



Flooding Oscillations

Bruno Rijsman, 16-Oct-2018, v1.2

Current Flooding Scope Rules

	South	North	East-West
Node S-TIE	Flood if level of originator is equal to this node.	Flood if level of originator is higher than this node.	Flood only if this node is not top-of-fabric.
Non-Node S-TIE	Flood self-originated only.	Flood only if neighbor is originator of TIE.	Flood only if self-originated and this node is not top-of-fabric.
All N-TIEs	Never flood.	Flood always.	Flood only if this node is top-of-fabric.
TIDE	Include at least TIEs in flooding scope		
TIRE	Include all N-TIEs and all peer's self-originated TIEs and all node S-TIEs	Flood only if neighbor is originator of TIE.	If this node is top-of-fabric then apply north scope rules, otherwise south scope rules

Introducing a Bit of Notation

```
in_flood_scope(from_node=X, to_node=Y, tie=T)
```

- Returns true if node X **MUST** flood TIE T to node Y
- Returns false if node X **MUST NOT** flood TIE T to node Y
- If false, node Y **MUST NOT** accept and reflood TIE T from node X
“Belt and suspenders”: Y evaluates function “from perspective of X”
- “Node Y is in the flooding scope of node X for TIE T”
- We are only talking about TIEs here
(Not discussing TIREs or TIDEs yet – will be discussed later)

Information needed to evaluate flood scope

```
in_flood_scope(from_node=X, to_node=Y, tie=T)
```

From node

- System ID
- Level
- Top of Fabric


To node

- System ID
- Level
for Direction (N,S,EW)
- Note: to_node is
always a neighbor of
from_node

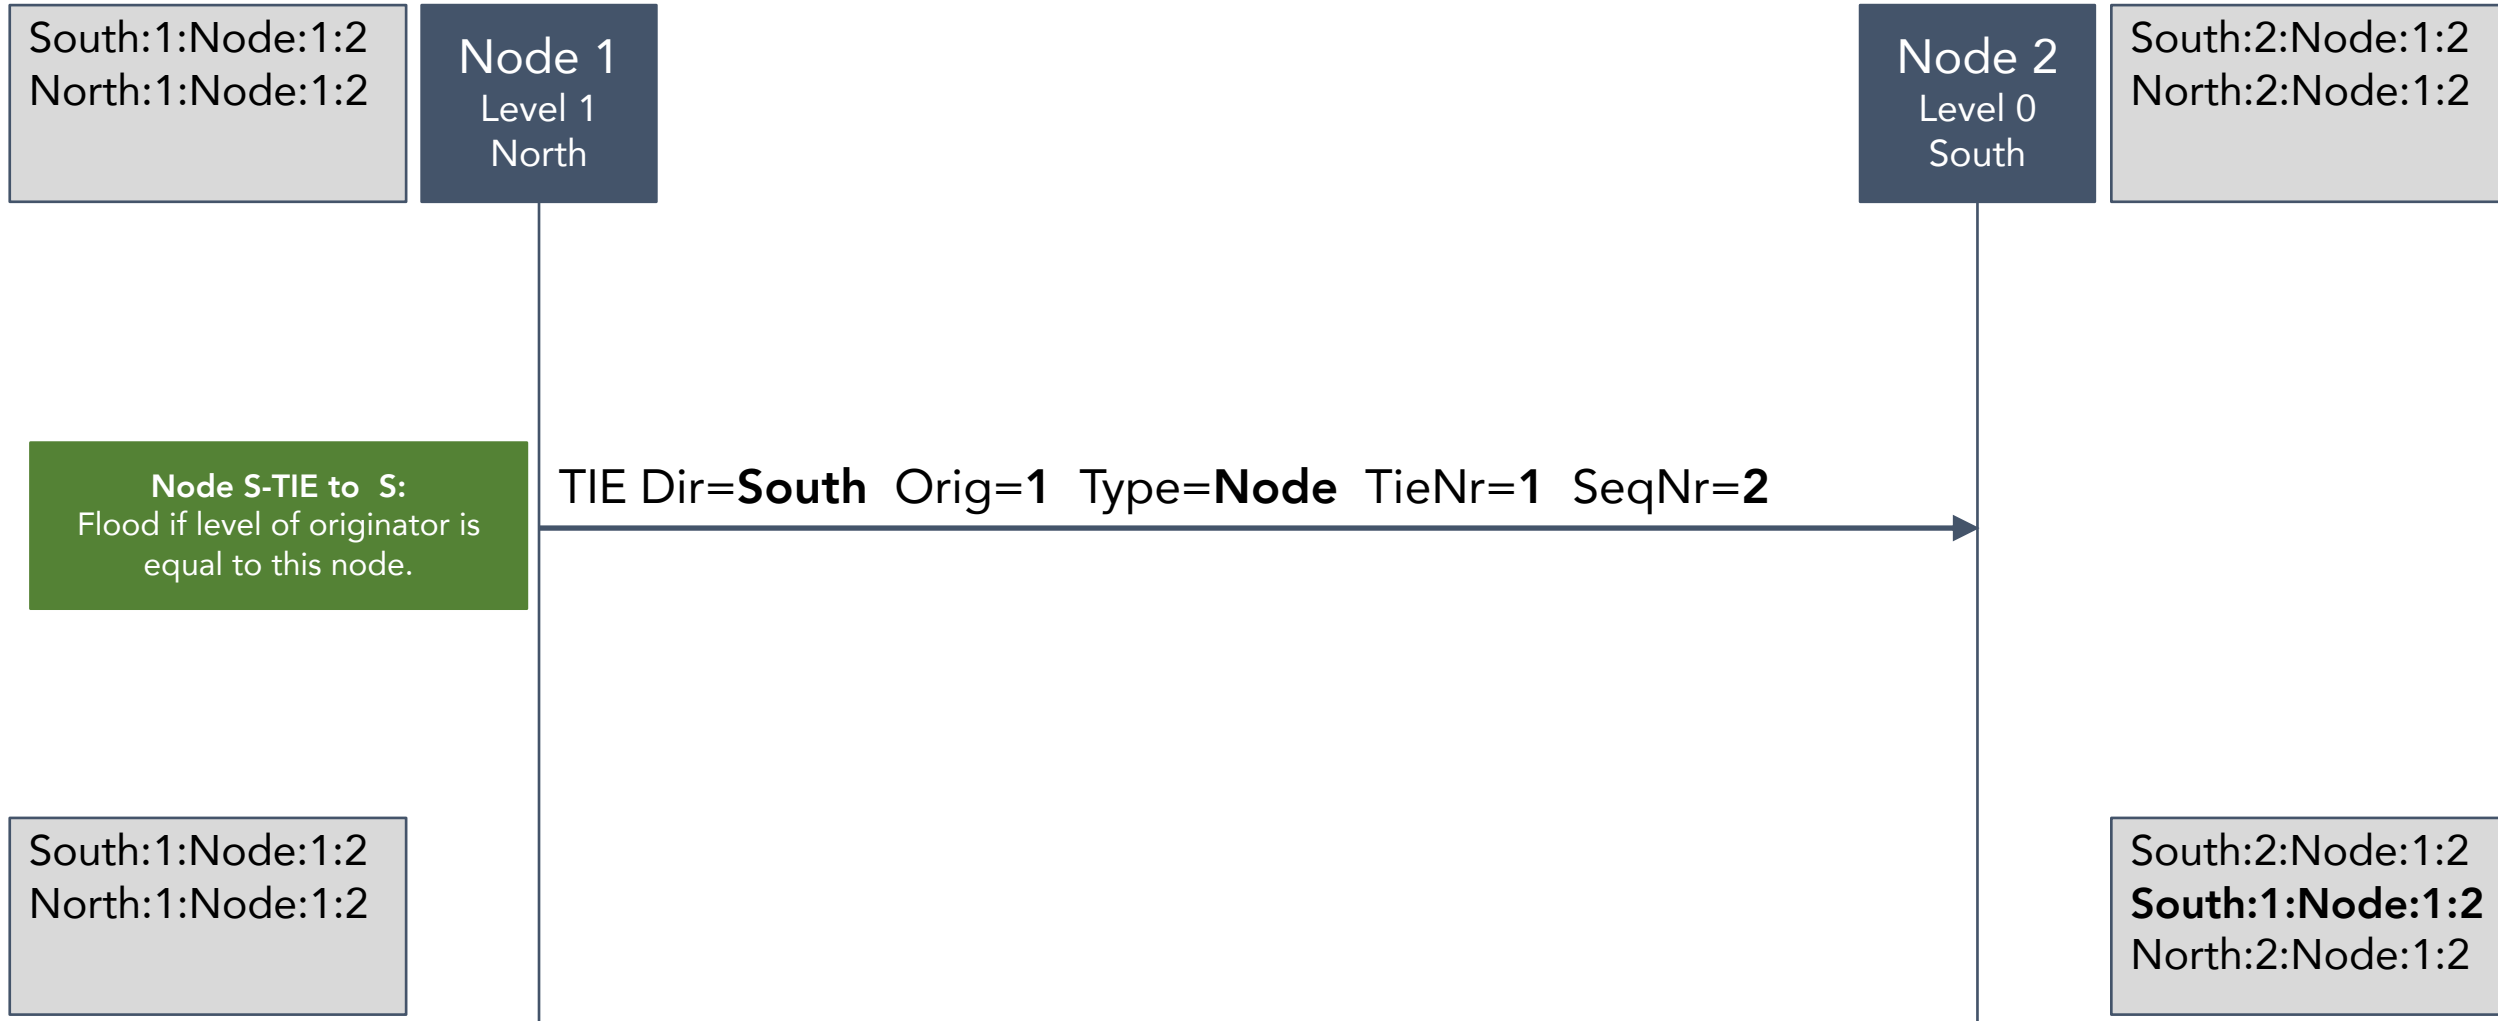
TIE

- Direction (N,S)
- Type
- Originator
for is self-originated flag
- Originator level
only for Node TIEs
only thing not in header


Oscillation #1



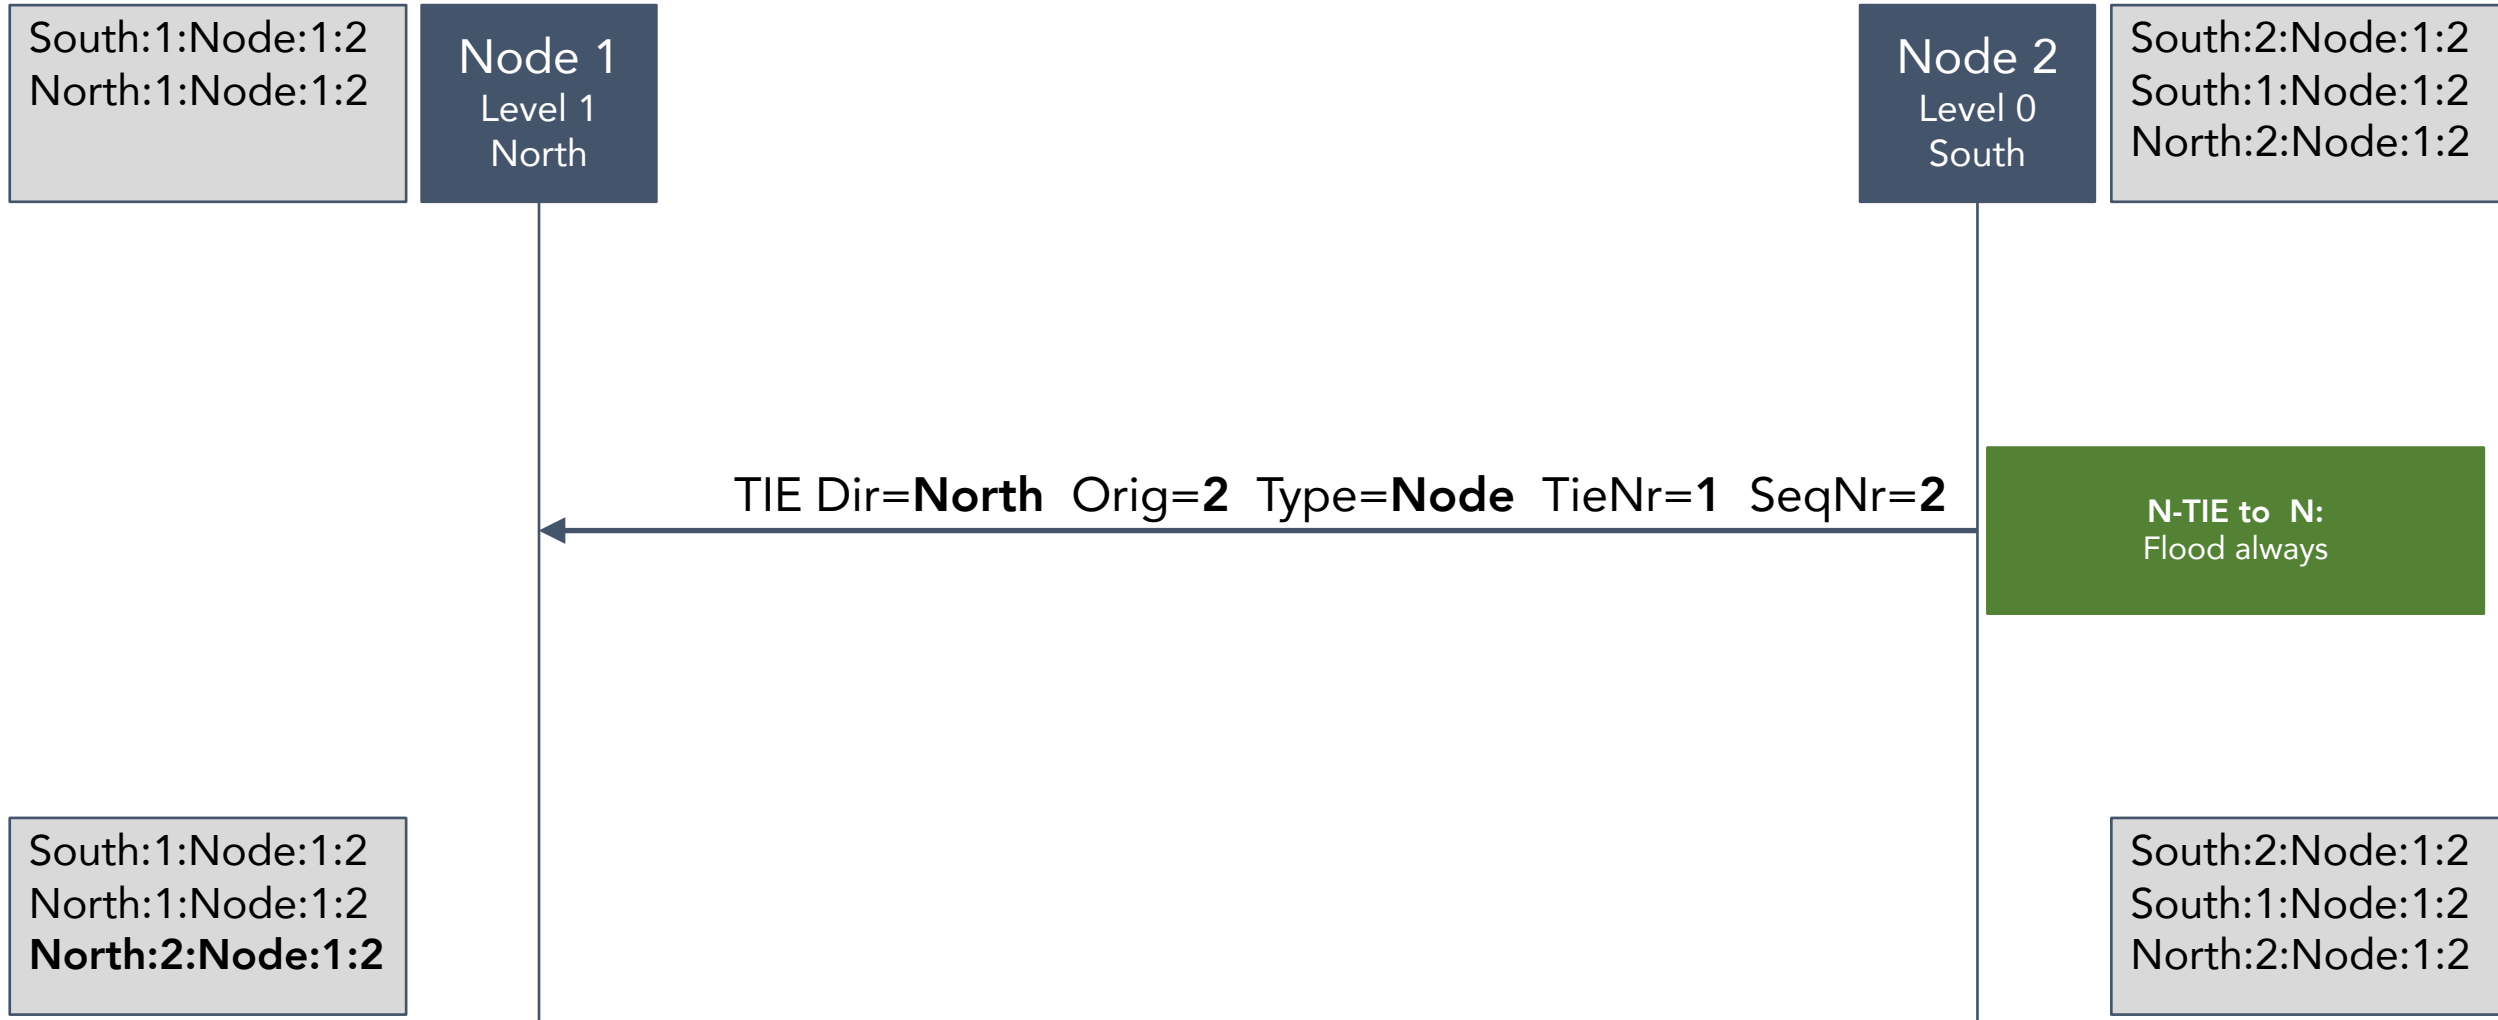
Node 1 sends its node TIE to node 2



Node 2 reflects the TIE

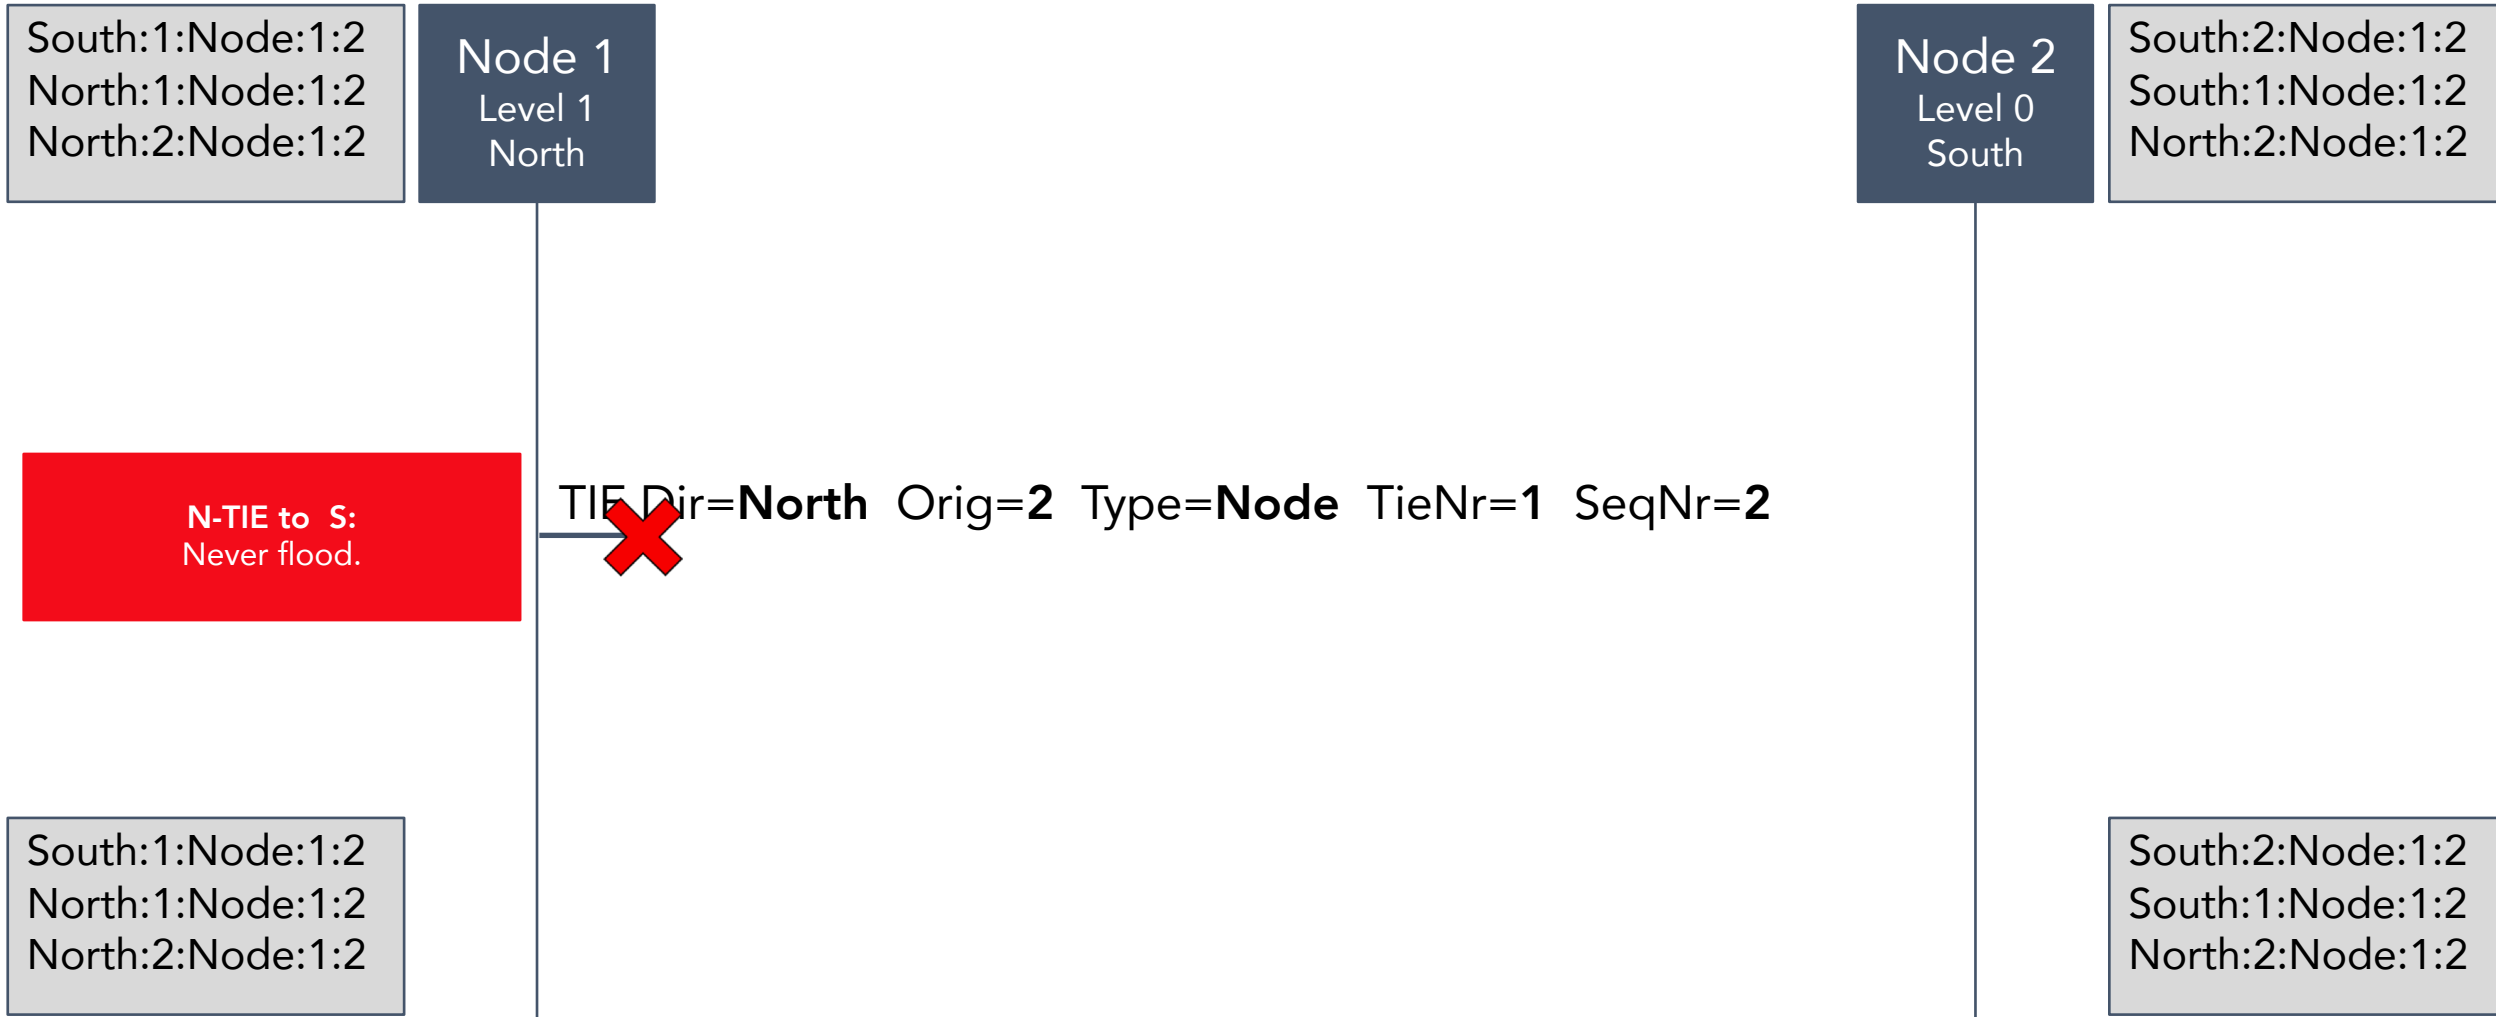


Node 2 sends its node TIE to node 1

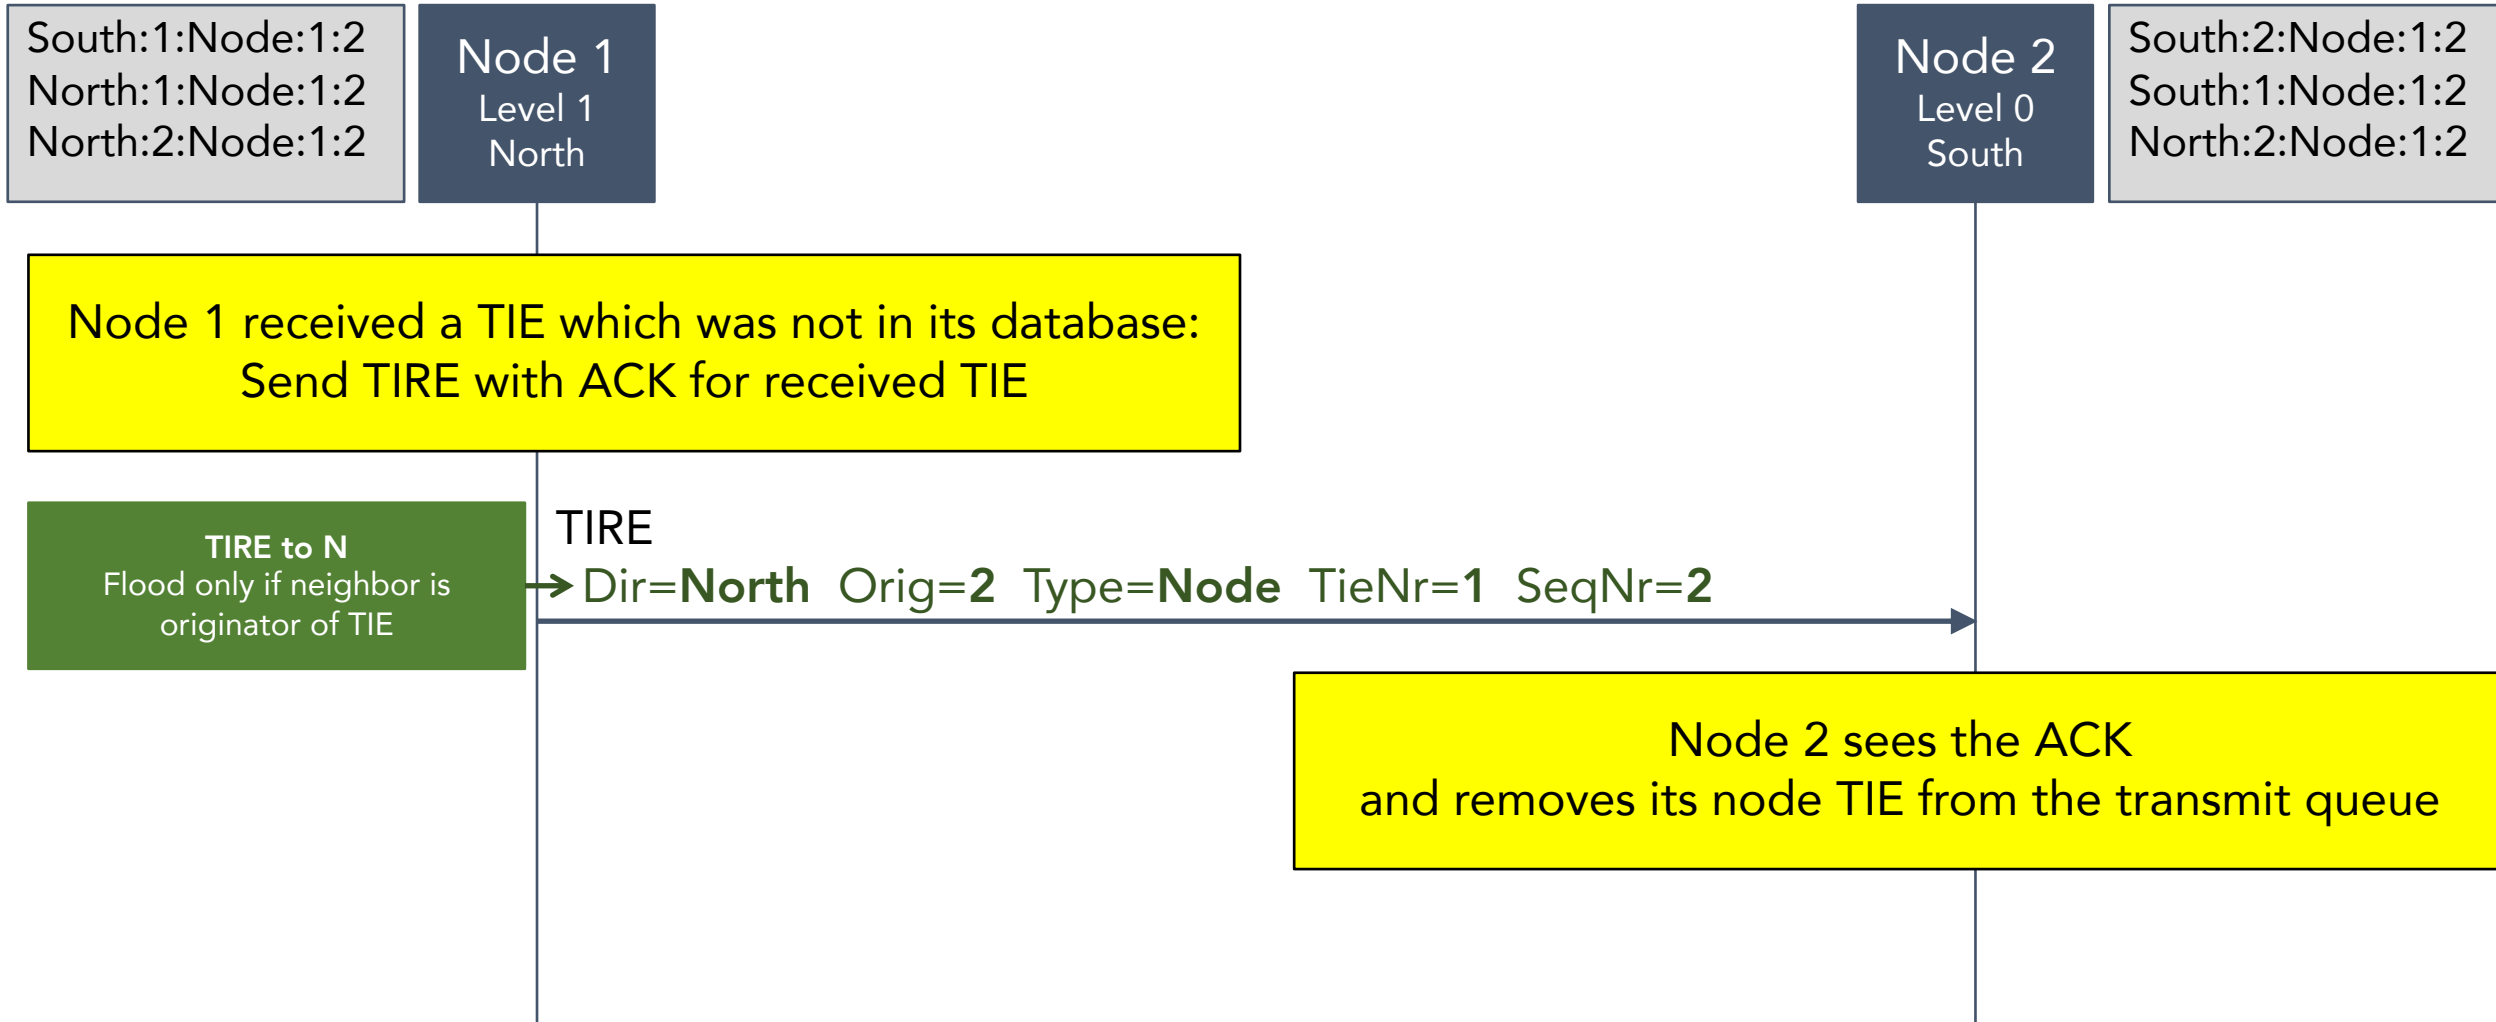


Node 1 does *not* reflect the TIE

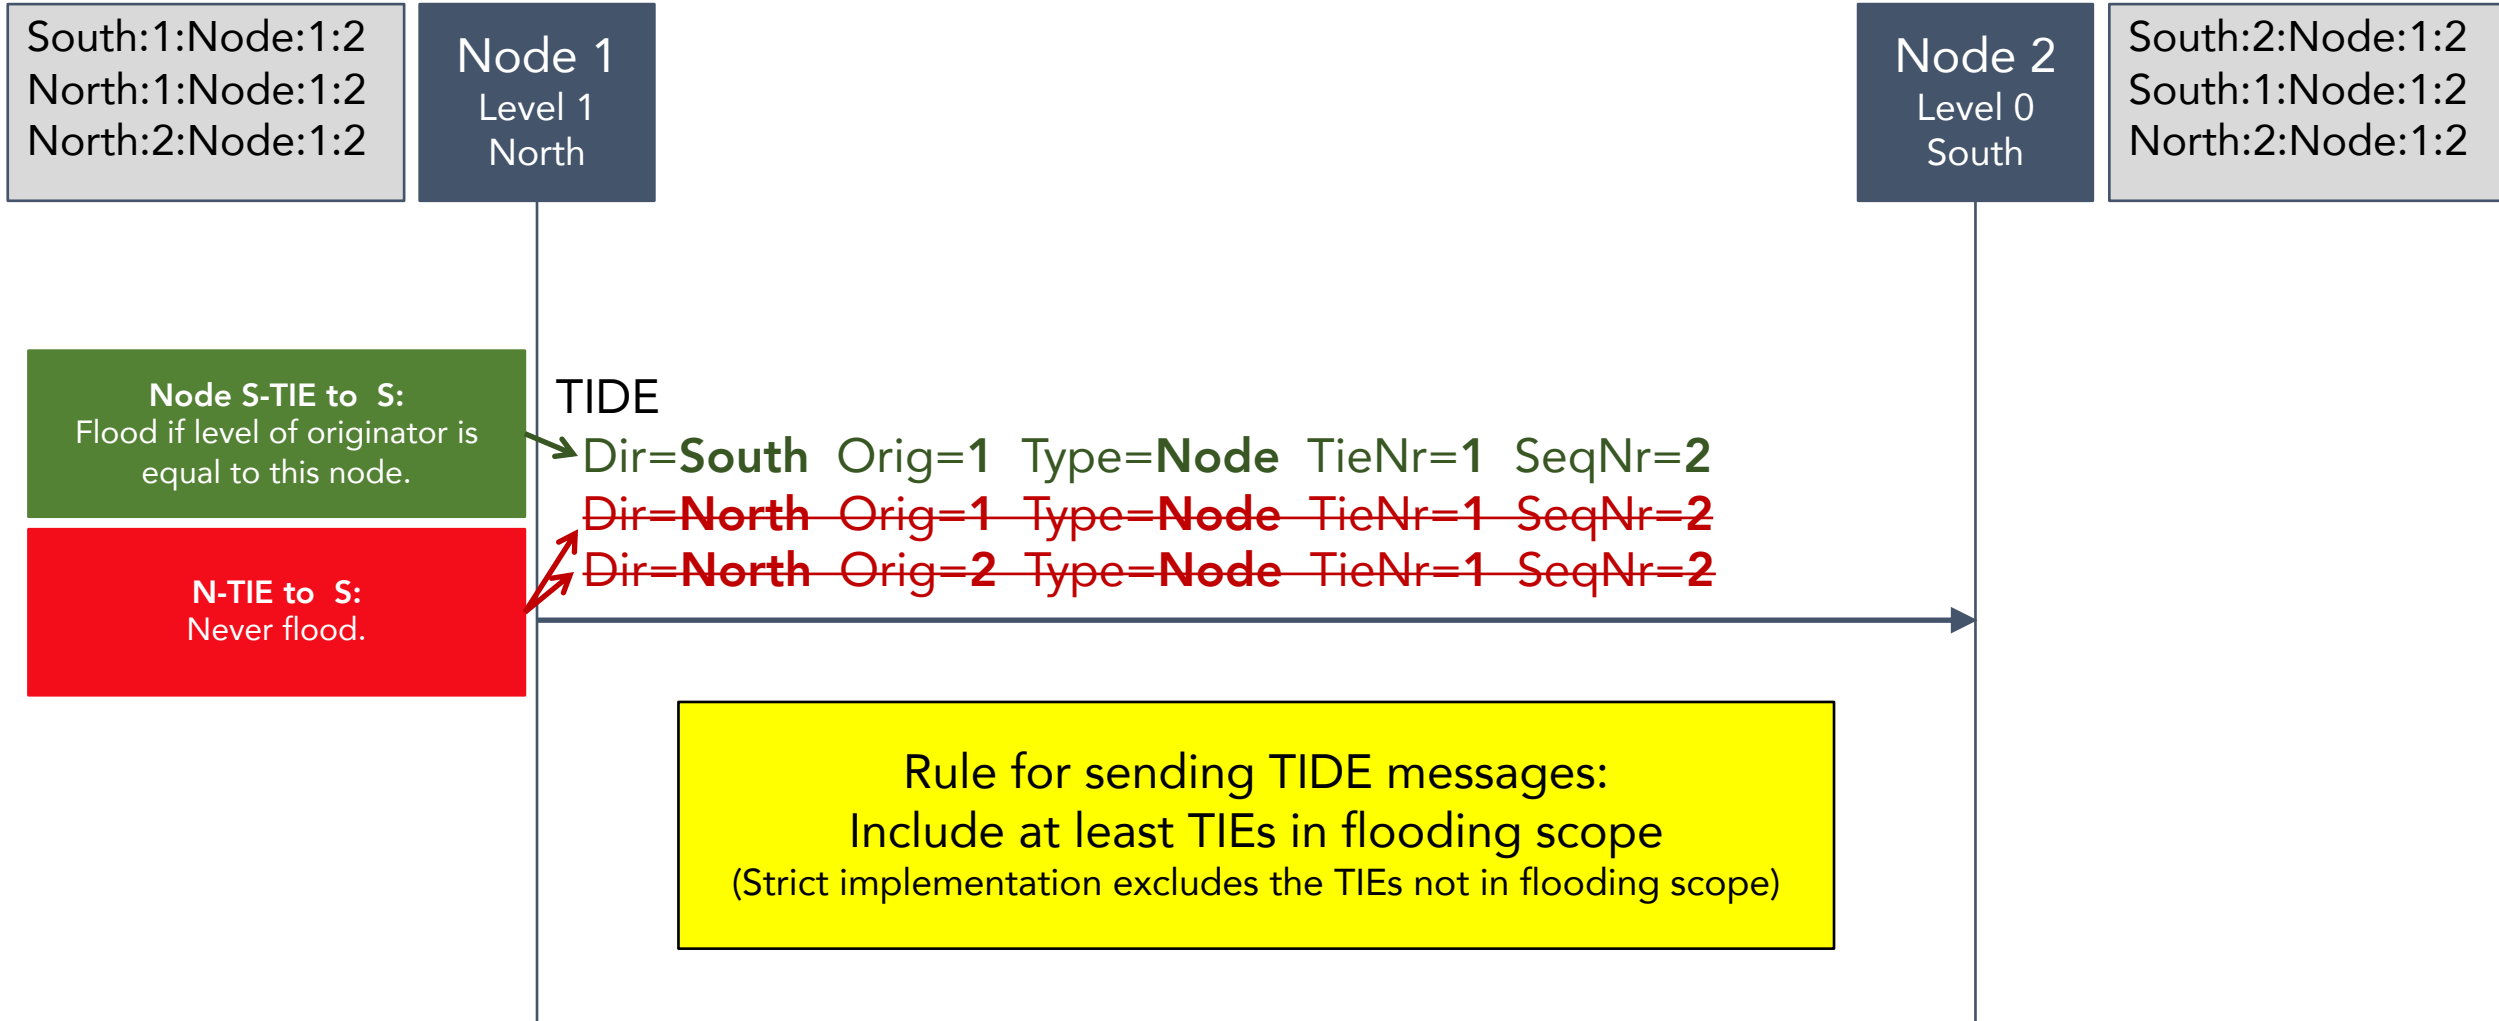
Because the flooding scope rules don't allow it.



TIRE from node 1 to node 2

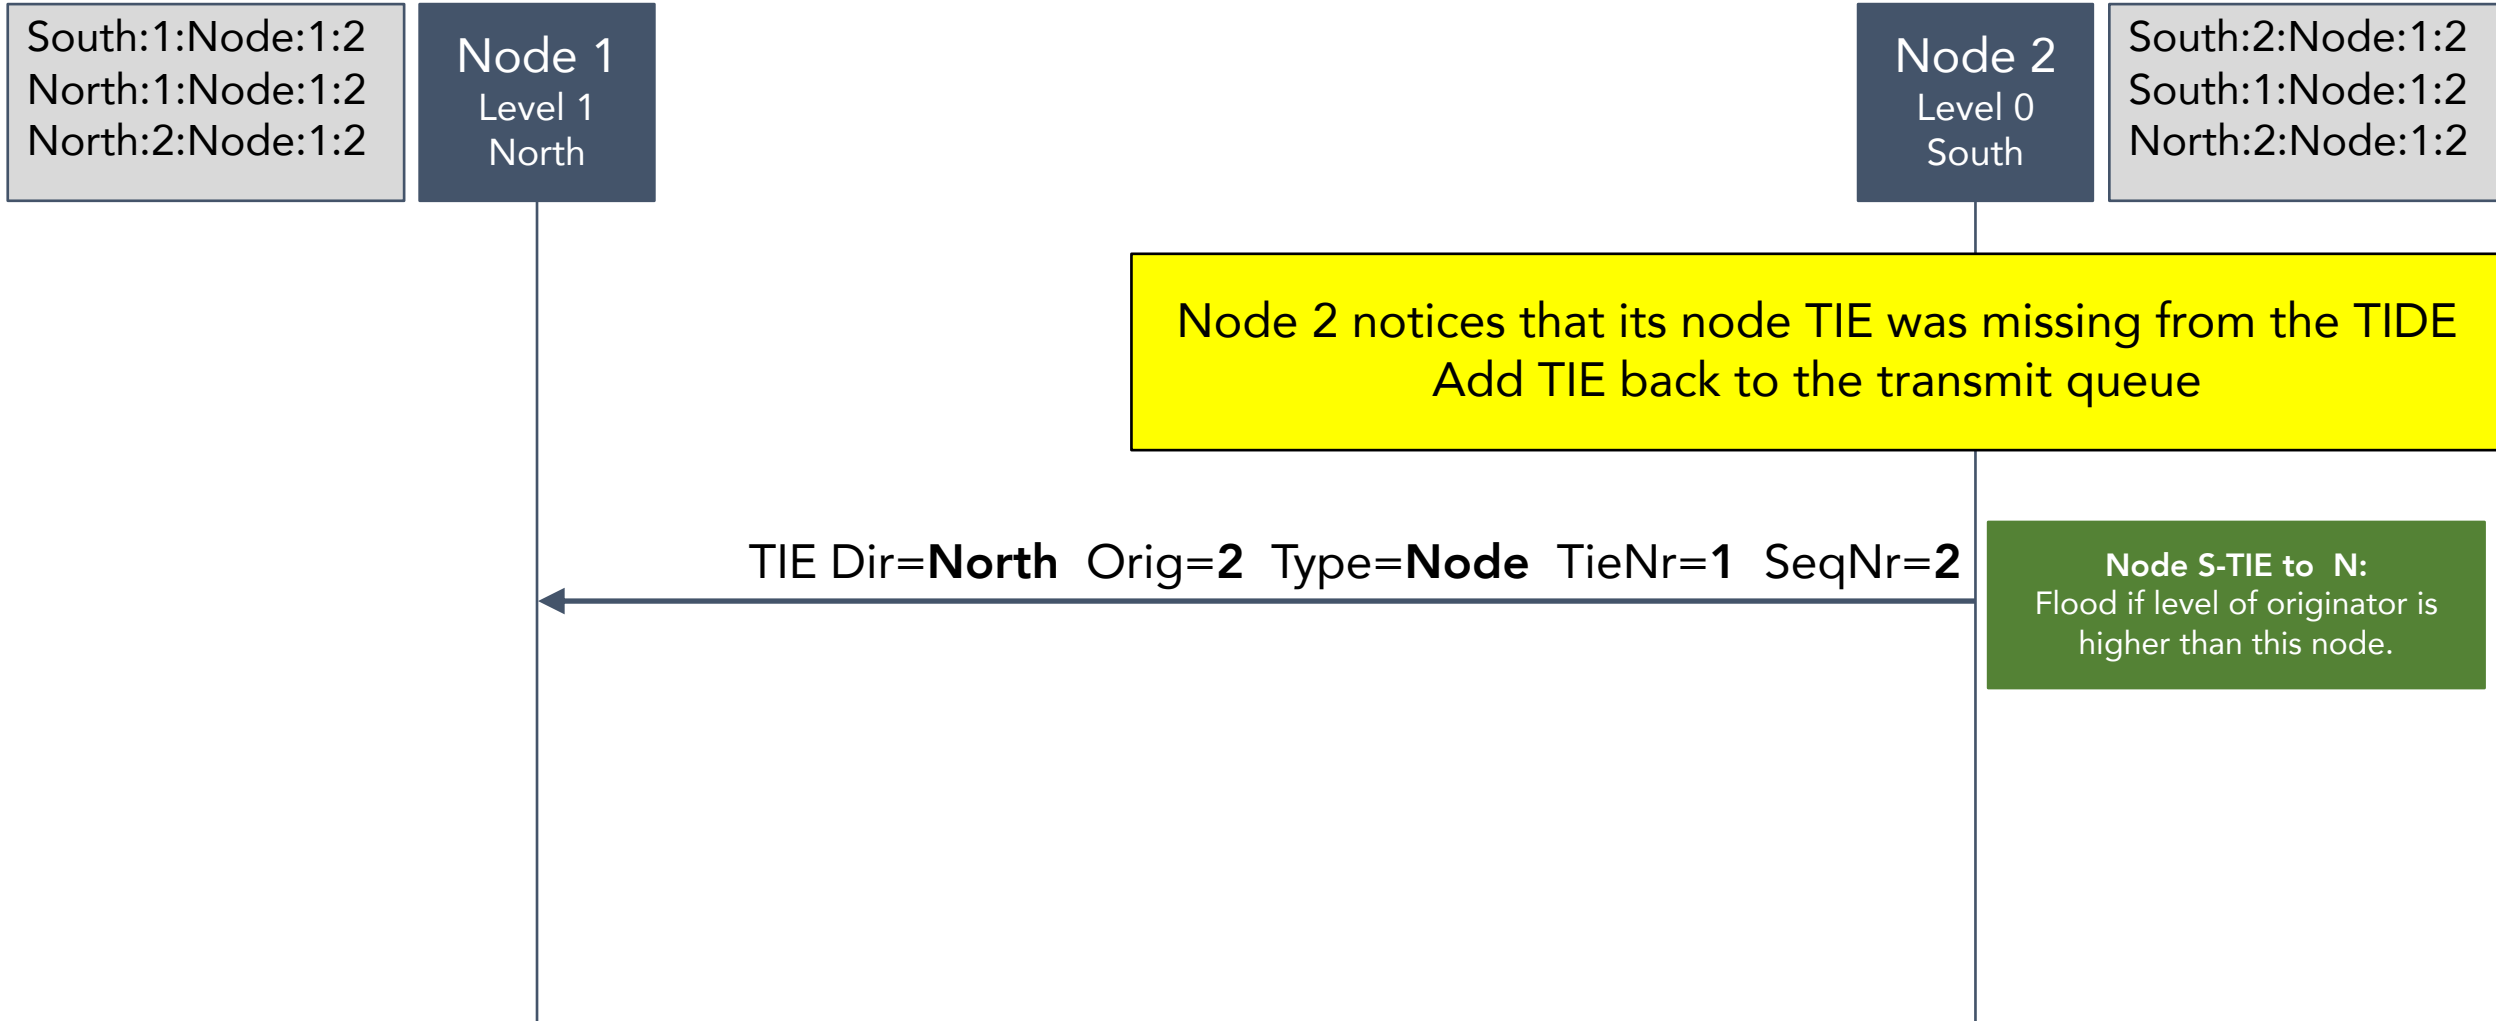


TIDE from node 1 to node 2



Node 2 resends its node TIE

Because it was missing in the TIDE received from node 1



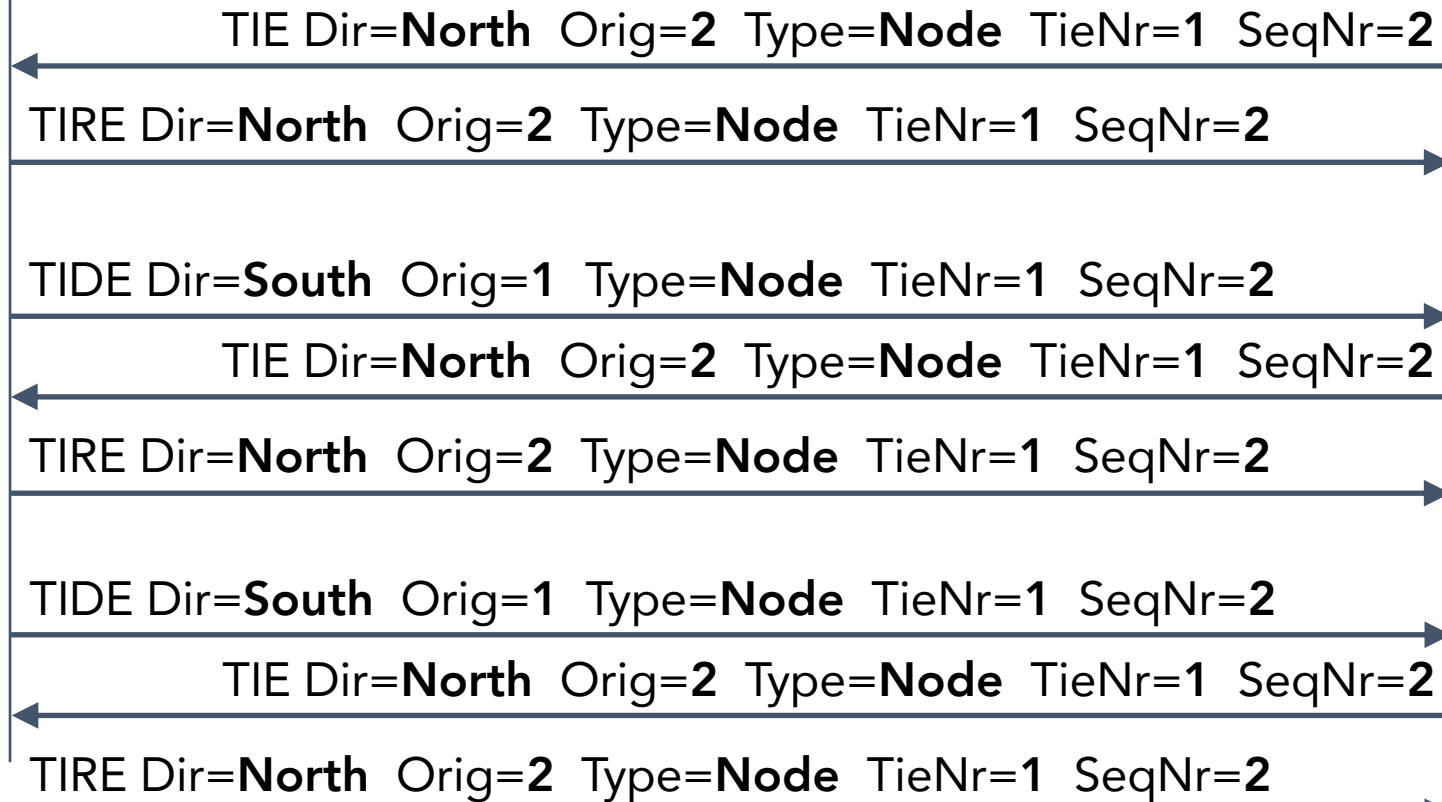
Oscillation

South:1:Node:1:2
North:1:Node:1:2
North:2:Node:1:2

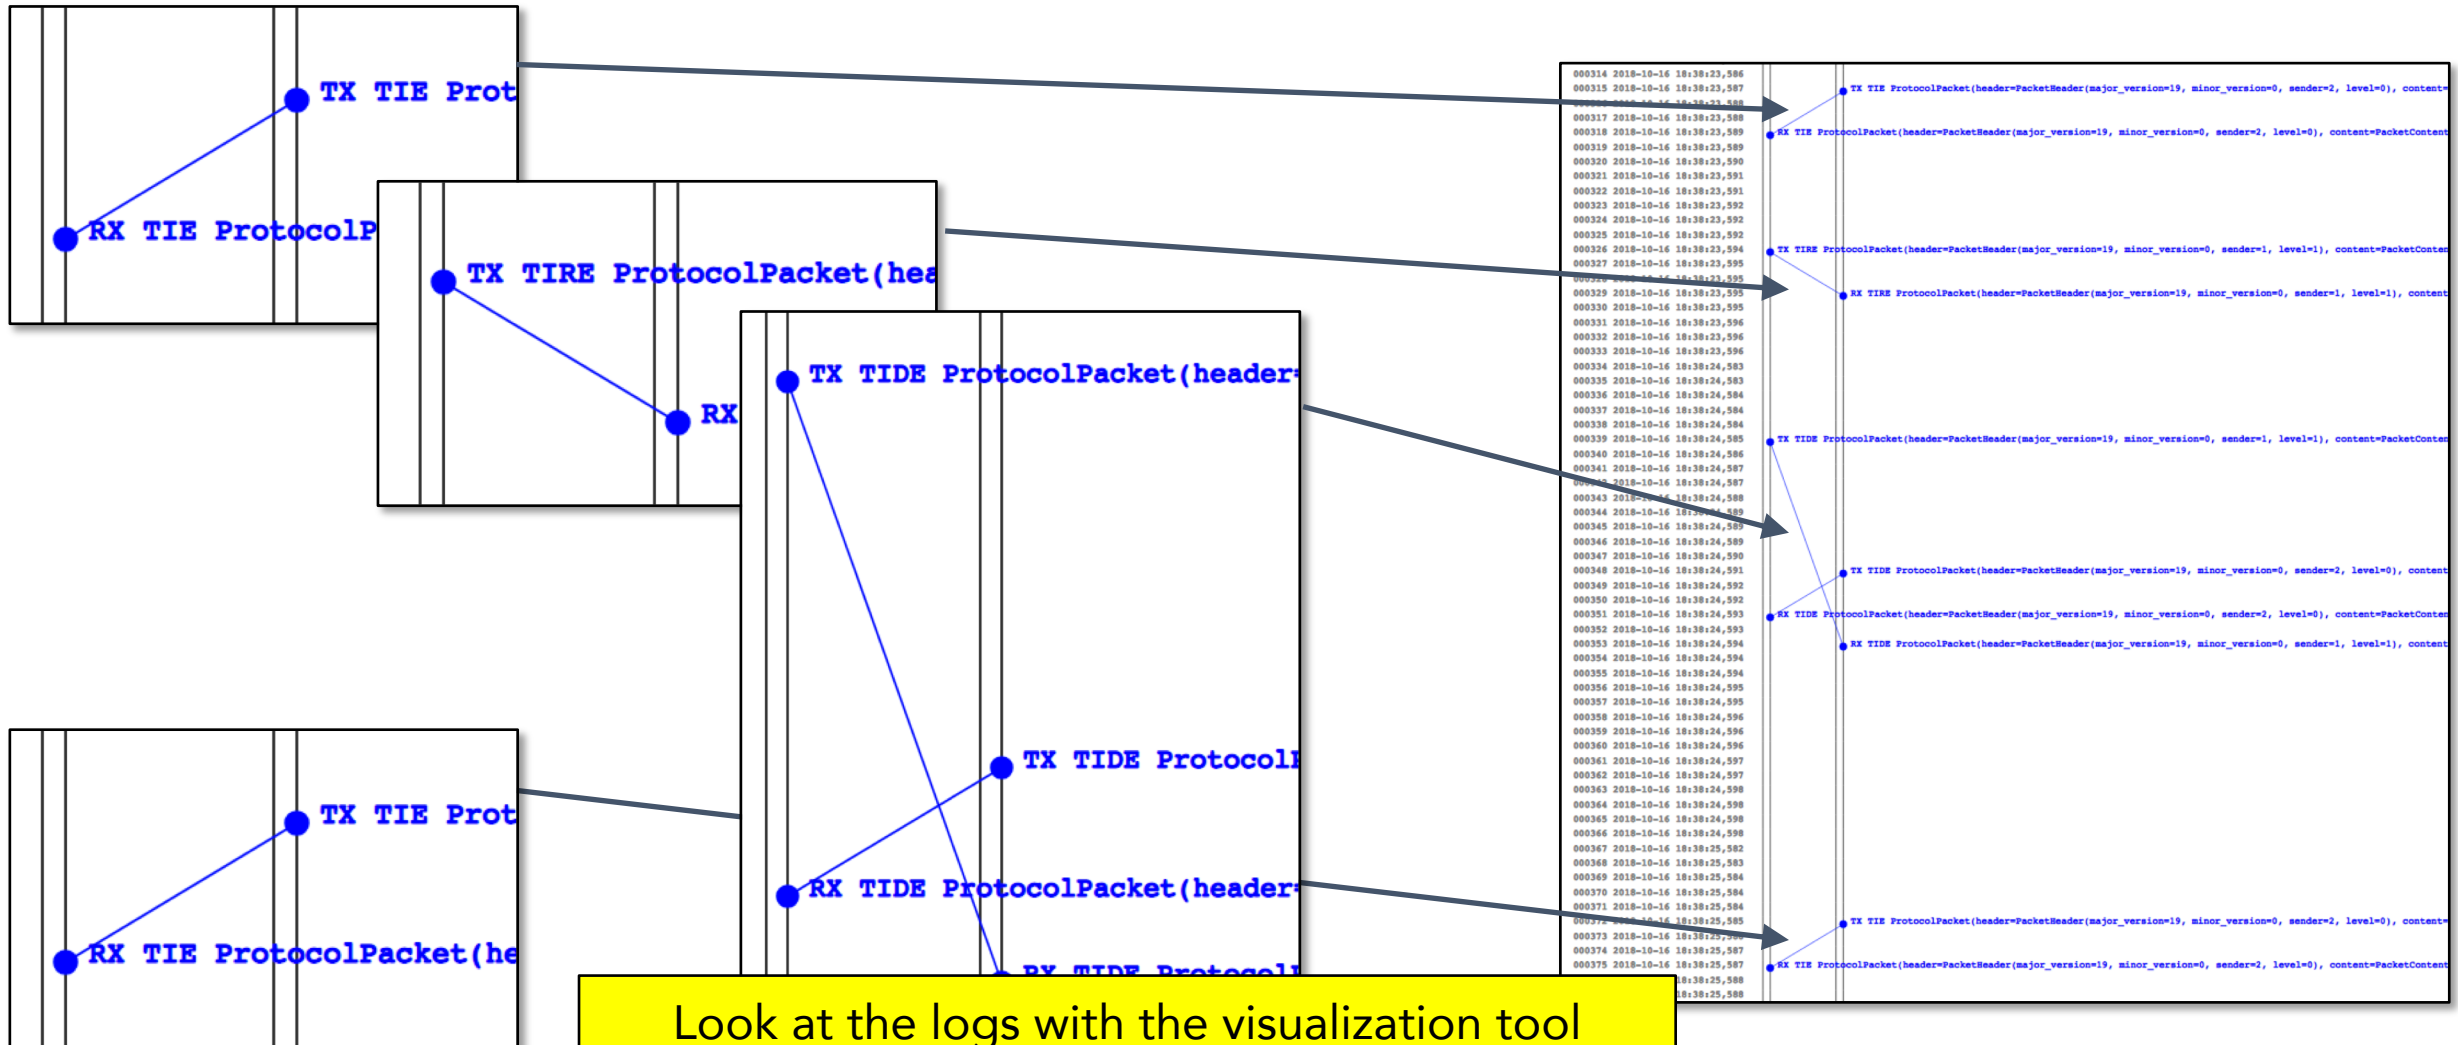
Node 1
Level 1
North

Node 2
Level 0
South

South:2:Node:1:2
South:1:Node:1:2
North:2:Node:1:2



What it looks like in "real life"



Look at the logs with the visualization tool
In a stable topology I should not see any TIEs or TIREs
It is easy to see oscillation patterns in the visualization

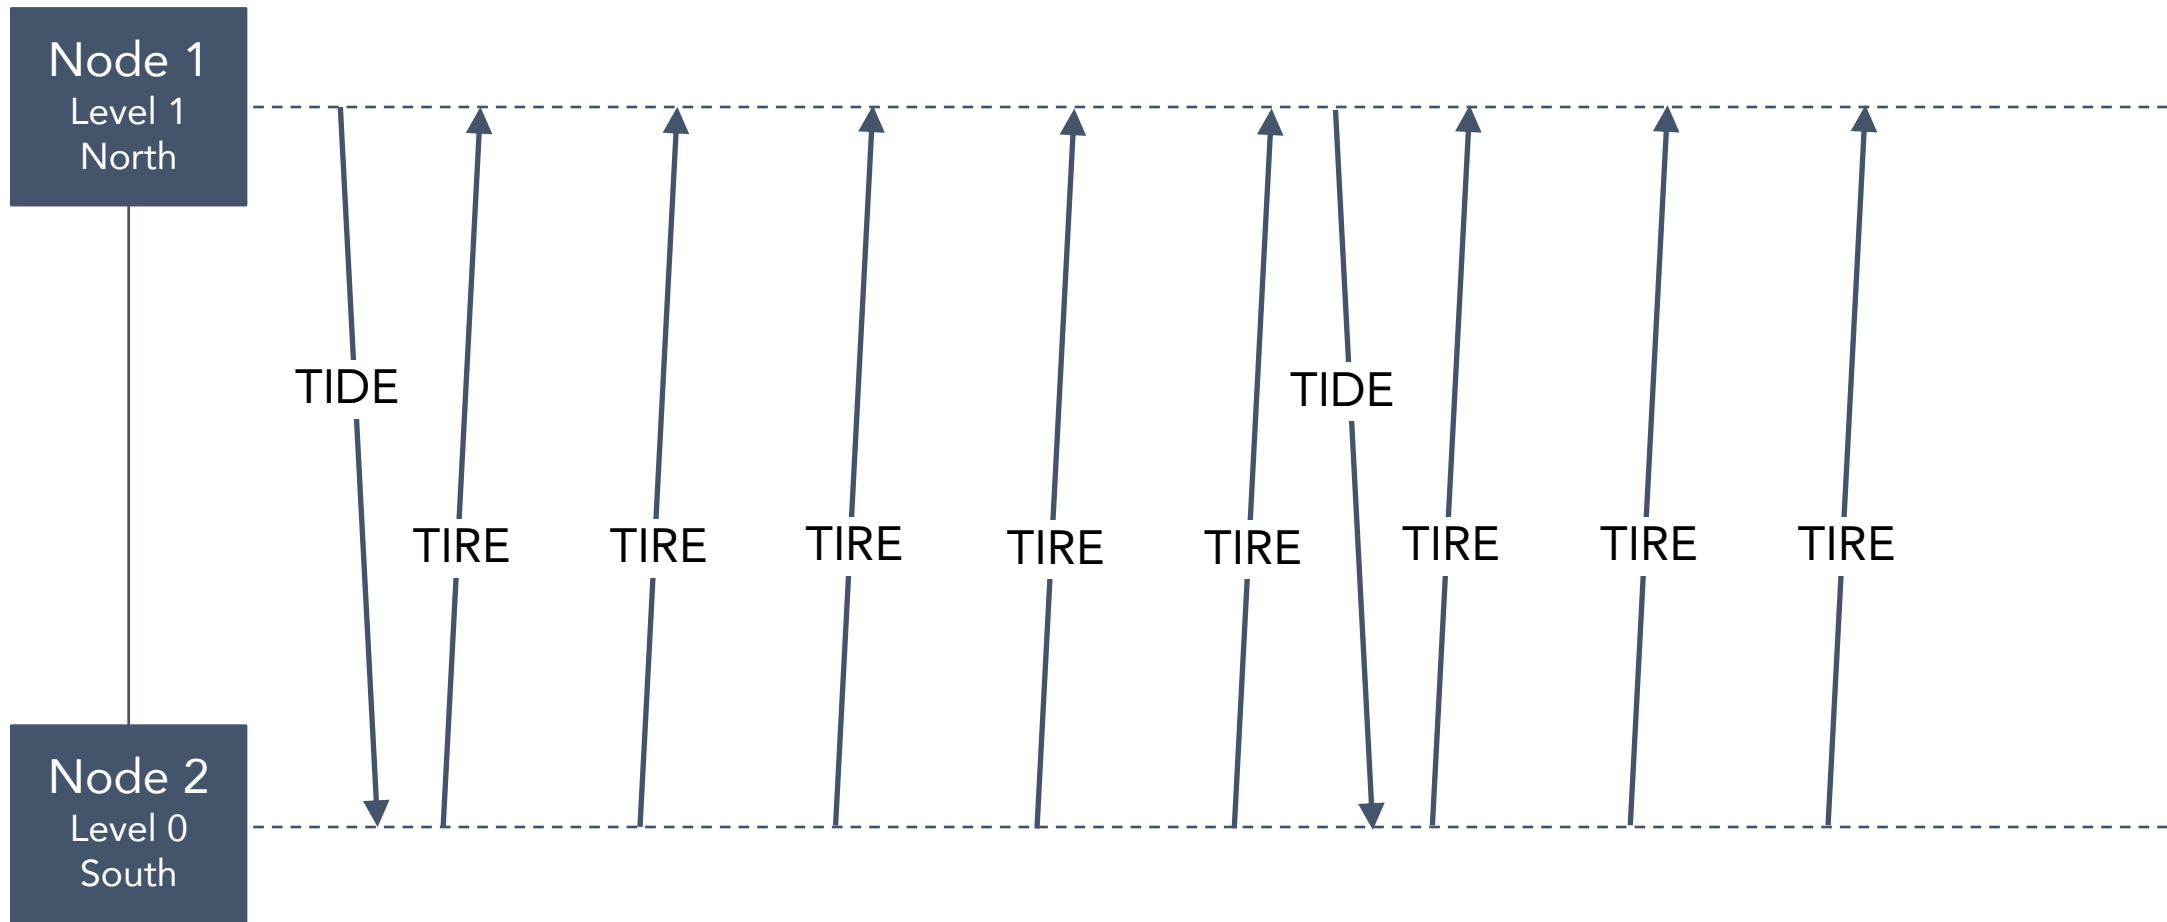
Root cause of the oscillation #1

- A TIDE message sent from node X to node Y serves two purposes:
- **Purpose 1: Announce TIEs that X wants to flood to Y**
 - TIDE **MUST** contain the TIEs that X **MUST** flood to Y
 - TIDE **MAY** contain additional TIEs that X **MUST NOT** flood to Y
 - To avoid having to encode a separate TIDE for each neighbor
 - Y must apply flooding scope rules, and ignore (“not accept”) extra TIEs in TIDE
- **Purpose 2: Acknowledge acceptance of TIEs that Y has flooded to X**
 - If Y has a TIE that it must flood to X and it is missing in the TIDE received from X, then Y will (re)send the TIE to X
 - The TIDE **MUST** contain all TIEs that X has received and accepted from Y (even if X has not intention of sending the TIE to Y, i.e. even if not in flooding scope)
- **The current RIFT specification only captures the first purpose**

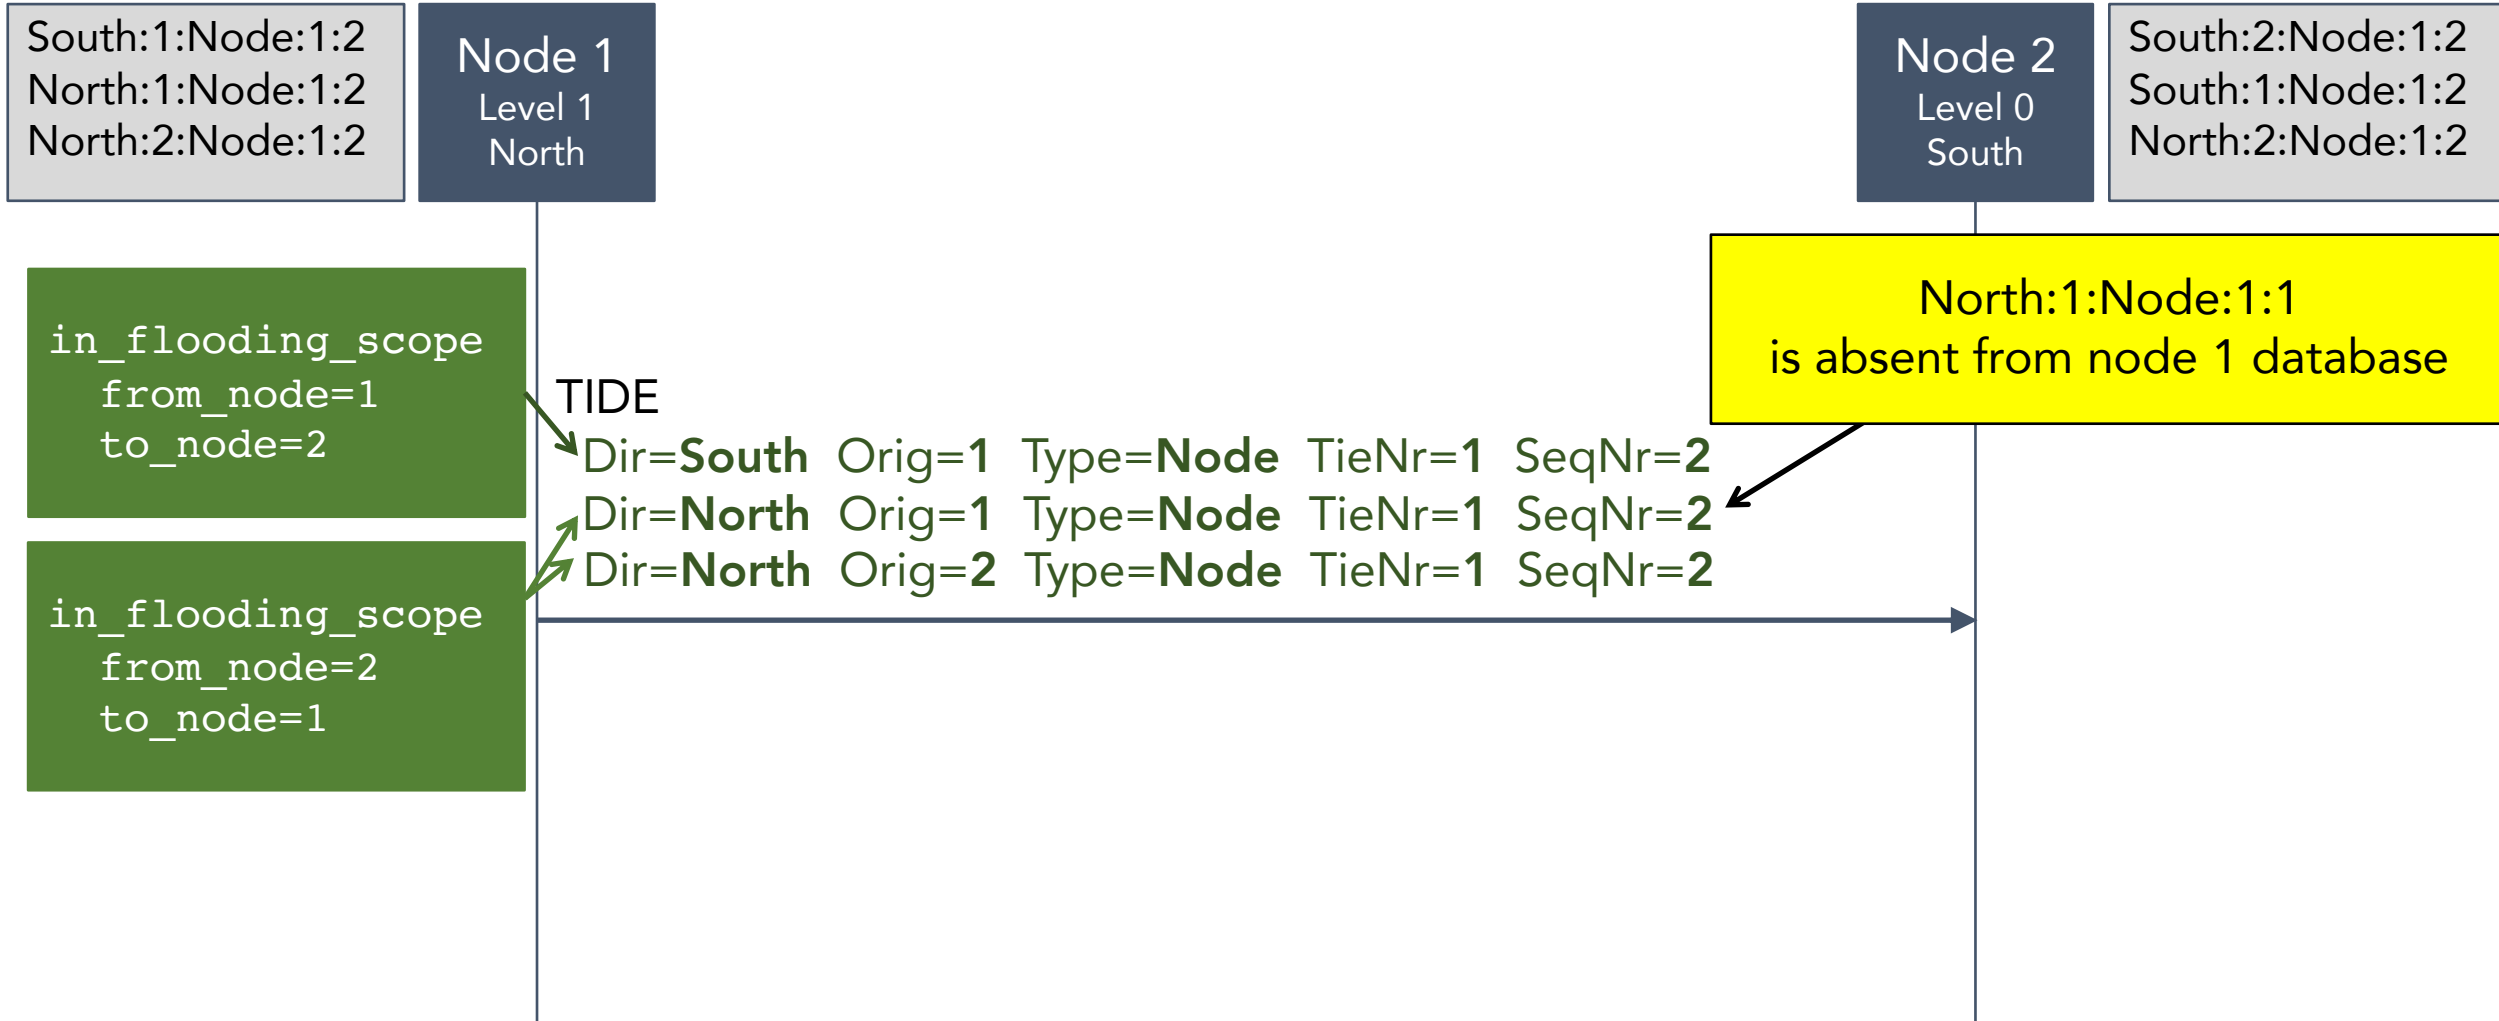
Short-term / ad-hoc solution for oscillation #1

- Update rule for sending TIDEs
- A TIDE message sent from node X to node Y:
 - Purpose 1: MUST include TIEs for which:
`in_flood_scope(from_node=X, to_node=Y, tie_header=T) == true`
 - Purpose 2: MUST include TIEs for which:
`in_flood_scope(from_node=Y, to_node=X, tie_header=T) == true`
 - MAY include additional TIEs

Oscillation #2 (after fixing oscillation #1)



TIDE from node 1 to node 2



Node 2 requests the missing TIE

South:1:Node:1:2
North:1:Node:1:2
North:2:Node:1:2

Node 1
Level 1
North

Node 2
Level 0
South

South:2:Node:1:2
South:1:Node:1:2
North:2:Node:1:2

Flooding scope rules for TIRE say that TIRE should include North:1:Node:1:2 because node 1 is originator of missing TIE

TIRE Dir=**North** Orig=**1** Type=**Node** TieNr=**1** SeqNr=**2**

TIRE to N:
Flood only if neighbor of originator

Node 1 does *not* send the requested TIE

Because the flooding scope rules don't allow it.

South:1:Node:1:2
North:1:Node:1:2
North:2:Node:1:2

Node 1
Level 1
North

Node 2
Level 0
South

South:2:Node:1:2
South:1:Node:1:2
North:2:Node:1:2

N-TIE to S:
Never flood.

TIE Dir=**North** Orig=1 Type=**Node** TieNr=1 SeqNr=2



Analysis of oscillation #2

Root cause of oscillation #2

- A TIRE message sent from node X to node Y serves two purposes:
- Purpose 1: X is requesting a missing TIEs it wants Y to send
- Purpose 2: X is acknowledging acceptance of TIEs it has received from Y
- The current TIRE flooding rule only captures the second purpose

Potential short-term / ad-hoc solution:

- Different TIRE rules for request missing / acknowledge
- Not (yet) implemented – want to step back and consider more drastic measures
- Note: so far we have only considered a trivial 2-node topology and not even looked at more complex topologies

A game of “whack-a-mole”

1. Find an oscillation scenario
2. Tweak the flooding scope rules to fix it.
3. Find a new oscillation scenario which is a result of the tweaked rules.
4. Go to step 2.



System behavior (oscillations)
extremely sensitive to rule details

Proposed long-term / fundamental solution

- Basic idea: encode target flooding scope into TIE header, e.g.:
 - Flood to "node 4 and direct south neighbors"
 - Flood to "node 18 and south-cone from there"
 - Flood to "level 0 and all north levels"
 - Flood to "level 2 and direct south level"
- Just a few bytes in the TIE header (from-where, direction, how-far)
- Advantages:
 - Explicitly signal intent, instead of trying to reverse-engineer intent from rules
 - I expect this to be simpler to implement and to understand behavior
 - Originator can control scope for individual TIEs (e.g. different keys in KV)
- **More detailed proposal and analysis to follow**