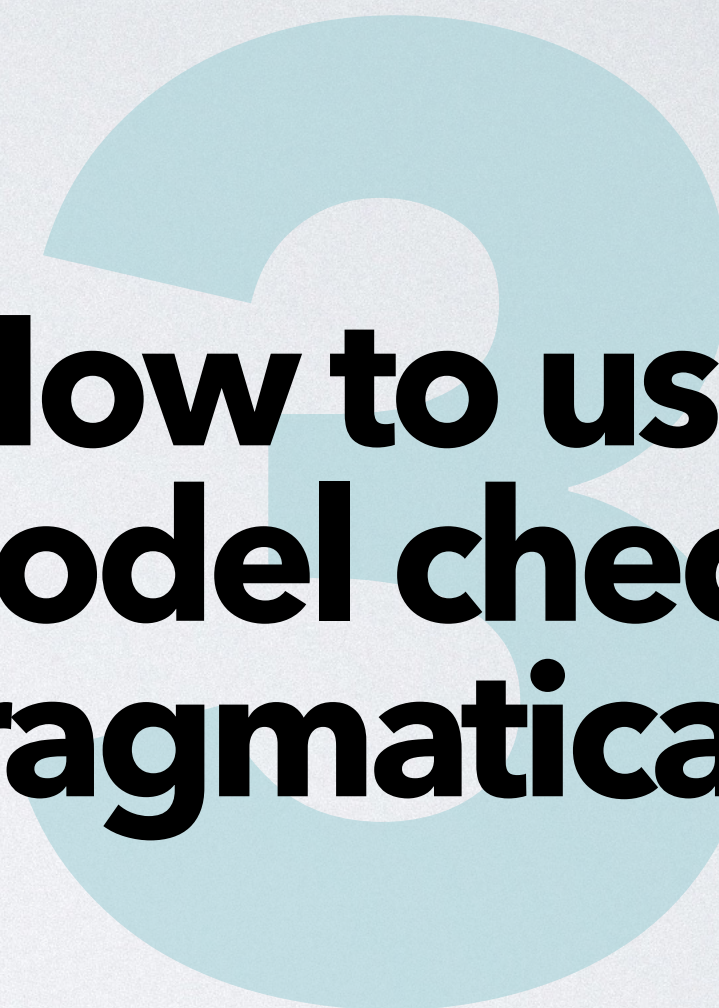




How does a
model checker
work?



**How to build an
effective model
checker?**



**How to use a
model checker
pragmatically?**

Key: Pragmatism

good news

model checking
is decidable

bad news

at a devastating
complexity

- in the end: we only optimize a small depth-first search...
... that is run billions of times
- we need to understand every aspect
- pragmatism is key: we want a result at all costs

The programming language: C++

manual memory
management

type system

low-level
optimization

preprocessor

optional
object orientation

threading



portability



Used anti-patterns

- heavy use of **preprocessor** (conditional compilation, compile-time decisions, architecture-dependent constants)
- **god objects**
- a lot of **global** variables
- no standard libraries/
generic algorithms (STL, Boost)
- remember: performance, not reusability

```
unsigned int depth_first() {
#ifdef DEPTH_FIRST
    ostream* graphstream = NULL;
    unsigned int i;
    State* NewState;
#endif
#ifdef CYCLE
    bool IsCyclic;
    unsigned int silentpath; // nr of consecuti
    Transition** fl;
#endif
    // init initial marking and hash table
    isbounded = 1;
#ifdef CYCLE
    // organize output file for -m/-M/-g/-G par
    if (Globals::gmflg) {
        graphstream = new ofstream(Globals::gra
        if (graphstream->fail()) {
            abort(4, "cannot open graph output
            _cfilename_(Globals::graphfile));
        }
    }
    if (Globals::GMflg) {
        graphstream = &std::cout;
    }
#else
    silentpath = 0;
#endif
#ifdef DISTRIBUTE
    if defined(SYMMETRY) && SYMMINTEGRATION==1
        Trace = new SearchTrace [Globals::Places[0]
#endif
    // initialize hash table
    for (i = 0; i < HASHSIZE; ++i) {
#ifdef BITHASH
        BitHashTable[i] = 0;
#else
        binHashTable[i] = NULL;
#endif
    }
#endif
#ifdef WITHFORMULA
```


Data structures

frontend (parser)

syntactic sugar

class hierarchies

generated code/
standard libraries

focus on simplicity

backend (verification)

canonic representation

flat C-style arrays

simple and user-specified
data types

focus on performance

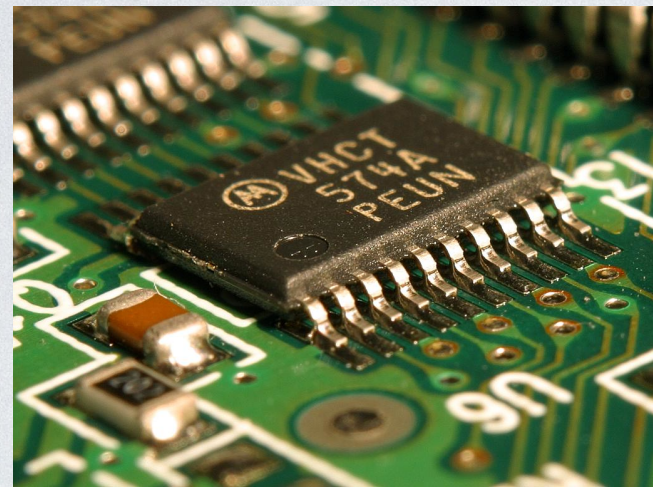
A remark on complexity

- While solving an EXPSPACE-complete problems, **don't be afraid of complexity!**
- Preprocessing and optimizations can have dramatic impact.
- NP-complete (SAT-solving) or NP-hard (integer linear programming) problems can be "feasible".

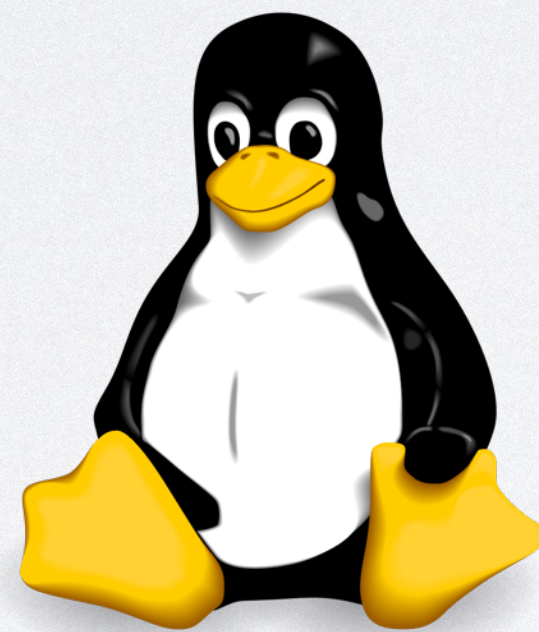
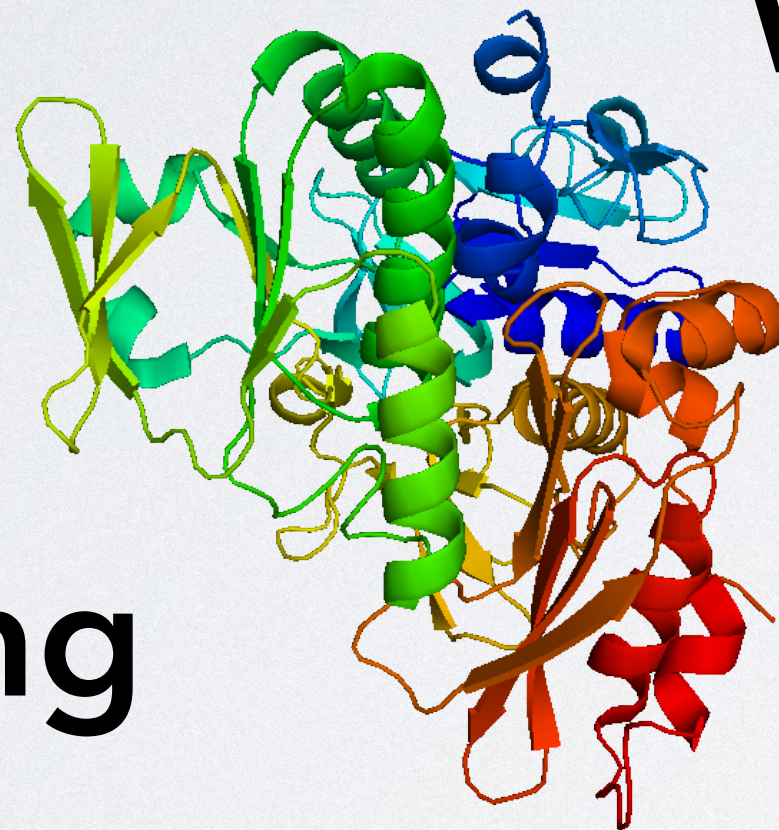
Formalism

modeling formalisms

WebSphere® software



WS-BPEL



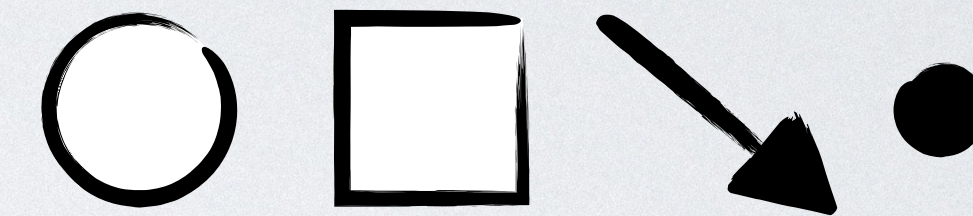
AI planning

- varying feature set
- domain dependent
- moving target, short lifespan

compilers



verification formalism



low-level Petri nets

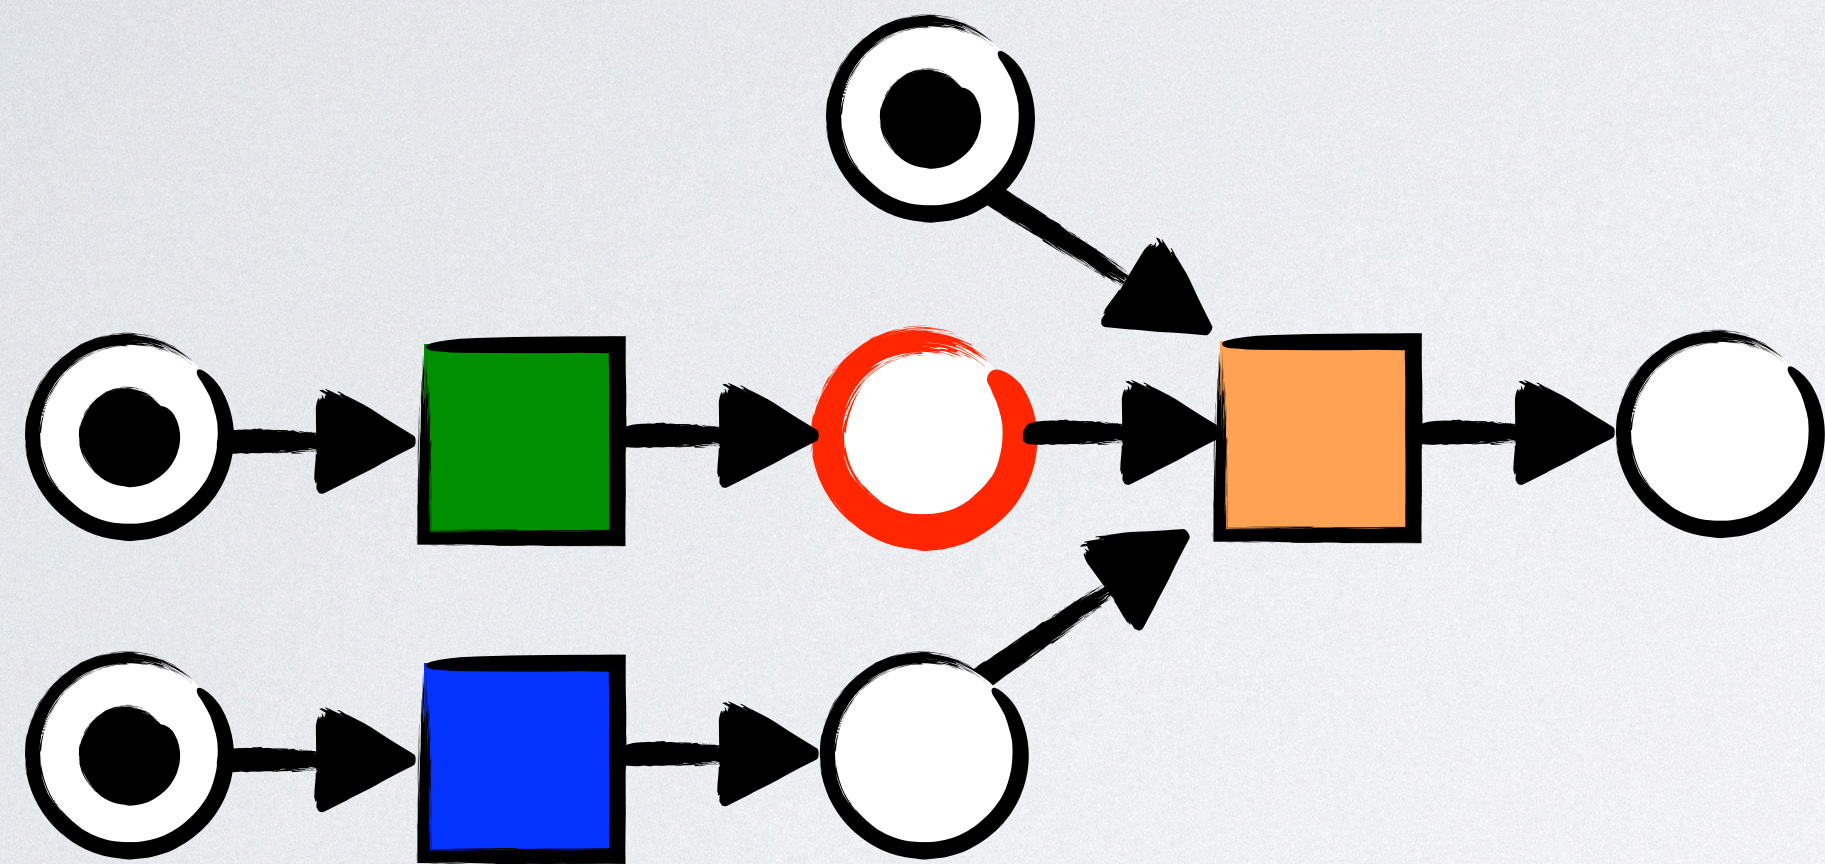
- no structural restrictions
- domain independent
- sound, mature theory

Petri net model checking

- exploit domain knowledge to optimize core functions
 - firing transitions
 - storing markings
- know lifecycle of concepts
- often contradicts object orientation

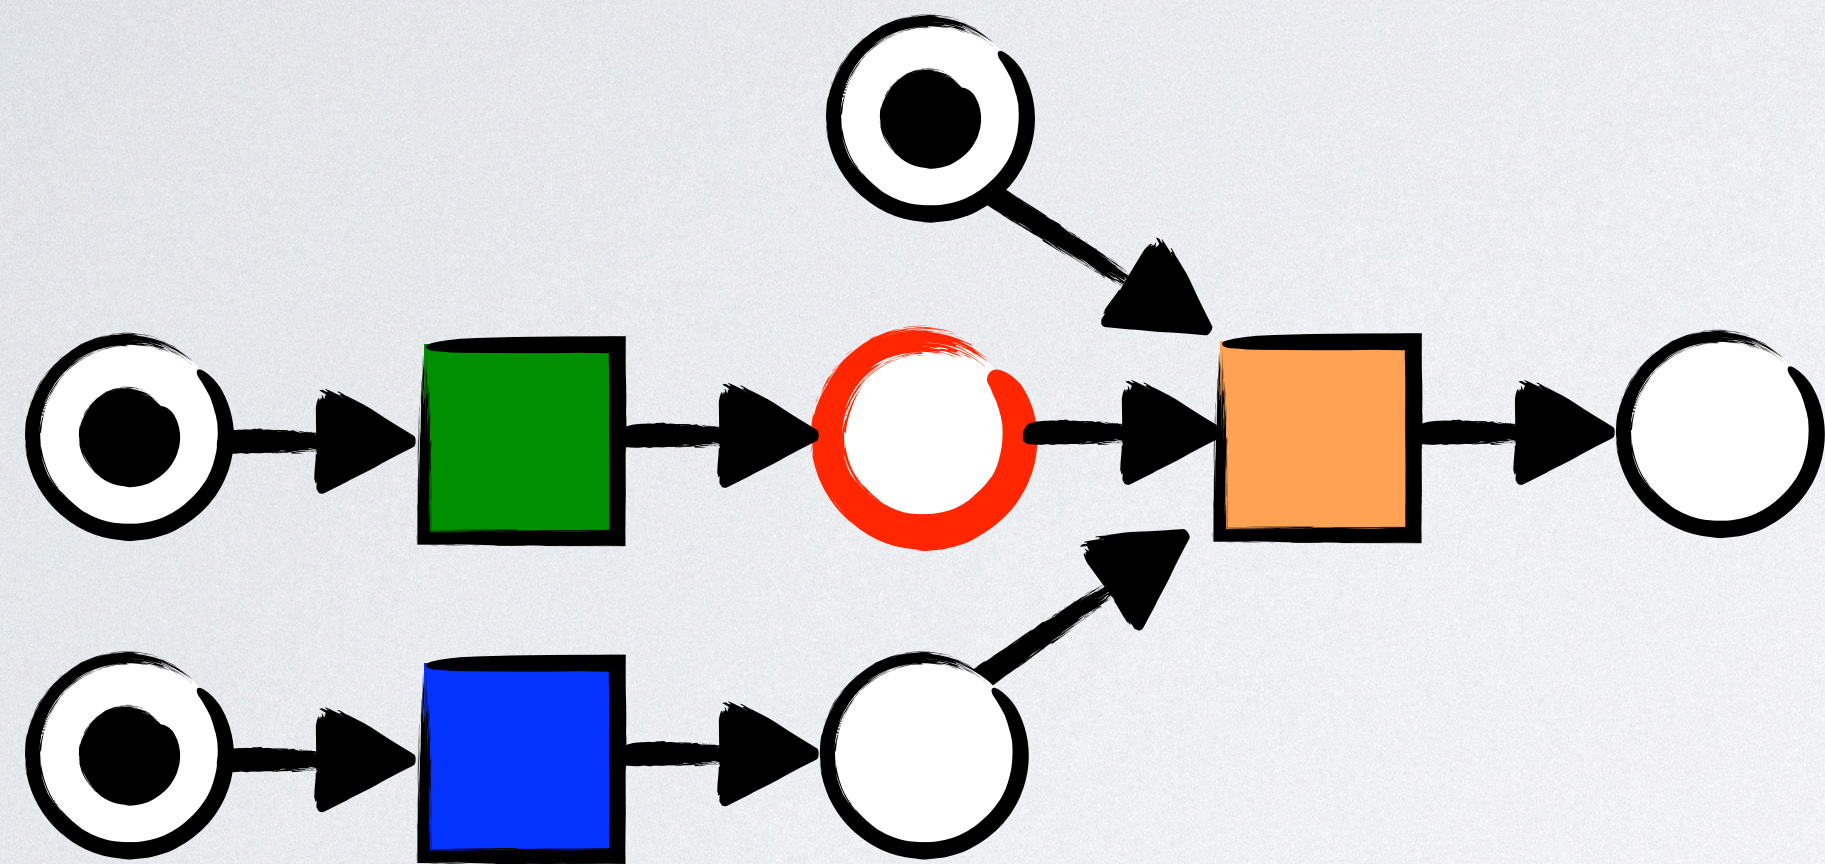
Checking enabledness

orange transition
is disabled



Checking enabledness

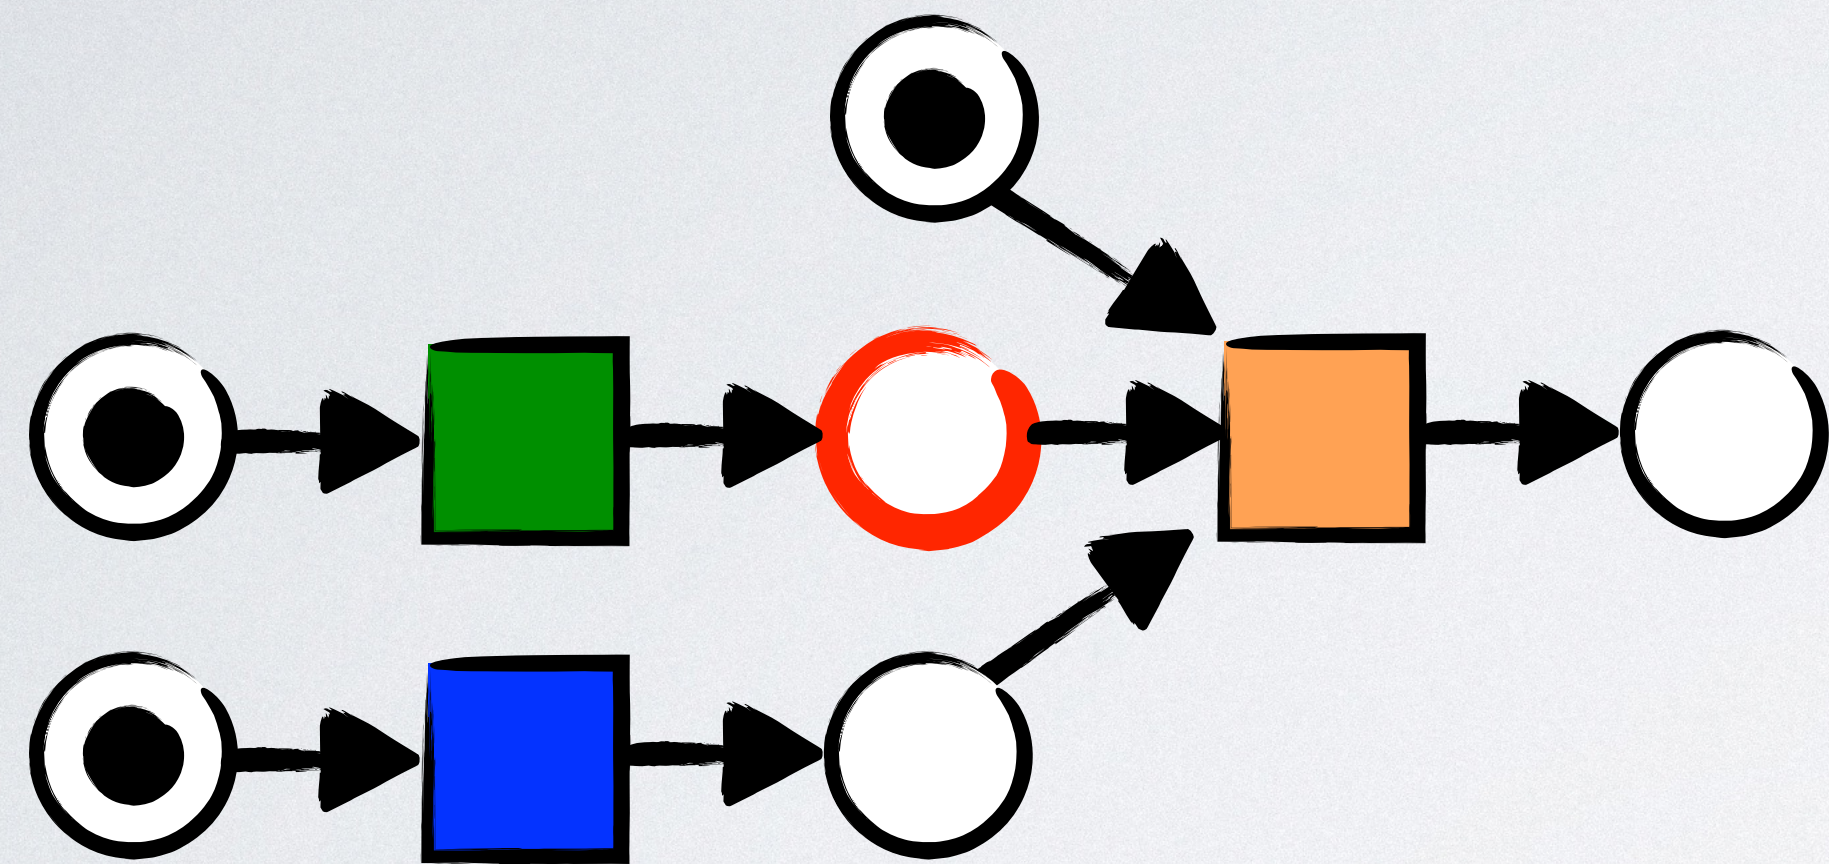
orange transition
is disabled



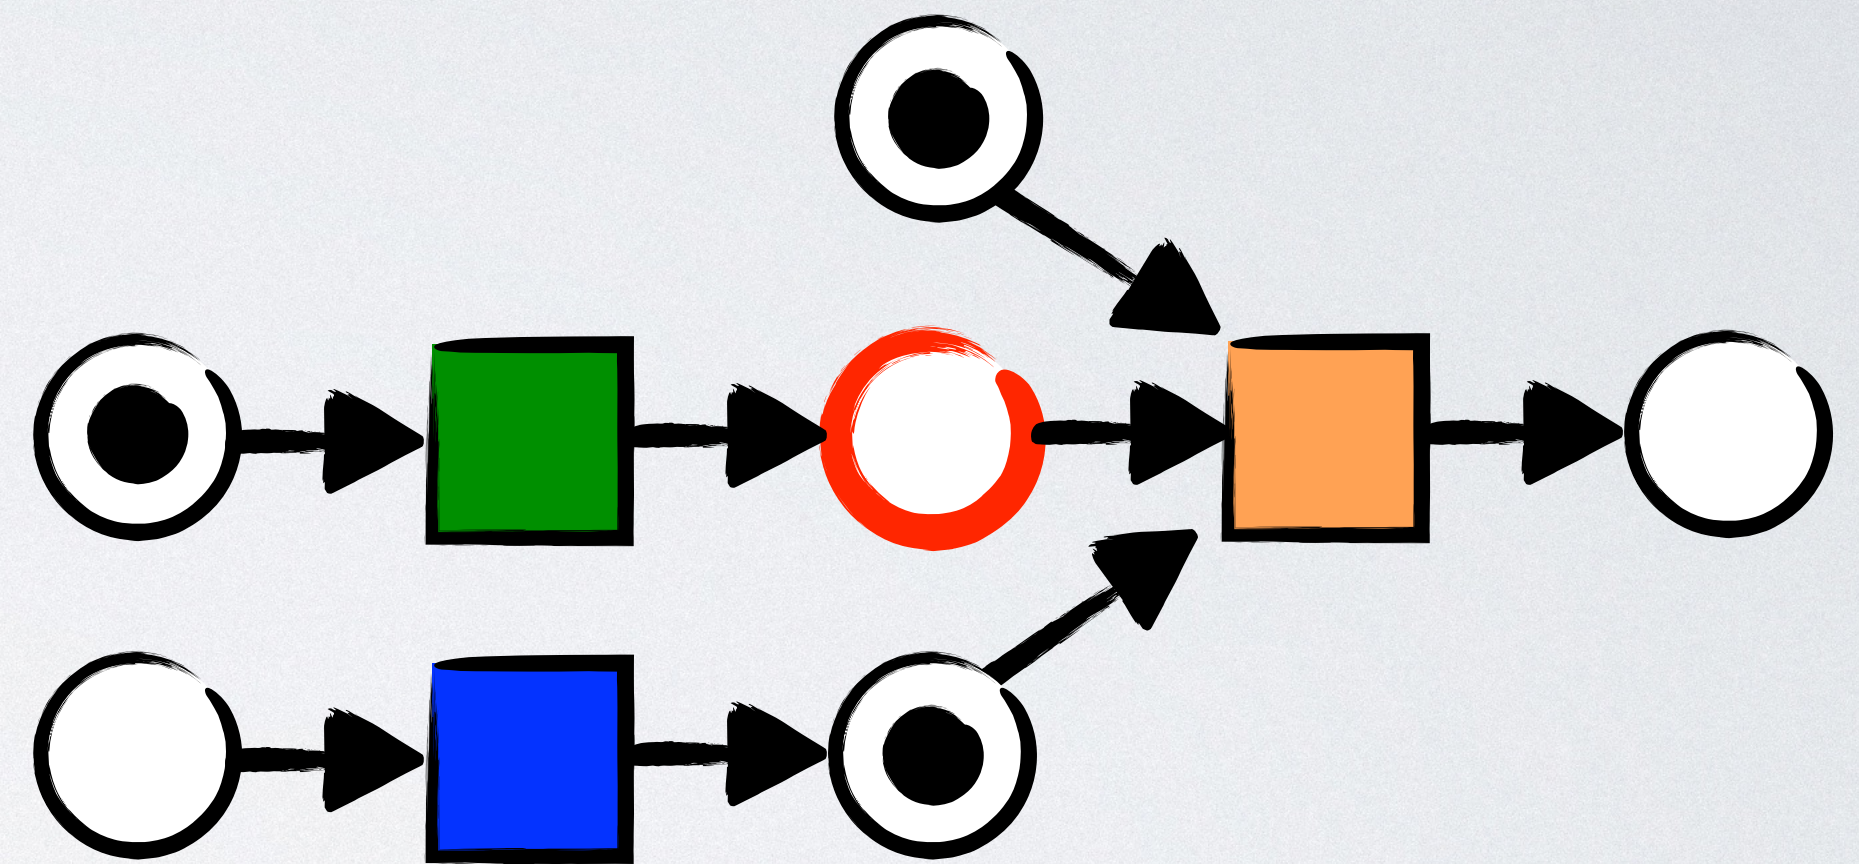
scapegoat
place

Checking enabledness

orange transition
is disabled



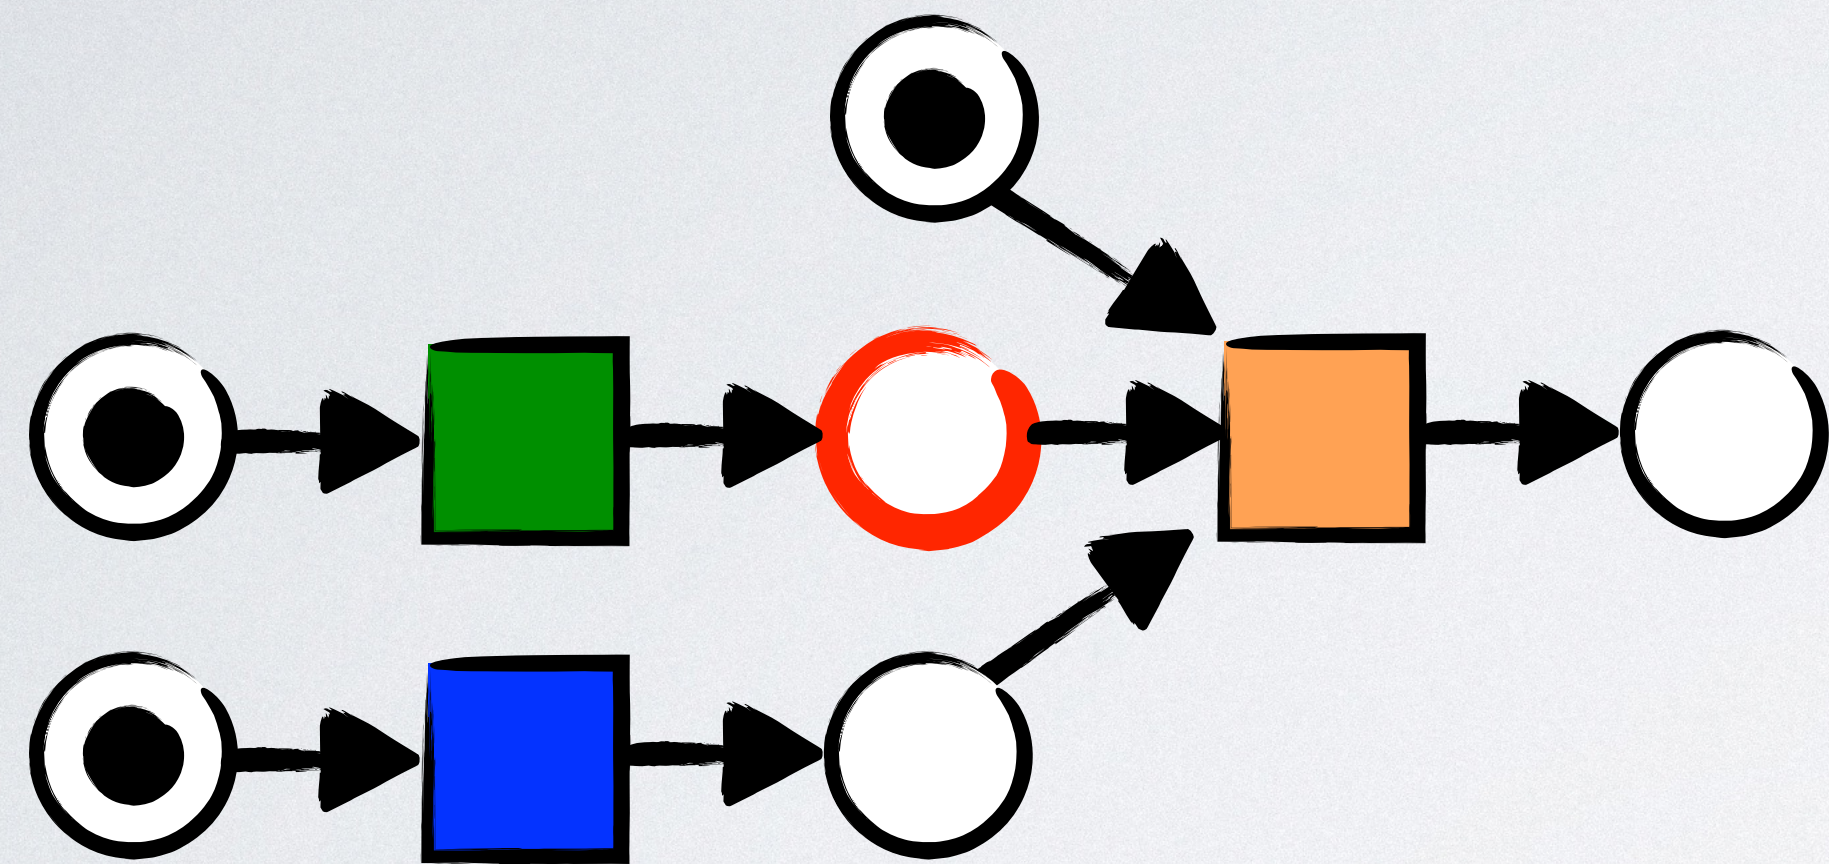
blue
transition
fires



scapegoat
place

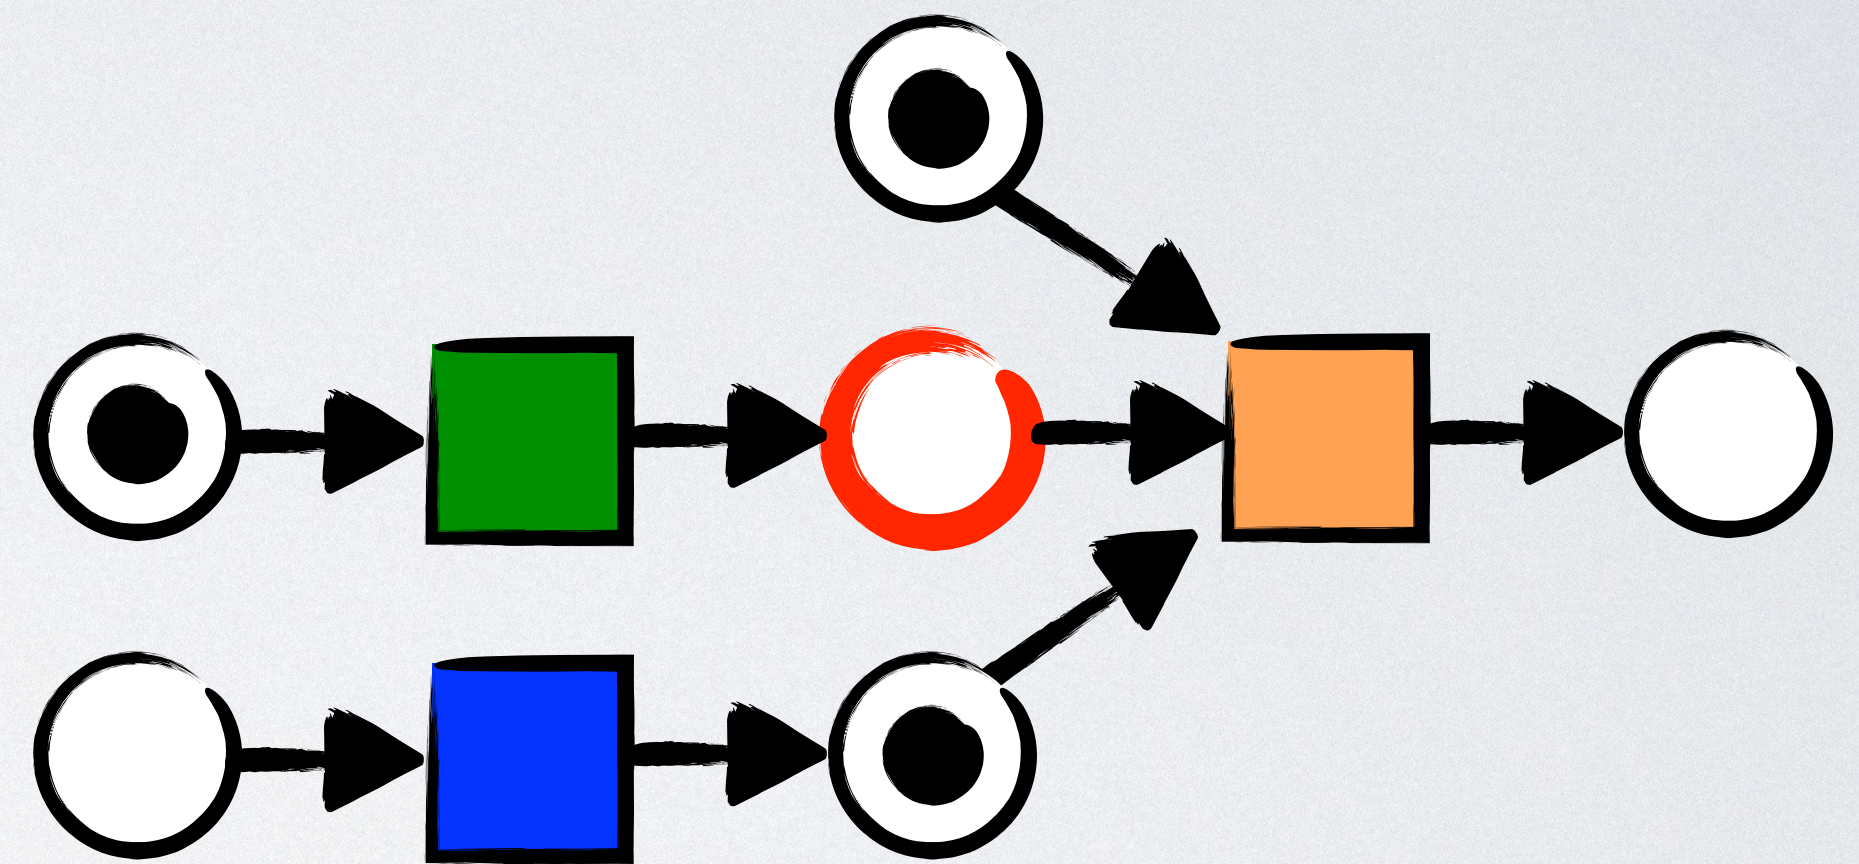
Checking enabledness

orange transition
is disabled



scapegoat
place

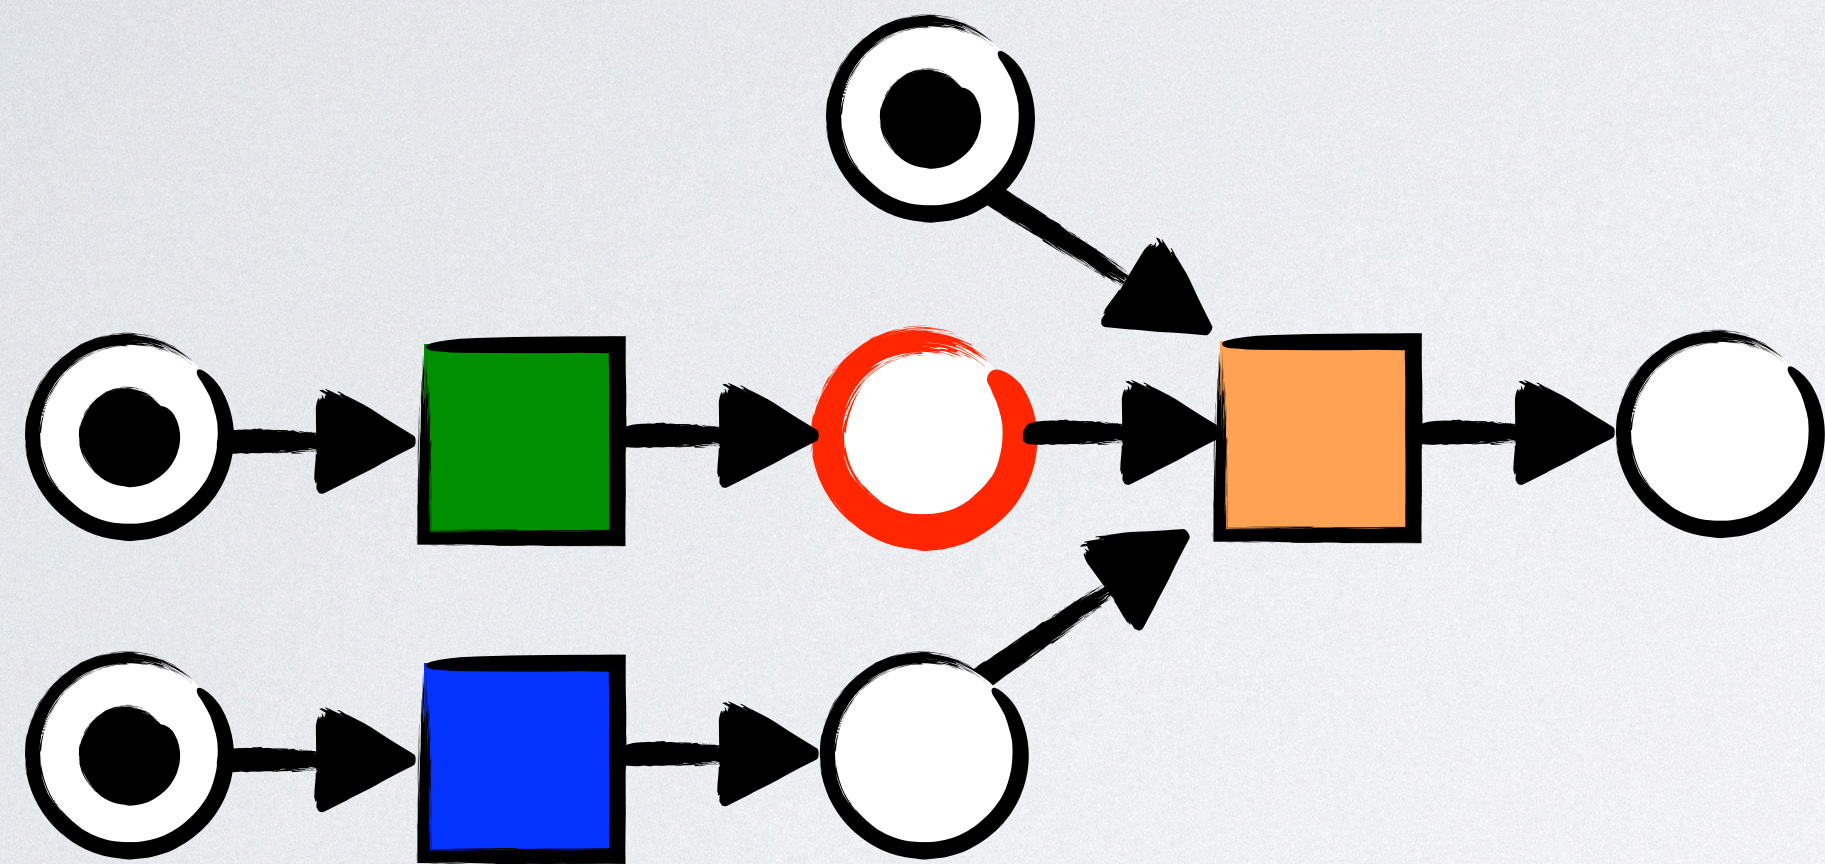
blue
transition
fires



scapegoat place
still unmarked

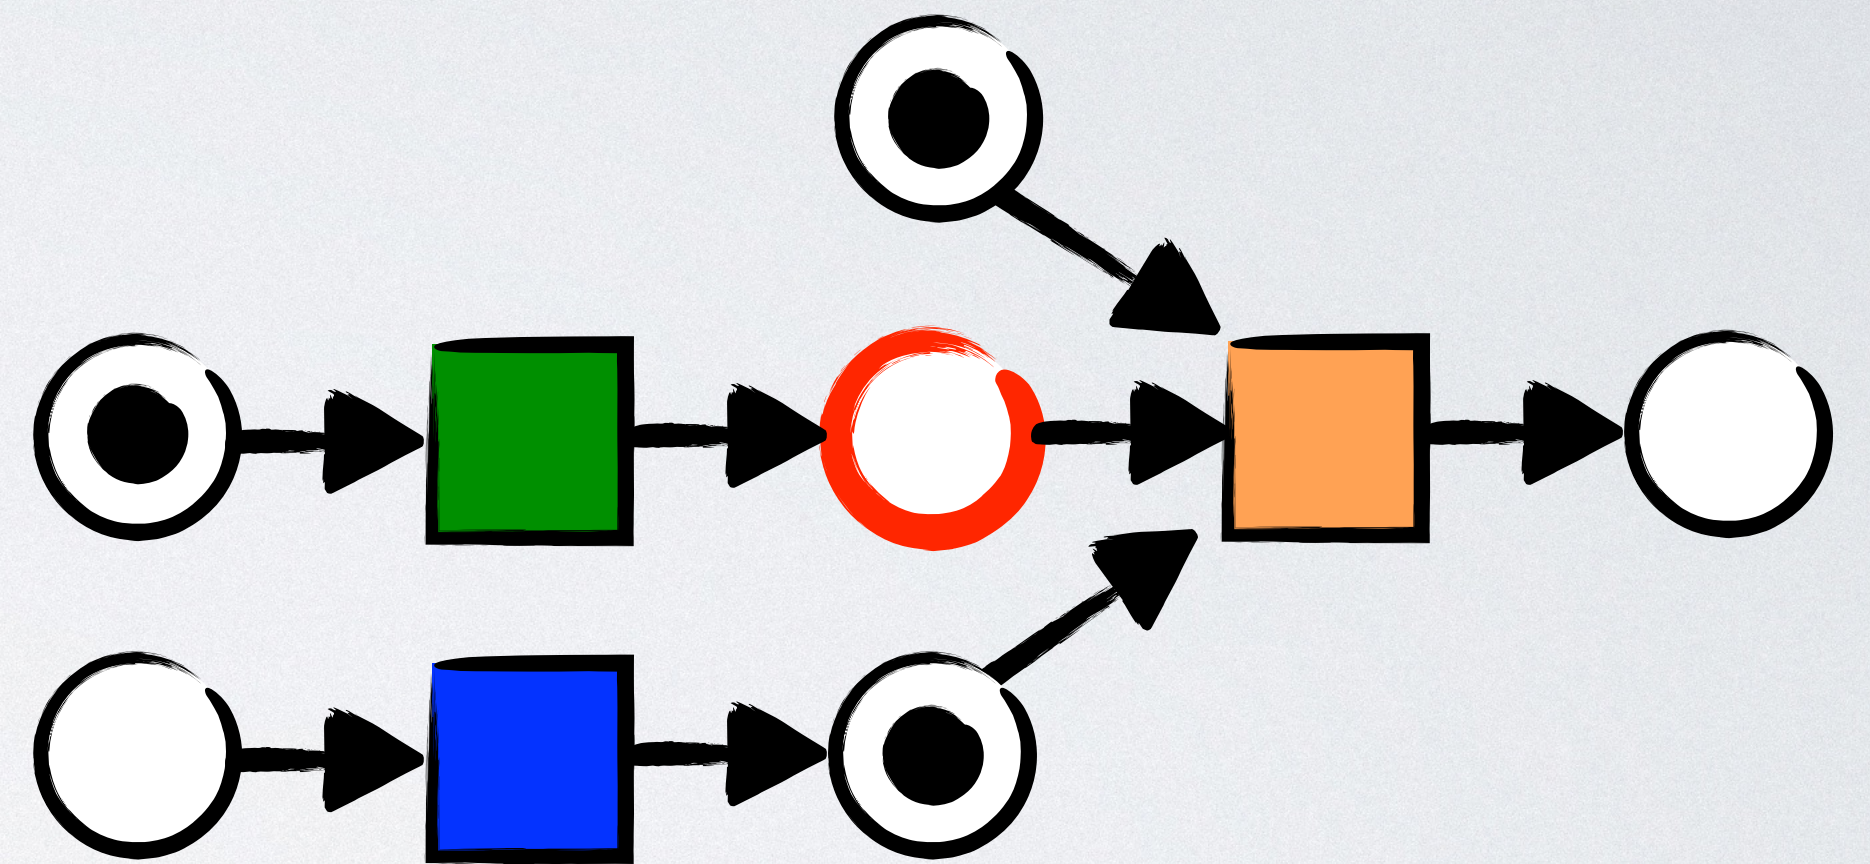
Checking enabledness

orange transition
is disabled



scapegoat
place

blue
transition
fires



scapegoat place
still unmarked



orange transition
is still disabled

Storing markings

markings

[0, 1, 3, 3, 1]

[0, 1, 3, 2, 0]

[0, 1, 3, 2, 1]

bit vectors

00 01 11 11 01

00 01 11 10 00

00 01 11 10 01

bin tree

00 01 11 11 01

00 01 11 10 00

00 01 11 10 01

Storing markings

markings

[0, 1, 3, 3, 1]

[0, 1, 3, 2, 0]

[0, 1, 3, 2, 1]

bit vectors

00 01 11 11 01

00 01 11 10 00

00 01 11 10 01

bin tree

00 01 11 11 01

00 01 11 10 00

00 01 11 10 01

Storing markings

markings

[0, 1, 3, 3, 1]

[0, 1, 3, 2, 0]

[0, 1, 3, 2, 1]

bit vectors

00 01 11 11 01

00 01 11 10 00

00 01 11 10 01

bin tree

1 01
00 01 11 10 00
00 01 11 10 01



Storing markings

markings

[0, 1, 3, 3, 1]

[0, 1, 3, 2, 0]

[0, 1, 3, 2, 1]

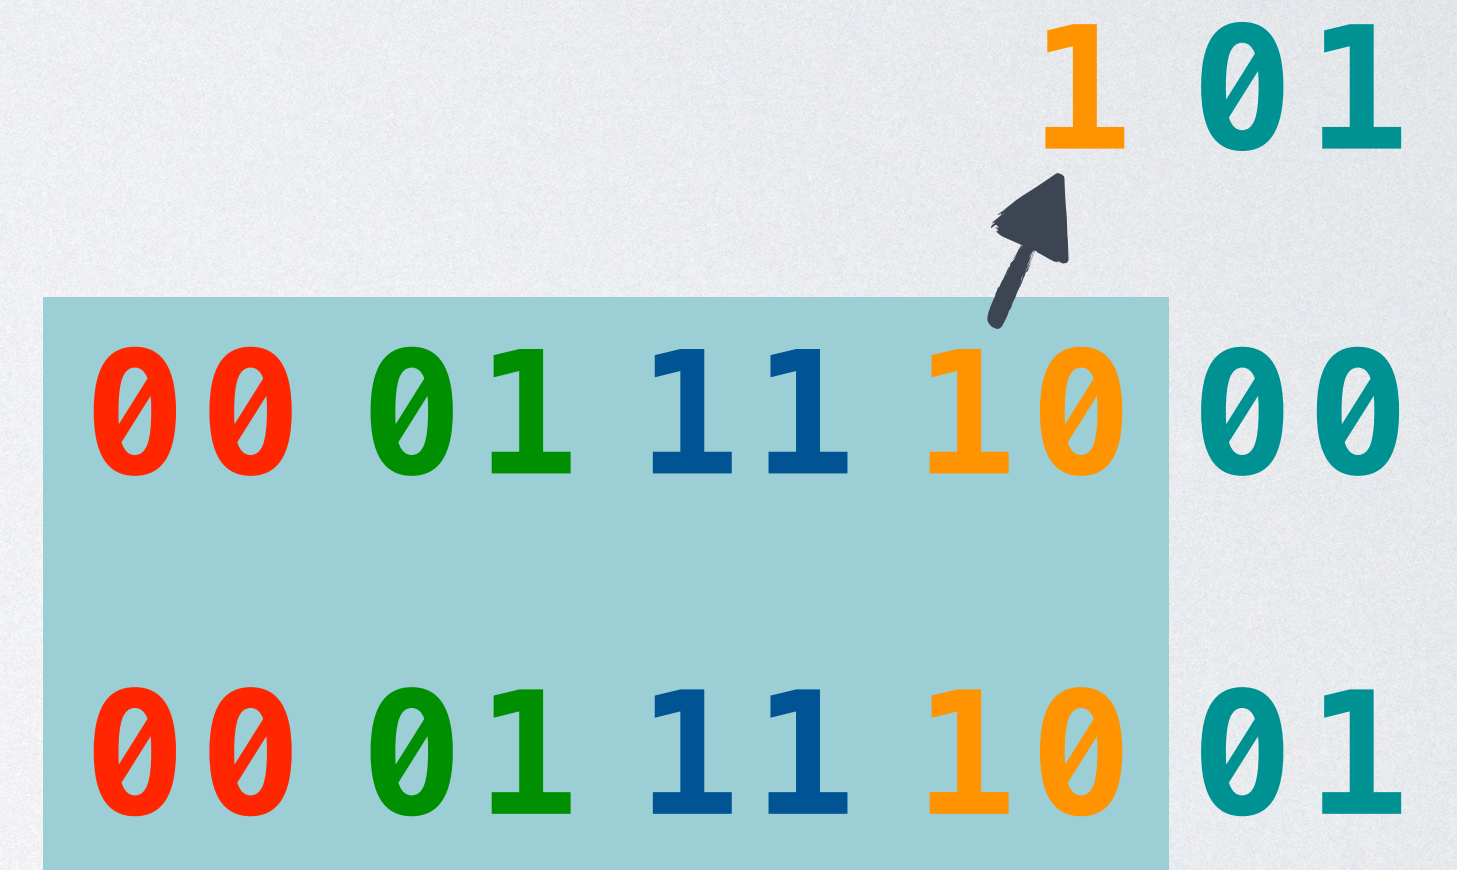
bit vectors

00 01 11 11 01

00 01 11 10 00

00 01 11 10 01

bin tree



Storing markings

markings

[0, 1, 3, 3, 1]

[0, 1, 3, 2, 0]

[0, 1, 3, 2, 1]

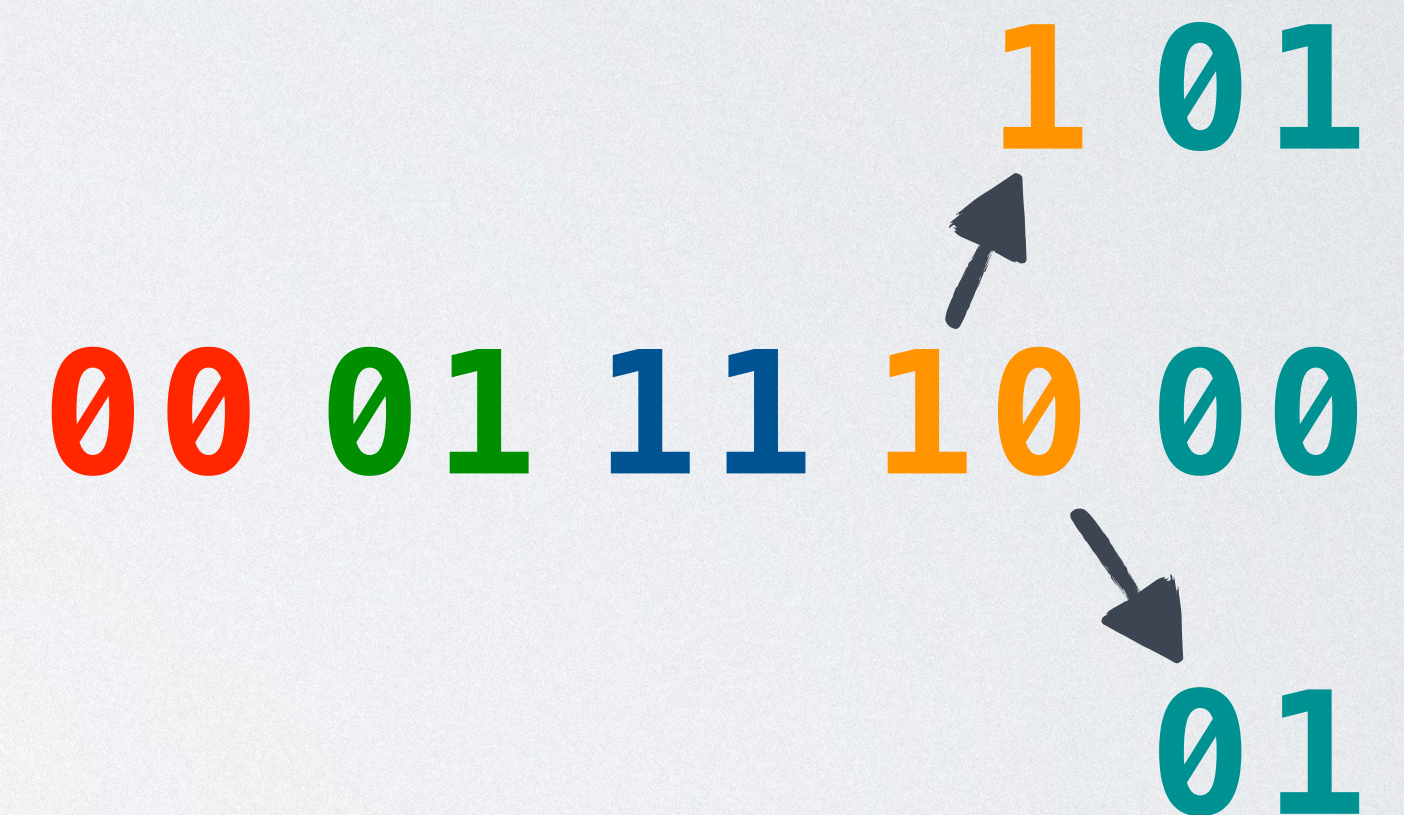
bit vectors

00 01 11 11 01

00 01 11 10 00

00 01 11 10 01

bin tree



Know what you need

- sometimes, only a special case is needed:
 - Tarjan algorithm where only TSCCs are needed
 - linear algebra (sparse matrix, only carrier is needed)
- no library can offer this

Optimizations

- profiling, low-level optimizations (caching)
- know your limits: (in LoLA: malloc)

Academic software design

hardly any tenure
programmers
seldom seen as important

coding is never
top priority
your thesis is!

hard to collect/
keep knowledge
people leave frequently and for good

definitely no coding
professionals
university cannot teach experience

maintenance is
not enforced
once the paper is out, nobody cares

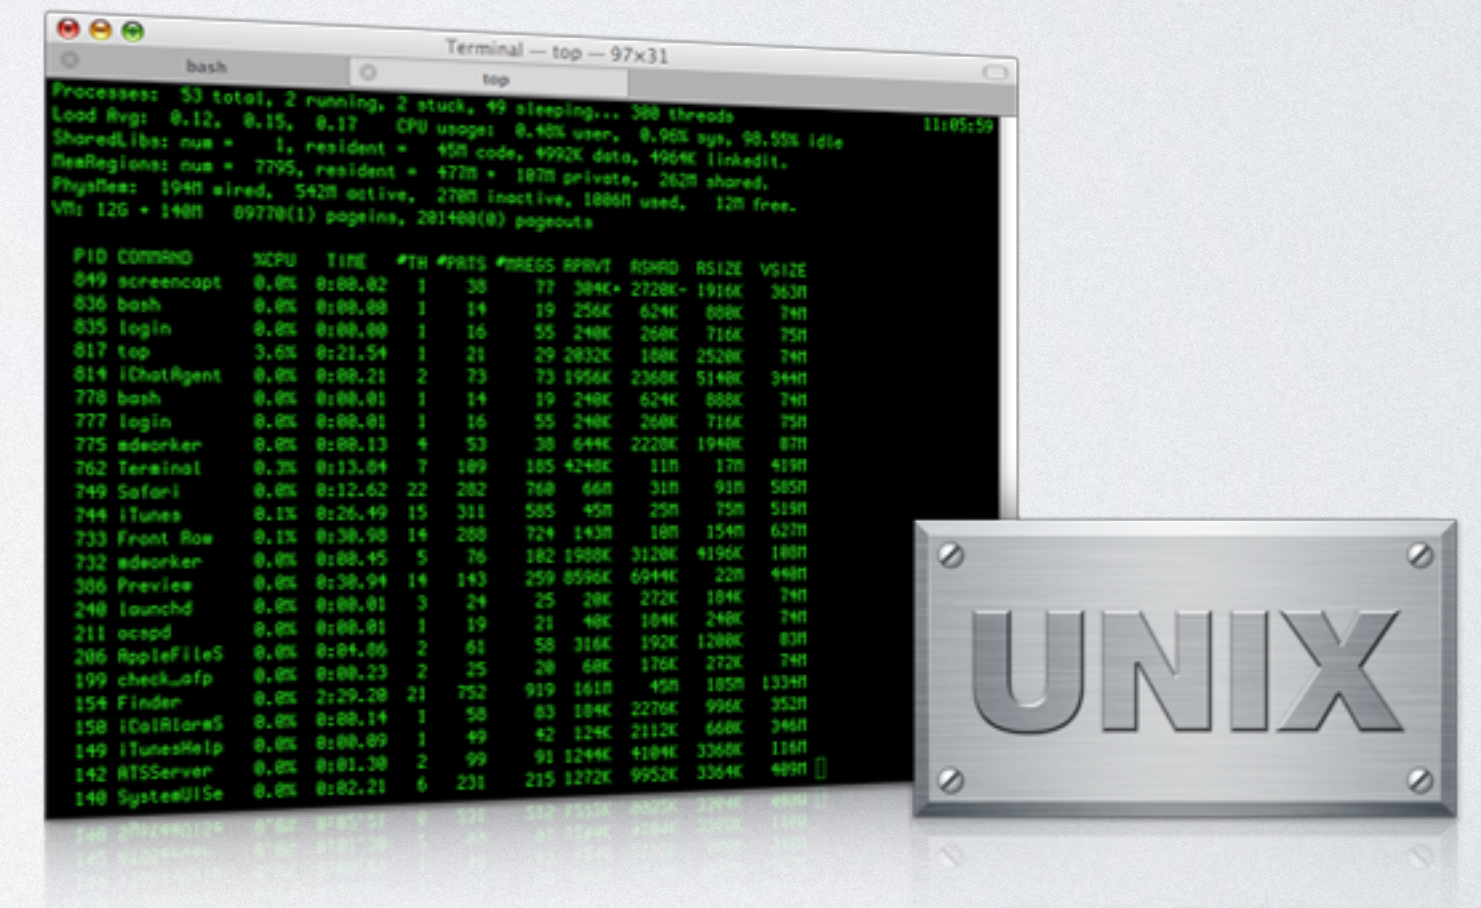
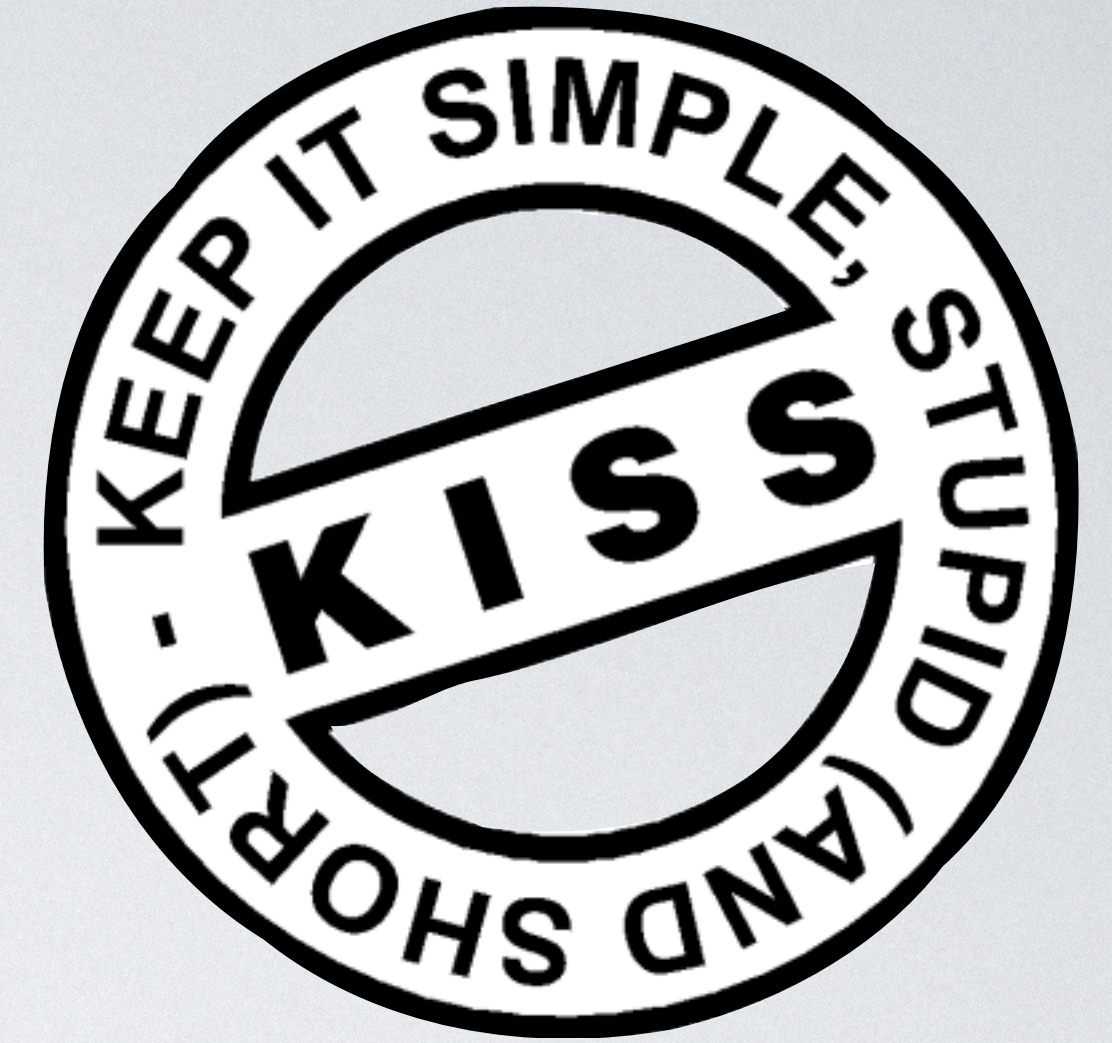
there is no
(paying) customer
no agreed feature set

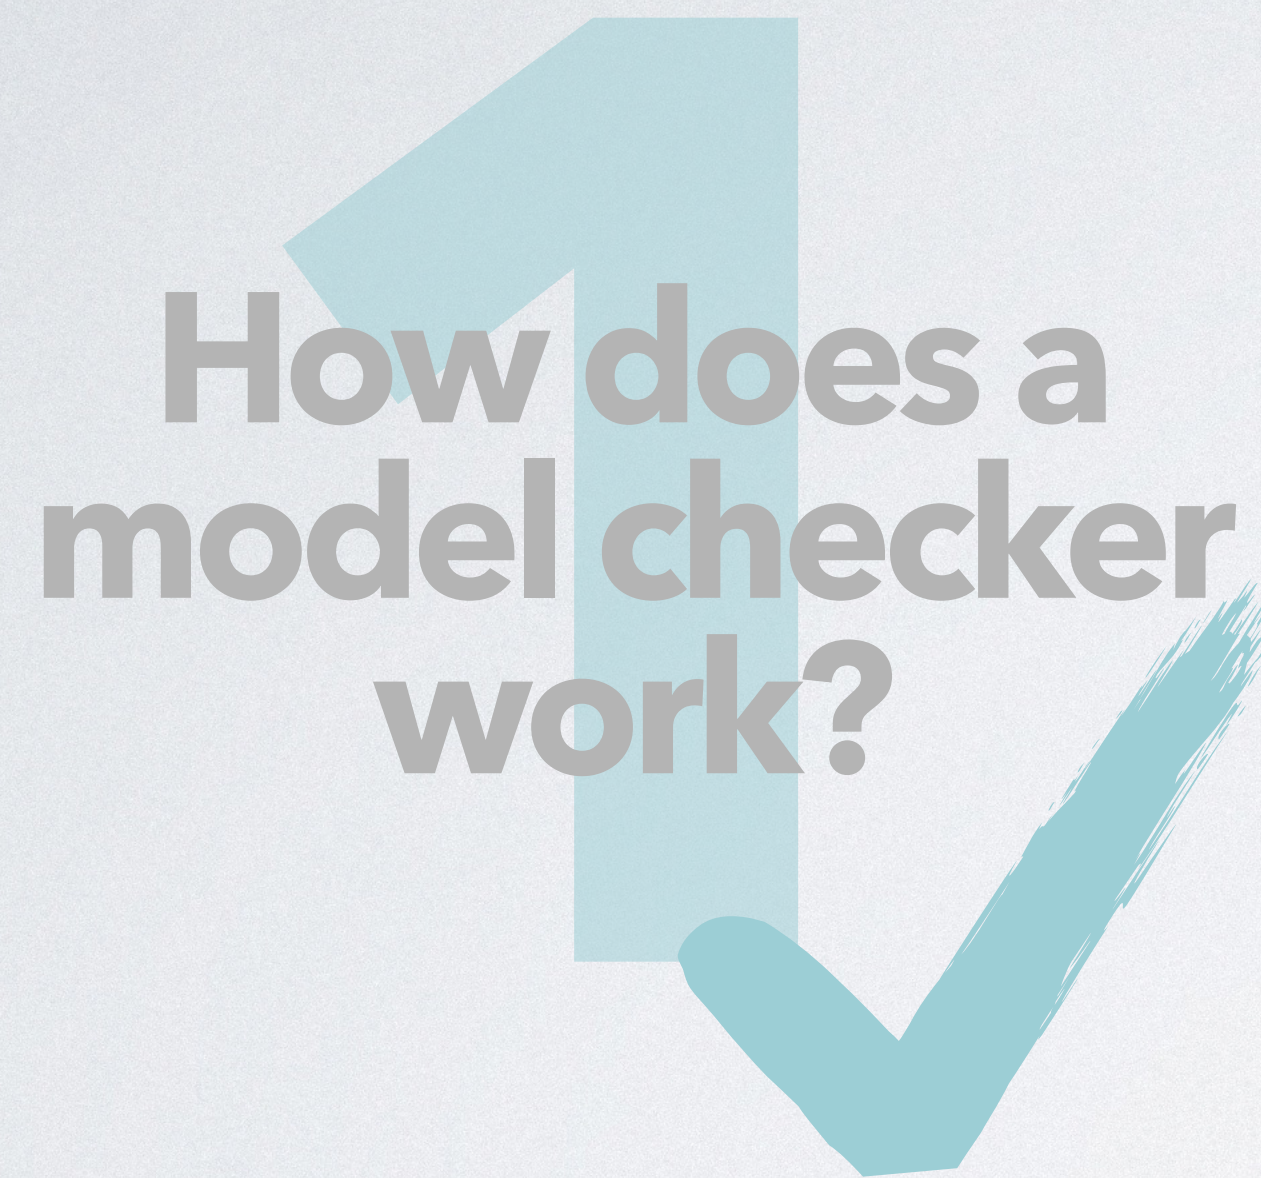
very fast
scientific progress
moving targets

frequently
changing staff
2-5 year contracts

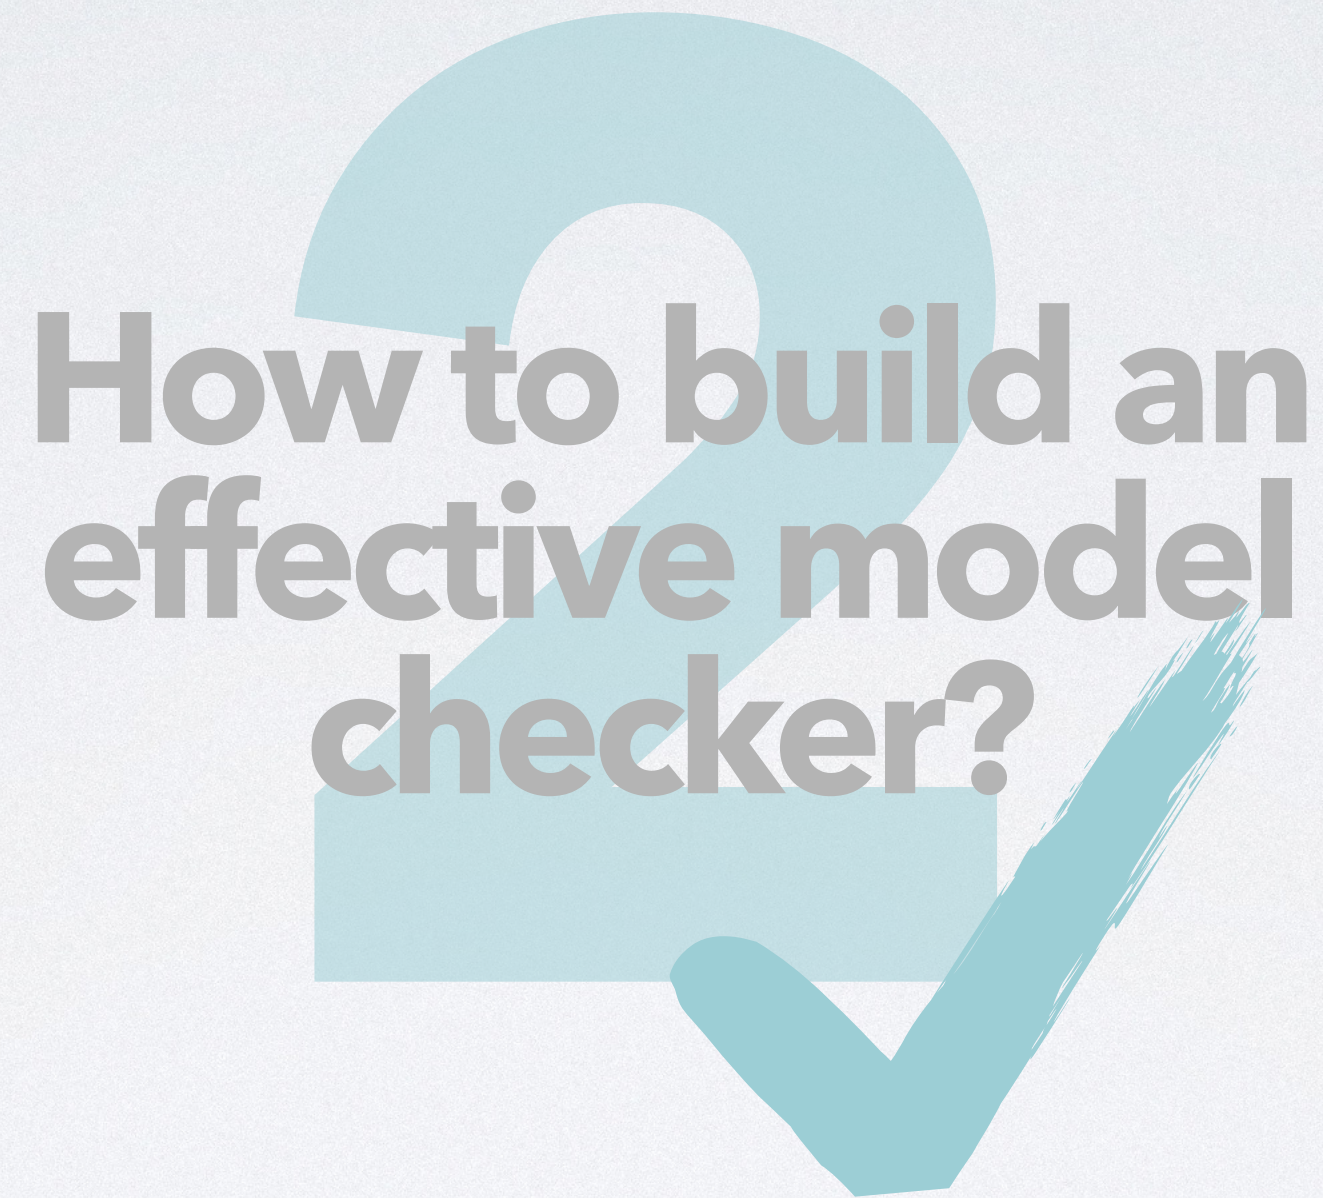
Lessons learned

- prototyping and check on real data
- KISS; few dependencies
- split large tools to smaller “brain-sized” units
- test coverage to avoid the fear of breaking everything
- goal orientation:
no UI, integration via streams

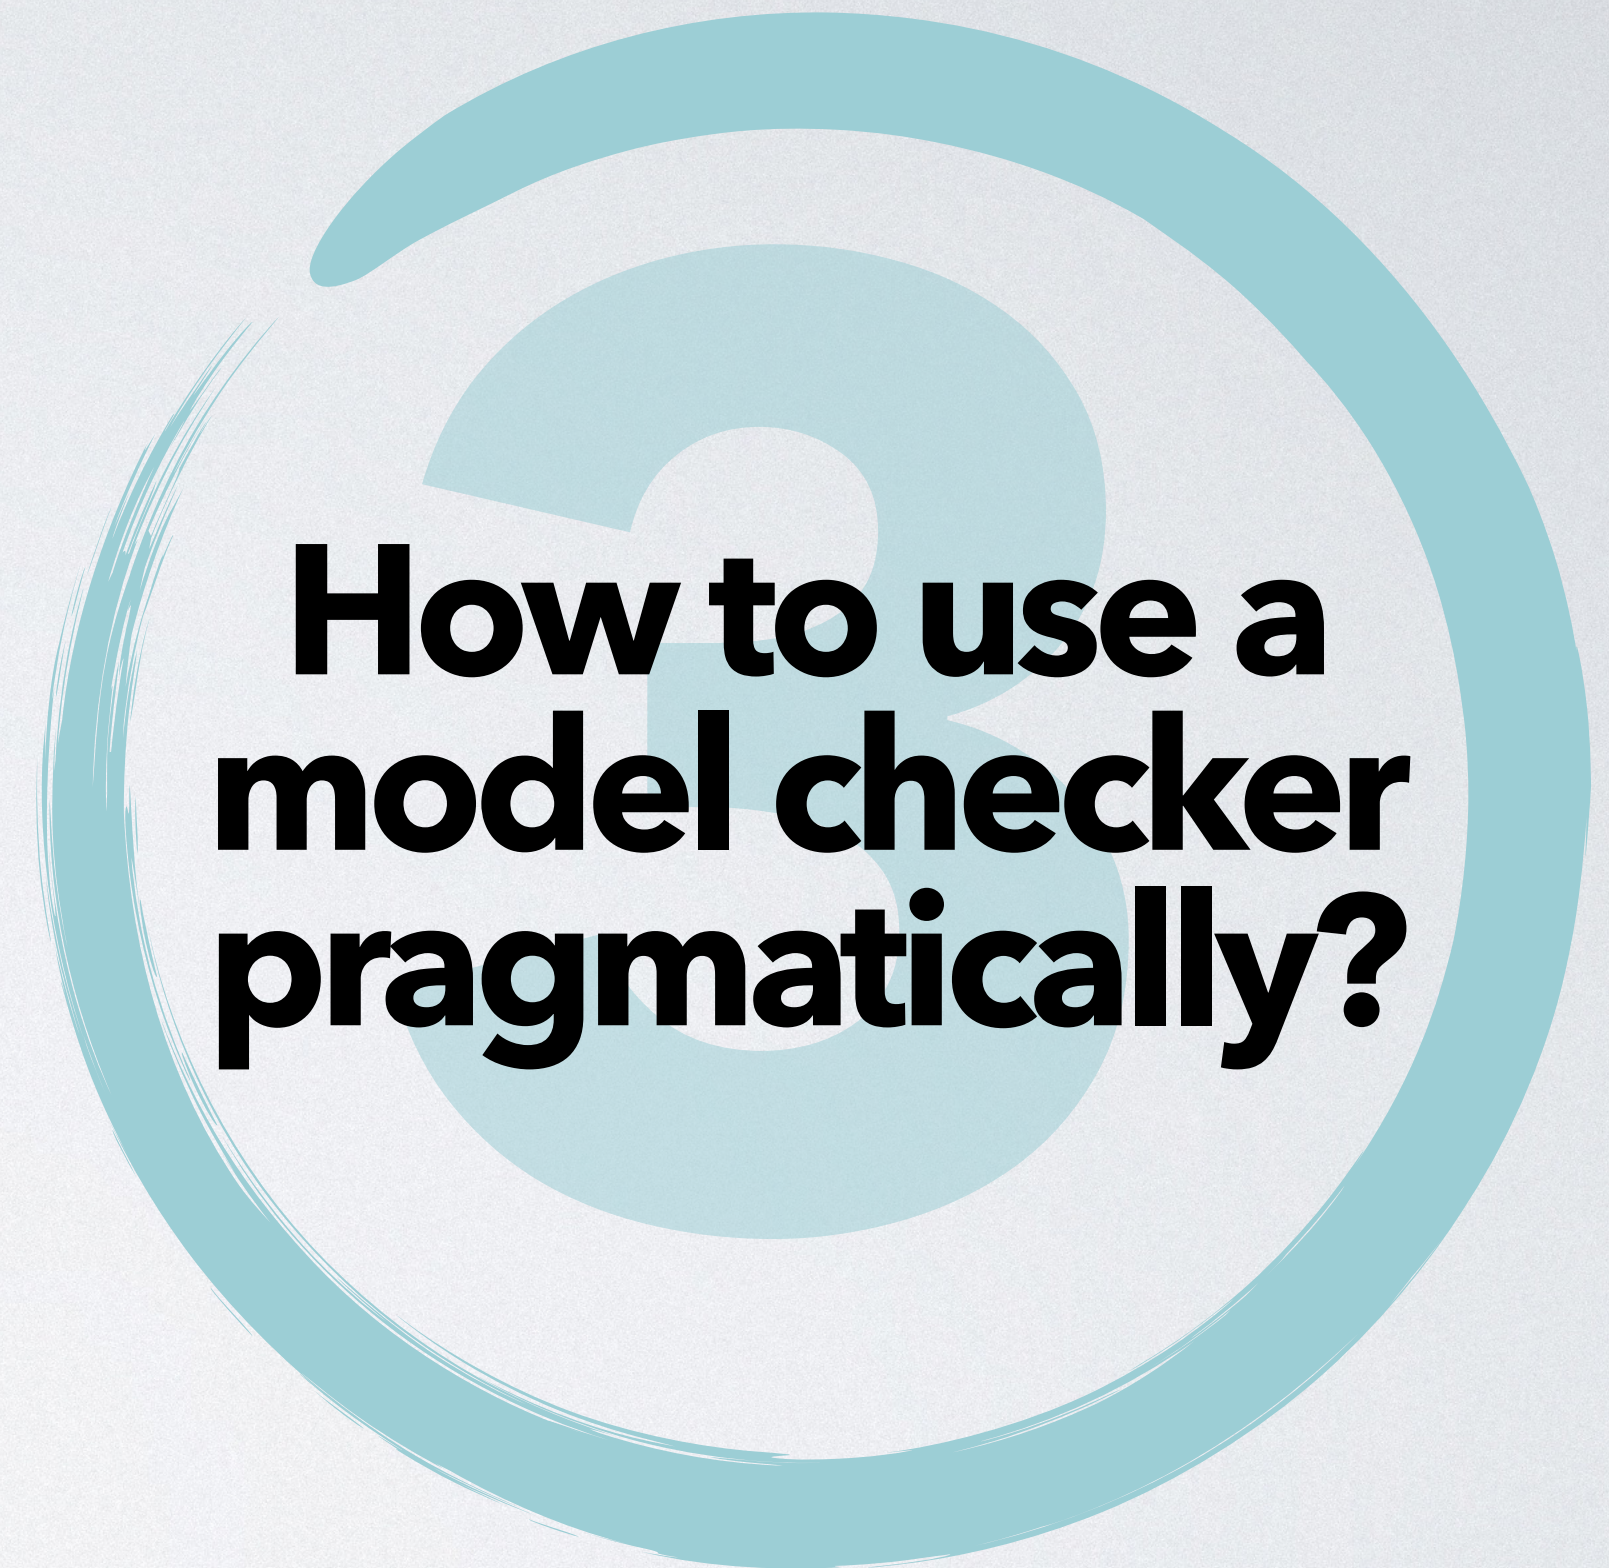




How does a
model checker
work?



How to build an
effective model
checker?



**How to use a
model checker
pragmatically?**

Verification questions: the don'ts

Don't ask for global states.

Usually, only a **few aspects** (marking of a few places) are relevant.

Don't ask two things at once.

If the properties affect different parts of the model, ask **separately**.

Don't use the X operator.

In **distributed systems**, it makes no sense to ask for "the" next state.

Don't order unordered things.

Usually, **all components** should be correct, not just component #1.

Verification questions: the dos

Ask simple questions.

Sometimes, you don't need **temporal logics** at all.

Make a verification model.

Manipulating the original model may help to ask simpler questions.

Use domain knowledge.

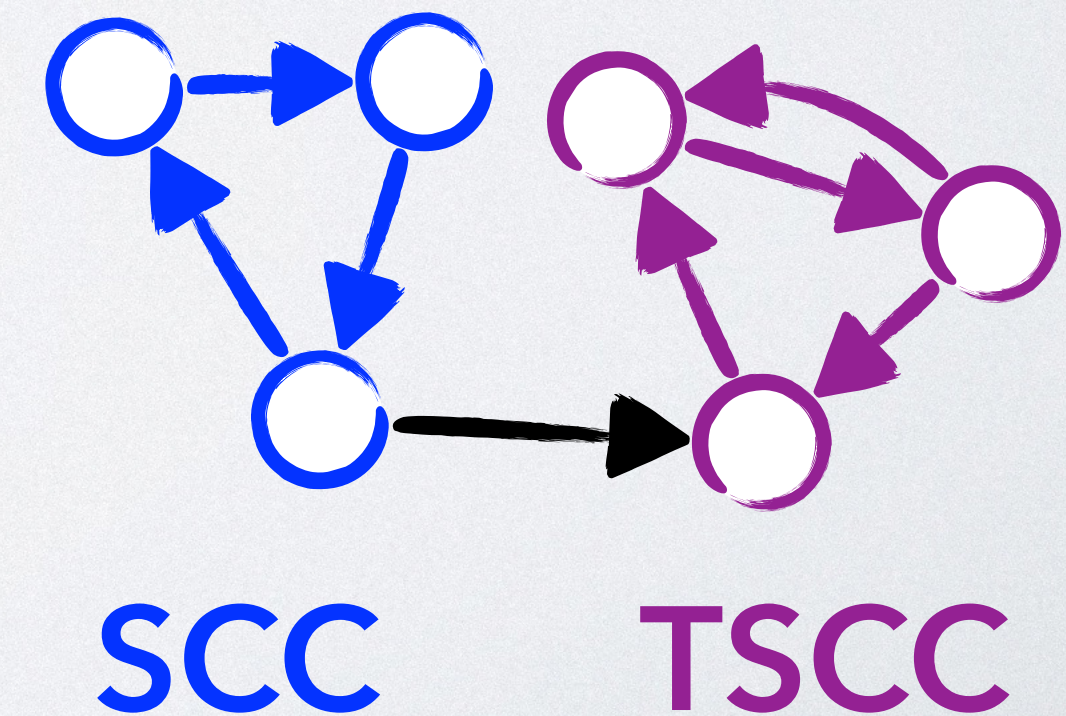
Exploit **implicit assumptions** about the model and the property.

Simple properties

complex	simple	reason
EF φ	reachability of φ	trivial check function
AG φ	\neg EF $\neg\varphi$	again reachability
AGEF φ	check for φ in all TSCC	reachability + TSCC detection

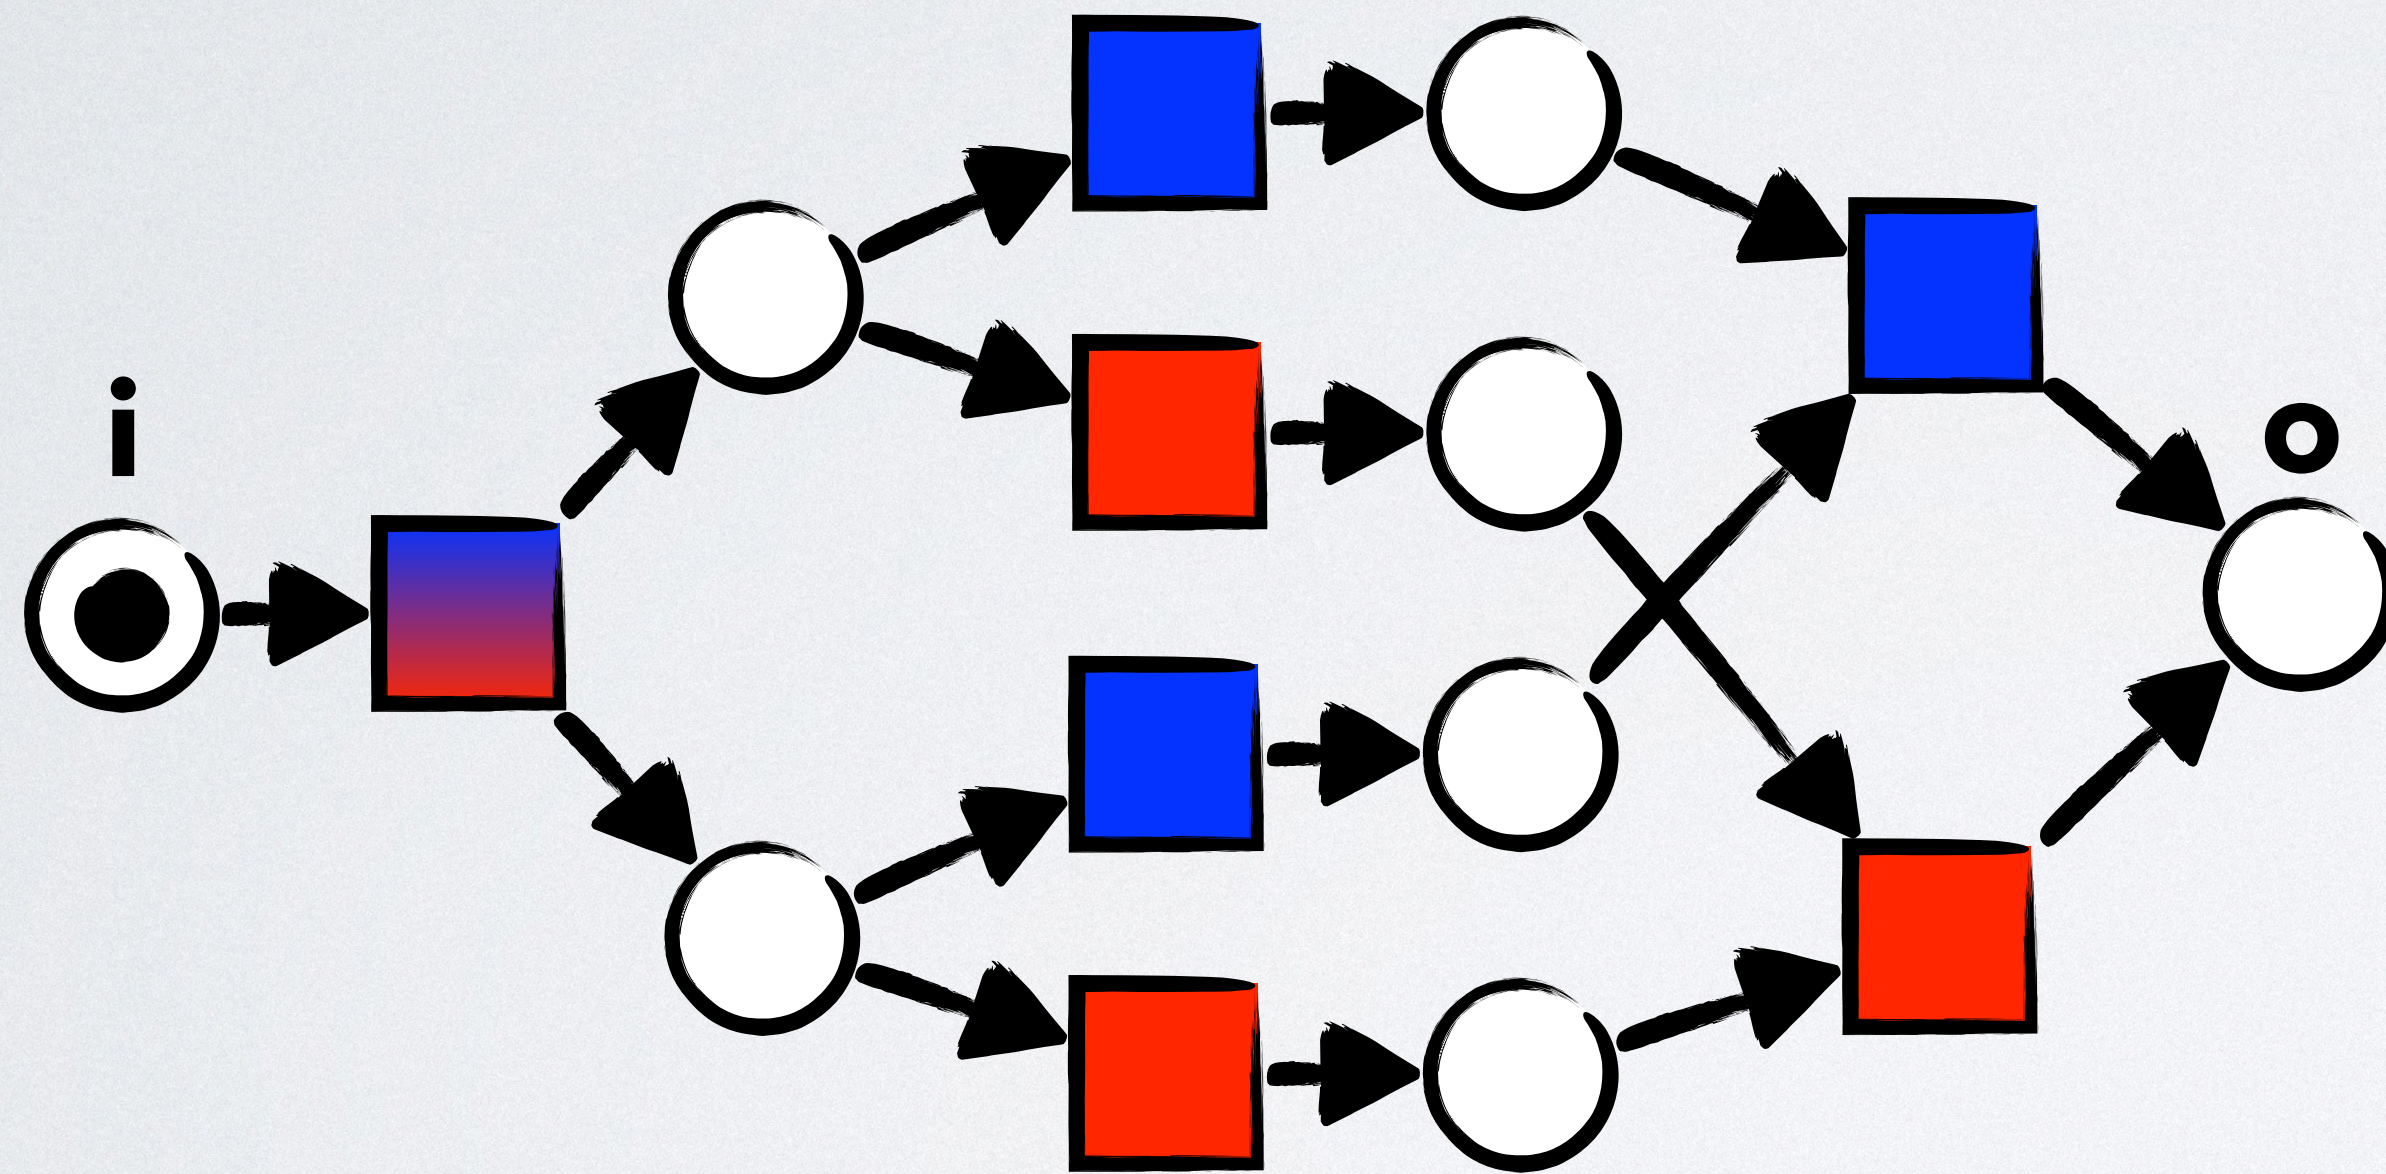
lesson: (hardly) no need
for complex CTL* formulae

compilers will help to find best match



Checking relaxed soundness

Definition: A workflow net is *relaxed sound* iff for all transitions exists a terminating firing sequence.



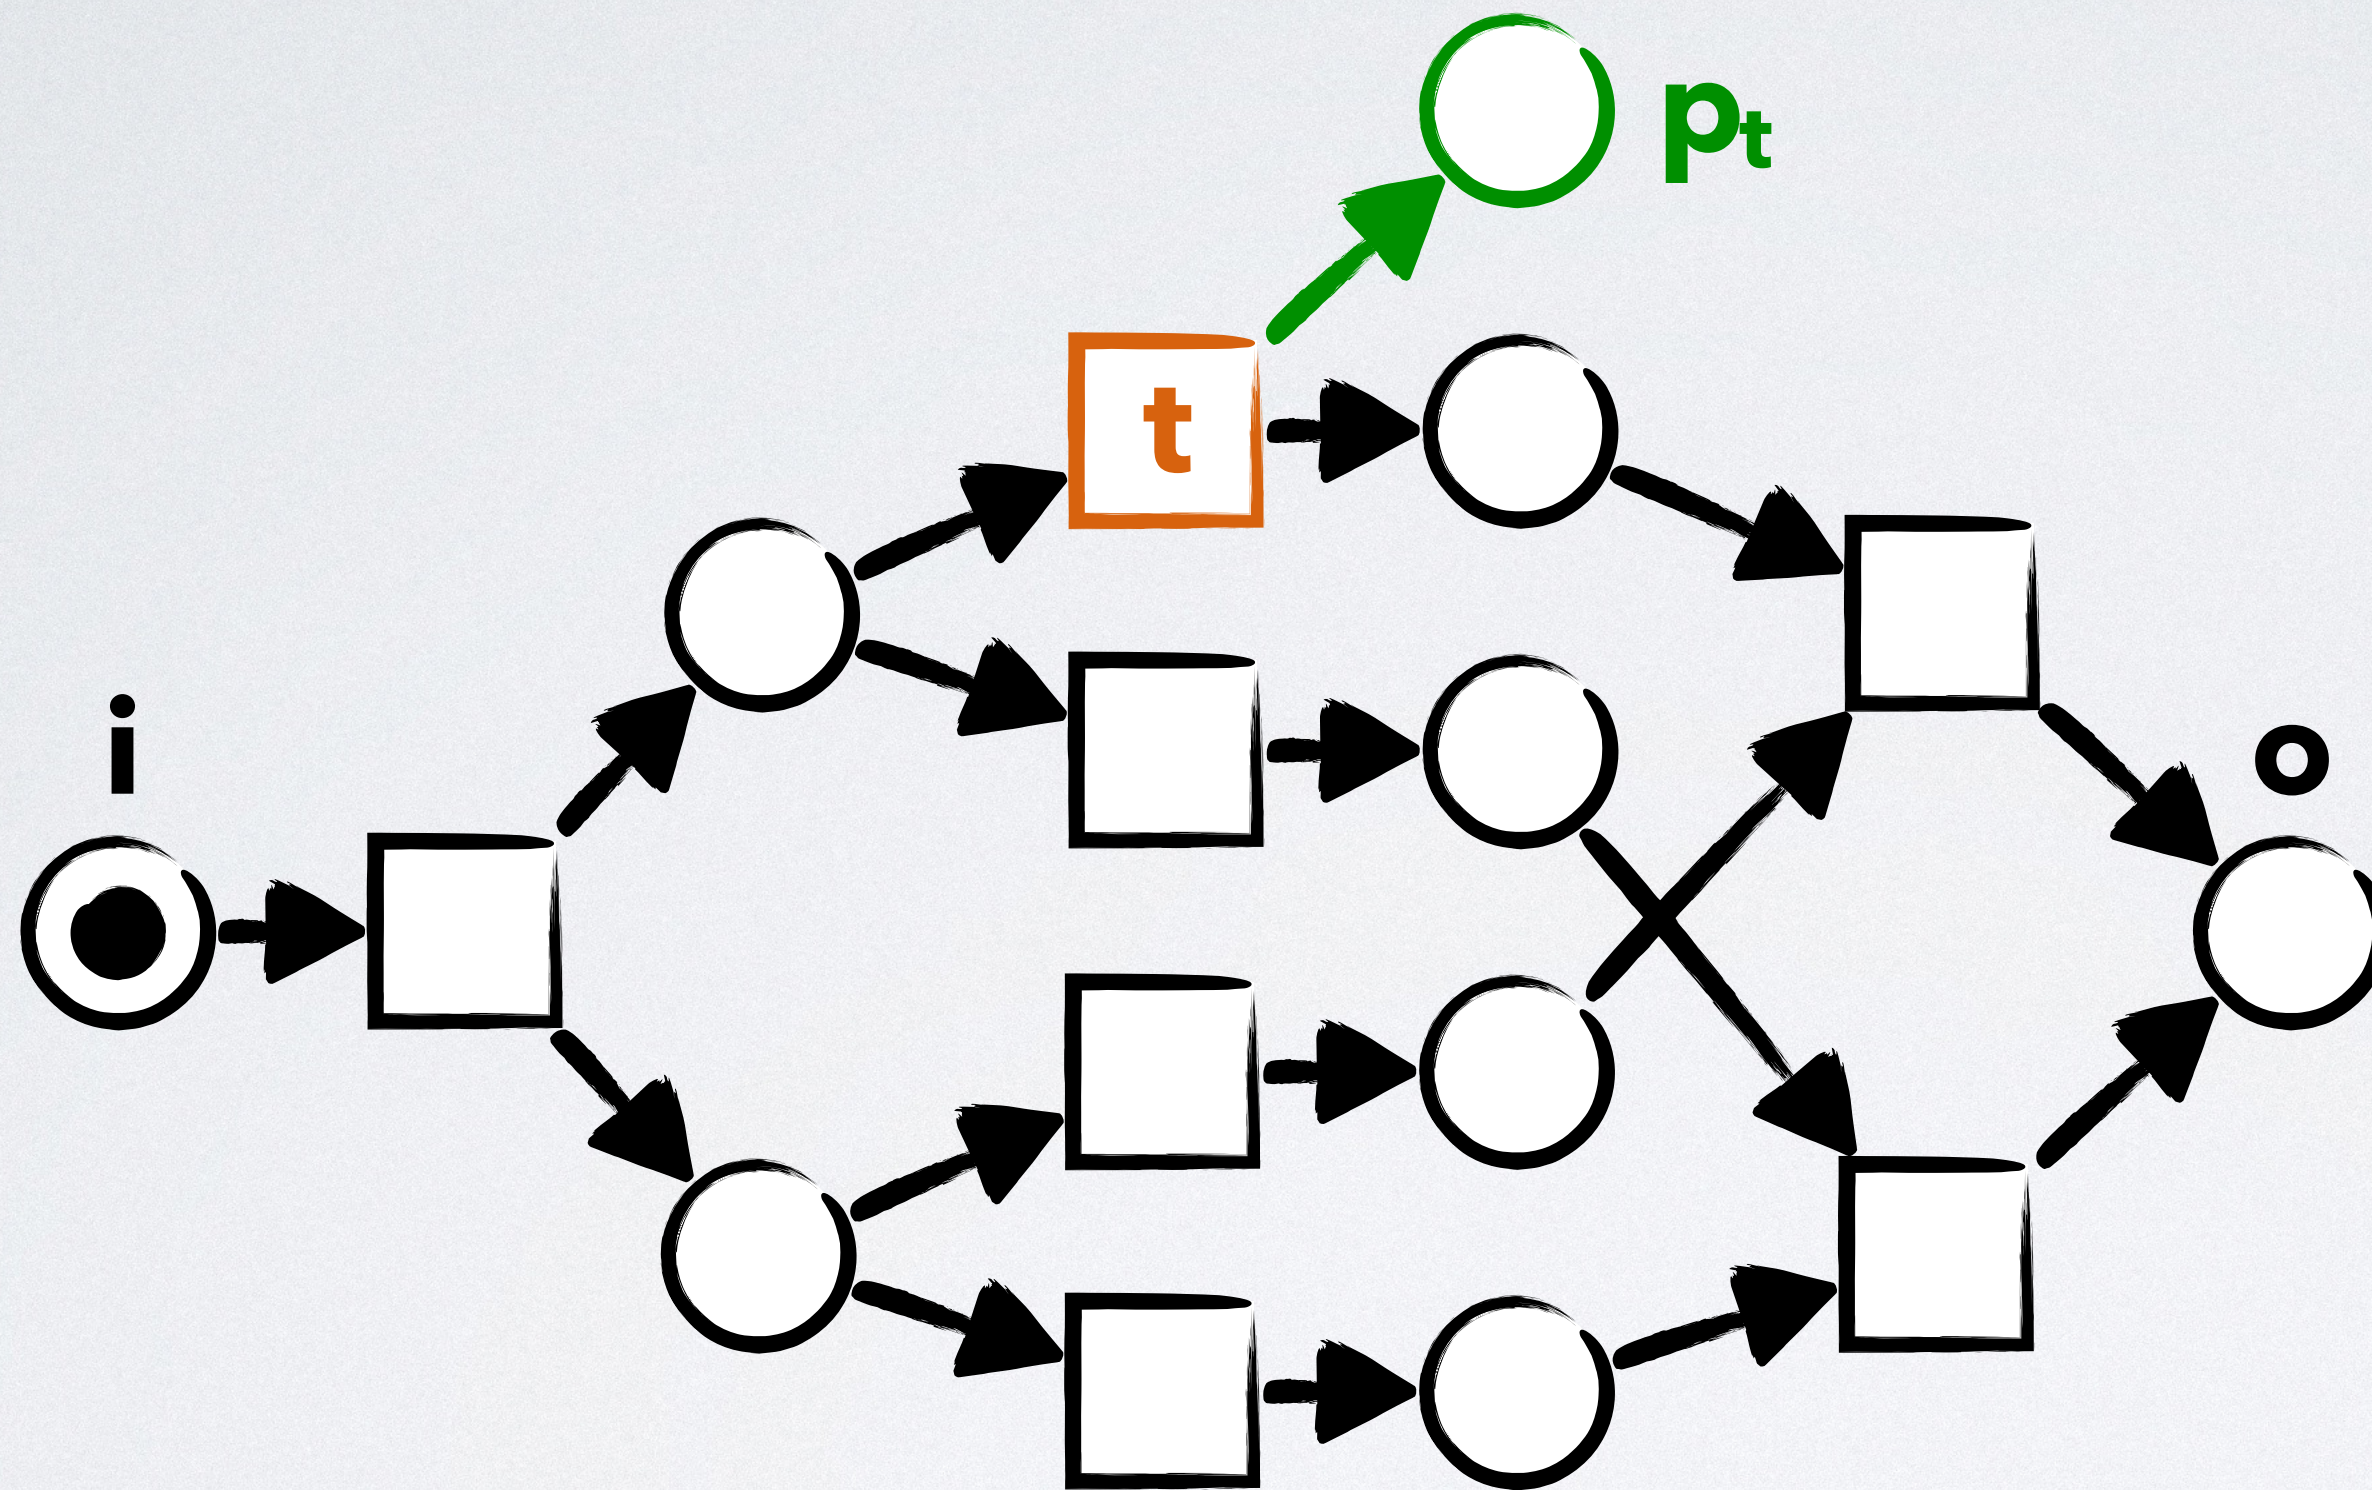
net is relaxed sound

if choices are not
synchronized, the net
deadlocks

but every transition
can fire in a
terminating sequence

Net manipulation: algorithm

1. for all transitions **t**: create a net with test place **p_t**:



2. check for **EF** (**p_t** > 0 \wedge o = 1)

Net manipulation: summary

- check 8 nets instead of 1
- + checking reachability is simpler than extended CTL
- + each state space is smaller than the original
- + one counterexample for each failure
- + parallelizable

Checking soundness

Definition: A workflow net is *sound* iff (1) the final marking $[o]$ is always reachable (2) $[o]$ is the only marking with tokens on place o , and (3) no transition is dead.

+ **domain knowledge:** nets are free-choice

Definition: A free-choice workflow net is *sound* iff (1) AGEF $[o]$ holds, (2) the net is 1-safe, and (3) no transition is dead.

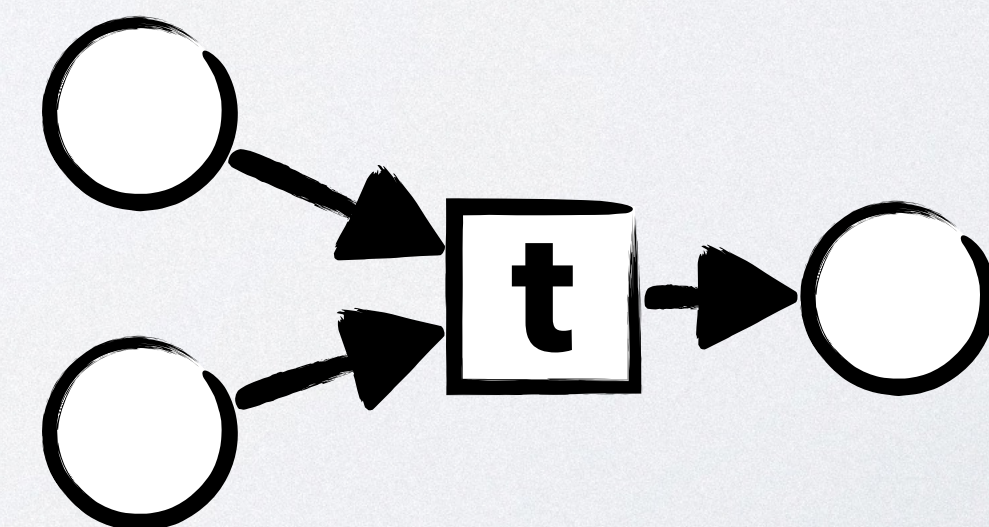
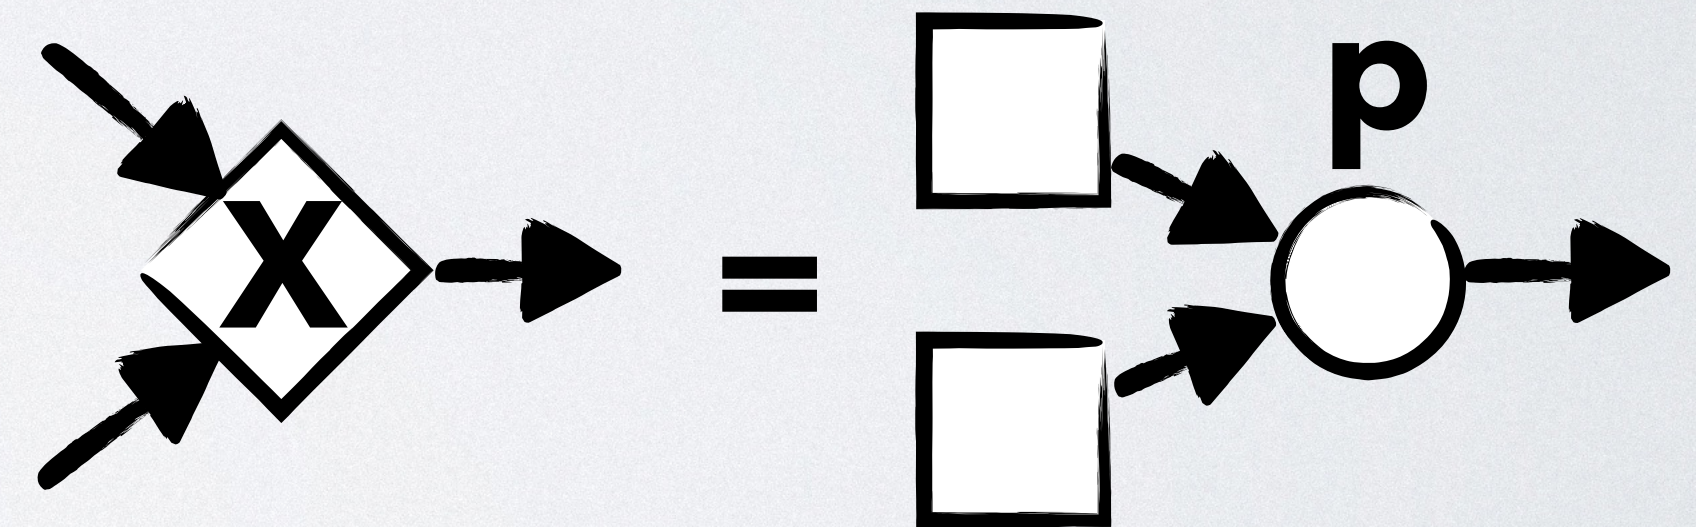
Checking soundness

Definition: A free-choice workflow net is *sound* iff (1) AGEF $[o]$ holds, (2) the net is 1-safe, and (3) no transition is dead.

(1) check if the marking $[o]$ is reachable in all TSCCs.

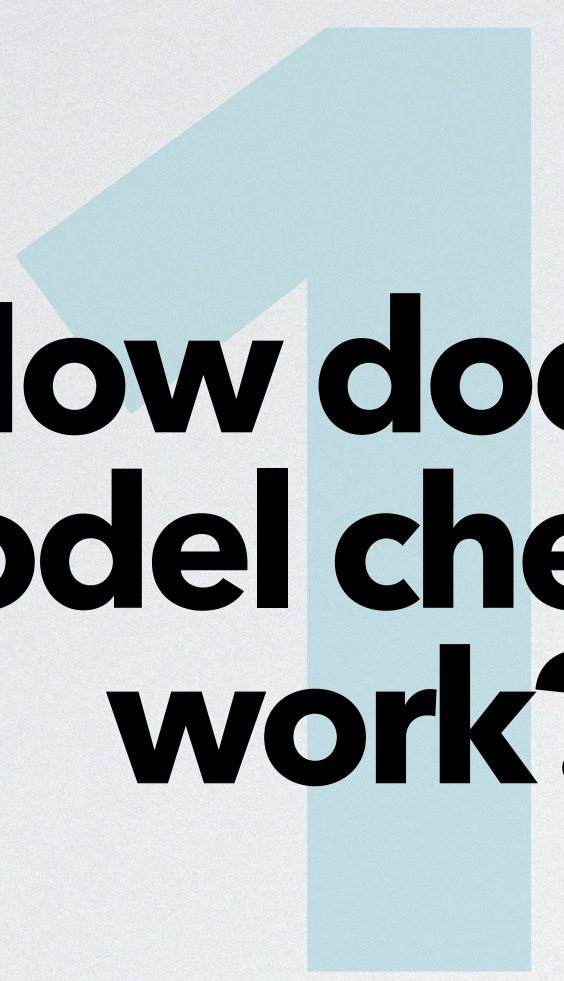
(2) for every join-place p , check if m with $m(p) > 1$ is reachable.

(3) for every transition t , check if $\bullet t$ is reachable



Pragmatic use: summary

- help the model checker help you
- reformulate your question
- many small state spaces are better than one large



**How does a
model checker
work?**



**How to build an
effective model
checker?**



**How to use a
model checker
pragmatically?**



Conclusions

Take home points

- model checking = theory + practice + **pragmatism**
- **academic software design** is a discipline on its own
- asking the **right question(s)** is crucial

Pragmatic model checking: from theory to implementations

Niels Lohmann

Universität
Rostock



Traditio et Innovatio

Copyrights

Atombombentest Romeo, public domain United States Department of Energy
http://commons.wikimedia.org/wiki/File:Castle_Romeo.jpg

Raspberry Pi, CC-BY-SA Jwrodgers
<http://commons.wikimedia.org/wiki/File:RaspberryPi.jpg>

apoptosis inducing factor, GPL
http://en.wikipedia.org/wiki/File:Apoptosis_inducing_factor.png

Tux, by lewing@isc.tamu.edu
<http://en.wikipedia.org/wiki/File:Tux.png>

Integrated Circuit, public domain
<http://commons.wikimedia.org/wiki/File:Chip.jpg>

Websphere software logo, public domain
http://commons.wikimedia.org/wiki/File:Websphere_logo.png

BPMN logo, copyright Object Management Group (OMG)