



**TECHNISCHE
UNIVERSITÄT
DRESDEN**

Faculty of Computer Science Institute of Systems Architecture, Operating Systems Group

MODELING DISTRIBUTED SYSTEMS

HERMANN HÄRTIG, DISTRIBUTED OPERATING SYSTEMS, SS2022

use models to analyze, prove, predict, ... properties of concrete systems AND

to **establish fundamental insights**

- abstract from details
- concentrate on functionality, properties, ... considered important for a specific system/application/question
- models in engineering disciplines very common, increasingly in CS as well

Purpose

- describe the timing requirements of an application
- describe available resources
- question:
can the application run on/use these resources such that timing requirements are fulfilled

Model elements:

- periodic tasks, deadlines, worst-case exec time, ...

Hopefully RTS class is offered in future (by my successor)

<u>Model</u>	<u>Objective/Question</u>
■ Failure Trees	are all failures and their combinations taken into account
■ statics models	does a house fall down (snow, quake) what kind of vehicles on a bridge
■ control laws	stability of controllers
■ Ohm's Law	behavior of circuits

WELL KNOWN EXAMPLES FOR MODELS



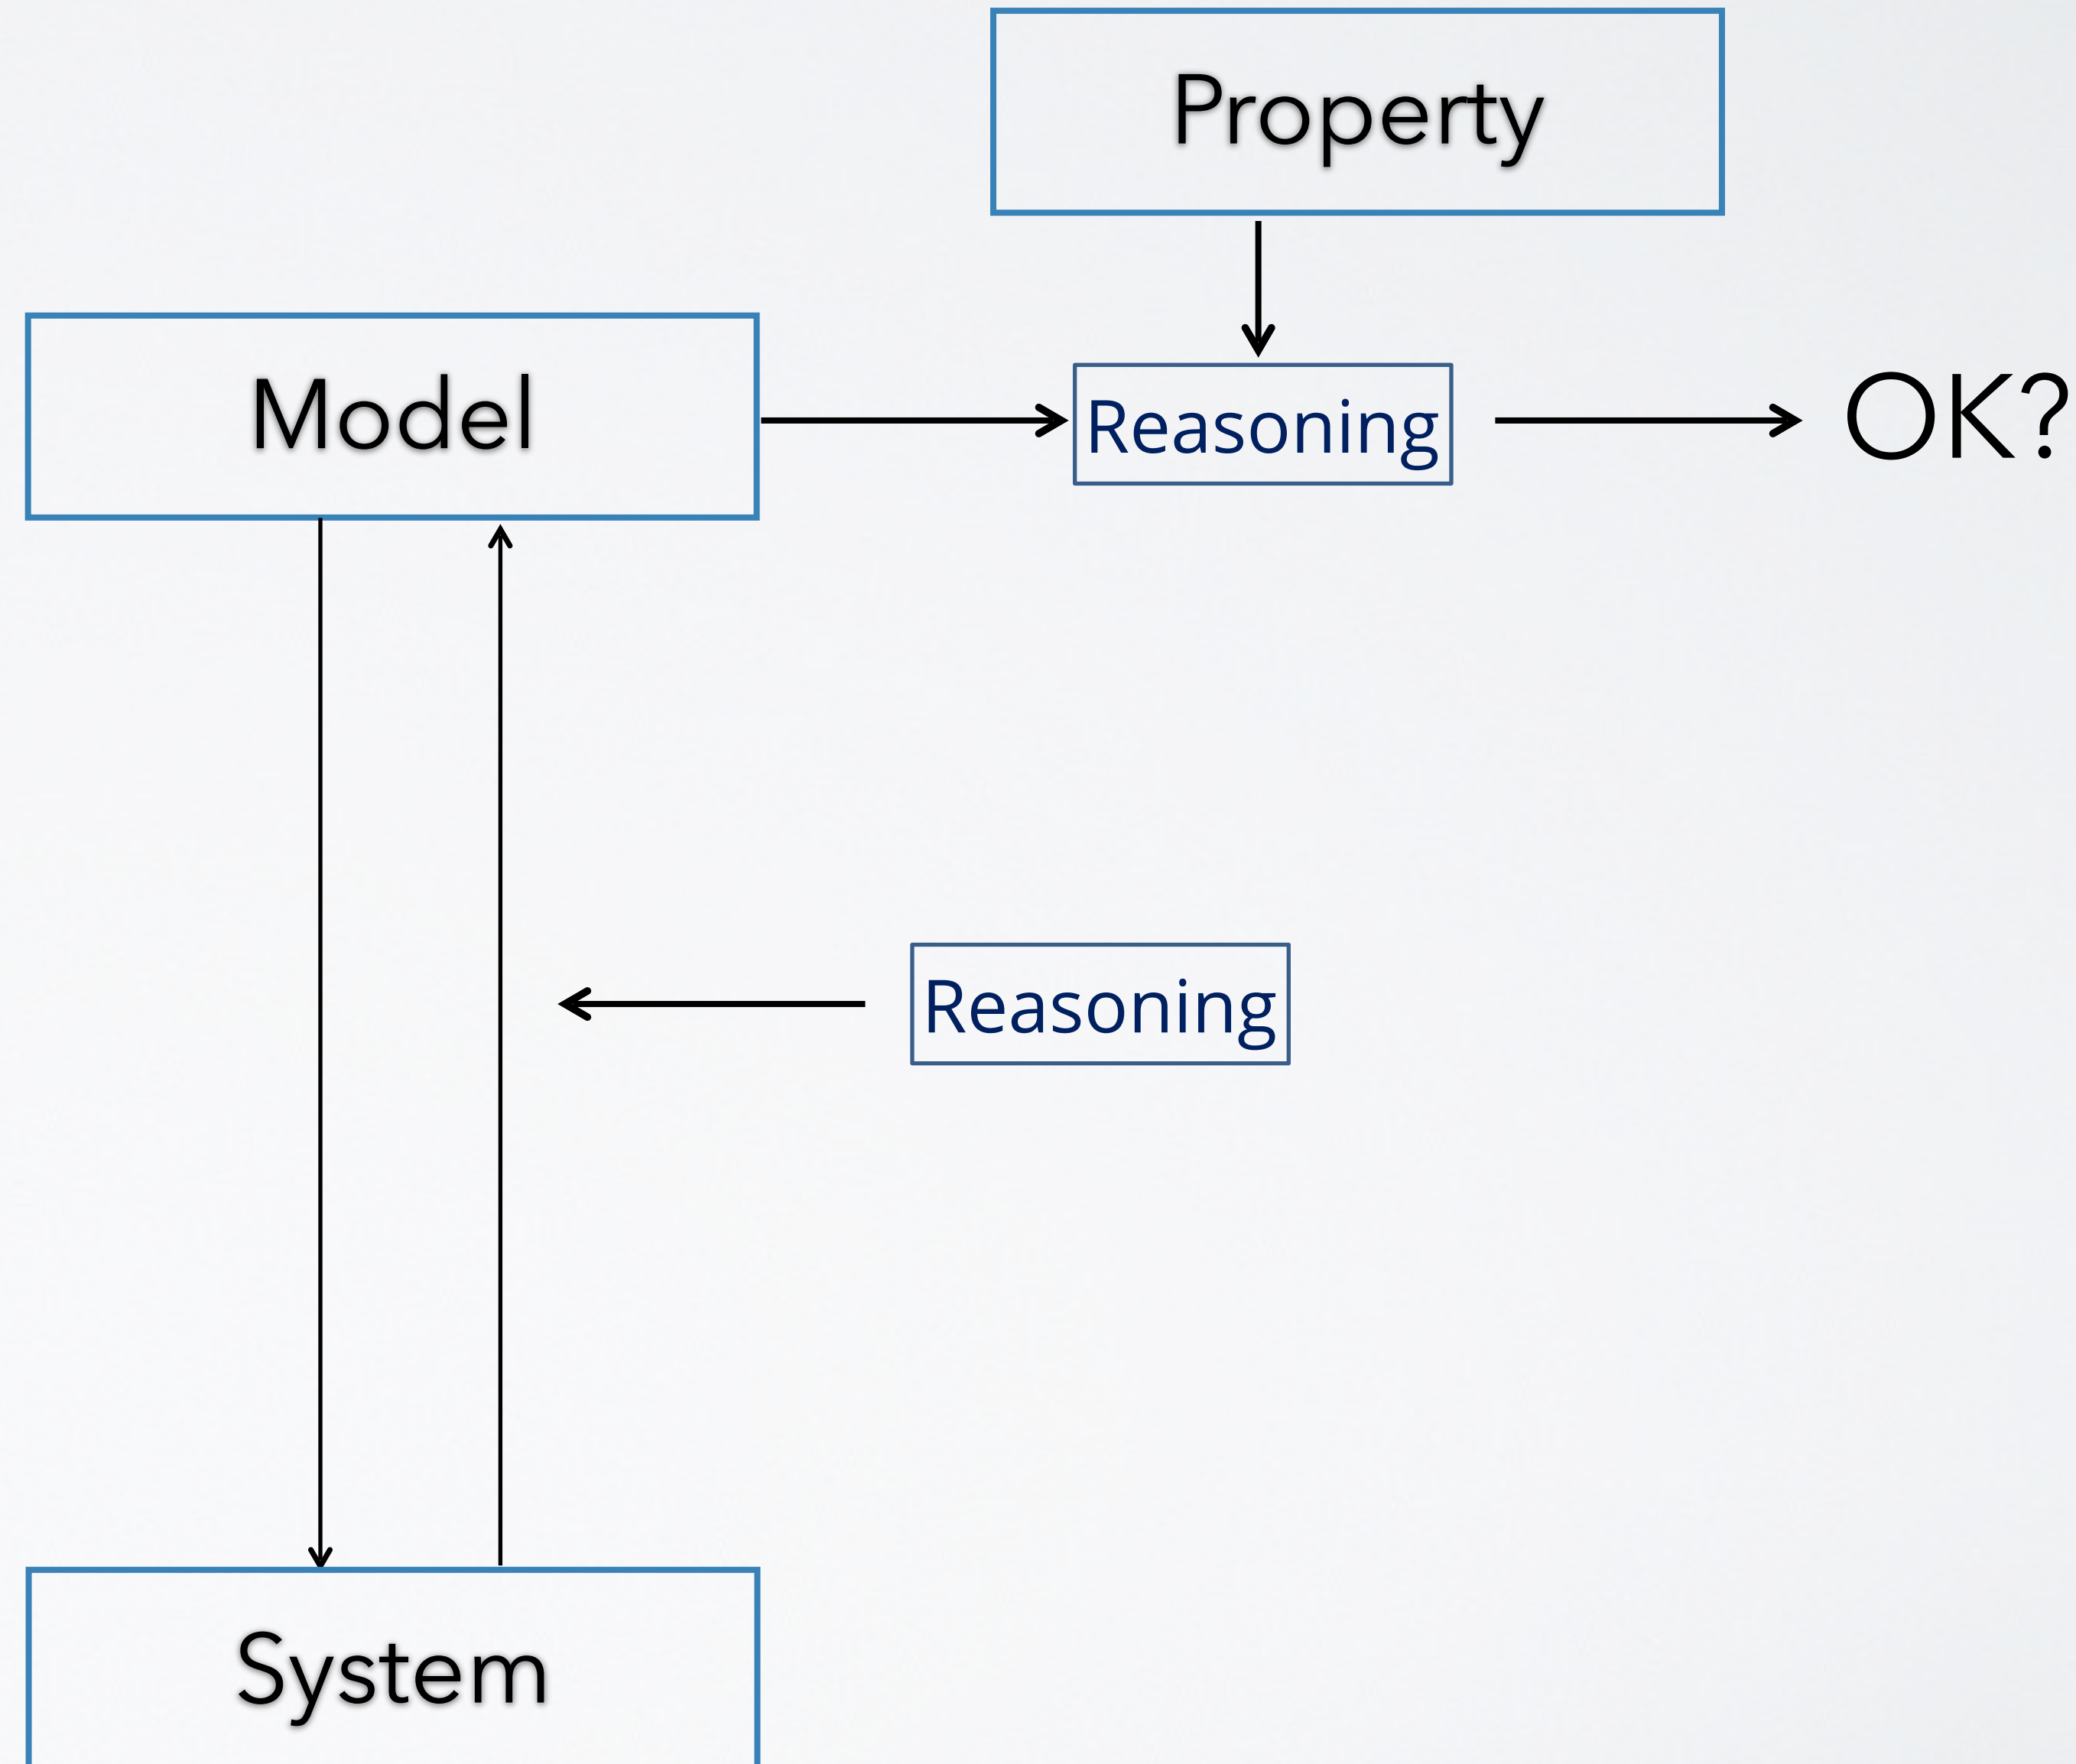
$$I = V/R$$

- Q1: Is it possible to build arbitrarily reliable Systems out of unreliable components?
- Q2: Can we achieve consensus in the presence of faults (consensus: all non-faulty components agree on action)?
- Q3: Is there an algorithm to determine for a system with a given setting of access control permissions, whether or not a Subject A can obtain a right on Object B?

2 Models per Question !

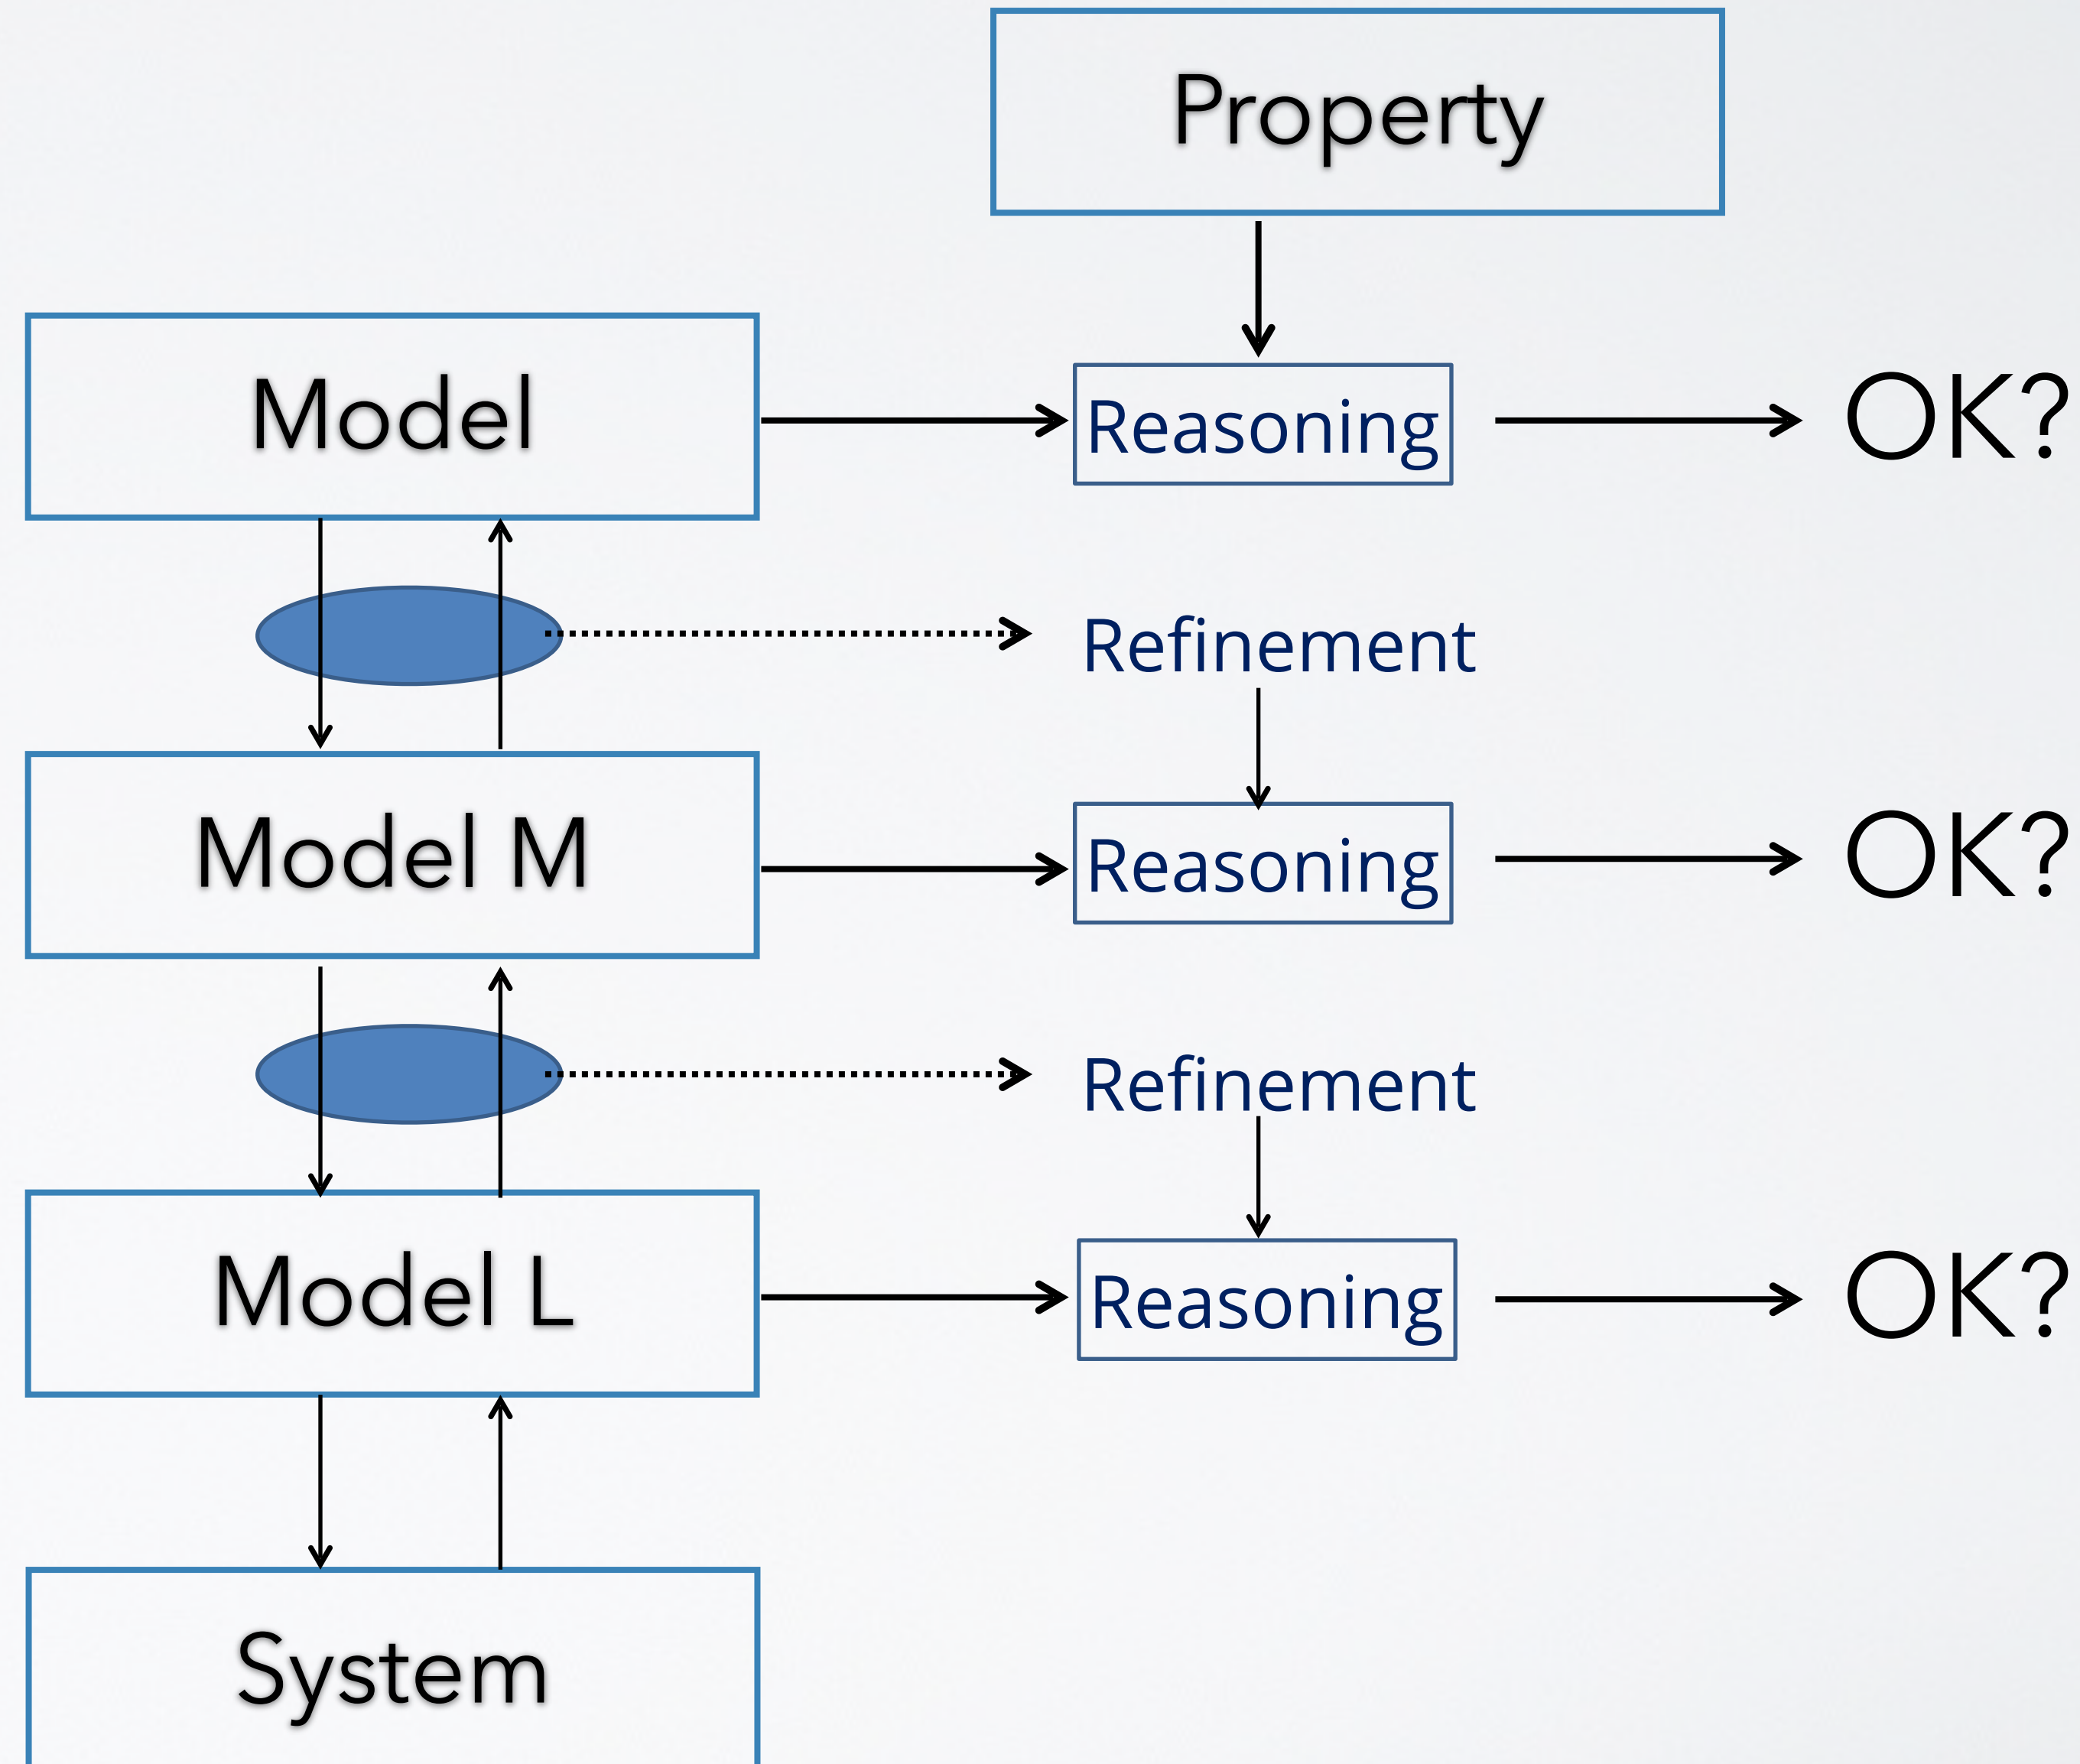
Reasoning:

- Common sense
- Formal Verification
- Careful Inspection
- Mathematics



Reasoning:

- Common sense
- Formal Verification
- Careful Inspection
- Mathematics ...
- **“Refinement”:**
 - Abstraction
 - Implementation
 - Formal Refinement



<u>Model</u>	<u>Objective/Question</u>
■ Amdahl's Law	Scalability
■ Turing Machine	Halting problem, Decidability
■ Logic	Correctness, Precision, ...

Speedup: $\frac{\text{original execution time}}{\text{enhanced execution time}}$

- P: section that can be parallelized
- 1-P: serial section
- N: number of CPUs

$$\text{Speedup}(P,N) = \frac{1}{\left(1 - P + \frac{P}{N}\right)}$$

- if N becomes VERY large, speedup approaches: $1/(1-P)$

- Objective of lecture:
understand the power of models and the need for their careful understanding
- models in detail, but math results by intuition not proofs
- Try to find answers to question Q1 ... Q3
BEFORE viewing the other pieces of the lecture and the full slide set

- Q1: Is it possible to build arbitrarily reliable Systems out of unreliable components?
- Q2: Can we achieve consensus in the presence of faults (consensus: all non-faulty components agree on action)?
- Q3: Is there an algorithm to determine for a system with a given setting of access control permissions, whether or not a Subject A can obtain a right on Object B?

2 Models per Question !

Q1: Can we build arbitrarily reliable Systems out of unreliable components ?

- How to build reliable systems from less reliable components
- Fault(Error, Failure, Fault,)
terminology in this lecture synonymously used for
“something goes wrong”
(more precise definitions and types of faults in SE)

Reliability:


- $R(t)$: probability for a system to survive time t

Availability:

- A : fraction of time a system works

- Fault detection and confinement
 - Recovery
 - Repair
-
- Redundancy
 - Information
 - time
 - structural
 - functional


John v. Neumann
Voter: *single point of failure*



Can we do better
→ distributed solutions?

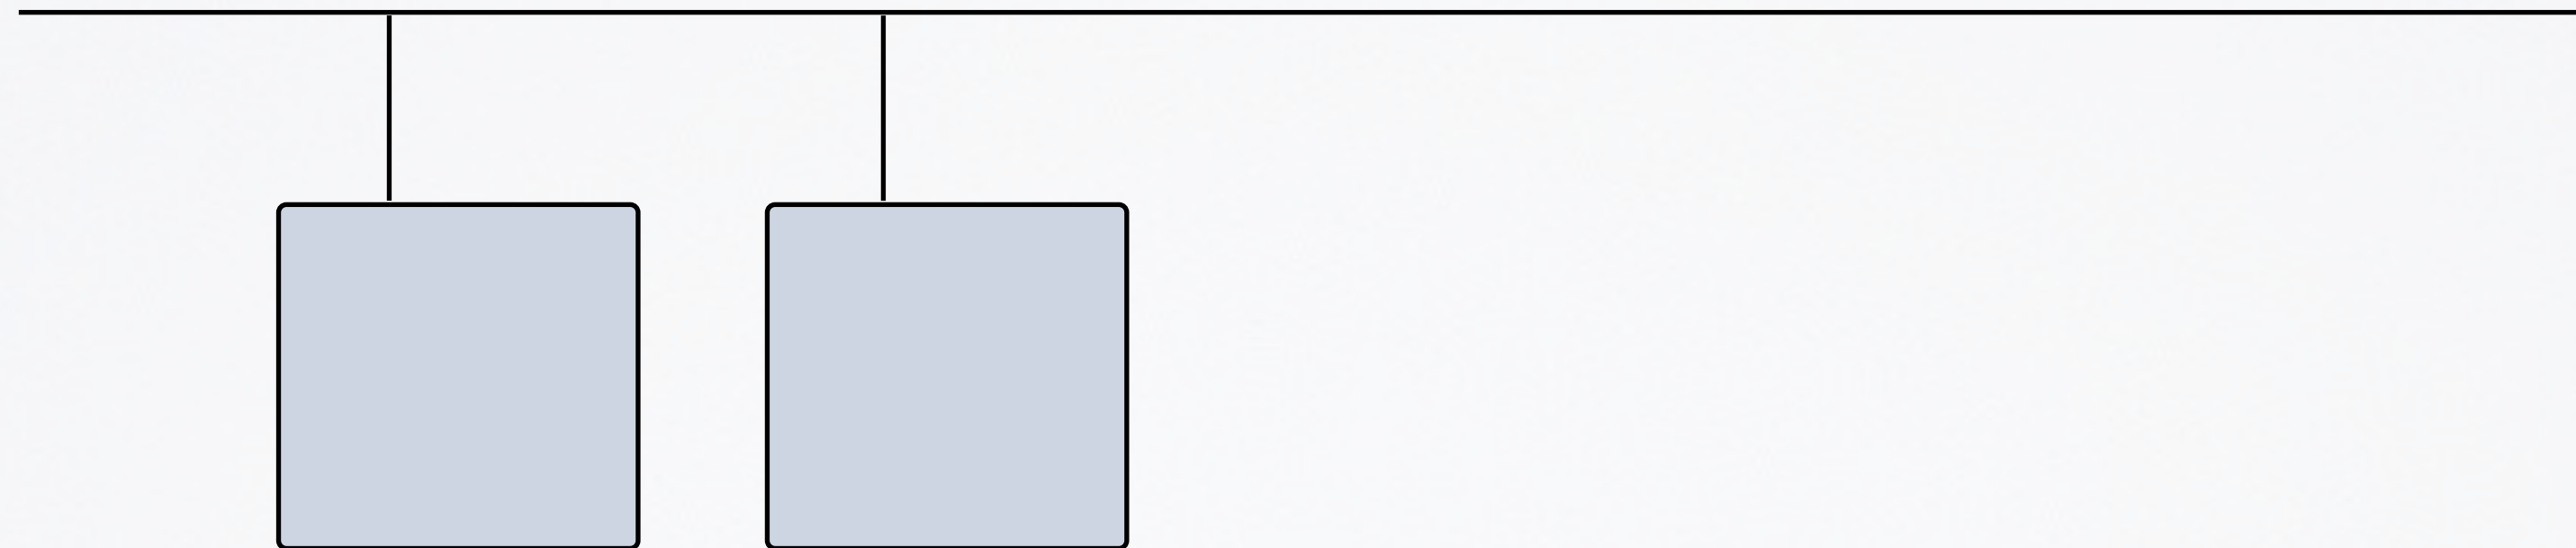
Parallel-Serial-Systems

(Pfitzmann/Härtig 1982)




Parallel-Serial-Systems

(Pfitzmann/Härtig 1982)




Parallel-Serial-Systems

(Pfitzmann/Härtig 1982)




Parallel-Serial-Systems

(Pfitzmann/Härtig 1982)




Serial-Systems



$$R_{whole} = \prod_{j=1}^m R_j$$

Each component must work for the whole system to work.

Parallel-Systems

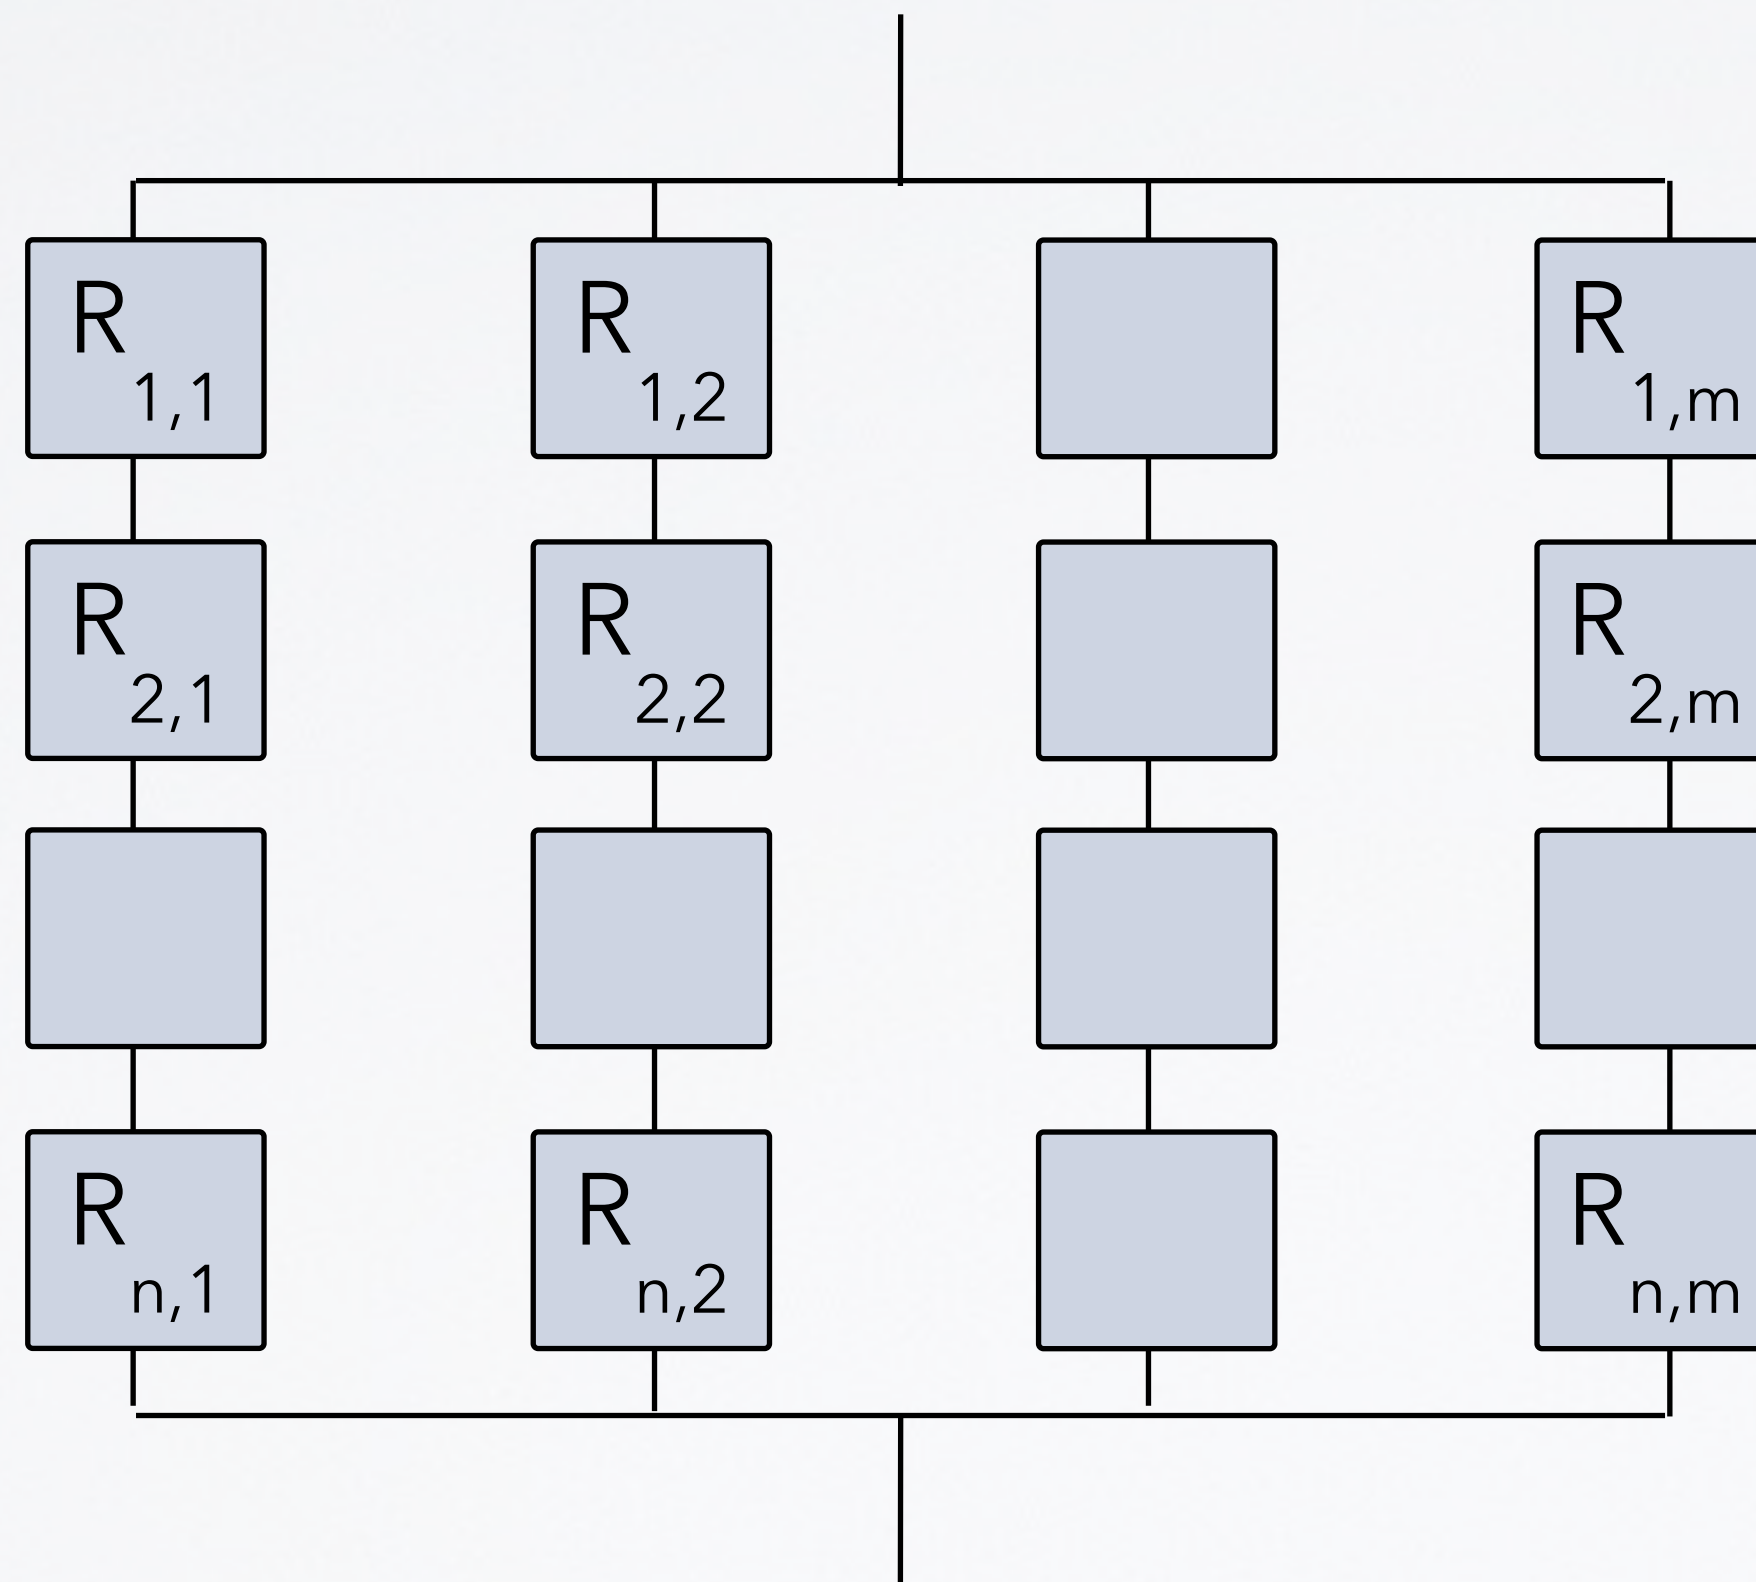


$$R_{whole} = 1 - \prod_{i=1}^m (1 - R_i)$$

One component must work for the whole system to work.

Each component must fail for the whole system to fail.


Serial-Parallel-Systems




$$R_{whole} = 1 - \prod_{j=1}^m \left(1 - \prod_{i=1}^n R_{i,j} \right)$$

Parallel-Serial-Systems

(Pfitzmann/Härtig 1982)





Fault Model

„Computer-Bus-Connector“

can fail such that Computer and/or Bus also fail

=>

conceptual separation of components into

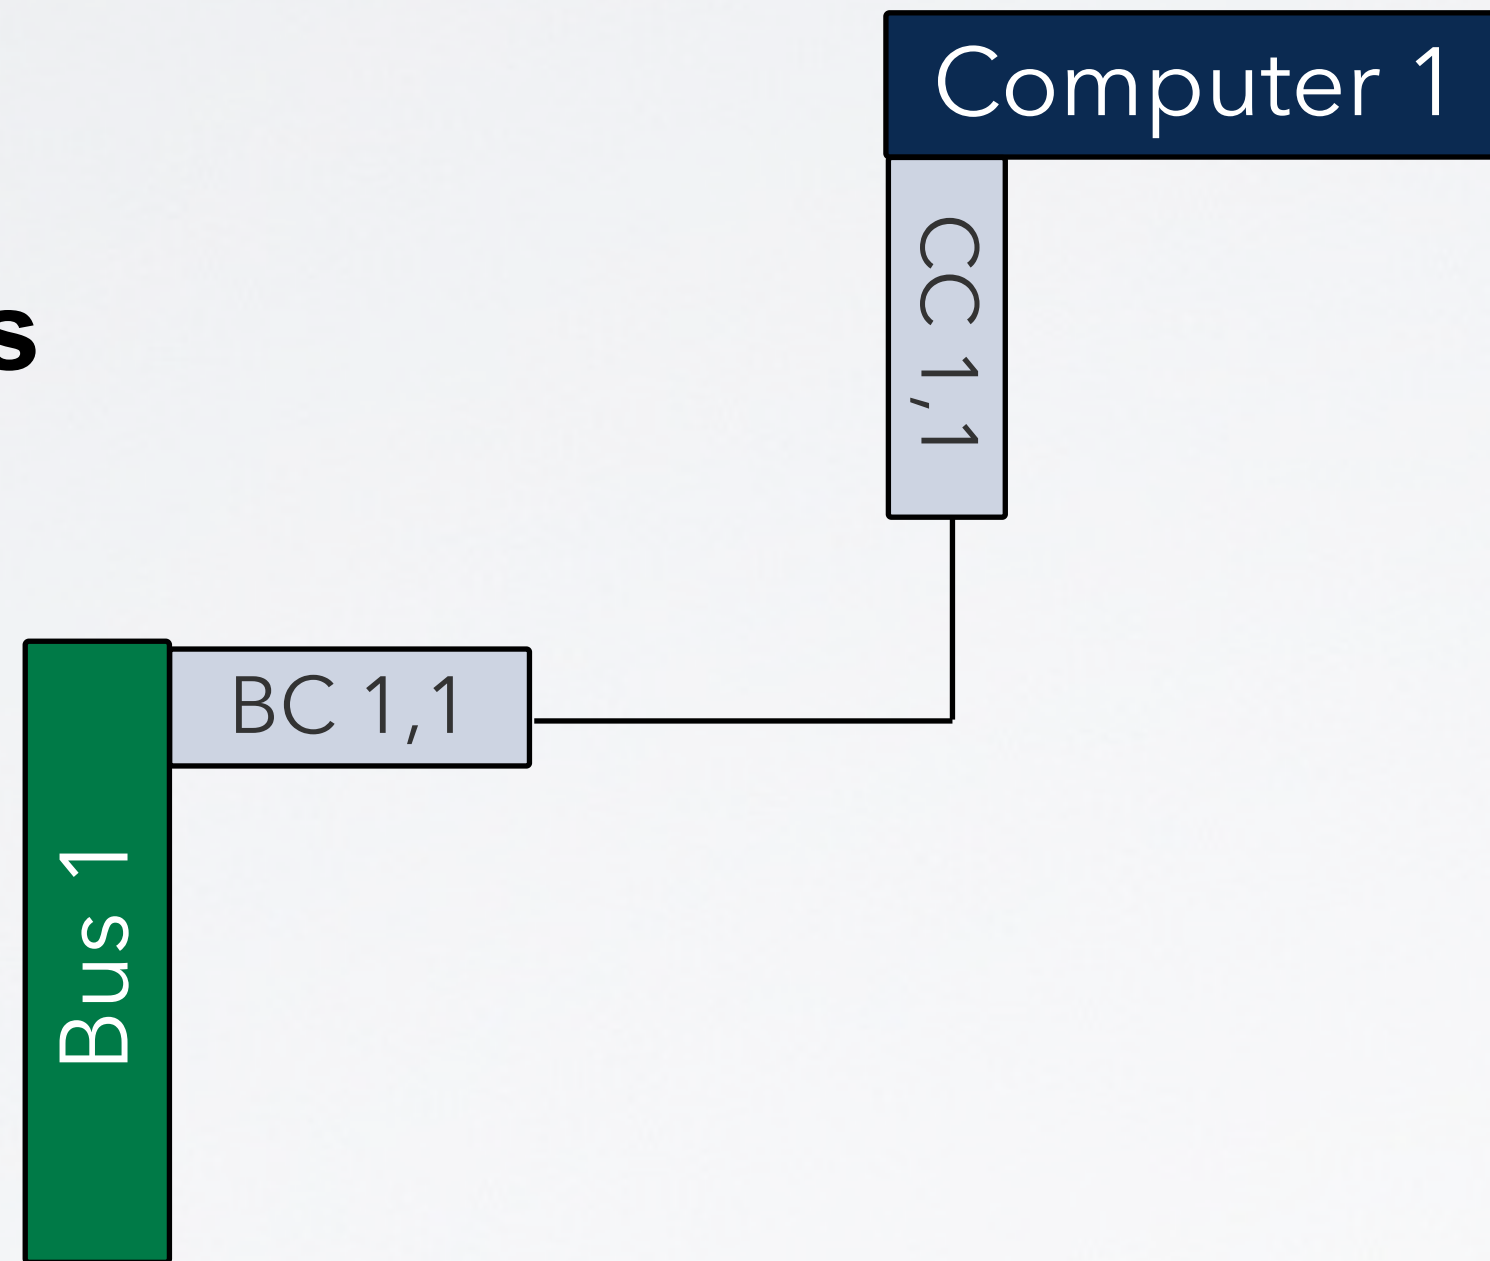
Computer, Bus: can fail per se

CC: Computer-Connector
fault also breaks the Computer

BC: Bus-Connector
fault also breaks Bus

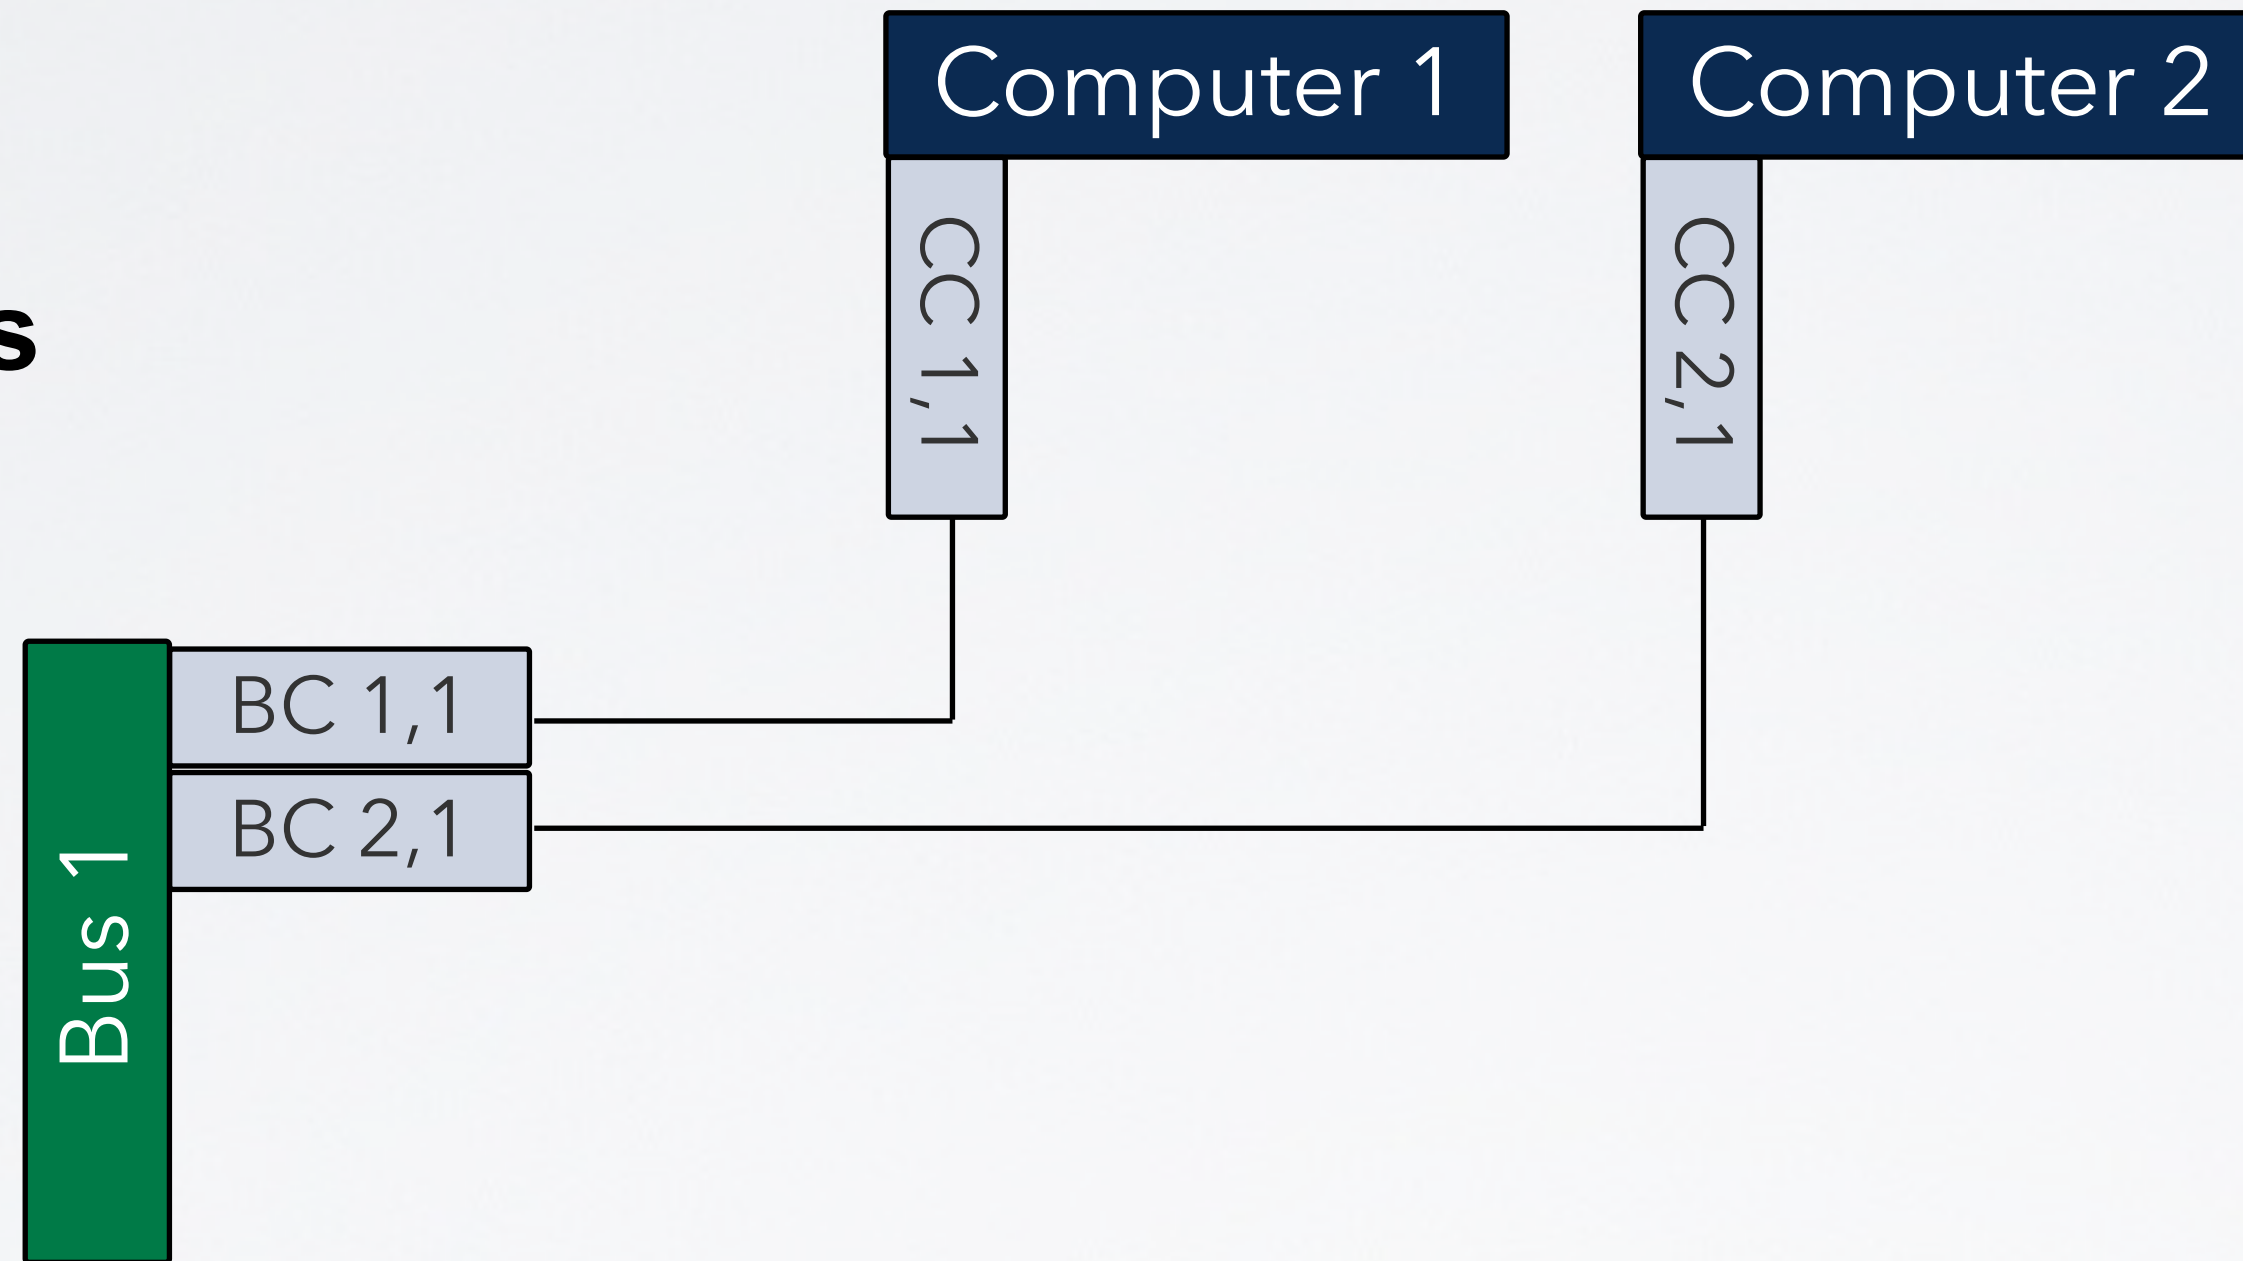
1 Buses

1 Computers




1 Buses

2 Computers

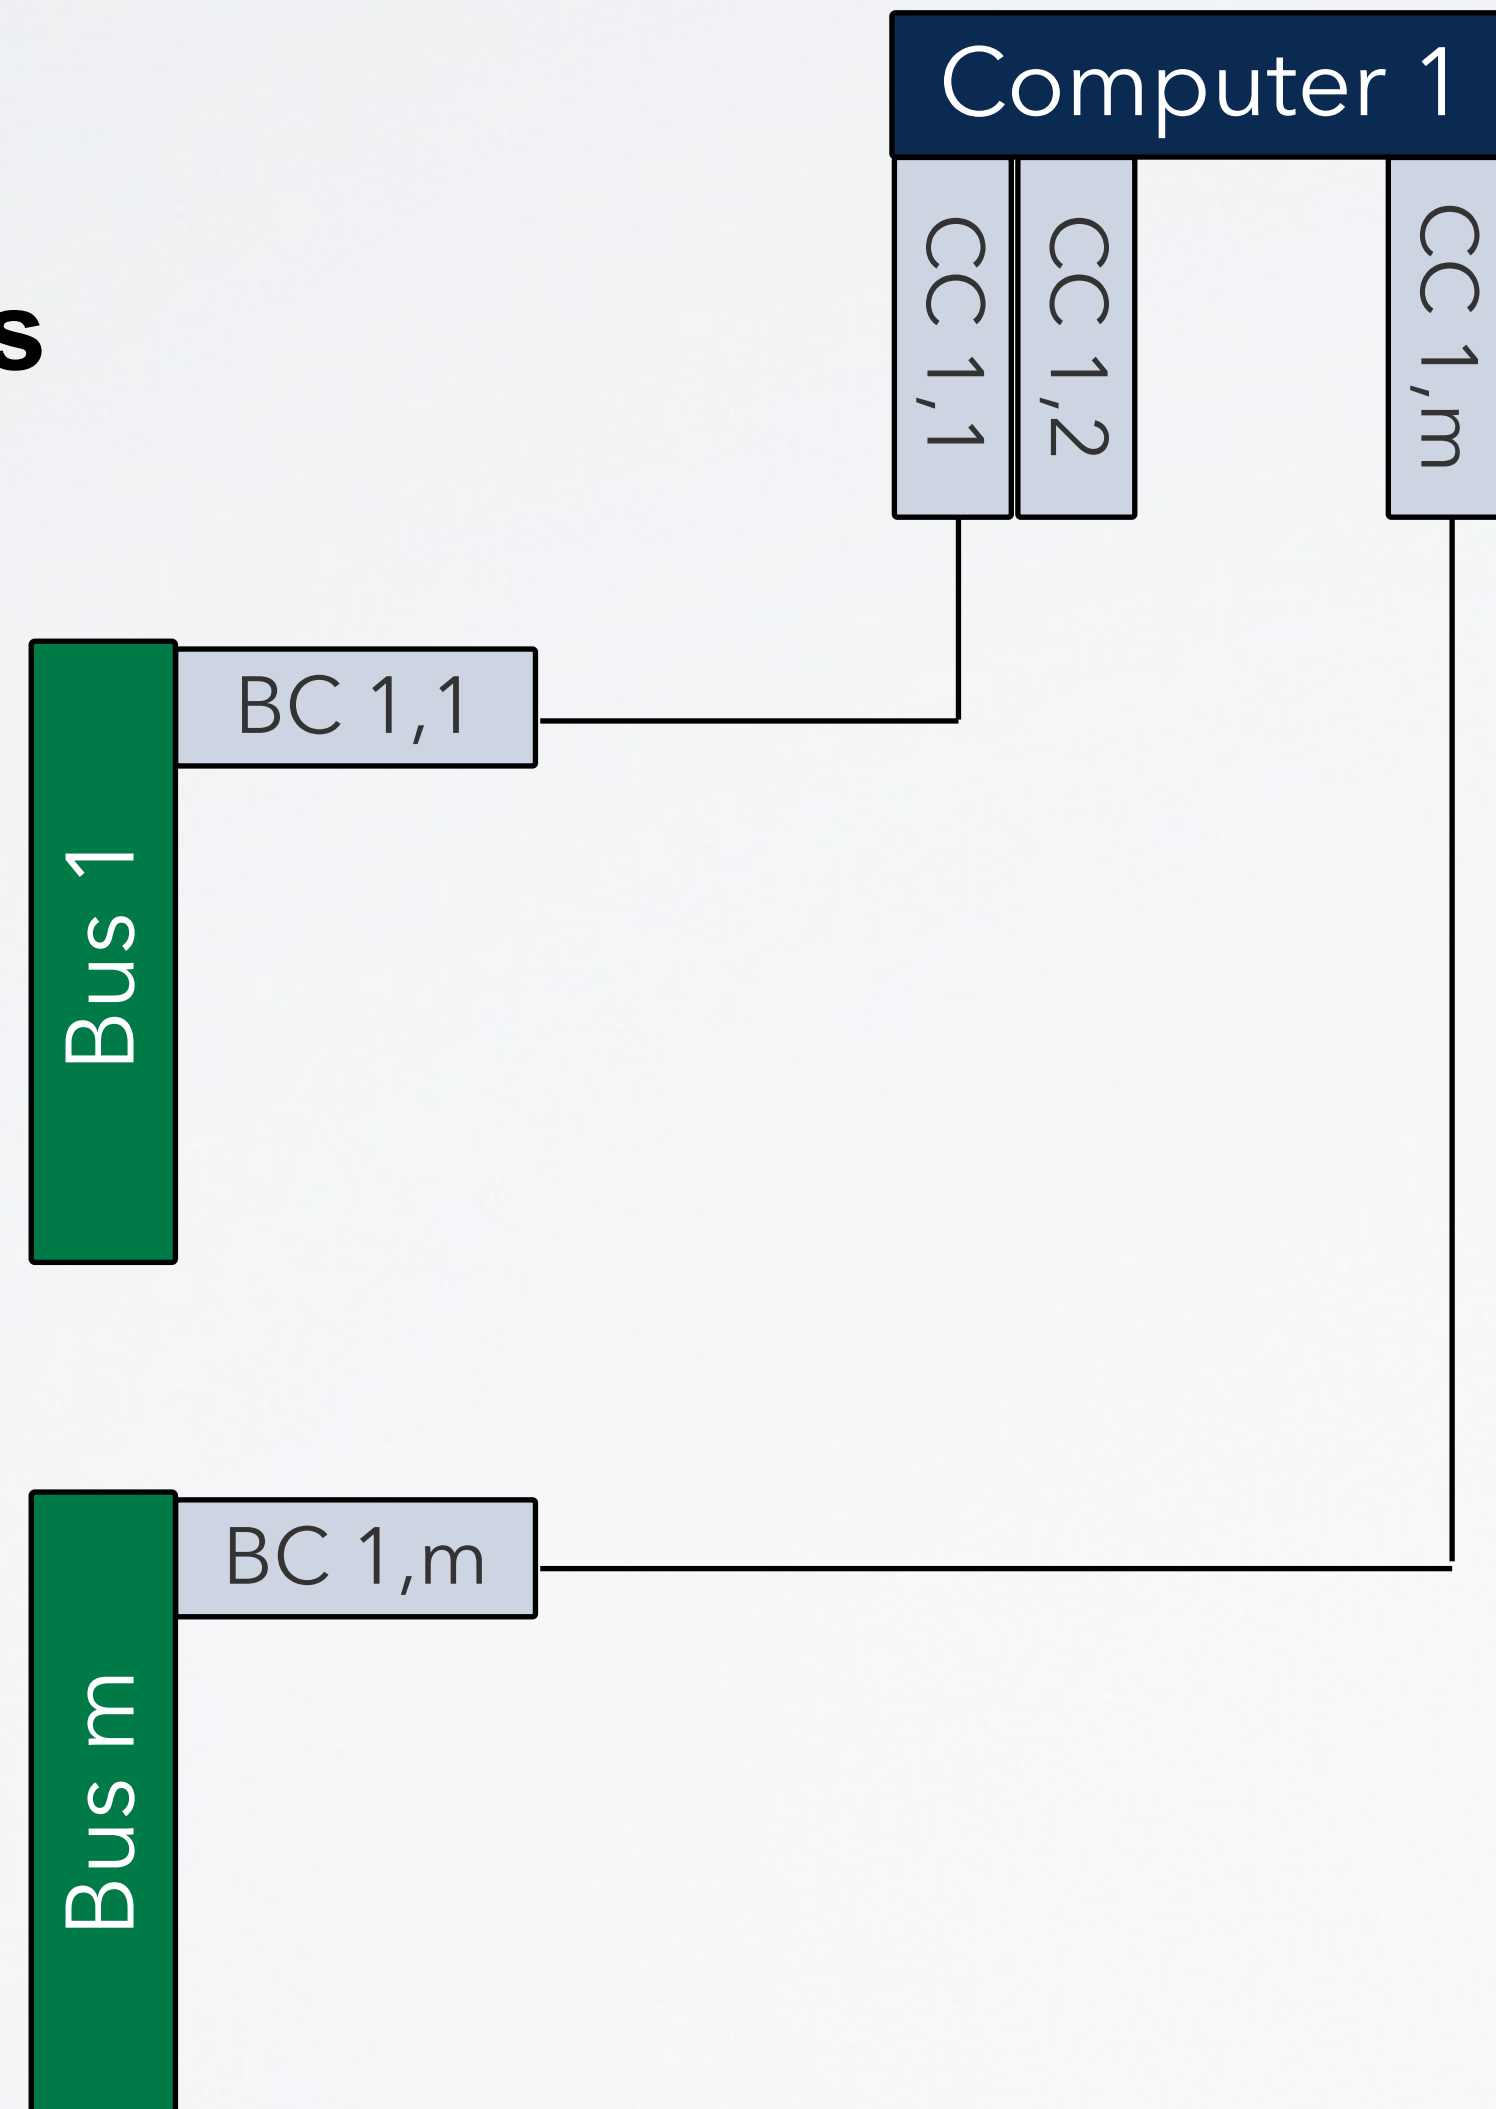


1 Buses

N Computers



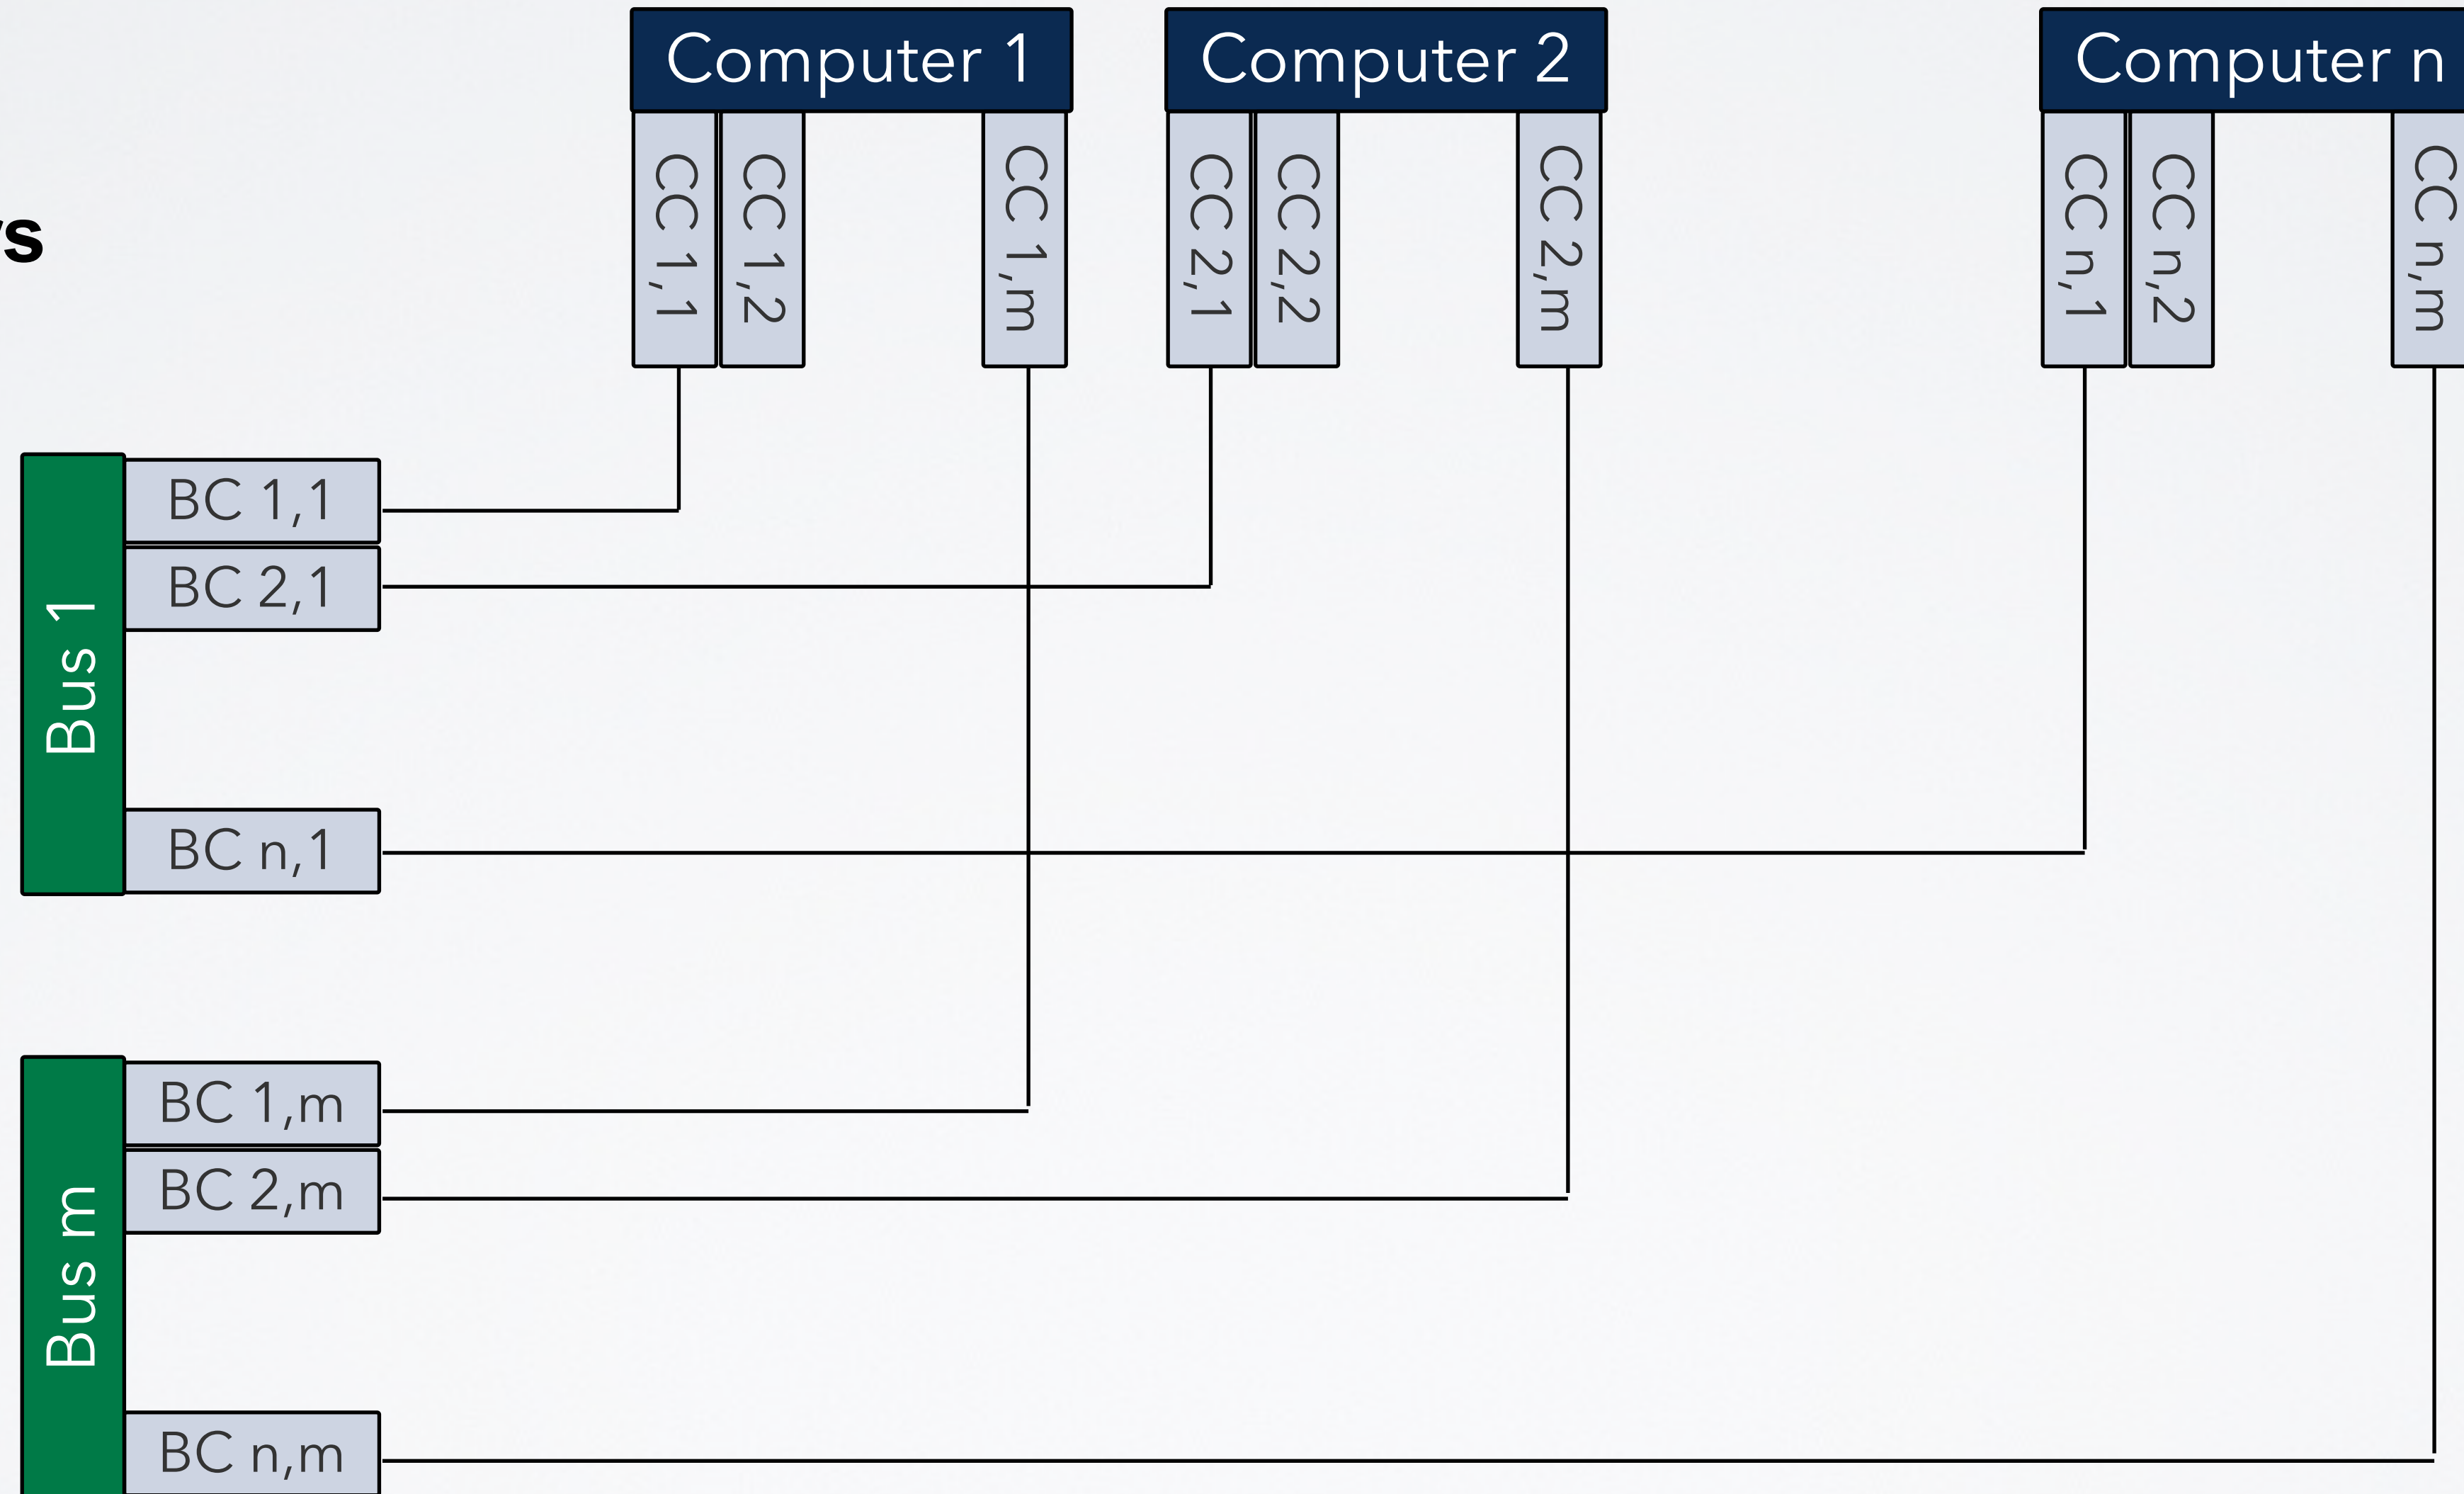
M Buses
1 Computers



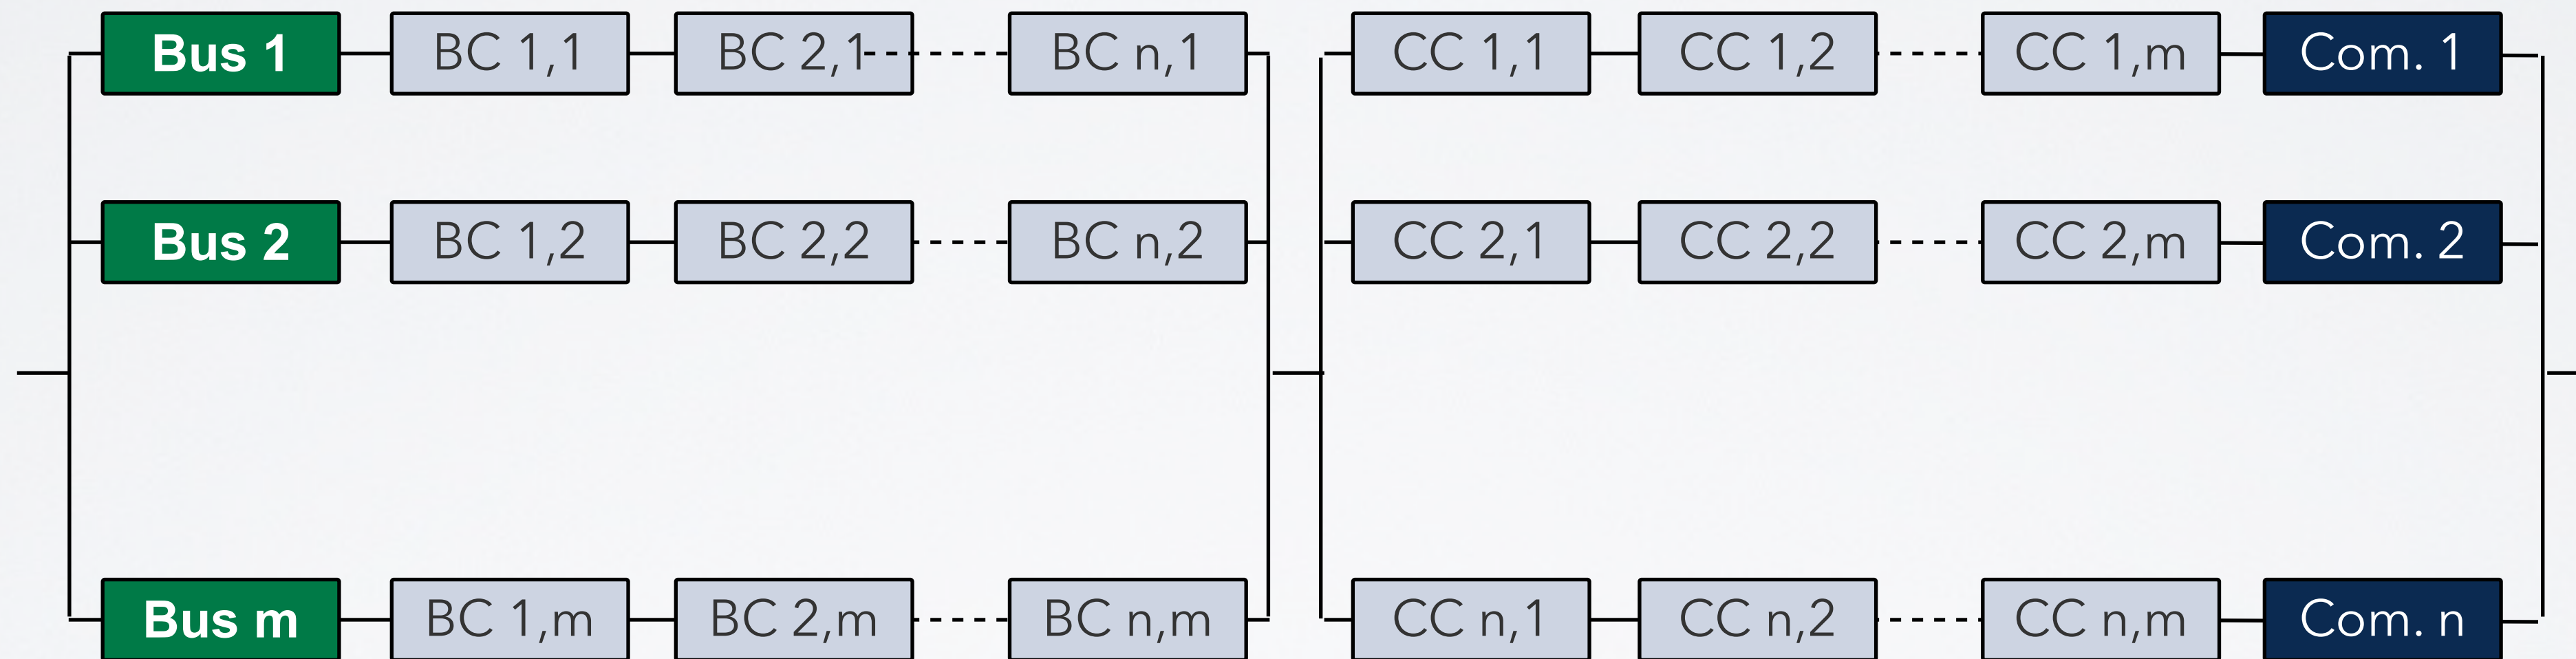
Q1/MODEL1: CONCRETE MODEL

M Buses

N Computers



Q1/MODEL1: CONCRETE MODEL FOR N,M



$$R_{whole}(n, m) = \left(1 - \left(1 - R_{Bus} \cdot R_{BC}^n\right)^m\right) \cdot \left(1 - \left(1 - R_{Computer} \cdot R_{CC}^m\right)^n\right)$$

$$\text{then: } R_{CC}, R_{BC} < 1: \lim_{n, m \rightarrow \infty} R(n, m) =$$

- System built of Synapses (John von Neumann, 1956)
- Computation and Fault Model :
 - Synapses deliver „0“ or „1“
 - Synapses deliver with $R > 0,5$:
 - with probability R correct result
 - with $(1-R)$ wrong result
- Then we can build systems that deliver correct result for any (arbitrarily high) probability R

- Q1: Is it possible to build arbitrarily reliable Systems out of unreliable components?
- Q2: Can we achieve consensus in the presence of faults (consensus: all non-faulty components agree on action)?
- Q3: Is there an algorithm to determine for a system with a given setting of access control permissions, whether or not a Subject A can obtain a right on Object B?

2 Models per Question !

Q2: Can we achieve consensus in the presence of faults
all non-faulty components agree on action?

- all correctly working units agree on result/action
- agreement non trivial (based on exchange of messages)

- p,q processes
 - communicate using messages
 - messages can get lost
 - no upper time for message delivery known
 - do not crash, do not cheat
- p,q to agree on action (e.g. attack, retreat, ...)
- how many messages needed ?
- first mentioned: Jim Gray 1978

Result: there is no protocol with finite messages

Prove by contradiction:


- assume there are finite protocols ($mp \dashrightarrow q, mq \dashrightarrow p$)*
- choose the shortest protocol MP,
- last message MX: $mp \dashrightarrow q$ or $mq \dashrightarrow p$
- MX can get lost
- \Rightarrow must not be relied upon \Rightarrow can be omitted
- \Rightarrow MP not the shortest protocol.
- \Rightarrow no finite protocol

n processes, f traitors, $n-f$ loyals


- communicate by reliable and timely messages (synchronous messages)
- traitors lie, also cheat on forwarding messages
- try to confuse loyals

Goal:

- loyals try to agree on non-trivial action (attack, retreat)
- non-trivial more specific:
 - one process is commander
 - if commander is loyal and gives an order, loyals follow the order otherwise loyals agree on arbitrary action




3 Processes: 1 traitor, 2 loyals



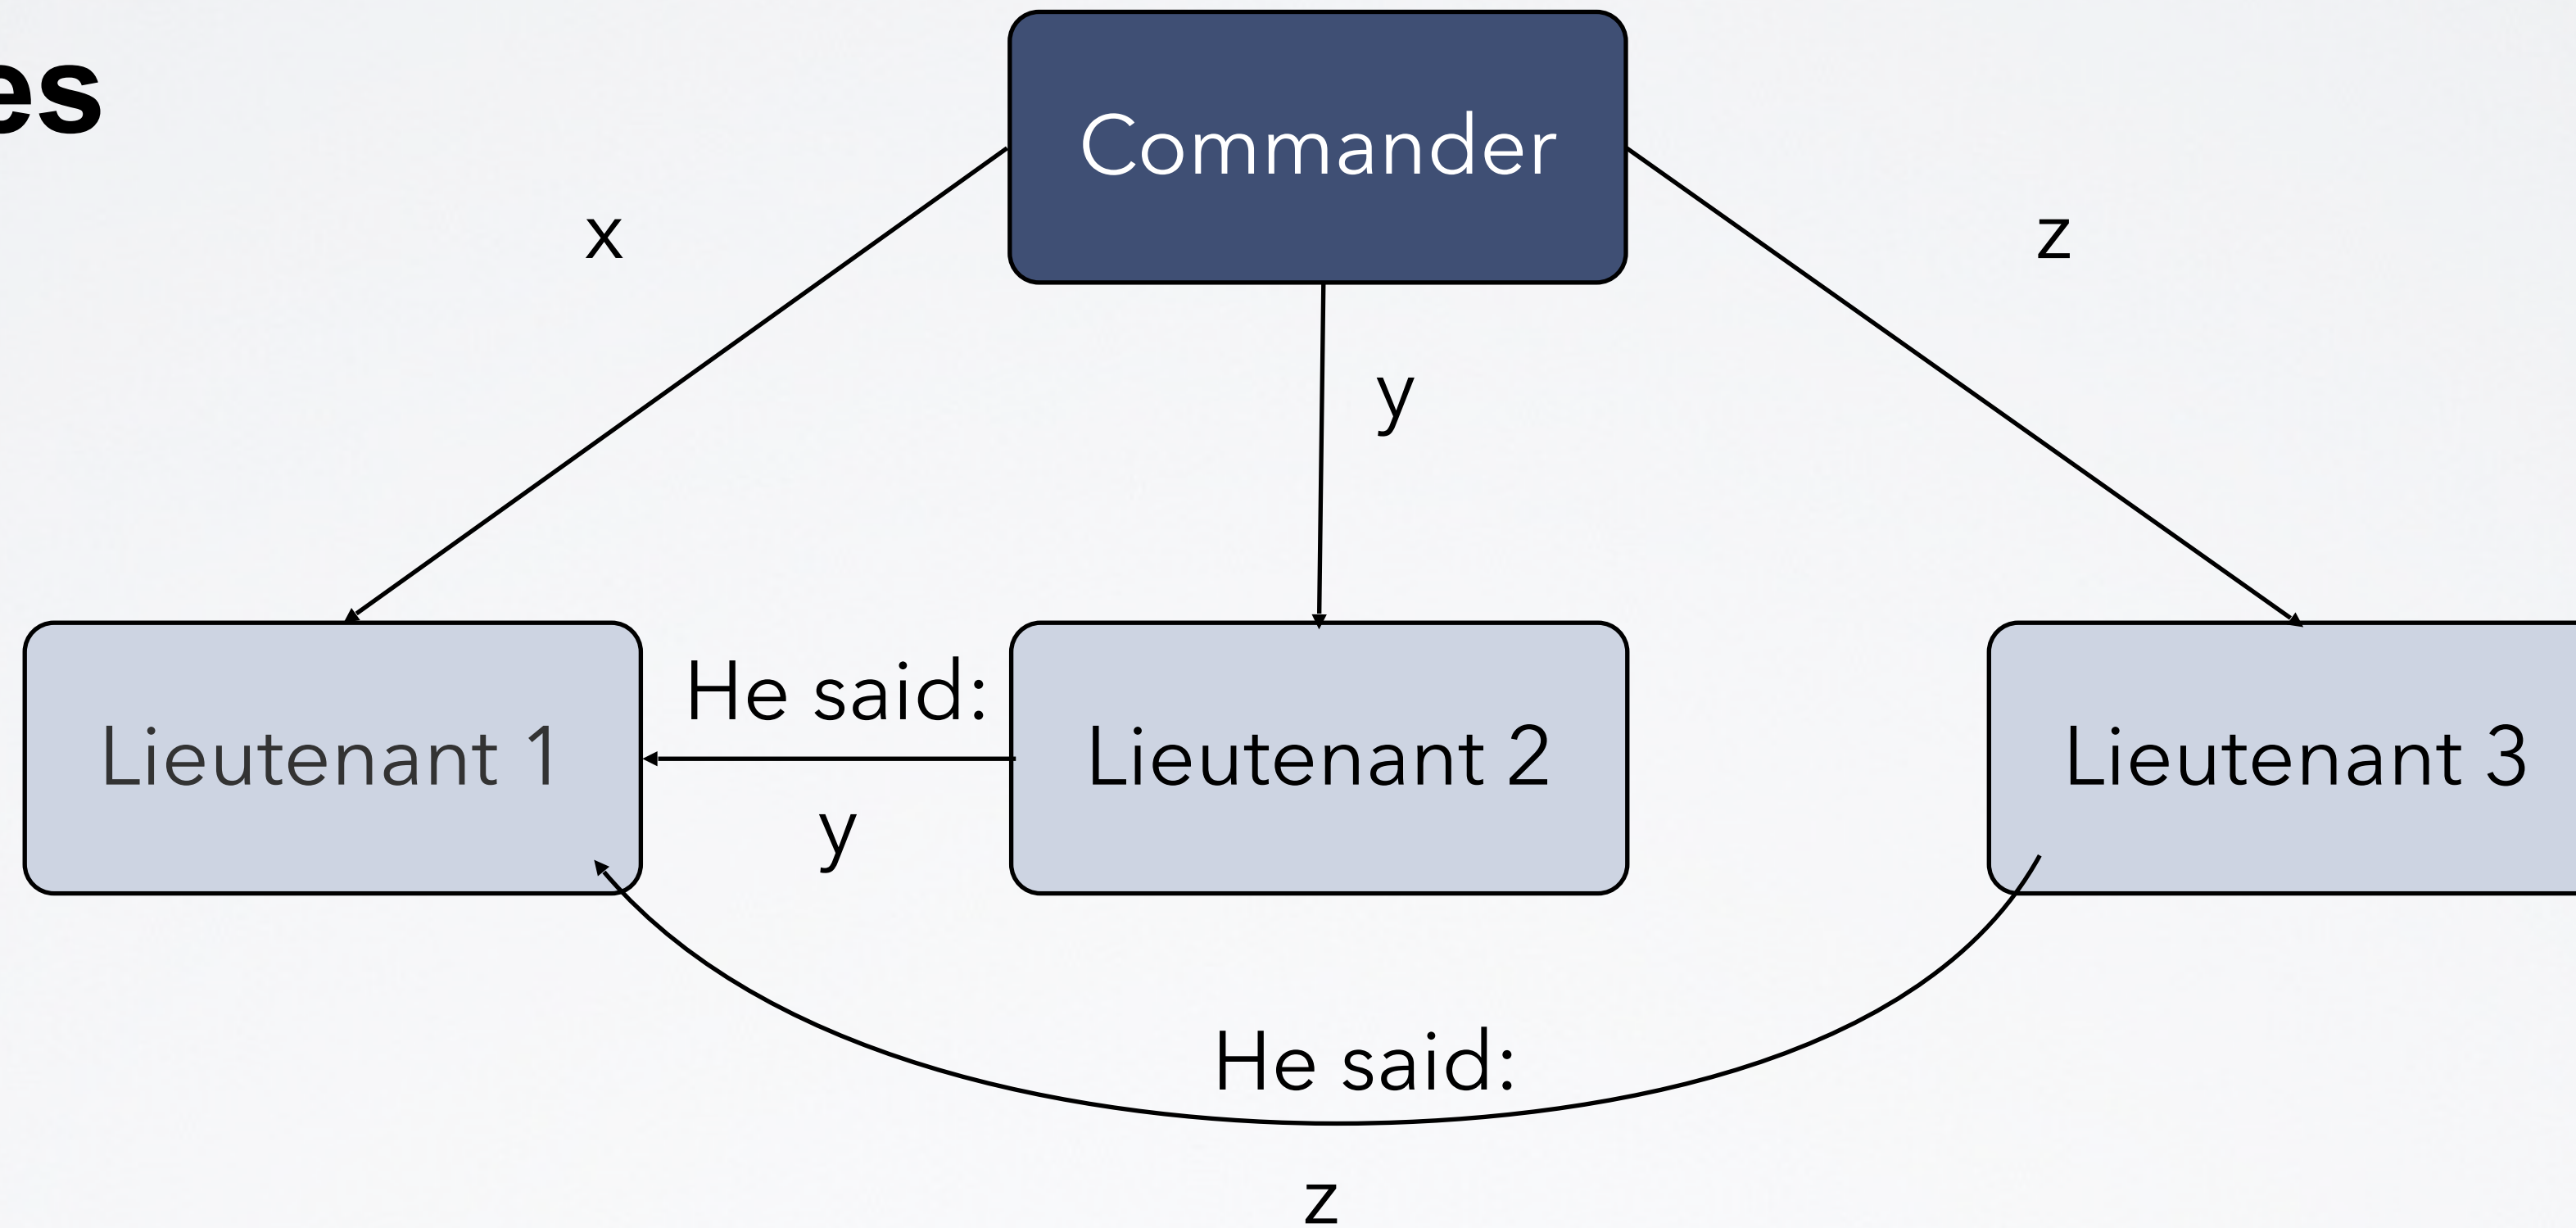
3 Processes: 1 traitor, 2 loyals

=> 3 processes not sufficient to tolerate 1 traitor

4 Processes



4 Processes



all lieutenant receive $x, y, z \Rightarrow$ can decide

General result: $3f + 1$ processes needed to tolerate f traitors

- Q1: Is it possible to build arbitrarily reliable Systems out of unreliable components?
- Q2: Can we achieve consensus in the presence of faults (consensus: all non-faulty components agree on action)?
- Q3: Is there an algorithm to determine for a system with a given setting of access control permissions, whether or not a Subject A can obtain a right on Object B?

2 Models per Question !

Q3: Is there an algorithm to determine for a system with a given setting of access control permissions, whether or not a Subject A can obtain a right on Object B?

- given a system of entities, acting as subjects and objects subjects performs operations on objects
- dynamic: subjects and objects are created and deleted
- access control permissions between entities can be changed according to some rules

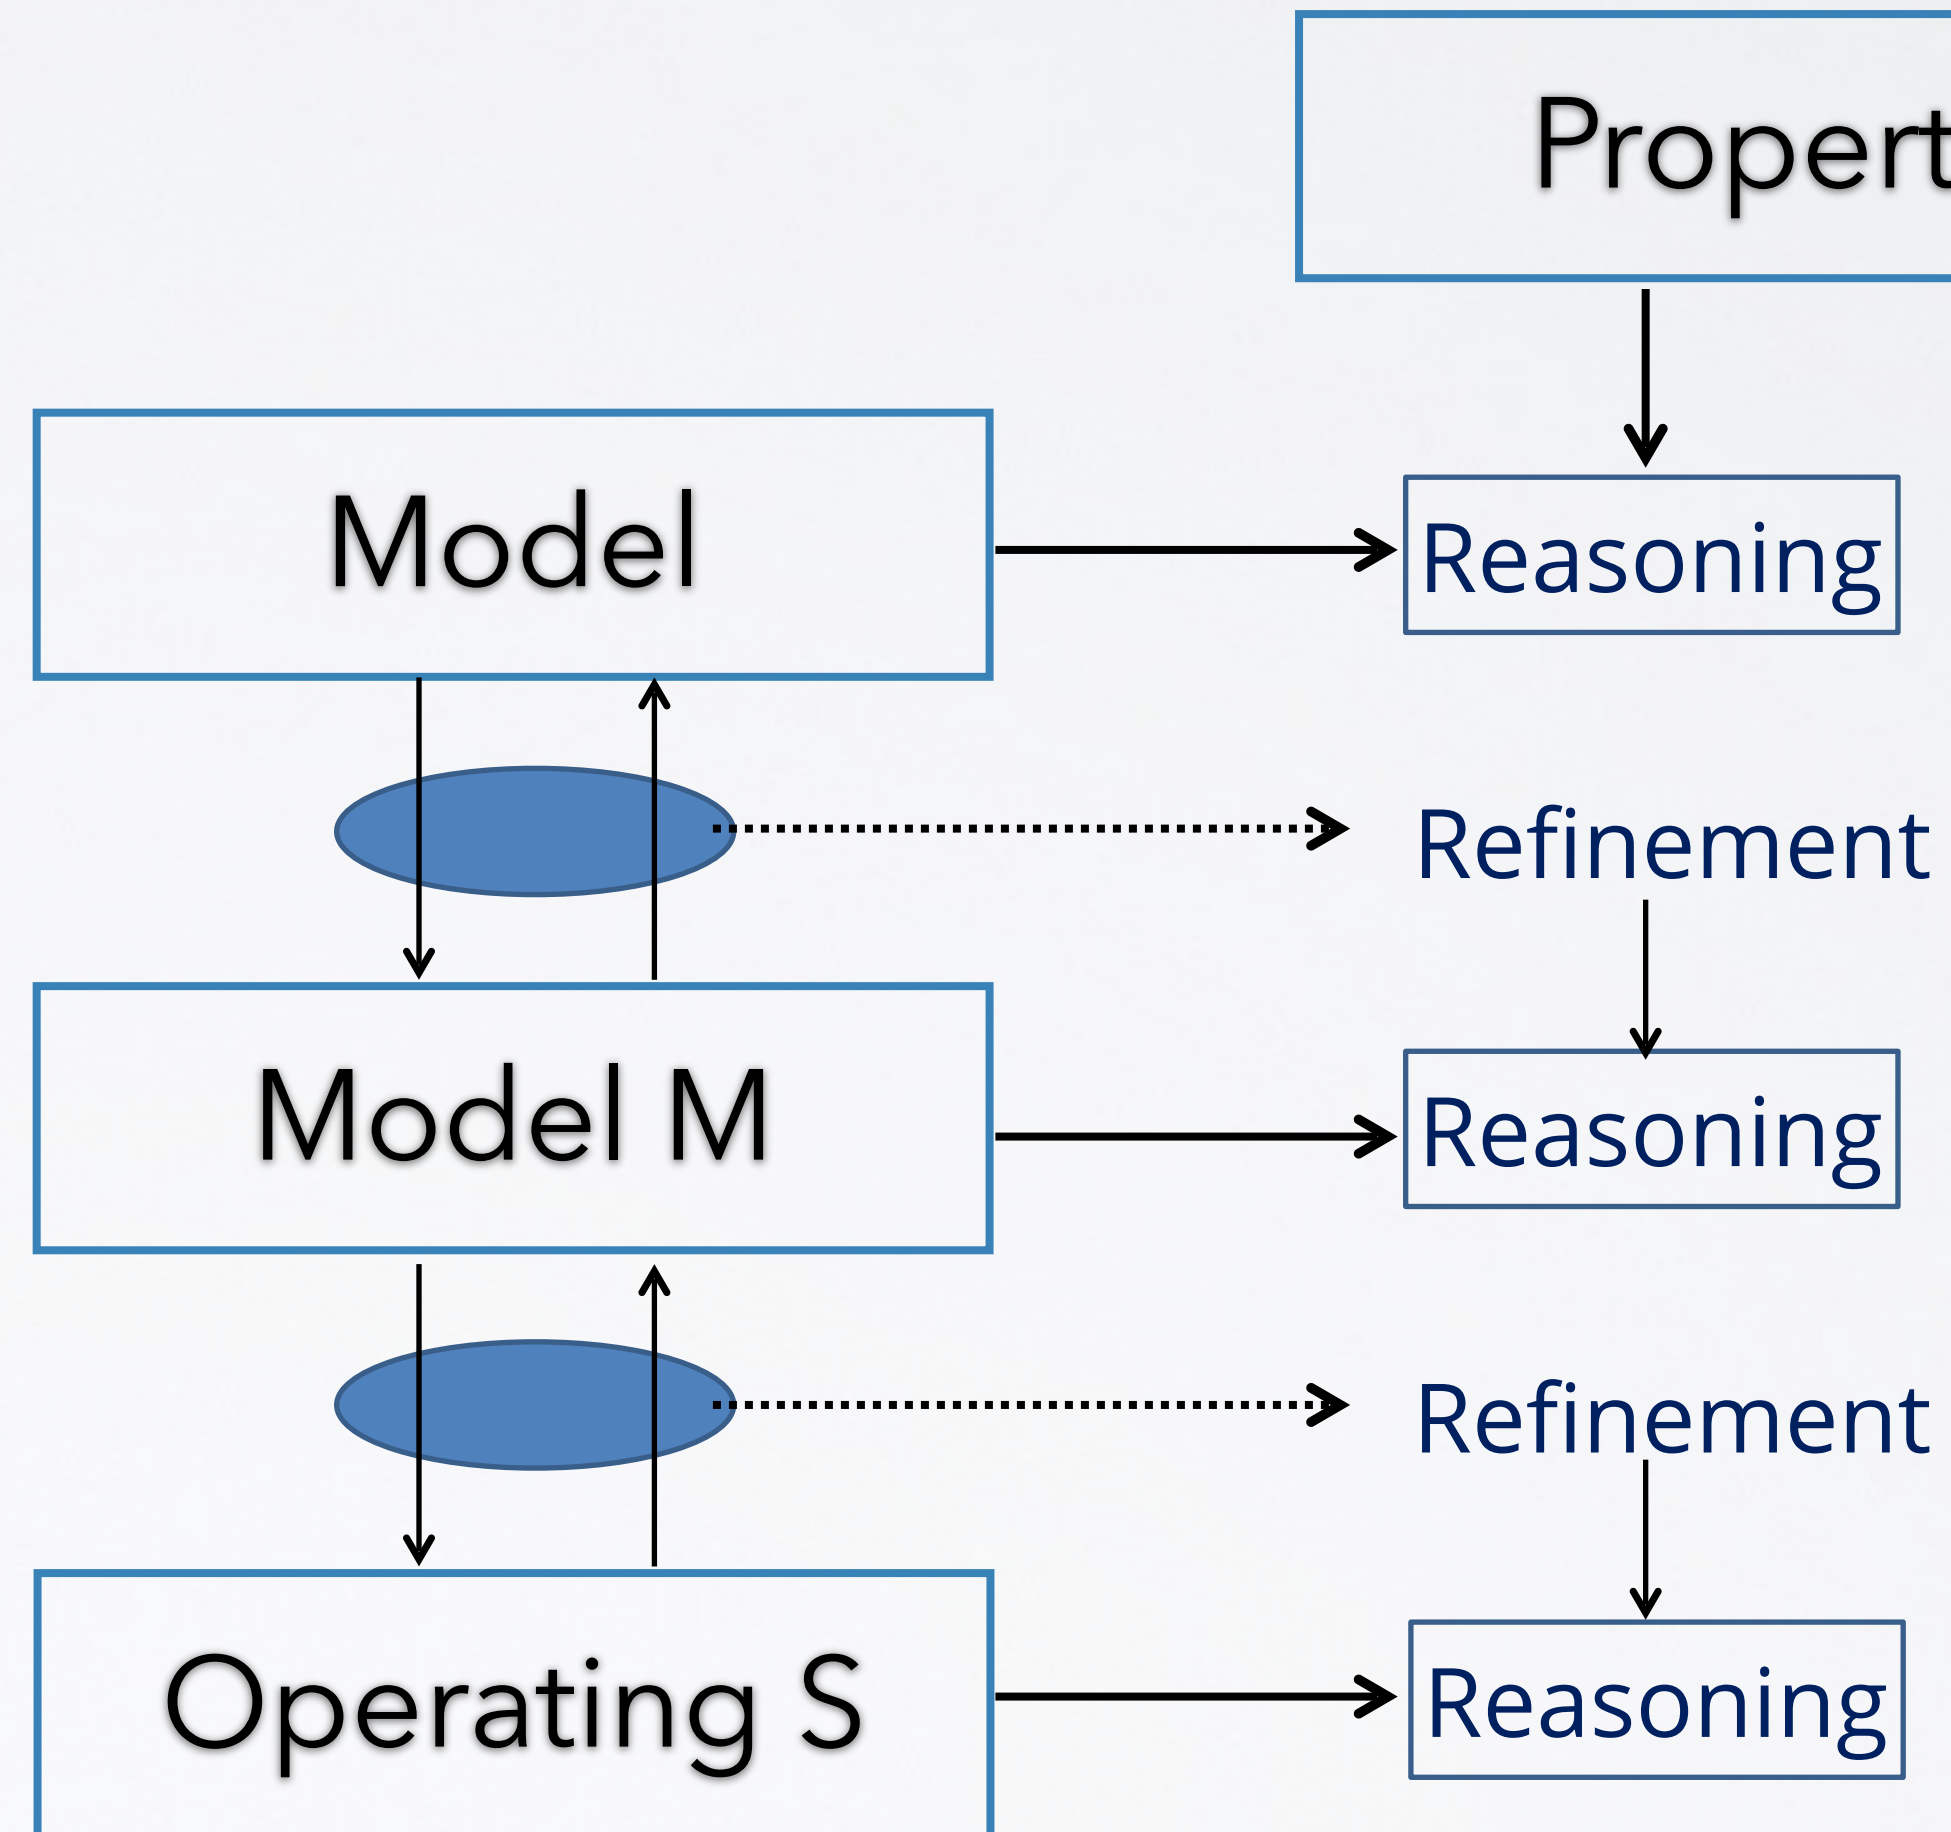
higher level models:

- Bell La Padula,
- Chinese wall

...

access control:

- 1) ACM-based operations
- 2) take grant



Subjects: S

Objects: O

Entities: $E = S \cup O$

Rights: {read, write, own,...}

Matrix: $S \times E \times R$

	O1	O2	S1	S2	S3
S1	r,w,own	r,w	r,w,own	--	r,w
S2	r,w	r,w,own	-	r,w,own	r
S3	r,w	r	w	--	r,w,own

Simple ACM Operations:

enter / delete R into cell (s,o)

create subject / object

destroy subject / object

ACM

- Access Control List (ACL)

	O1	O2	S1	S2	S3
S1	r,w,own	r,w	r,w,own	--	r,w
S2	r,w	r,w,own	-	r,w,own	r
S3	r,w	r	w	--	r,w,own

- Capabilities

	O1	O2	S1	S2	S3
S1	r,w,own	r,w	r,w,own	--	r,w
S2	r,w	r,w,own	-	r,w,own	r
S3	r,w	r	w	--	r,w,own

- Define Protection Mechanisms of an Operating System in terms of sequences of simple ACM operations
- only such defined mechanism provided by the OS can be used to manipulate ACM

- "Leakage":
 - an access right is placed into S/O that has not been there before
 - it does not matter whether or not that is allowed
- Is leakage decidable ?

Define OS-Mechanisms by simple ACM-Operations:

example:

UNIX create file (S1,F)

create object

enter own into A(S1,F)

enter read into A(S1,F)

enter write into A(S1,F)

	O1	O2	S1	S2	F
S1	r,w,own	r,w	r,w,own	--	r,w, own
S2	r,w	r,w,own	-	r,w,own	-
S3	r,w	r	w	--	-

Example:

UNIX chmod -w (S2,F)

if own $\in A(\text{caller}, F)$

then delete w in $A(S2, F)$

	O1	O2	S1	S2	F
S1	r,w,own	r,w	r,w,own	--	r,w, own
S2	r,w	r,w,own	-	r,w,own	r,-
S3	r,w	r	w	--	-

Q3:

Given an OS with a ACM-based description of protection mechanisms is "Leakage" decidable for any R in $A(x,y)$?

Decidable

- or
- if no subjects/objects can be created
 - if only **one** primitive ACM operation per OS-Mechanism
- by exhaustive search !

Q3 in general:

- undecidable (proof: reduction to Turing machine)

ACM

- Access Control List (ACL)

	O1	O2	S1	S2	S3
S1	r,w,own	r,w	r,w,own	--	r,w
S2	r,w	r,w,own	-	r,w,own	r
S3	r,w	r	w	--	r,w,own

- Capabilities

	O1	O2	S1	S2	S3
S1	r,w,own	r,w	r,w,own	--	r,w
S2	r,w	r,w,own	-	r,w,own	r
S3	r,w	r	w	--	r,w,own


Directed Graph:

Subjects: ●

Objects: ○

Either S or O: ⊗


x has capability on Y
with set of rights α on y:



t take right

x has cap with set of rights


τ that includes t



g grant right

x has cap with set of rights


γ that includes **g**




Rules:

take rule ($\alpha \subseteq \beta$)

a takes (α to y) from z




grant rule ($\alpha \subseteq \beta$)



Rules:


create rule

x create (α to new vertex) y



remove rule

x removes (α to) y



CanShare(α, x, y, G_0):


there exists a sequence of $G_0 \dots G_n$ with $G_0 \vdash^* G_n$

and there is an edge in G_n : 

Q3/ 2: CAREFUL: LEMMA


take rule ($\alpha \subset \beta$)

a takes (α to y) from z




grant rule ($\alpha \subset \beta$)

z grants (α to y) to




Question:



Q3/ 2: CAREFUL: LEMMA

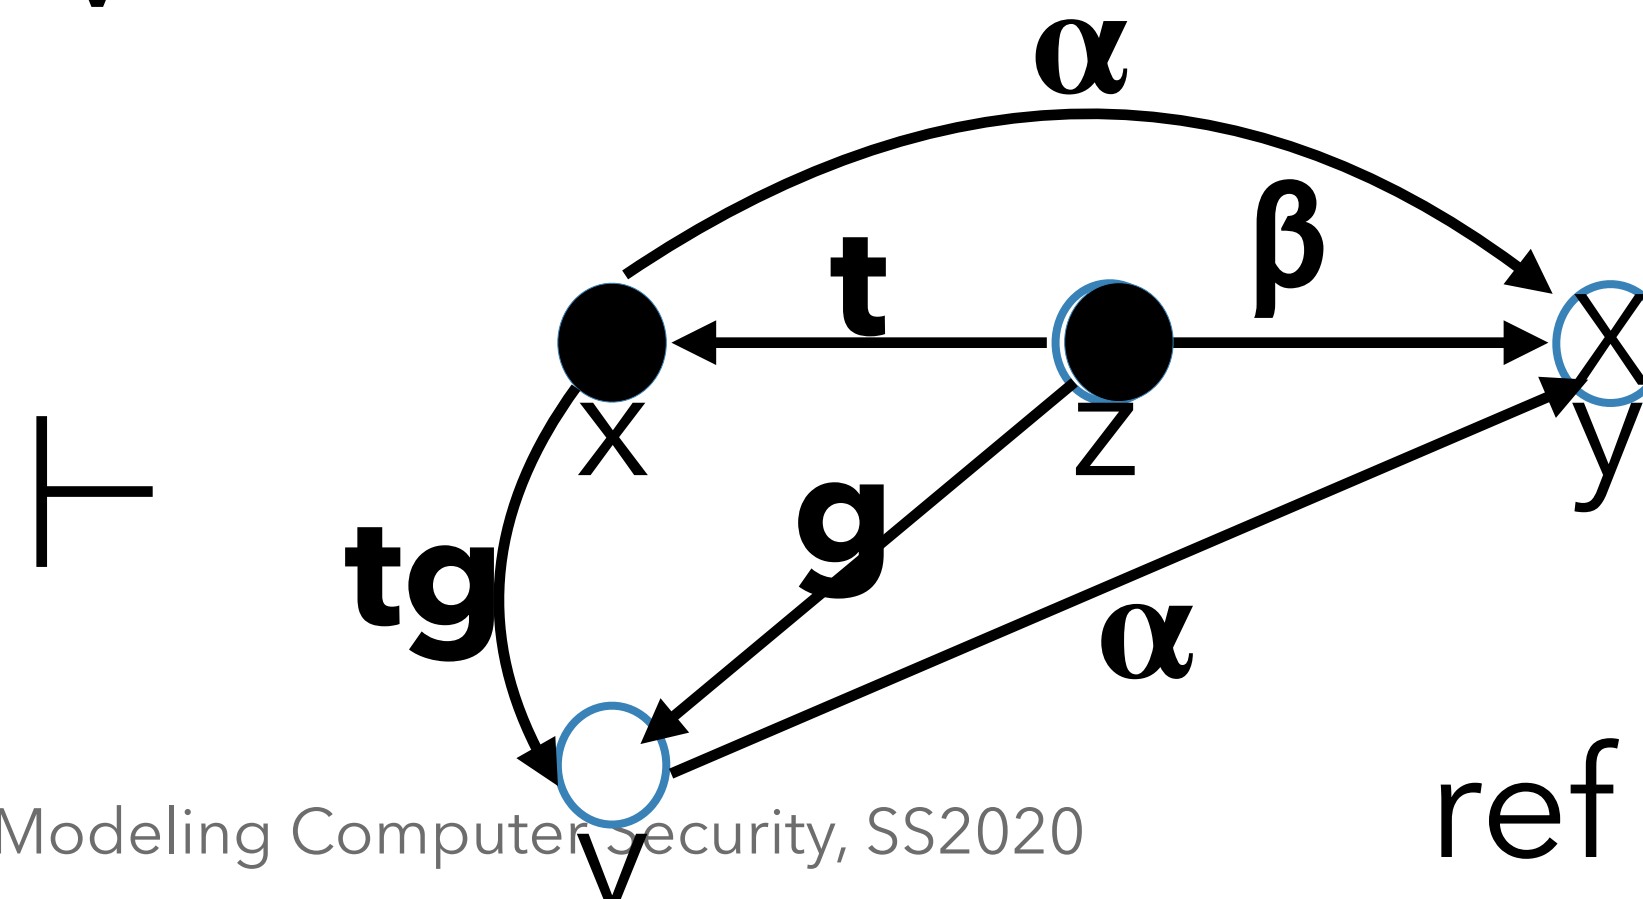

$(\alpha \subset \beta)$

create rule



z takes (g to v) from x

z grants (α to y) to v



CanShare(α, x, y, G_0):

there exists a sequence of $G_0 \dots G_n$ with $G_0 \vdash^* G_n$

and there is an edge: 

CanShare decidable in linear time !

- three questions, 2 models per question, different answers !!!
- modeling is powerful
- need to look extremely carefully into understanding
models !!!

- Q1/M1:

Pfitzmann A., Härtig H. (1982) Grenzwerte der Zuverlässigkeit von Parallel-Serien-Systemen. In: Nett E., Schwärtzel H. (eds) Fehlertolerierende Rechnersysteme. Informatik-Fachberichte, vol 54. Springer, Berlin, Heidelberg (in German only)

- Q1/M2:

John v. Neuman, PROBABILISTIC LOGICS AND THE SYNTHESIS OF RELIABLE. ORGANISMS FROM UNRELIABLE COMPONENTS.

- Q2: most textbooks on distributed systems

- Q3: textbook: Matt Bishop, Computer Security, Art and Science, Addison Wesley 2002