When is CAN the weakest link?

A bound on Failures-In-Time in CAN-Based Distributed Real-Time Systems

Arpan Gujarati
Björn B. Brandenburg

Max Planck Institute for Software Systems

Failures due to Transient Faults

Harsh environments

- Spark plugs
- Hard radiation
- High-power machinery

Failures due to Transient Faults

Harsh environments

- ➡ Spark plugs
- ➡ Hard radiation
- ➡ High-power machinery

Electromagnetic Interference (EMI)

- ➡ **Bit-flips** in the hosts
- ➡ … and in the network

Failures due to Transient Faults

Harsh environments

- ➡ Spark plugs
- ➡ Hard radiation
- ➡ High-power machinery

Electromagnetic Interference (EMI)

- ➡ **Bit-flips** in the hosts
- ➡ … and in the network

EMI-induced transient faults

- ➡ Manifest as **program-visible failures**

Failures due to Transient Faults

Transmission failures
(faults on the wire)

Commission failures
(bit-flips in the memory buffers)

Crash failures
(due to fault-induced exceptions)

Failures due to Transient Faults

Transmission failures
(faults on the wire)

Tolerated by **error detection** and **retransmissions**

Commission failures
(bit-flips in the memory buffers)

Crash failures
(due to fault-induced exceptions)

Failures due to Transient Faults

Transmission failures
(faults on the wire)

Tolerated by **error detection** and **retransmissions**

Commission failures
(bit-flips in the memory buffers)

Crash failures
(due to fault-induced exceptions)

Tolerated by active **replication of tasks** on independent hosts

Failures due to Transient Faults

Transmission failures
(faults on the wire)

Tolerated by **error detection** and **retransmissions**

Commission failures
(bit-flips in the memory buffers)

Crash failures
(due to fault-induced exceptions)

Tolerated by active **replication of tasks** on independent hosts

- How to decide the **best replication strategy**?
  - ➡ Is Triple Modular Redundancy (TMR) enough? or is Quadruple Modular Redundancy (QMR) required?
  - ➡ Would you replicate only the high-frequency tasks? or only the high-criticality tasks?

Retransmissions vs. Replication Tradeoff

Retransmissions vs. Replication Tradeoff

For tolerating retransmissions-induced delays

- ➡ (Ensure no deadline violations)
- ➡ **The more slack, the better!**

Retransmissions vs. Replication Tradeoff

For tolerating retransmissions-induced delays
- ➡ (Ensure no deadline violations)
- ➡ **The more slack, the better!**

versus

Active replication of tasks
- ➡ **Reduced slack** in the schedule

Retransmissions vs. Replication Tradeoff

For tolerating retransmissions-induced delays

- ➡ (Ensure no deadline violations)
- ➡ **The more slack, the better!**

versus

Active replication of tasks

- ➡ **Reduced slack** in the schedule

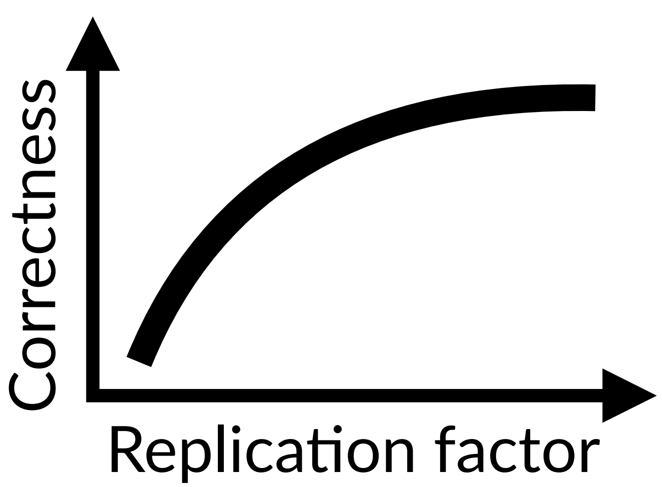

Retransmissions vs. Replication Tradeoff

For tolerating retransmissions-induced delays
- ➡ (Ensure no deadline violations)
- ➡ **The more slack, the better!**

versus

Active replication of tasks
- ➡ **Reduced slack** in the schedule

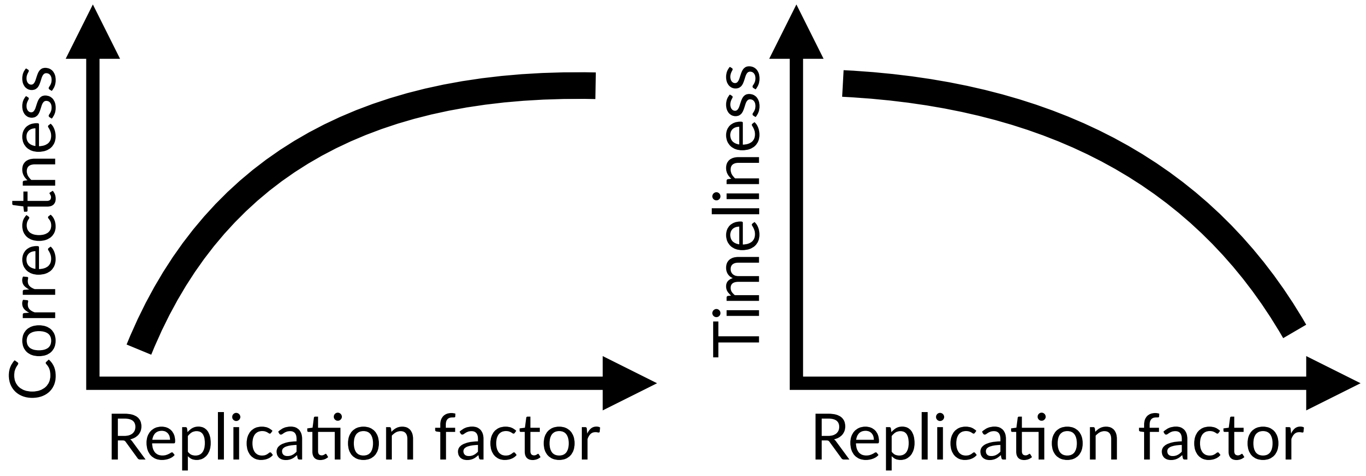

Retransmissions vs. Replication Tradeoff

For tolerating retransmissions-induced delays

- ➡ (Ensure no deadline violations)
- ➡ **The more slack, the better!**

versus

Active replication of tasks

- ➡ **Reduced slack** in the schedule

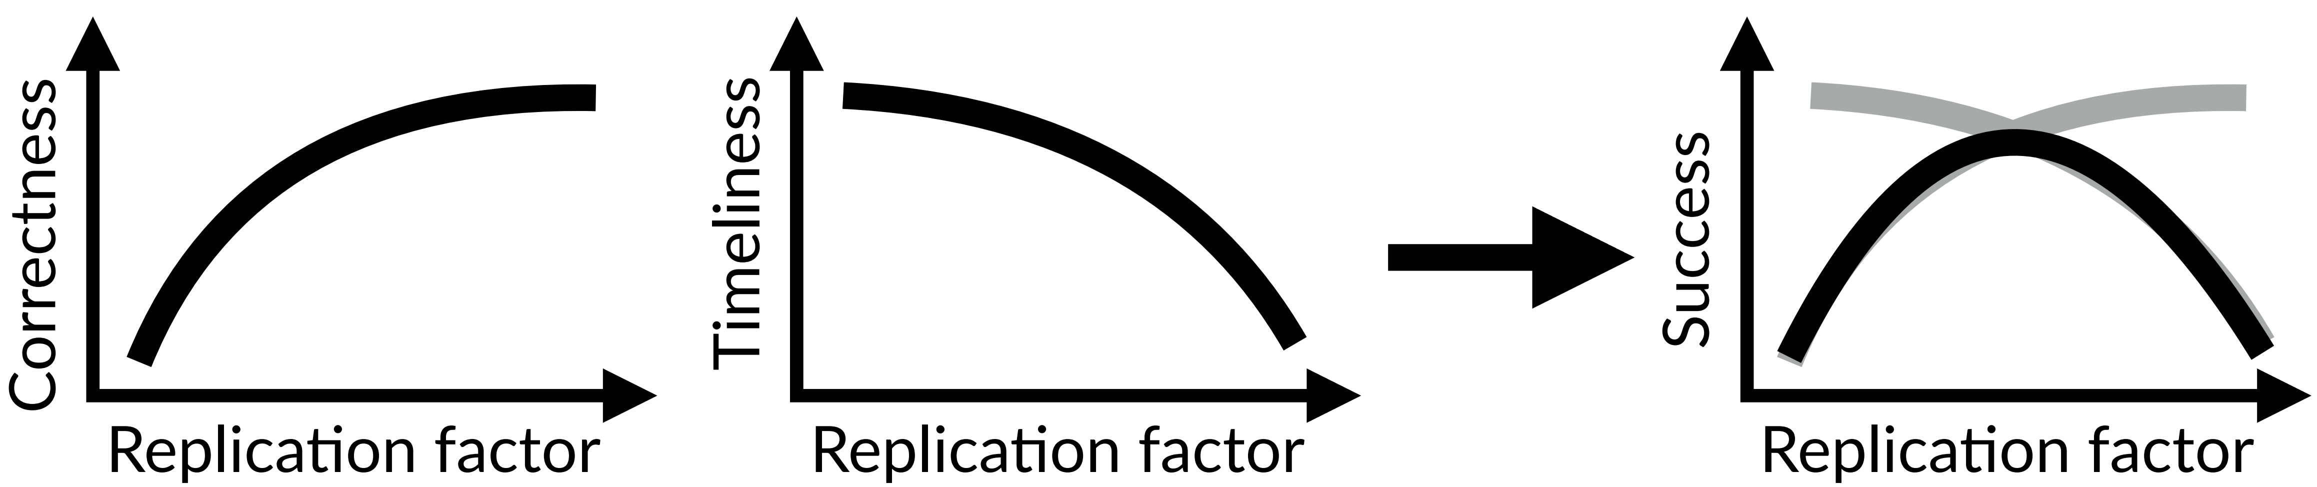

Retransmissions vs. Replication Tradeoff

For tolerating retransmissions-induced delays
- ➡ (Ensure no deadline violations)
- ➡ **The more slack, the better!**

versus

Active replication of tasks
- ➡ **Reduced slack** in the schedule

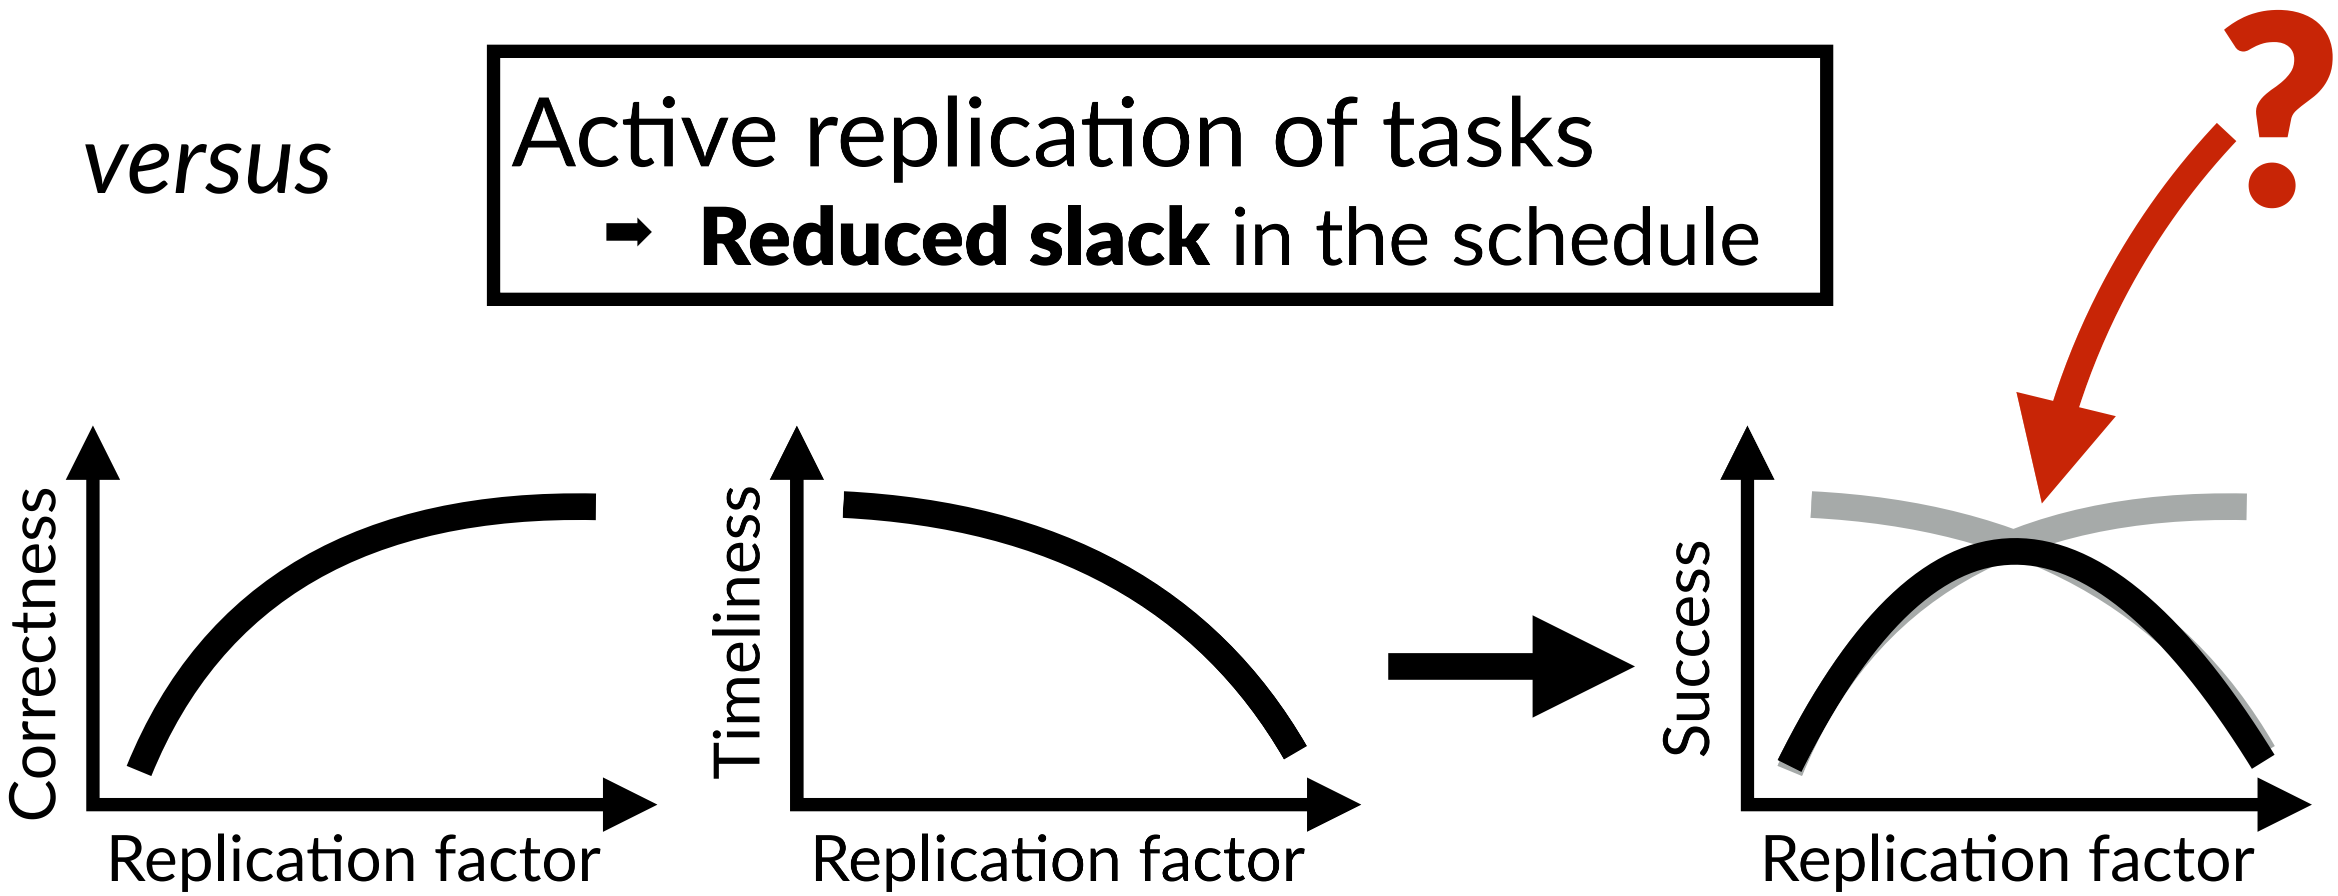

How to statically determine the **optimal replication factor?**

This Work

For CAN-based distributed real-time systems…

- **Probabilistic analysis**
  - ➡ Quantify the replication vs. retransmissions tradeoff

The Larger Picture

The CAN-based system is just one component in a safety-critical system…

The Larger Picture

The CAN-based system is just one component in a safety-critical system…

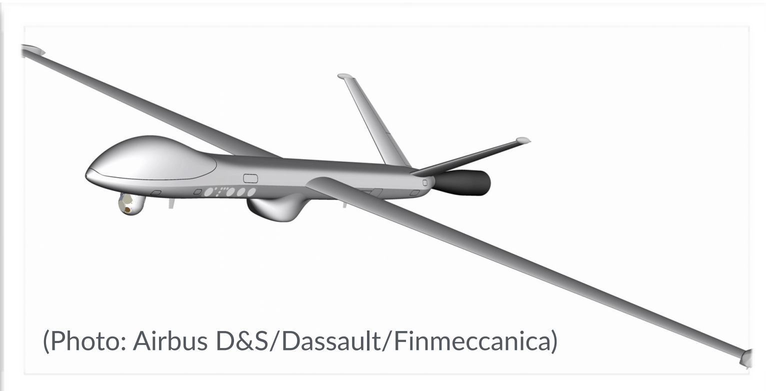

(Photo: Airbus D&S/Dassault/Finmeccanica)

UAV

=

Power

+

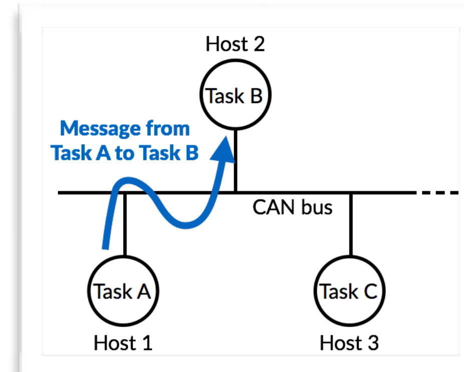

CAN

+ …

The Larger Picture

The CAN-based system is just one component in a safety-critical system…

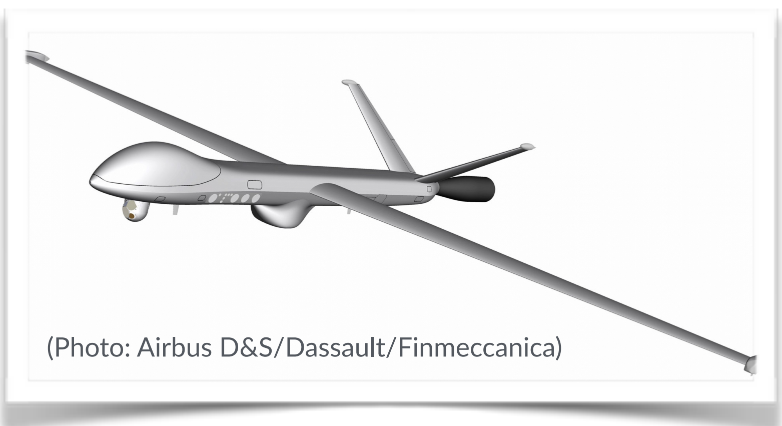

UAV

=

Power

+

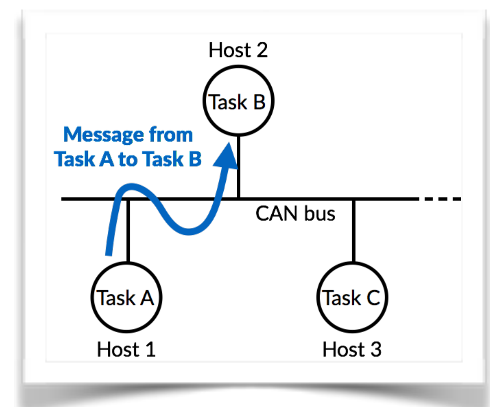

CAN

+ …

Replicate tasks, add more ECUs to the CAN subsystem

The Larger Picture

The CAN-based system is just one component in a safety-critical system…

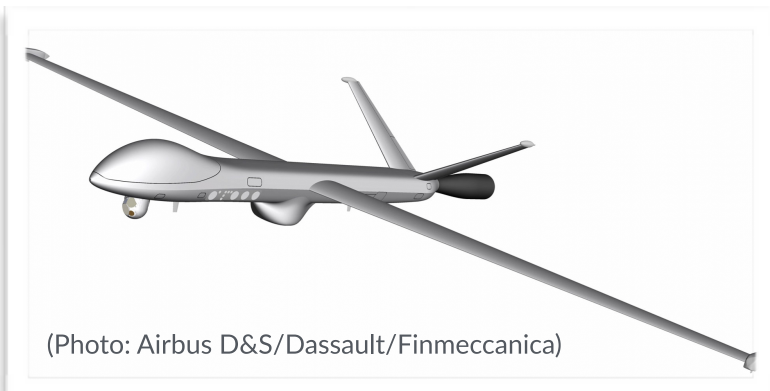
(Photo: Airbus D&S/Dassault/Finmeccanica)

UAV

=

Power

+

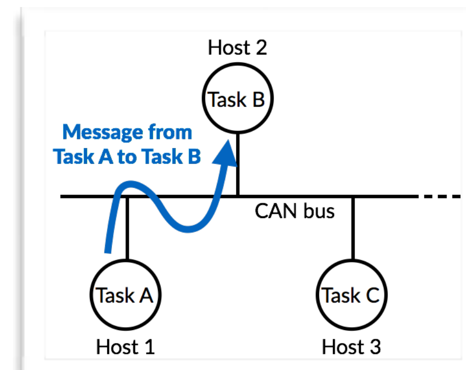

CAN

+ …

Add a heat-sink to the power supply unit

Replicate tasks, add more ECUs to the CAN subsystem

The Larger Picture

The CAN-based system is just one component in a safety-critical system...

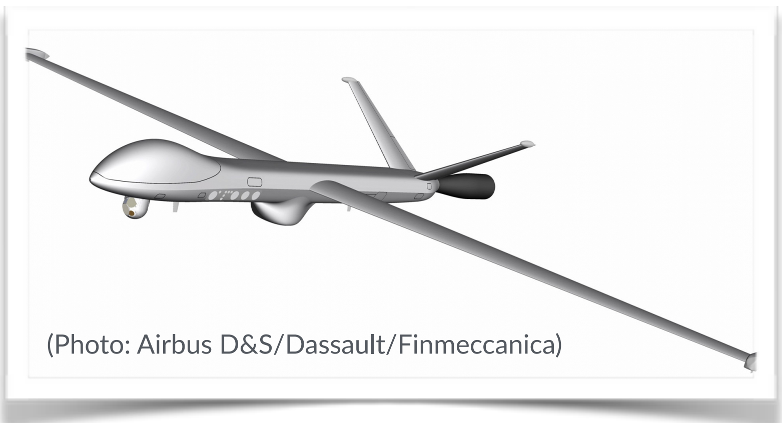

UAV

=

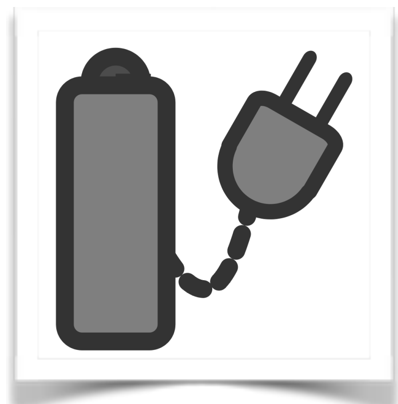

Power

+

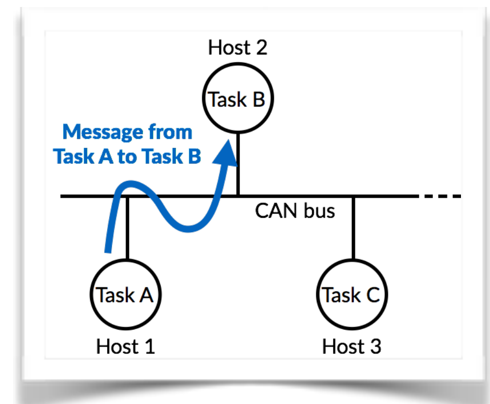

CAN

+ •••

Add a heat-sink to the power supply unit

Replicate tasks, add more ECUs to the CAN subsystem

What if the UAV has **strict weight constraints**?

- and you can either add the heat sink or the additional ECUs
- How do you decide the **best choice**?

Failures-In-Time (FIT) Rate

Expected #failures in one billion operating hours

- ➡ e.g., 1M UAVs flying for 1K hours each

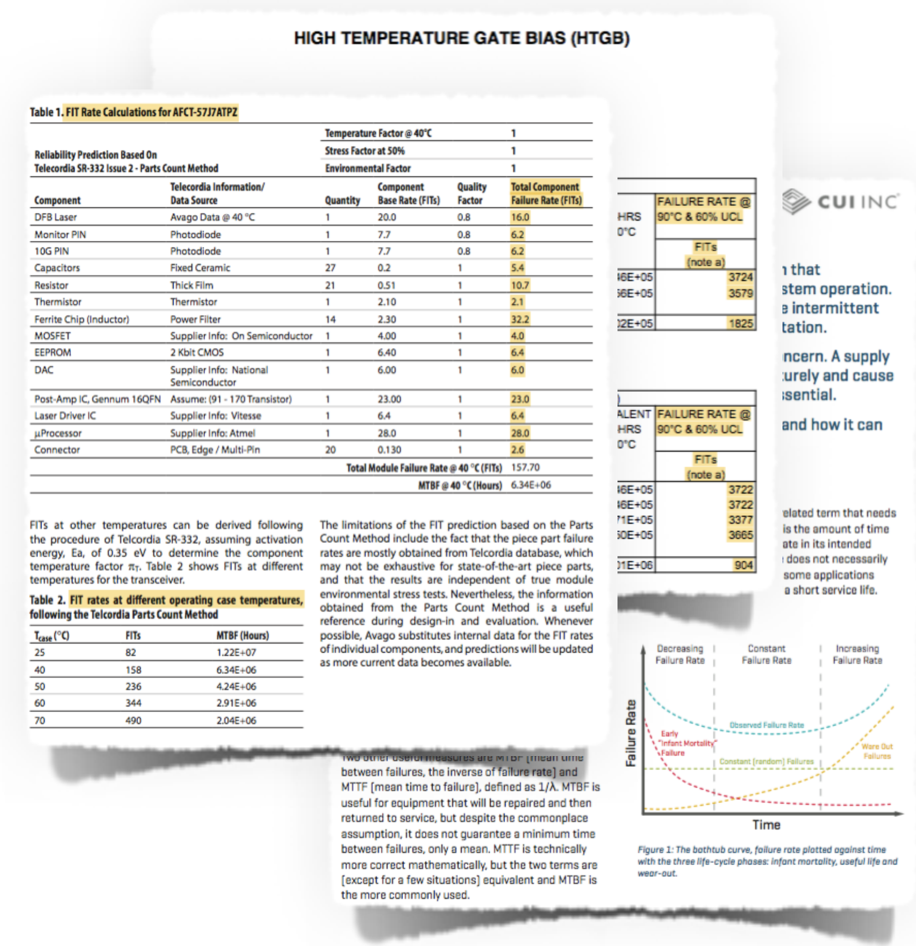

FIT rates are widely used in the industry

Failures-In-Time (FIT) Rate

Expected #failures in one billion operating hours

- ➡ e.g., 1M UAVs flying for 1K hours each

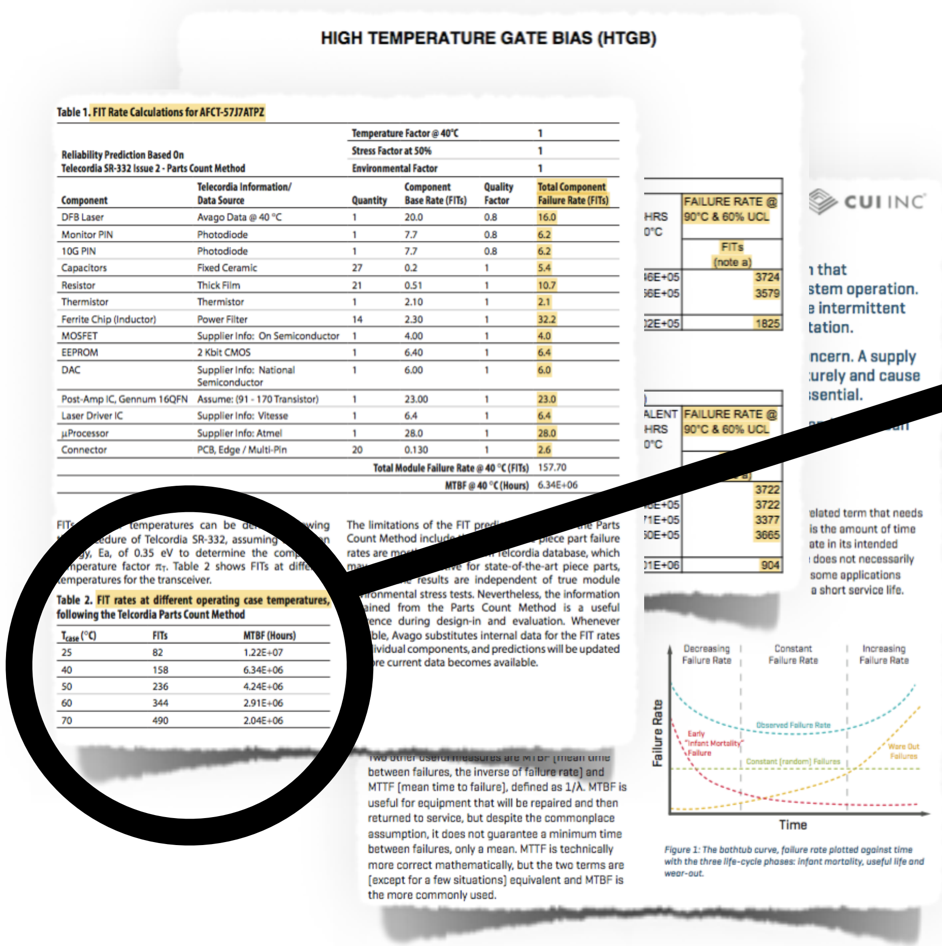

FIT rates are widely used in the industry

Table 2. FIT rates at different operating case temperatures, following the Telcordia Parts Count Method

T_{case} (°C)	FITs	MTBF (Hours)
25	82	1.22E+07
40	158	6.34E+06
50	236	4.24E+06
60	344	2.91E+06
70	490	2.04E+06

Failures-In-Time (FIT) Rate

Expected #failures in one billion operating hours

- e.g., 1M UAVs flying for 1K hours each

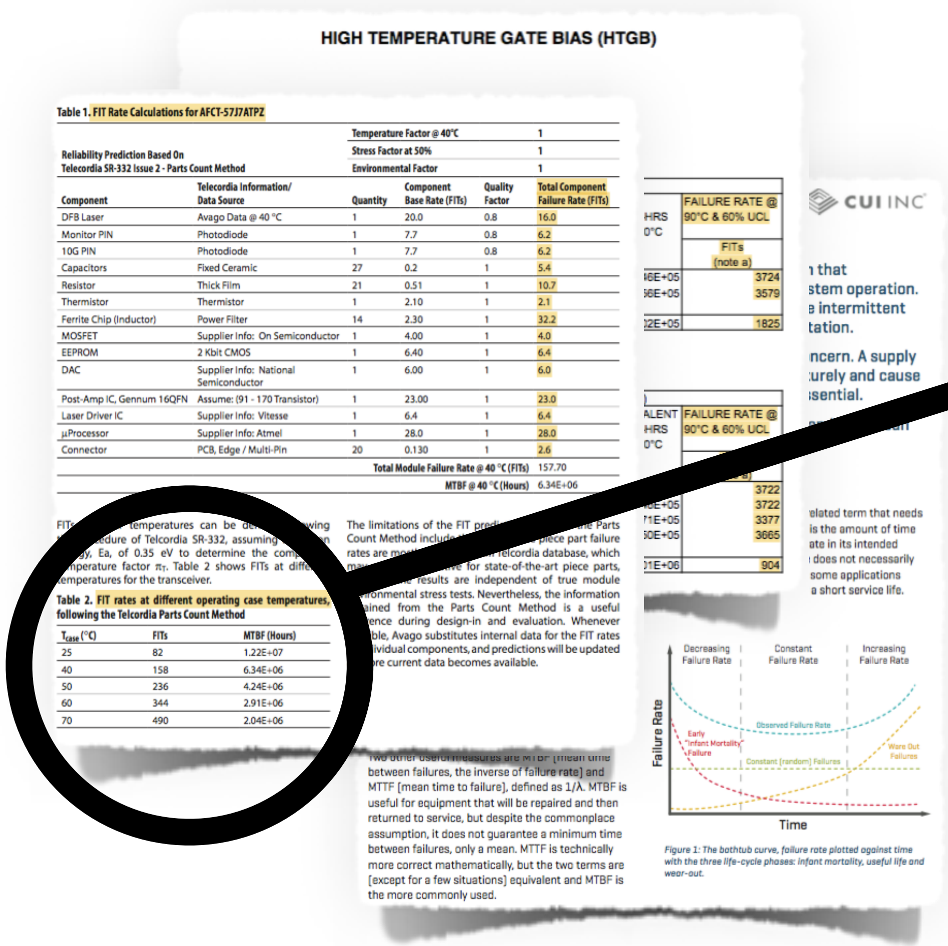

FIT rates are widely used in the industry

Table 2. FIT rates at different operating case temperatures, following the Telcordia Parts Count Method

T_{case} (°C)	FITs	MTBF (Hours)
25	82	1.22E+07
40	158	6.34E+06
50	236	4.24E+06
60	344	2.91E+06
70	490	2.04E+06

When is the CAN-based distributed real-time system the **weakest link in the system**?

This Work

For CAN-based distributed real-time systems…

- **Probabilistic analysis**
  - ➡ Quantify the replication vs. retransmissions tradeoff

- **FIT rate analysis**
  - ➡ Builds upon the proposed probabilistic analysis

Overview

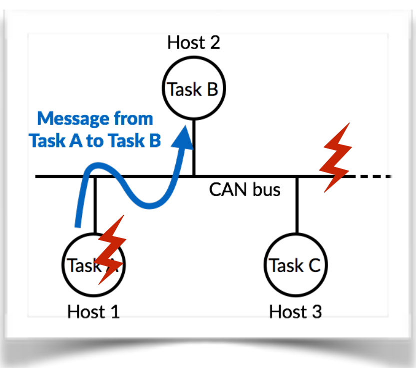

Model

$$\sum_{\mathbb{H}'\subseteq\mathbb{H}} \Phi^{\mathbb{H}'}_{crash} \cdot \sum_{\mathbb{M}'_1\subseteq\mathbb{M}_1} \left(\Phi^{\mathbb{H}',\,\mathbb{M}'_1}_{timely} \cdot \Phi^{\mathbb{H}',\,\mathbb{M}'_1}_{correct}\right)$$

Analysis

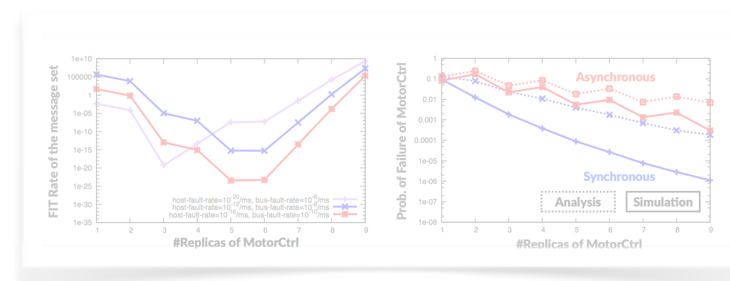

Evaluation

Fault Abstraction & Modeling

Transmission failures
(faults on the wire)

Commission failures
(bit-flips in the memory buffers)

Crash failures
(due to fault-induced exceptions)

Fault Abstraction & Modeling

Transmission failures
(faults on the wire)

Commission failures
(bit-flips in the memory buffers)

Crash failures
(due to fault-induced exceptions)

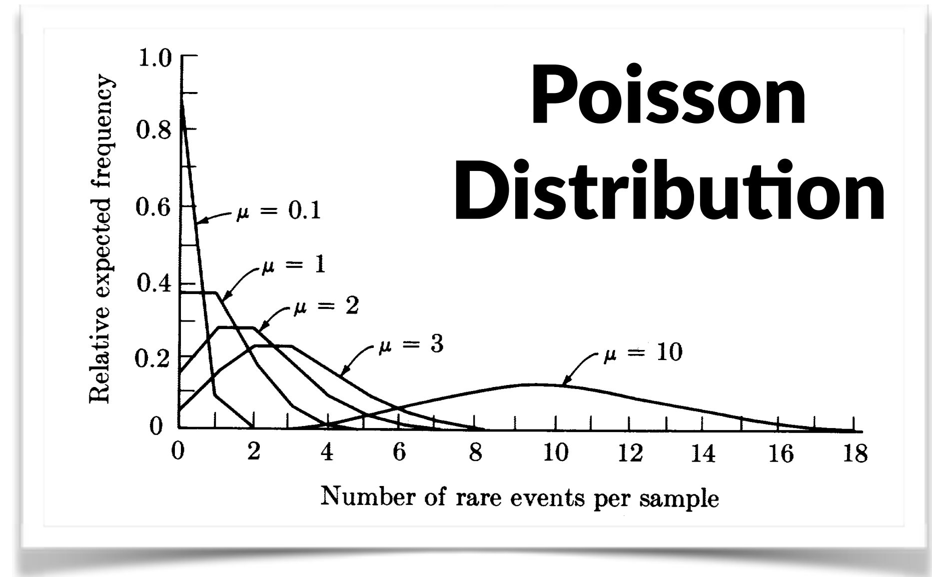

Fault Abstraction & Modeling

Transmission failures
(faults on the wire)

Commission failures
(bit-flips in the memory buffers)

Crash failures
(due to fault-induced exceptions)

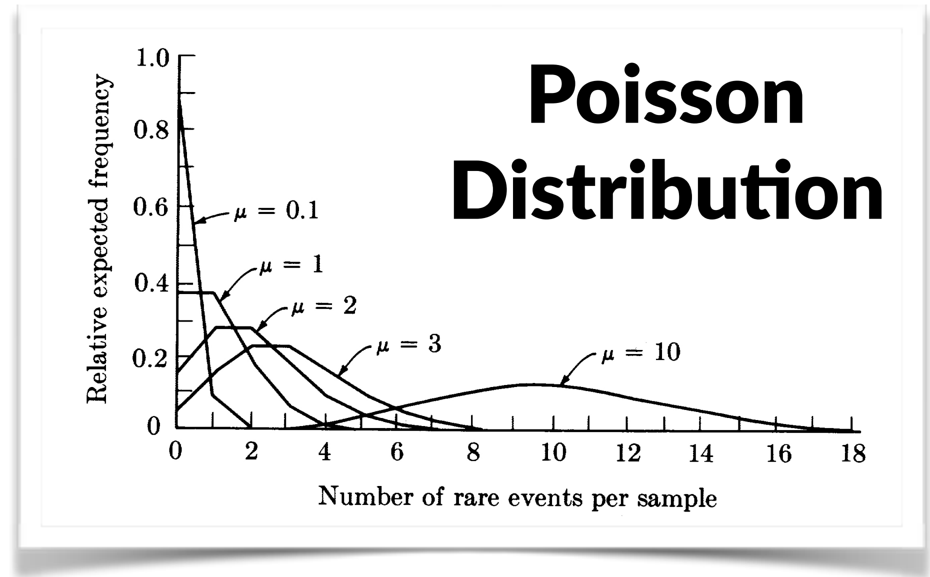

Probability that each message is omitted / corrupted / retransmitted

Fault Abstraction & Modeling

Transmission failures
(faults on the wire)

Commission failures
(bit-flips in the memory buffers)

Crash failures
(due to fault-induced exceptions)

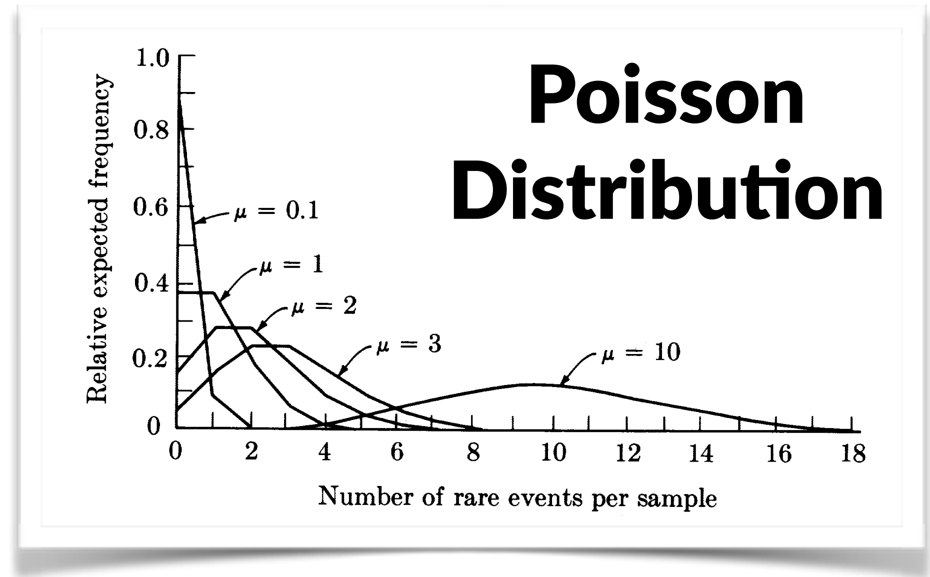

Probability that each message is omitted / corrupted / retransmitted

We do not consider **software defects**…

System Model

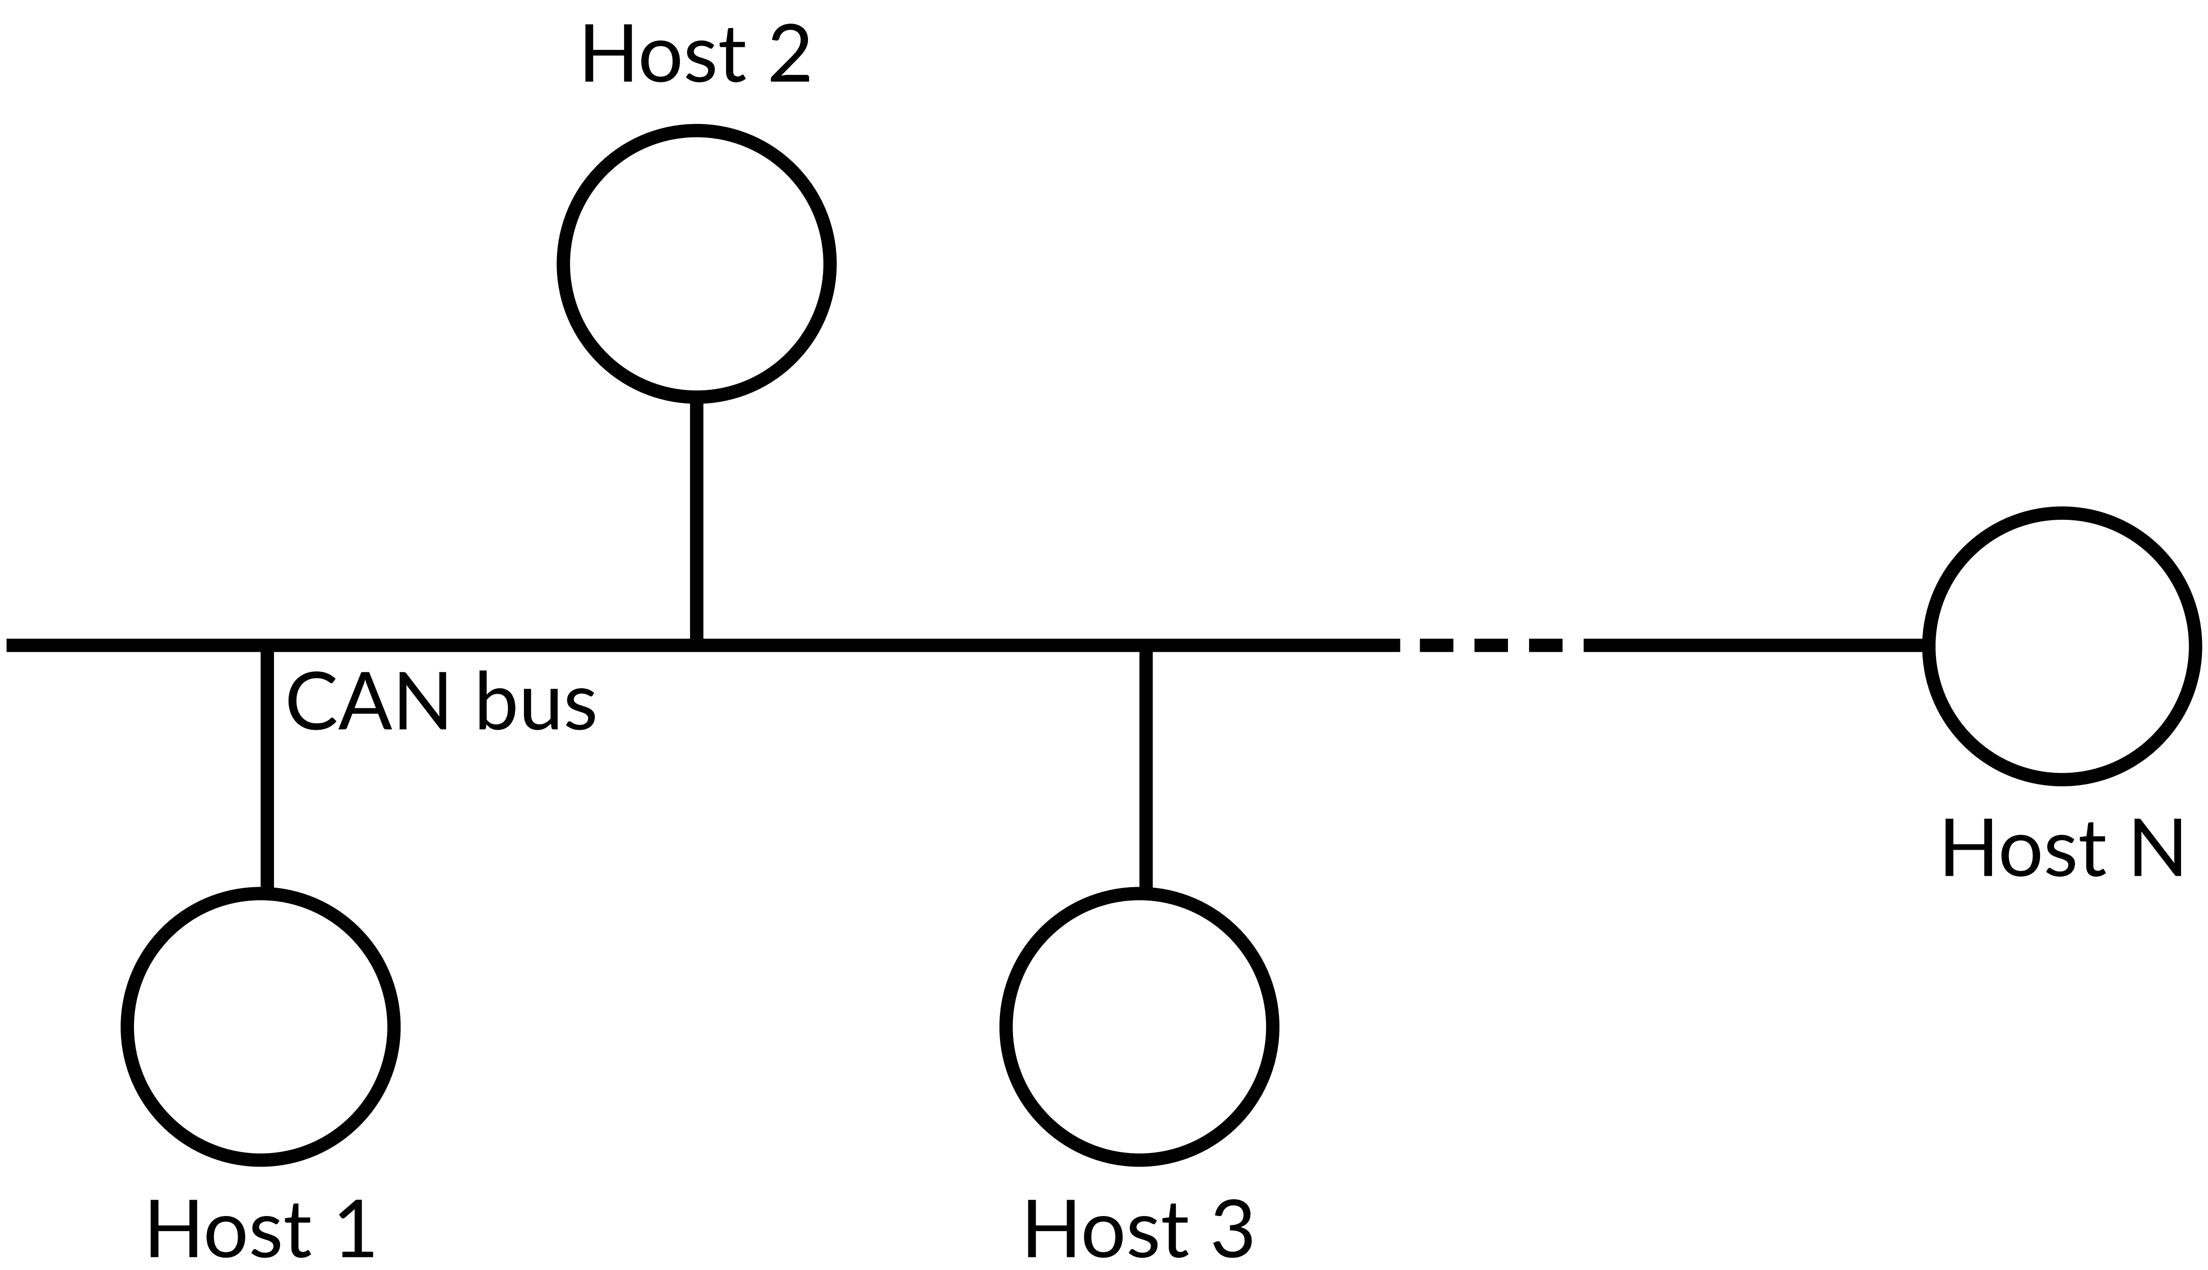

System Model

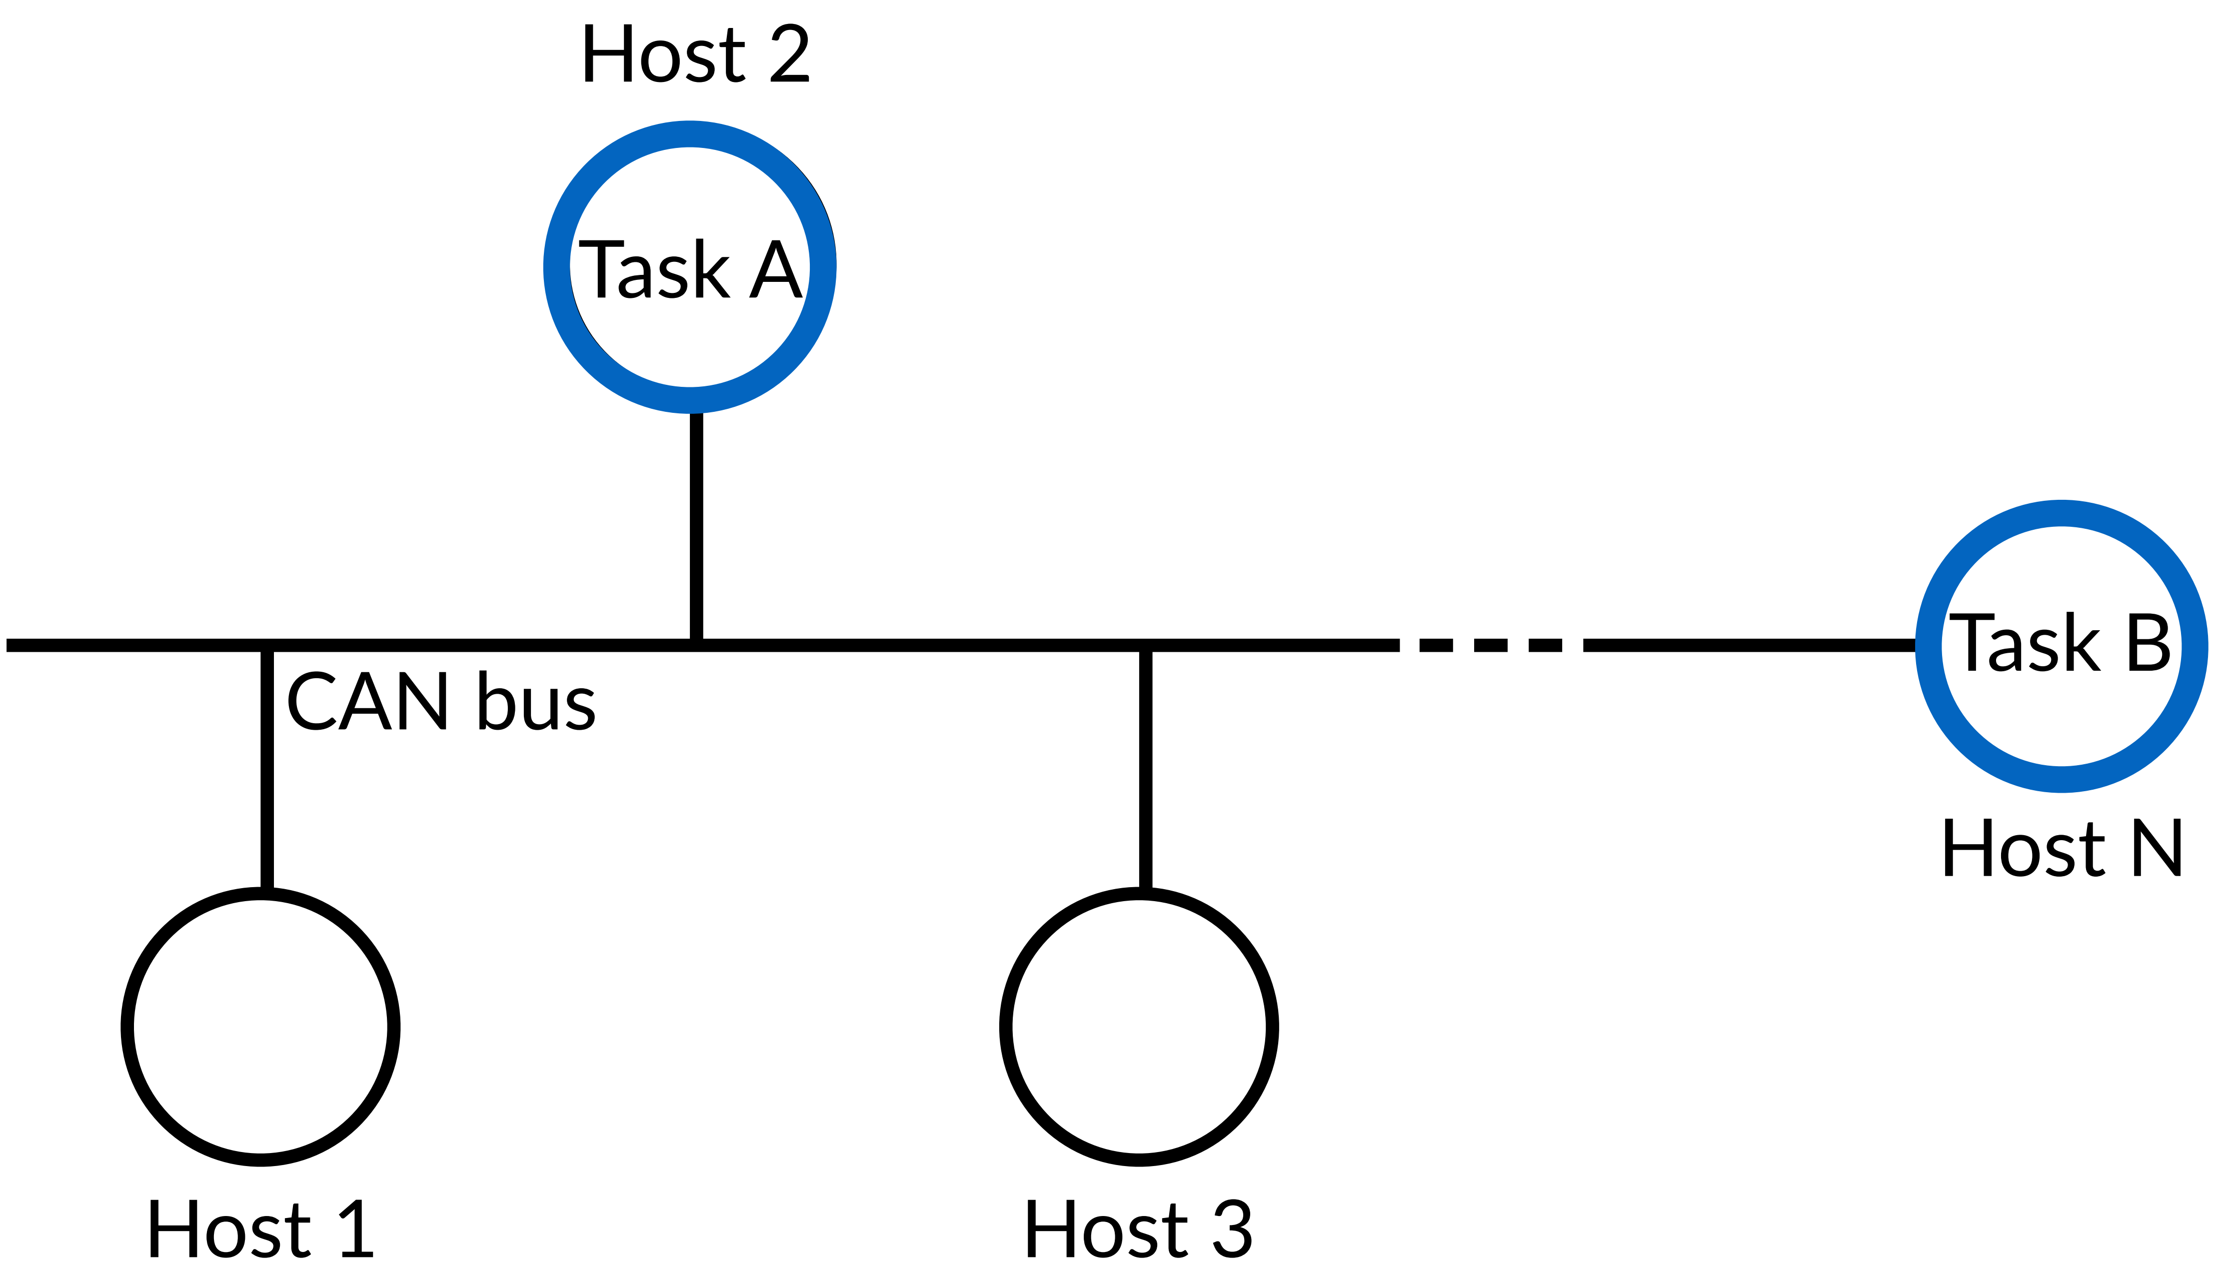

System Model

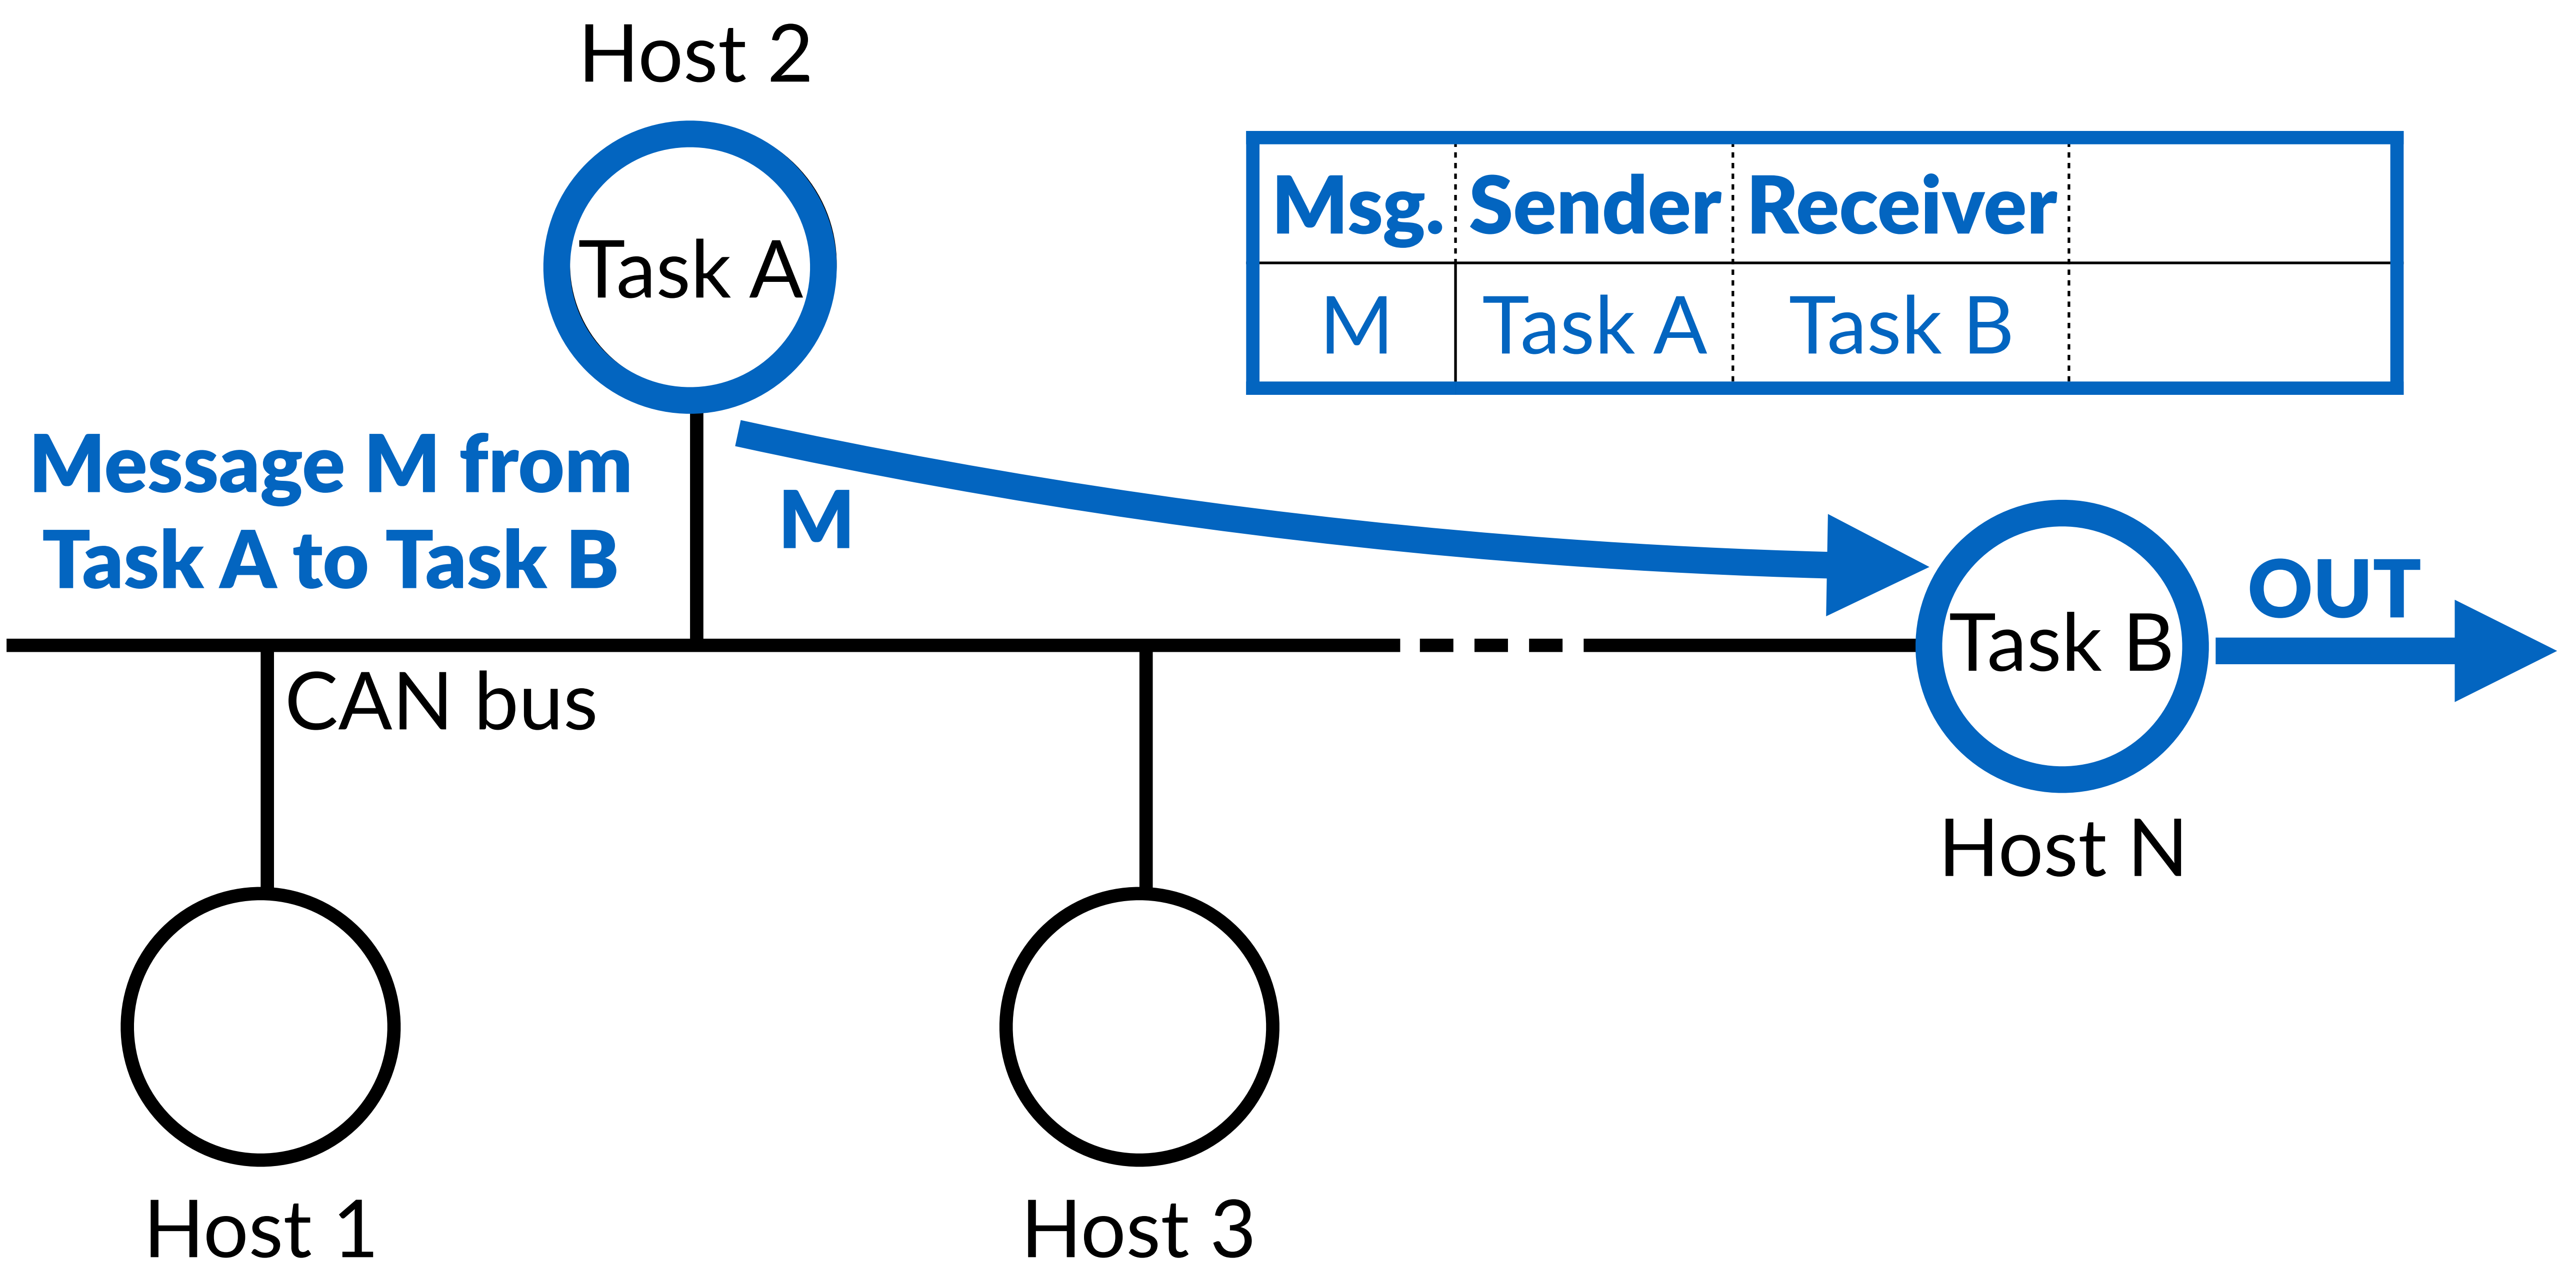

Msg.	Sender	Receiver	
M	Task A	Task B	

System Model

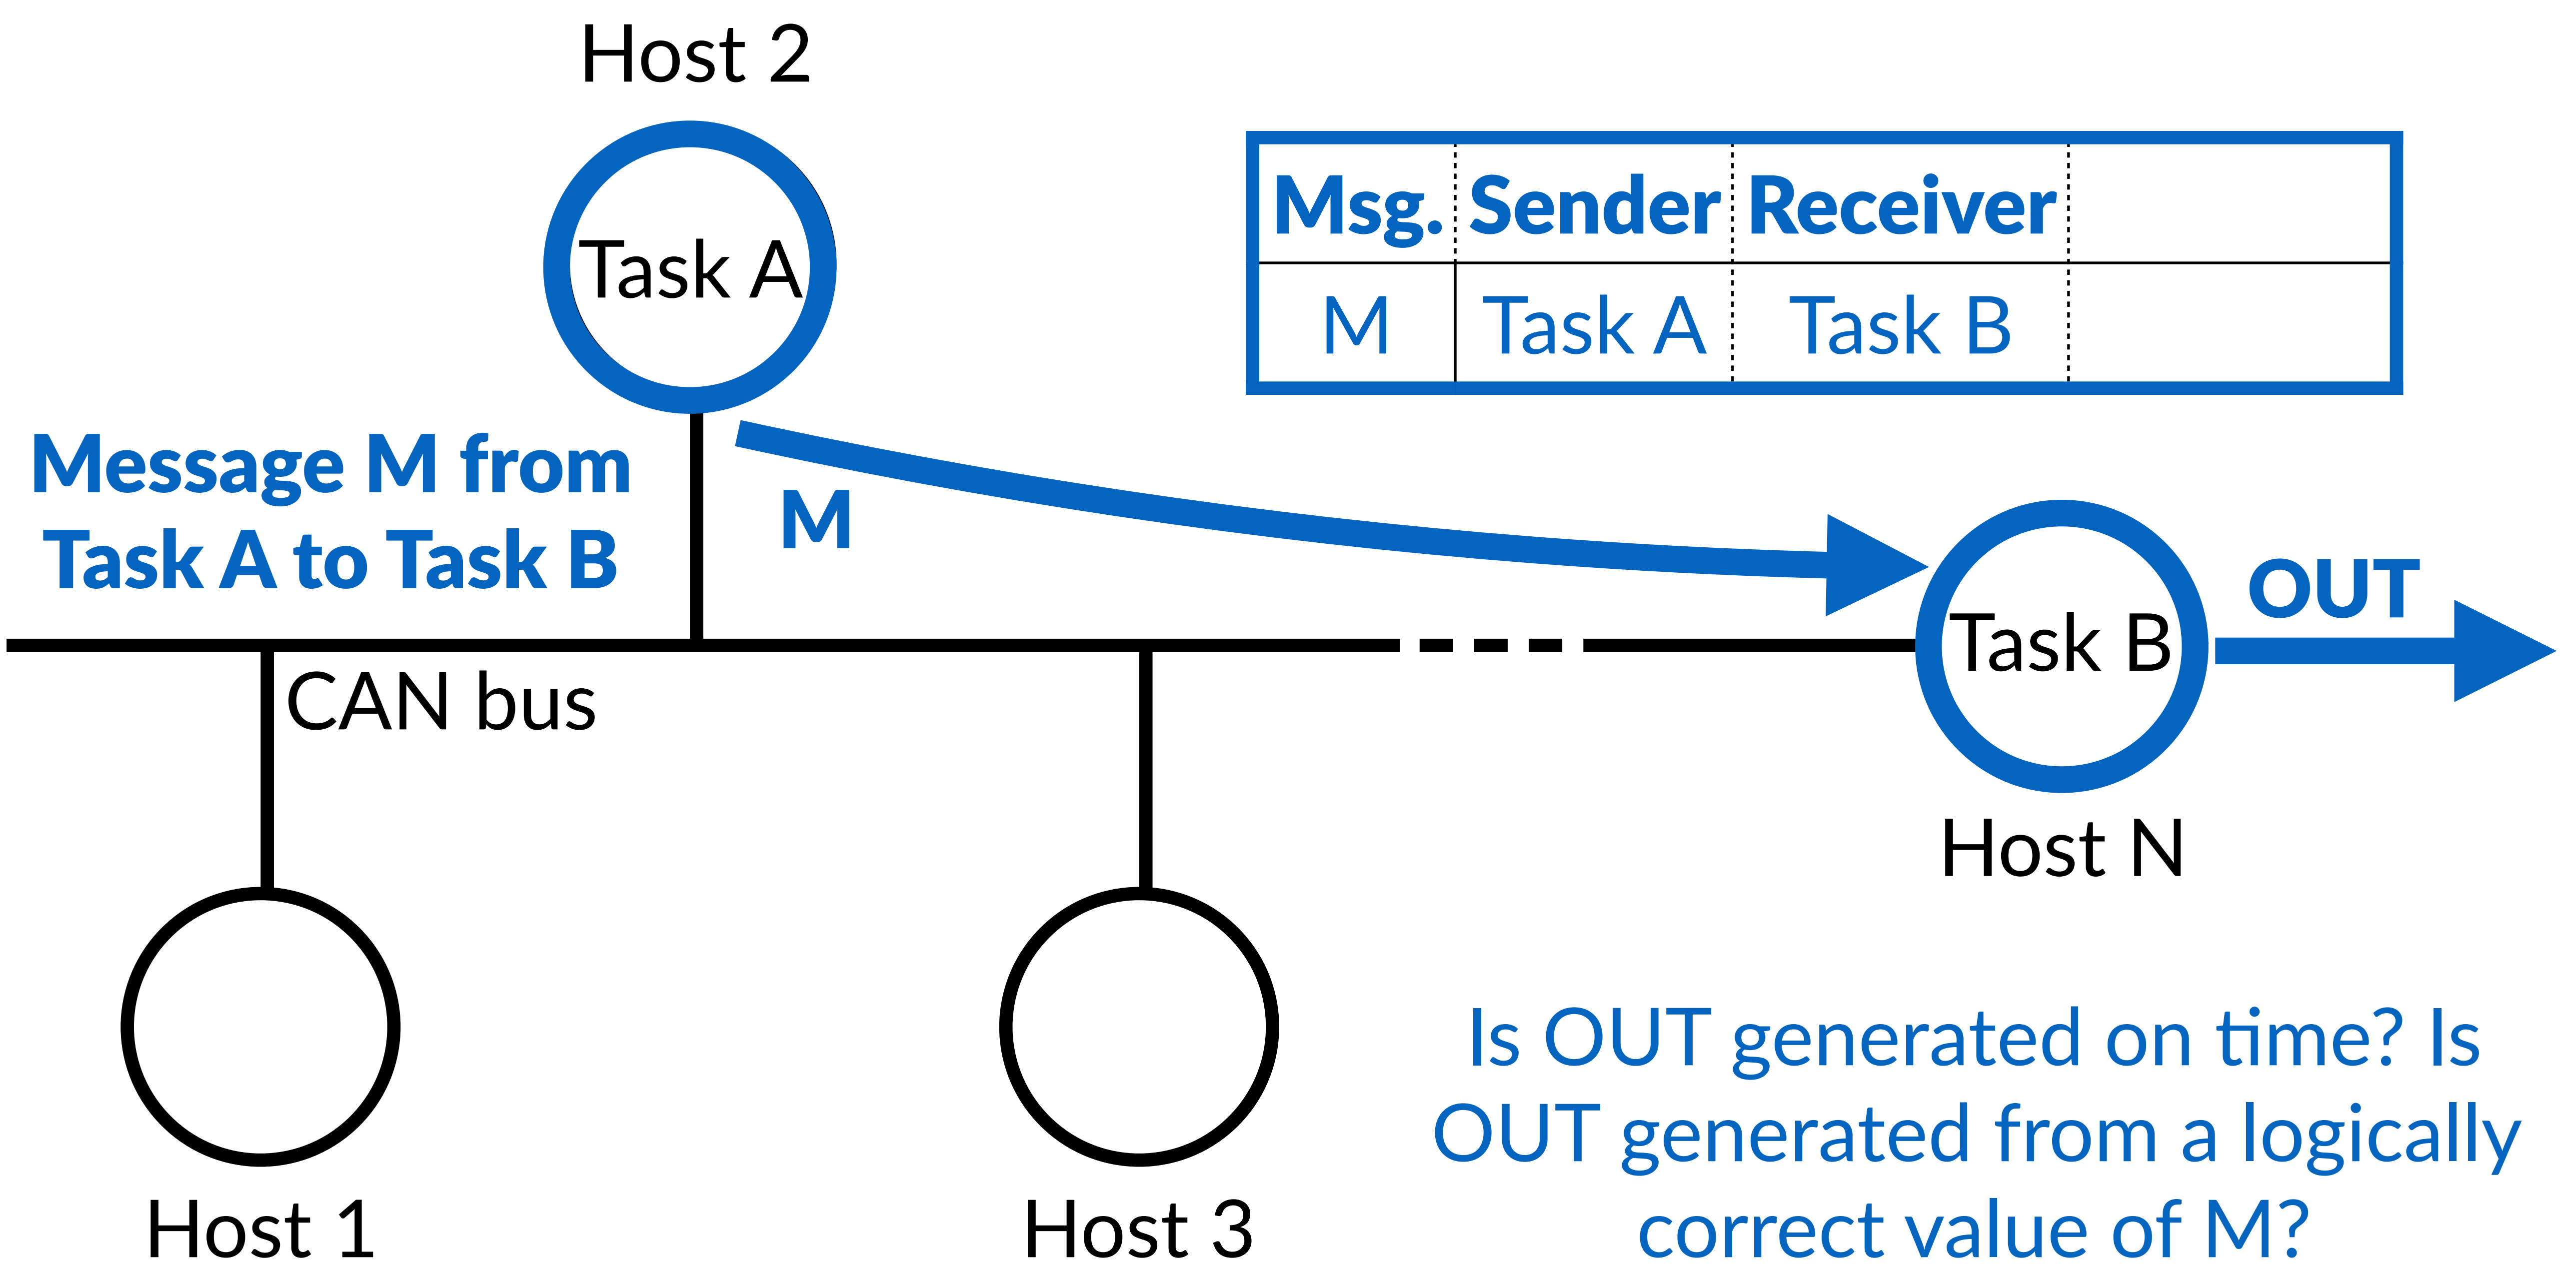

System Model

with Task Replication

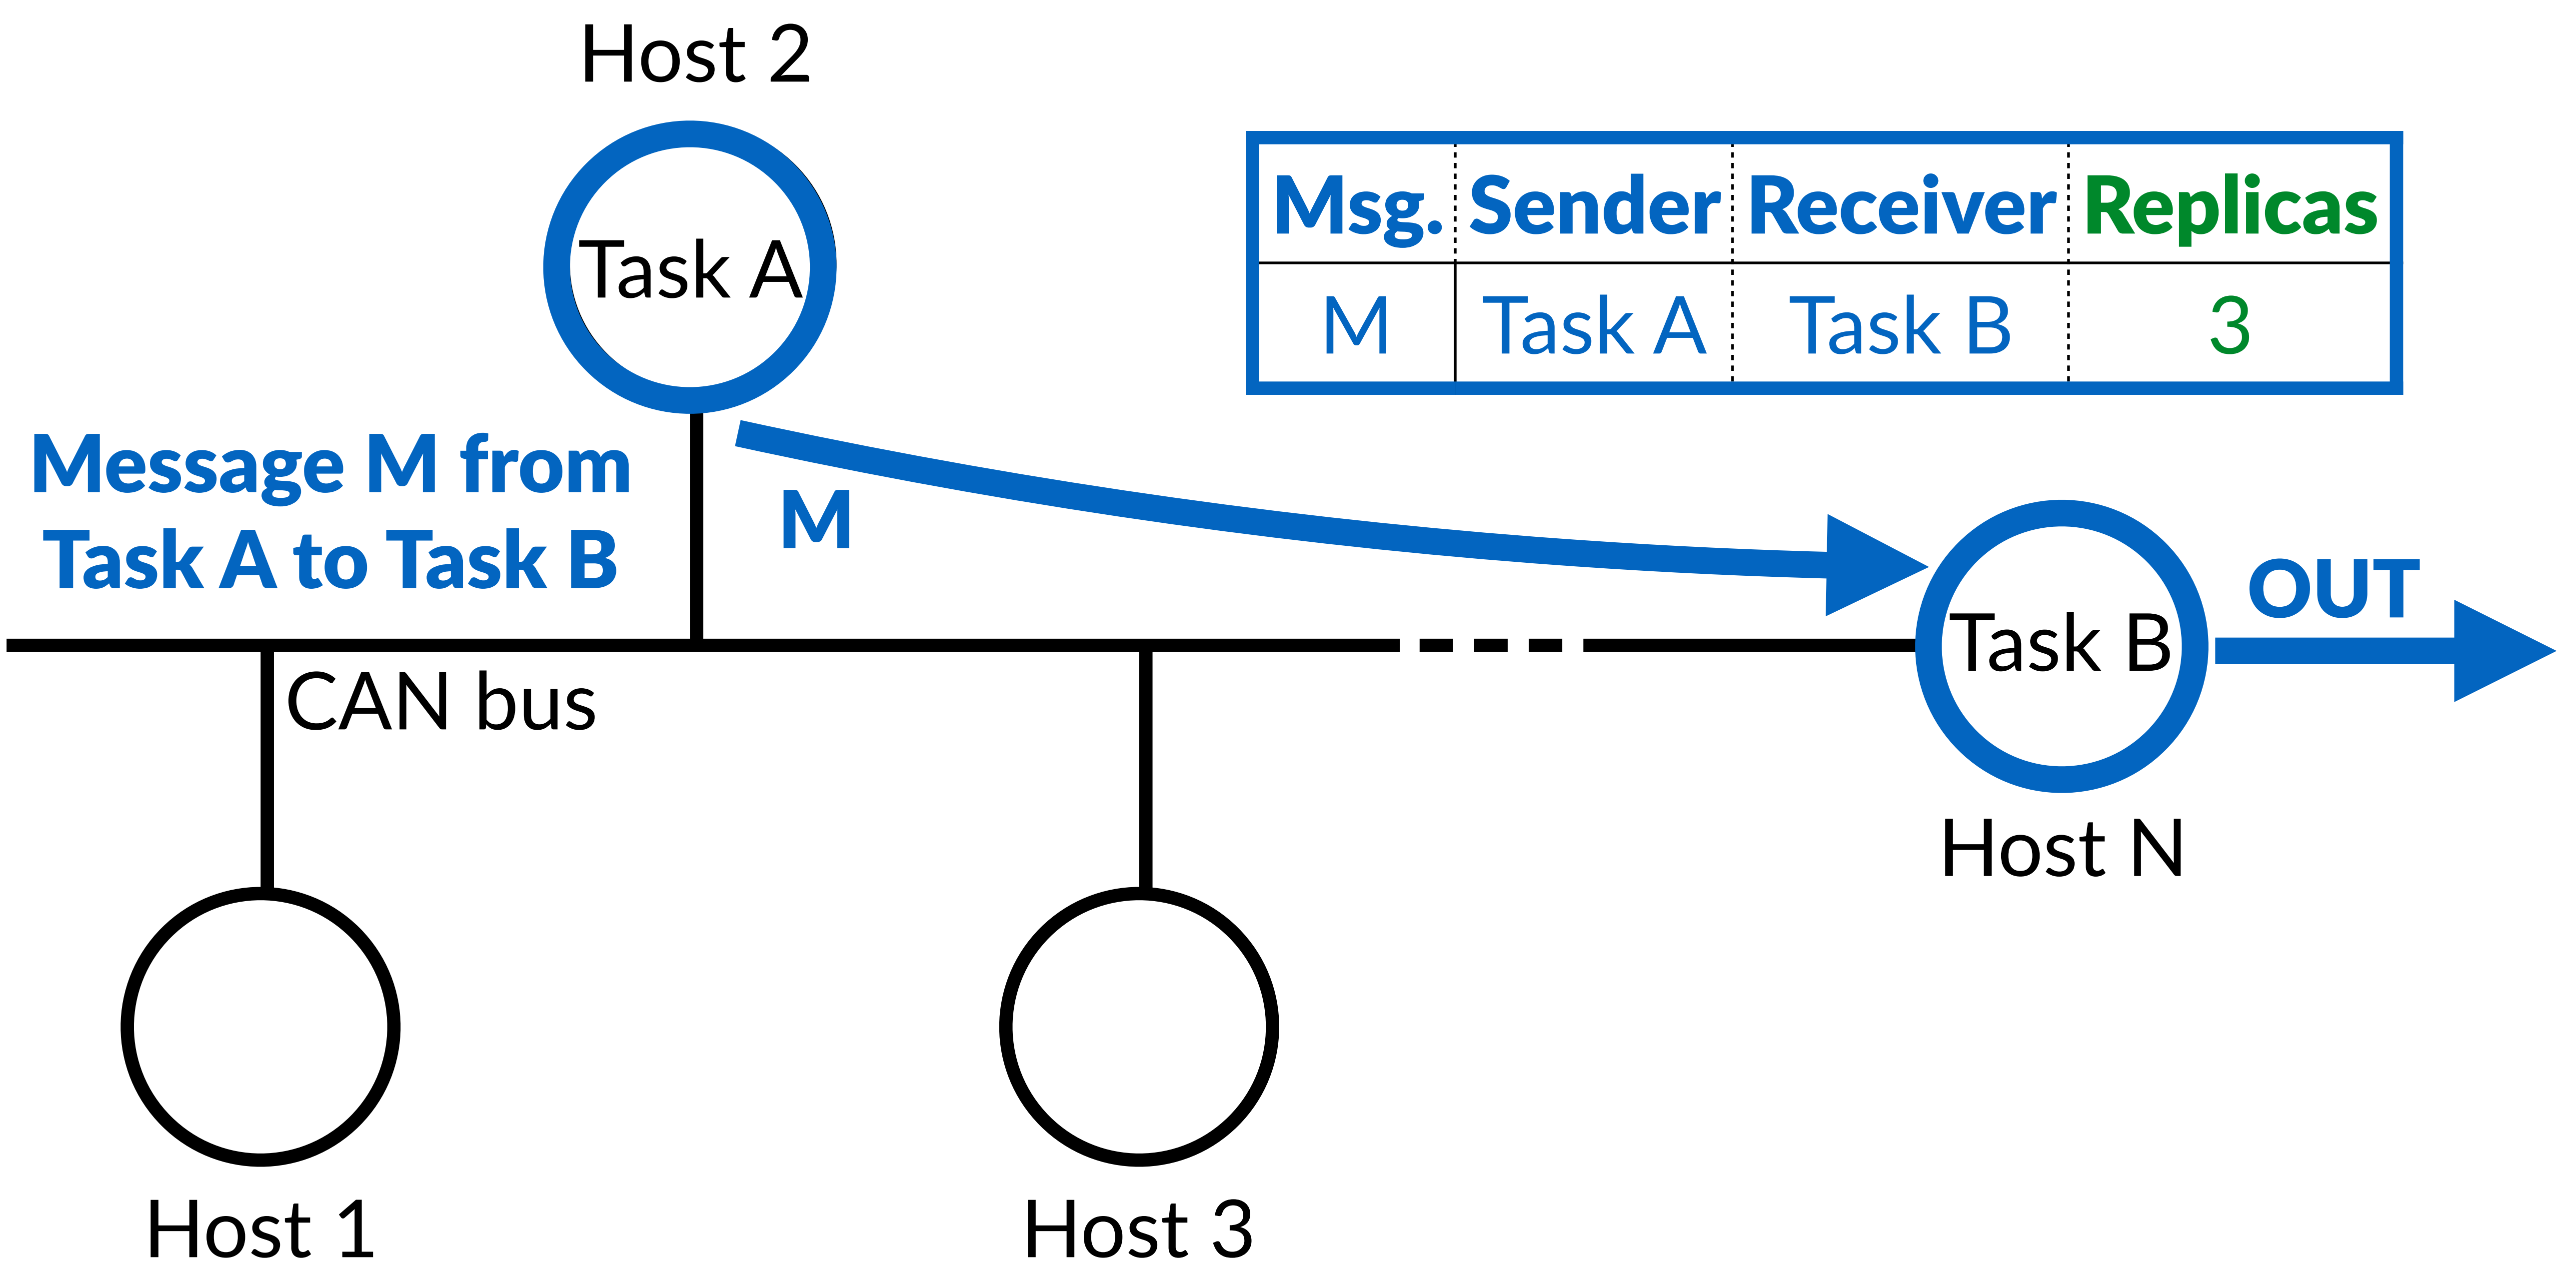

System Model

with Task Replication

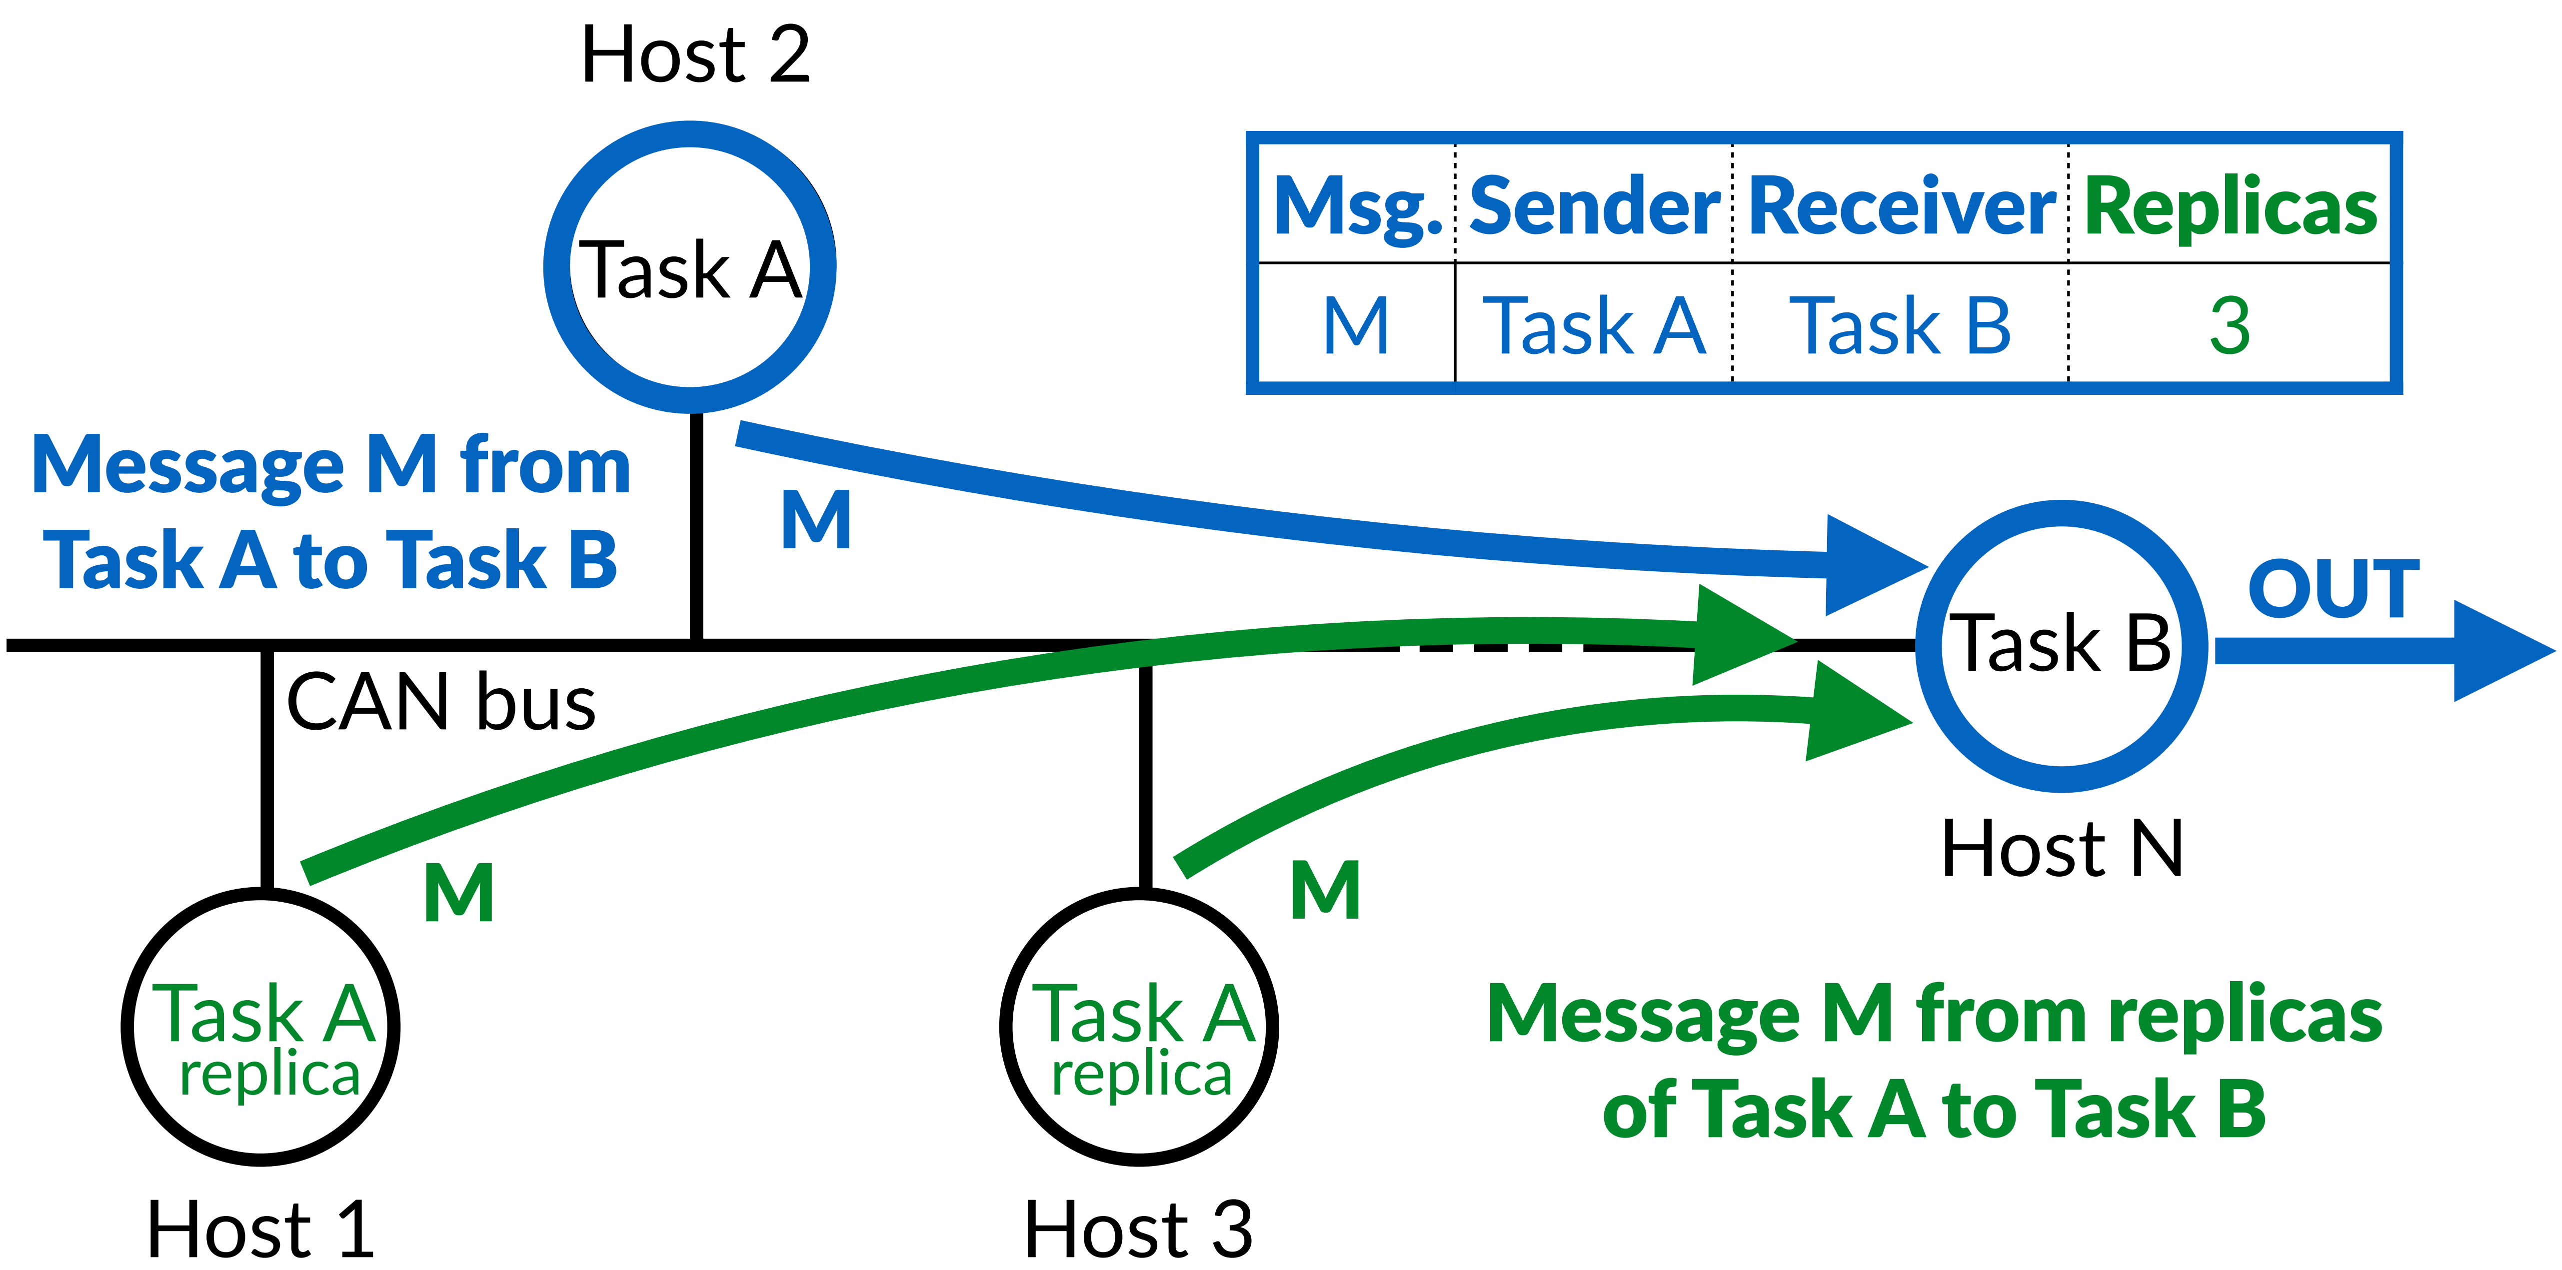

Aggregating the replicated messages

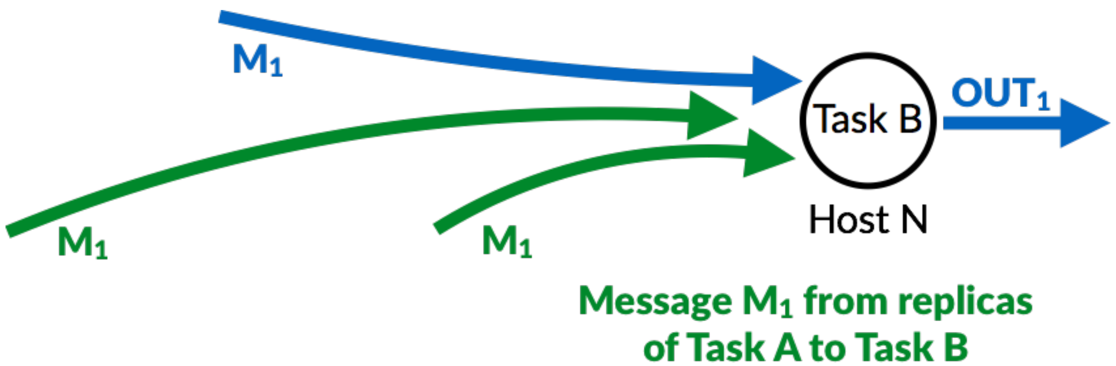

How & when to compute OUT
from multiple copies of M?

Aggregating the replicated messages

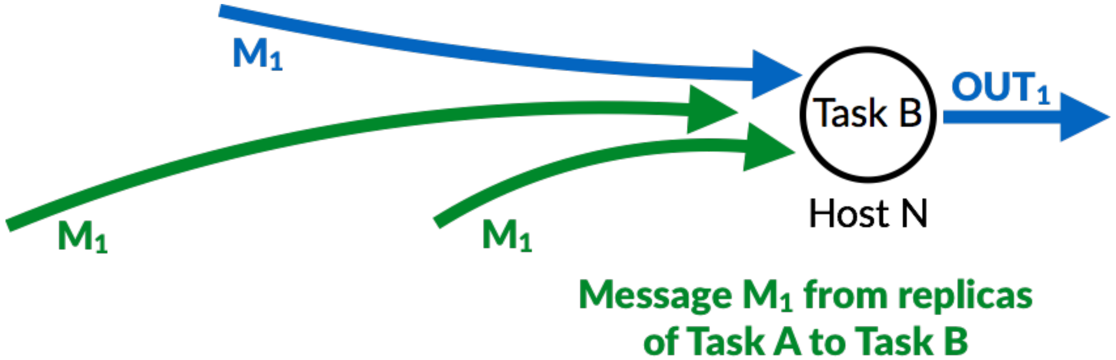

How & when to compute OUT
from multiple copies of M?

- Case 1: Synchronous Systems
  - ➡ **Common** global time base
  - ➡ e.g. majority value **at the absolute deadline**

Aggregating the replicated messages

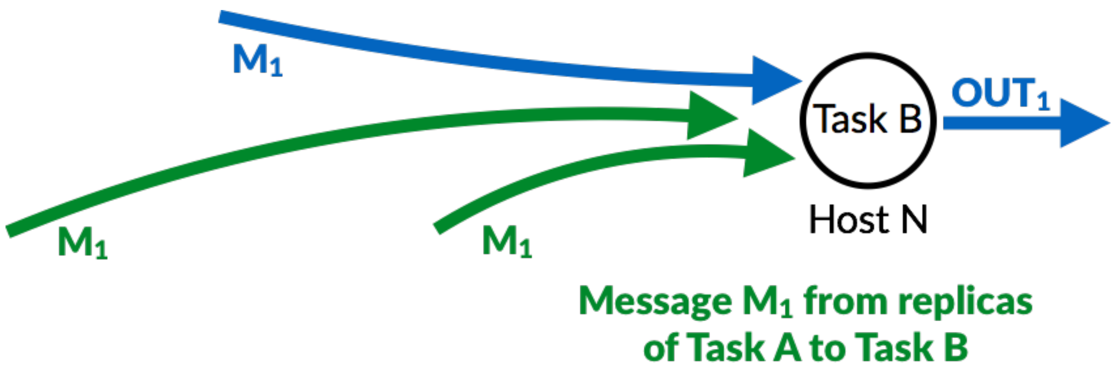

How & when to compute OUT
from multiple copies of M?

- Case 1: Synchronous Systems
  - ➡ **Common** global time base
  - ➡ e.g. majority value **at the absolute deadline**
- Case 2: Asynchronous Systems
  - ➡ **No** global time base
  - ➡ e.g. majority value **after “enough” copies have been received**

Overview

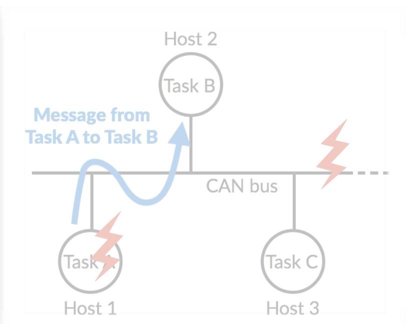

Model

$$\sum_{\mathbb{H}'\subseteq\mathbb{H}} \Phi^{\mathbb{H}'}_{crash}\cdot \sum_{\mathbb{M}_1'\subseteq\mathbb{M}_1}\left(\Phi^{\mathbb{H}',\,\mathbb{M}_1'}_{timely}\cdot\Phi^{\mathbb{H}',\,\mathbb{M}_1'}_{correct}\right)$$

Analysis

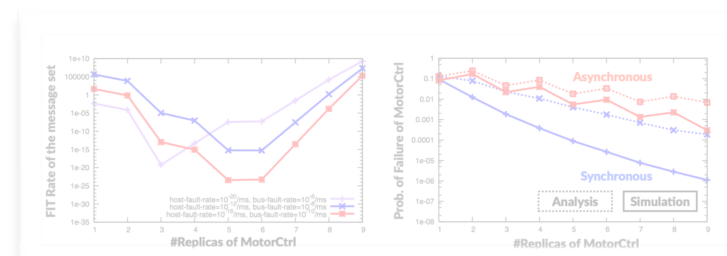

Evaluation

The Larger Picture…

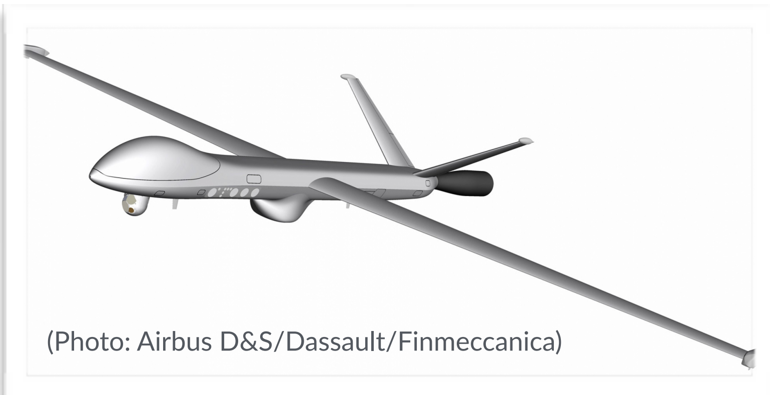
(Photo: Airbus D&S/Dassault/Finmeccanica)

=

+

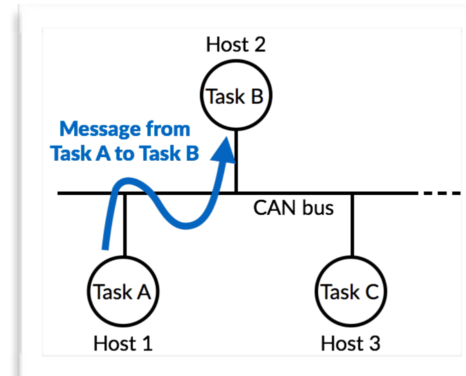

+ …

The Larger Picture…

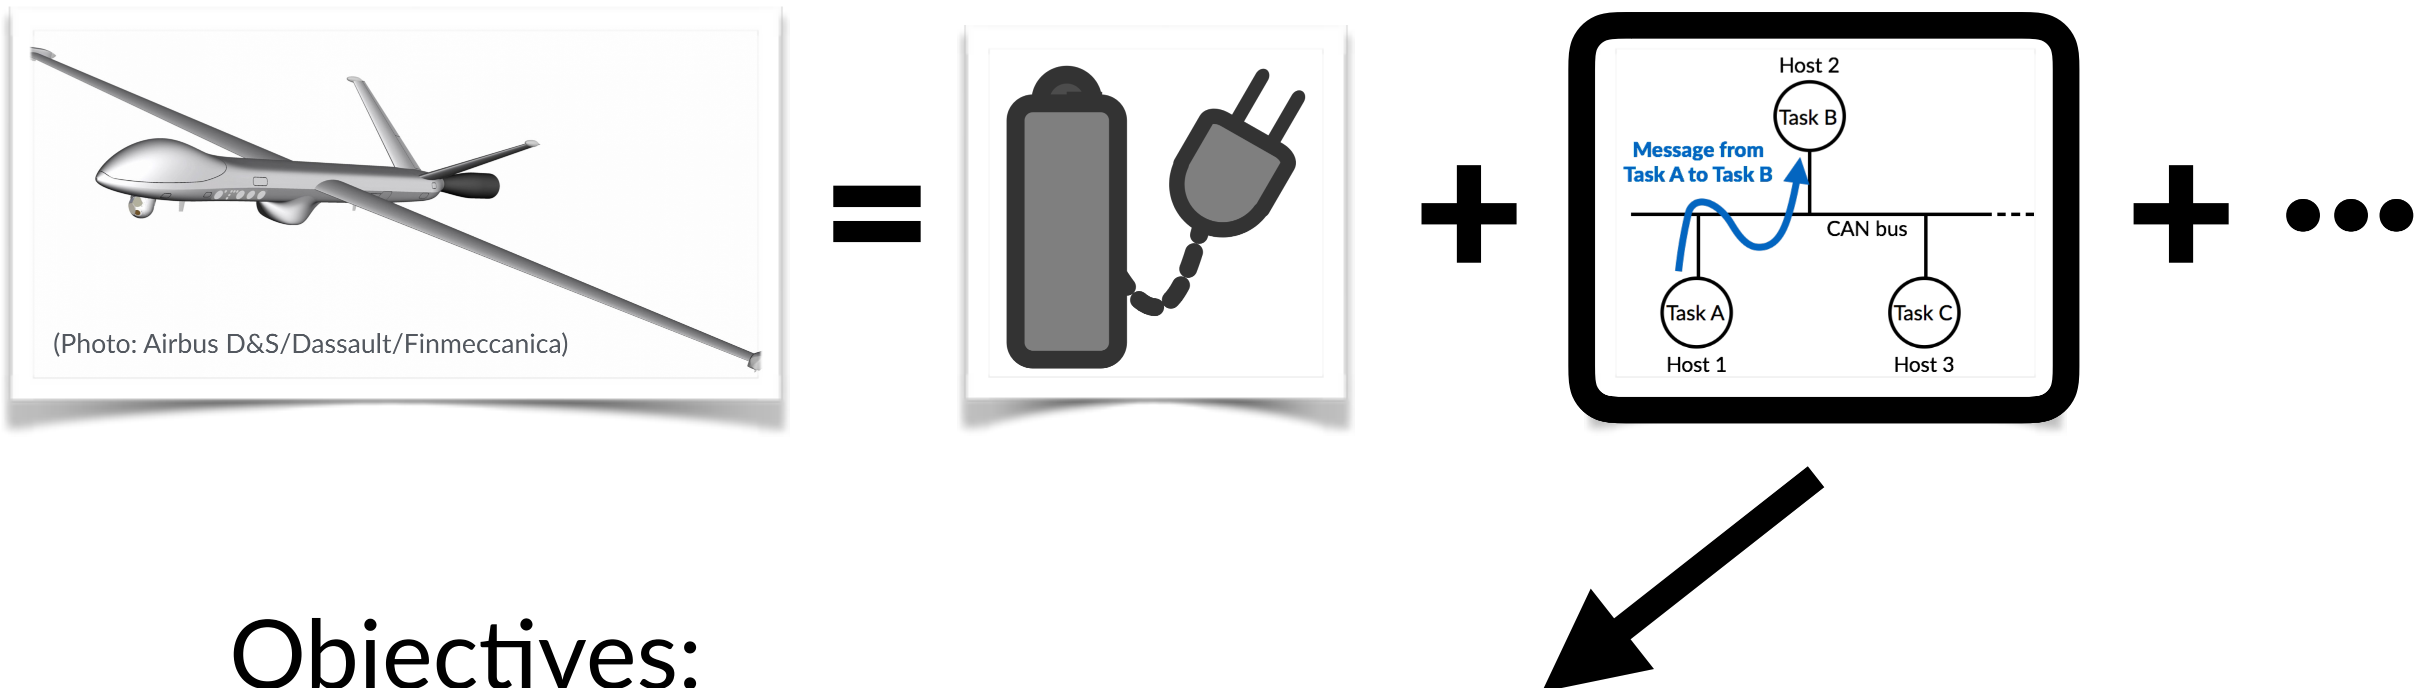

Objectives:

- ➡ **A good replication strategy** for the CAN-based system
- ➡ **Compare** the reliability of the CAN-based system **with other components in the safety-critical system**

The Larger Picture…

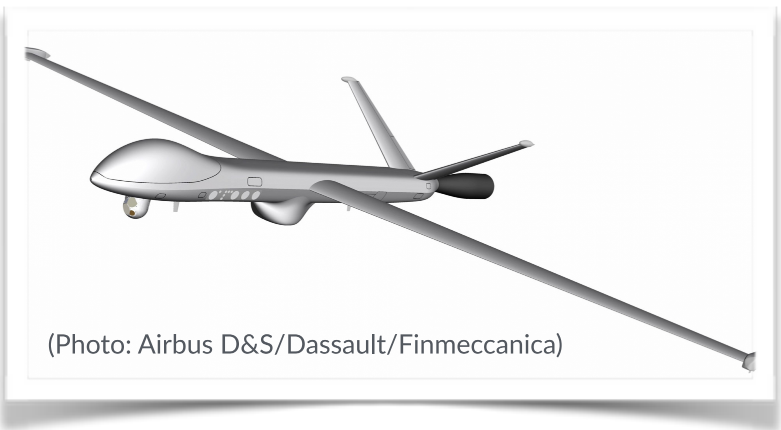

=

+

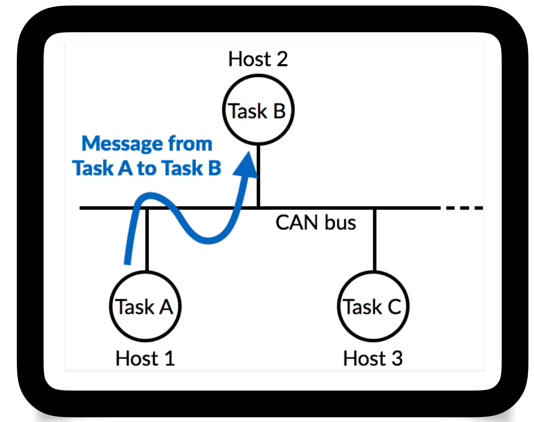

+ •••

Objectives:

- ➡ **A good replication strategy** for the CAN-based system
- ➡ **Compare** the reliability of the CAN-based system **with other components in the safety-critical system**

Solution: **FIT rate analysis**

- ➡ Using the **probabilistic analysis**

FIT Rate Analysis of the System

FIT rate of the CAN subsystem

FIT Rate Analysis of the System

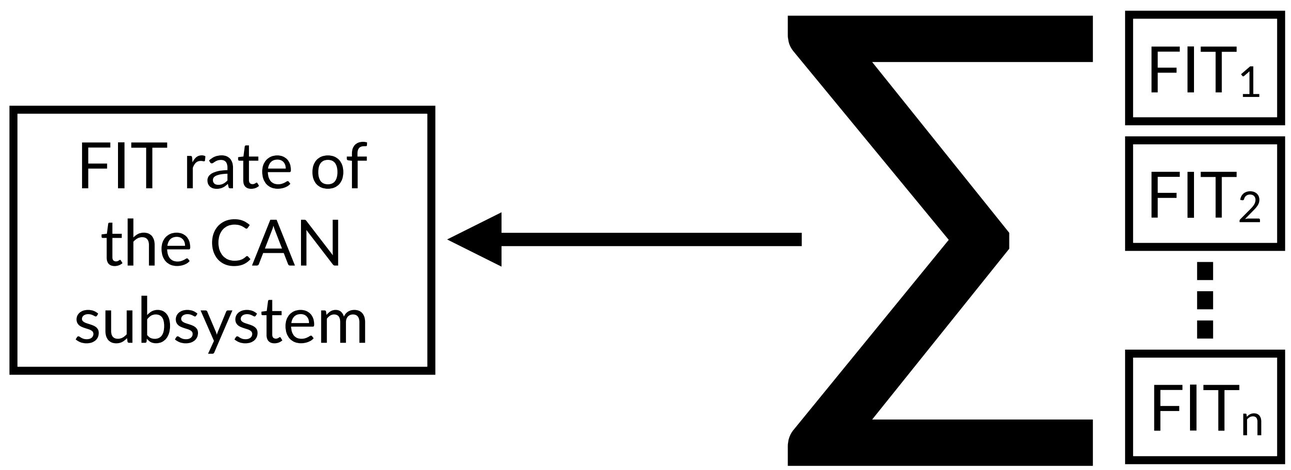

FIT Rate Analysis of the System

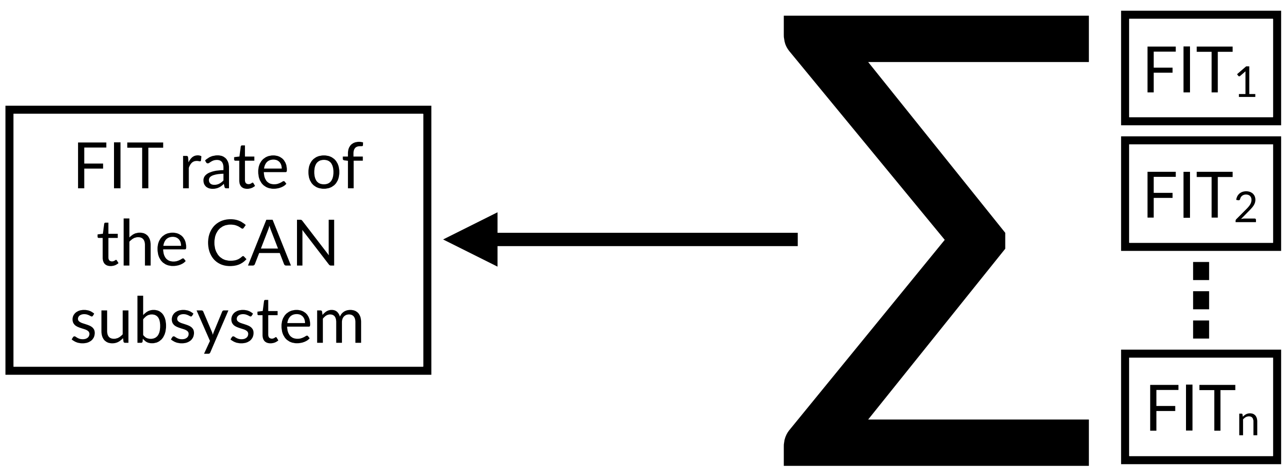

Lower Bound on the Probability of Successful Transmission of M_1

FIT Rate Analysis of the System

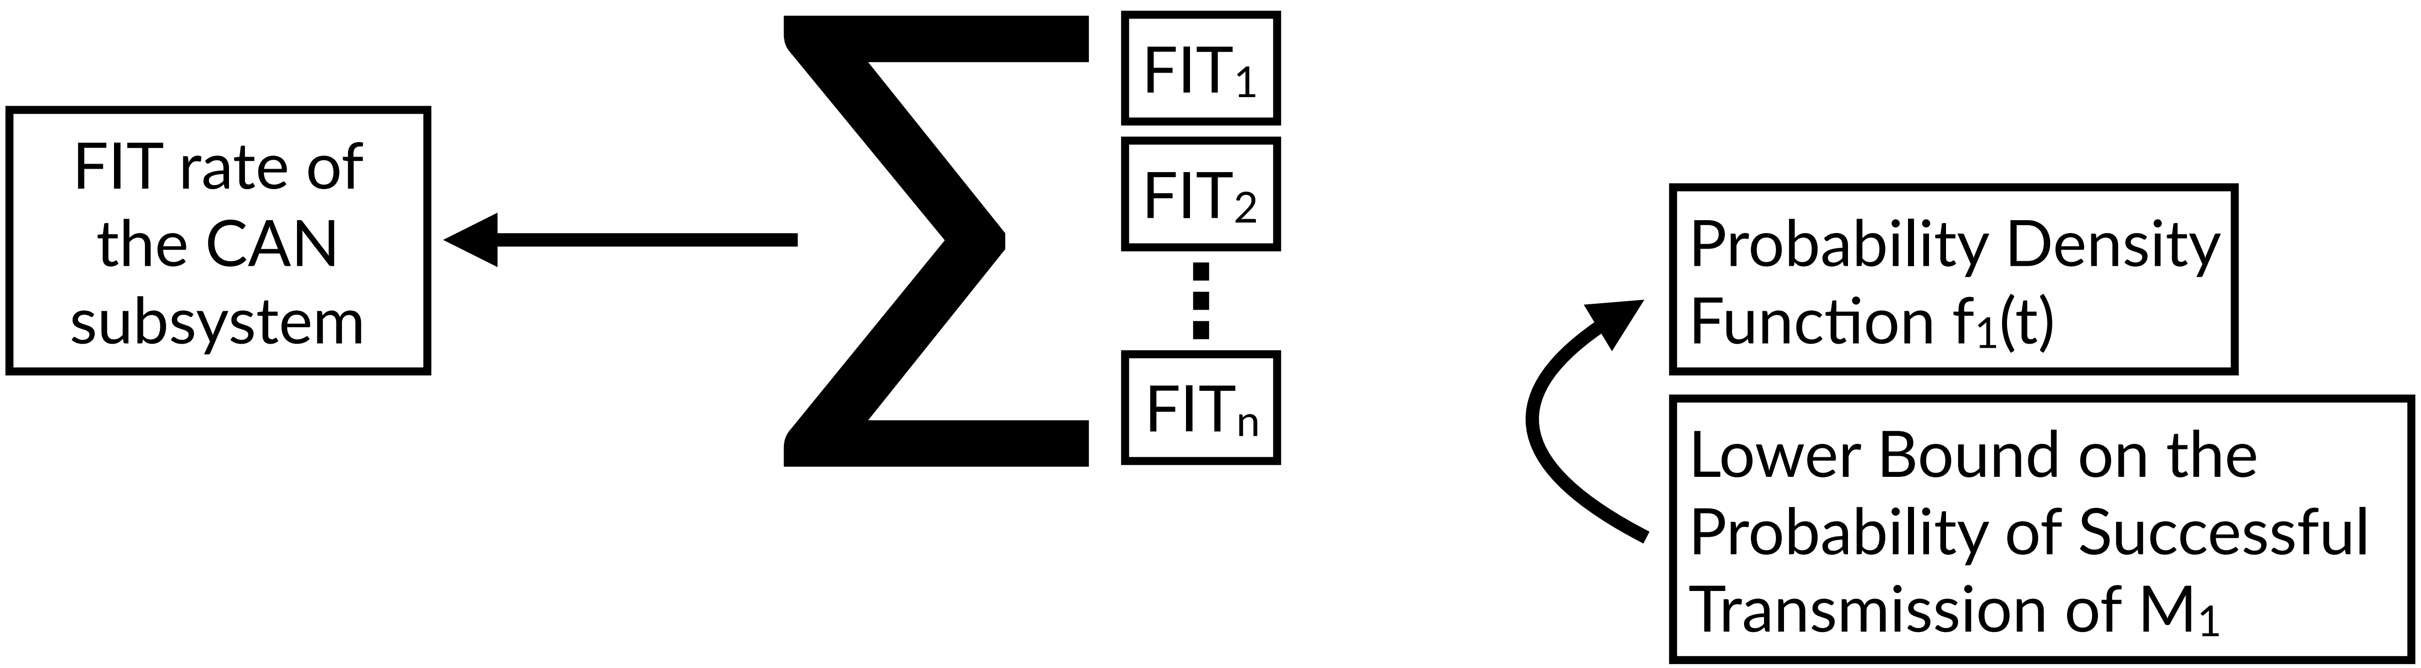

FIT Rate Analysis of the System

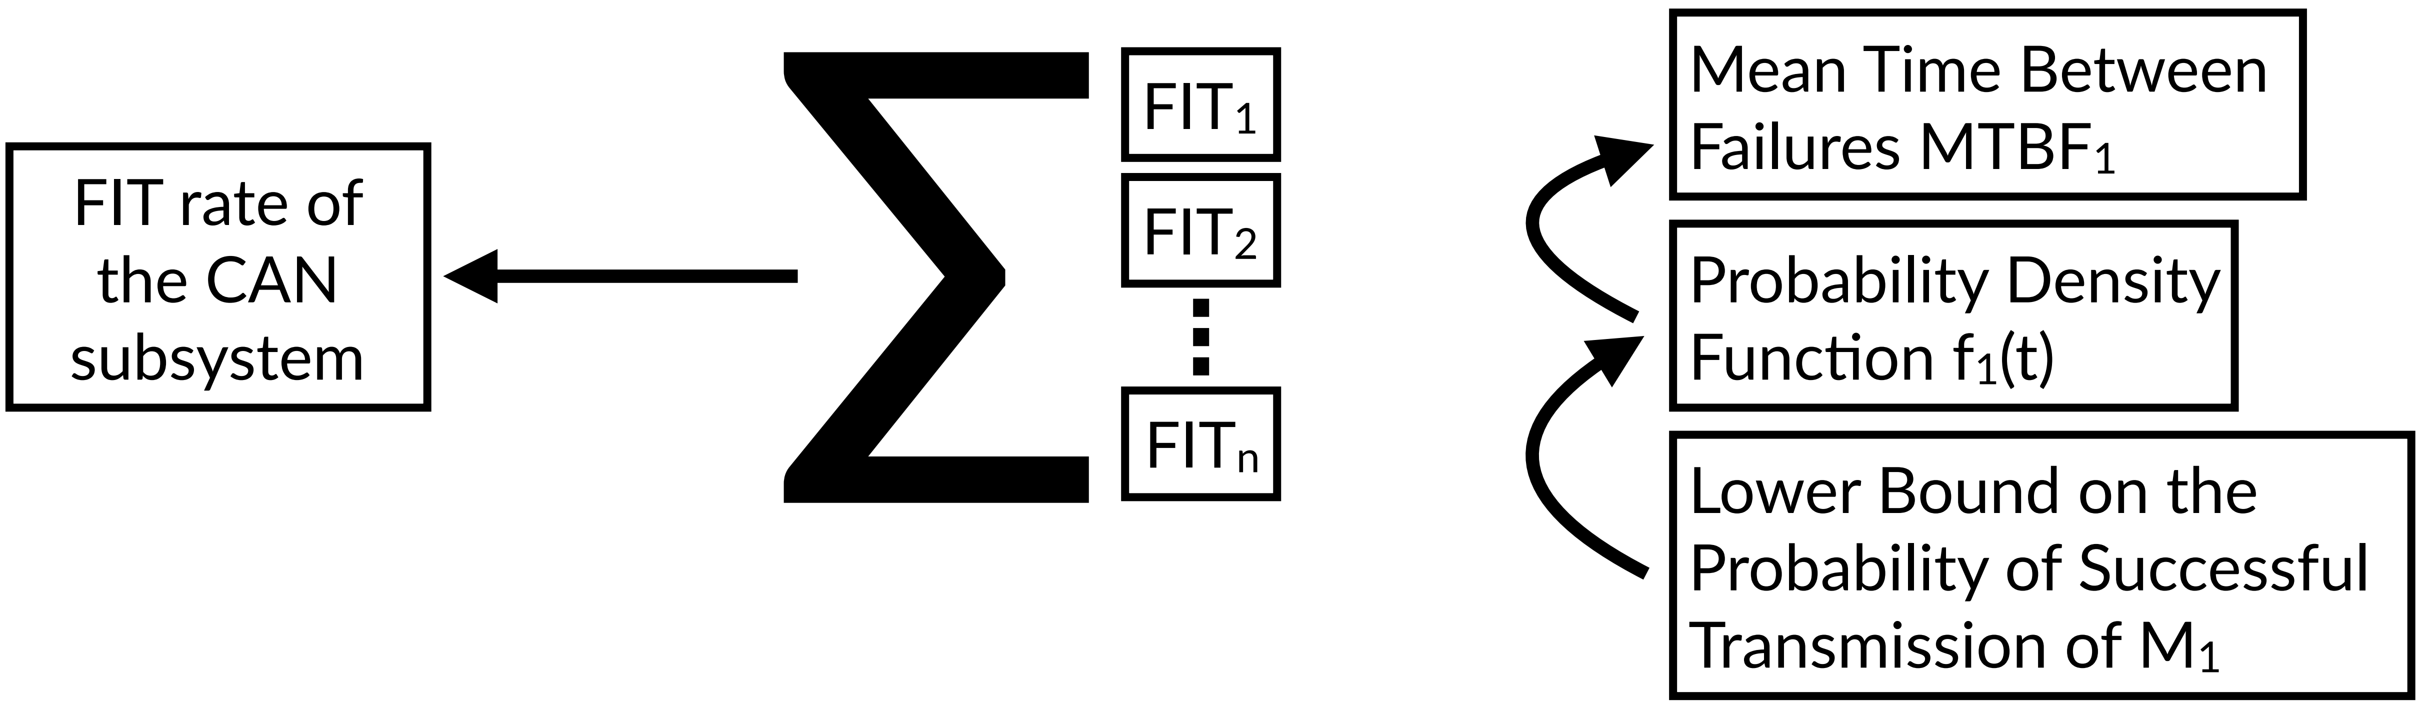

FIT Rate Analysis of the System

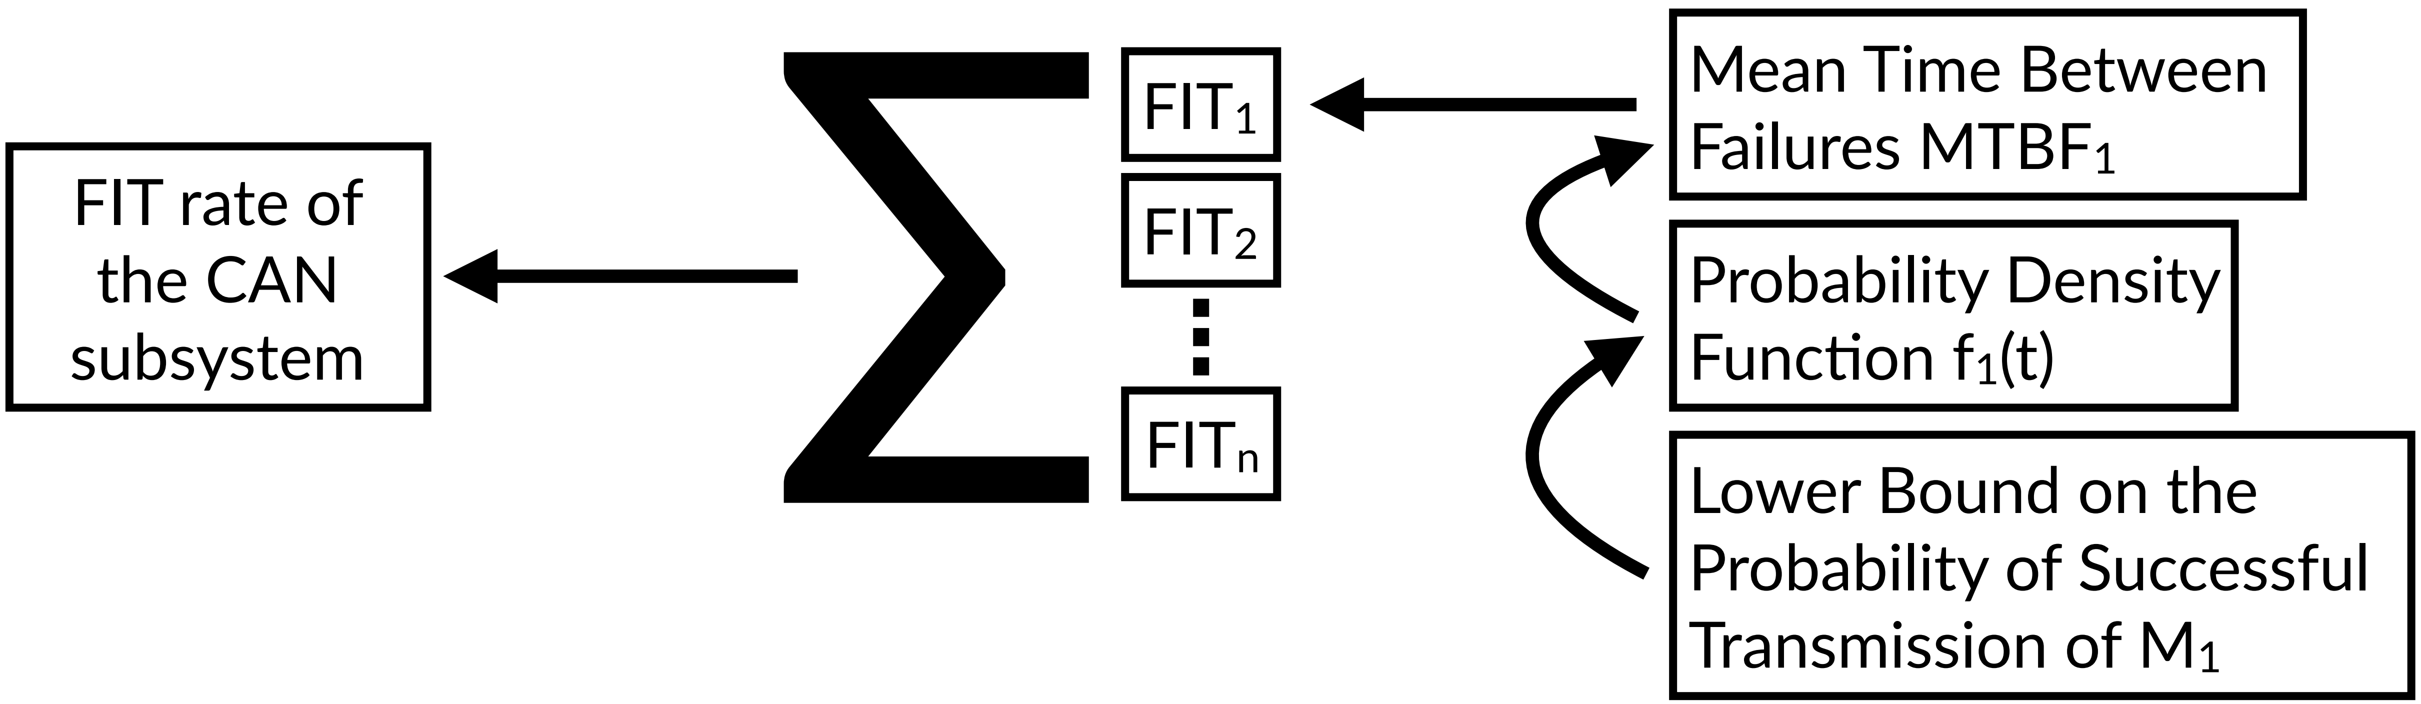

FIT Rate Analysis of the System

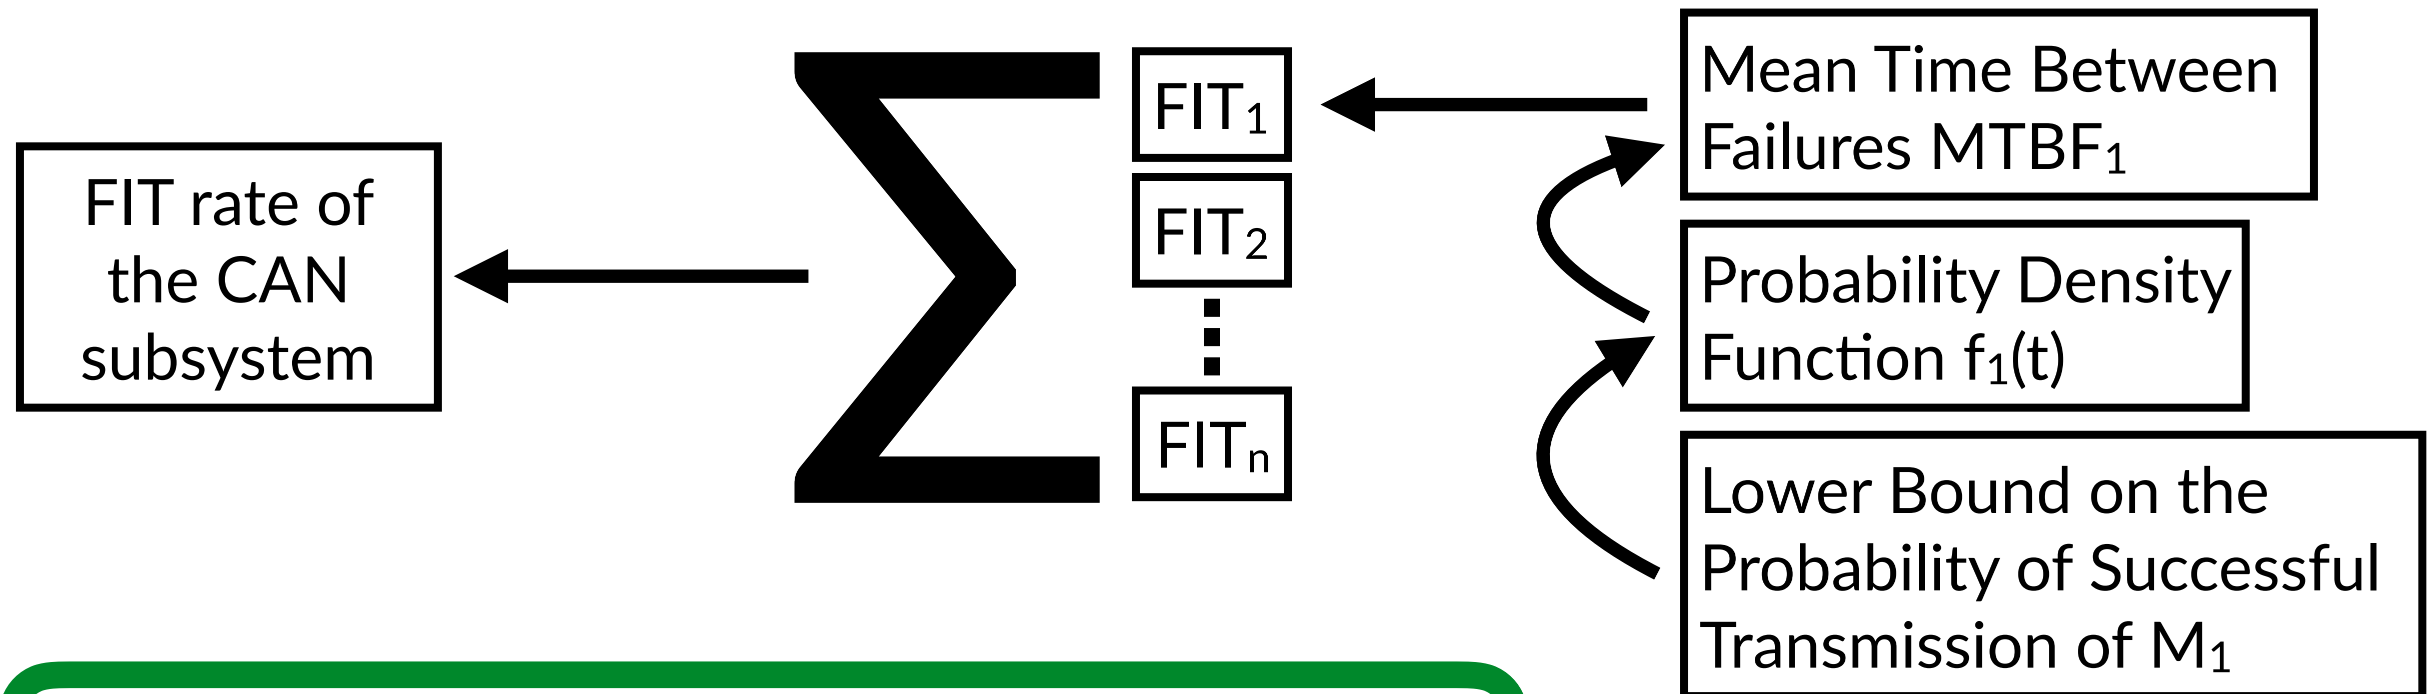

Standard procedure to compute FIT rates given the failure probabilities, but **tailored for real-time workloads**

FIT Rate Analysis of the System

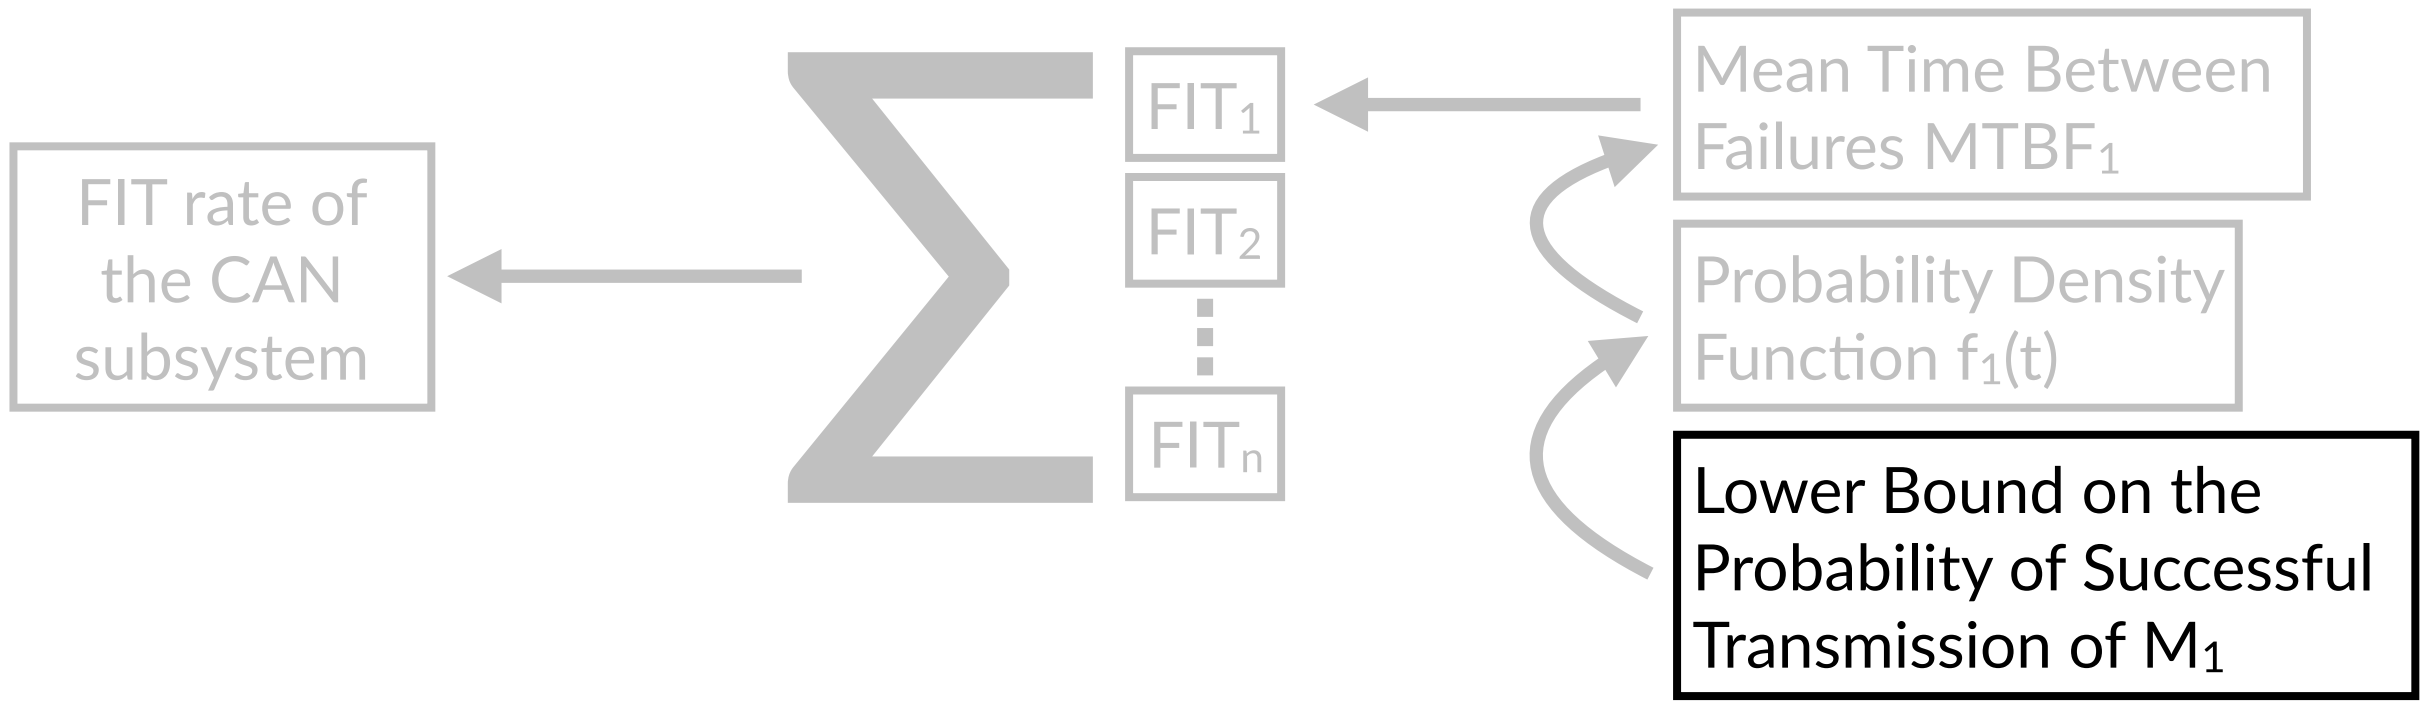

FIT Rate Analysis of the System

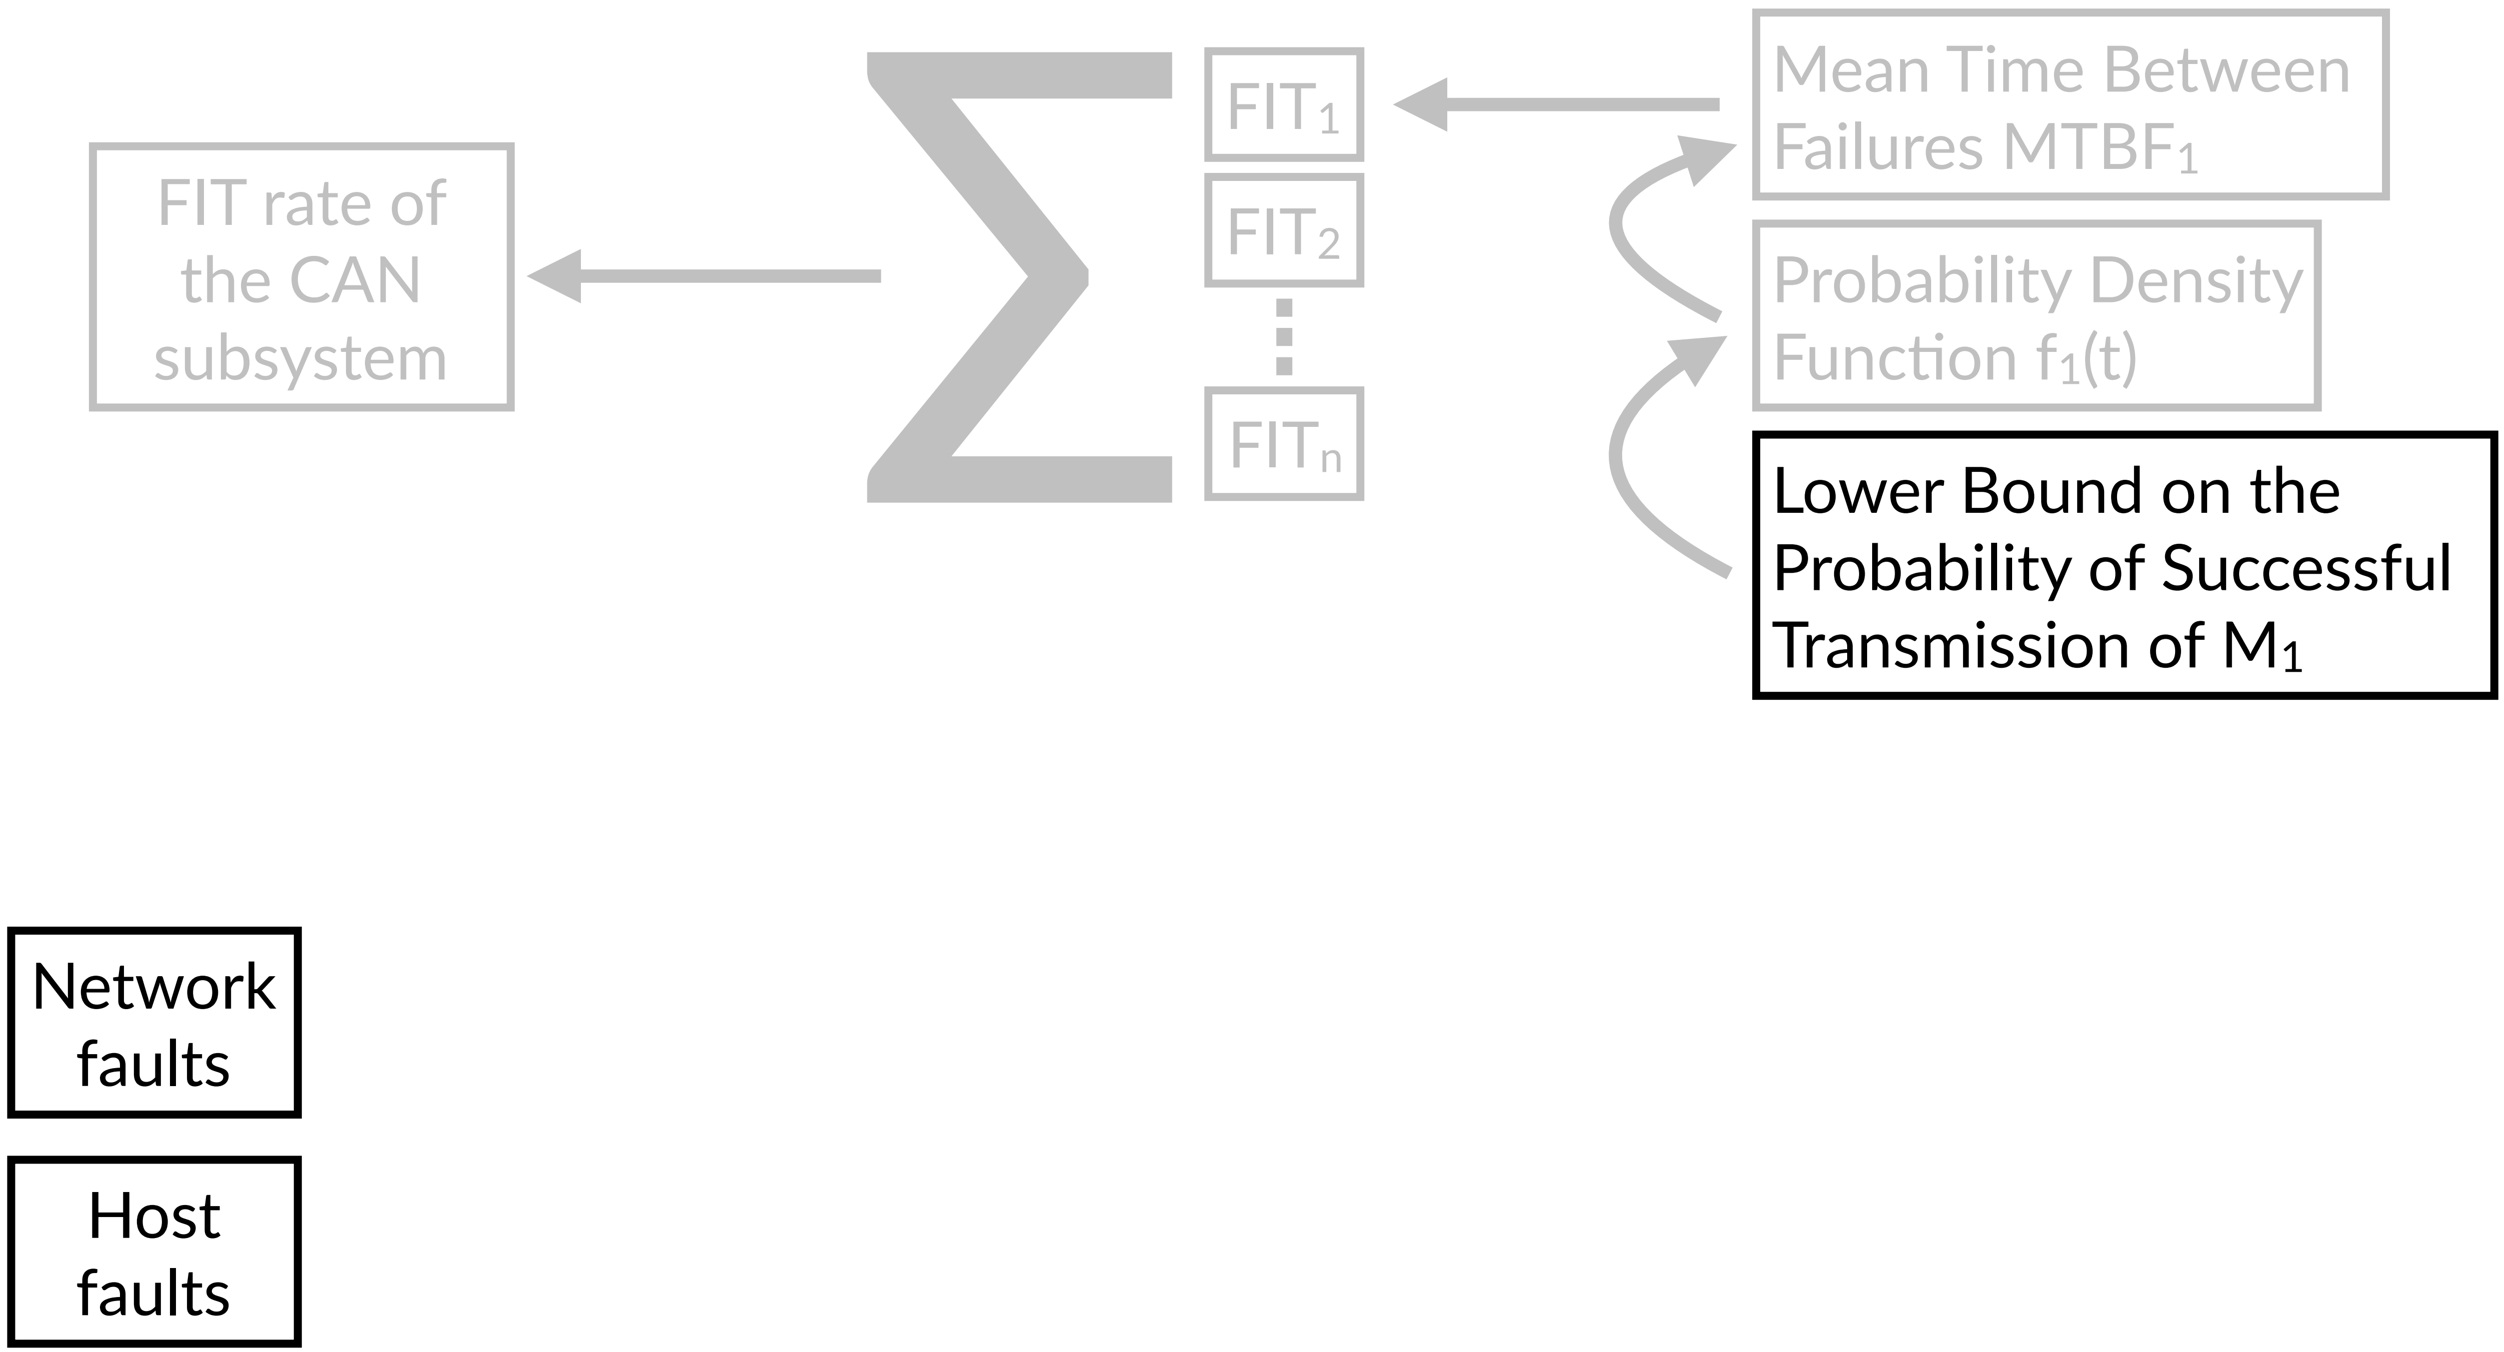

FIT Rate Analysis of the System

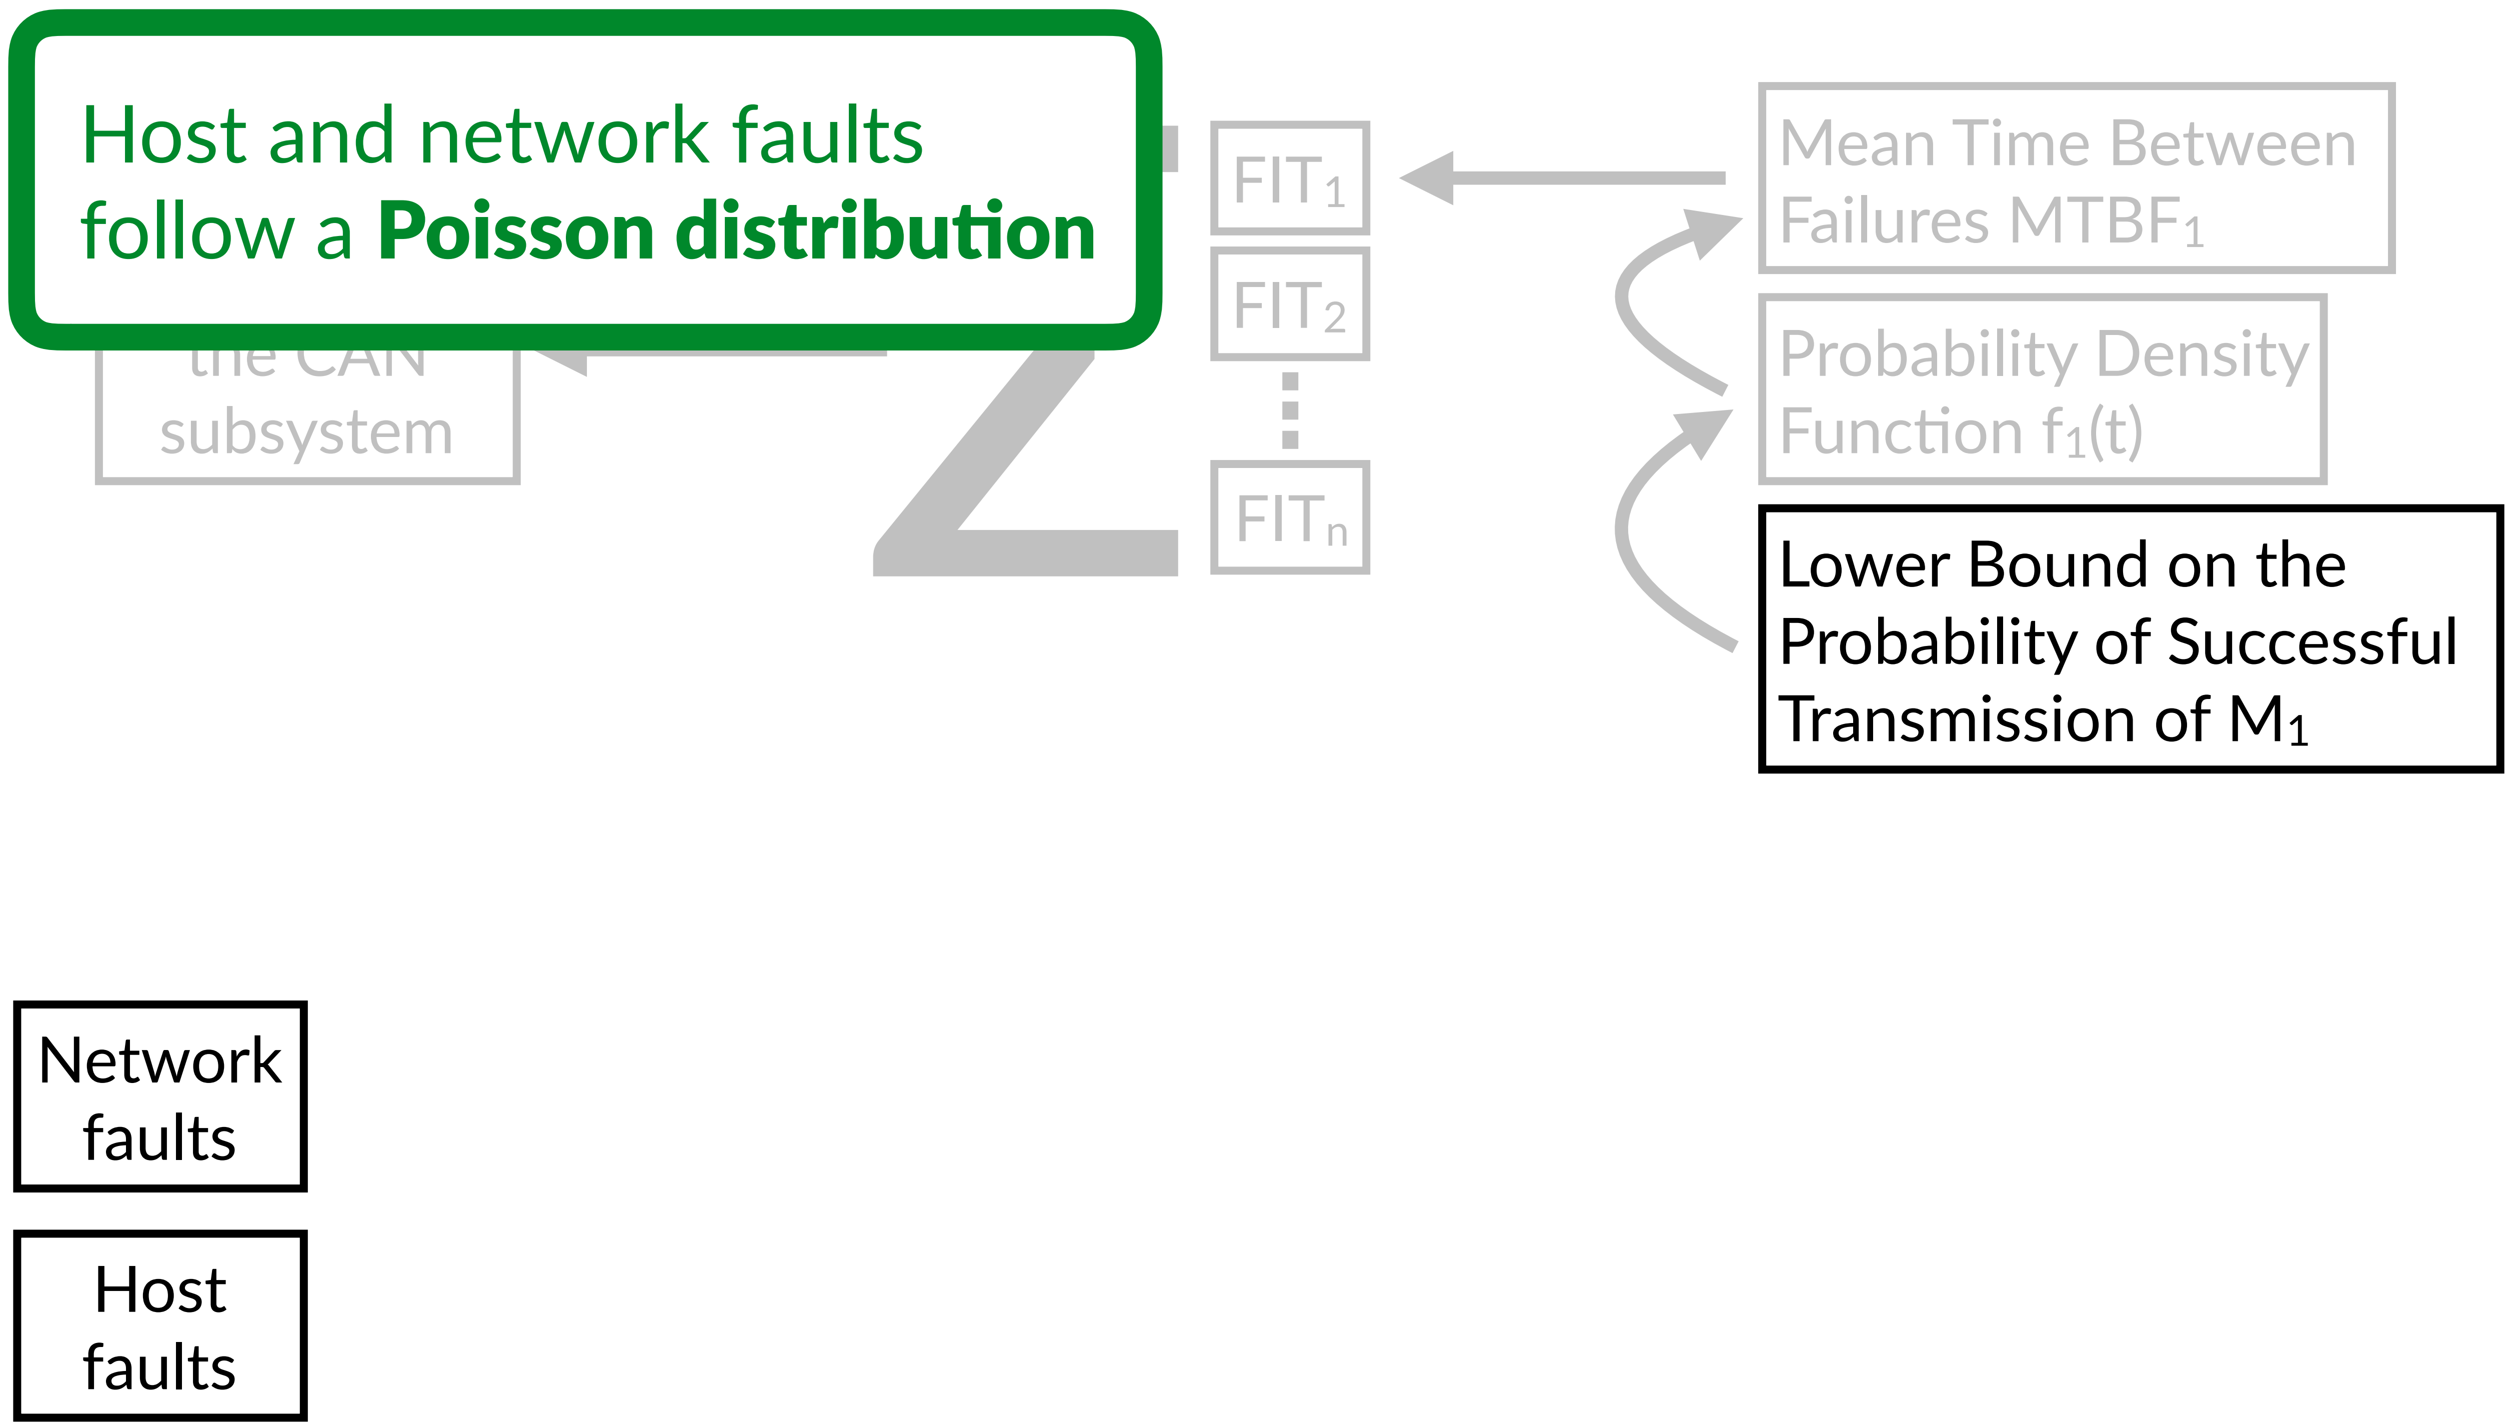

FIT Rate Analysis of the System

Host and network faults follow a **Poisson distribution**

Probabilities that each message:

- retransmitted due to **transmission failures**
- omitted due to **crash failures**
- corrupted due to **commission failures**

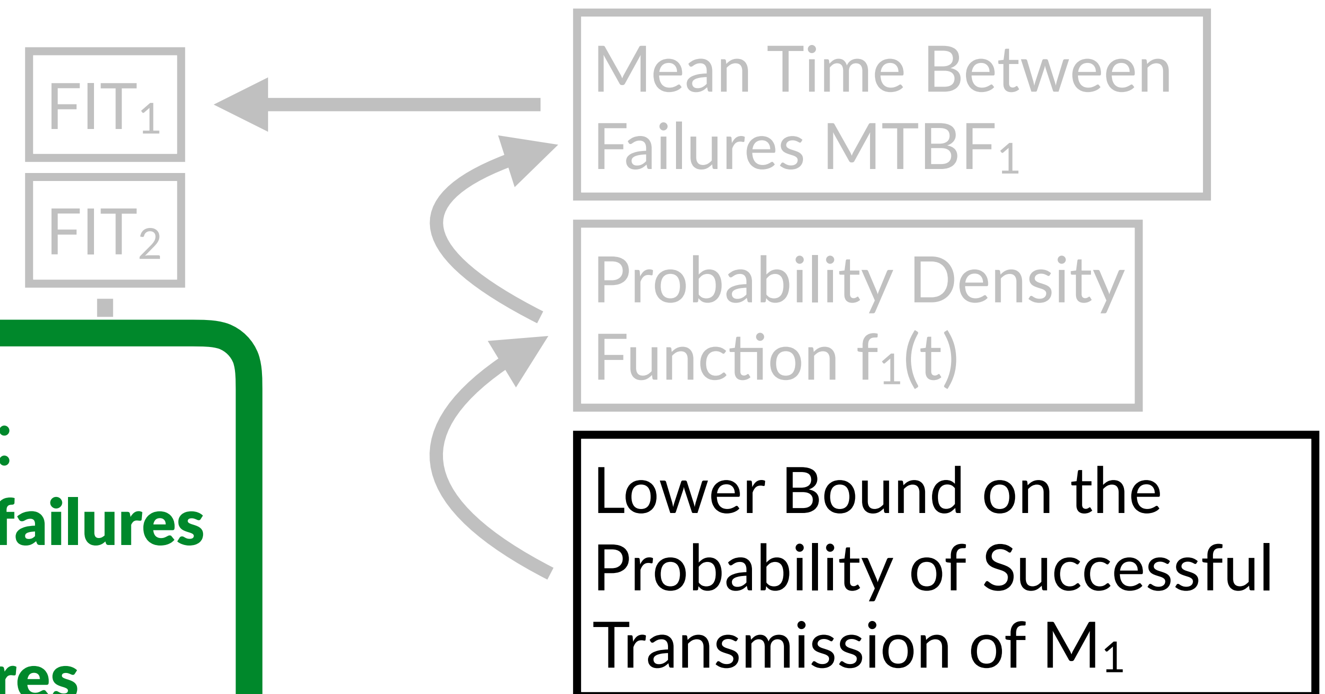

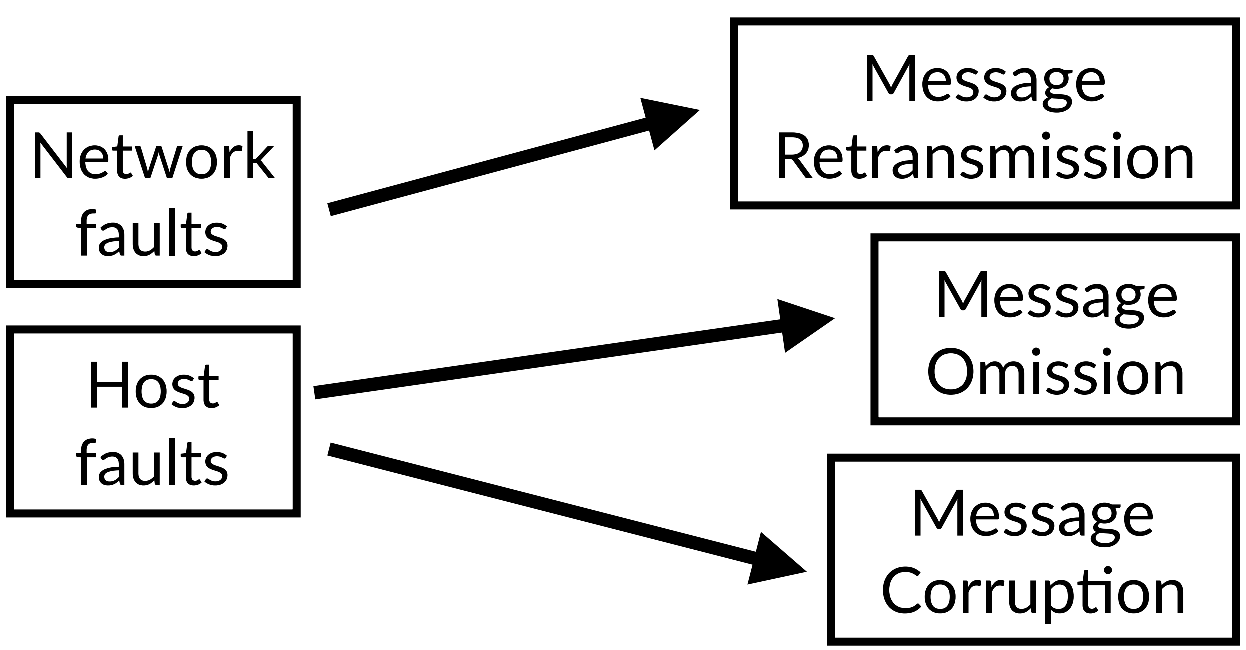

ystem

- Broster et al.'s **probabilistic response-time analysis***

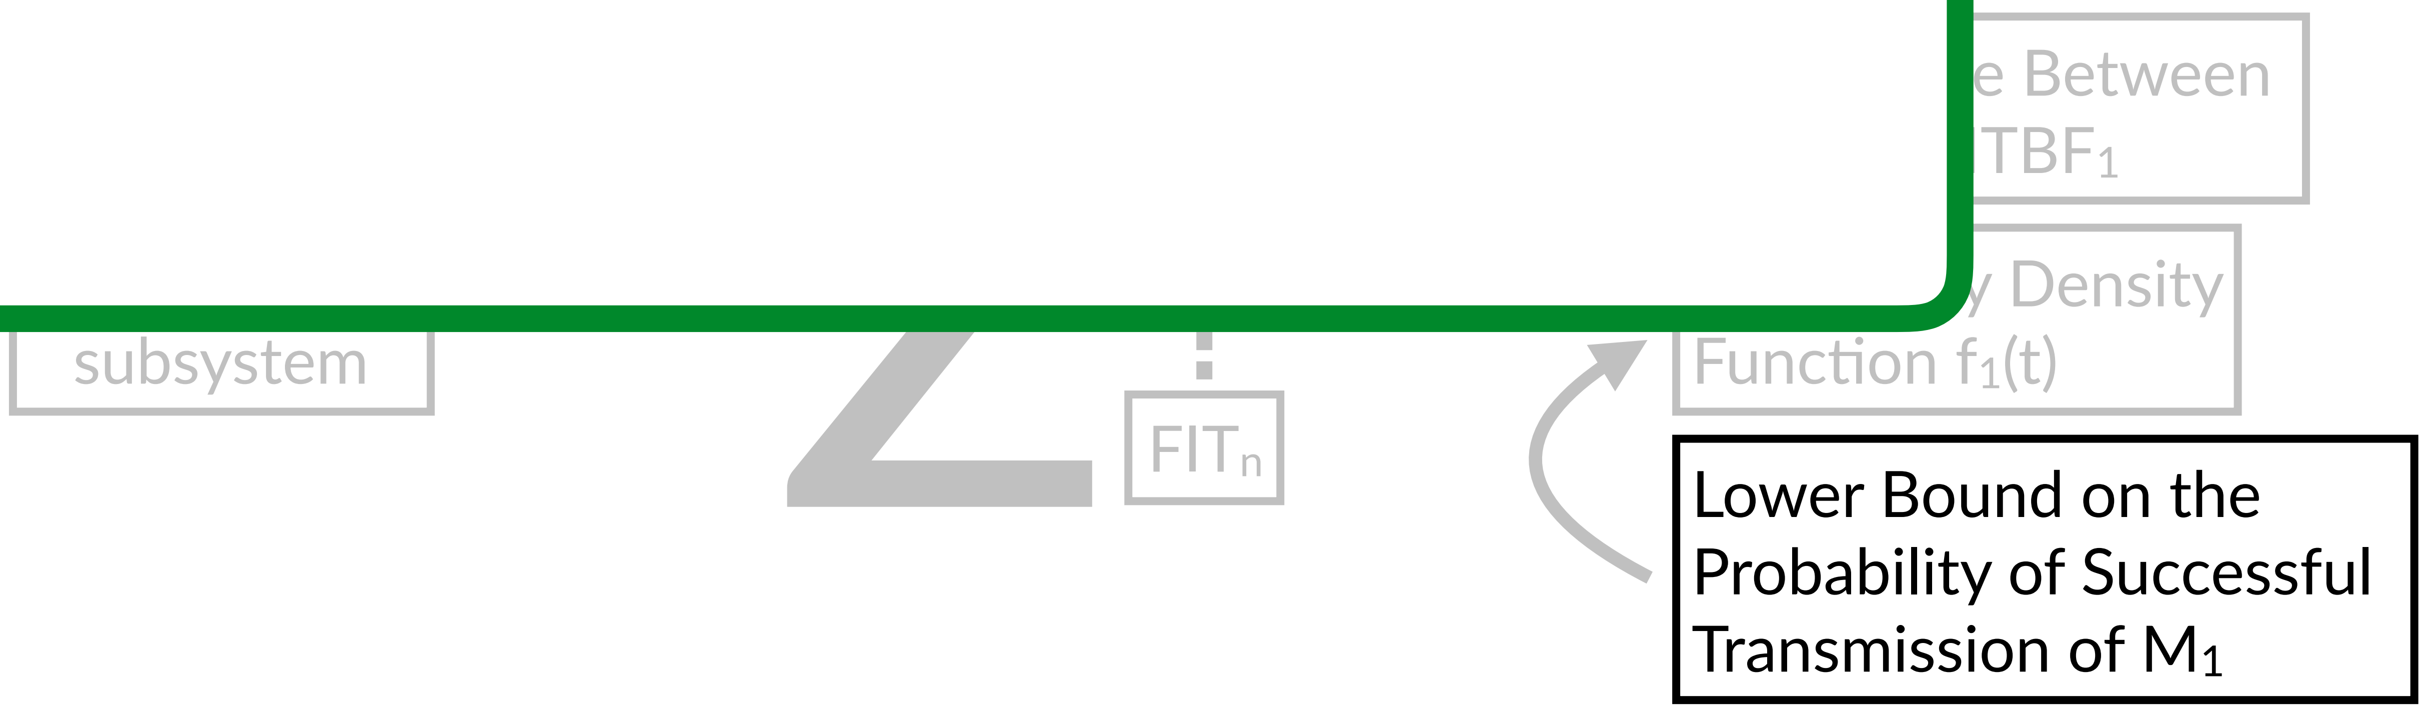

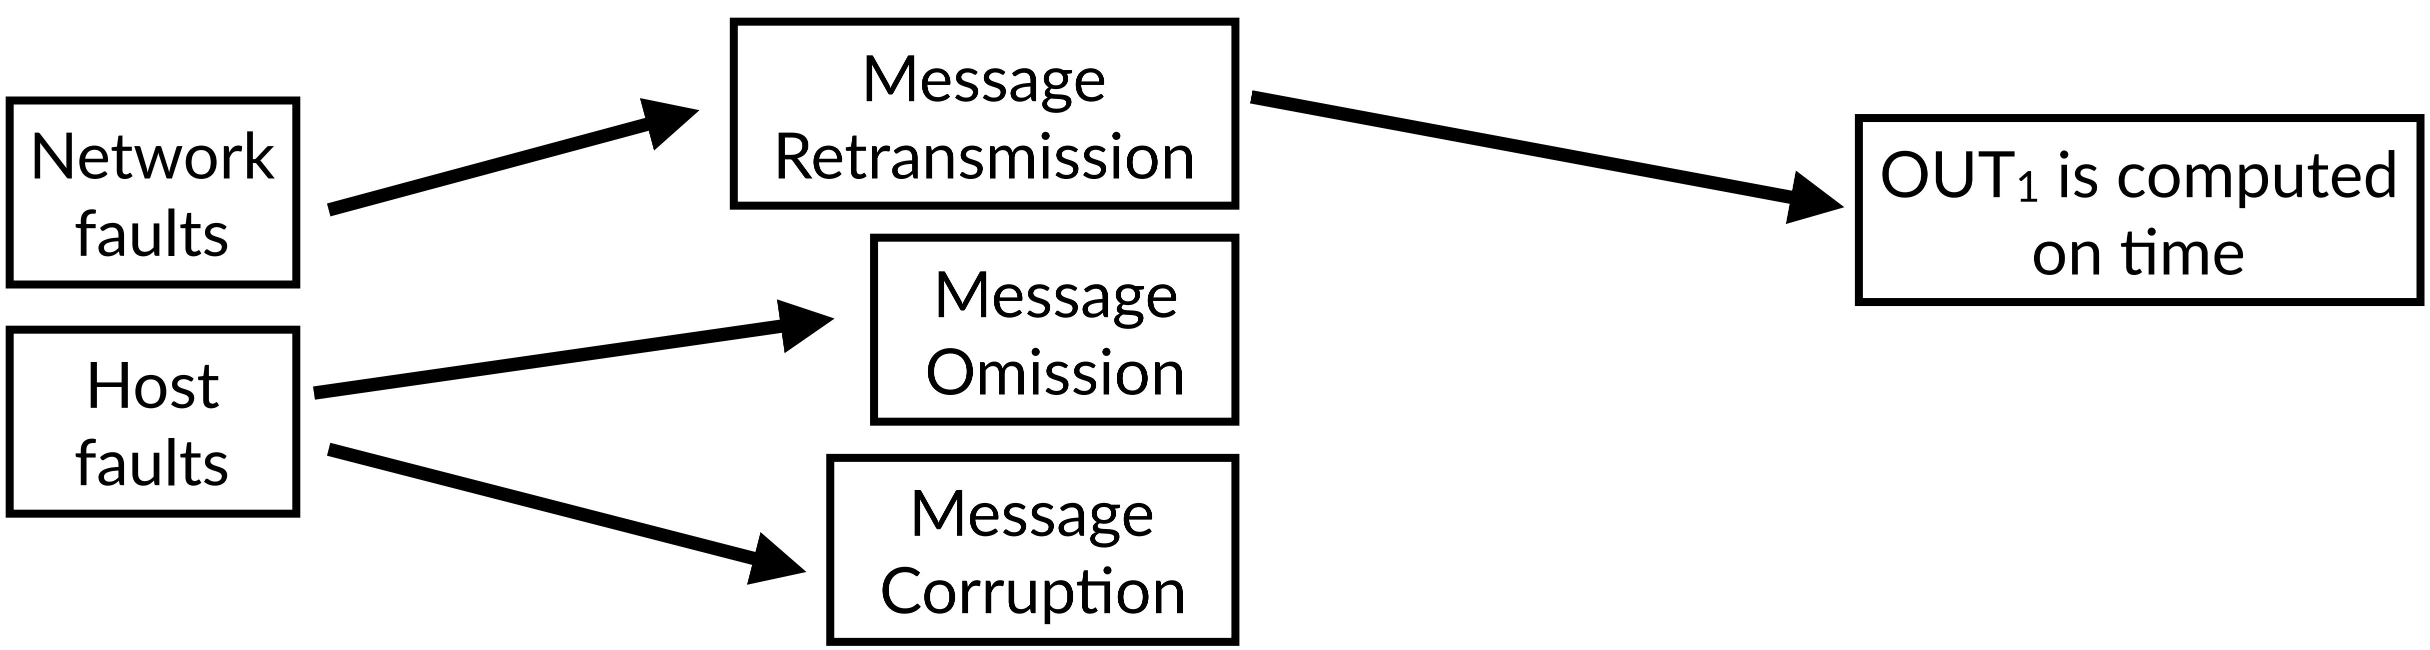

*Broster, Ian, Alan Burns, and Guillermo Rodriguez-Navas. **"Timing analysis of real-time communication under electromagnetic interference."** Real-Time Systems 30.1-2 (2005): 55-81.

- Broster et al.'s **probabilistic response-time analysis***
- We extend the analysis for a **set of message replicas**
  - ➡ E.g., any 1 out of 3 message replicas are transmitted on time

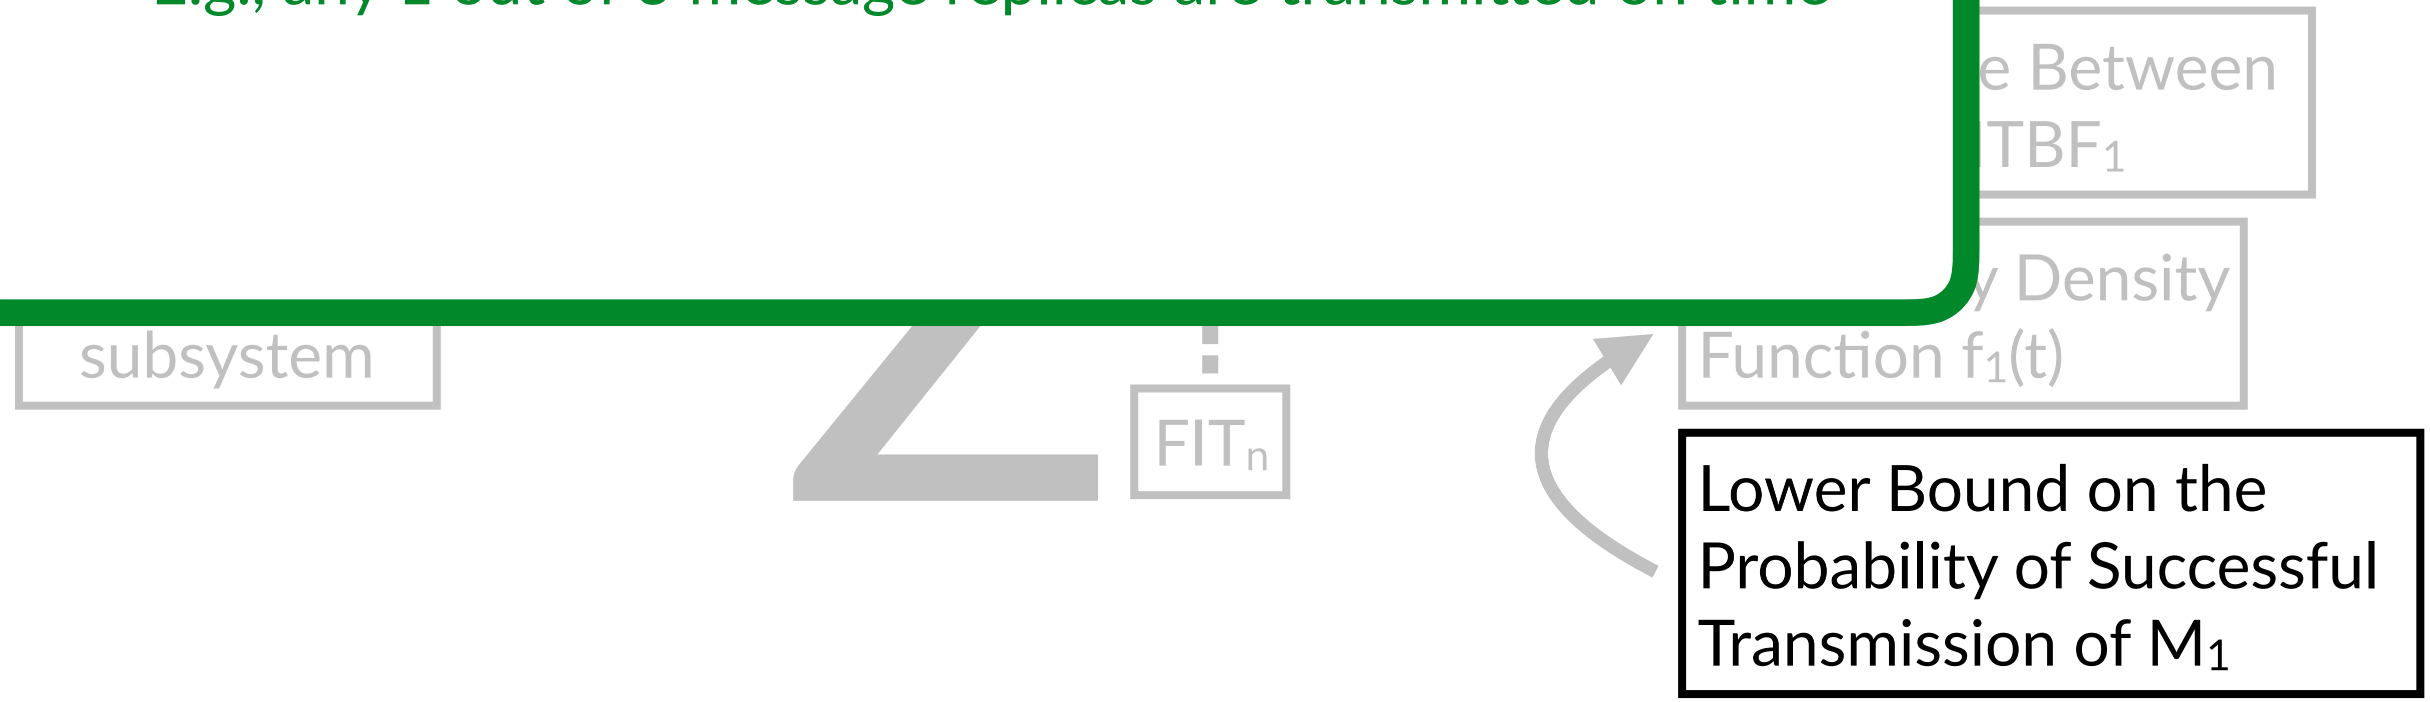

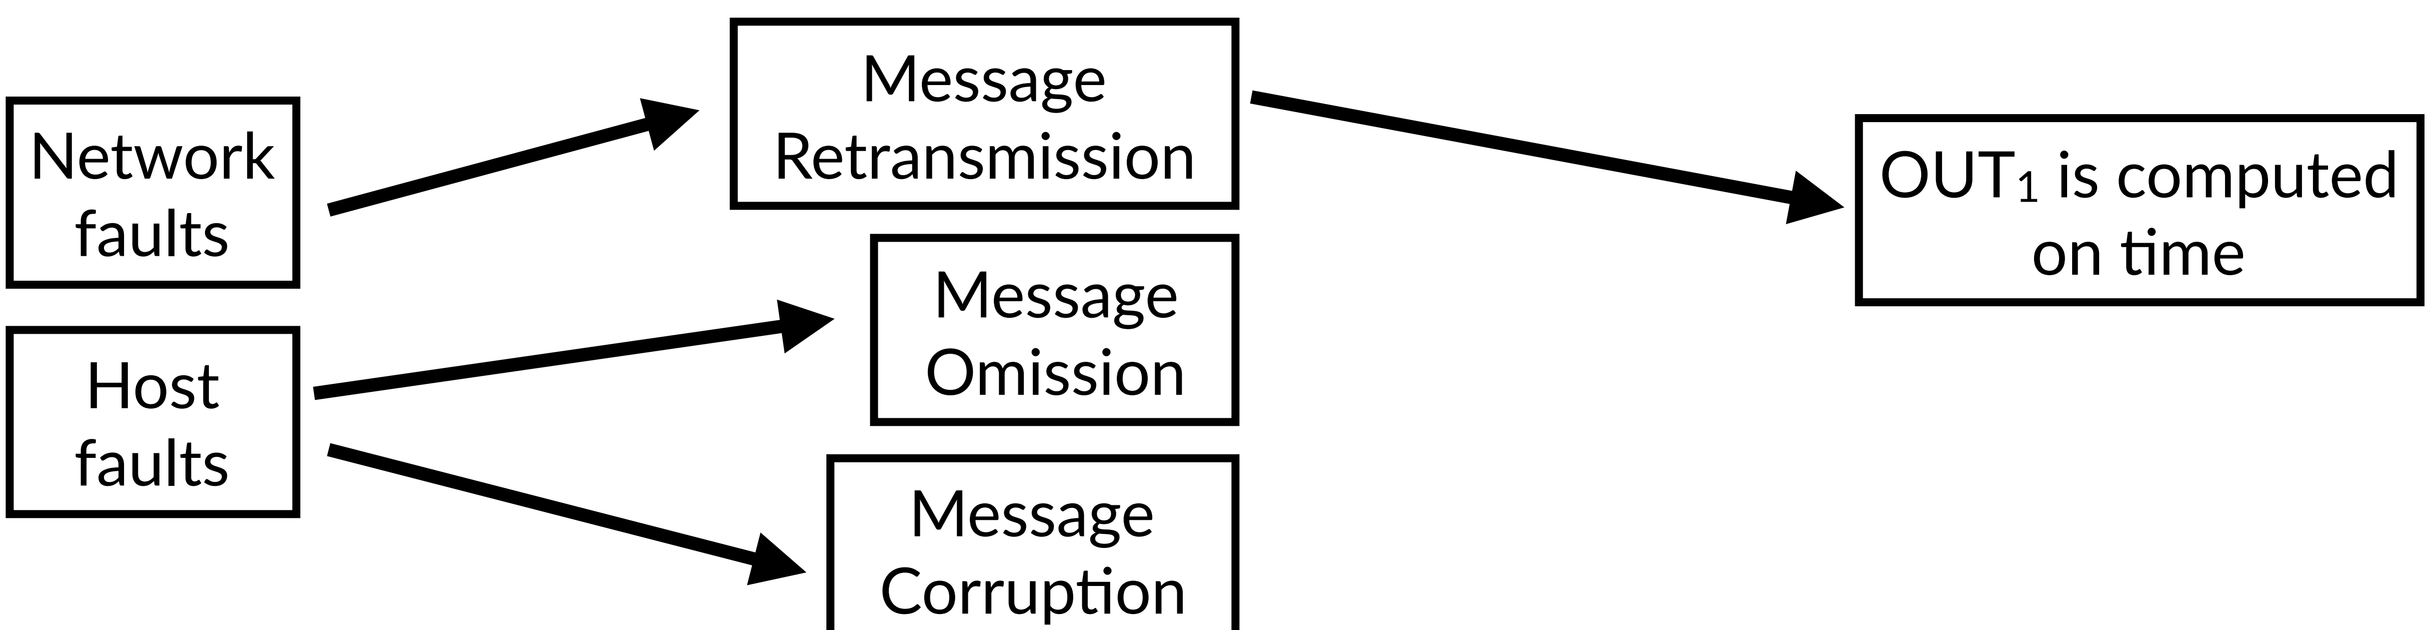

*Broster, Ian, Alan Burns, and Guillermo Rodriguez-Navas. **"Timing analysis of real-time communication under electromagnetic interference."** Real-Time Systems 30.1-2 (2005): 55-81.

ystem

- Broster et al.'s **probabilistic response-time analysis***
- We extend the analysis for a **set of message replicas**
  - ➡ E.g., any 1 out of 3 message replicas are transmitted on time
- Deadline violations due to **crash failures**

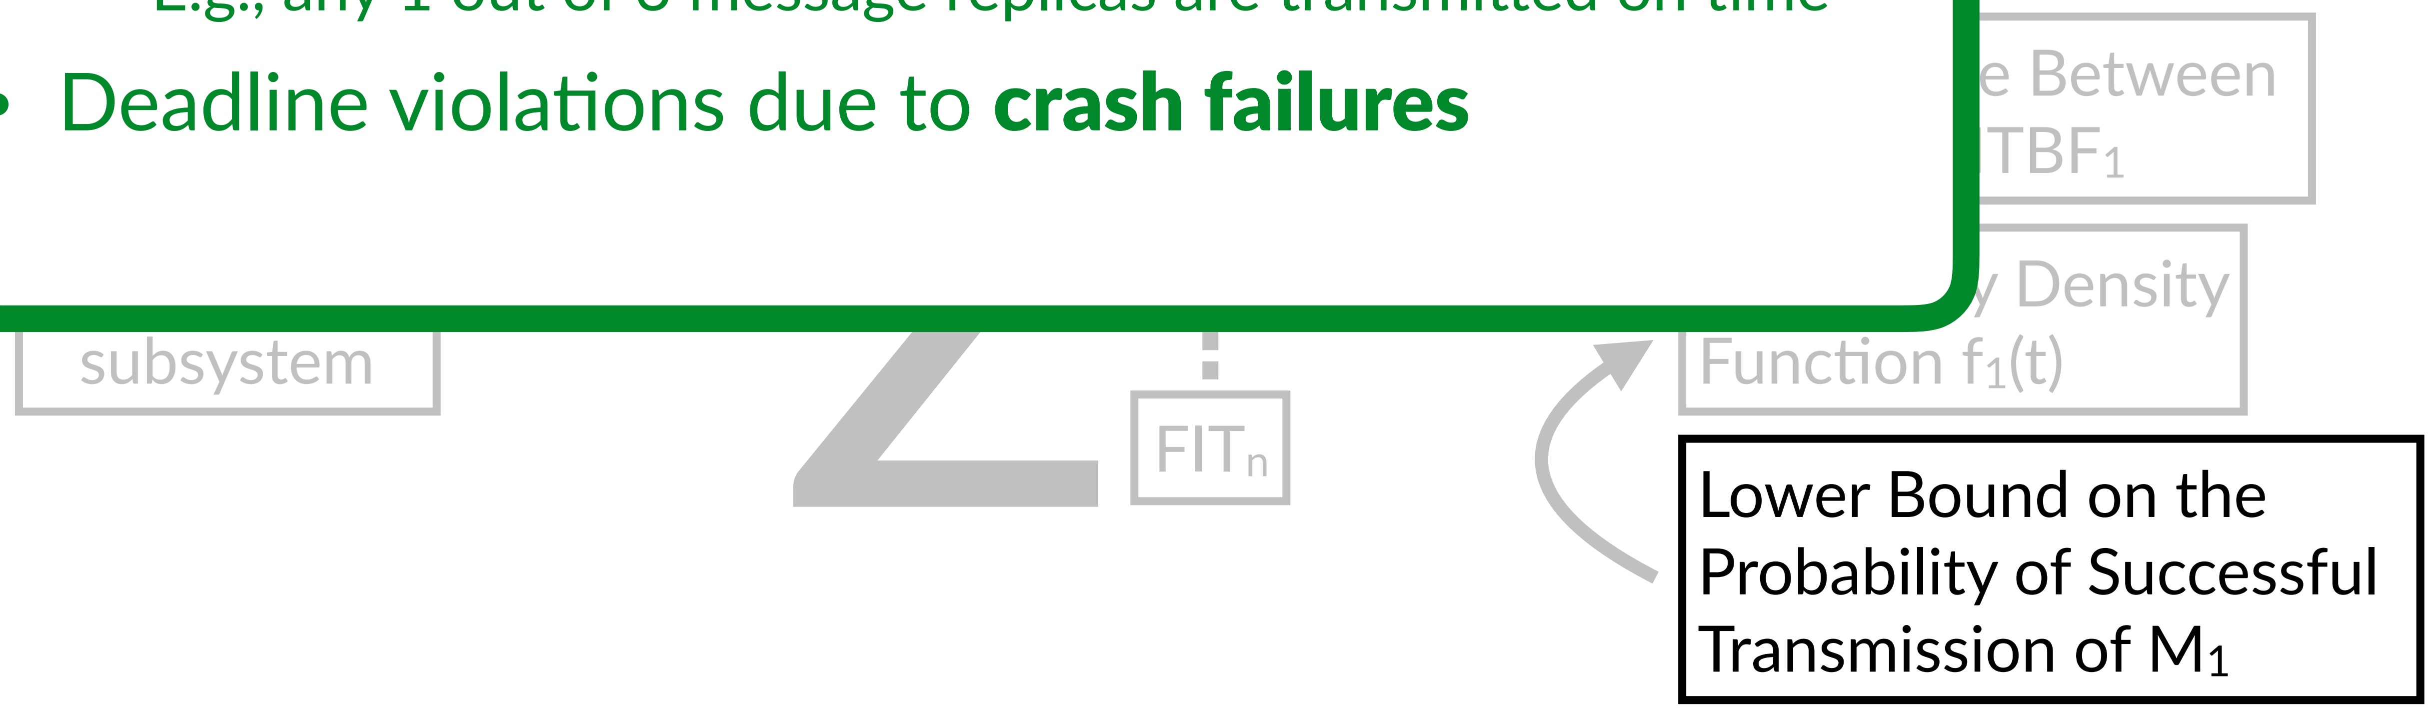

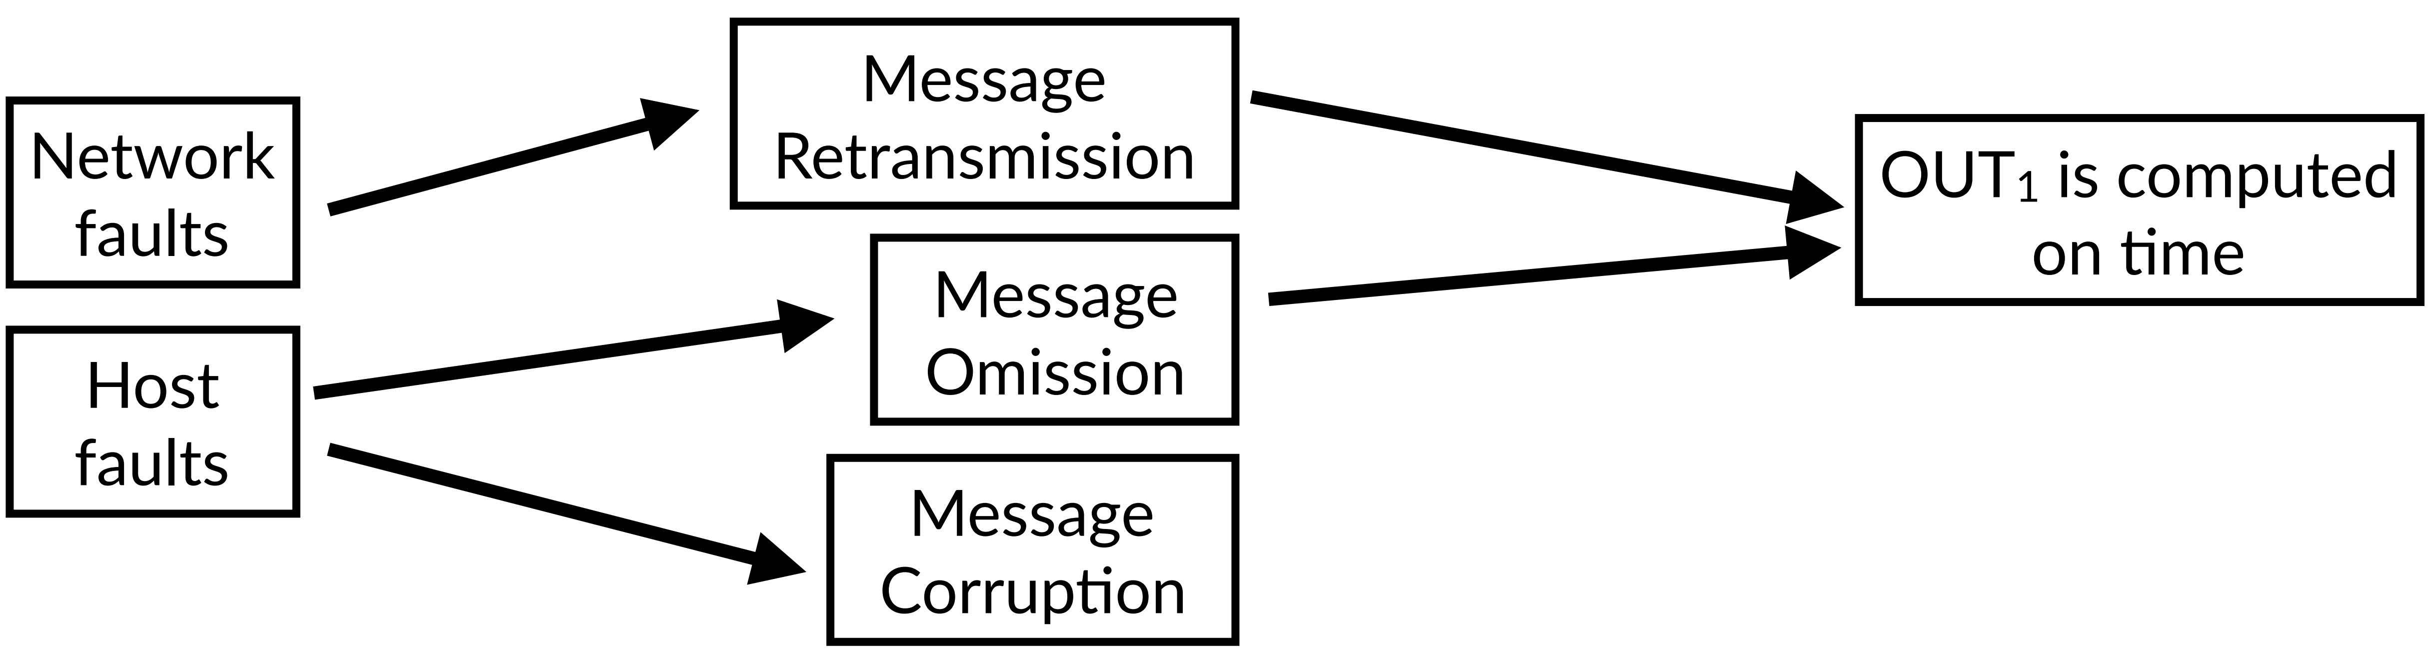

*Broster, Ian, Alan Burns, and Guillermo Rodriguez-Navas. **"Timing analysis of real-time communication under electromagnetic interference."** Real-Time Systems 30.1-2 (2005): 55-81.

- Broster et al.'s **probabilistic response-time analysis***
- We extend the analysis for a **set of message replicas**
  - ➡ E.g., any 1 out of 3 message replicas are transmitted on time
- Deadline violations due to **crash failures**
- Incorrect output due to **commission failures**

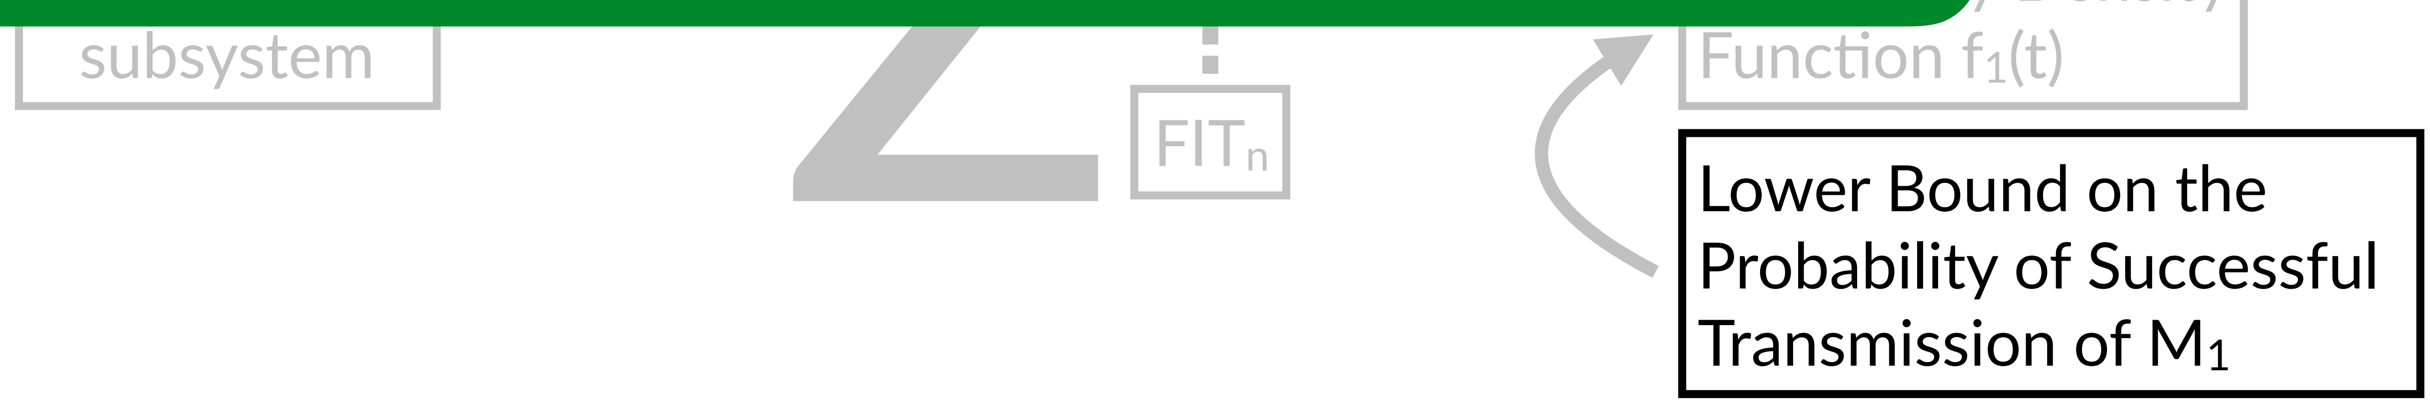

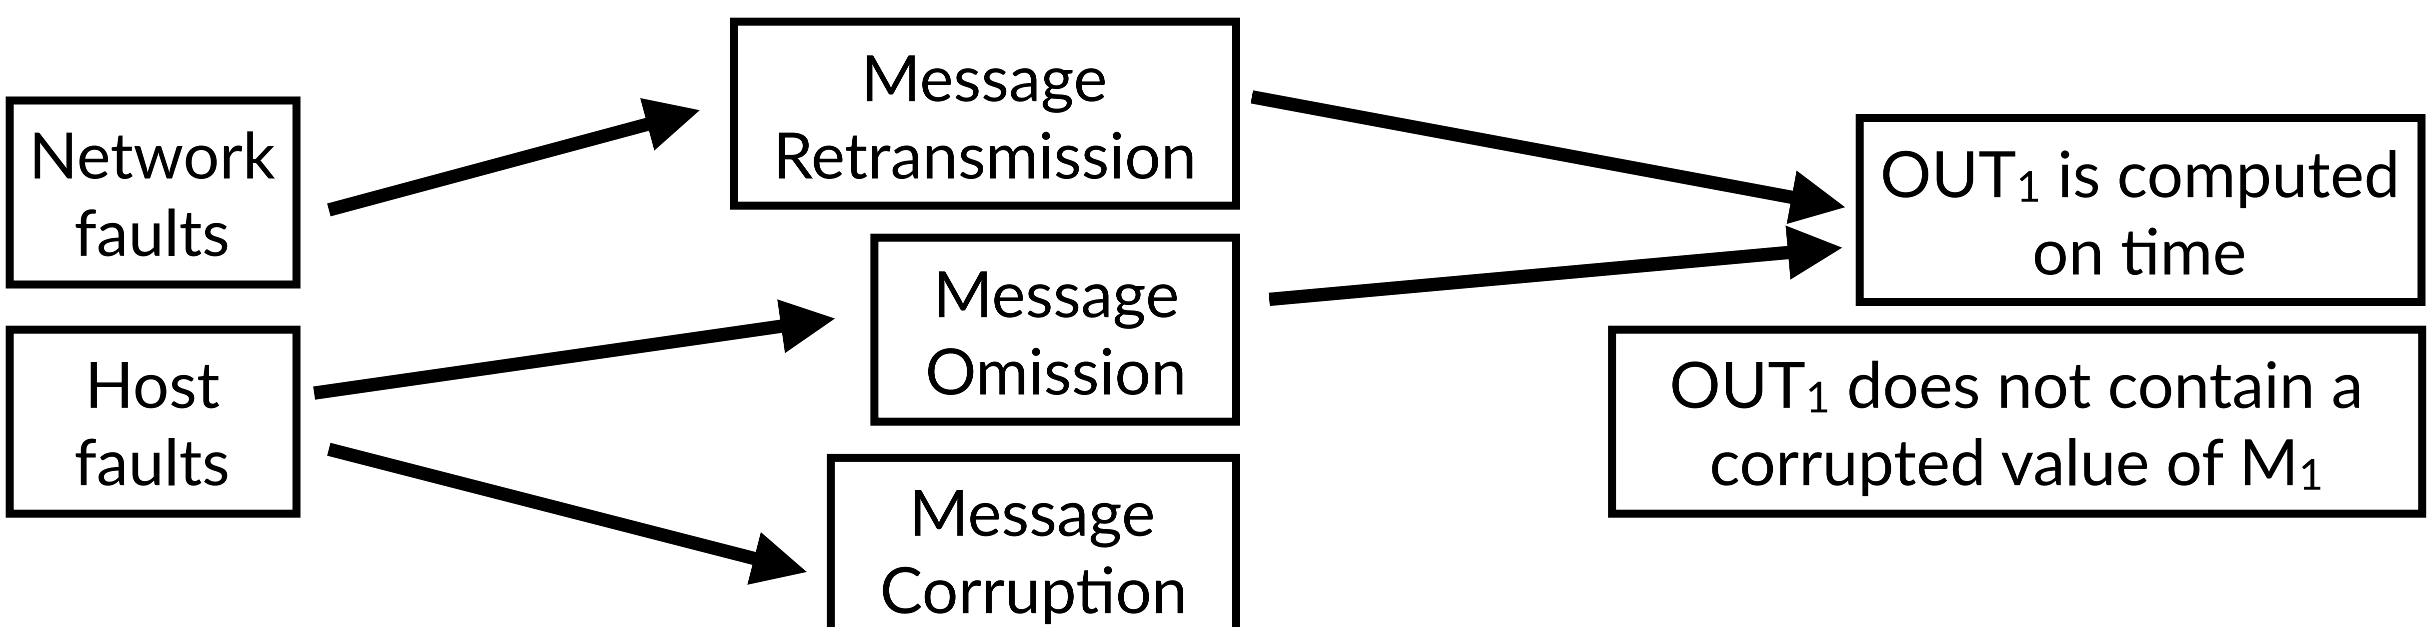

*Broster, Ian, Alan Burns, and Guillermo Rodriguez-Navas. **"Timing analysis of real-time communication under electromagnetic interference."** Real-Time Systems 30.1-2 (2005): 55-81.

- Broster et al.'s **probabilistic response-time analysis***
- We extend the analysis for a **set of message replicas**
  - ➡ E.g., any 1 out of 3 message replicas are transmitted on time
- Deadline violations due to **crash failures**
- Incorrect output due to **commission failures**

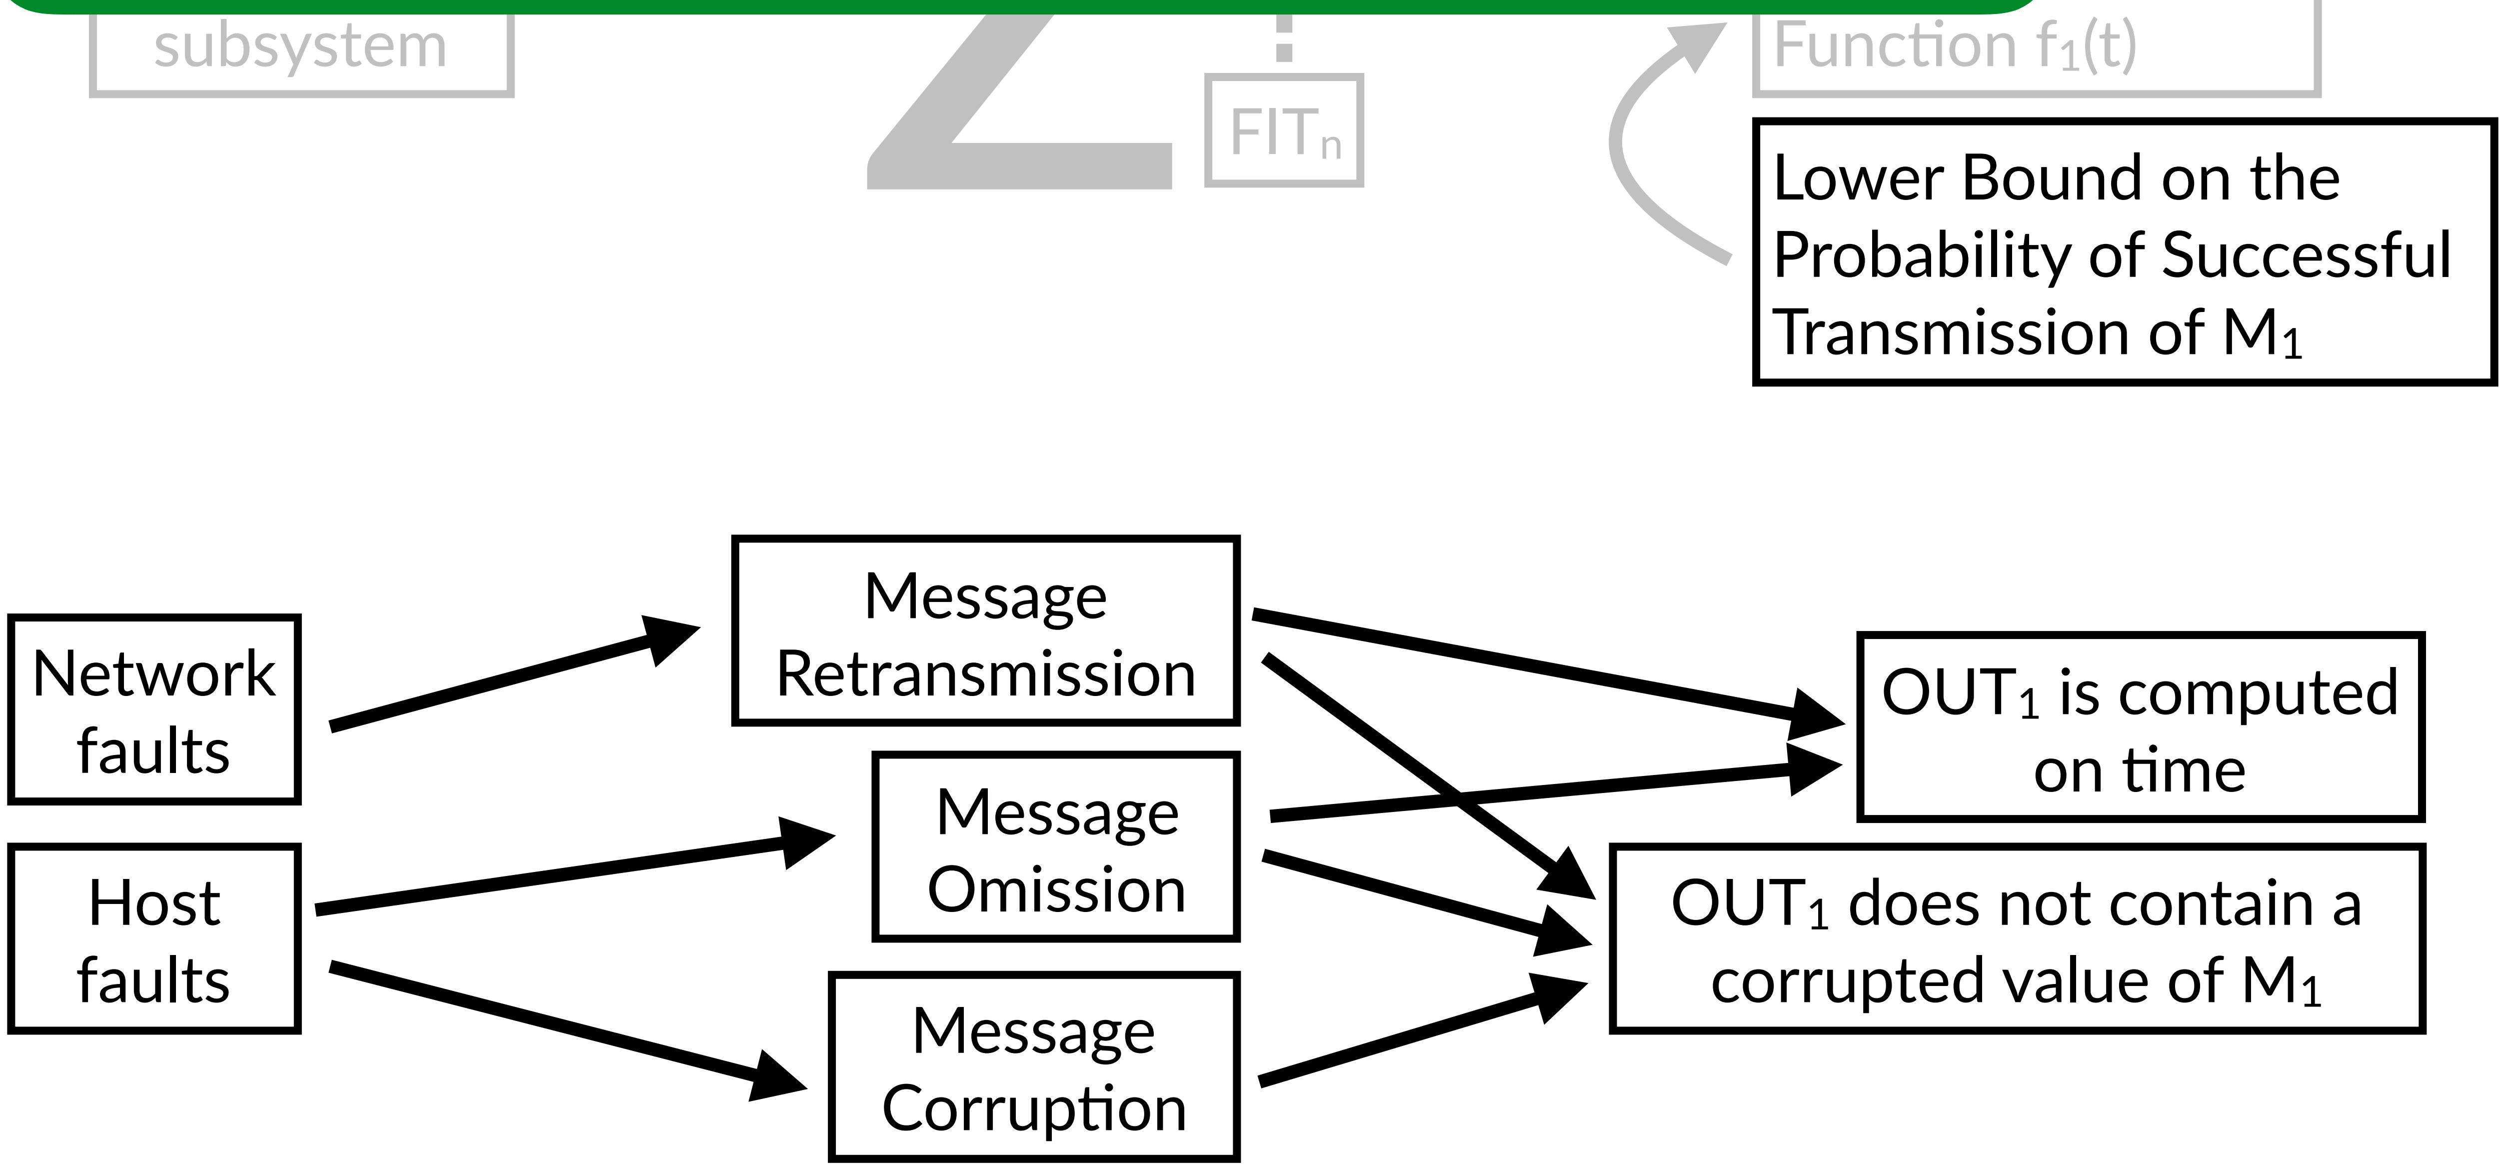

*Broster, Ian, Alan Burns, and Guillermo Rodriguez-Navas. **"Timing analysis of real-time communication under electromagnetic interference."** Real-Time Systems 30.1-2 (2005): 55-81.

- Broster et al.'s **probabilistic response-time analysis***
- We extend the analysis for a **set of message replicas**
  - ➡ E.g., any 1 out of 3 message replicas are transmitted on time
- Deadline violations due to **crash failures**
- Incorrect output due to **commission failures**

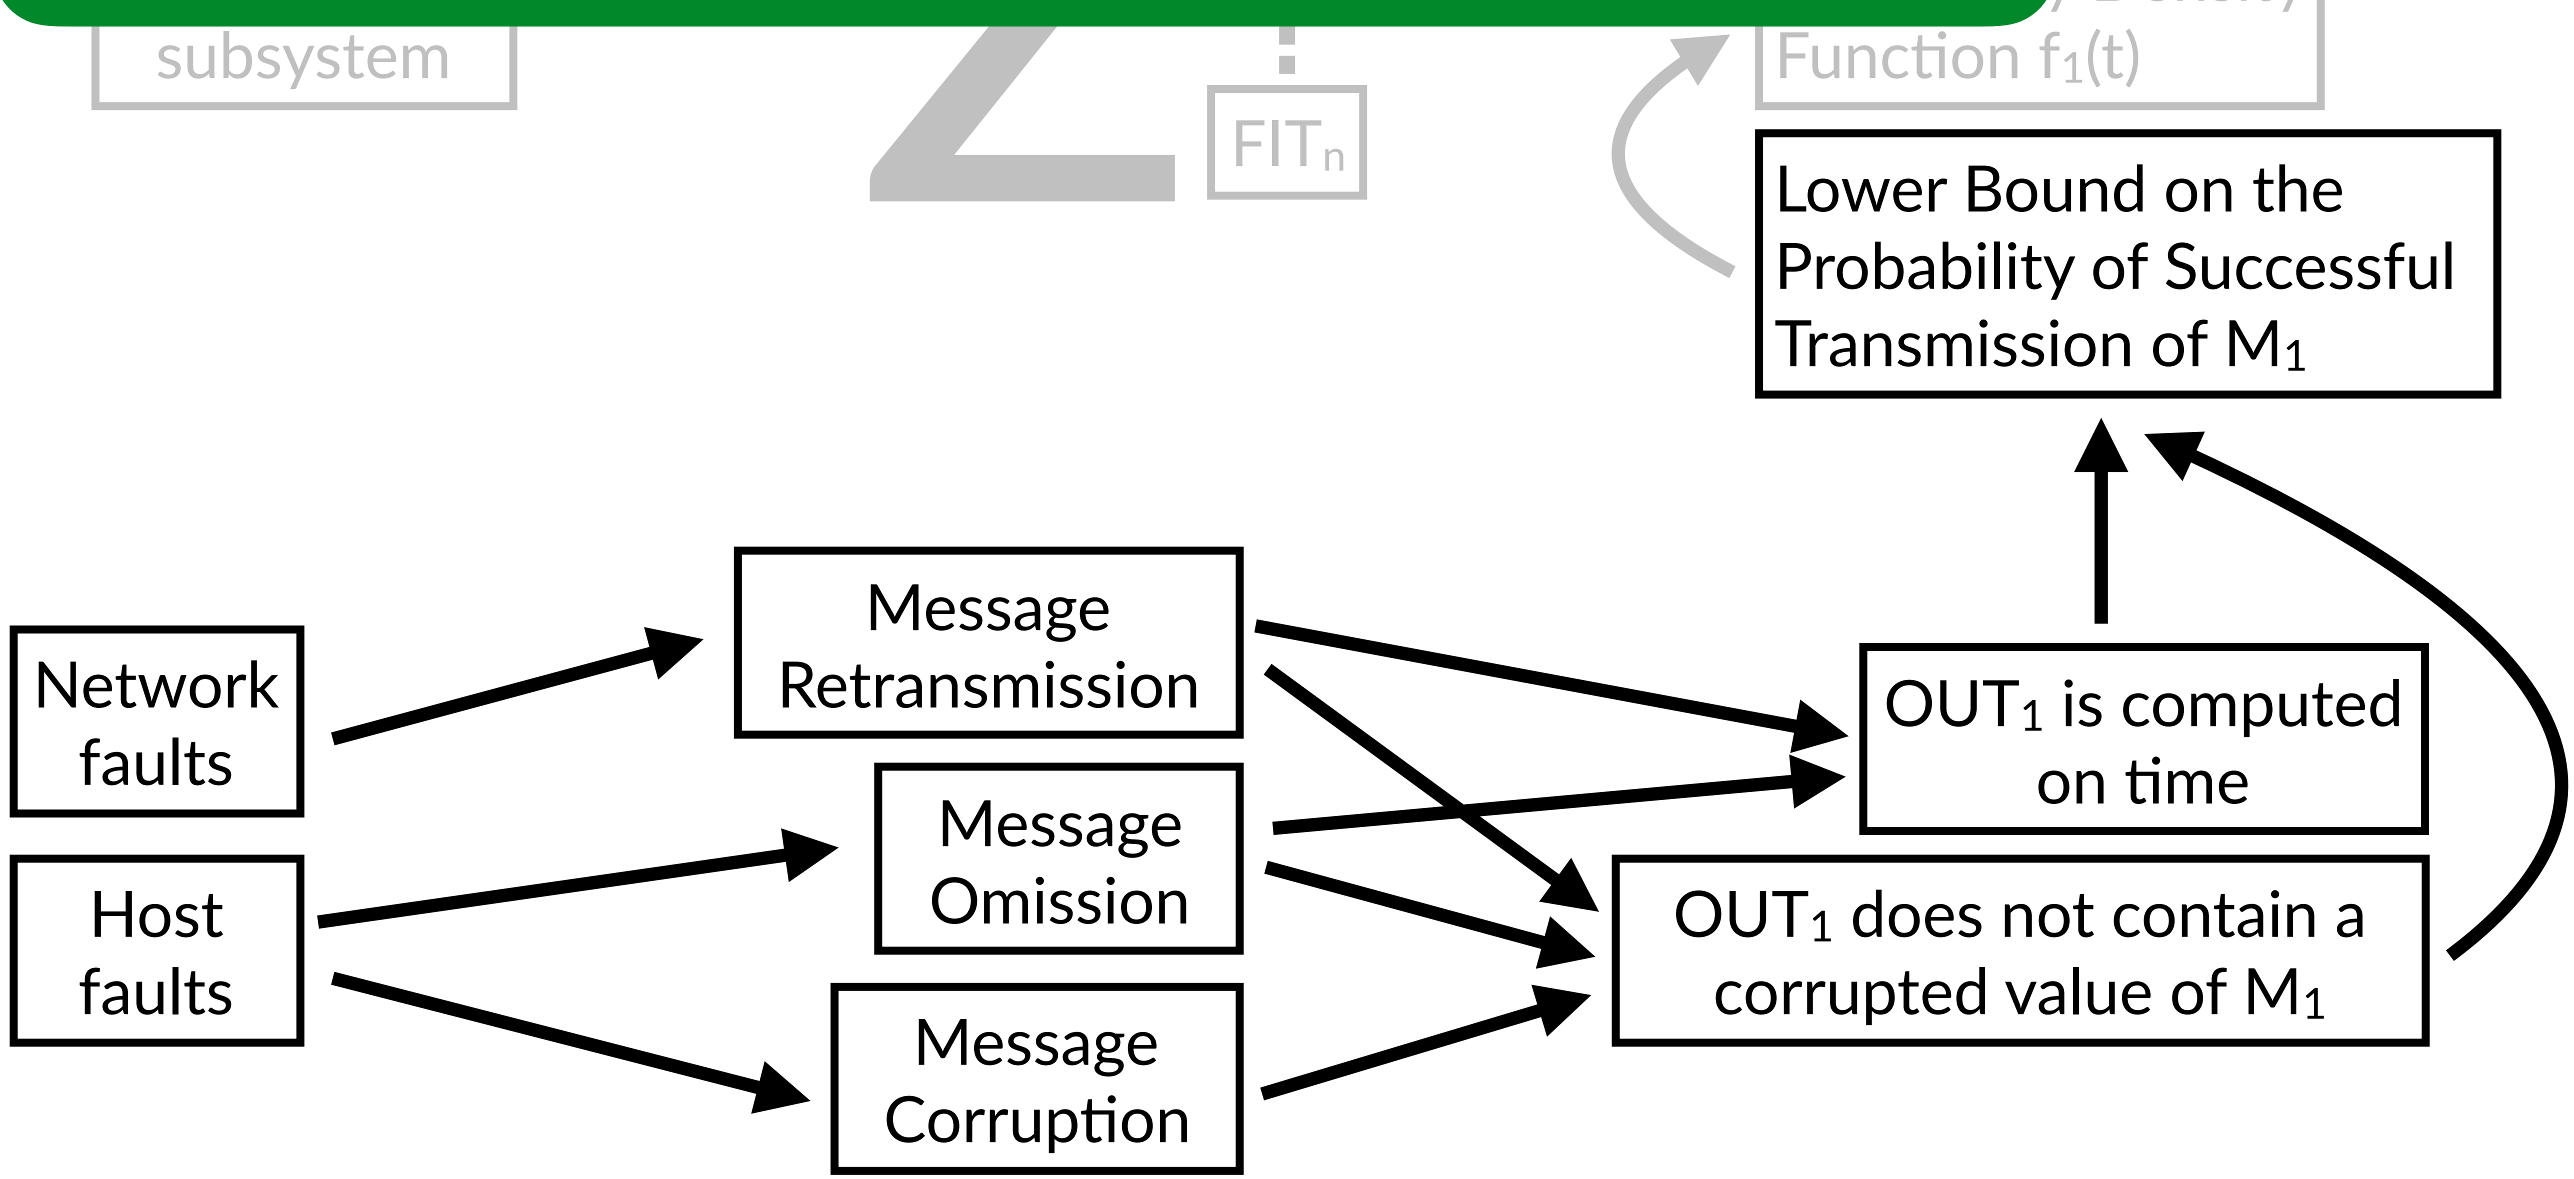

*Broster, Ian, Alan Burns, and Guillermo Rodriguez-Navas. **"Timing analysis of real-time communication under electromagnetic interference."** Real-Time Systems 30.1-2 (2005): 55-81.

FIT Rate Analysis of the System

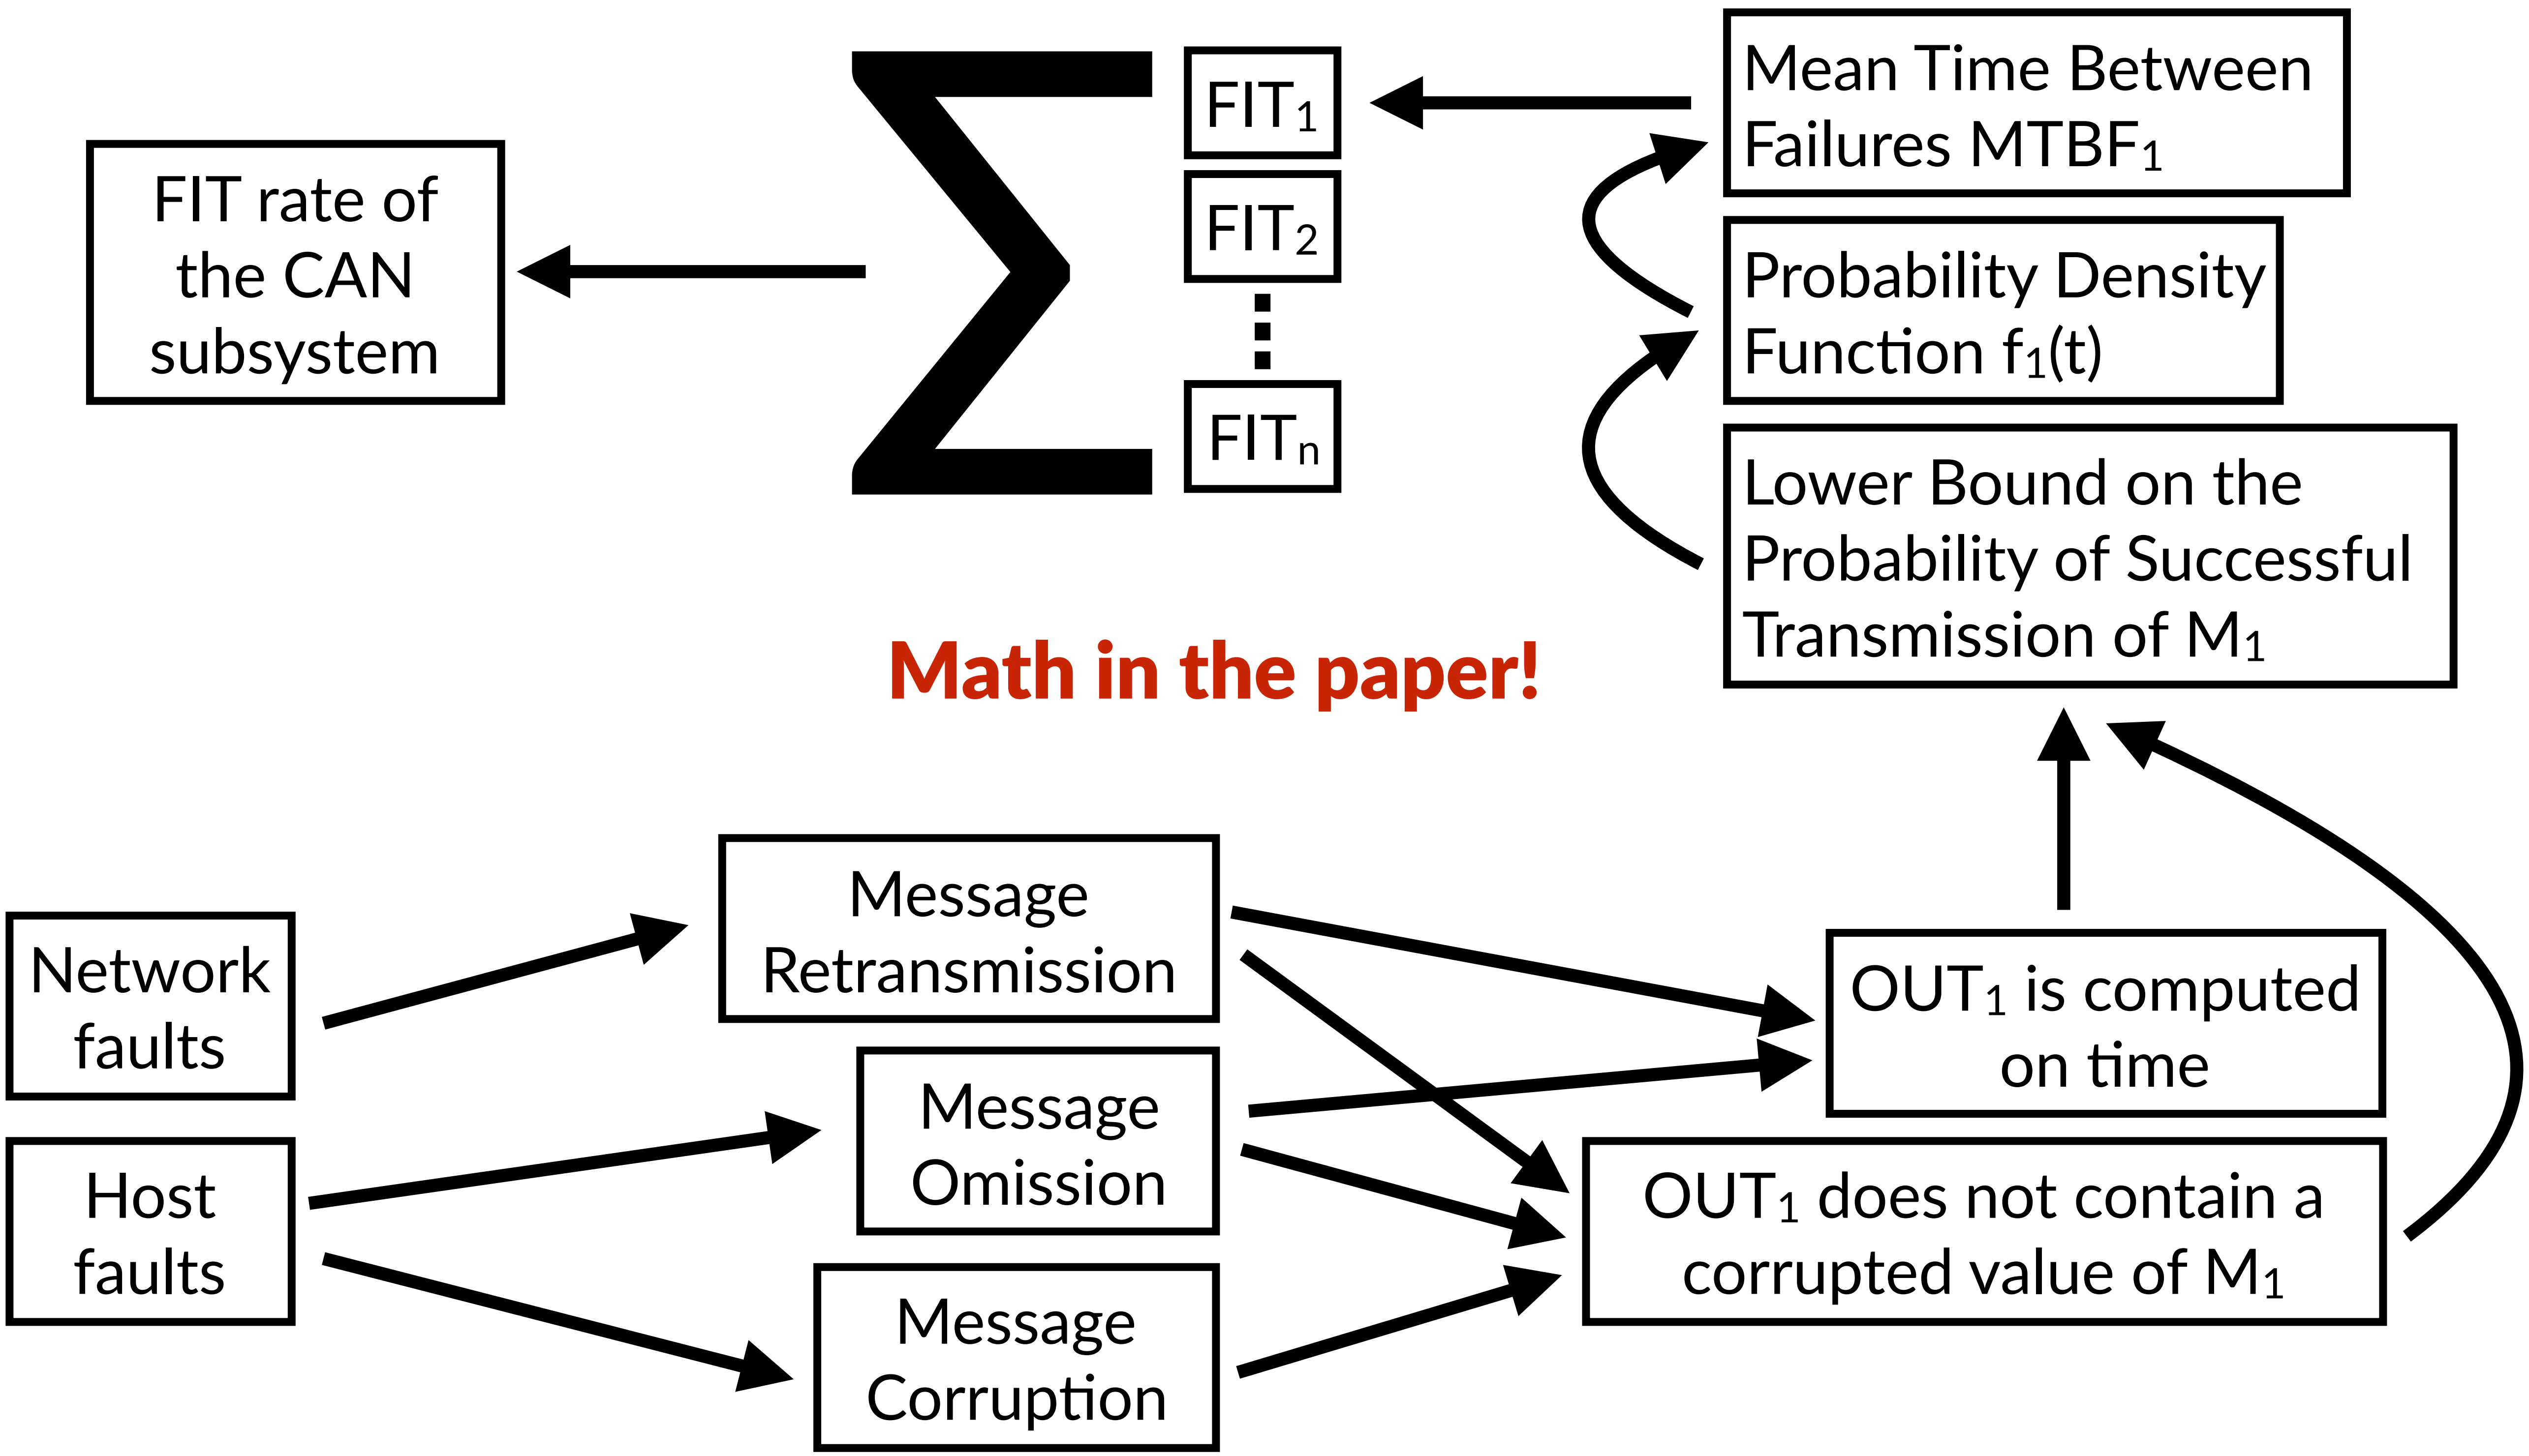

Math in the paper!

Overview

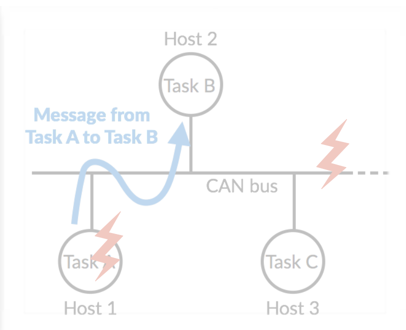

Model

$$\sum_{\mathbb{H}'\subseteq\mathbb{H}} \Phi^{\mathbb{H}'}_{crash} \cdot \sum_{\mathbb{M}'_1\subseteq\mathbb{M}_1} \left(\Phi^{\mathbb{H}',\,\mathbb{M}'_1}_{timely} \cdot \Phi^{\mathbb{H}',\,\mathbb{M}'_1}_{correct}\right)$$

Analysis

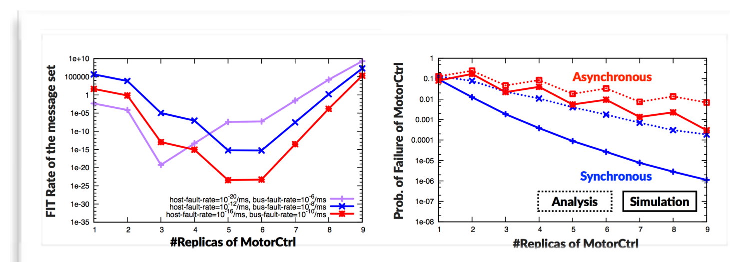

Evaluation

Mobile Robot Workload*

Task Name	Length (bytes)	Period (ms)	Deadline (ms)
MotorCtrl	2	2	2
Wheel1	3	4	4
Wheel2	3	4	4
RadioIn	8	8	8
Proximity	1	12	12
Logging	8	240	240

Broster, Ian, Alan Burns, and Guillermo Rodriguez-Navas. **"Comparing real-time communication under electromagnetic interference."** Real-Time Systems, 2004. ECRTS 2004. Proceedings. 16th Euromicro Conference on. IEEE, 2004.

Mobile Robot Workload*

Only the **MotorCtrl task is replicated** (#replicas vary from 1 to 9)

Task Name	Length (bytes)	Period (ms)	Deadline (ms)
MotorCtrl	2	2	2
Wheel1	3	4	4
Wheel2	3	4	4
RadioIn	8	8	8
Proximity	1	12	12
Logging	8	240	240

Broster, Ian, Alan Burns, and Guillermo Rodriguez-Navas. **"Comparing real-time communication under electromagnetic interference."** Real-Time Systems, 2004. ECRTS 2004. Proceedings. 16th Euromicro Conference on. IEEE, 2004.

Evaluation

- Assess the proposed FIT rate derivation
  - ➡ Comparison with results from CAN bus simulation

Evaluation

- Assess the proposed FIT rate derivation
  ➡ Comparison with results from CAN bus simulation

- Is the FIT rate analysis too coarse-grained?
  ➡ Analysis for various fault rates

Analysis versus Simulation for MotorCtrl

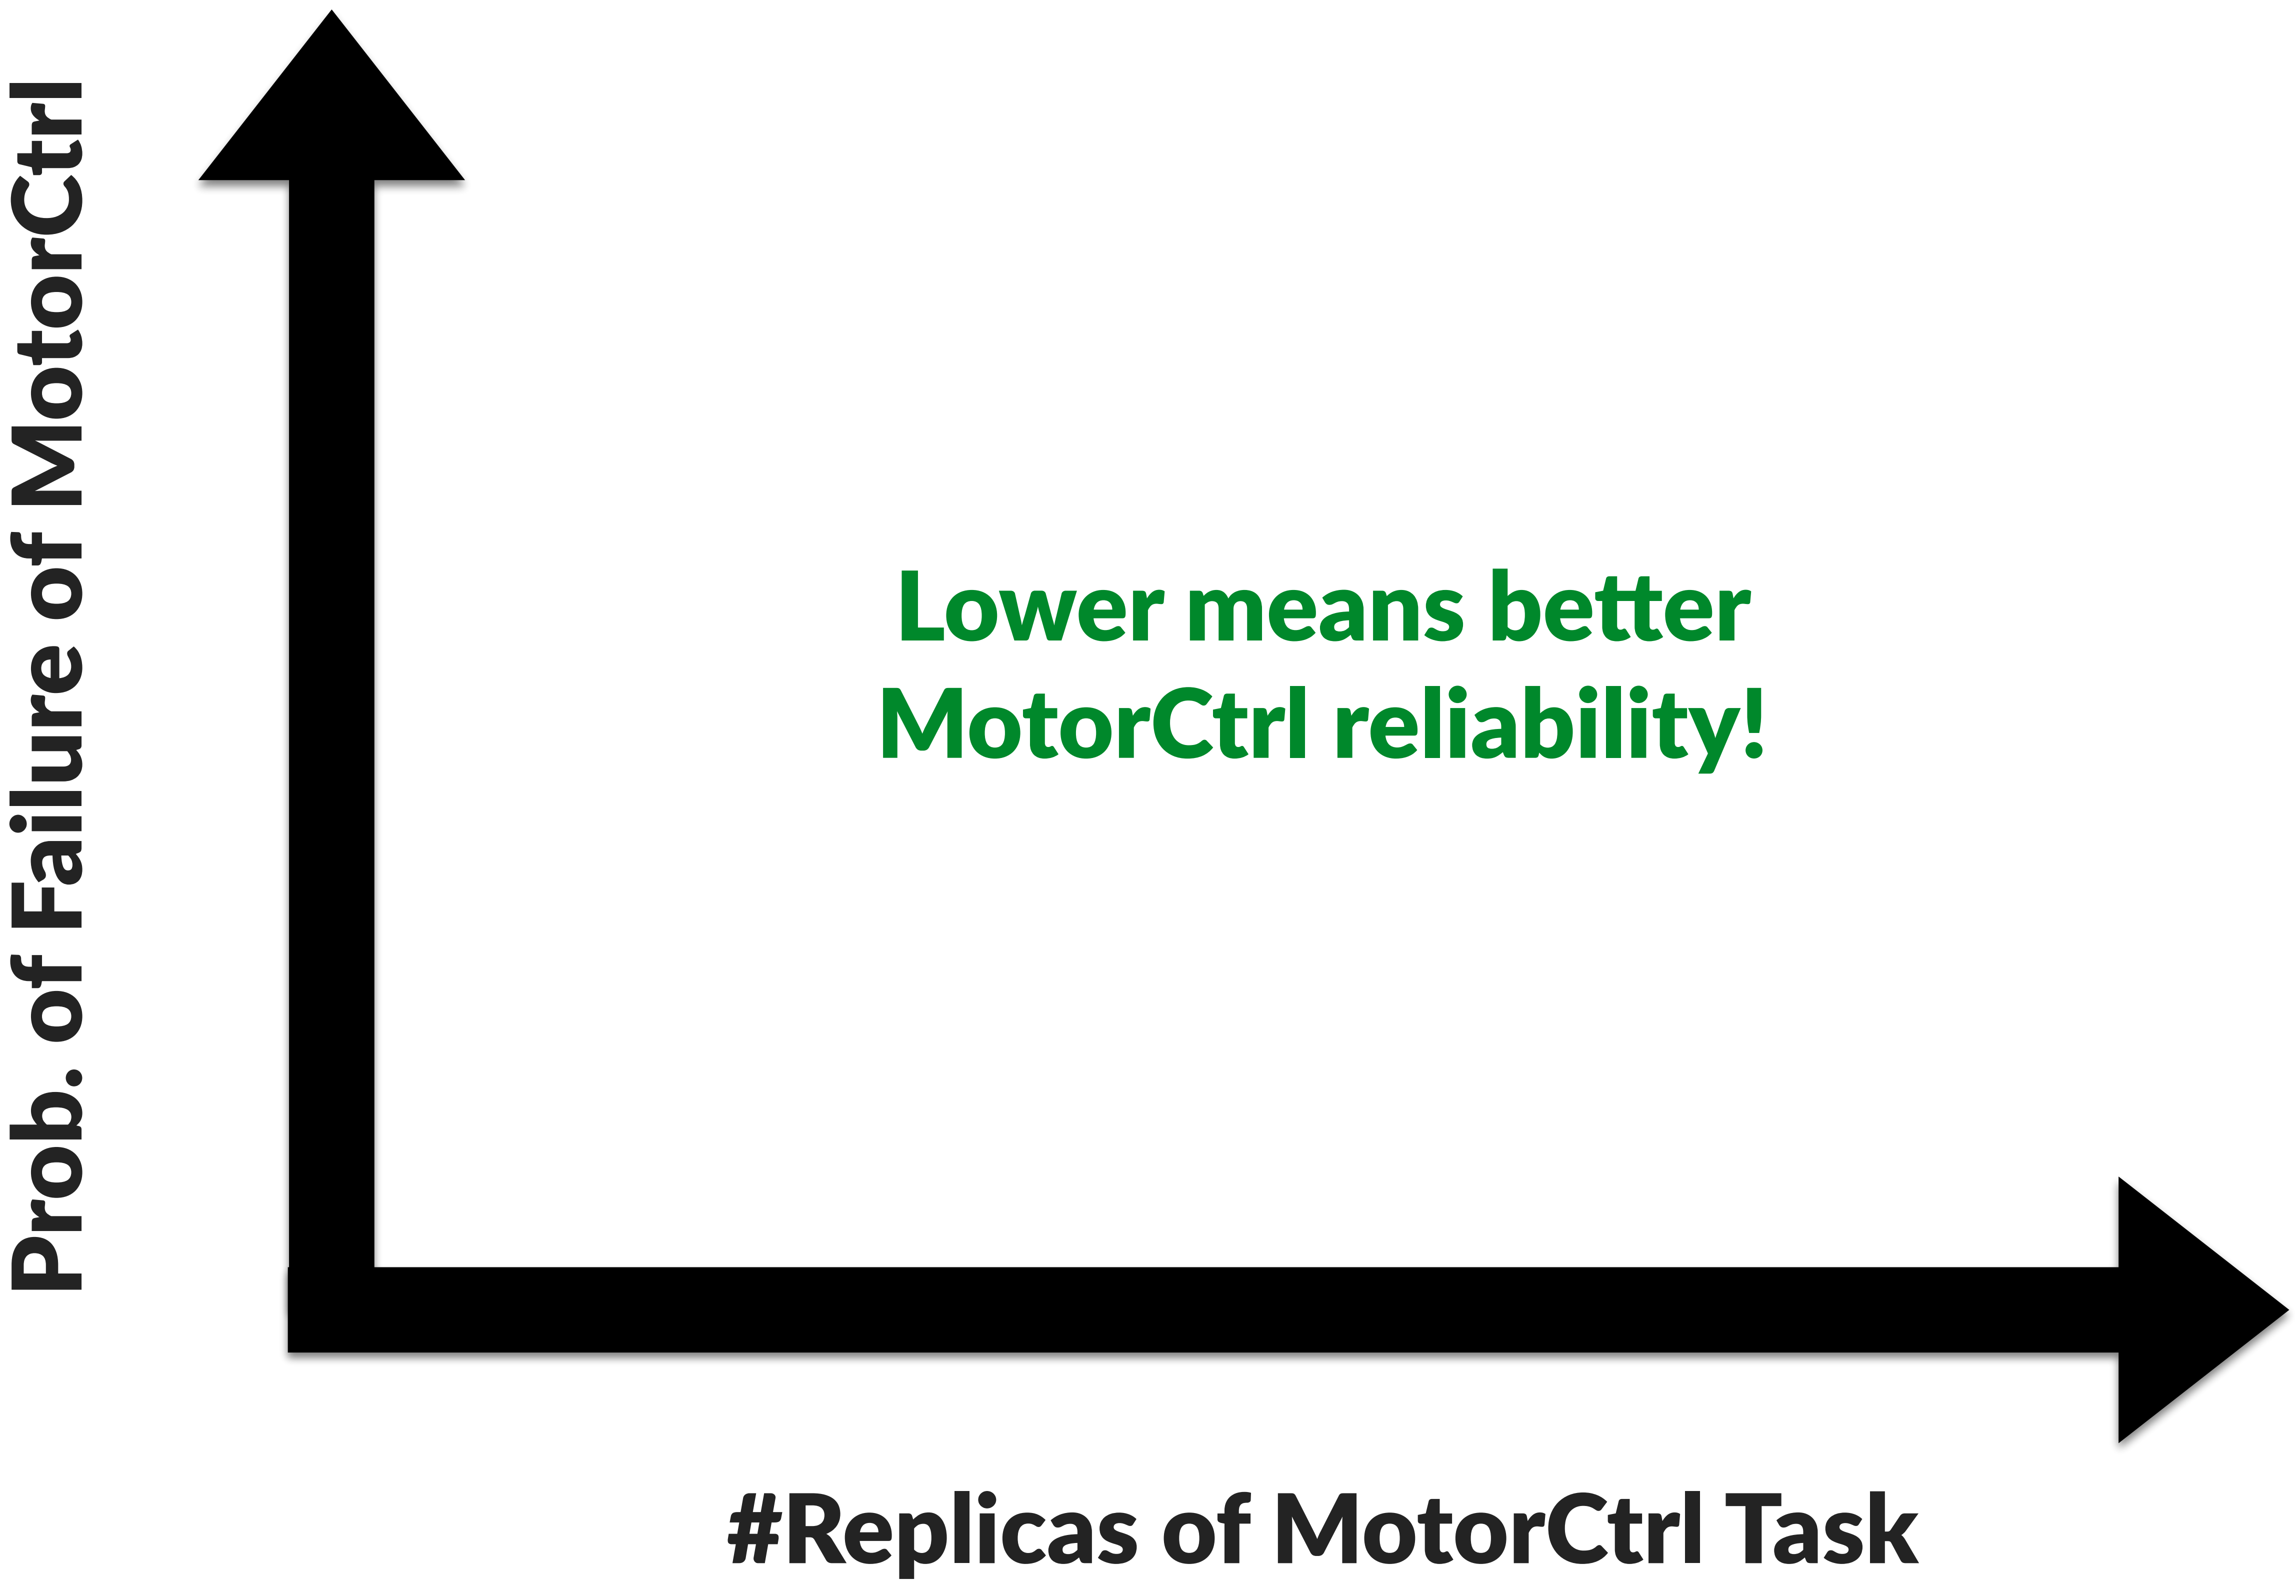

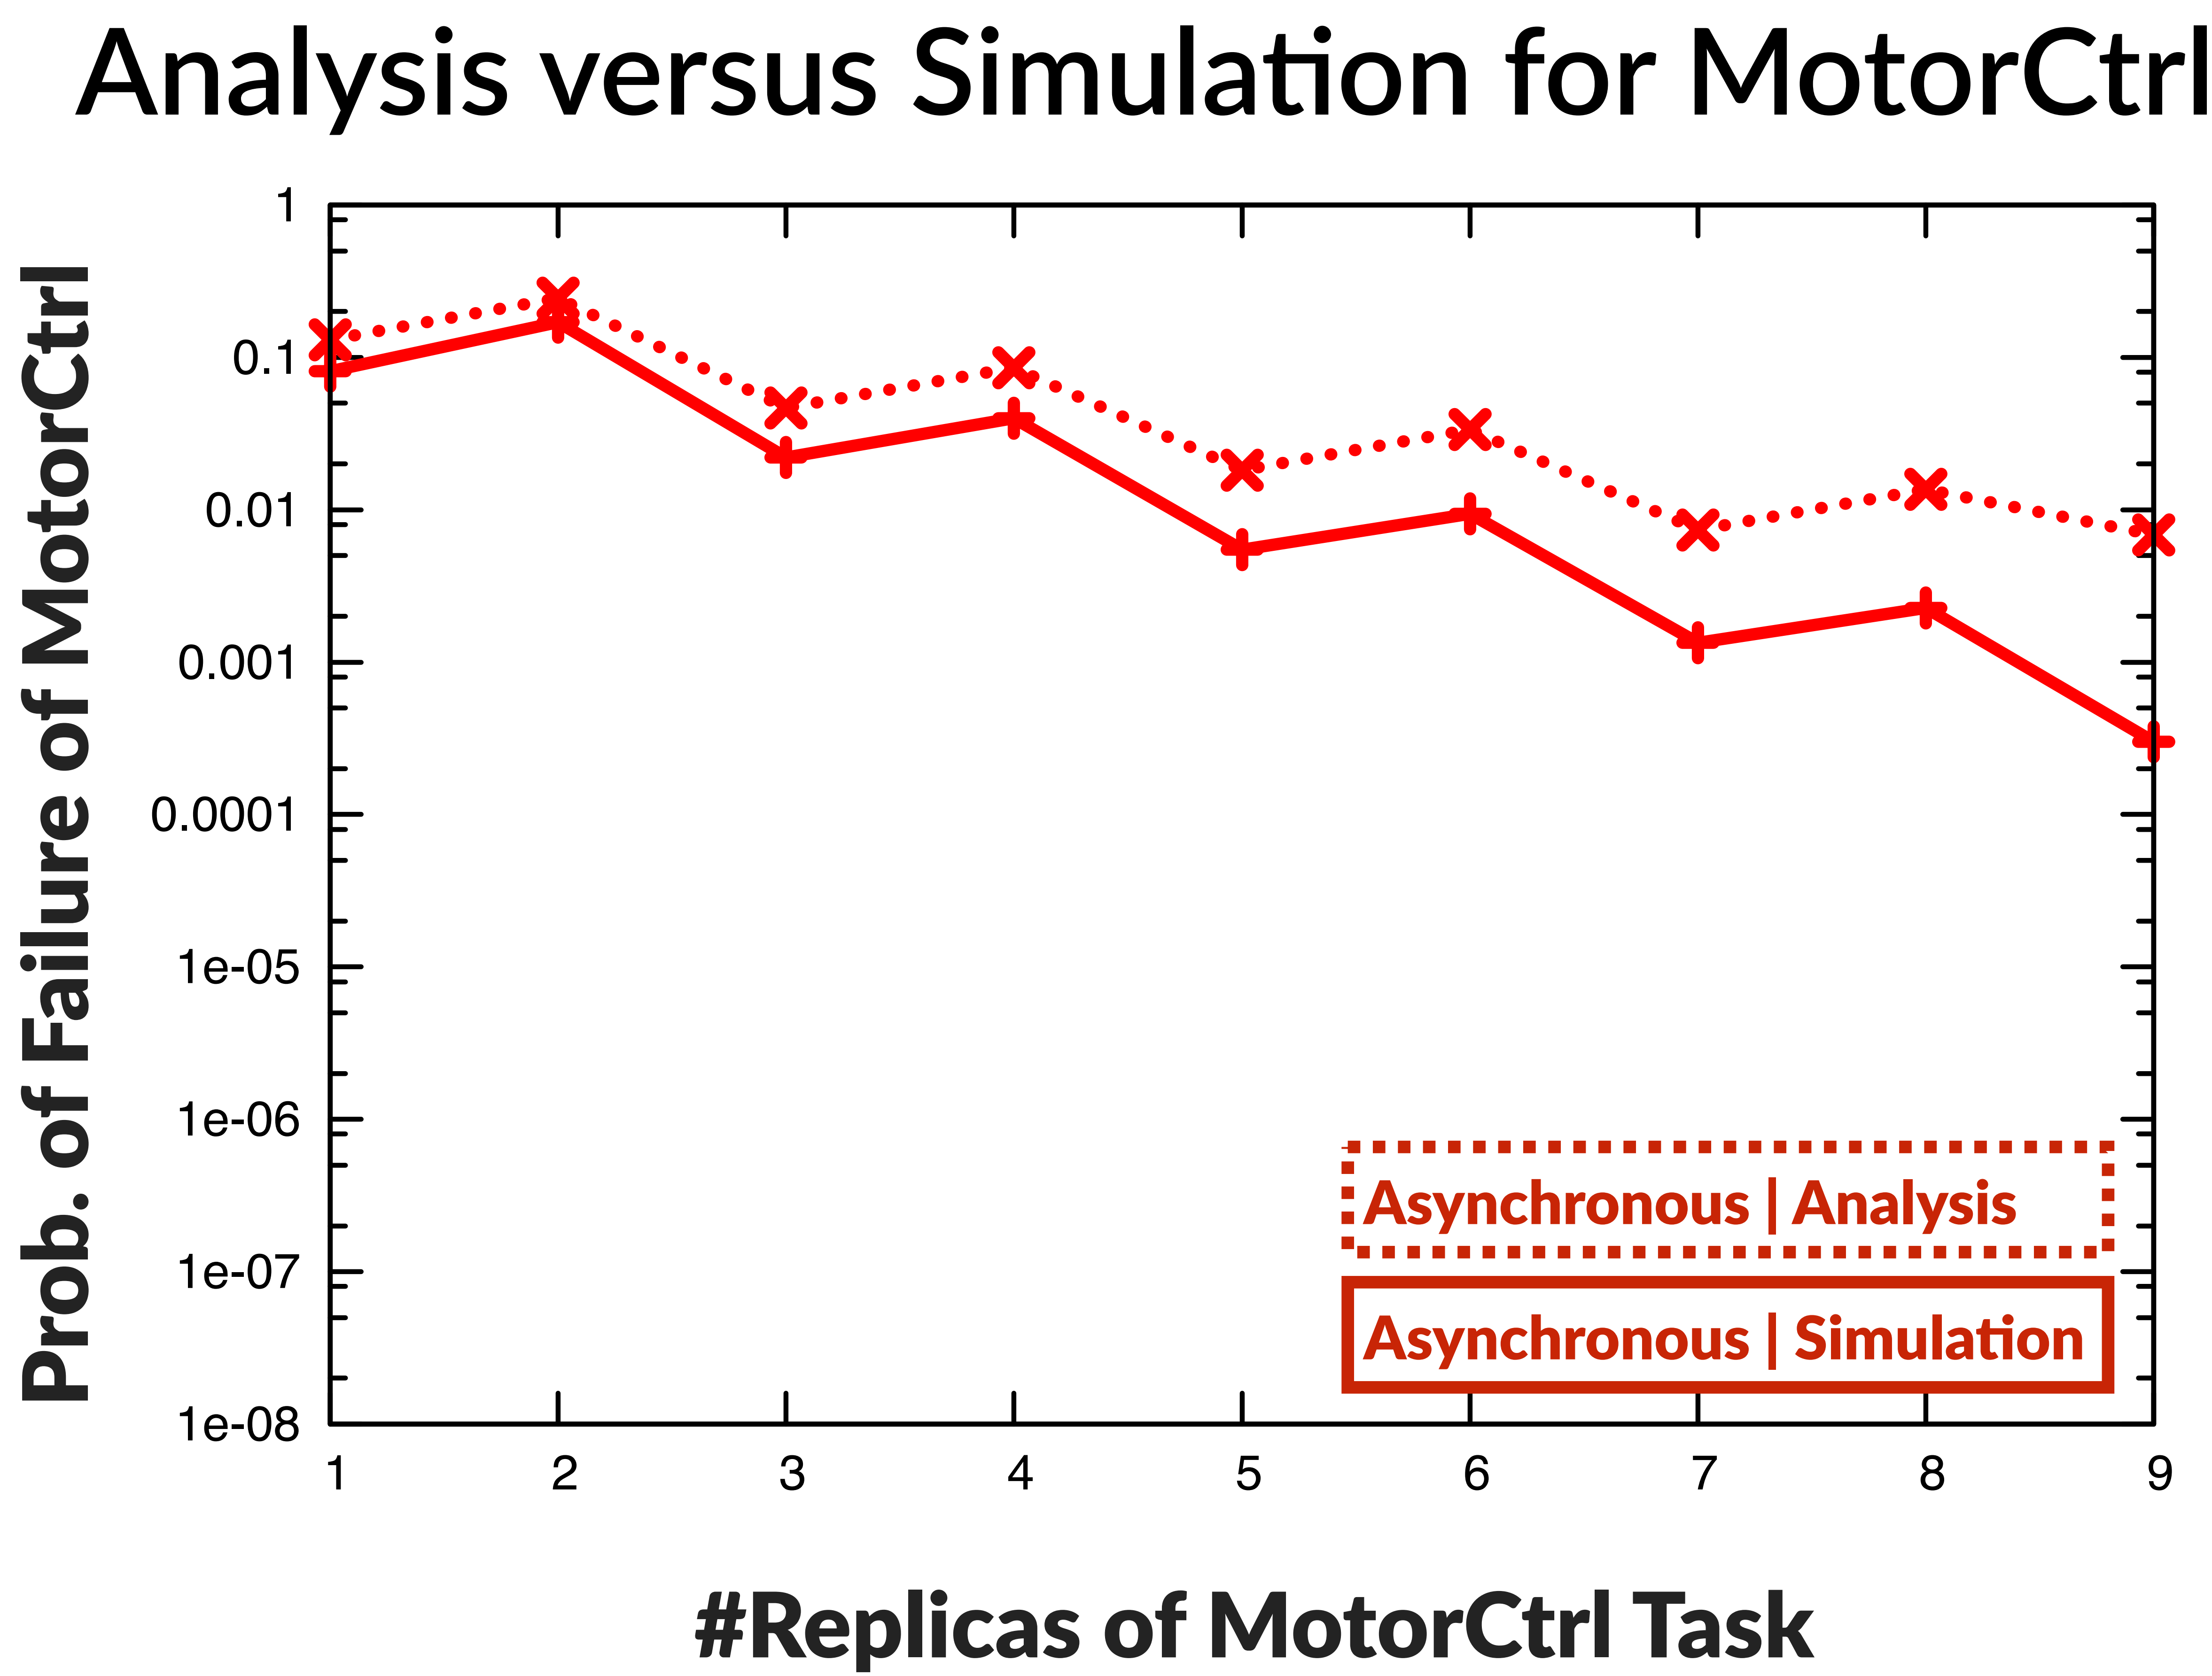
Analysis versus Simulation for MotorCtrl
Prob. of Failure of MotorCtrl
1
0.1
0.01
0.001
0.0001
1e-05
1e-06
1e-07
1e-08
1
2
3
4
5
6
7
8
9
#Replicas of MotorCtrl Task
Asynchronous | Analysis
Asynchronous | Simulation

Analysis versus Simulation for MotorCtrl

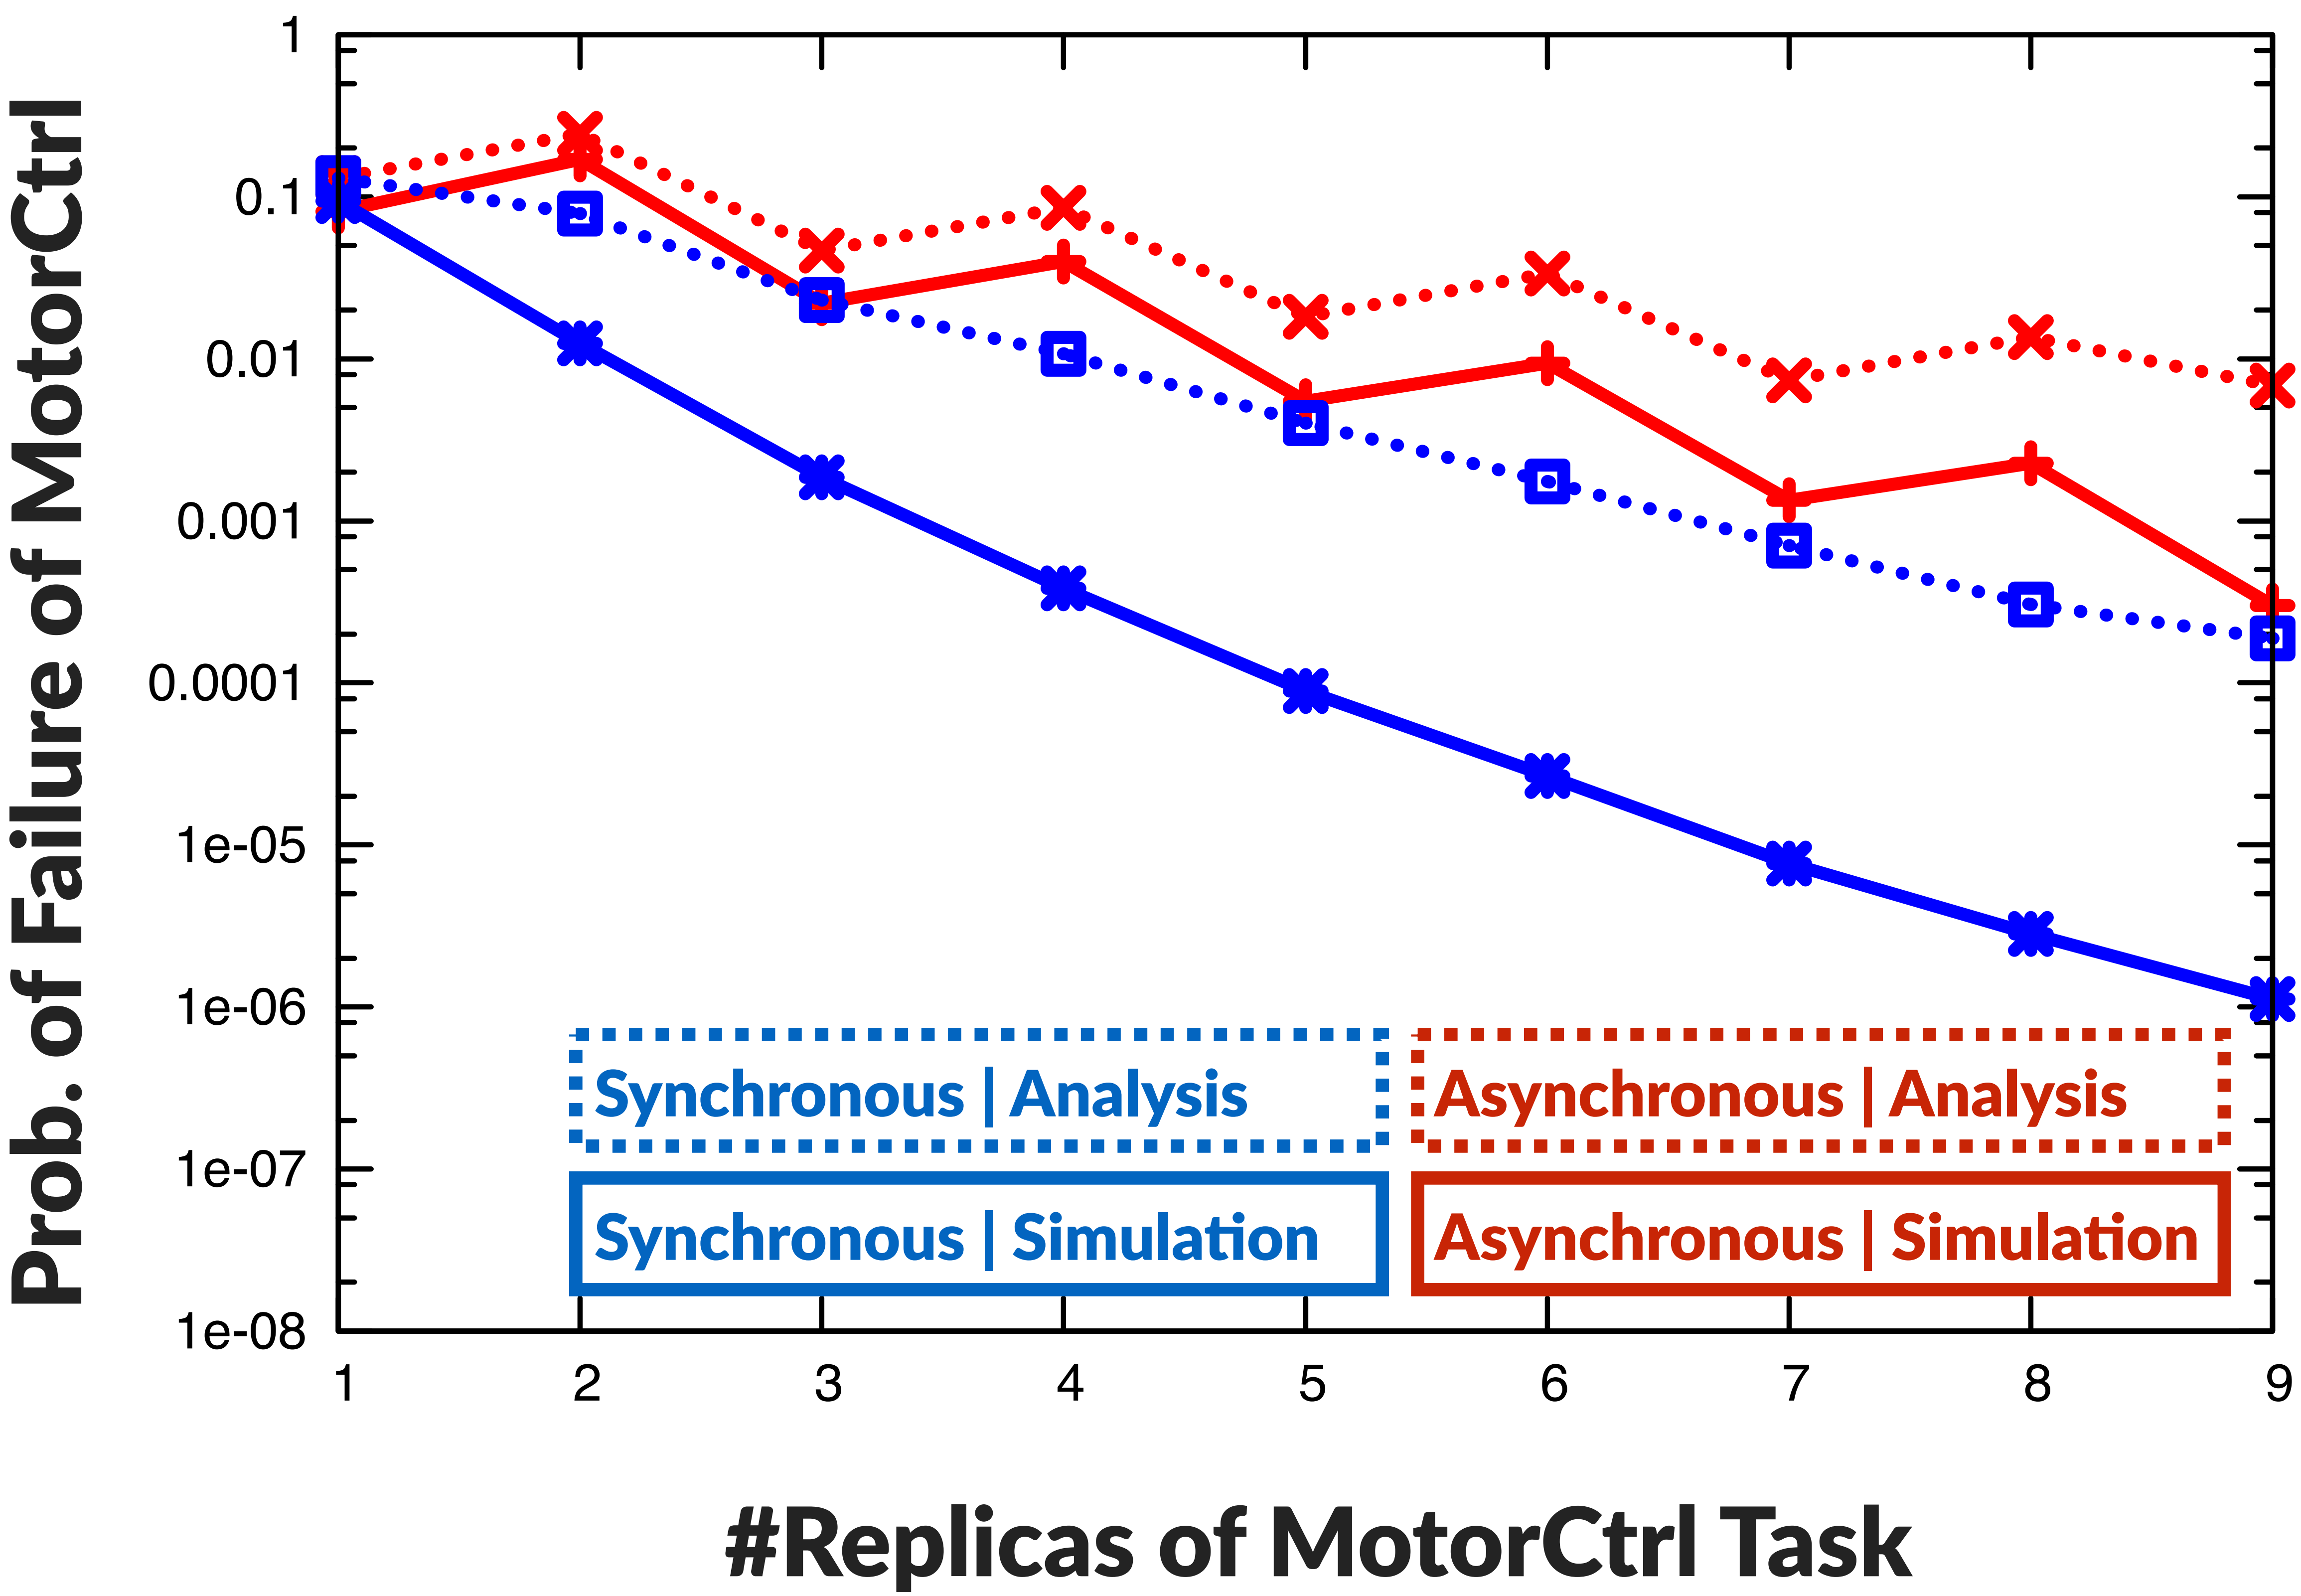

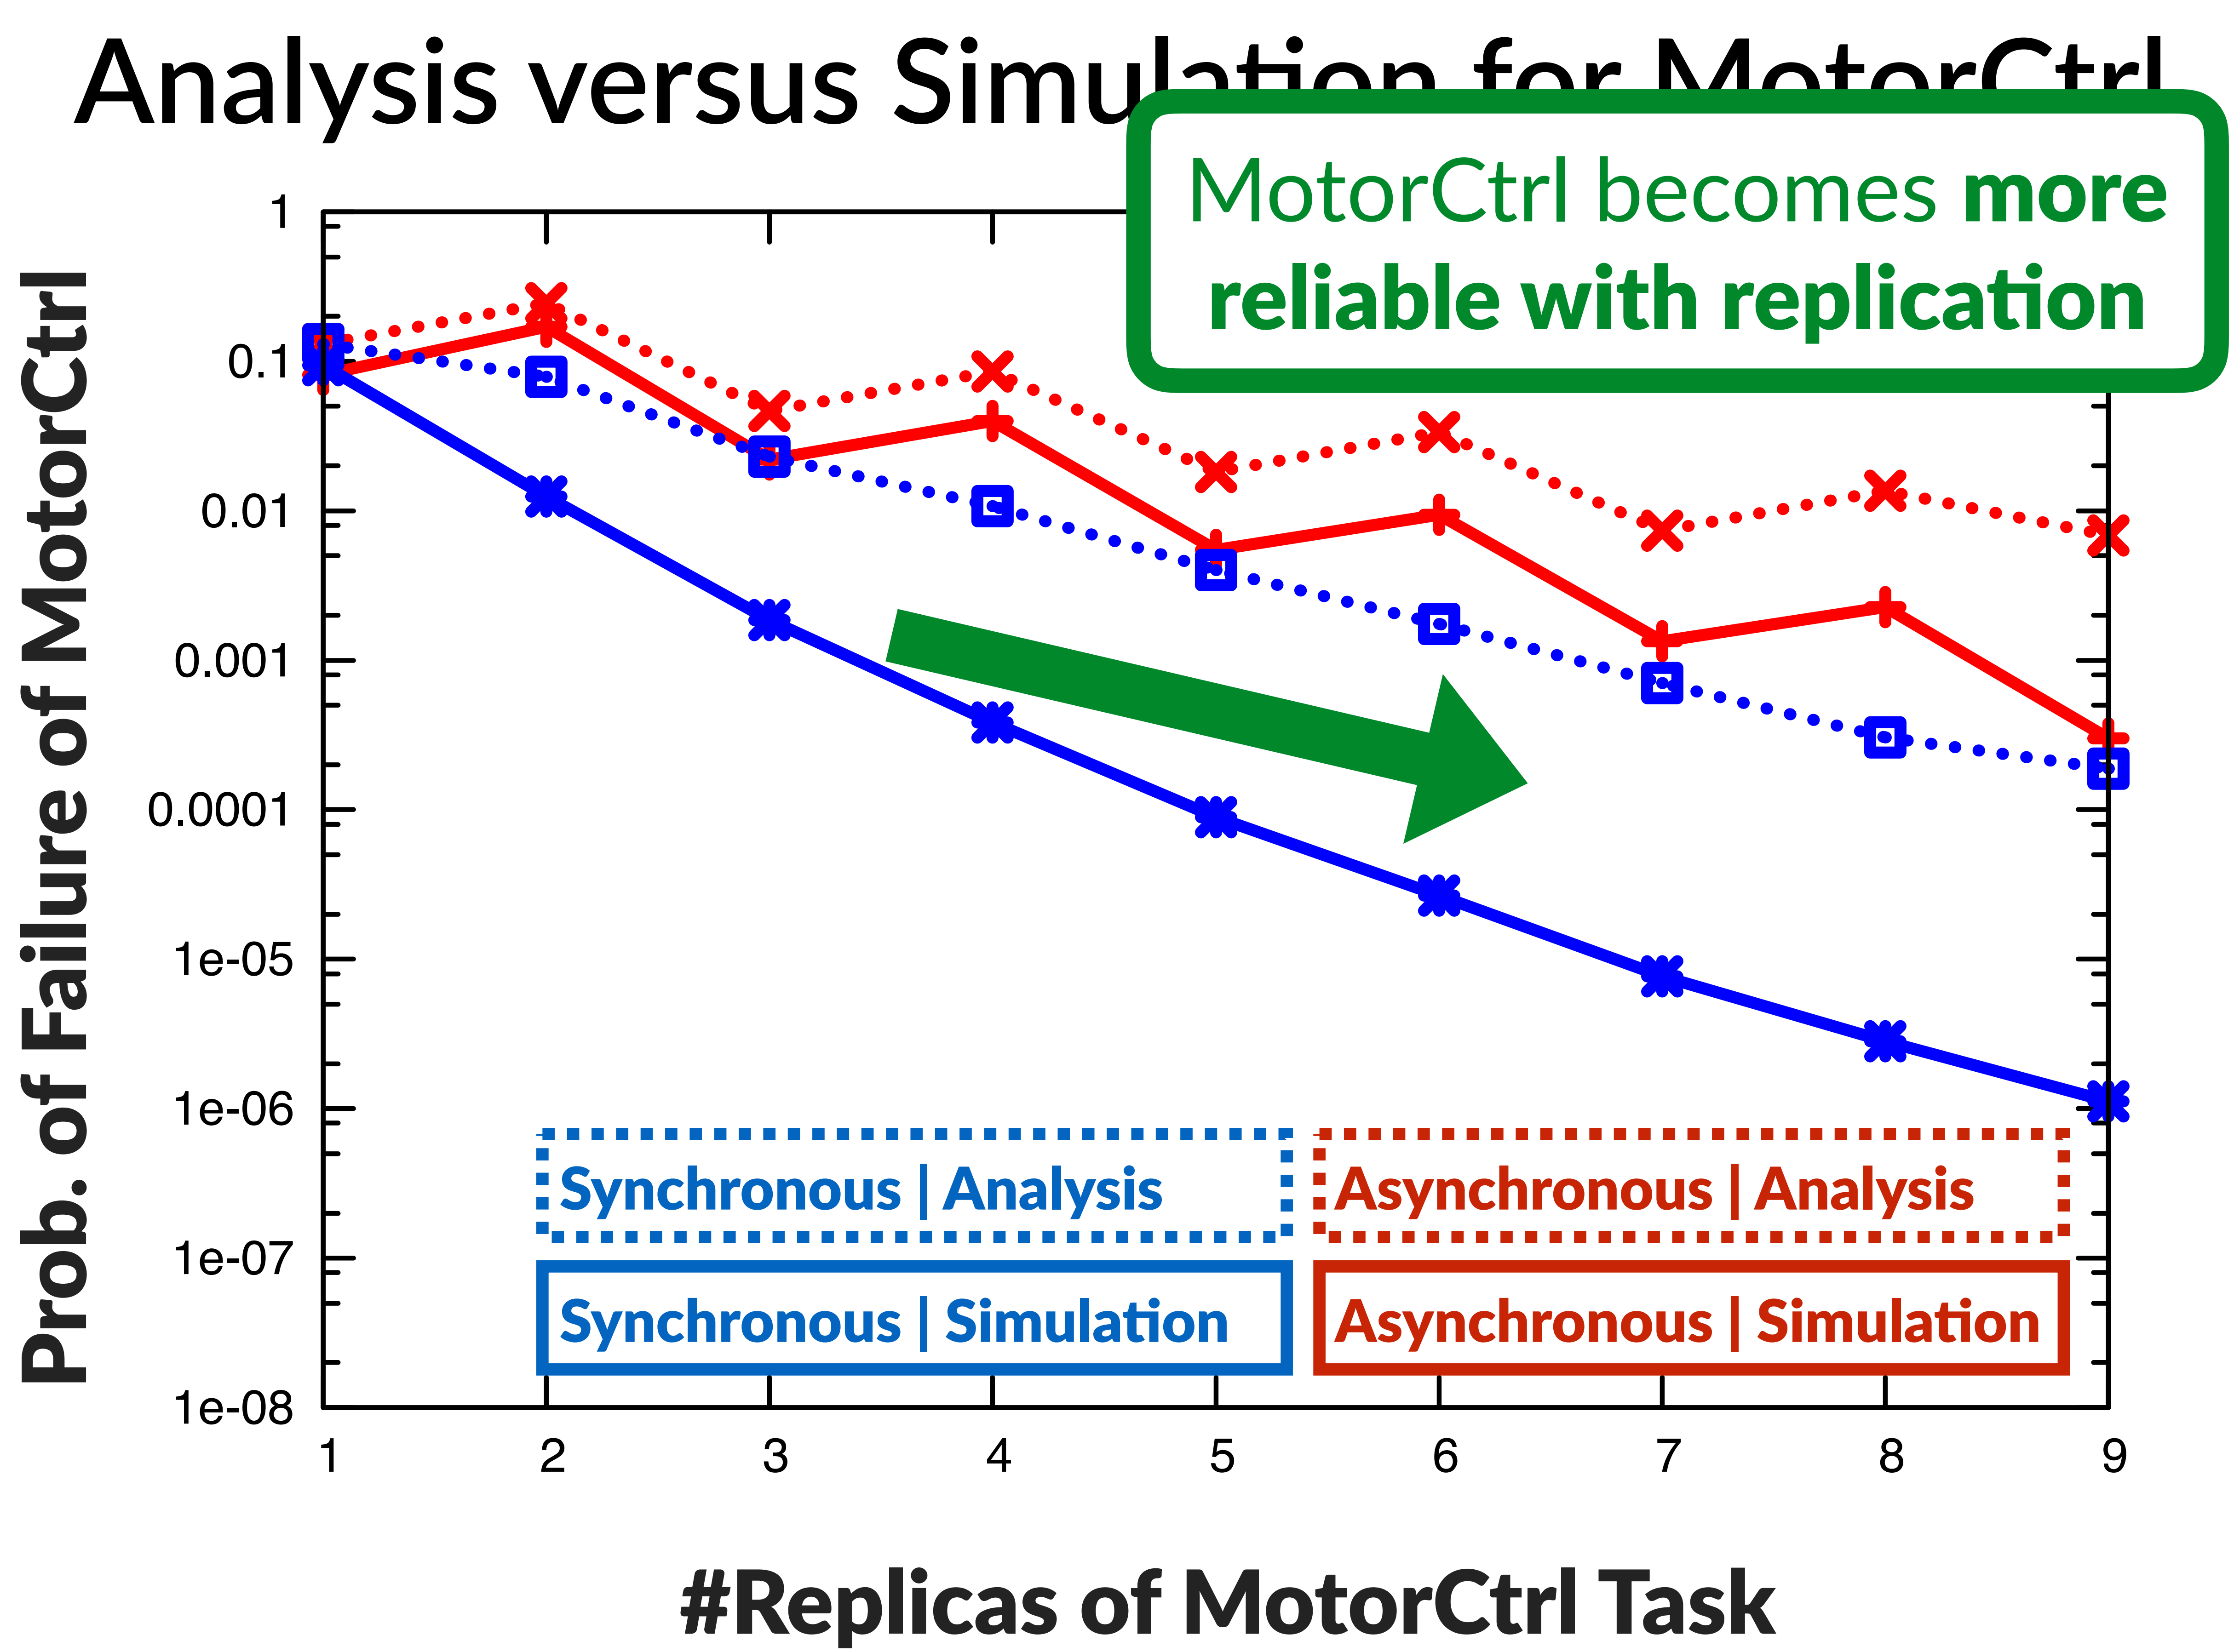

Analysis versus Simulation for MotorCtrl
MotorCtrl becomes more reliable with replication
Prob. of Failure of MotorCtrl
1
0.1
0.01
0.001
0.0001
1e-05
1e-06
1e-07
1e-08
1
2
3
4
5
6
7
8
9
#Replicas of MotorCtrl Task
Synchronous | Analysis
Asynchronous | Analysis
Synchronous | Simulation
Asynchronous | Simulation

Analysis versus Simulation for MotorCtrl

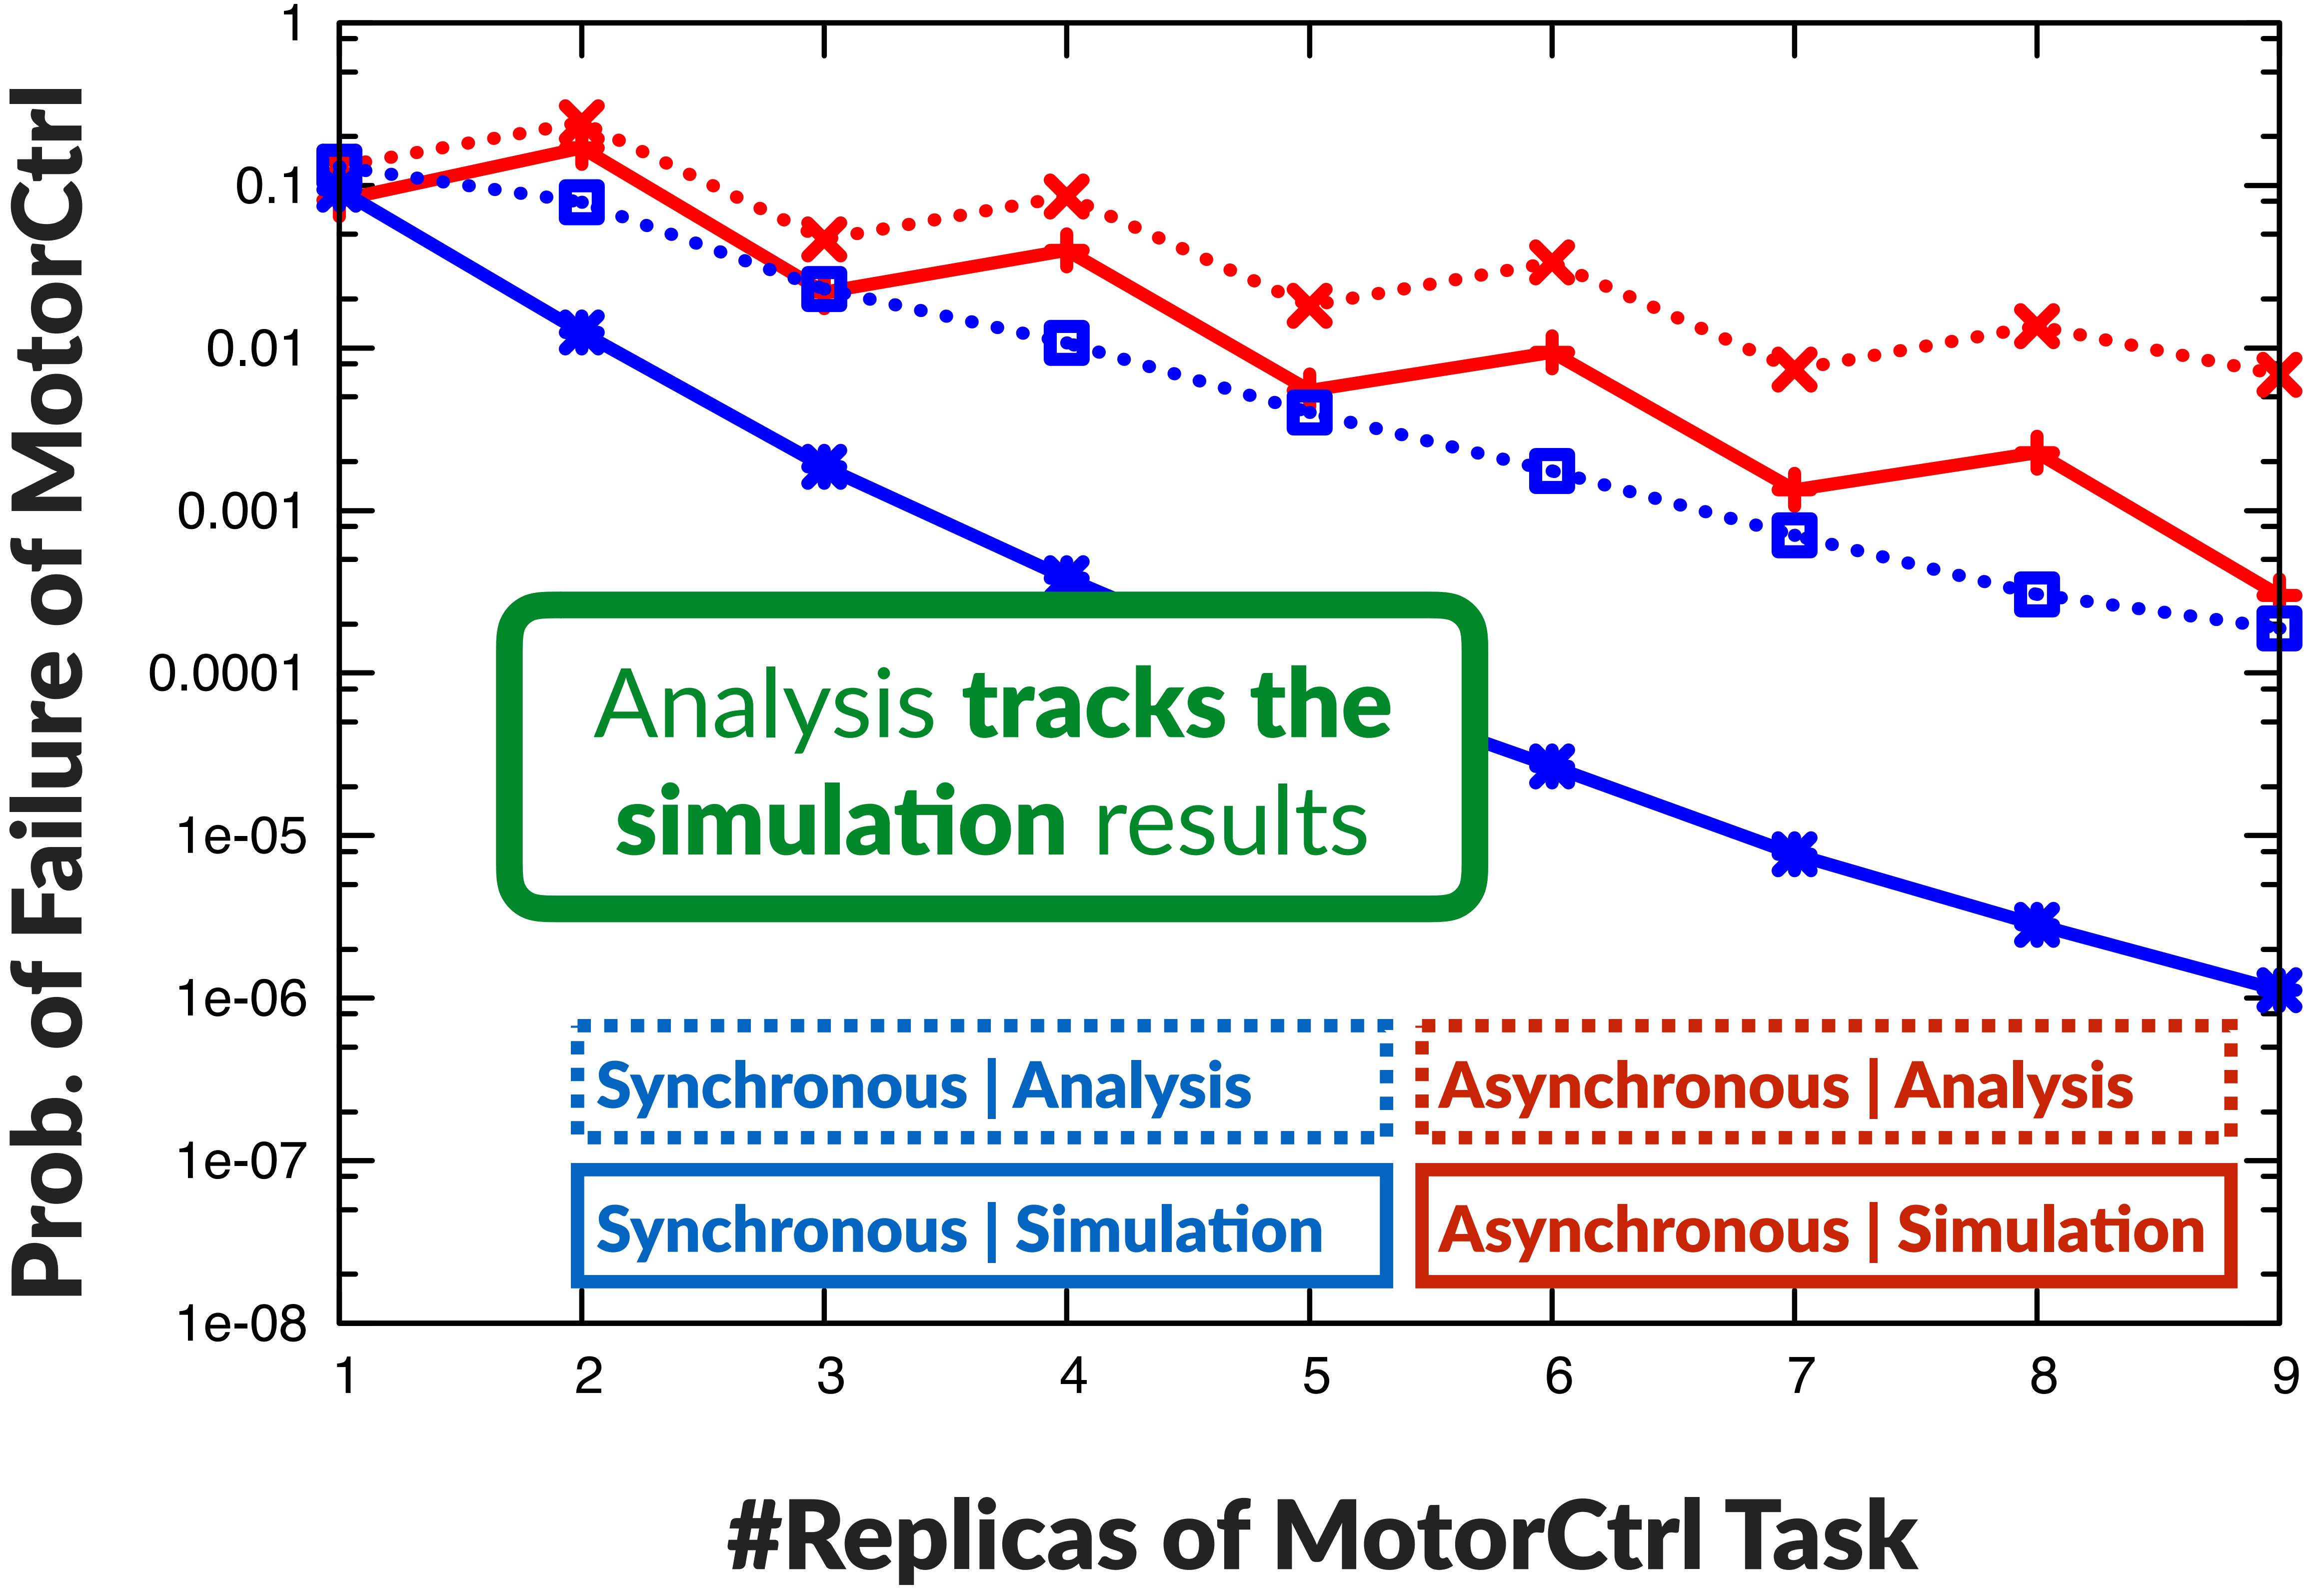

FIT Rate Analysis of the CAN Subsystem

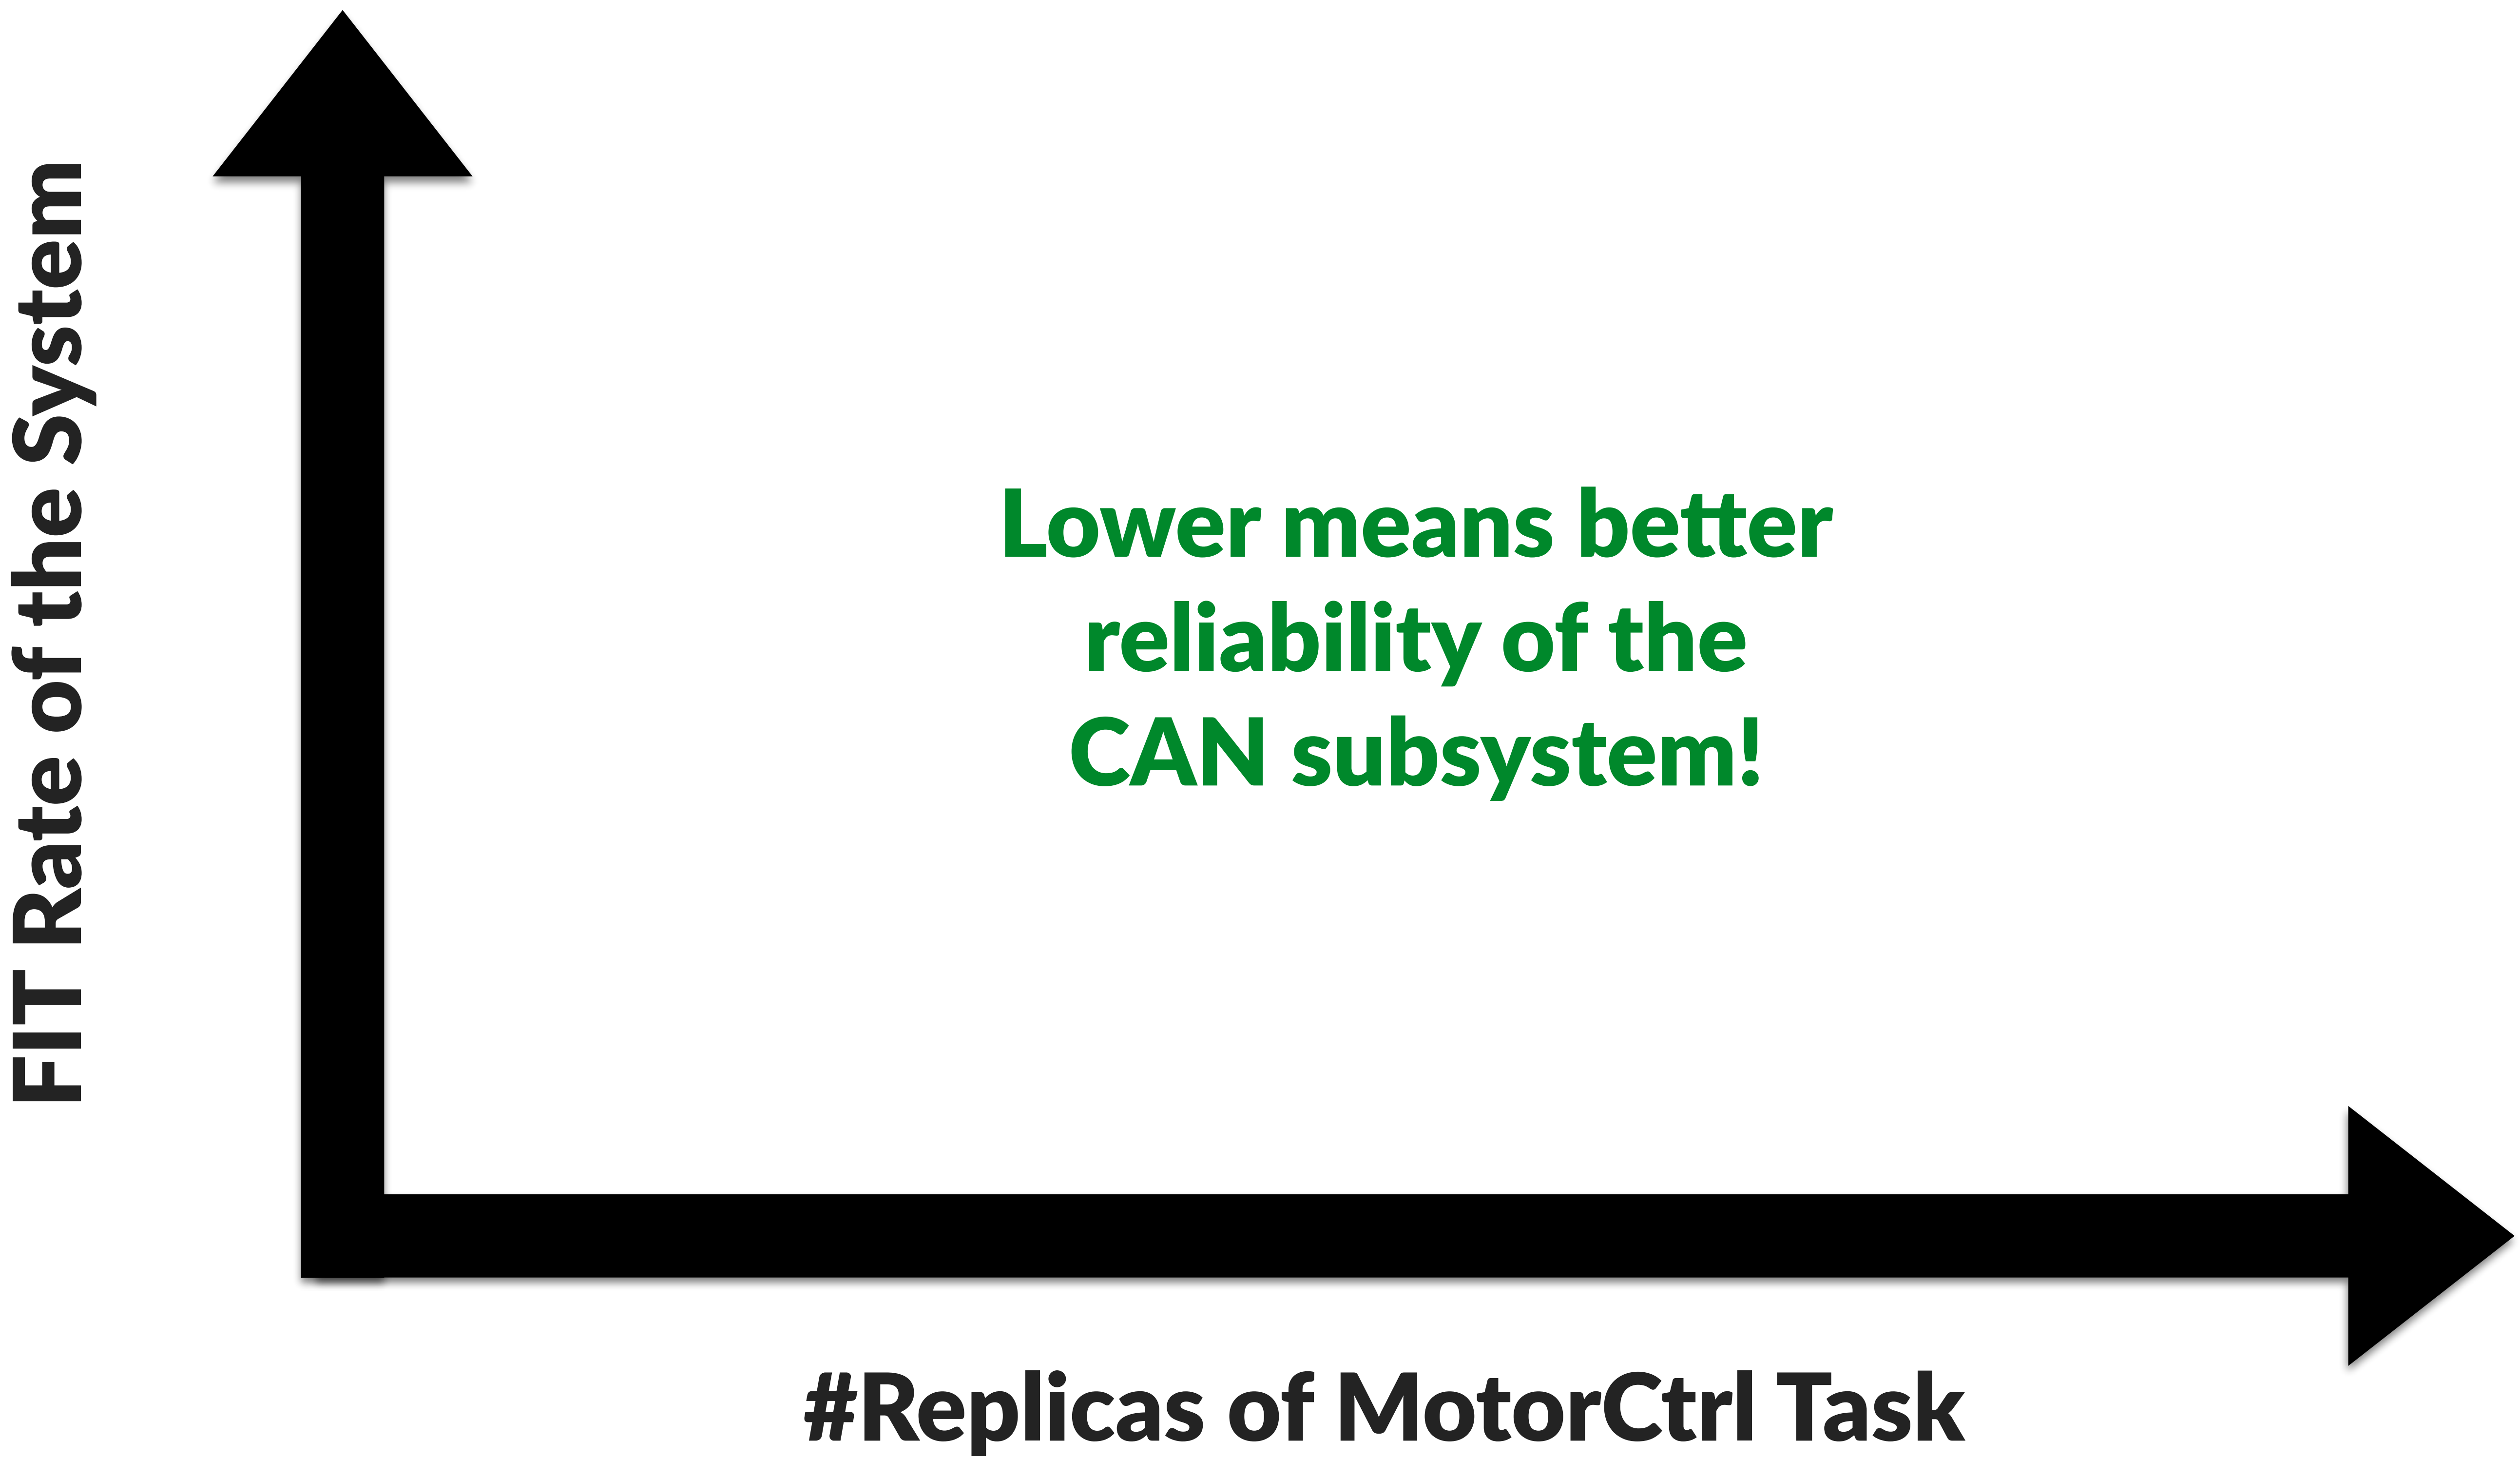

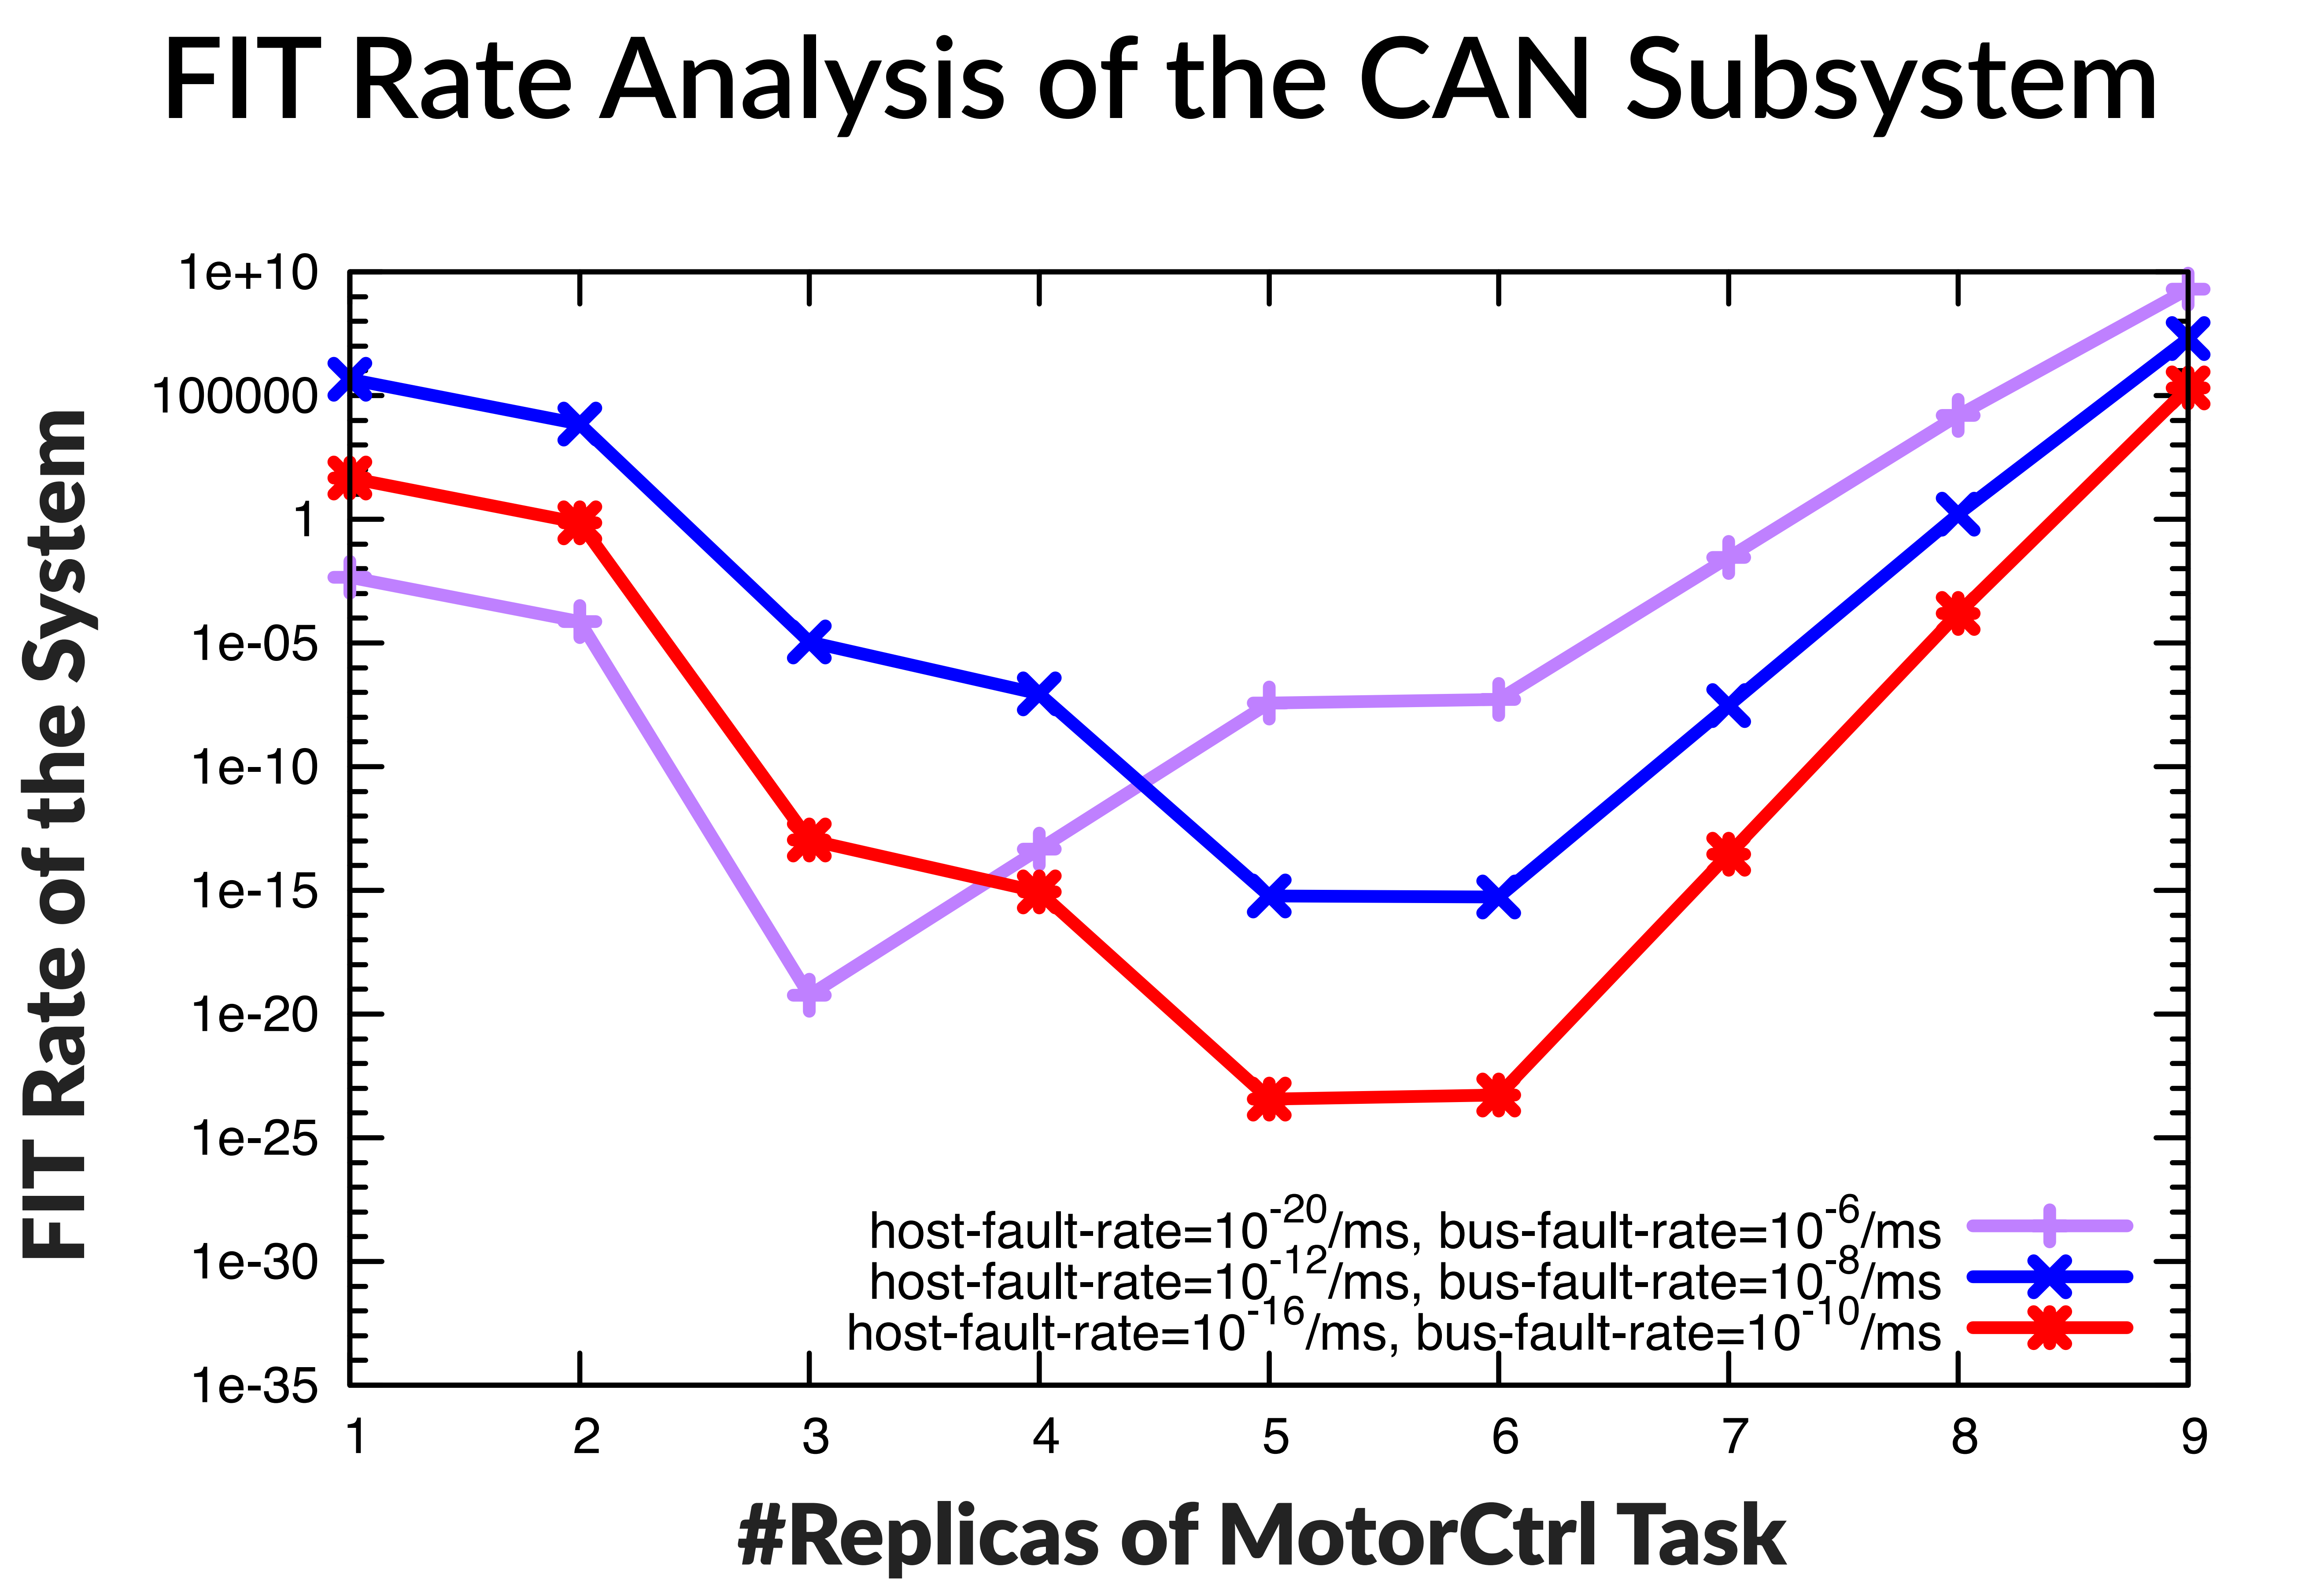

FIT Rate Analysis of the CAN Subsystem
FIT Rate of the System
#Replicas of MotorCtrl Task
1e+10
100000
1
1e-05
1e-10
1e-15
1e-20
1e-25
1e-30
1e-35
1
2
3
4
5
6
7
8
9
host-fault-rate=10^-20/ms, bus-fault-rate=10^-6/ms
host-fault-rate=10^-12/ms, bus-fault-rate=10^-8/ms
host-fault-rate=10^-16/ms, bus-fault-rate=10^-10/ms

FIT Rate Analysis of the CAN Subsystem

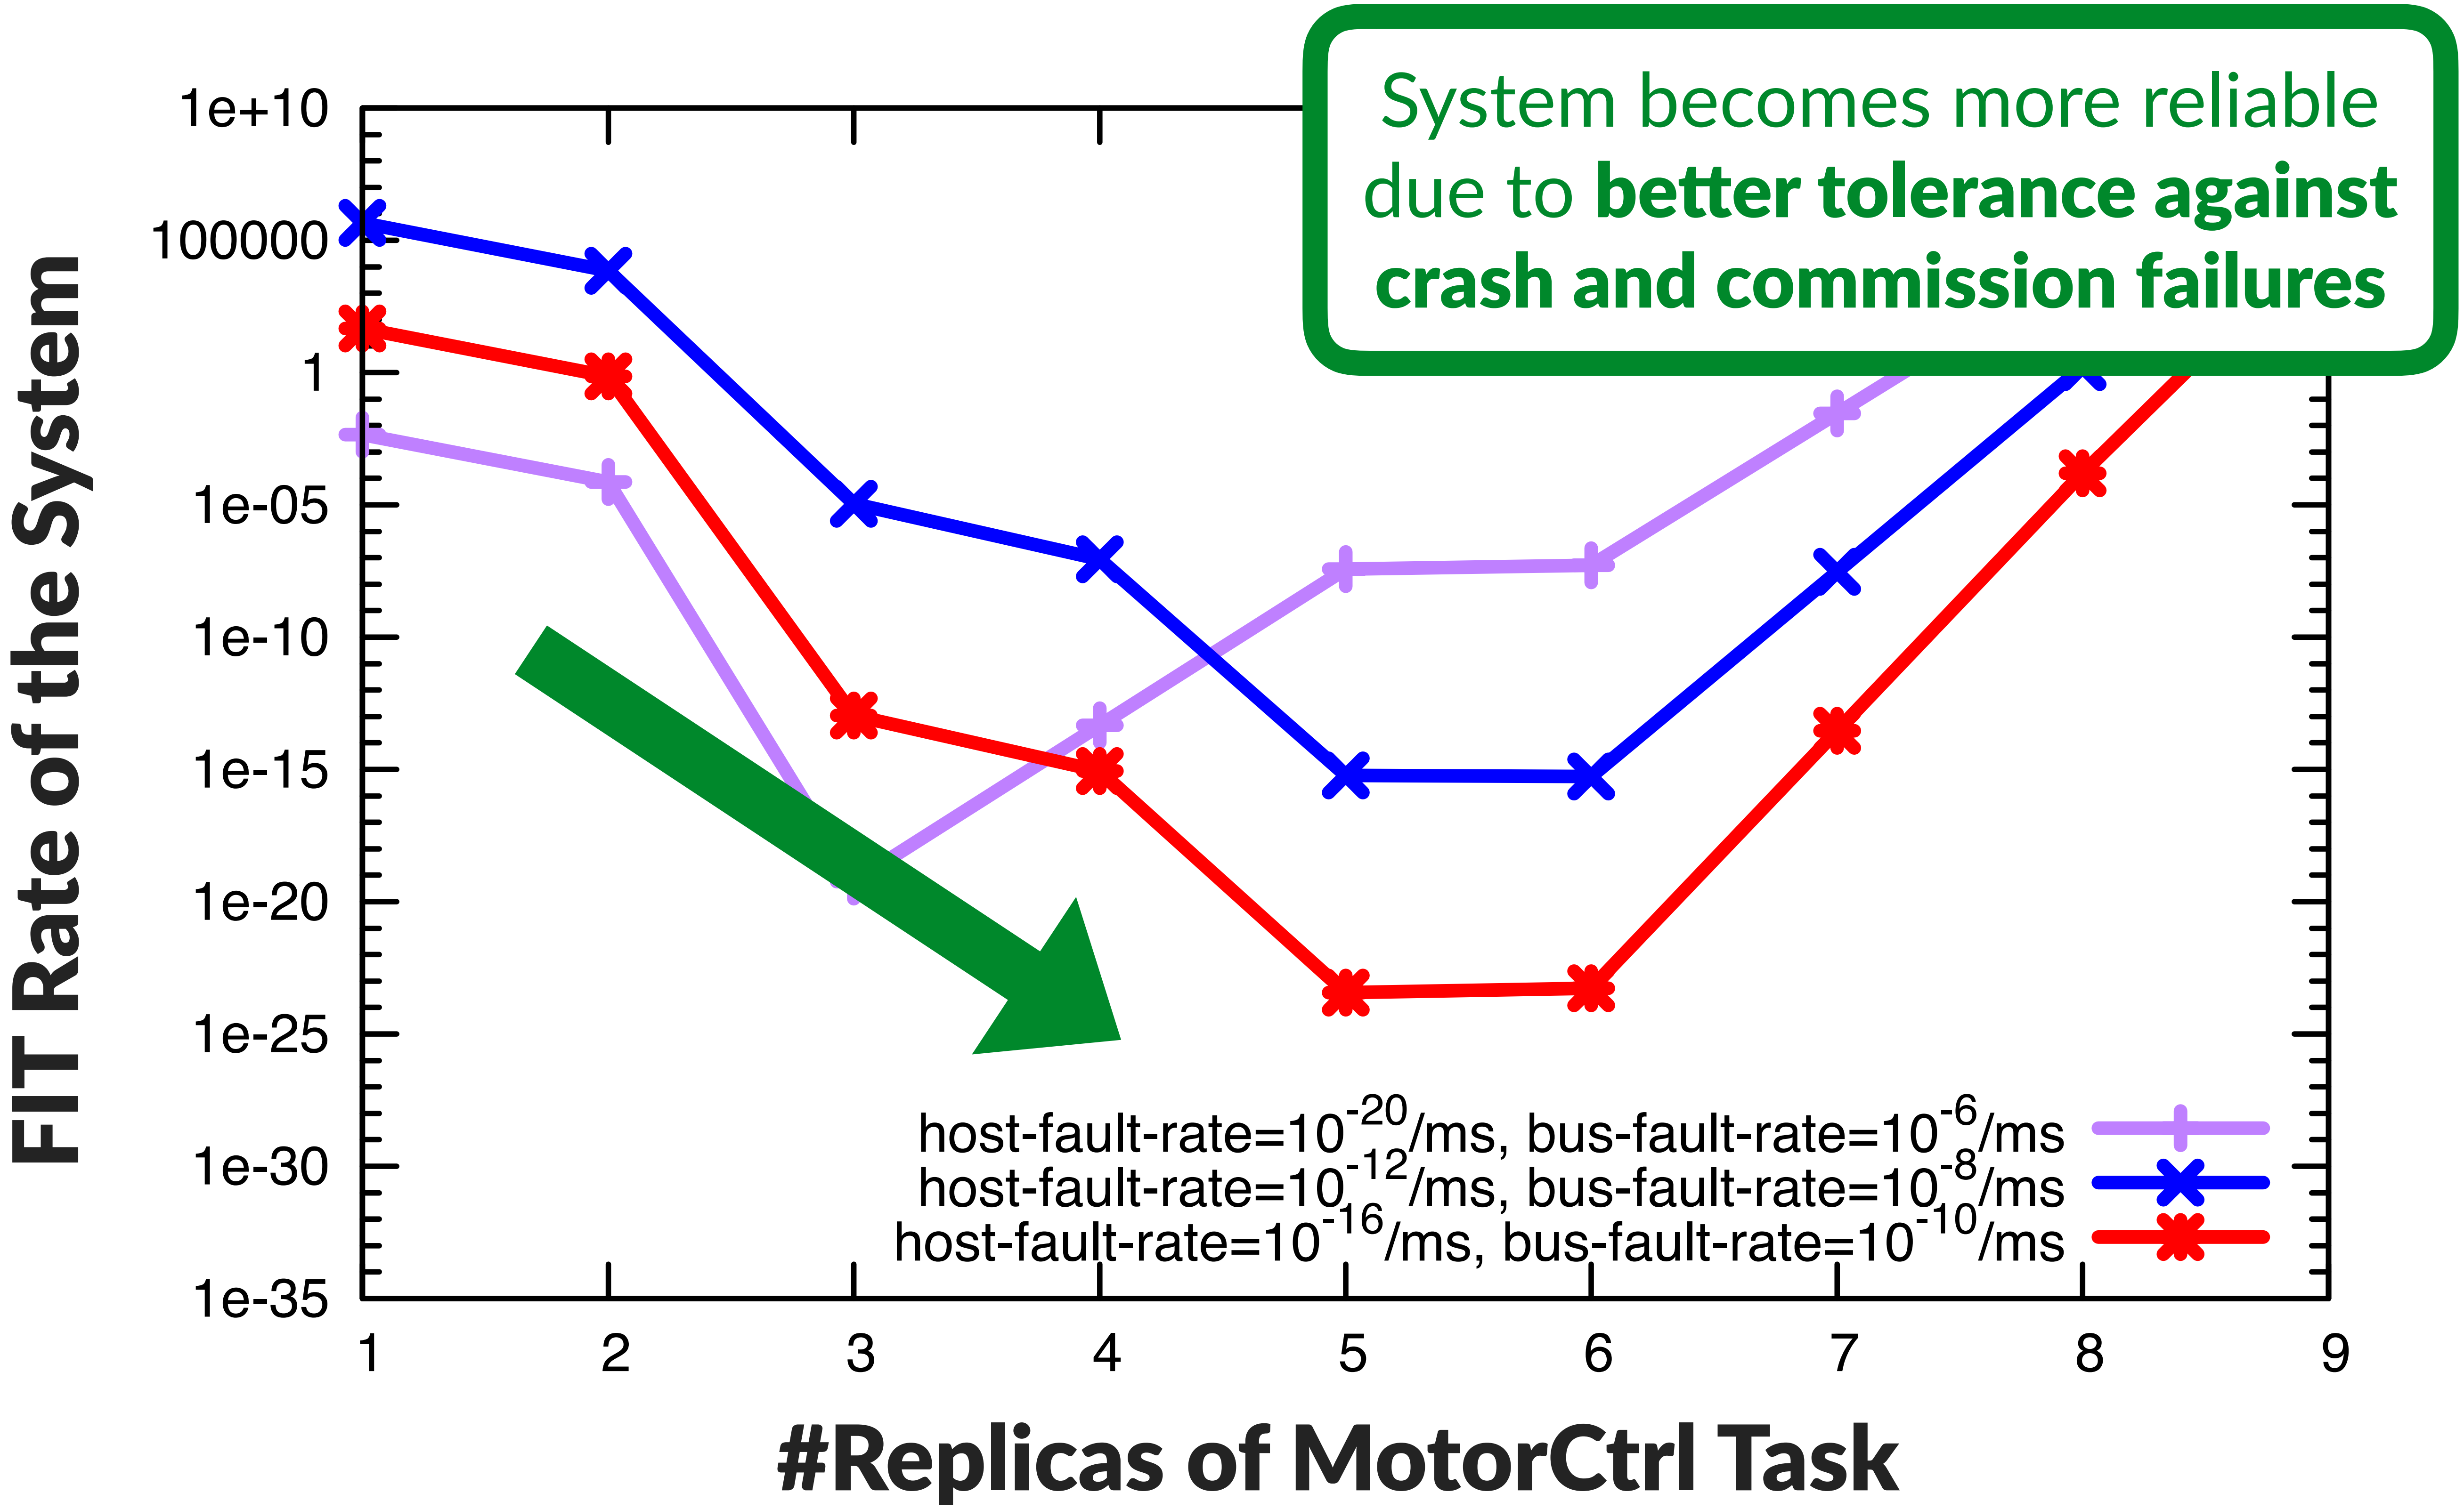

FIT Rate Analysis of the CAN Subsystem

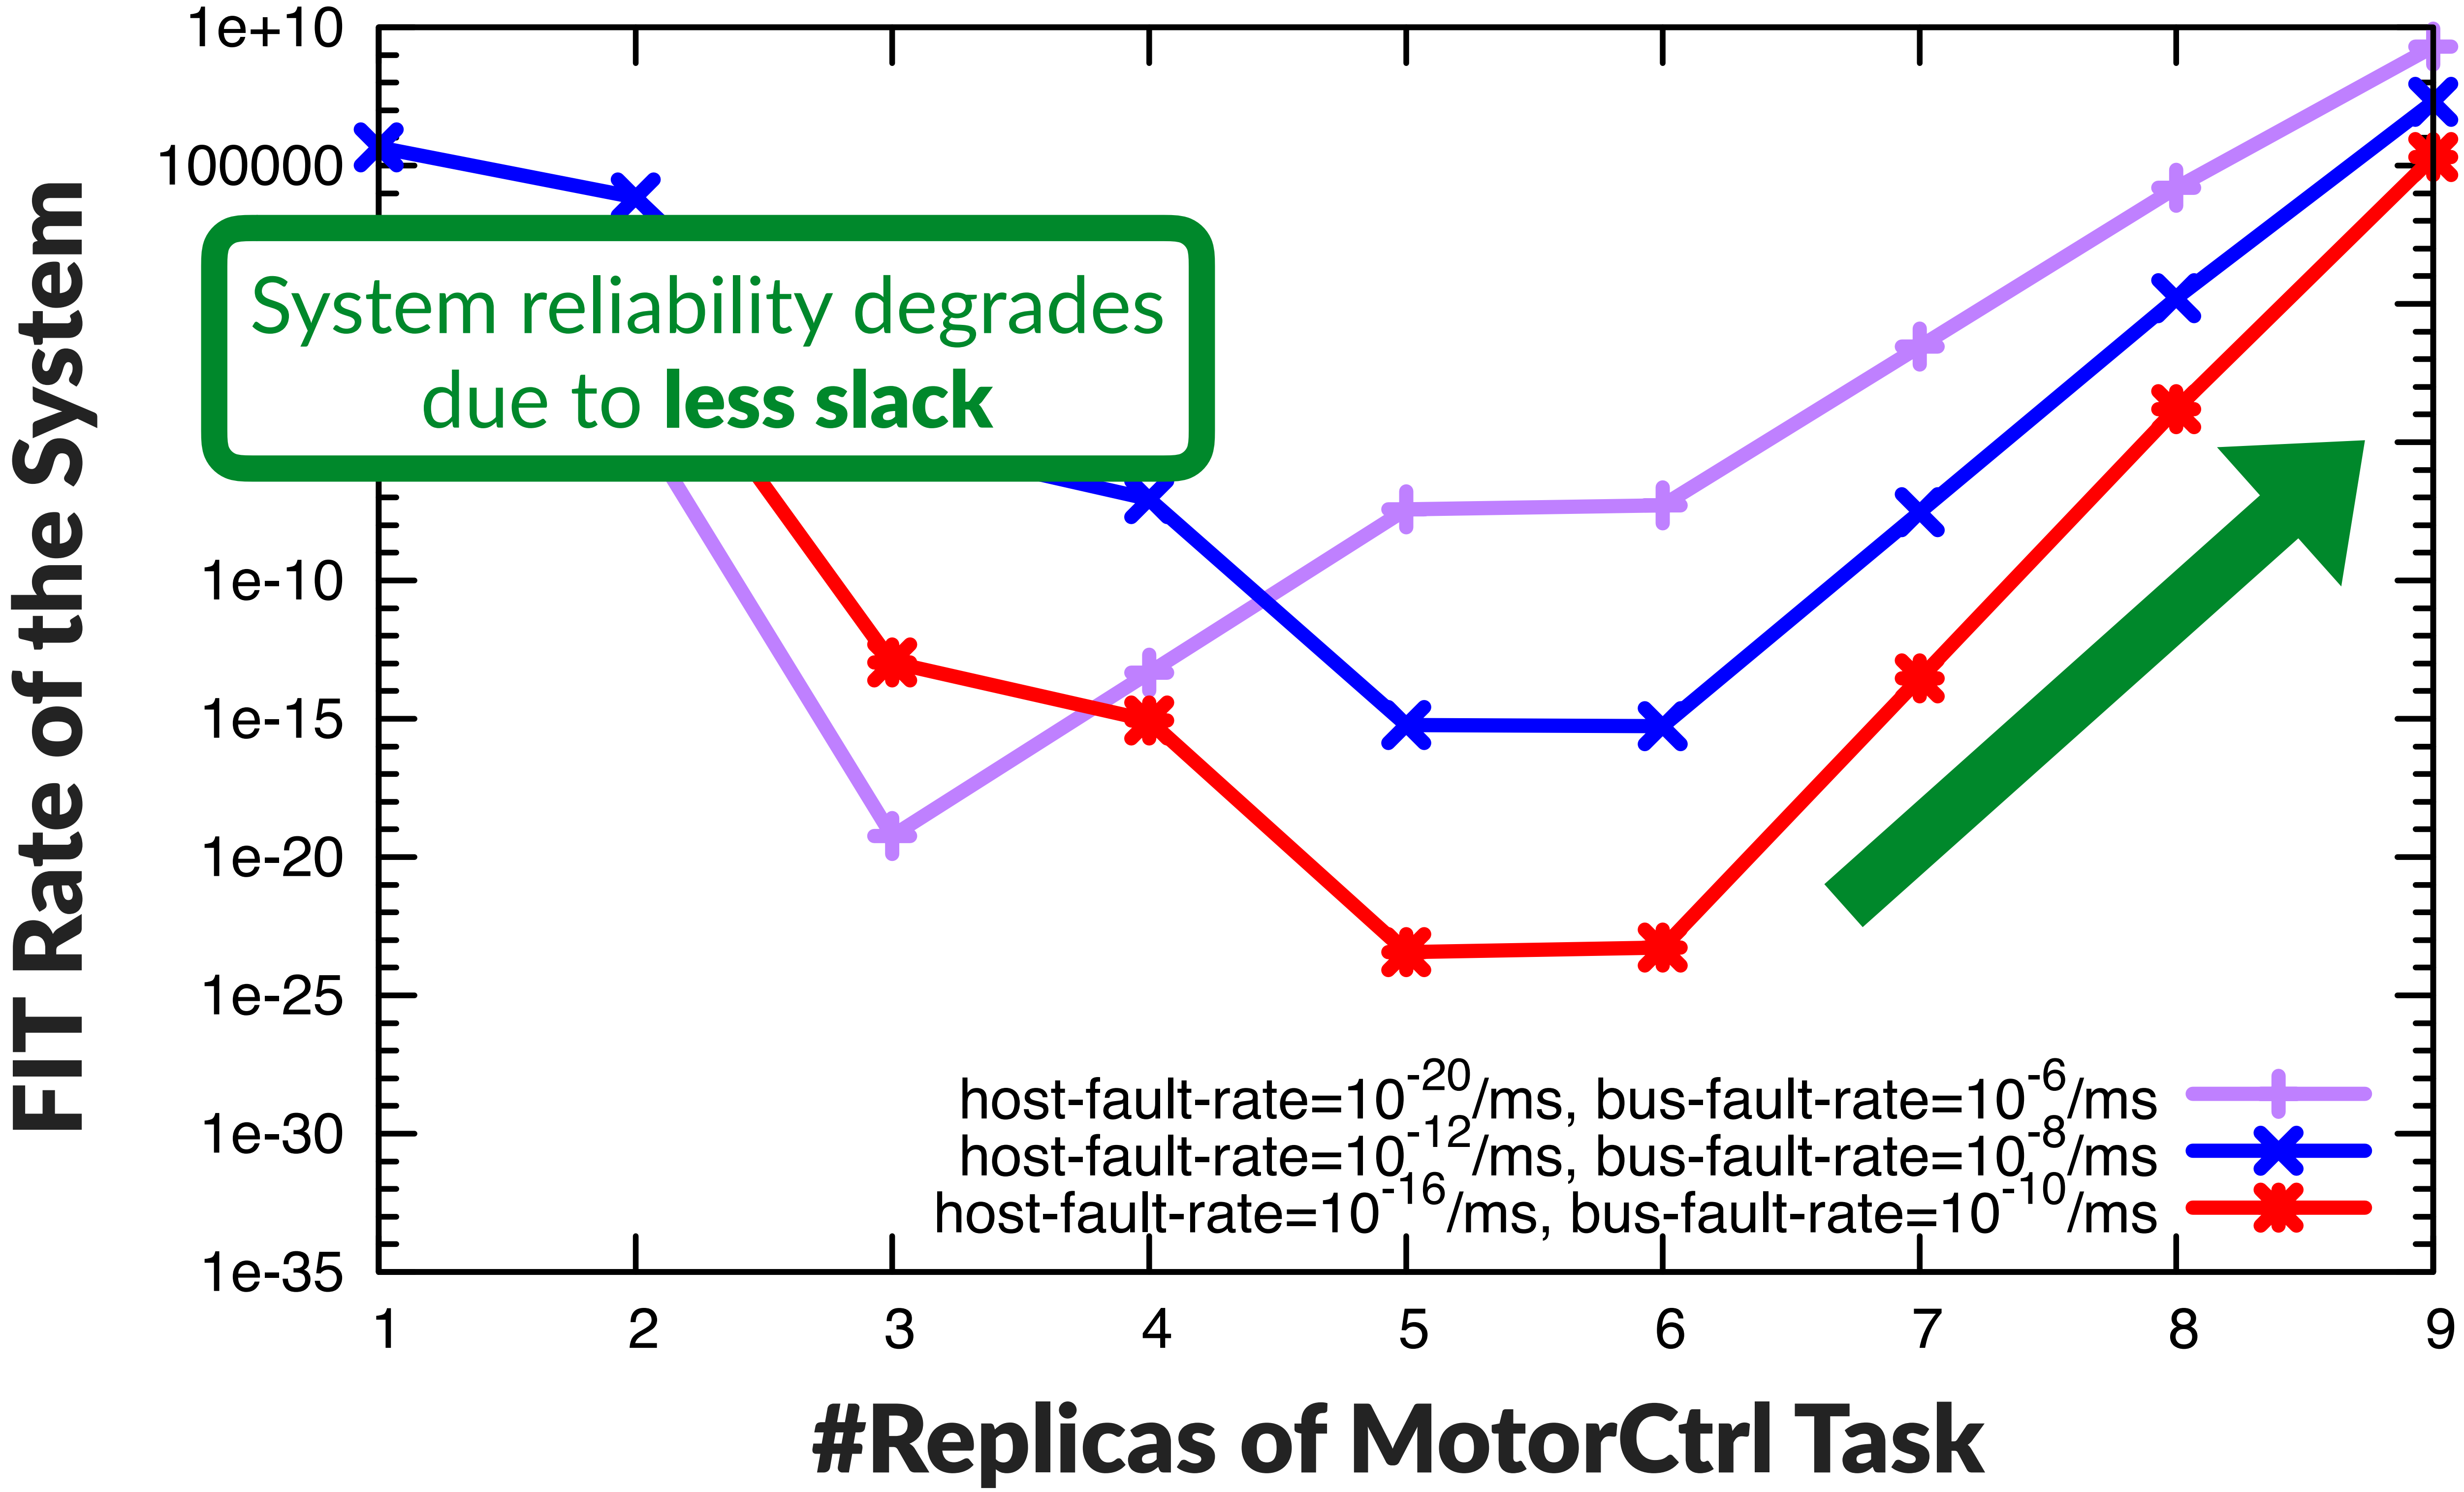

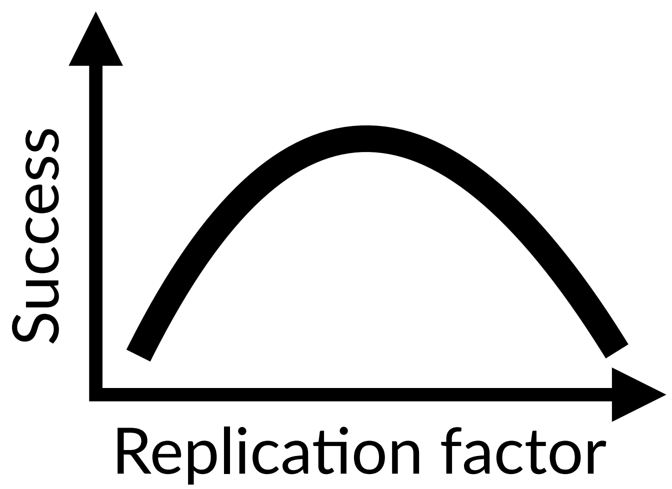
Success
Replication factor

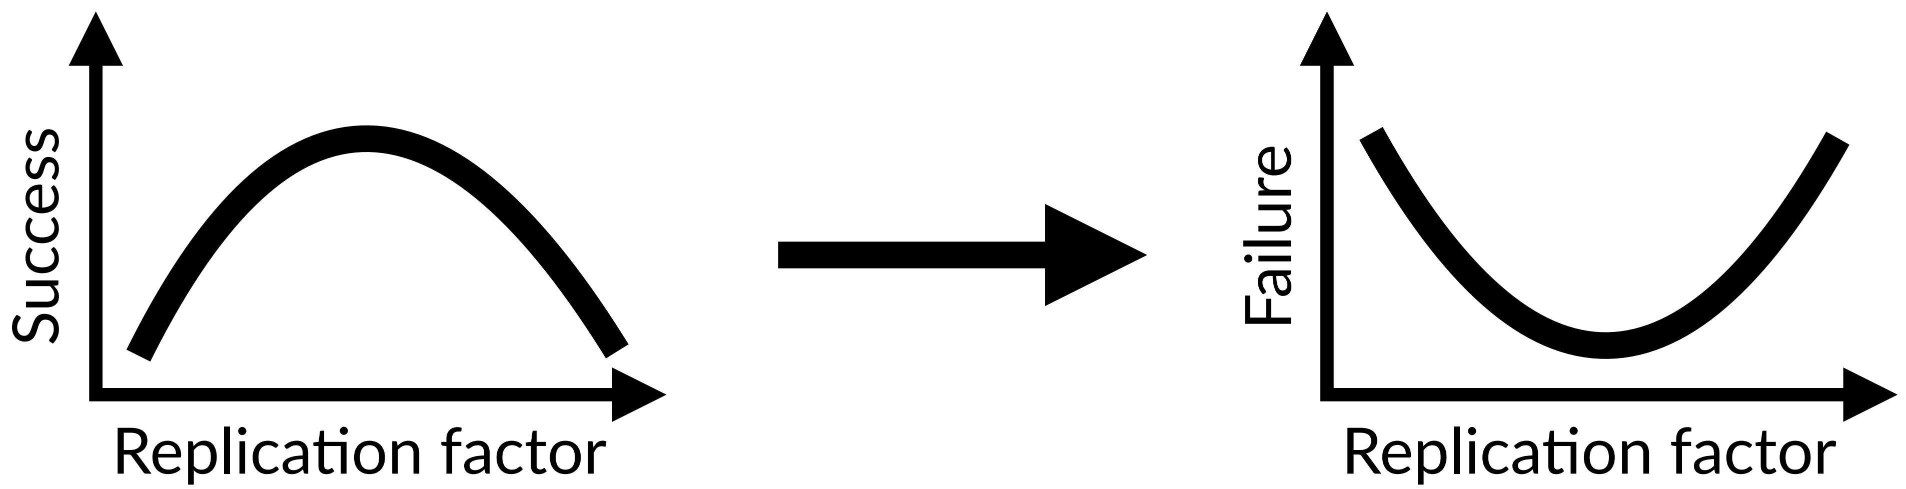

Success
Replication factor
Failure
Replication factor

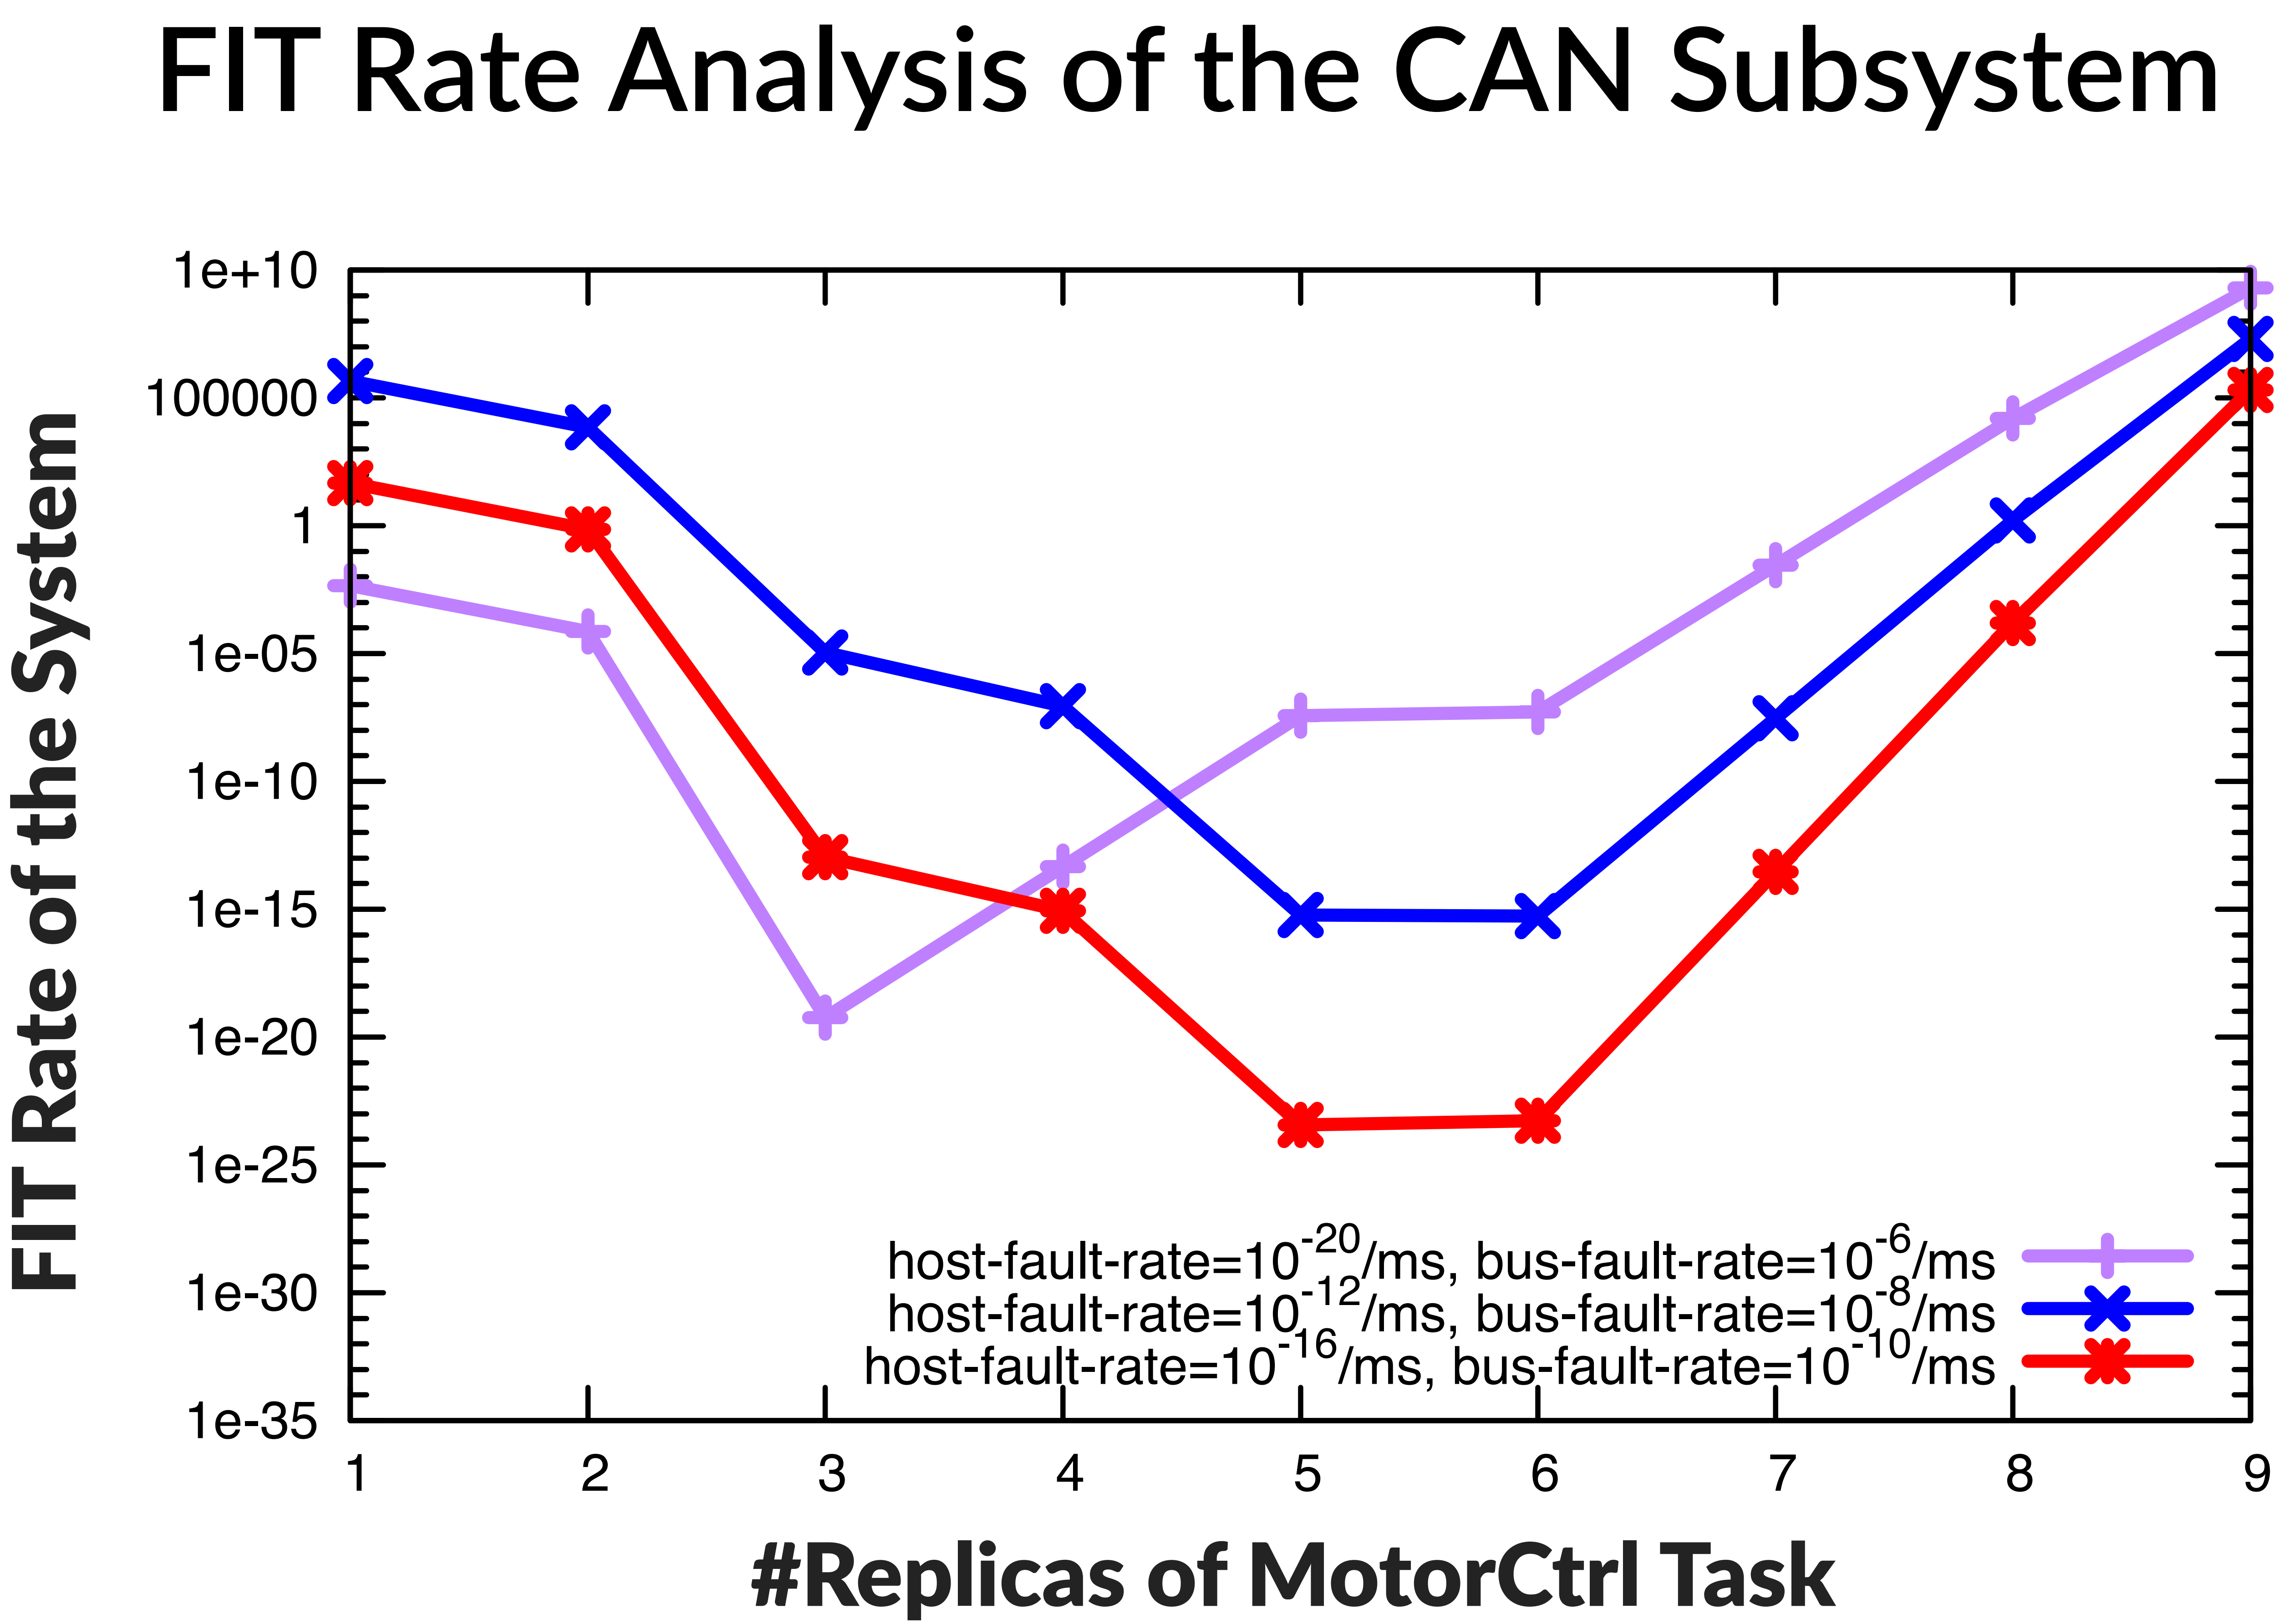

FIT Rate Analysis of the CAN Subsystem
FIT Rate of the System
#Replicas of MotorCtrl Task
1e+10
100000
1
1e-05
1e-10
1e-15
1e-20
1e-25
1e-30
1e-35
1
2
3
4
5
6
7
8
9
host-fault-rate=10-20/ms, bus-fault-rate=10-6/ms
host-fault-rate=10-12/ms, bus-fault-rate=10-8/ms
host-fault-rate=10-16/ms, bus-fault-rate=10-10/ms

FIT Rate Analysis of the CAN Subsystem

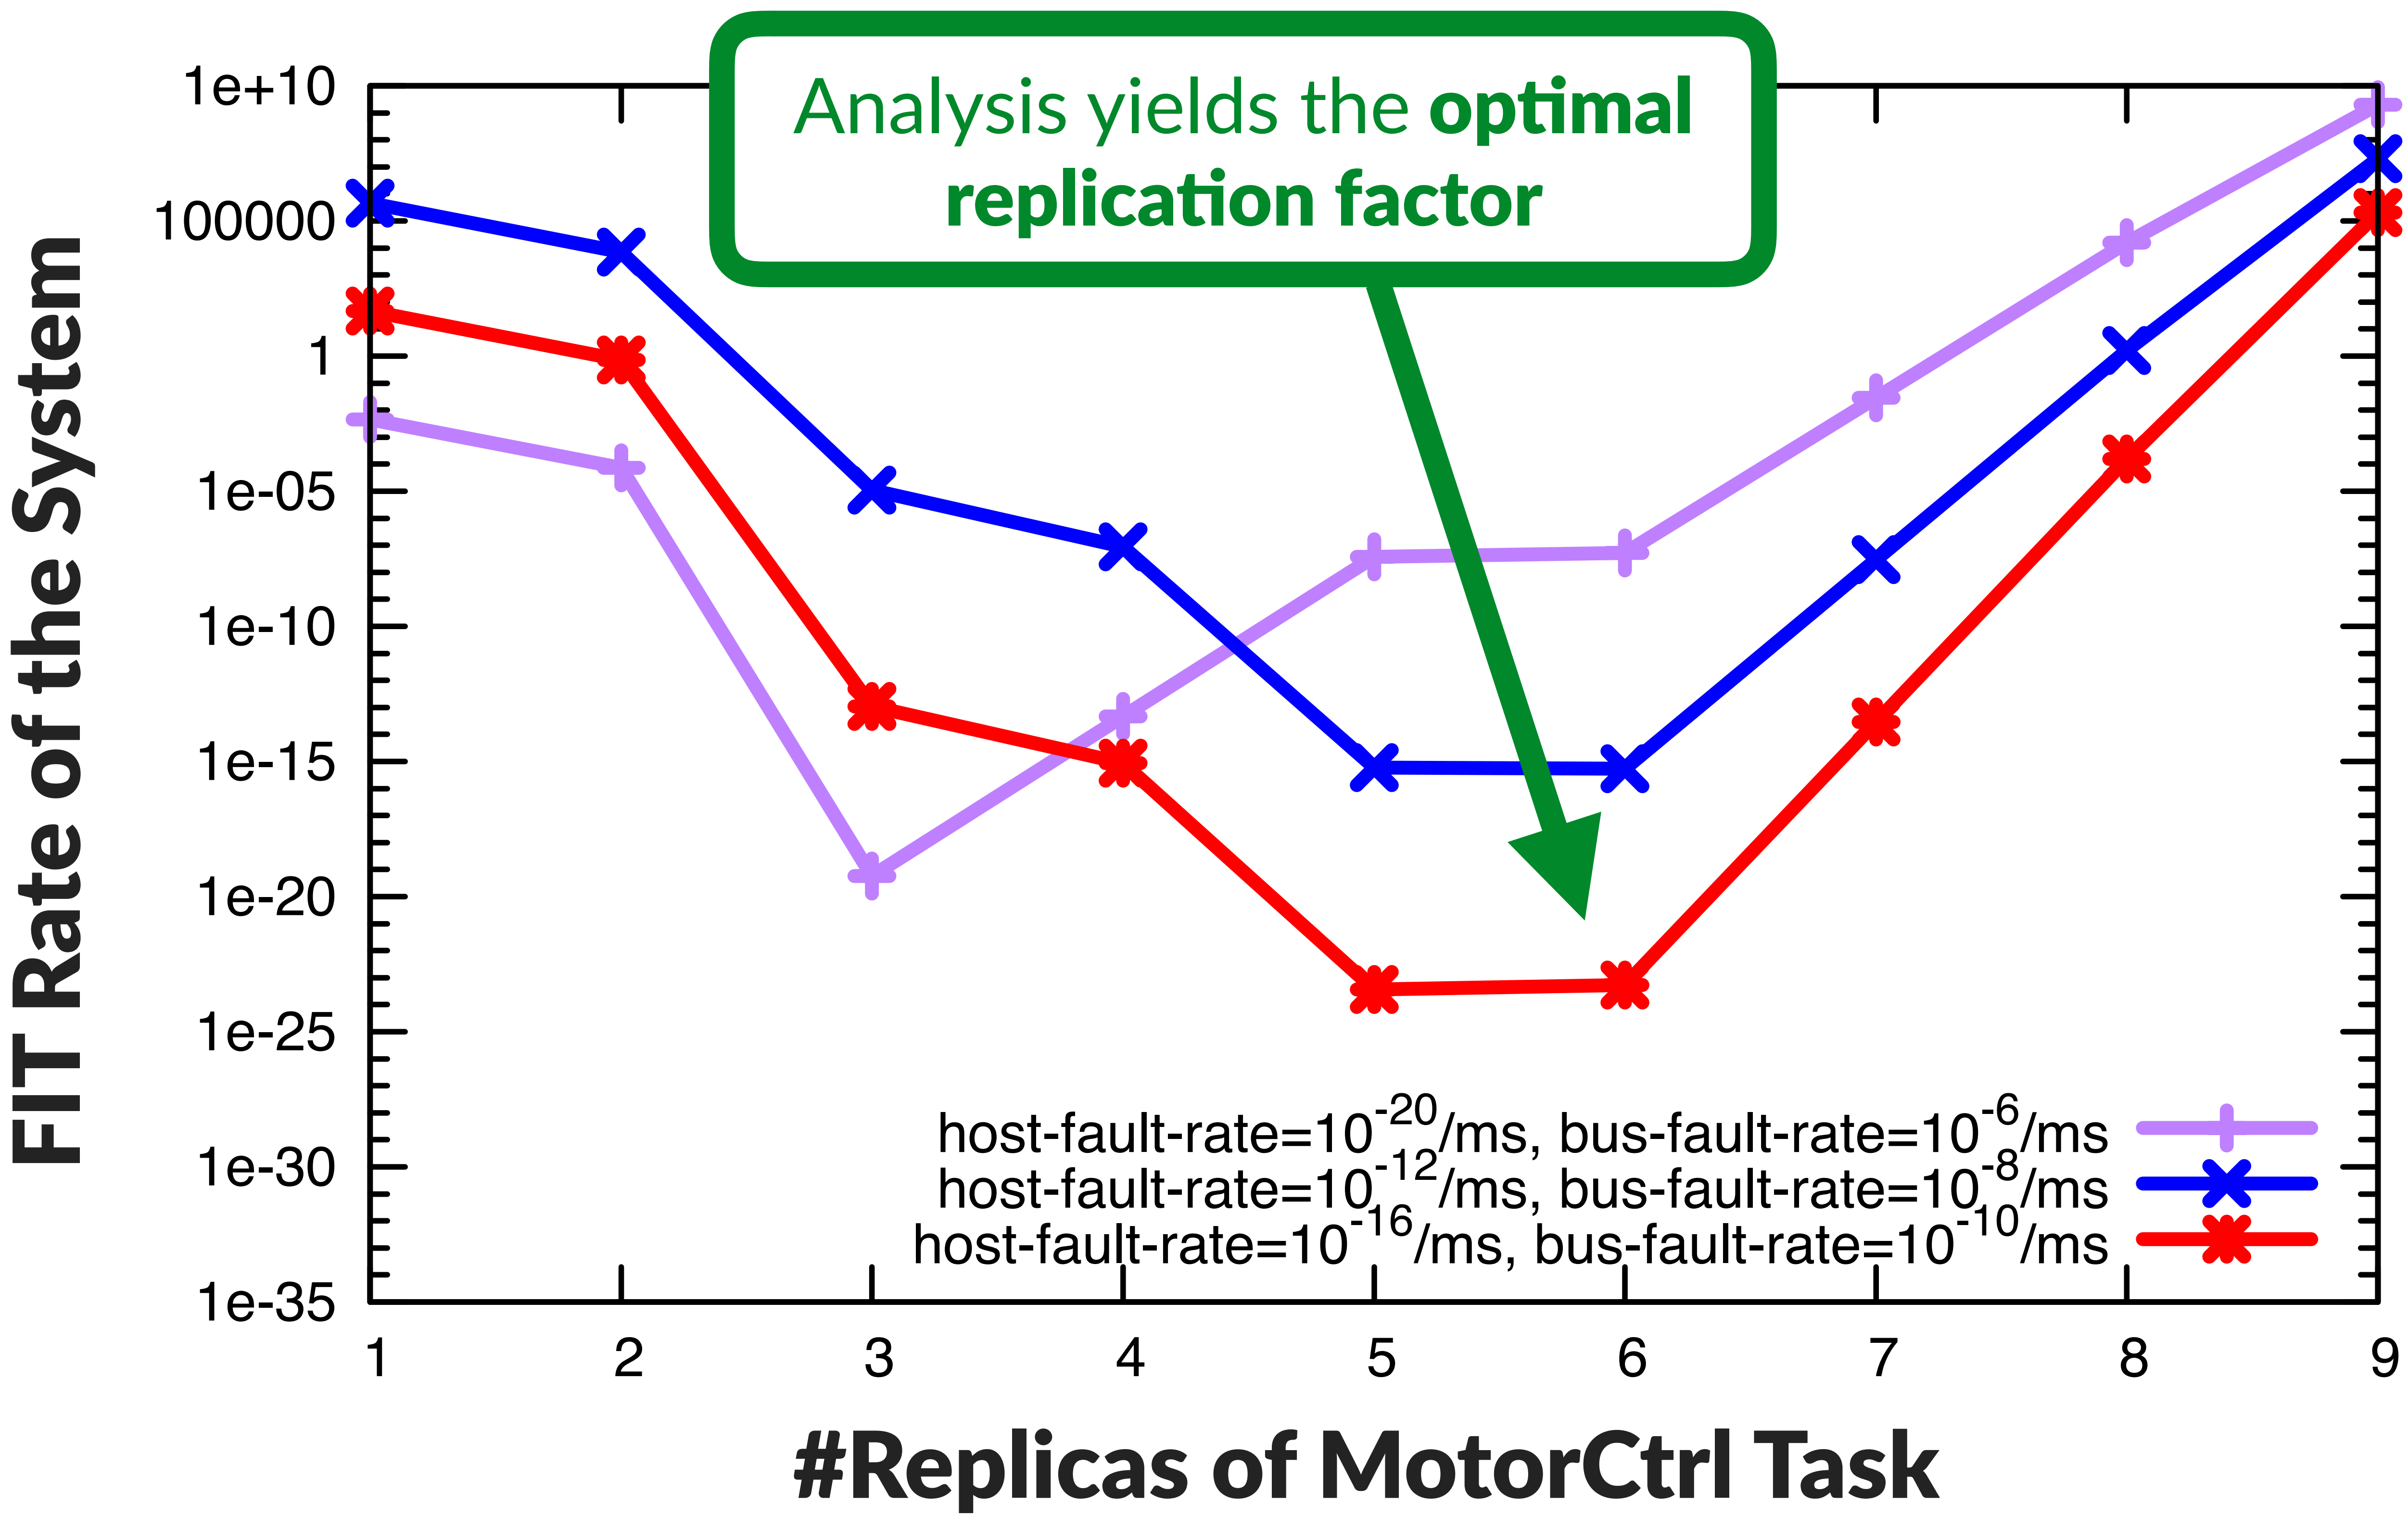

FIT Rate Analysis of the CAN Subsystem

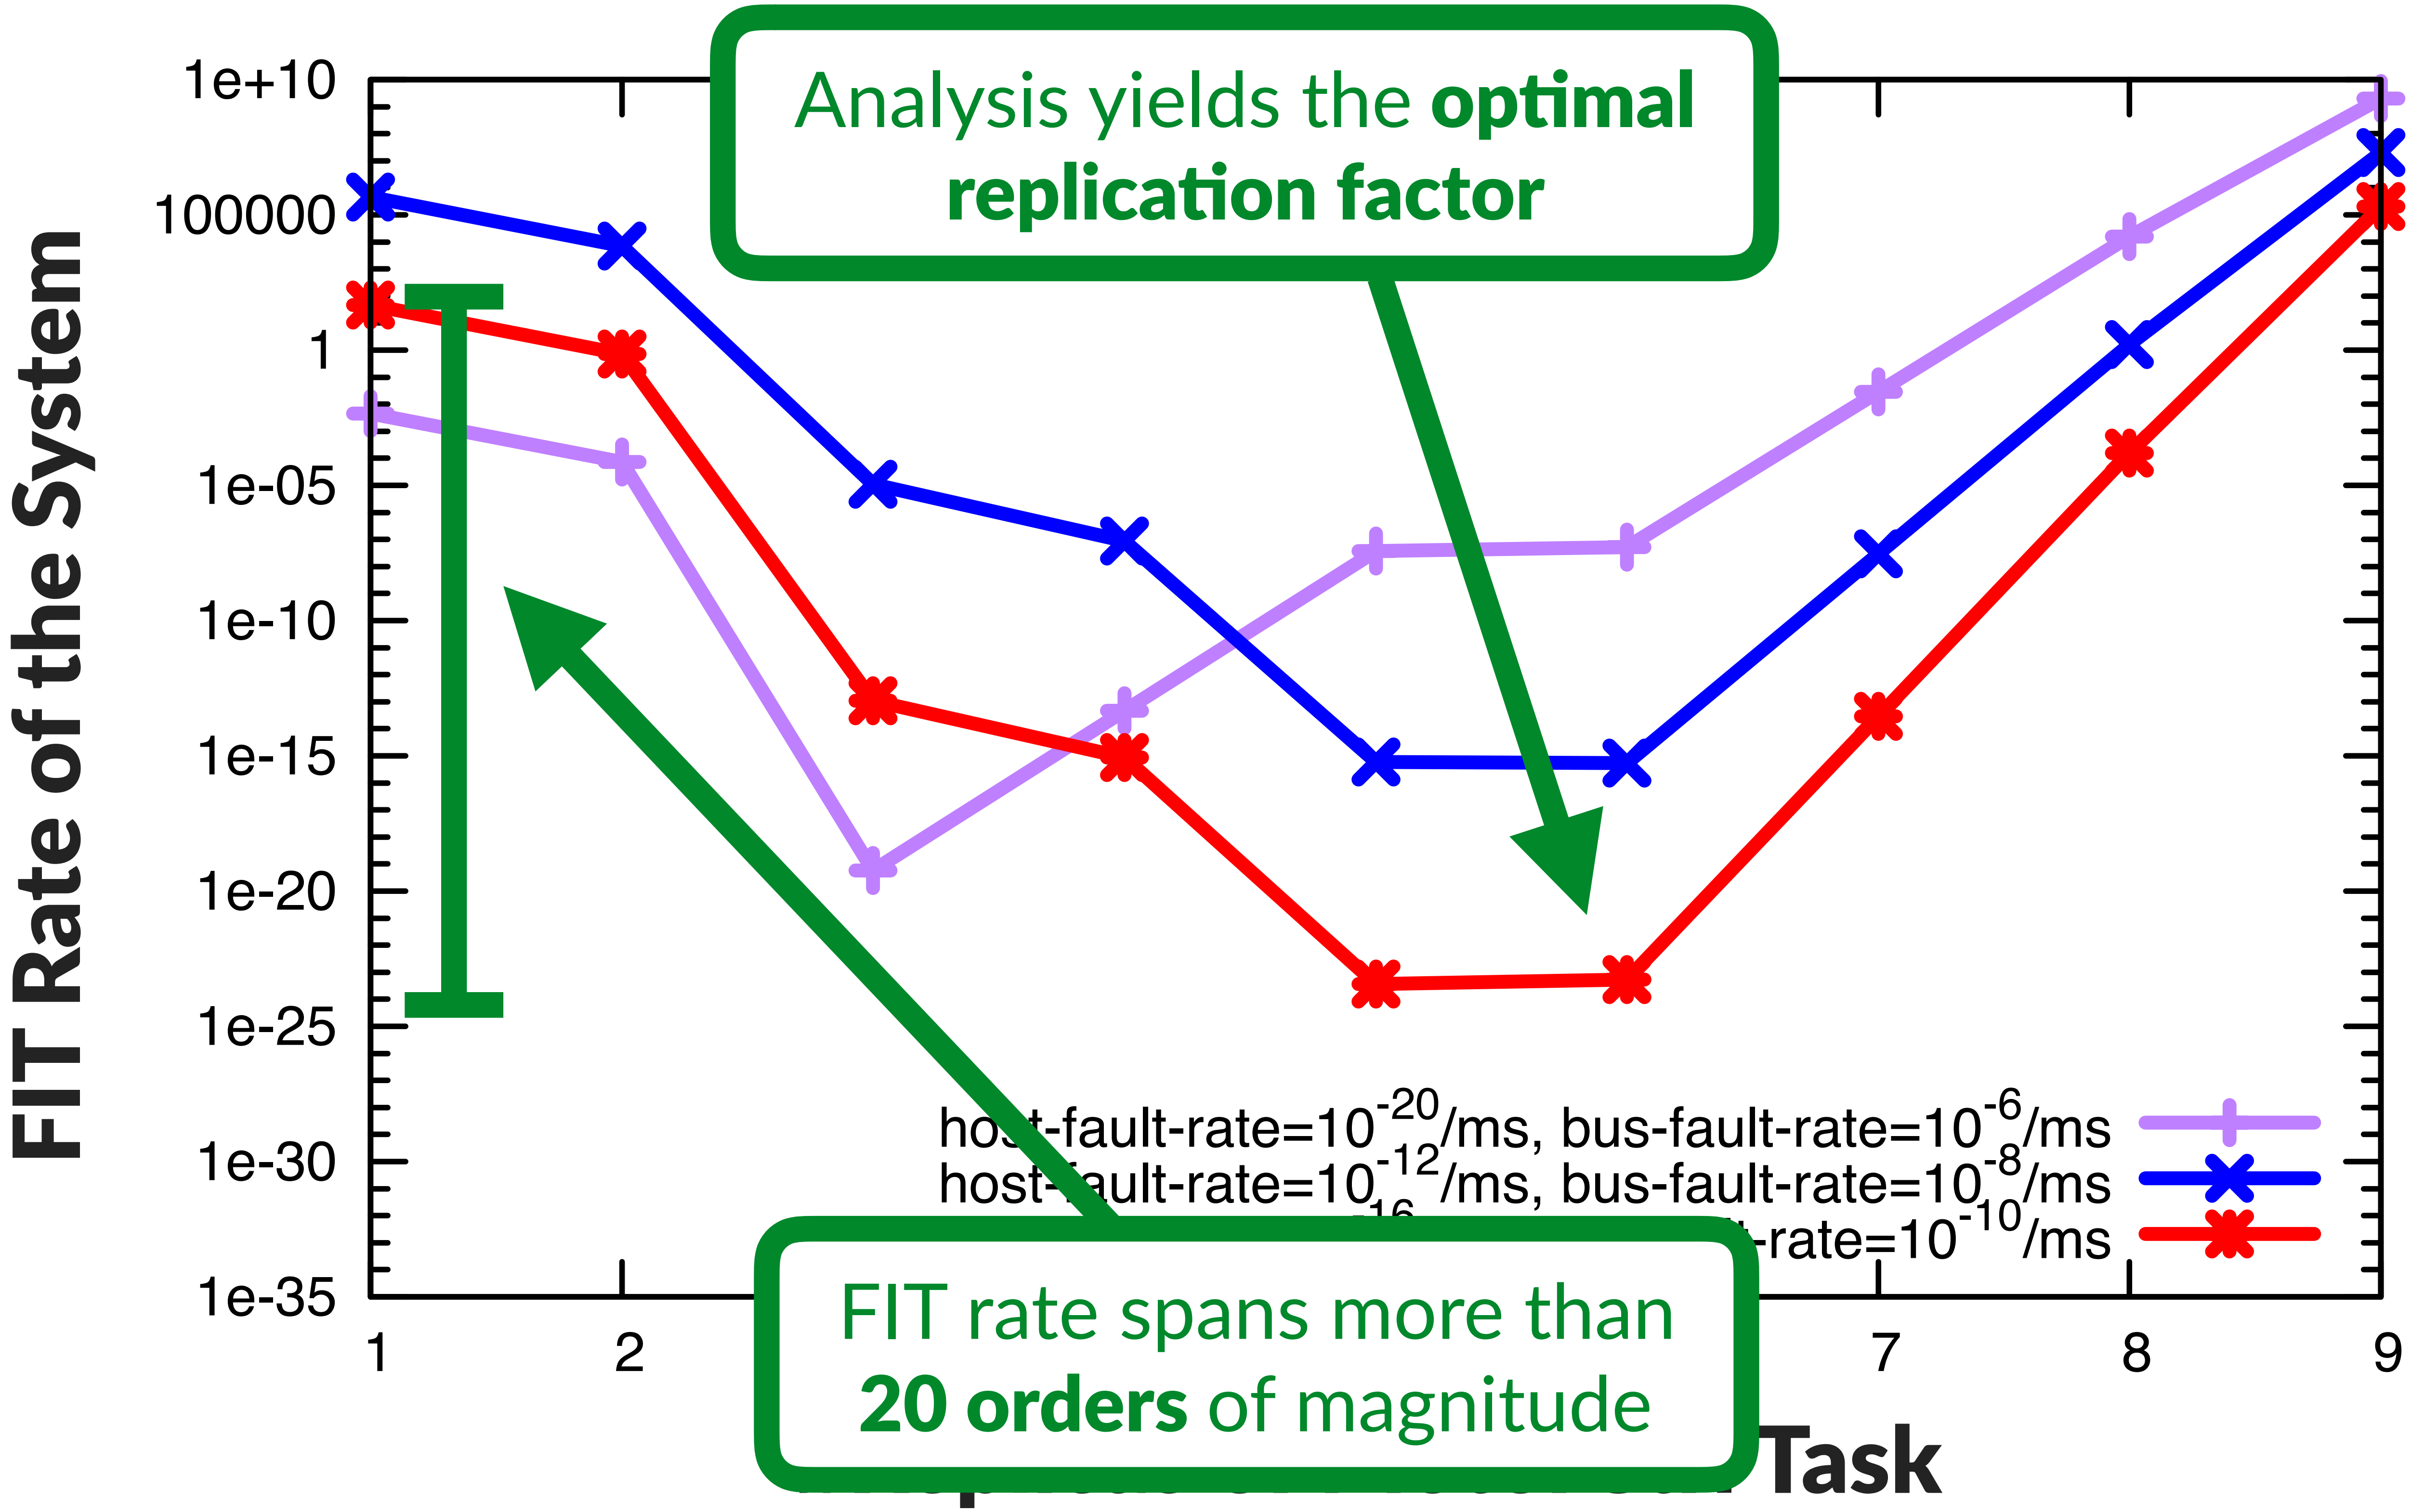

Summary

Summary

Find the best replication strategy for CAN-based systems

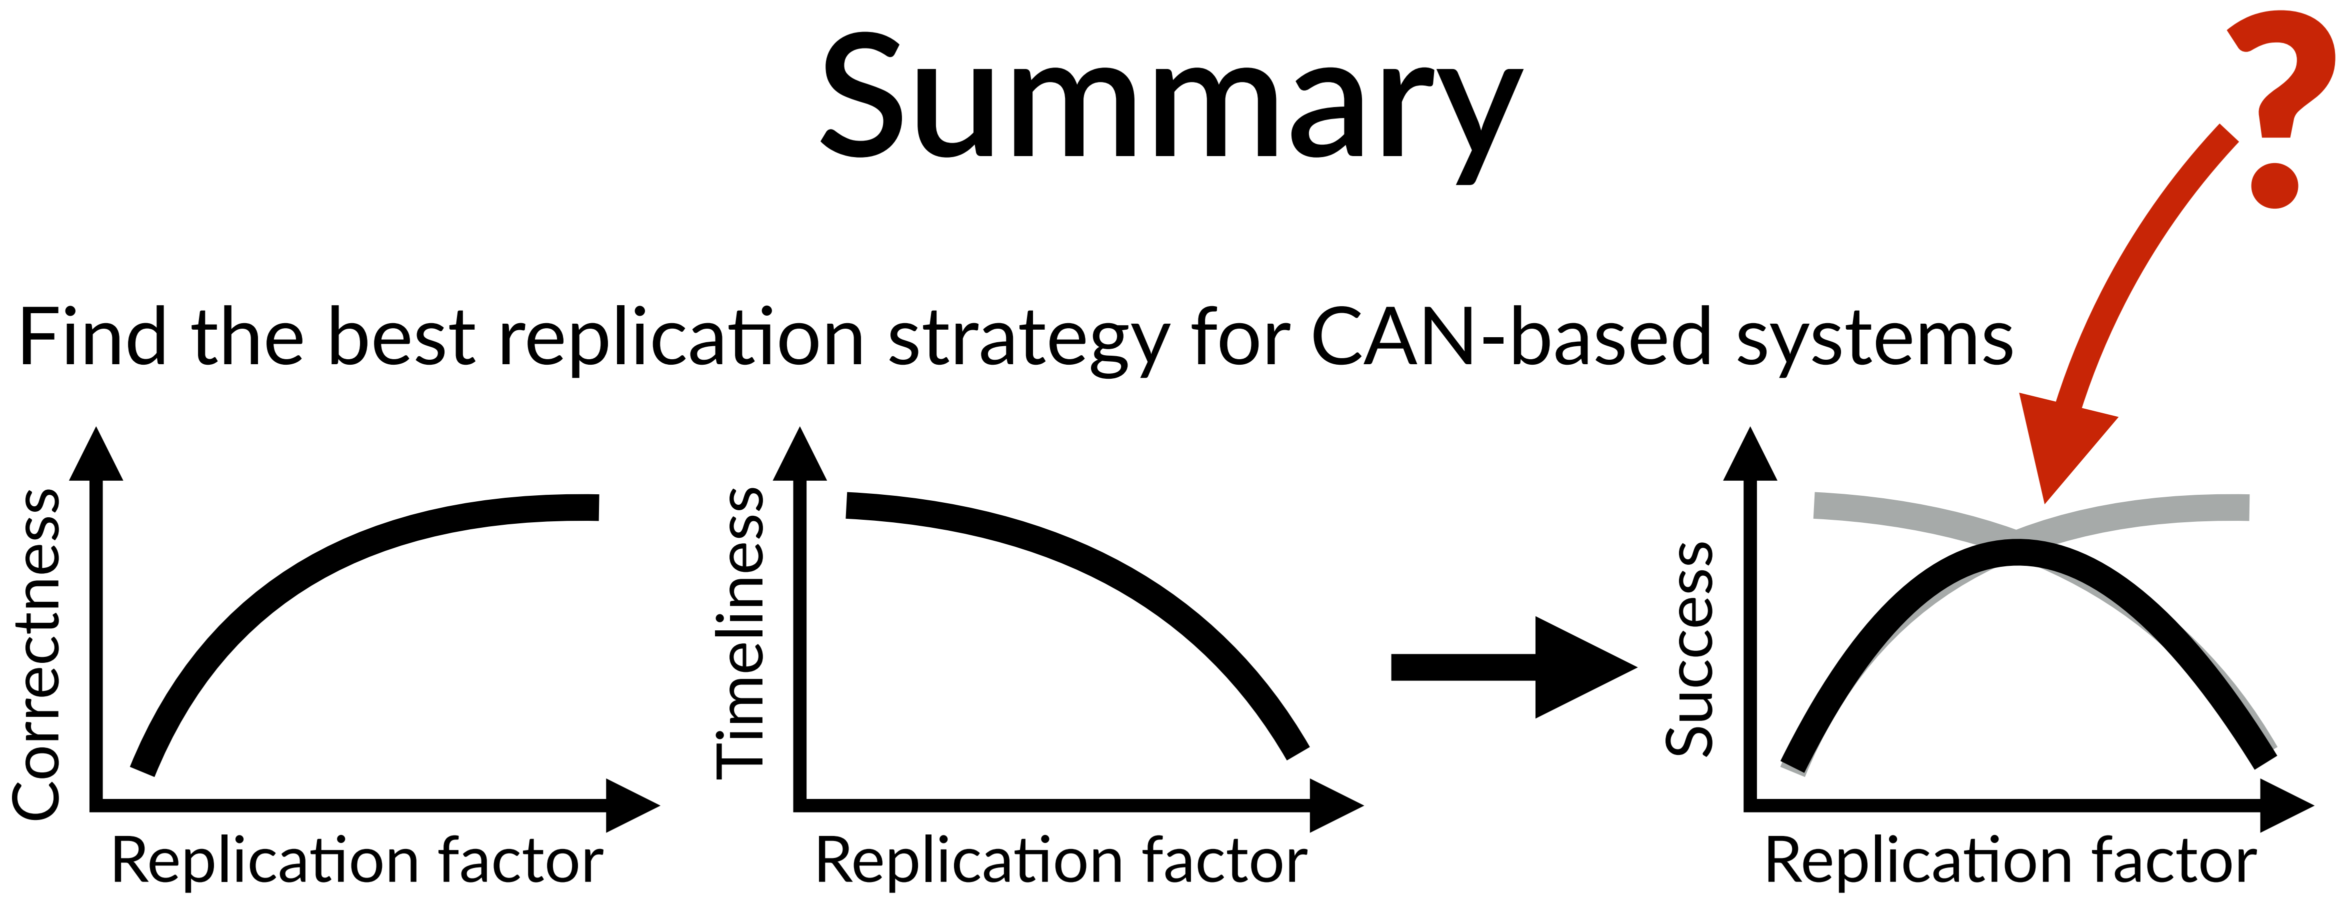

Summary

Find the best replication strategy for CAN-based systems

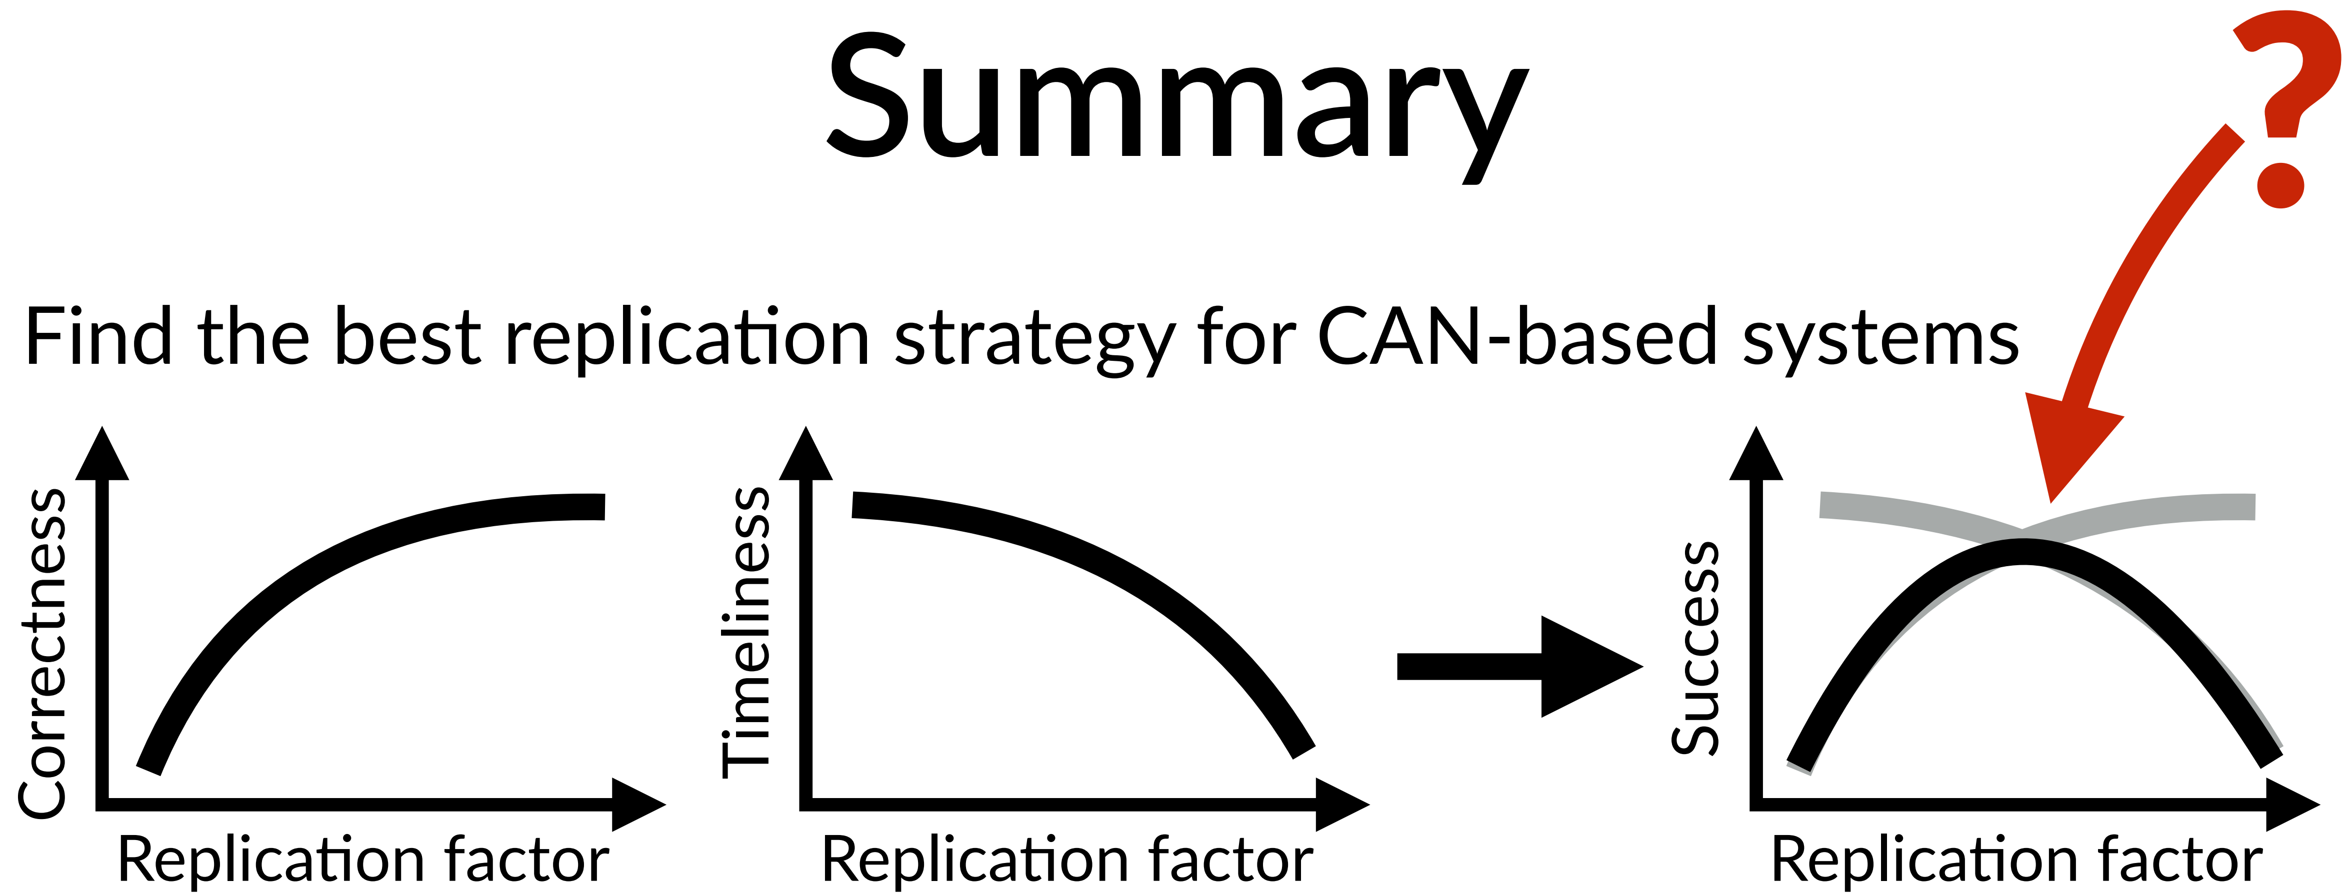

Compare reliability of the CAN-bus subsystem in the context of the larger safety-critical system

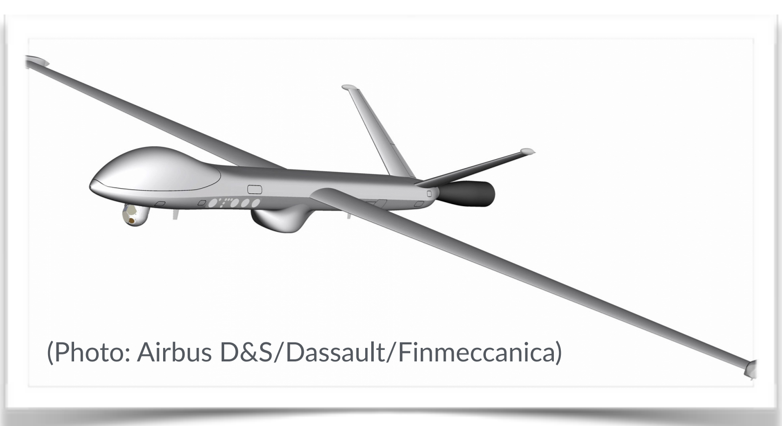

UAV

=

Power

+

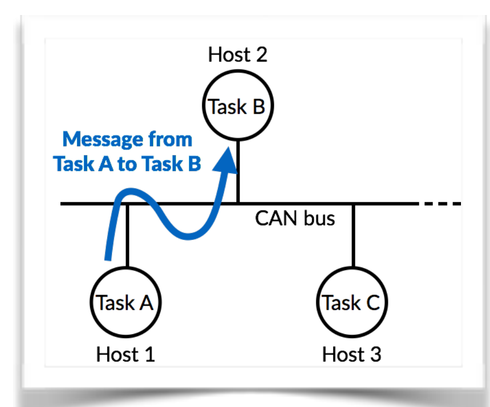

CAN

+ •••

Future Work

- More complex system models
  - ➡ CAN-based systems bridged together
  - ➡ Sporadic DAG models

Future Work

- More complex system models
  - ➡ CAN-based systems bridged together
  - ➡ Sporadic DAG models
- Study of other technologies, e.g., Real Time Ethernet

Future Work

- More complex system models
  - ➡ CAN-based systems bridged together
  - ➡ Sporadic DAG models
- Study of other technologies, e.g., Real Time Ethernet
- **http://www.mpi-sws.org/~bbb/projects/schedcat**