Netsukuku
Peer to Peer Network

Luca Dionisi
27 11 2009

- The goals of Netsukuku
  - Sustainable Infrastructure
  - Accessibility, freedom of expression and privacy
- How it works
  - Why a new distributed protocol
  - Operational Overview
  - Hierarchical Structure
  - Architectural weaknesses and solutions
- What it achieves

- The goals of Netsukuku
  - Sustainable Infrastructure
  - Accessibility, freedom of expression and privacy
- How it works
  - Why a new distributed protocol
  - Operational Overview
  - Hierarchical Structure
  - Architectural weaknesses and solutions
- What it achieves

Netsukuku

a definition

Netsukuku is a P2P network system designed to handle massive numbers of nodes with minimal consumption of CPU and memory resources.

It can be used to build a world-wide distributed, fault-tolerant, anonymous, and censorship-immune network, fully independent from the Internet.

Netsukuku

a definition

Netsukuku is a P2P network system designed to handle massive numbers of nodes with minimal consumption of CPU and memory resources.

It can be used to build a **world-wide** distributed, fault-tolerant, anonymous, and censorship-immune network, **fully independent from the Internet**.

Netsukuku

a definition

Why do we need a **world-wide** network,
fully independent from the Internet ?

1. price
2. sustainable infrastructure
3. accessibility, freedom of expression and privacy

- The goals of Netsukuku
  - Sustainable Infrastructure
  - Accessibility, freedom of expression and privacy
- How it works
  - Why a new distributed protocol
  - Operational Overview
  - Hierarchical Structure
  - Architectural weaknesses and solutions
- What it achieves

Sustainable Infrastructure

going to improve benefit/cost ratio

–

Sustainable Infrastructure

going to improve benefit/cost ratio

Mesh network community

- we own the infrastructure (devices and media)

Sustainable Infrastructure

going to improve benefit/cost ratio

Mesh network community

- we own the infrastructure (devices and media)
- we pay for them una-tantum

Sustainable Infrastructure

going to improve benefit/cost ratio

Mesh network community

- we own the infrastructure (devices and media)
- we pay for them una-tantum, on a voluntary basis

Sustainable Infrastructure

going to improve benefit/cost ratio

Internet

- Infrastructure is owned by telecom companies
- They are controlled by big corporations

Sustainable Infrastructure

going to improve benefit/cost ratio

Telecommunication companies

http://en.wikipedia.org/wiki/Telecom#Society_and_telecommunication
Telecommunication has a significant social, cultural and economic impact on modern society. In 2006, estimates placed the telecommunication industry's revenue at $1.2 trillion (ndr 1,200,000,000,000) or just under 3% of the gross world product (official exchange rate).

Sustainable Infrastructure

going to improve benefit/cost ratio

Telecommunication companies

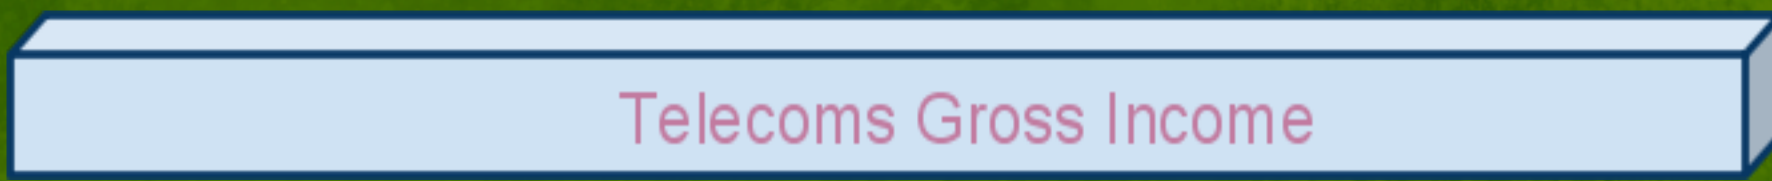

Sustainable Infrastructure

going to improve benefit/cost ratio

Telecommunication companies

Sustainable Infrastructure

going to improve benefit/cost ratio

Telecommunication companies

Sustainable Infrastructure

going to improve benefit/cost ratio

Telecommunication companies

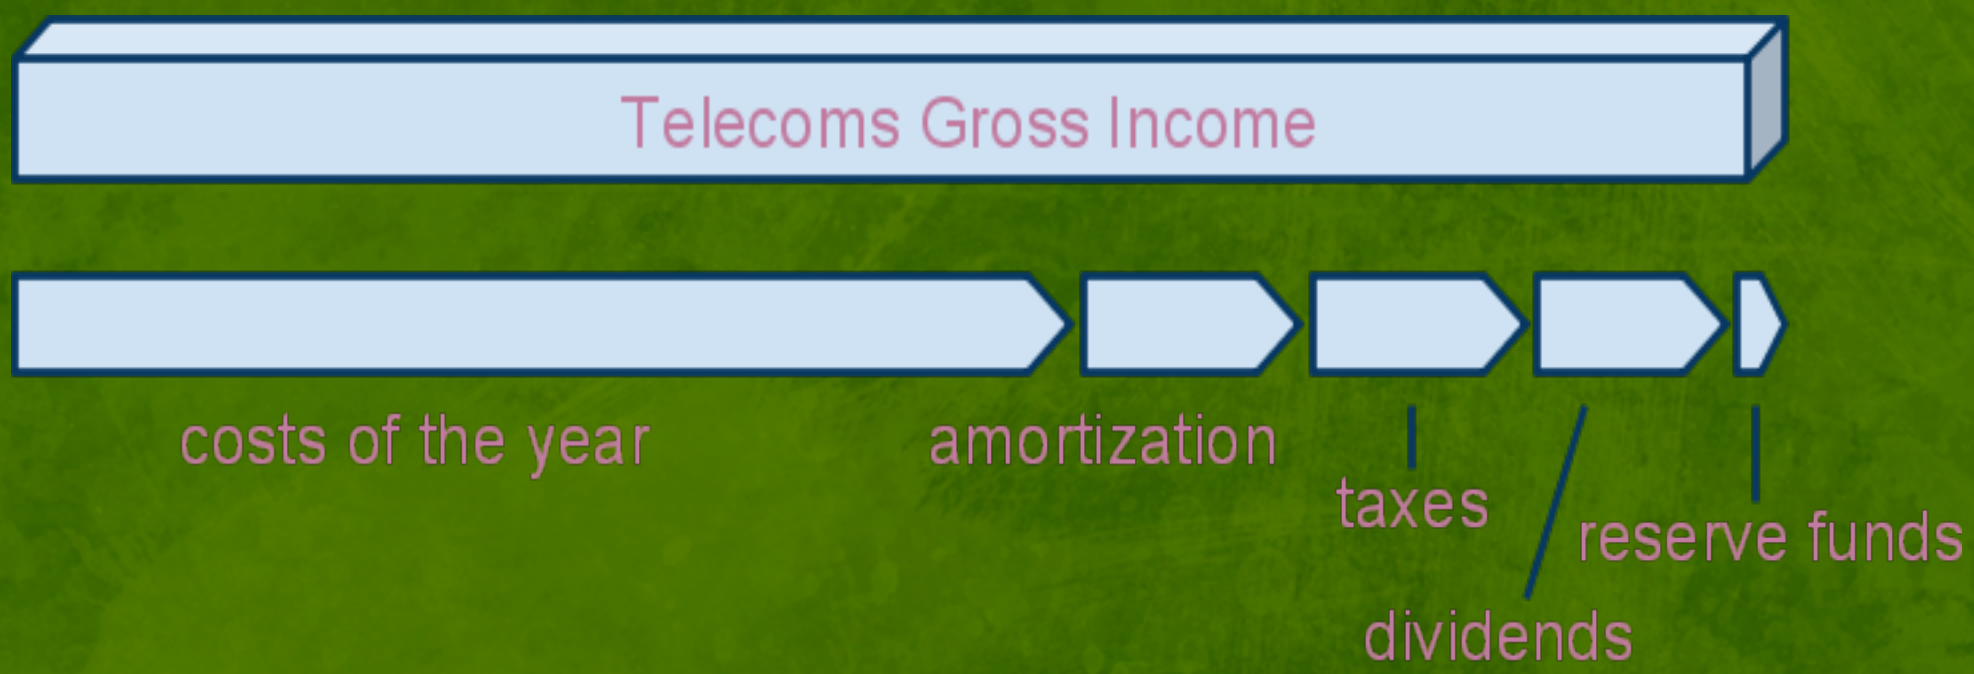

Sustainable Infrastructure

going to improve benefit/cost ratio

Internet

- Easy prediction: prices are going to grow up, infrastructure is going to worsen.

Sustainable Infrastructure

going to improve benefit/cost ratio

Internet

- Easy prediction: prices are going to grow up, infrastructure is going to worsen.
- QoS: a means to do this without much complaints from public opinion.

- The goals of Netsukuku
  - Sustainable Infrastructure
  - Accessibility, freedom of expression and privacy
- How it works
  - Why a new distributed protocol
  - Operational Overview
  - Hierarchical Structure
  - Architectural weaknesses and solutions
- What it achieves

Accessibility, Freedom and Privacy

utopistic in a centralised network

Internet

- Was designed to be a secure, distributed, and failure-resistant communication system.

Accessibility, Freedom and Privacy

utopistic in a centralised network

Internet

- Was designed to be a secure, distributed, and failure-resistant communication system.
- Today paradoxically it is such centralised and hierarchic that ISPs have a complete power, even to cut off entire nations.

Accessibility, Freedom and Privacy

utopistic in a centralised network

Internet

- Was designed to be a secure, distributed, and failure-resistant communication system.
- Today paradoxically it is such centralised and hierarchic that ISPs have a complete power, even to cut off entire nations.
- Today, people can join the Internet only in accordance with the restrictive conditions and terms imposed by the ISPs, and subject to the filters and restrictions placed upon them, often even without their knowledge.

Accessibility

utopistic in a centralised network

Internet

- Today more than 1 billion people can connect to the Internet.

Accessibility

utopistic in a centralised network

Internet

- Today more than 1 billion people can connect to the Internet.
- The remaining 5 billion lacks the economic resources.

Accessibility

utopistic in a centralised network

Internet

- Today more than 1 billion people can connect to the Internet.
- The remaining 5 billion lacks the economic resources.
- They are still waiting for the big ISPs to provide a service within their reach.
  The ISPs don't feel this an economic opportunity for them.

Freedom of expression and privacy

utopistic in a centralised network

Internet

- The centralization also enables other abuses of power, such as censorship of all dissident thought.

Freedom of expression and privacy

utopistic in a centralised network

Internet

- The centralization also enables other abuses of power, such as censorship of all dissident thought.
- Authorities with access to ISP data may use and have used the system as a means to identify and persecute locals who voice dissent.

Freedom of expression and privacy

utopistic in a centralised network

Internet

- The centralization also enables other abuses of power, such as censorship of all dissident thought.
- Authorities with access to ISP data may use and have used the system as a means to identify and persecute locals who voice dissent.
- In this regard, other efforts such as Tor, Freenet and the like, are doomed to fail, since if we use them we still rely on the ISP's backbones, which the ISP controls and monitors.

Accessibility, Freedom and Privacy

more realistic in a peer-2-peer network

P2P

- To regain our rights to accessibility, freedom of expression and privacy in our network (because in the end WE are the network) we need a world-wide network that has at least the potential to be fully independent from the Internet.

“ All I know is that first you've got to get mad! ...get up out of your chairs, open the window, stick your head out and yell: I'm as mad as hell, and I'm not gonna take this anymore!”

Network

- The goals of Netsukuku
  - Sustainable Infrastructure
  - Accessibility, freedom of expression and privacy
- How it works
  - Why a new distributed protocol
  - Operational Overview
  - Hierarchical Structure
  - Architectural weaknesses and solutions
- What it achieves

Netsukuku - how it works

An infrastructure needs hardware and software

- The Netsukuku network is composed of nearby nodes (computers, routers, embedded devices, ...) directly linked each other.

Netsukuku - how it works

An infrastructure needs hardware and software

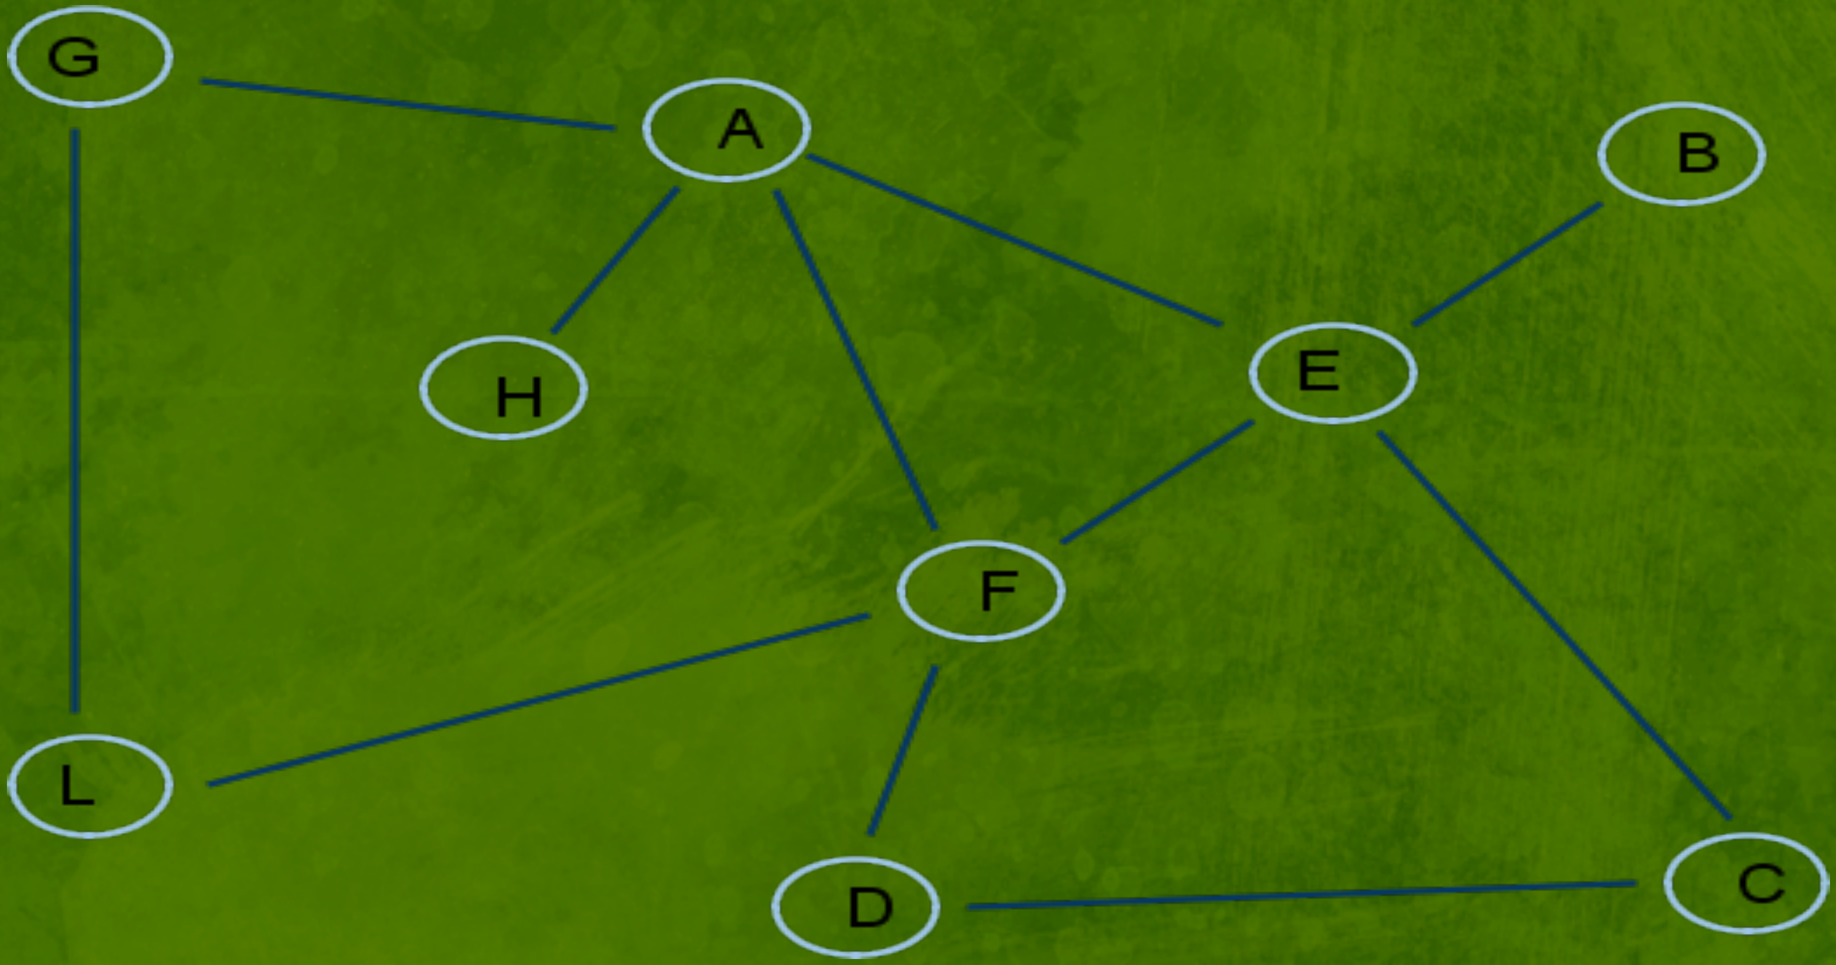

Netsukuku - how it works

An infrastructure needs hardware and software

- The Netsukuku network is composed of nearby nodes (computers, routers, embedded devices, ...) directly linked each other.
- Each node executes the `ntkd` daemon.
  This daemon augments level 3 of the OSI model with its own true distributed routing protocol.

- The goals of Netsukuku
  - Sustainable Infrastructure
  - Accessibility, freedom of expression and privacy
- How it works
  - **Why a new distributed protocol**
  - Operational Overview
  - Hierarchical Structure
  - Architectural weaknesses and solutions
- What it achieves

Why a new routing protocol?

Existing distributed routing protocols

- Interior Gateway
  - Distance Vector
    - RIP, IGRP, EIGRP (flat topology)
  - Link State
    - OSPF, IS-IS (hierarchical system)
- Exterior Gateawy
  - Path Vector
    - BGP (flat topology)

Why a new routing protocol?

Existing distributed routing protocols

An AS is a IP prefix under the control of one entity.
It presents itself to the Internet with a certain routing policy.

- Interior Gateway
  - Between routers of a AS.
- Exterior Gateawy
  - Between Various AS.

Why a new routing protocol?

Existing distributed routing protocols

Distance Vector

- Each node, for each destination, knows:
  - next hop to use to reach it
  - distance
- It suffers the problem of looping paths and count-to infinity
- Uses flat topology. Does not scale.

Why a new routing protocol?

Existing distributed routing protocols

Link State

- Each router, for each other router, knows all their direct neighbours.
- It calculate for itself the best path for each destination.
- Requires a lot of CPU. Does not scale.

Why a new routing protocol?

Existing distributed routing protocols

Path Vector

- Each node, for each destination, knows:
  - next hop to use to reach it
  - all the other hops
  - distance
- It doesn't suffer the problem of looping paths and count-to infinity
- Uses flat topology. Does not scale.

Why a new routing protocol?

Existing distributed routing protocols

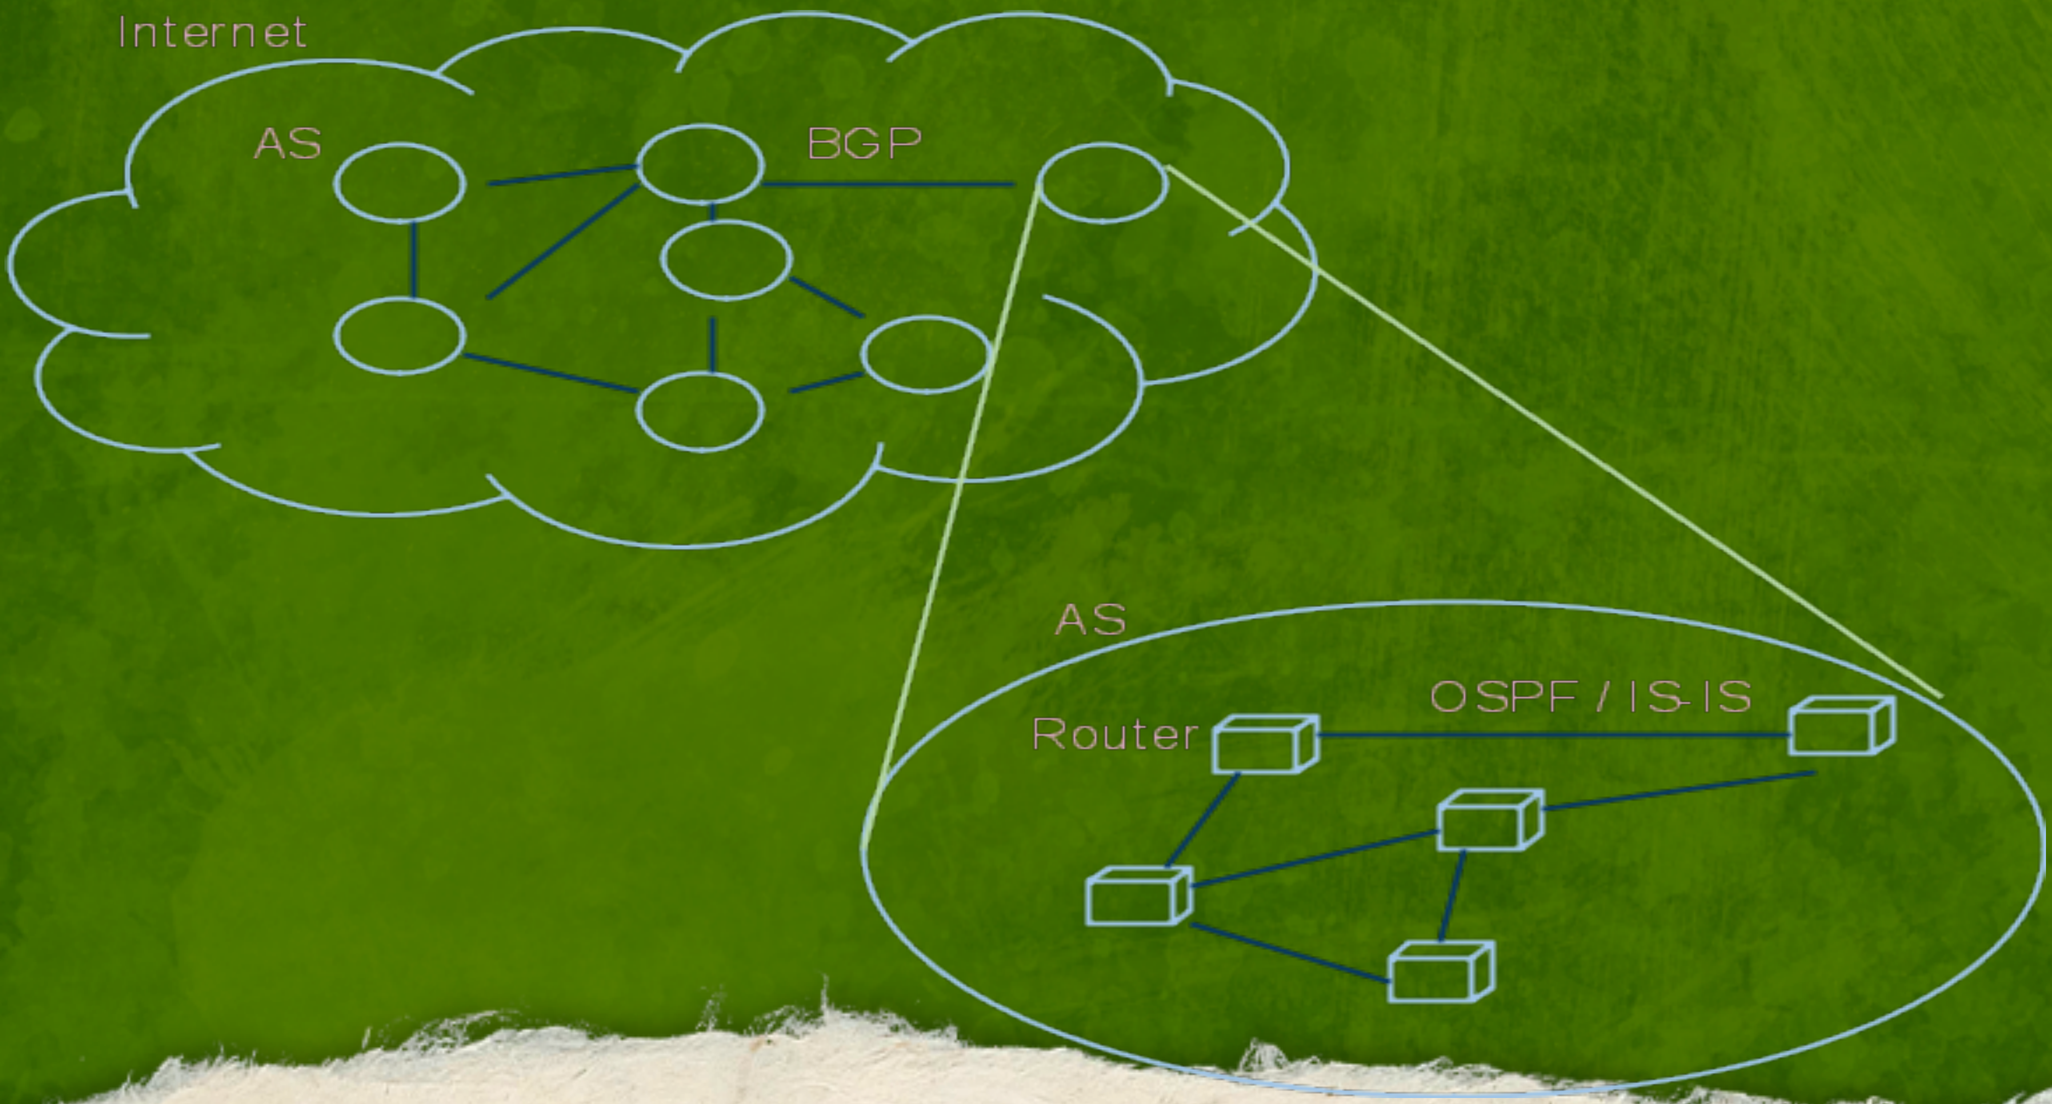

Netsukuku

A new distributed routing protocol

- It is a Path Vector protocol that uses a Hierarchical Structure.
- It is able to run in a distributed manner on nodes with very small computing resources. Also on embedded devices.

- The goals of Netsukuku
  - Sustainable Infrastructure
  - Accessibility, freedom of expression and privacy
- How it works
  - Why a new distributed protocol
  - Operational Overview
  - Hierarchical Structure
  - Architectural weaknesses and solutions
- What it achieves

Operational Overview

Recall the network graph

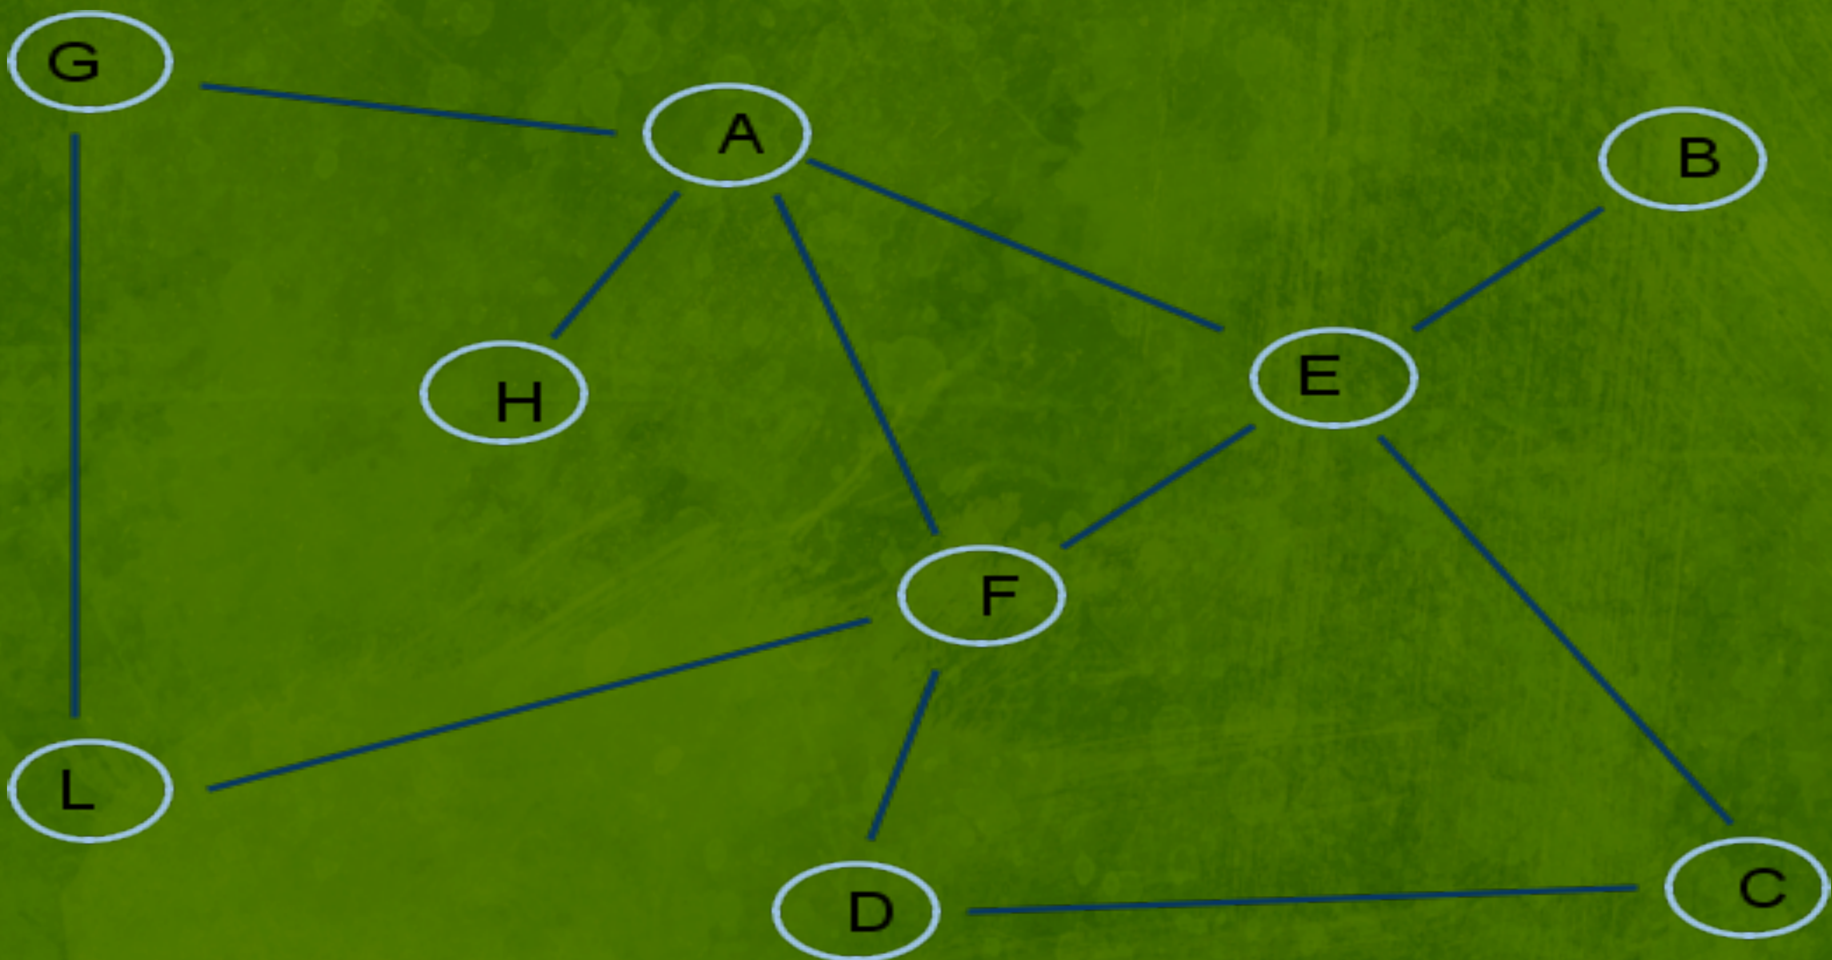

Operational Overview

Stability of the network

At a given point in time, node A has this knowledge:

Operational Overview

Stability of the network

At a given point in time, node A has this knowledge:

- Existence of all other nodes

Operational Overview

Stability of the network

At a given point in time, node A has this knowledge:

- Existence of all other nodes
- For each node D of the network, a subset of all the paths potentially going from A to D

Operational Overview

Recall the network graph

Operational Overview

Recall the network graph

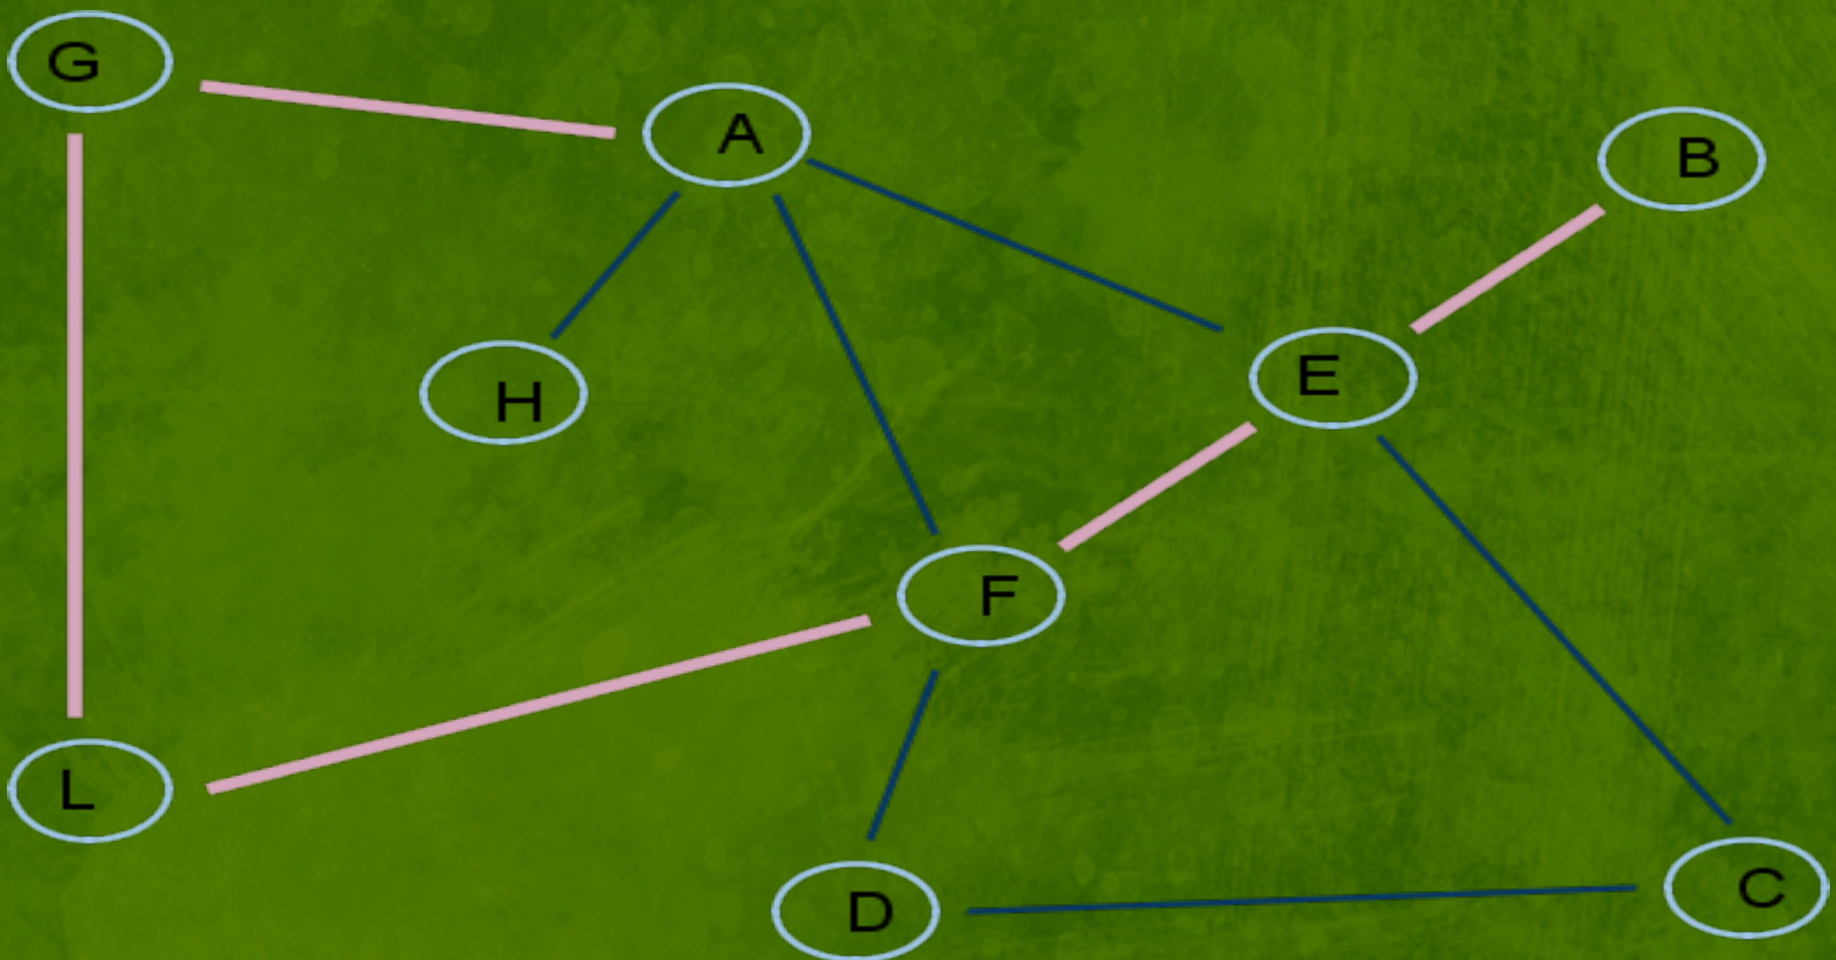

Operational Overview

Recall the network graph

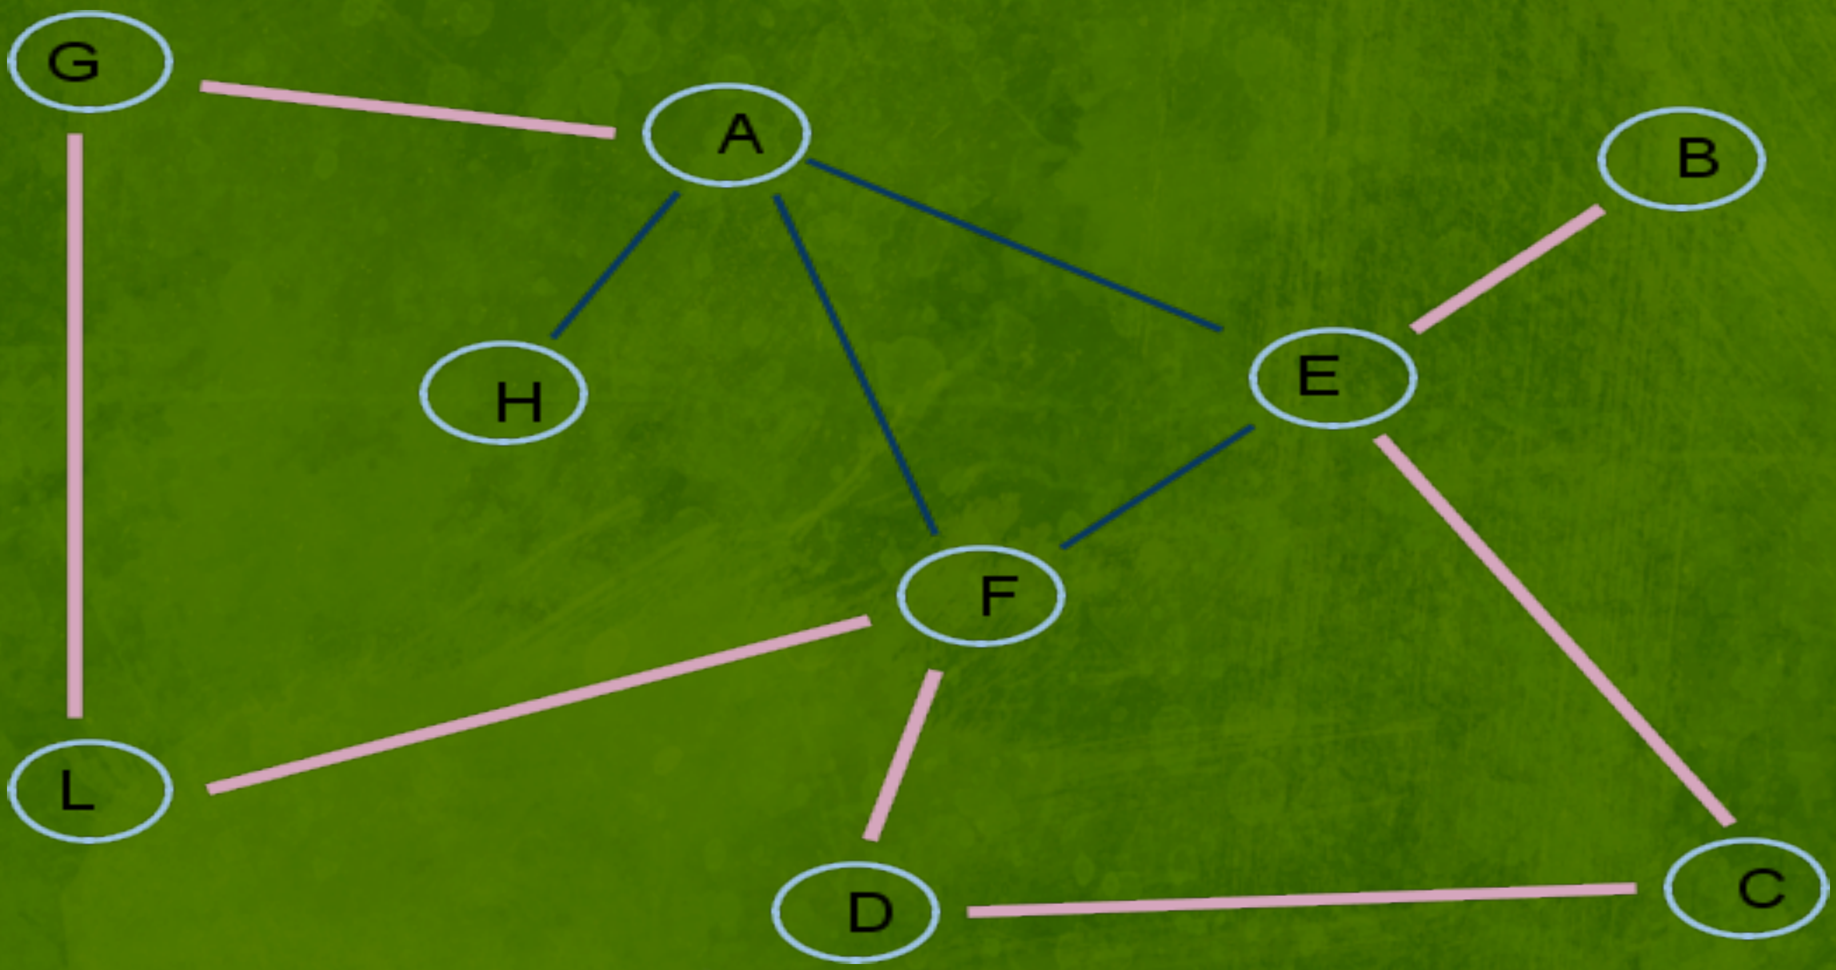

Operational Overview

Recall the network graph

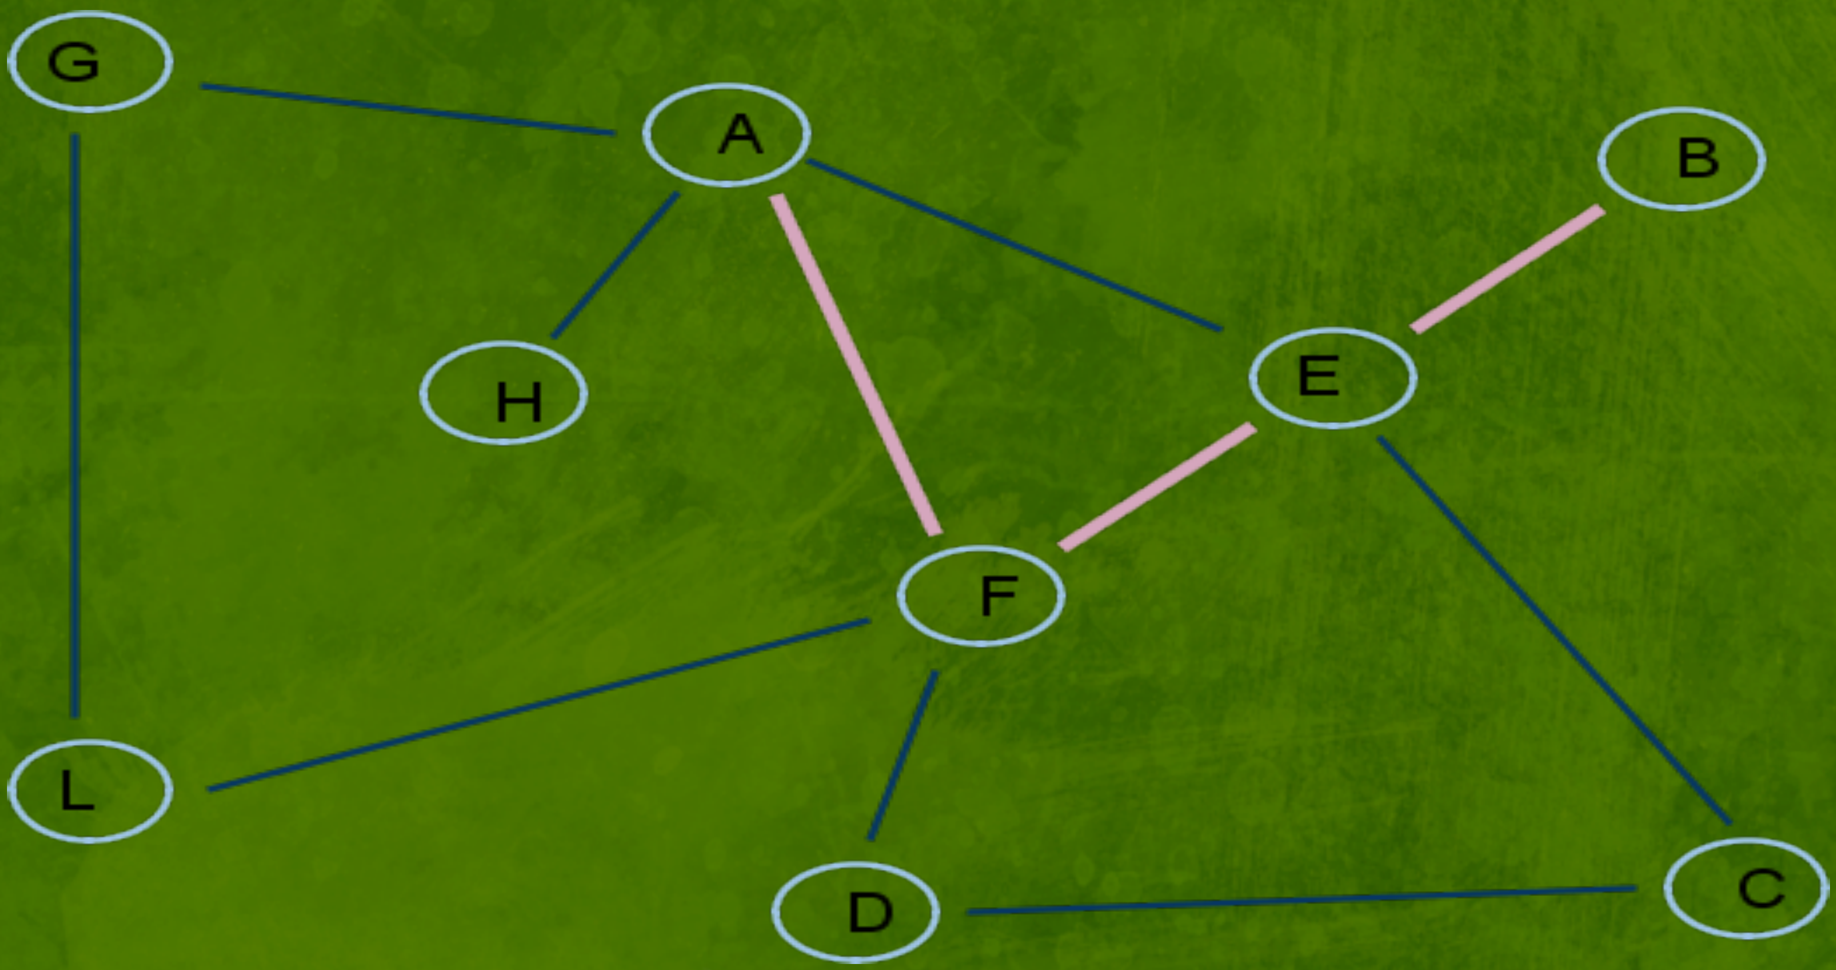

Operational Overview

Recall the network graph

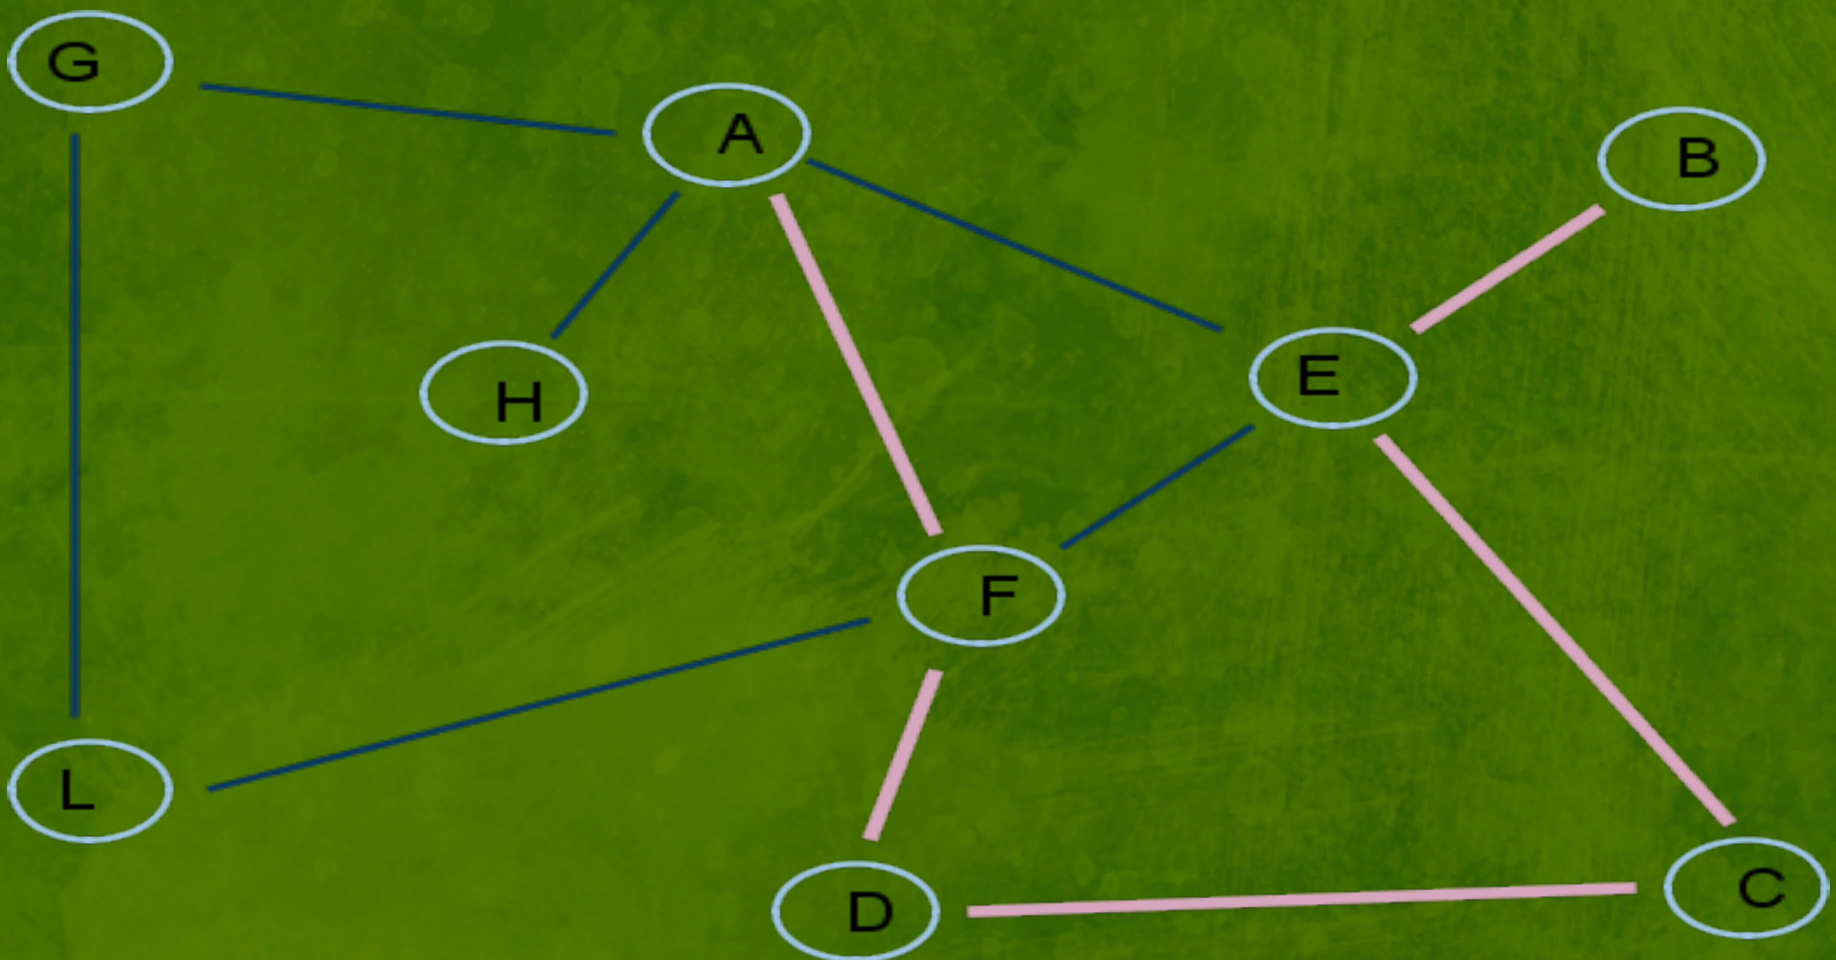

Operational Overview

Recall the network graph

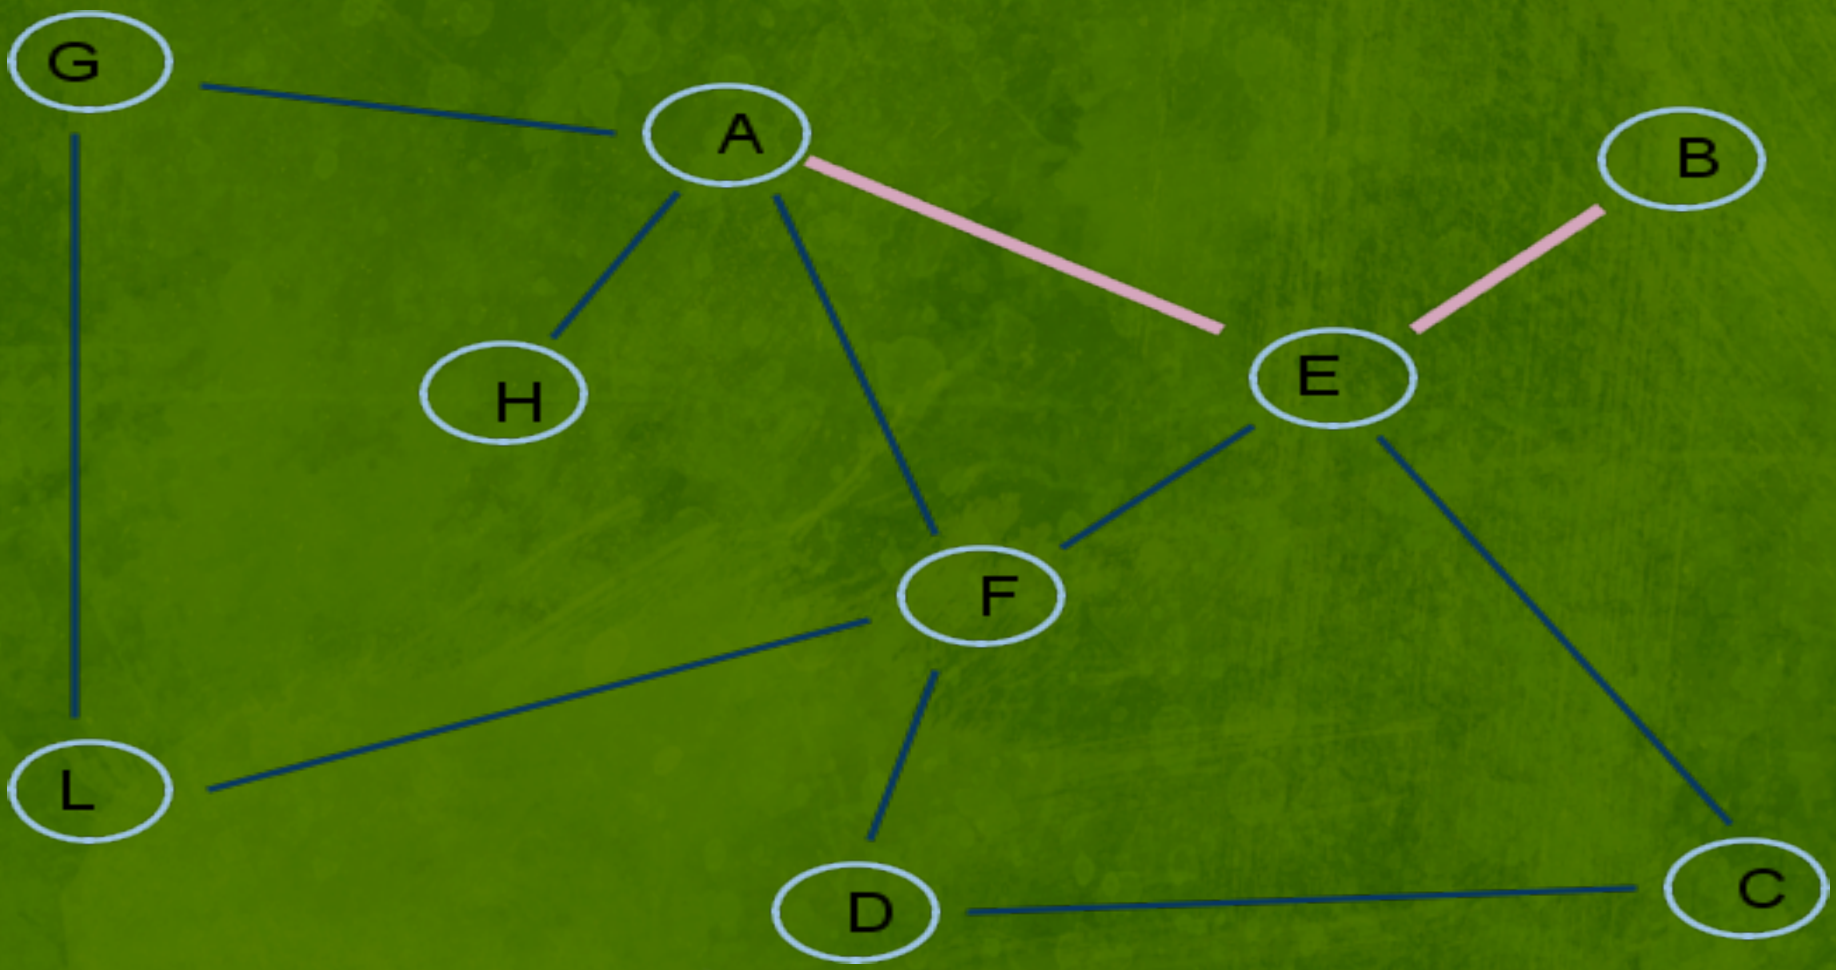

Operational Overview

Recall the network graph

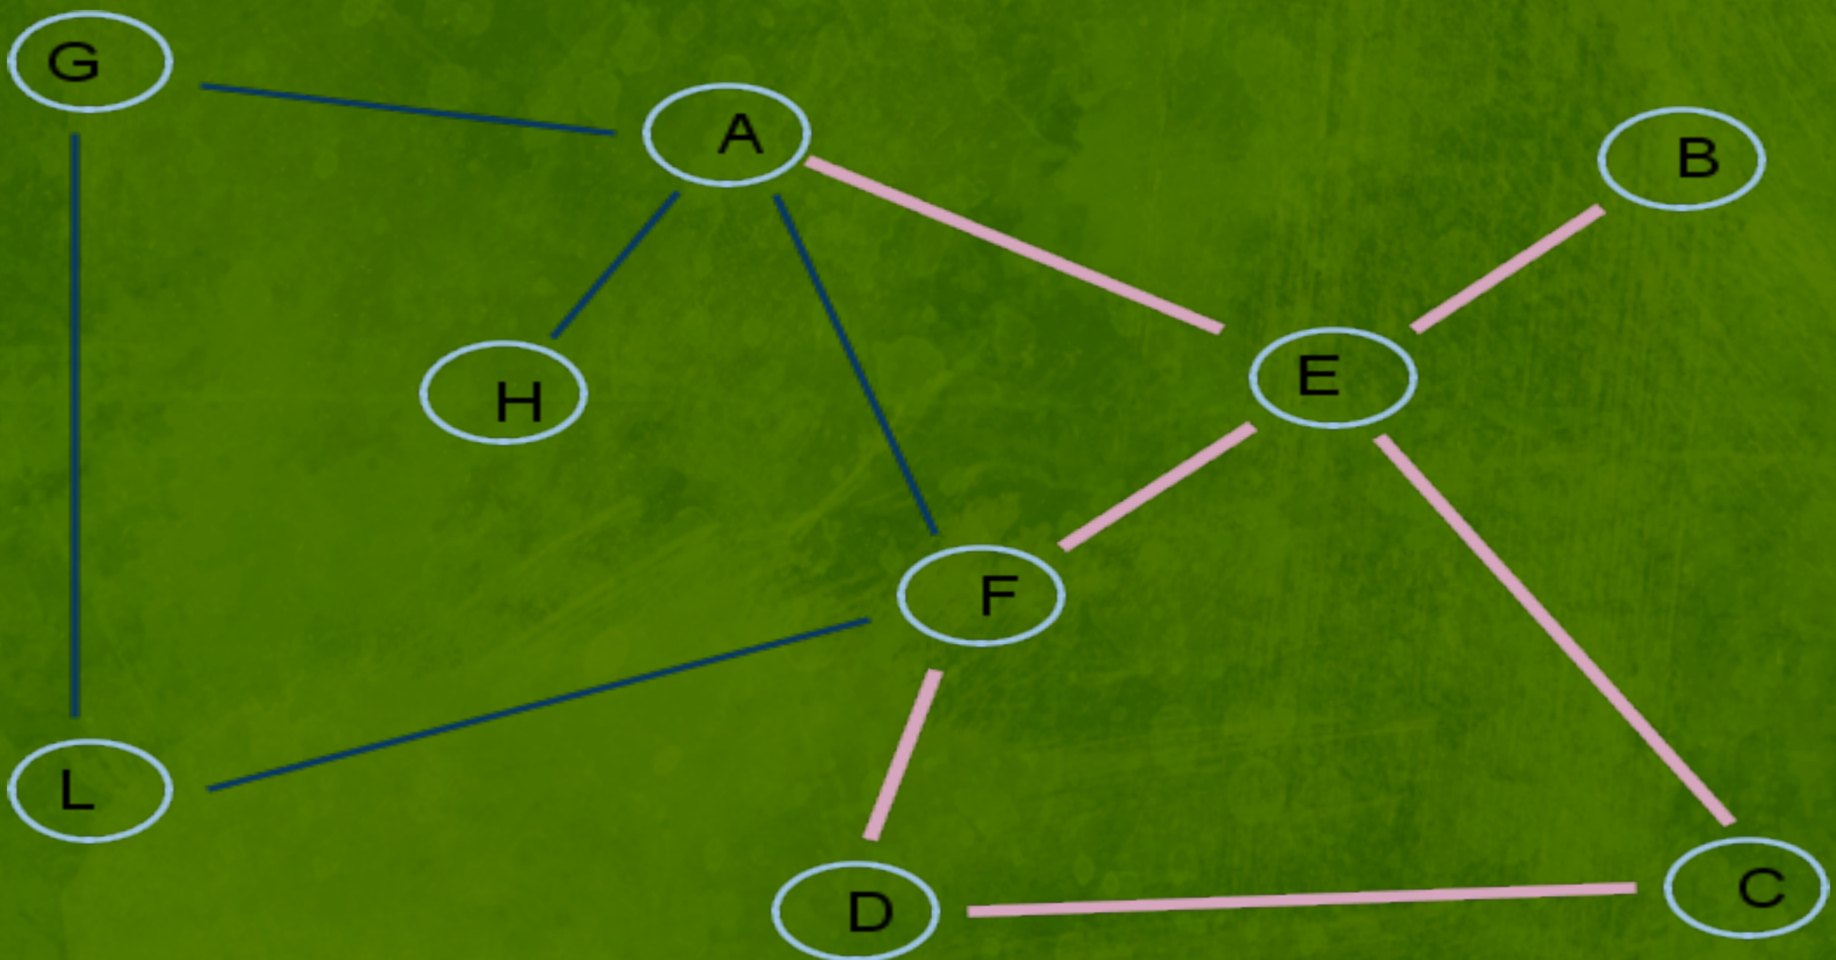

Operational Overview

Stability of the network

At a given point in time, node A has this knowledge:

- Existence of all other nodes
- For each node D of the network, a subset of all the paths potentially going from A to D
    - for each neighbour B, just the best path from B to D

Operational Overview

Recall the network graph

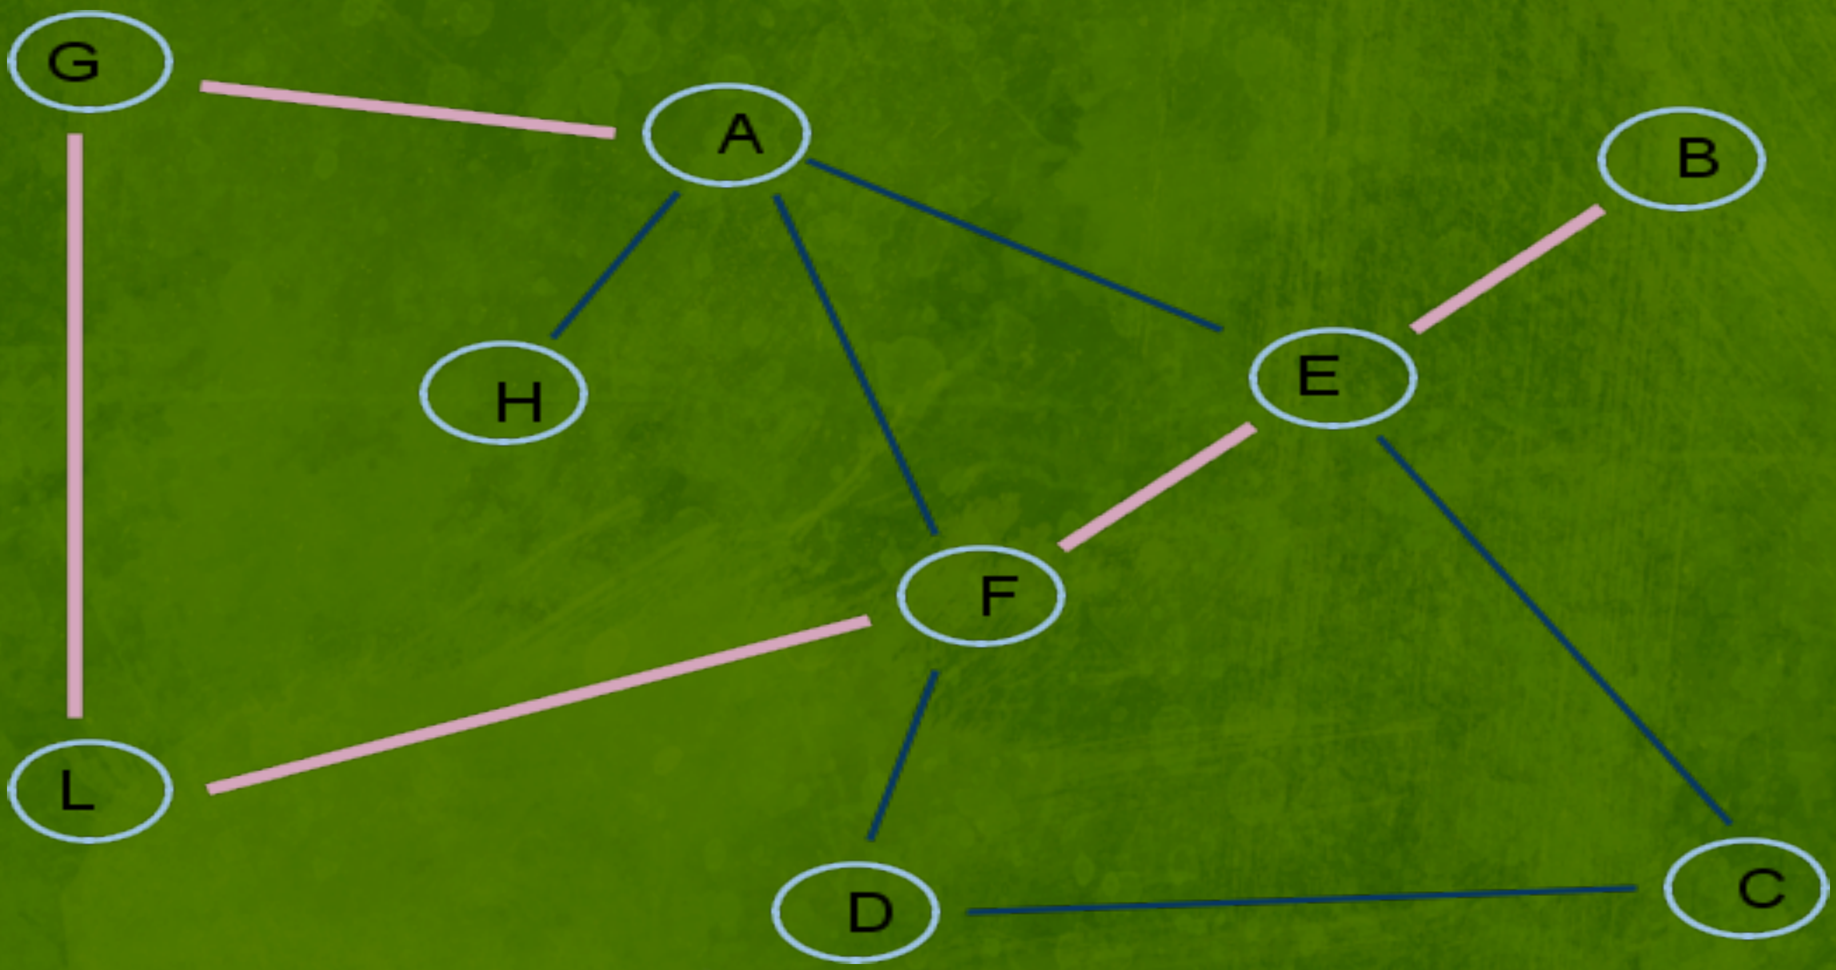

Operational Overview

Recall the network graph

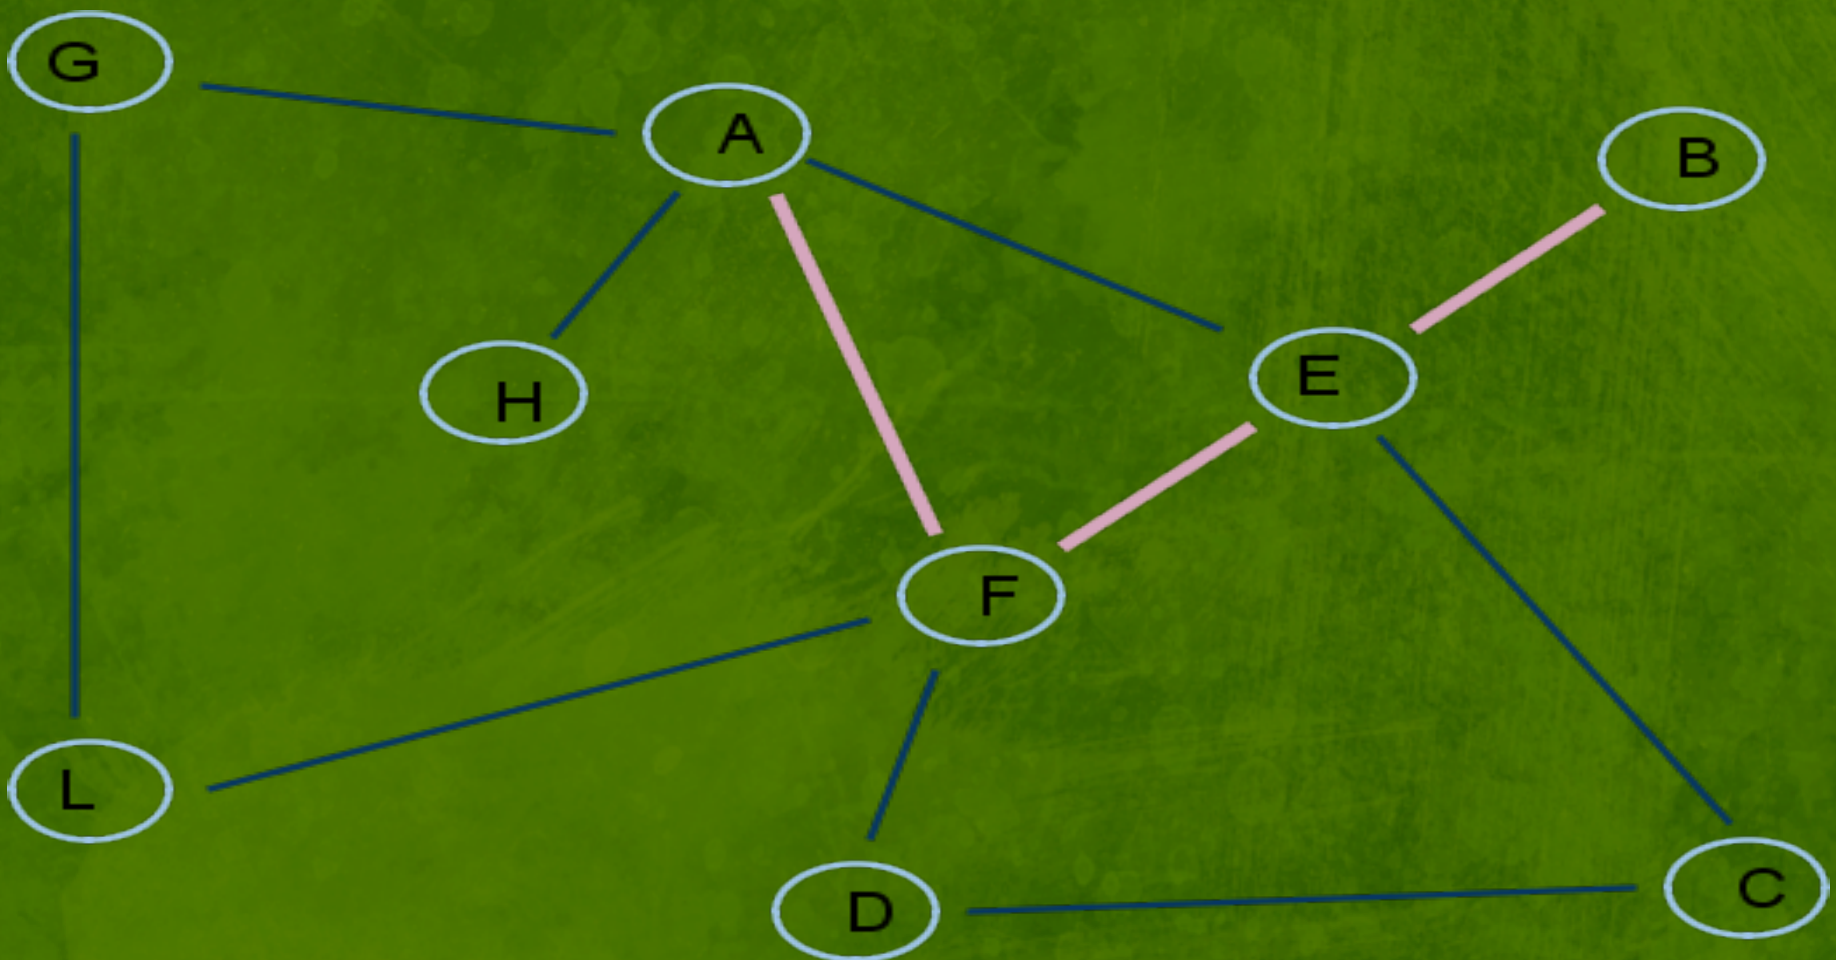

Operational Overview

Recall the network graph

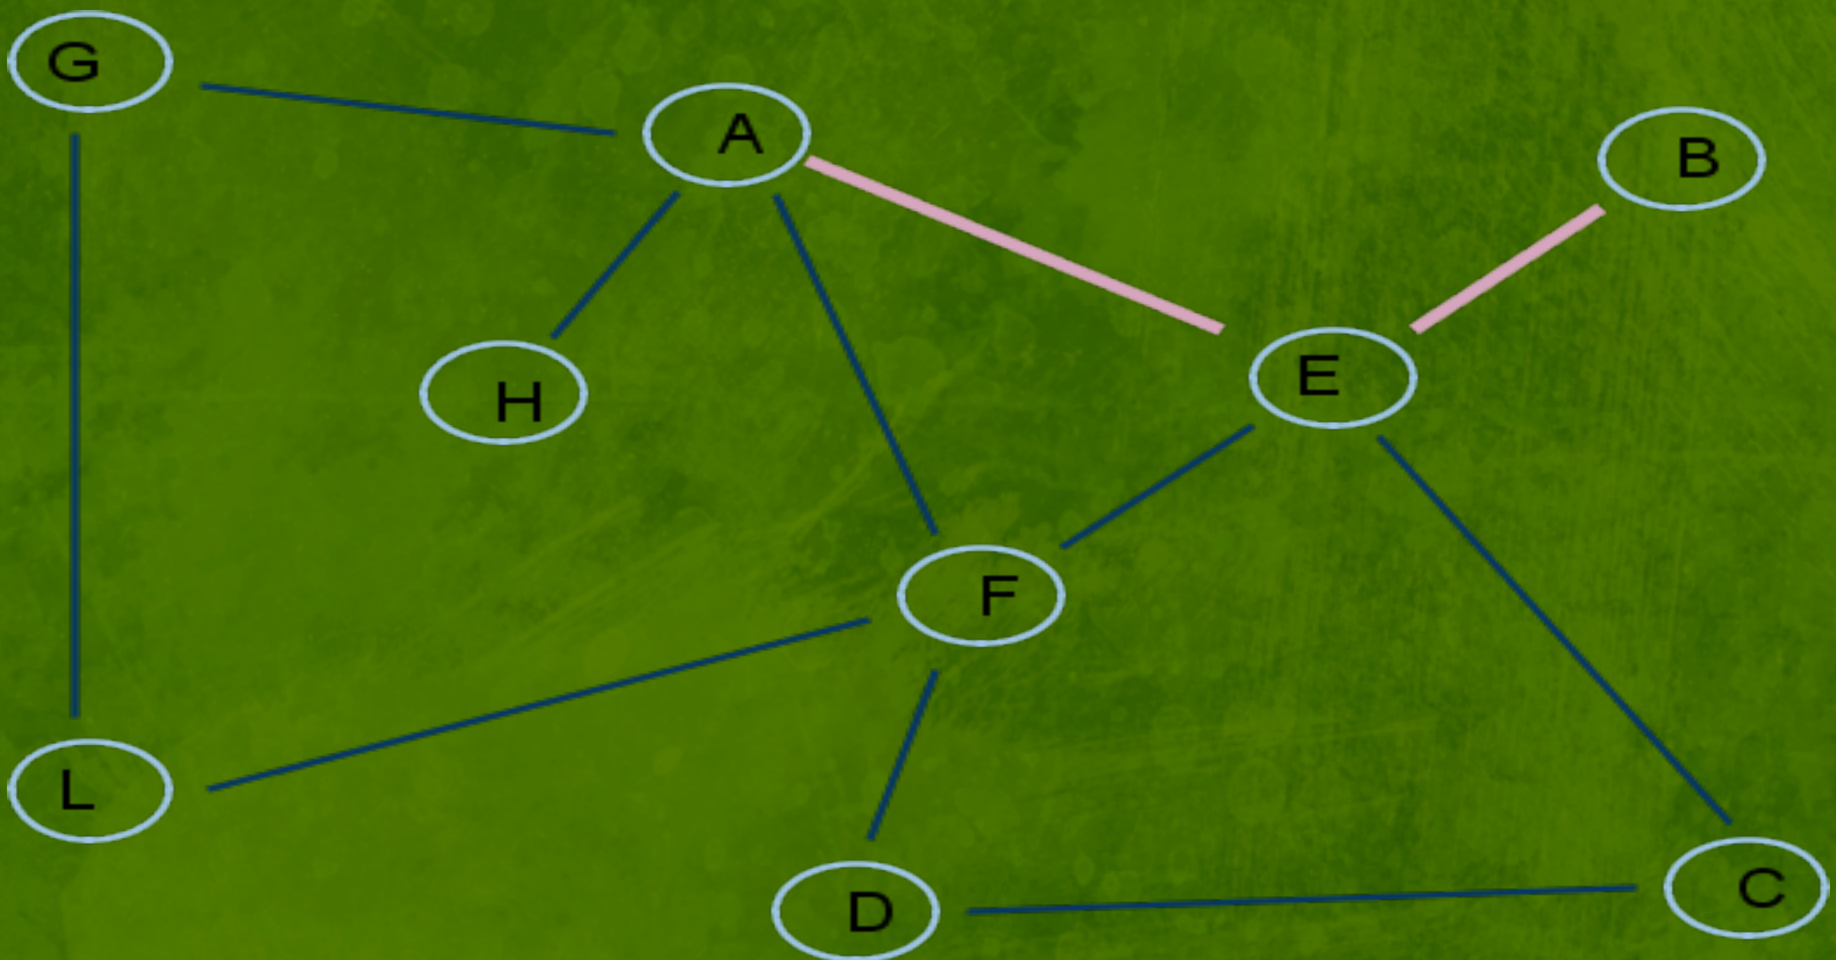

Operational Overview

Stability of the network

At a given point in time, node A has this knowledge:

- Existence of all other nodes
- For each node D of the network, a subset of all the paths potentially going from A to D
  - for each neighbour B, just the best path from B to D
  - no paths that contain a loop

Operational Overview

Another network graph

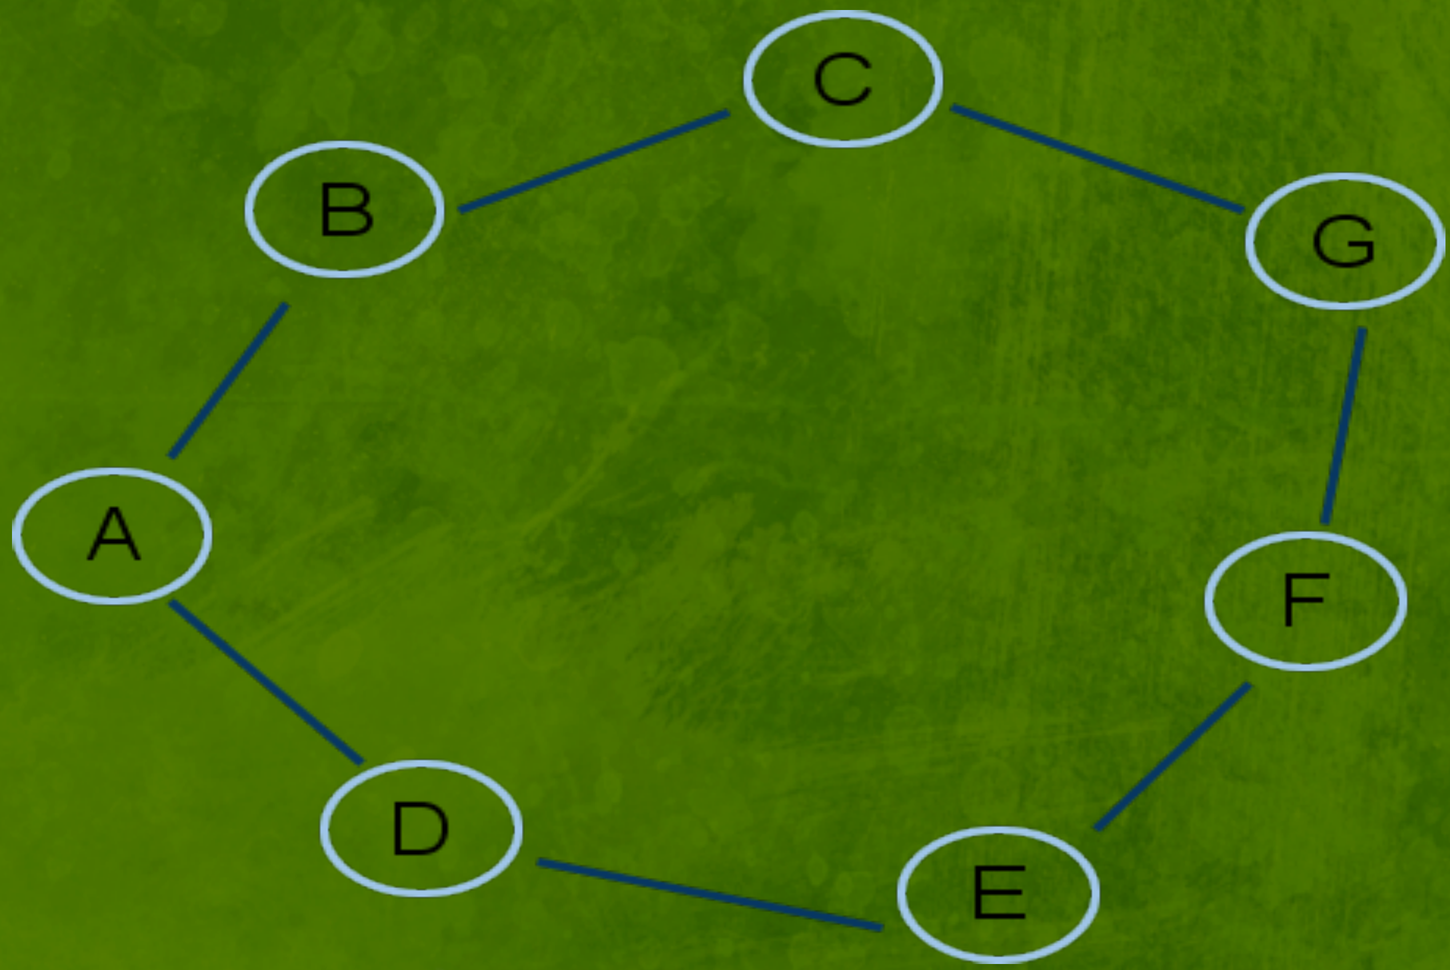

Operational Overview

Recall the network graph

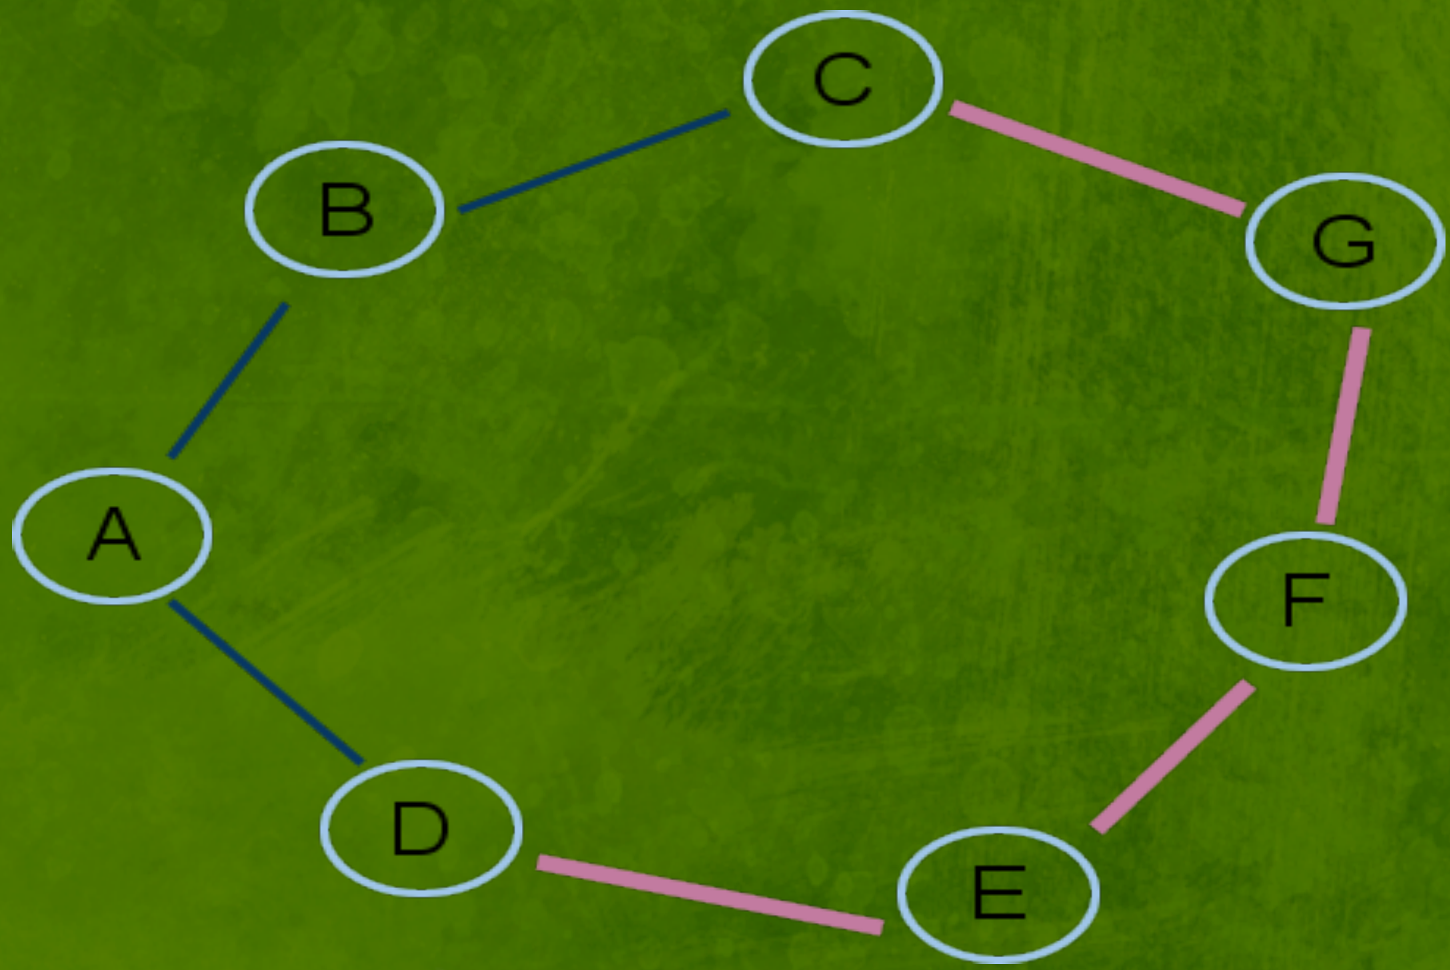

Operational Overview

Recall the network graph

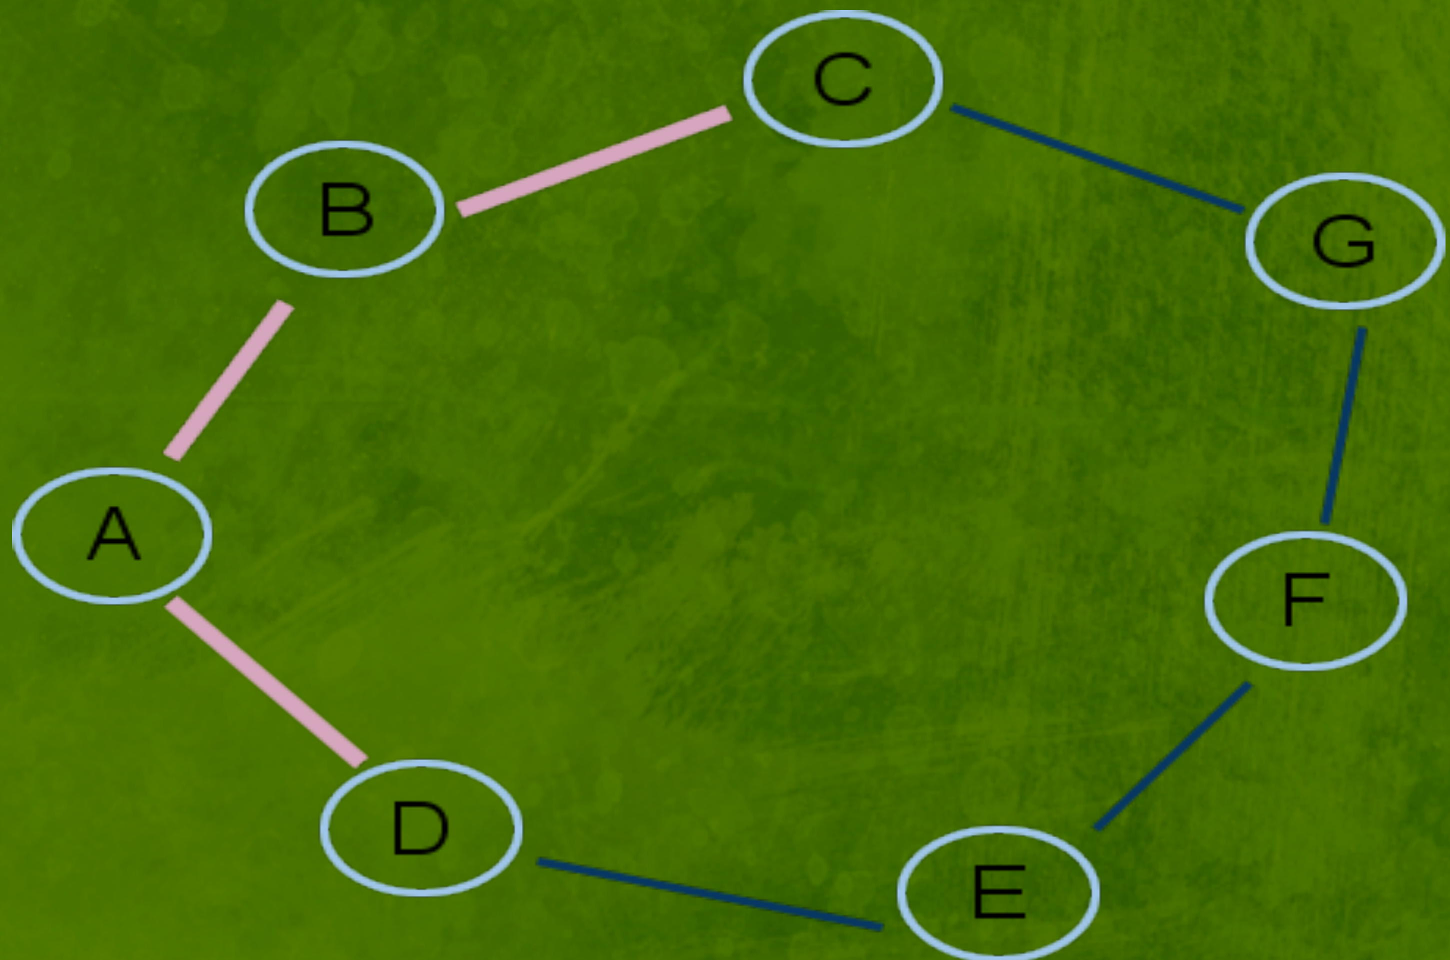

Operational Overview

Recall the network graph

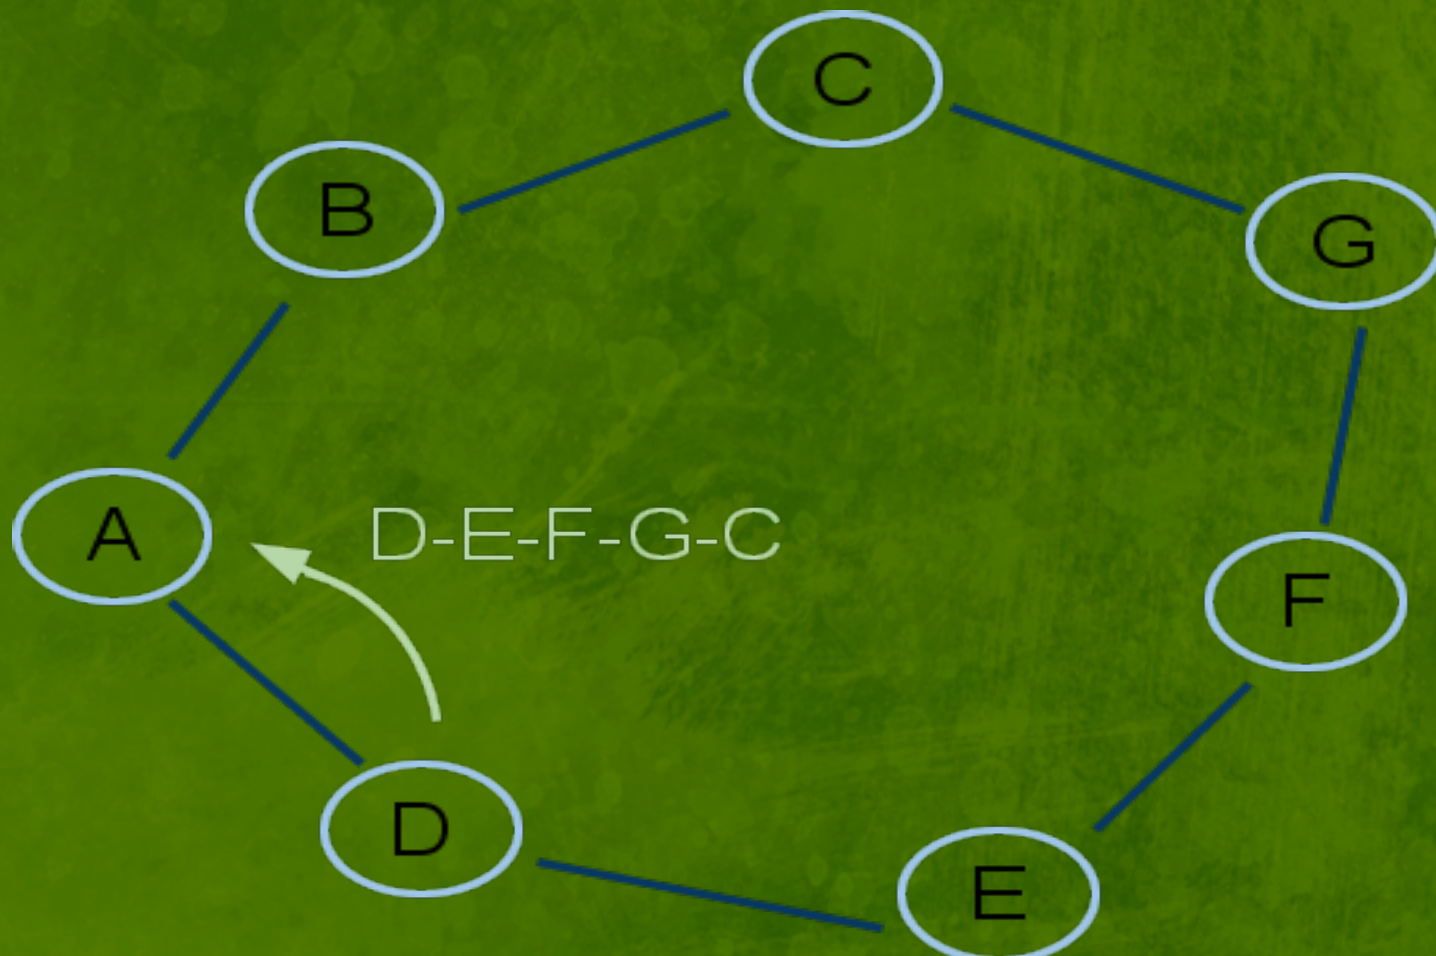

Operational Overview

Recall the network graph

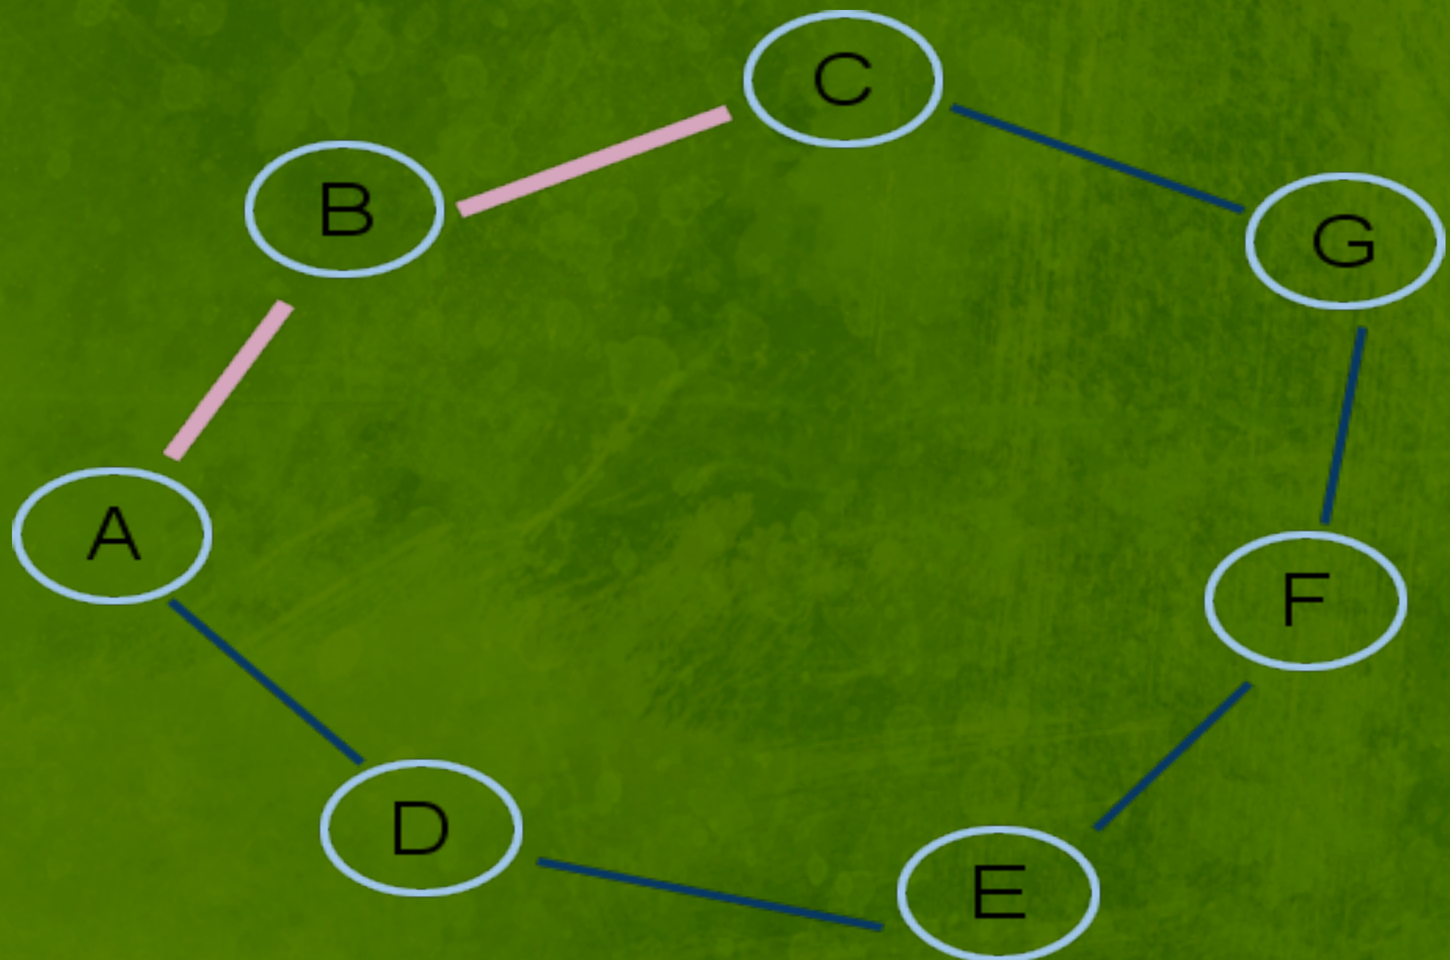

Operational Overview

Recall the network graph

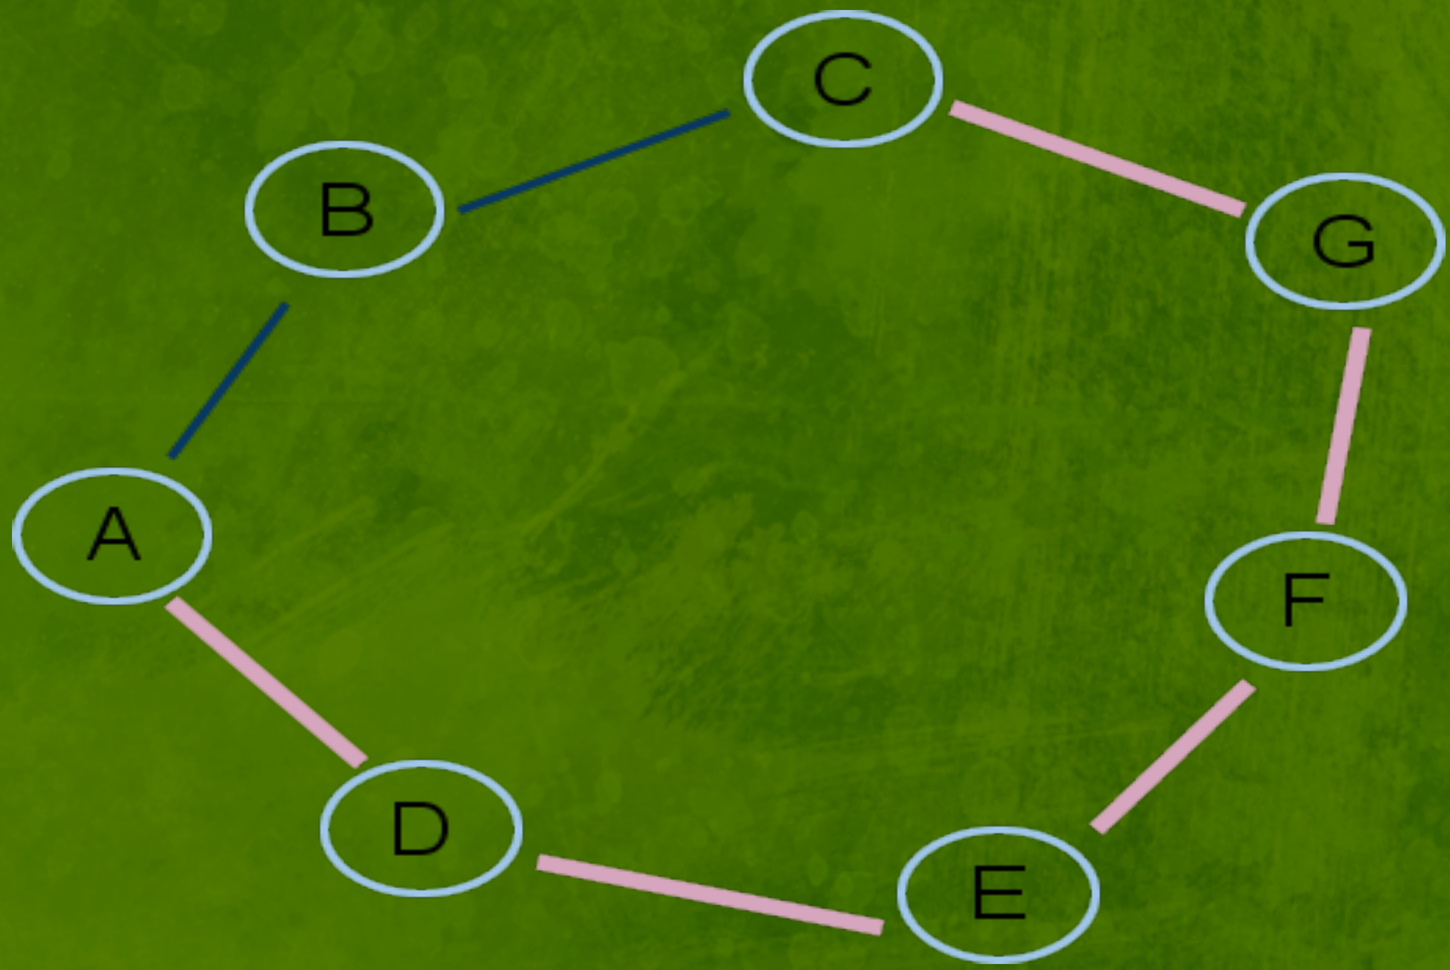

Operational Overview

Recall the network graph

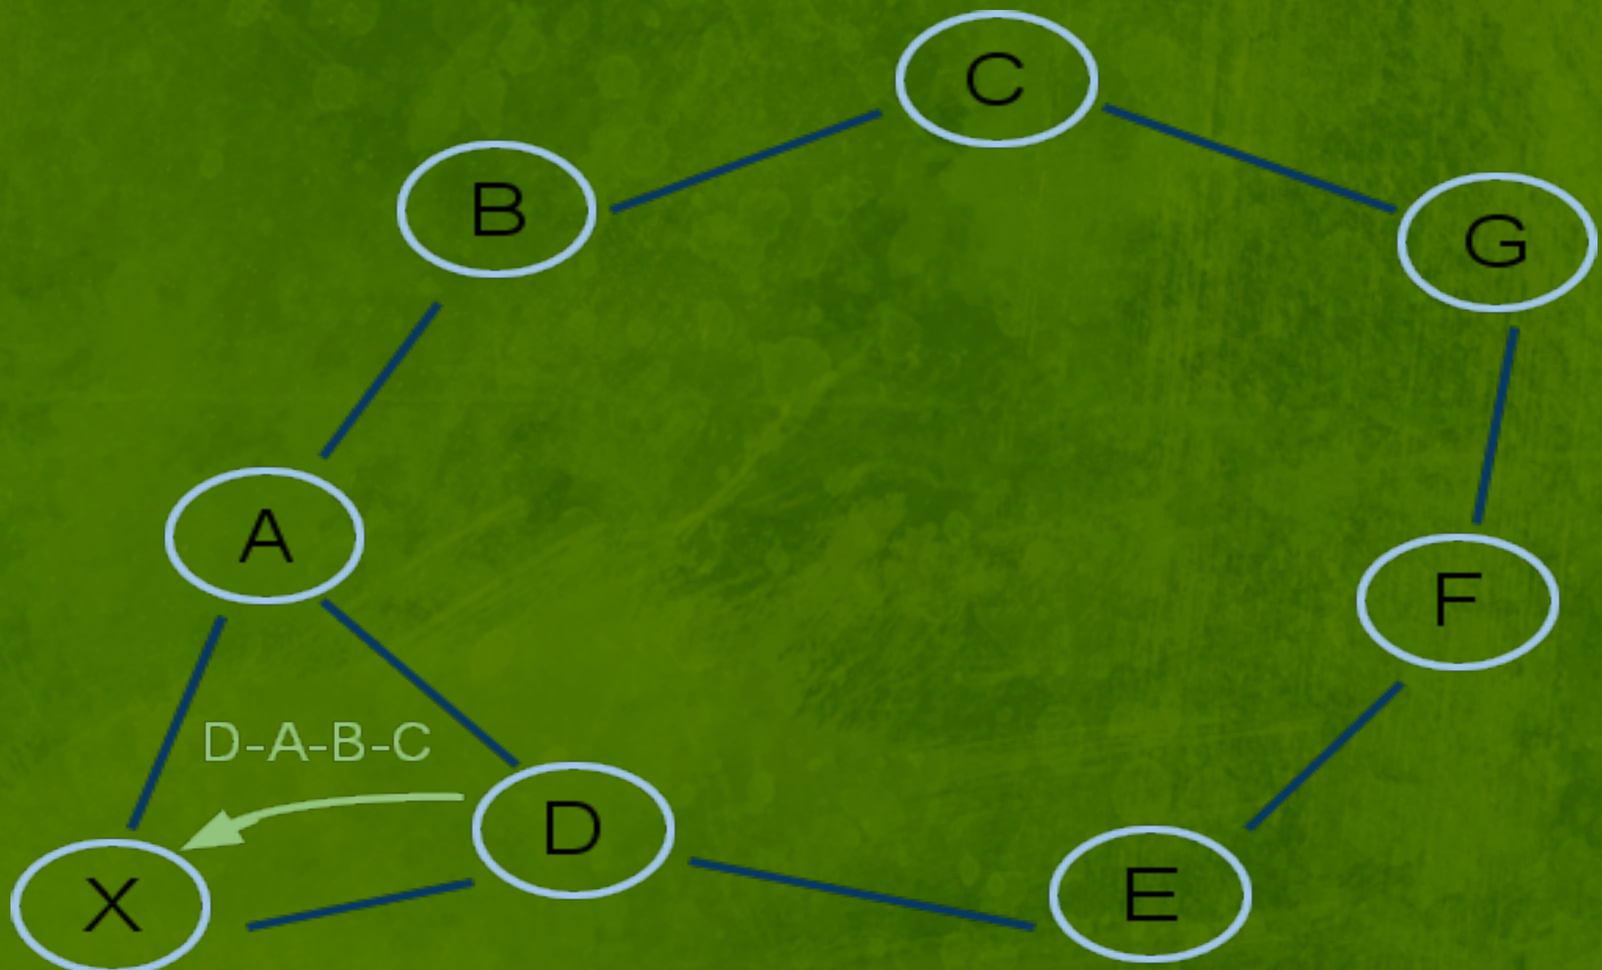

Operational Overview

Recall the network graph

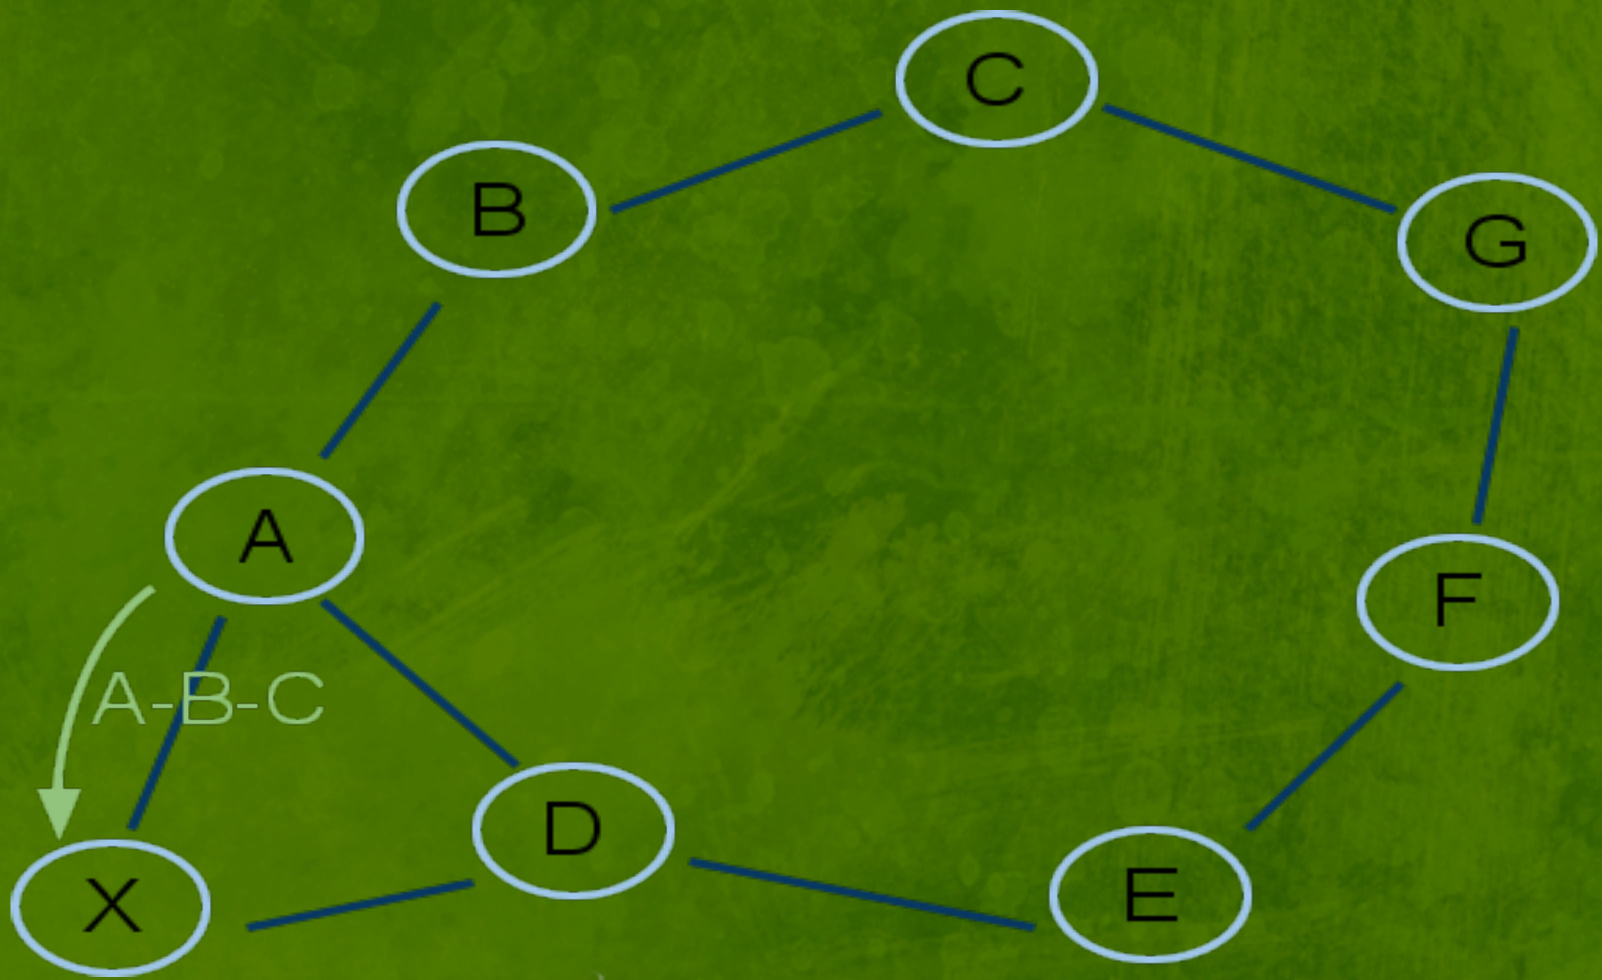

Operational Overview

Recall the network graph

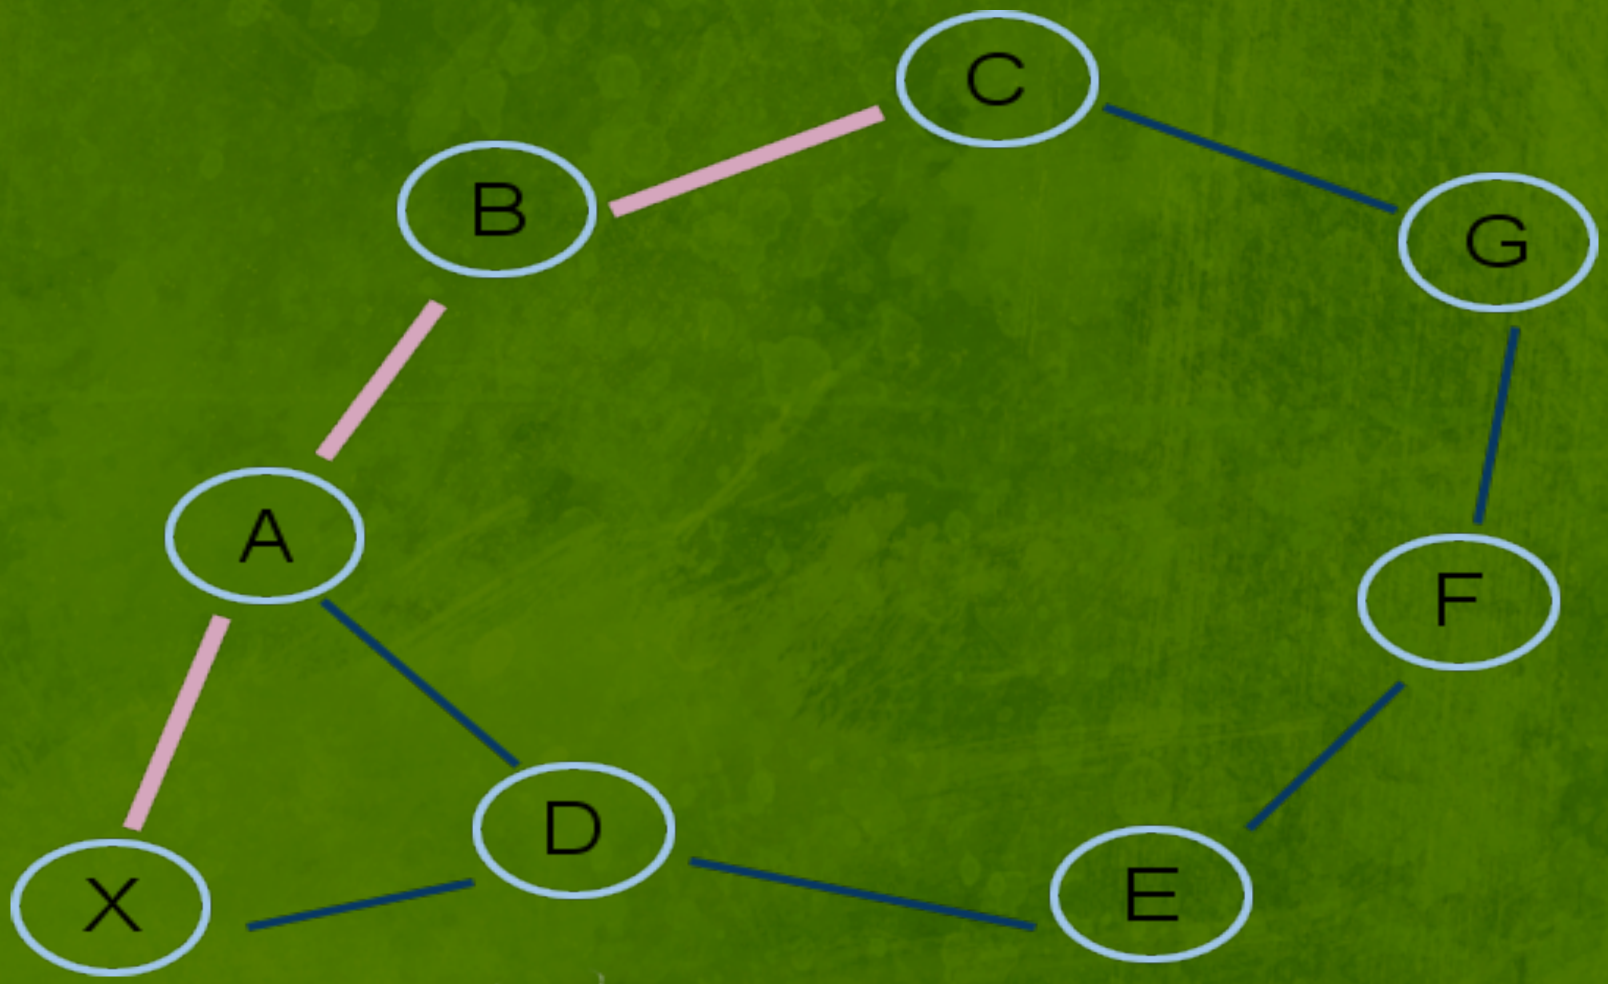

Operational Overview

Recall the network graph

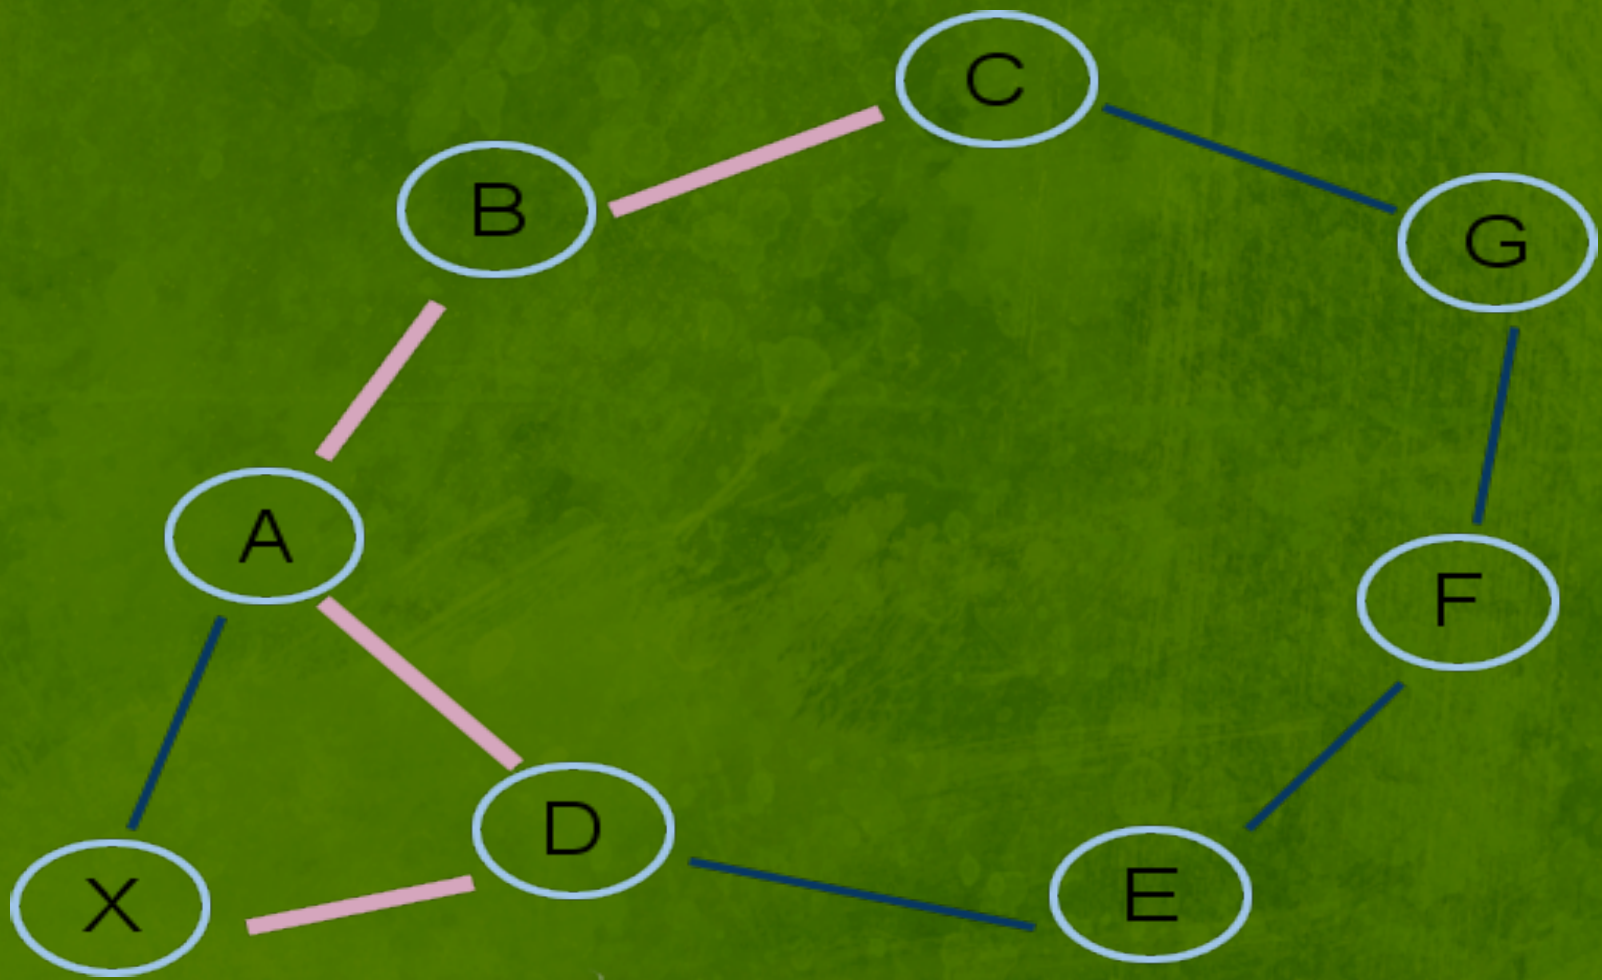

Operational Overview

Recall the network graph

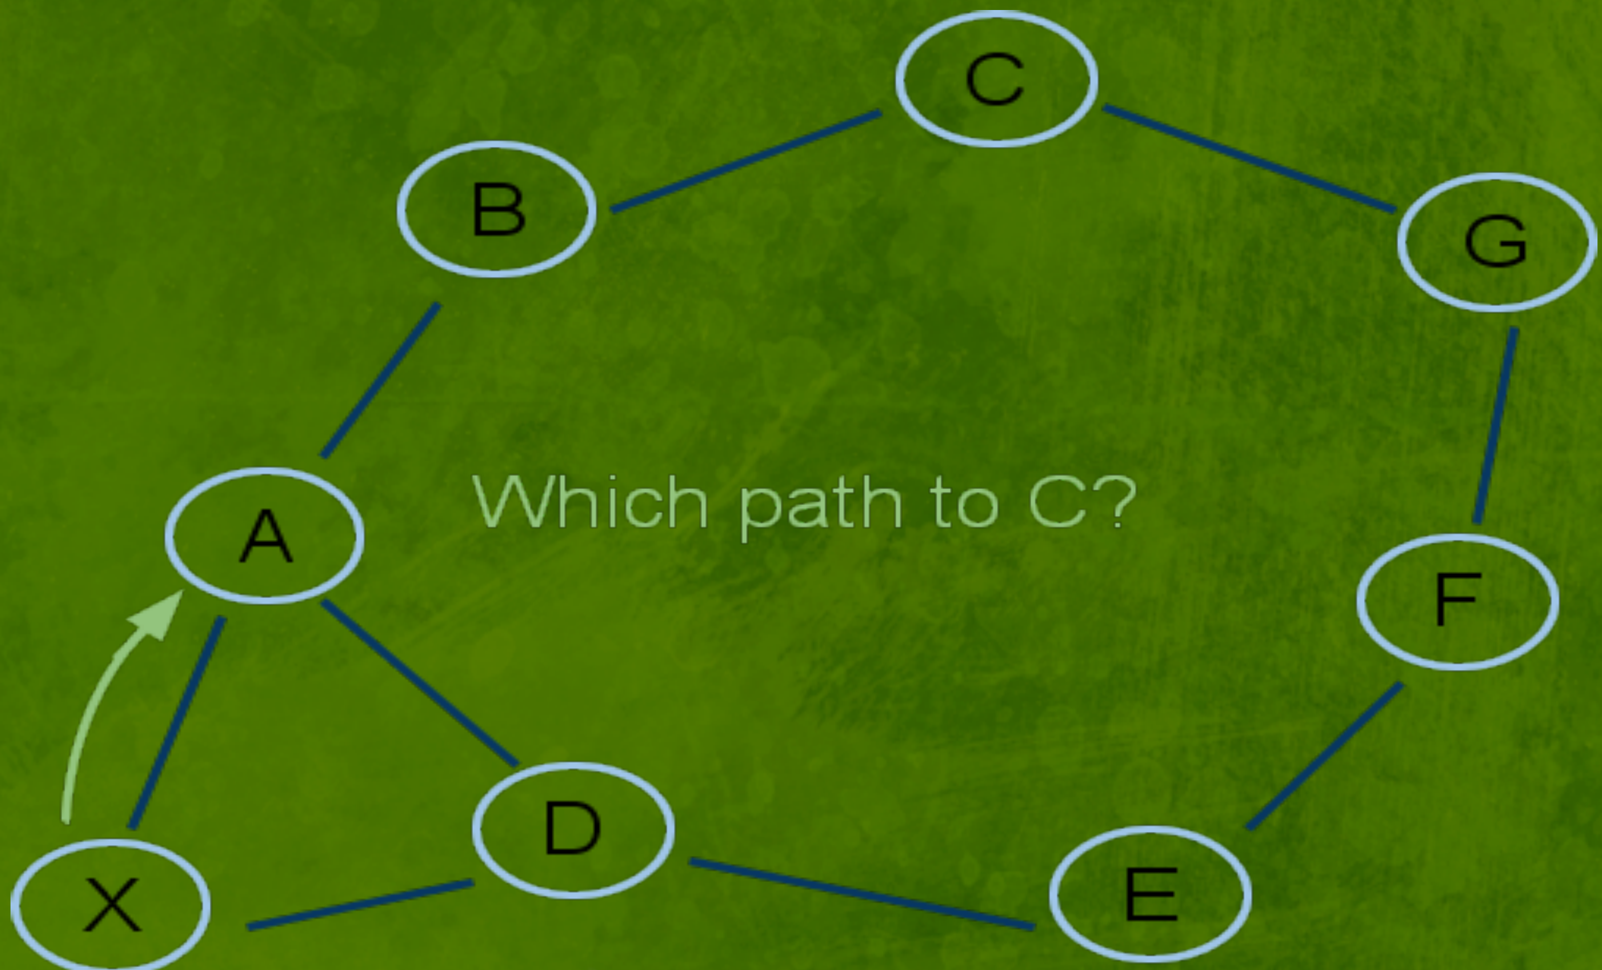

Operational Overview

Recall the network graph

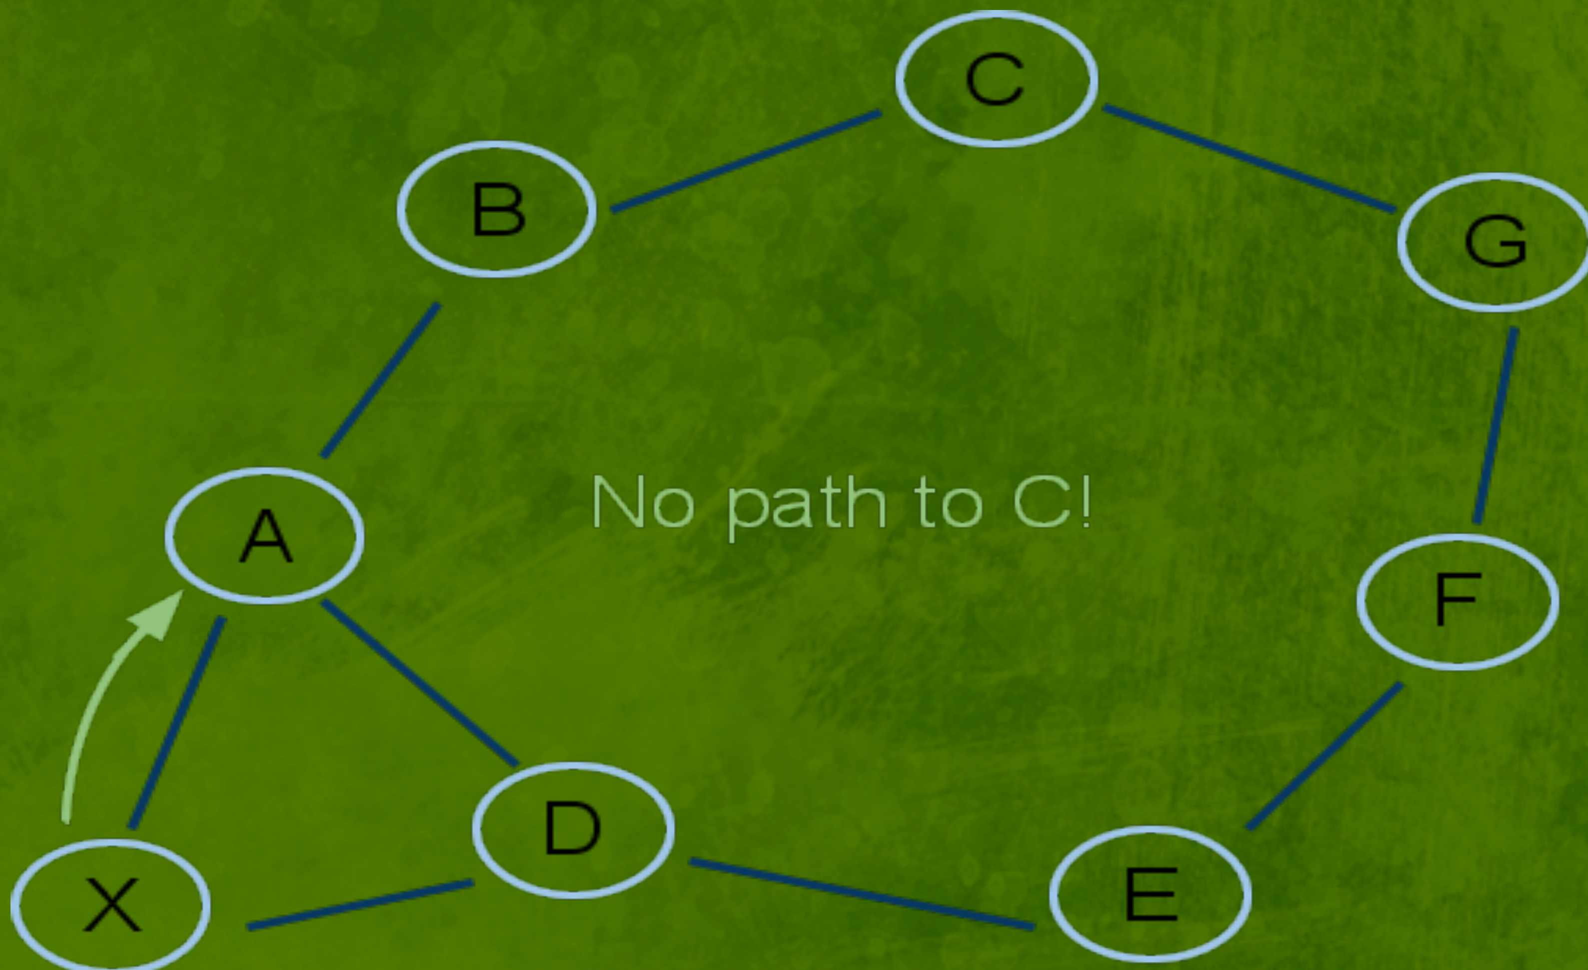

Operational Overview

Stability of the network

At a given point in time, node A has this knowledge:

- Existence of all other nodes
- For each node D of the network, a subset of all the paths potentially going from A to D
  - for each neighbour B, just the best path from B to D
  - no paths that contain a loop
- For each path, the REM and all the hops

Operational Overview

From stability to new stability

From a state of stability, these events may lead to inconsistent state:

- new node
- new link
- changed rem of a link
- dead link
- dead node

Operational Overview

From stability to new stability

Any node directly linked notices the change and updates its map.

Operational Overview

From stability to new stability

Any node directly linked notices the change and updates its map.

Then sends to neighbours all the updated paths.

Operational Overview

From stability to new stability

Any node directly linked notices the change and updates its map.

Then sends to neighbours all the updated paths.

A neighbour receives and updates its map, **if it is the case**.

Operational Overview

From stability to new stability

Any node directly linked notices the change and updates its map.

Then sends to neighbours all the updated paths.

A neighbour receives and updates its map, **if it is the case**.

Then sends to neighbours all the updated paths, **if there are**.

Operational Overview

From stability to new stability

Any node directly linked notices the change and updates its map.

Then sends to neighbours all the updated paths.

A neighbour receives and updates its map, **if it is the case**.

Then sends to neighbours all the updated paths, **if there are**.

When propagation finishes, the net reaches new stability.

Operational Overview

From stability to new stability

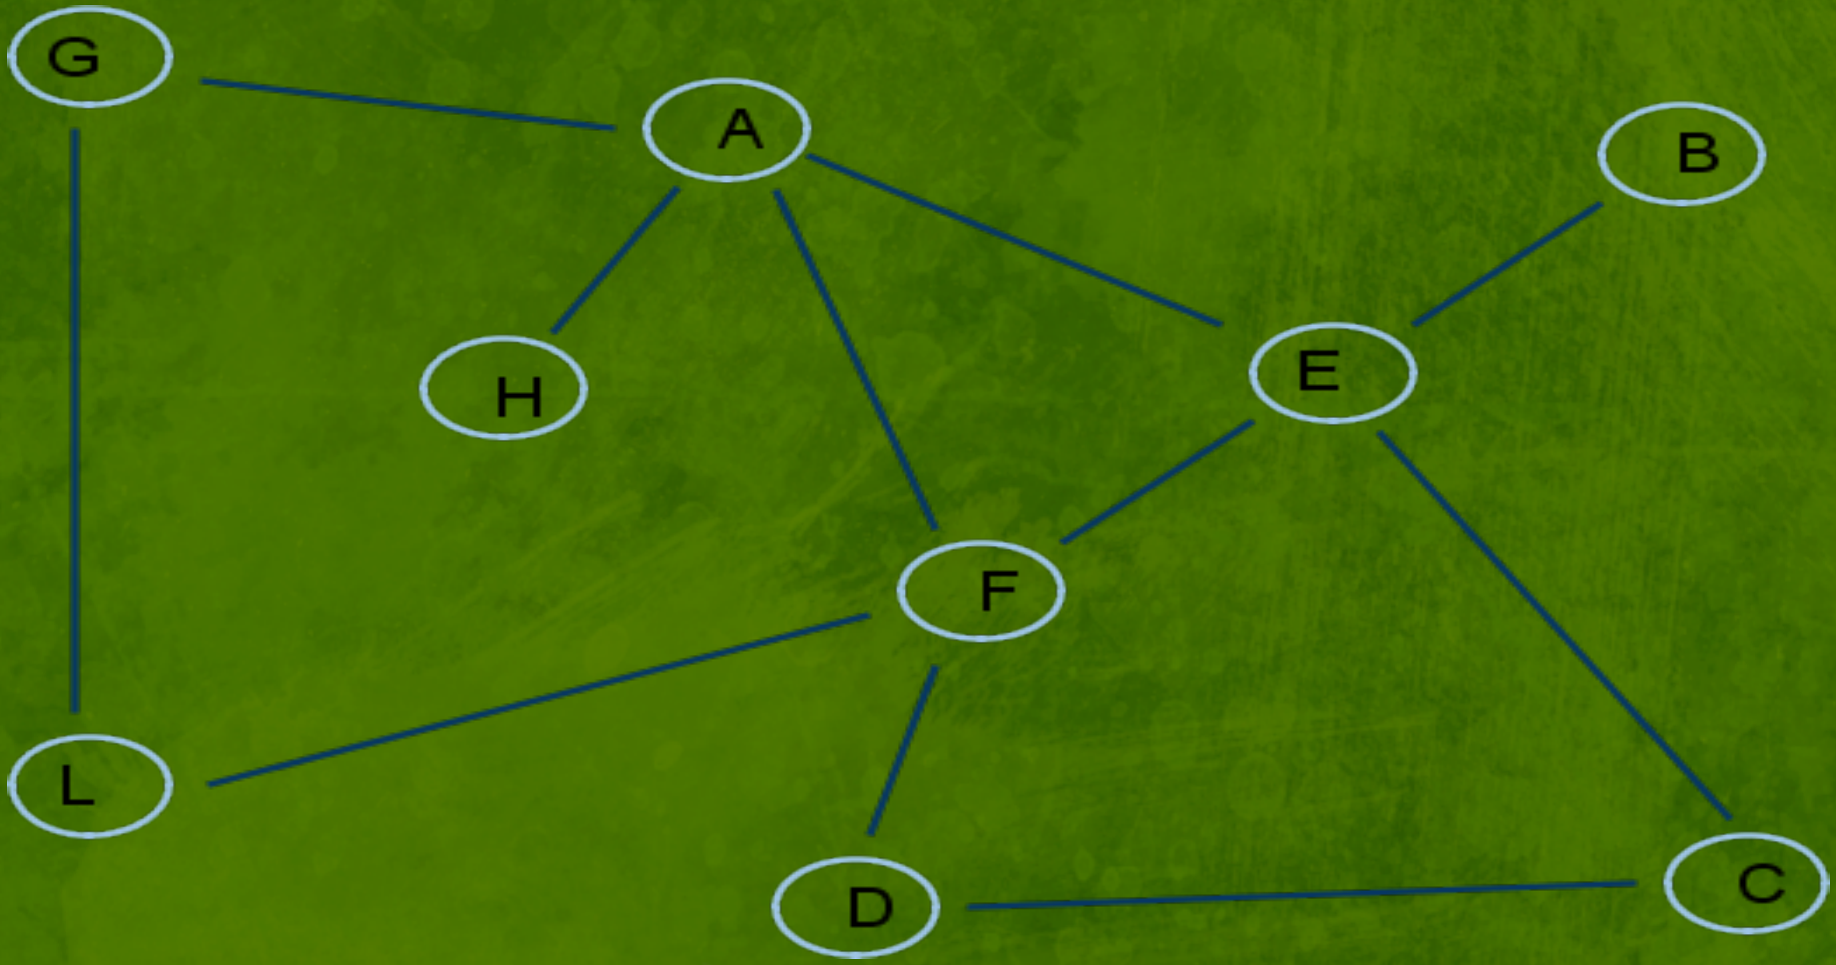

Operational Overview

From stability to new stability

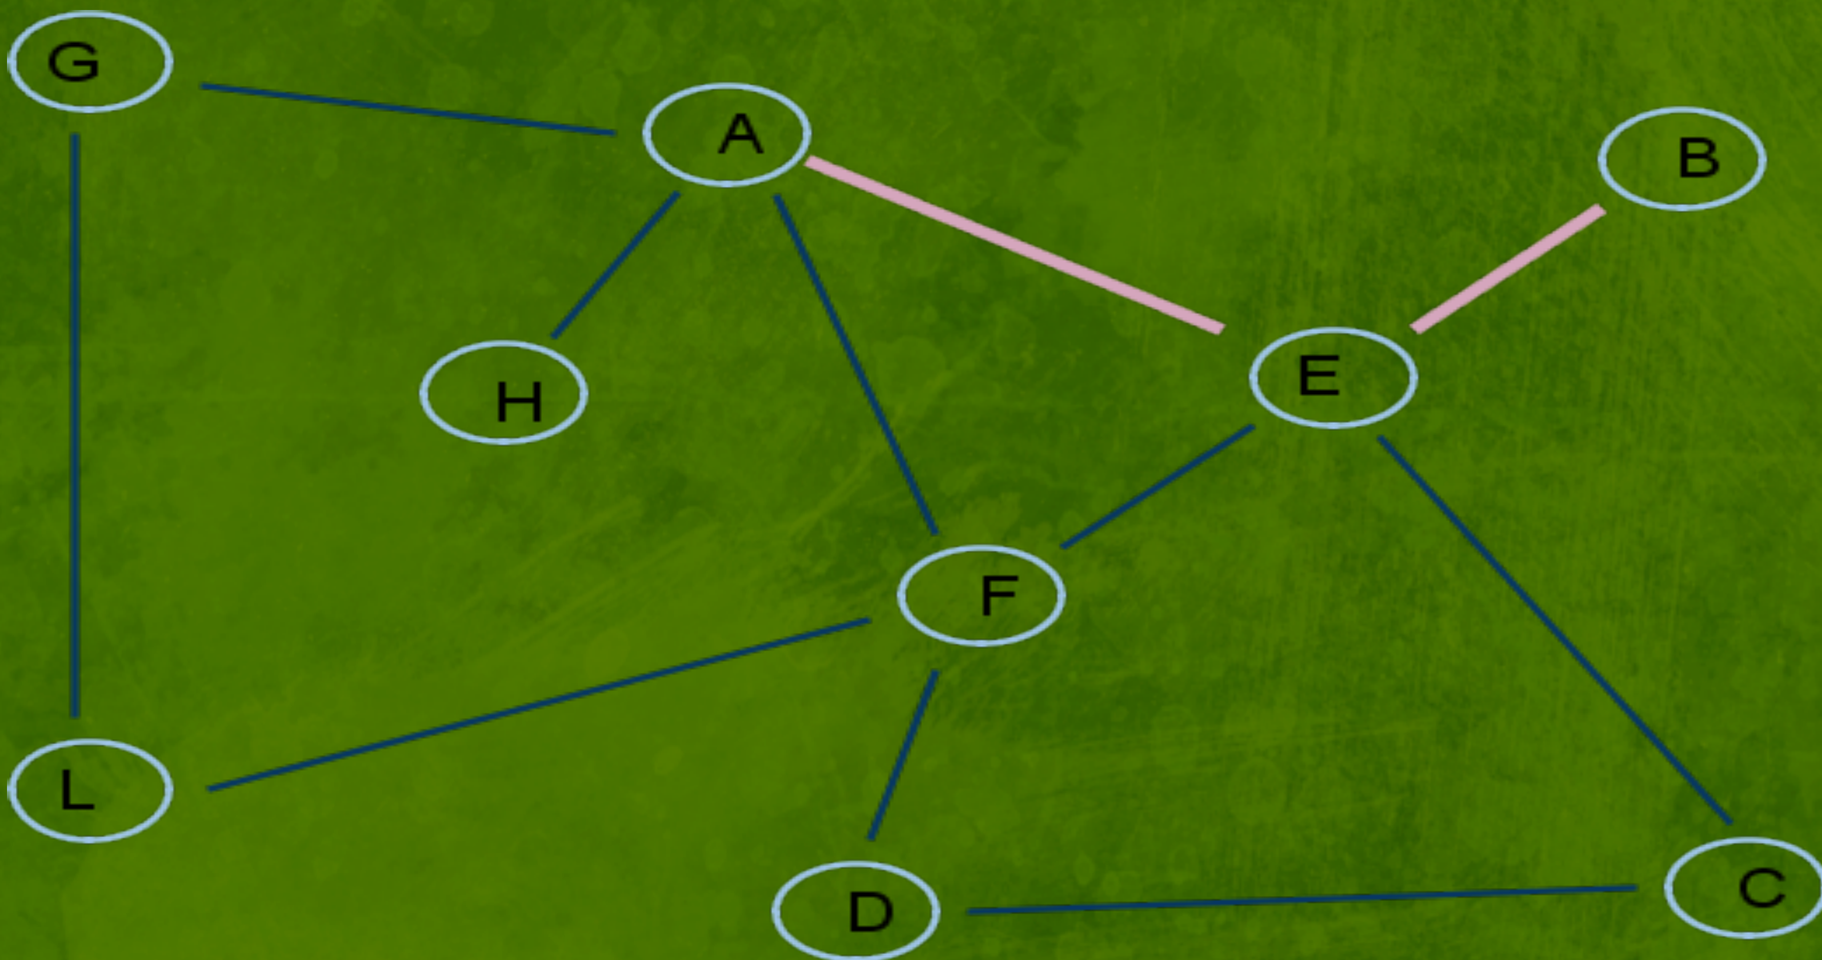

Operational Overview

From stability to new stability

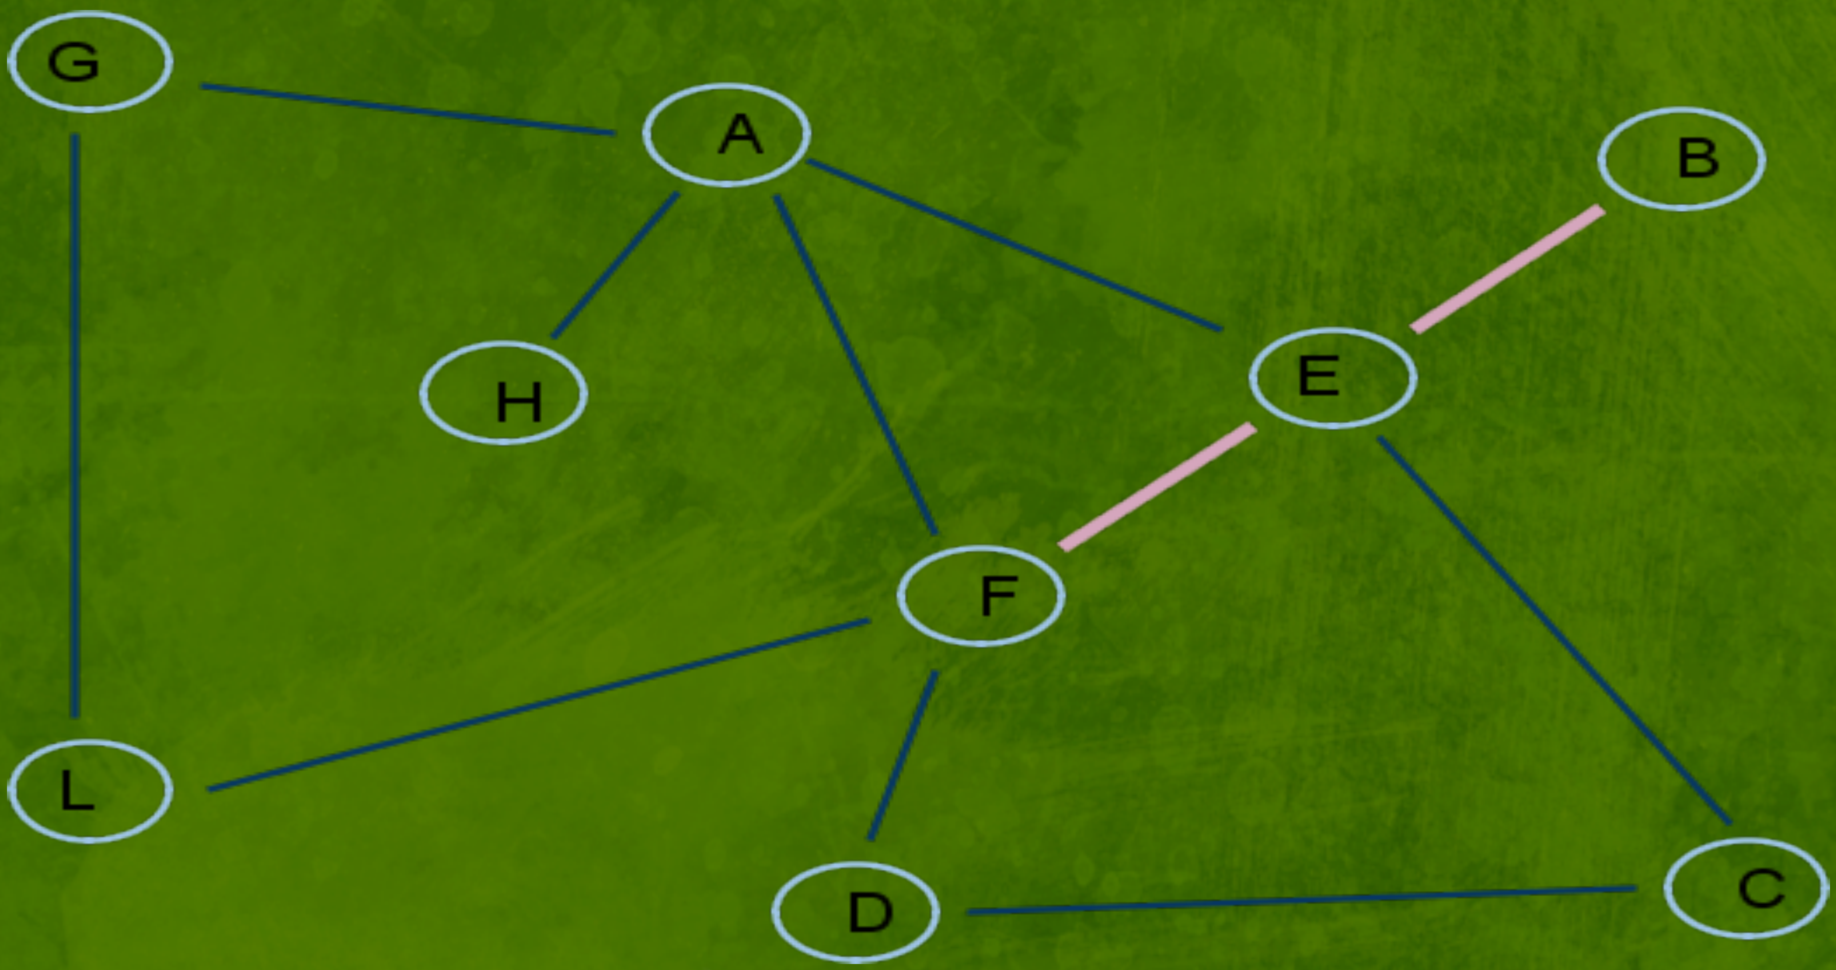

Operational Overview

From stability to new stability

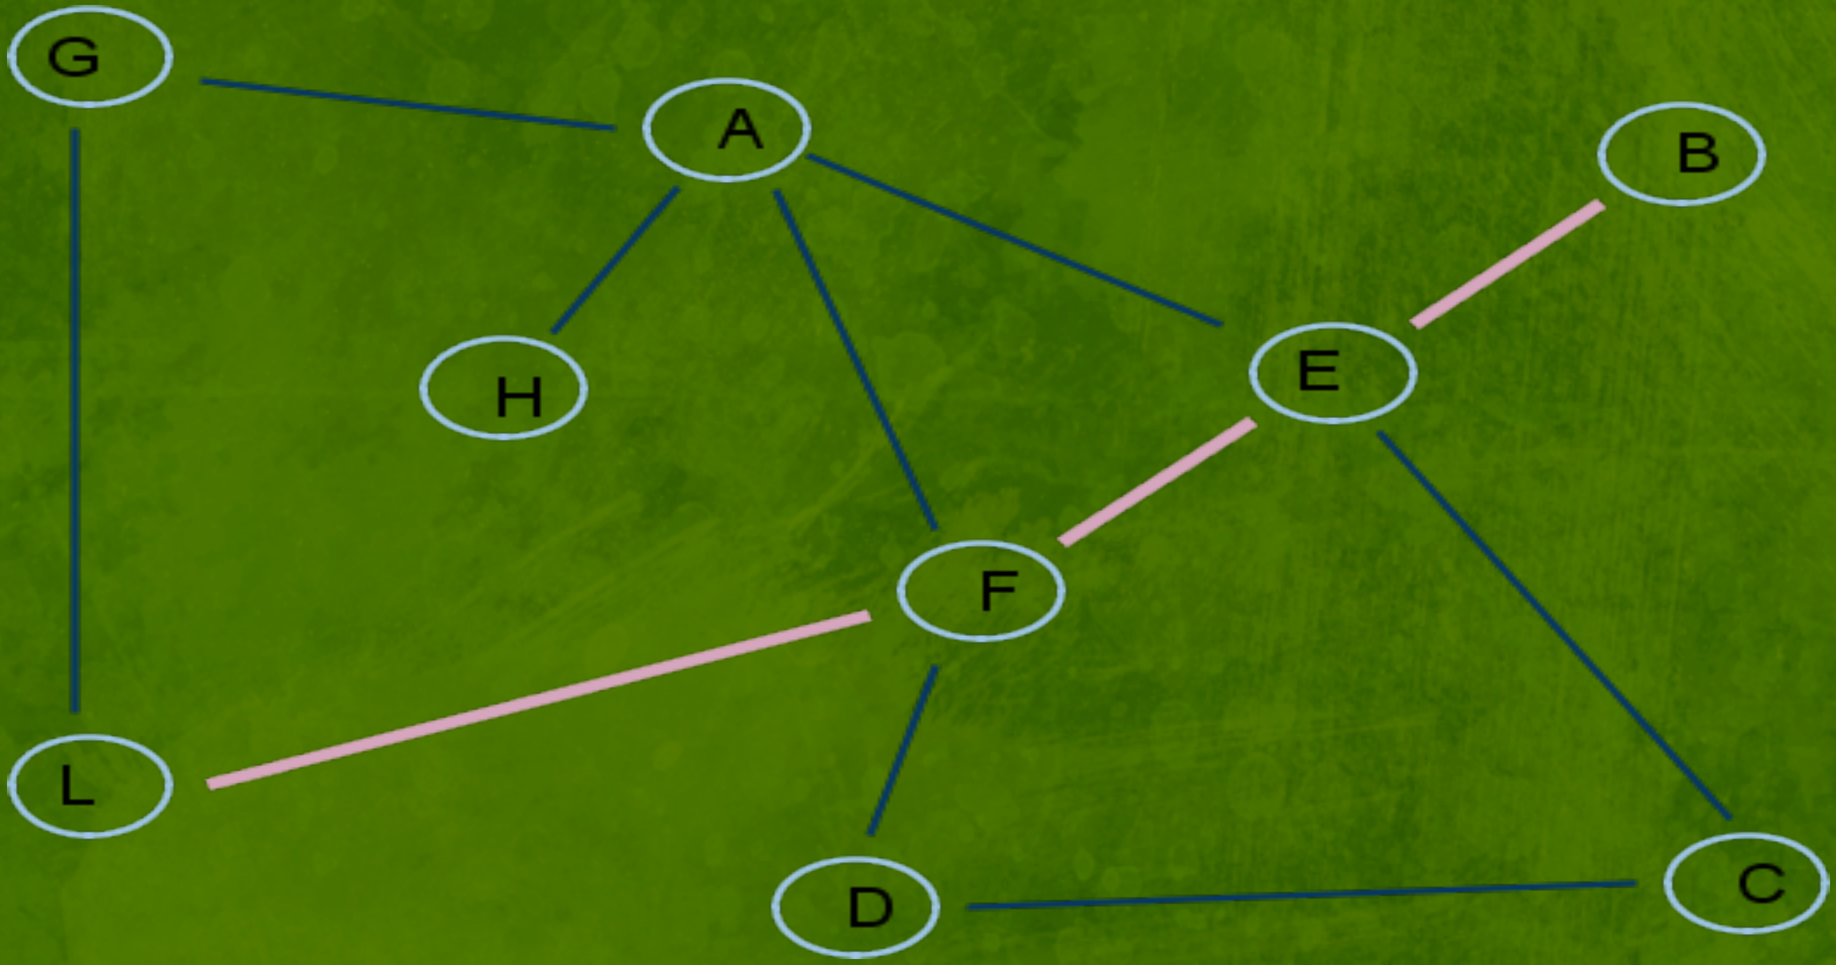

Operational Overview

From stability to new stability

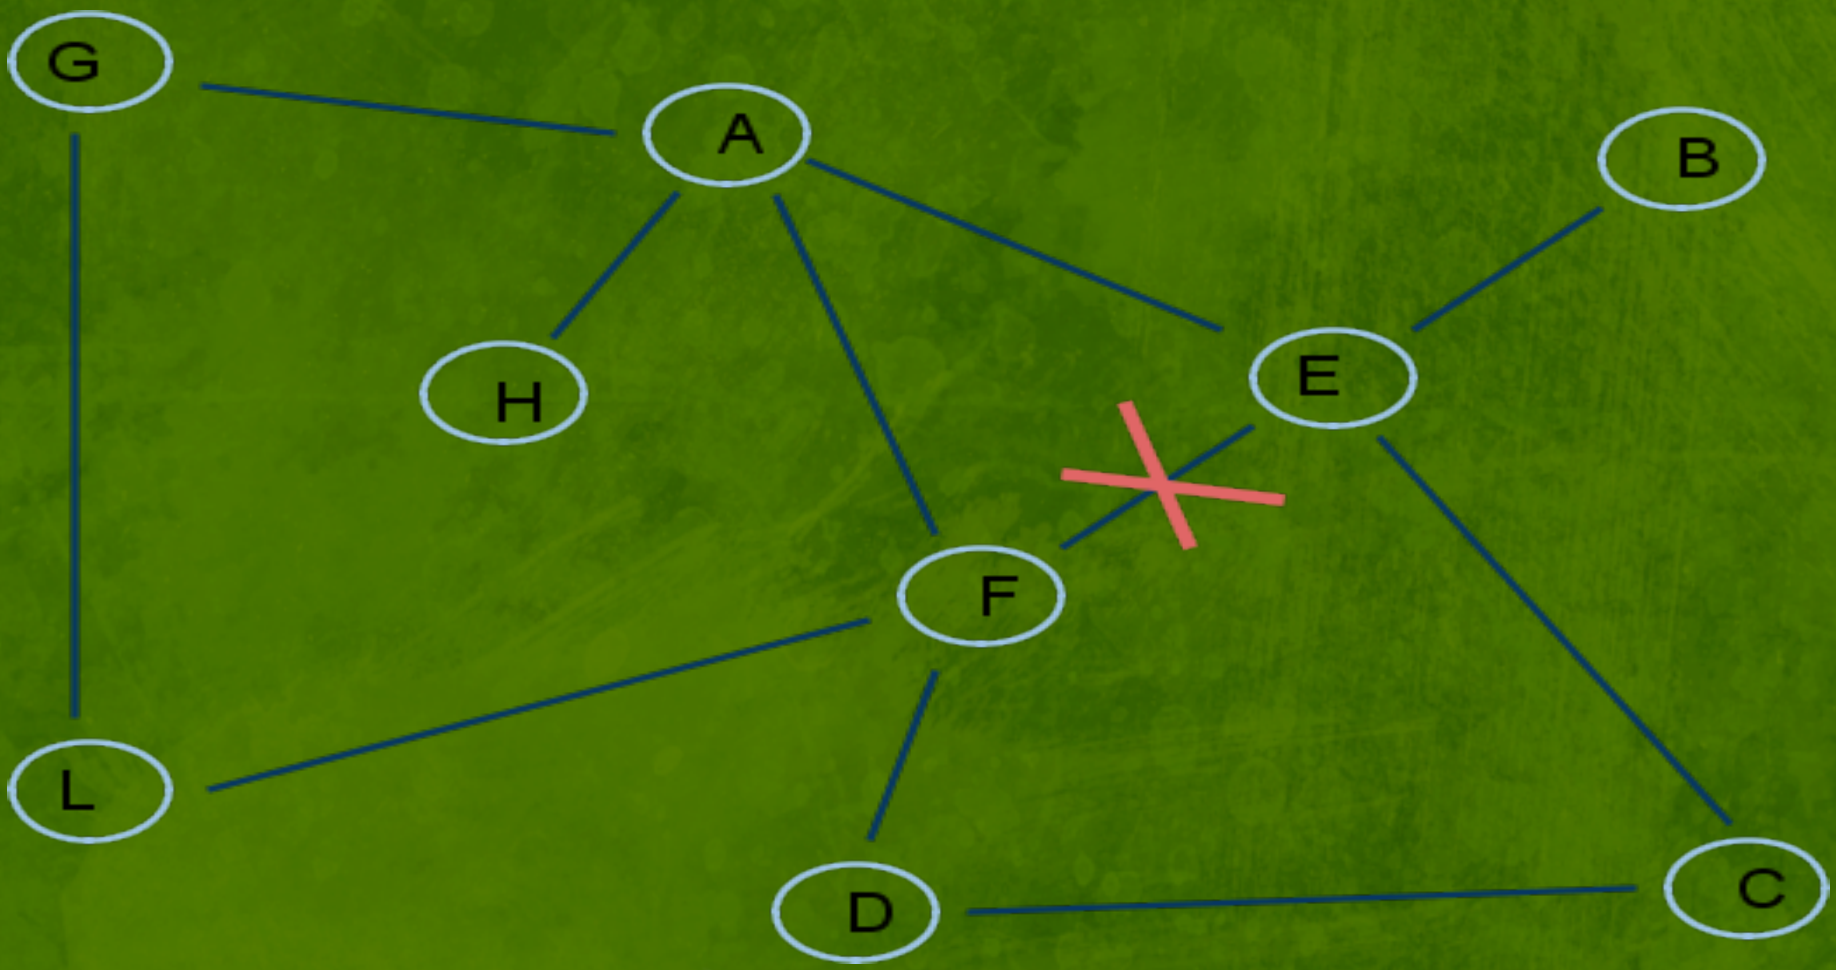

Operational Overview

From stability to new stability

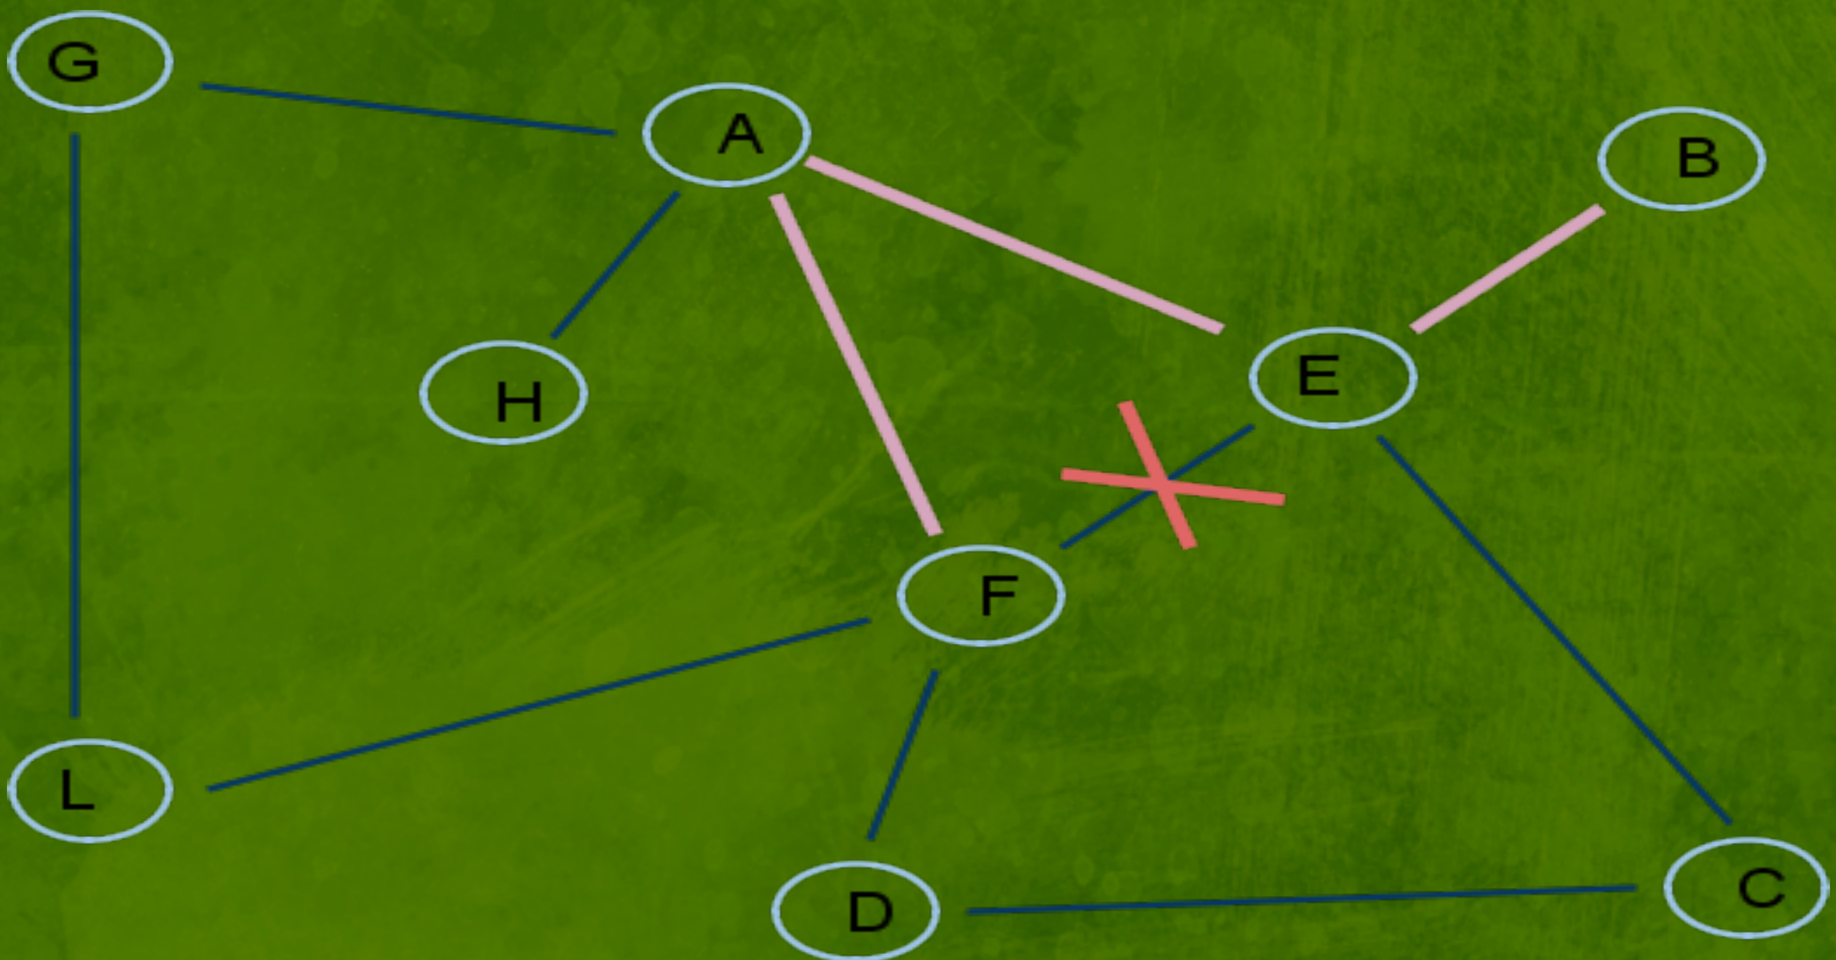

Operational Overview

From stability to new stability

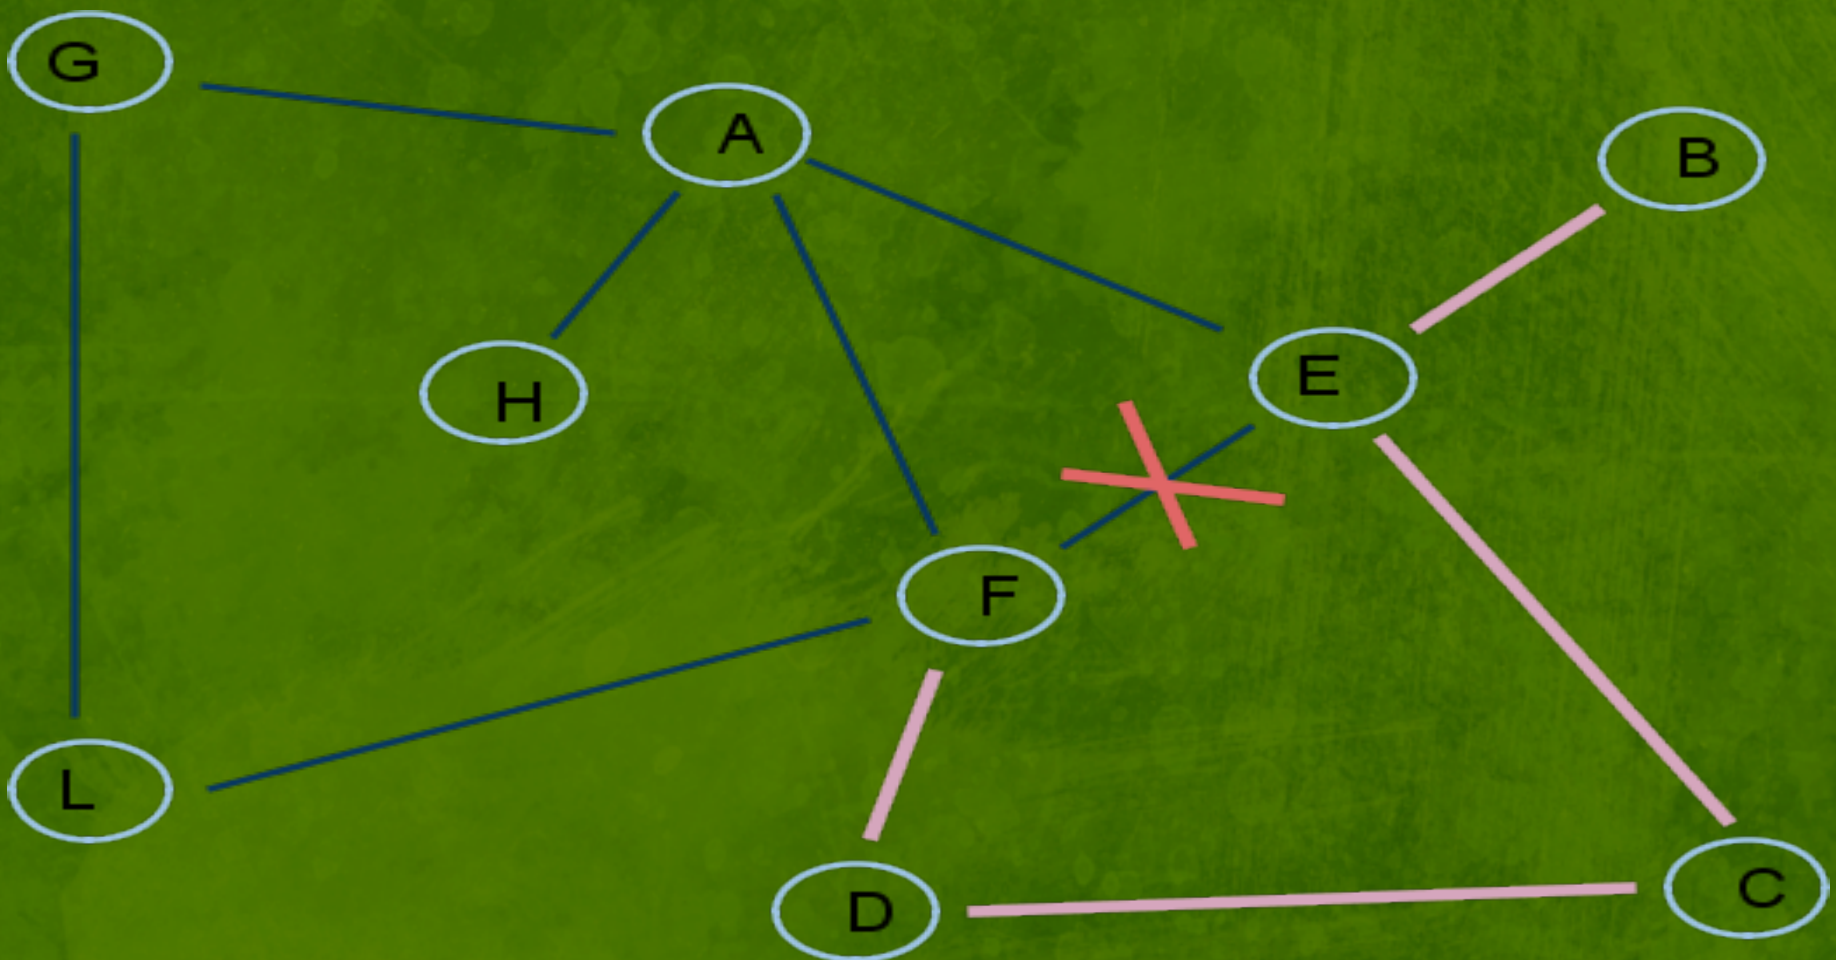

Operational Overview

From stability to new stability

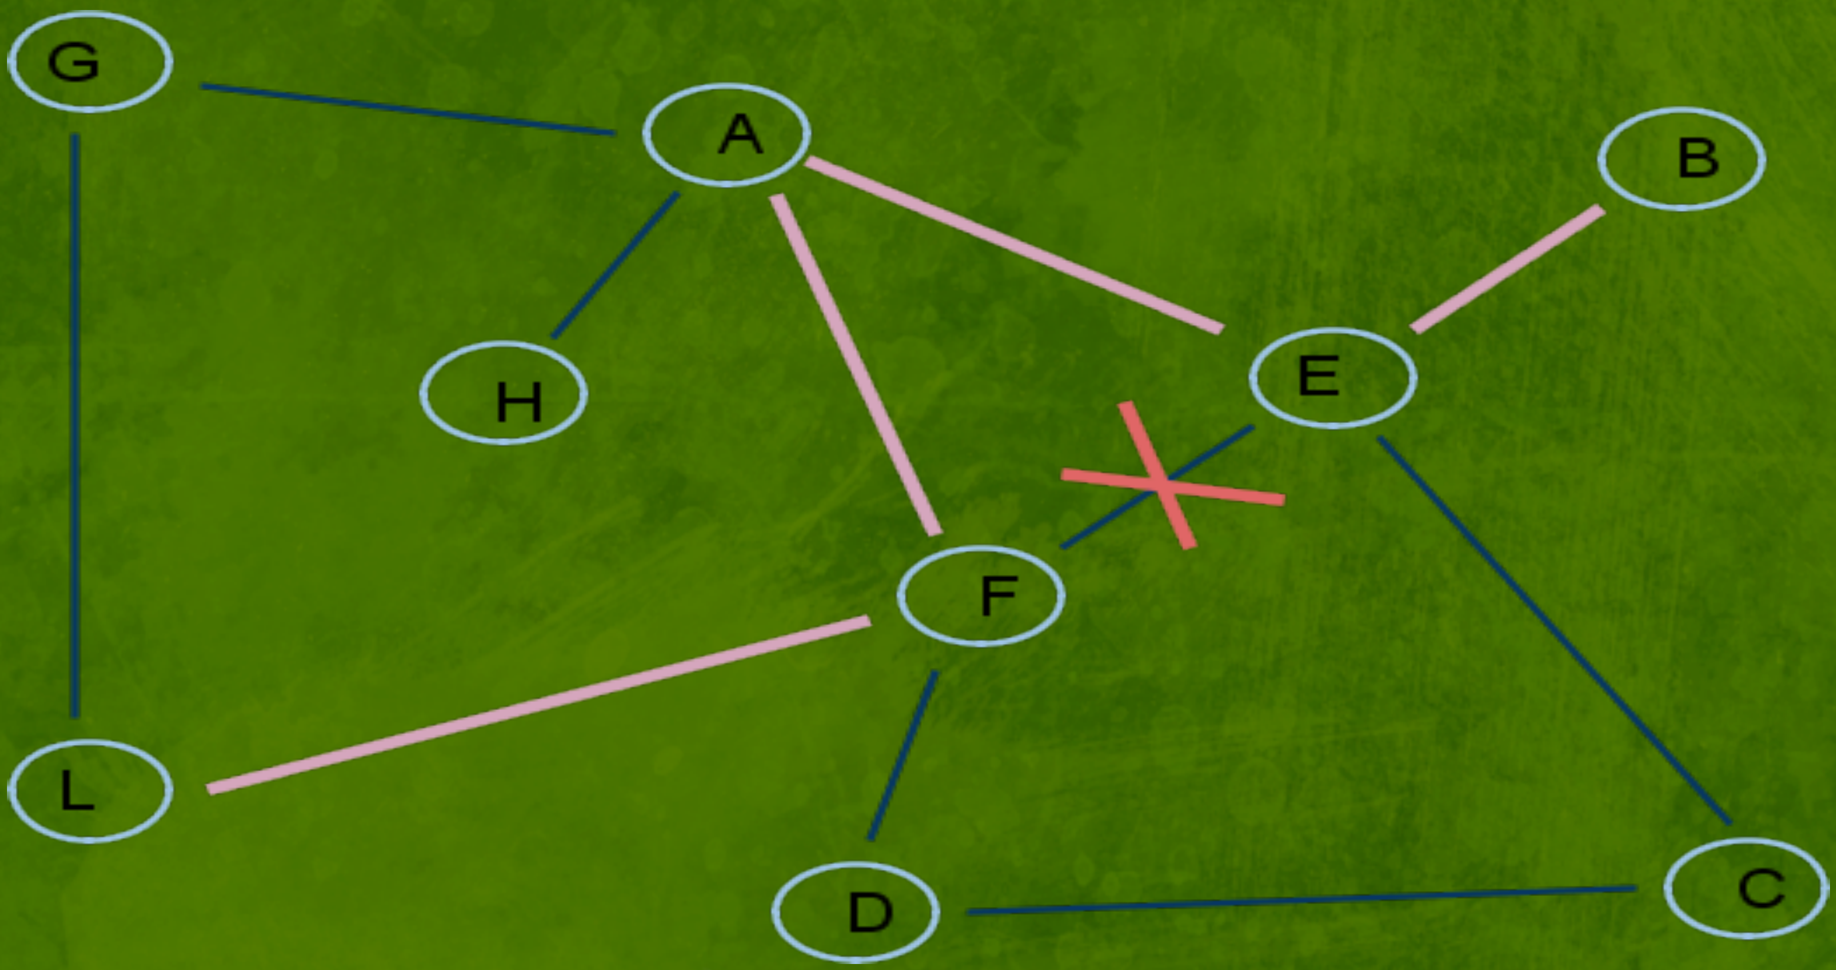

Operational Overview

From stability to new stability

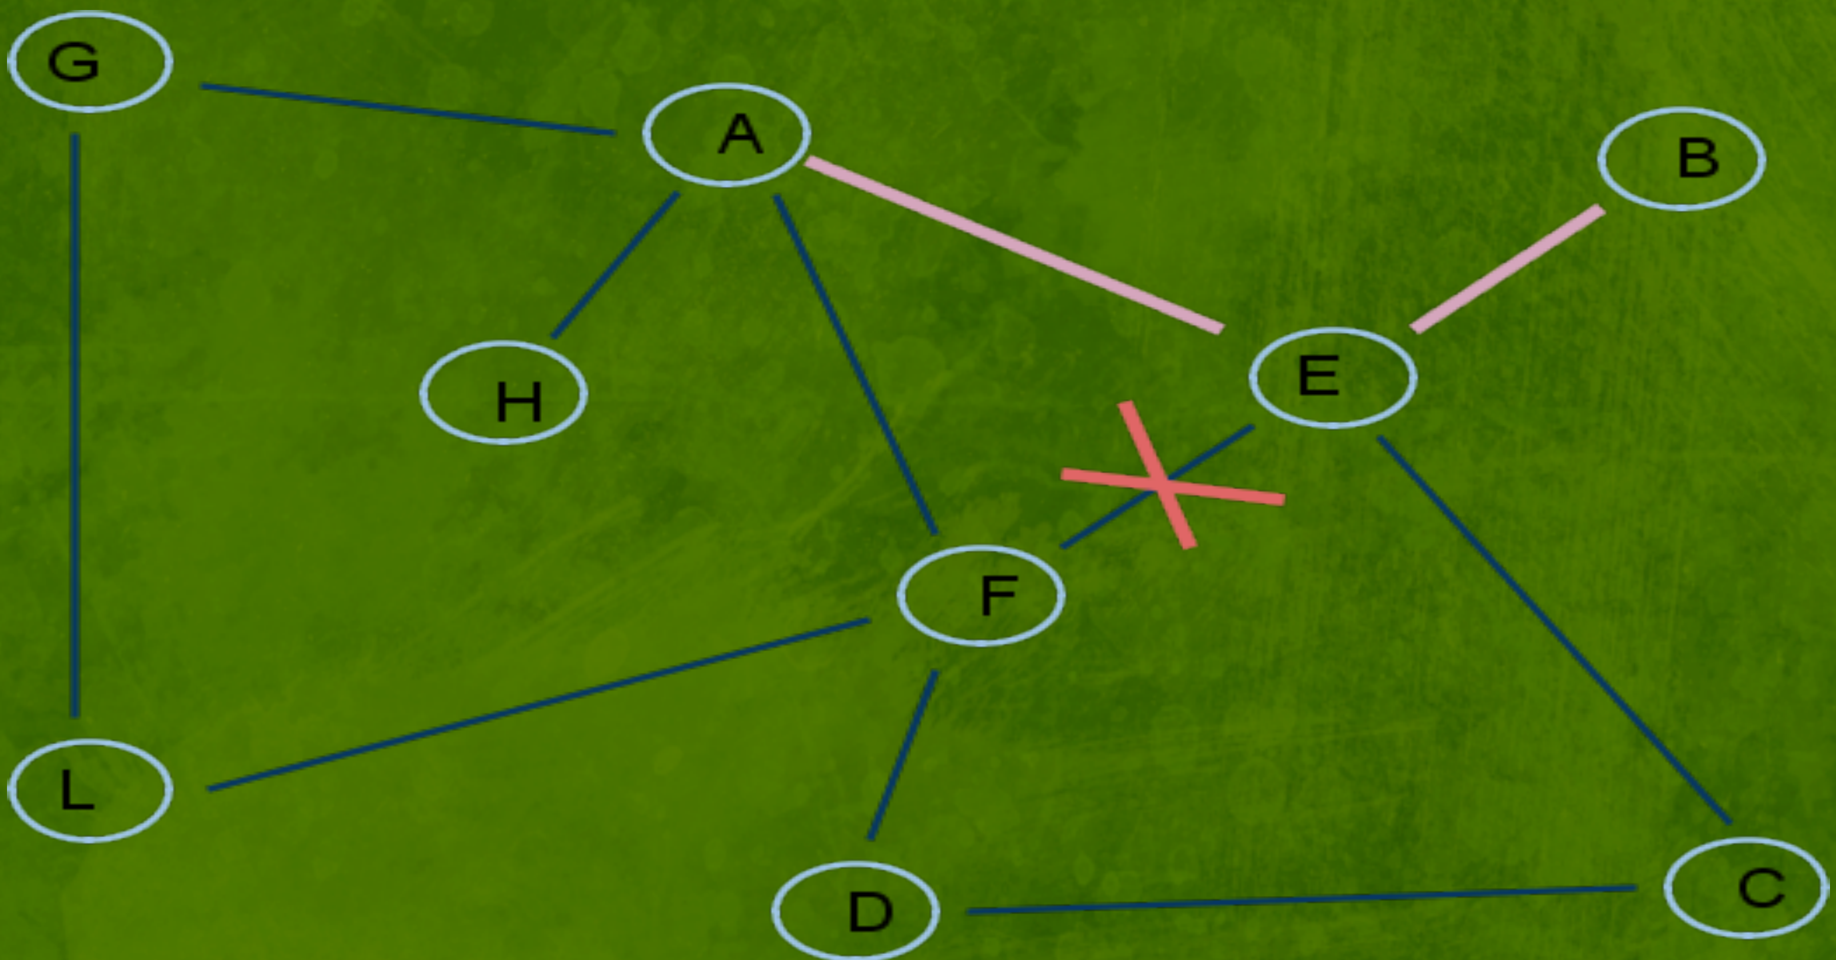

- The goals of Netsukuku
    - Sustainable Infrastructure
    - Accessibility, freedom of expression and privacy
- How it works
    - Why a new distributed protocol
    - Operational Overview
    - Hierarchical Structure
    - Architectural weaknesses and solutions
- What it achieves

Hierarchical Structure

Groups of nodes at n levels

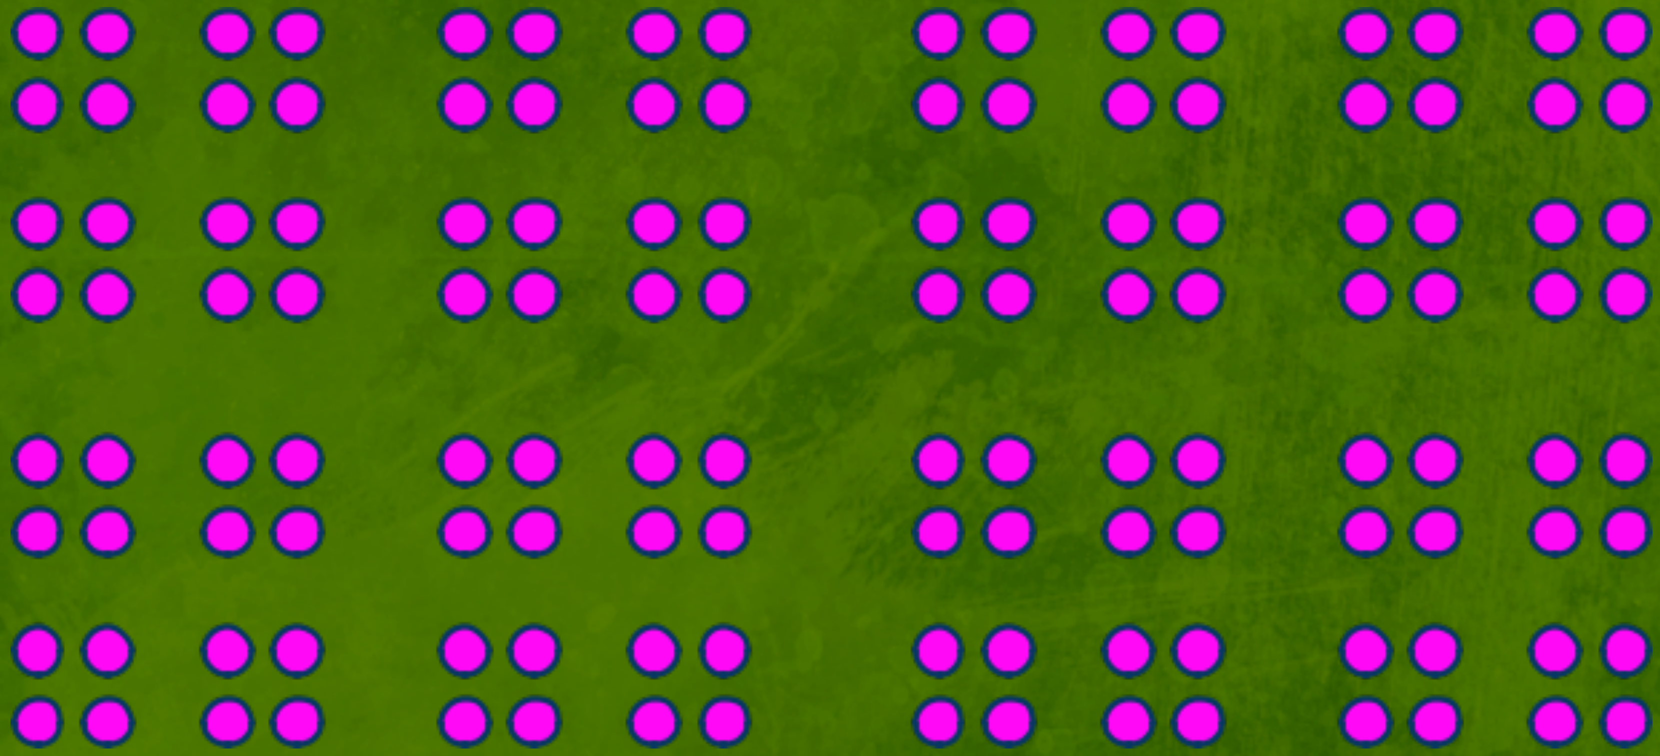

Hierarchical Structure

Groups of nodes at n levels

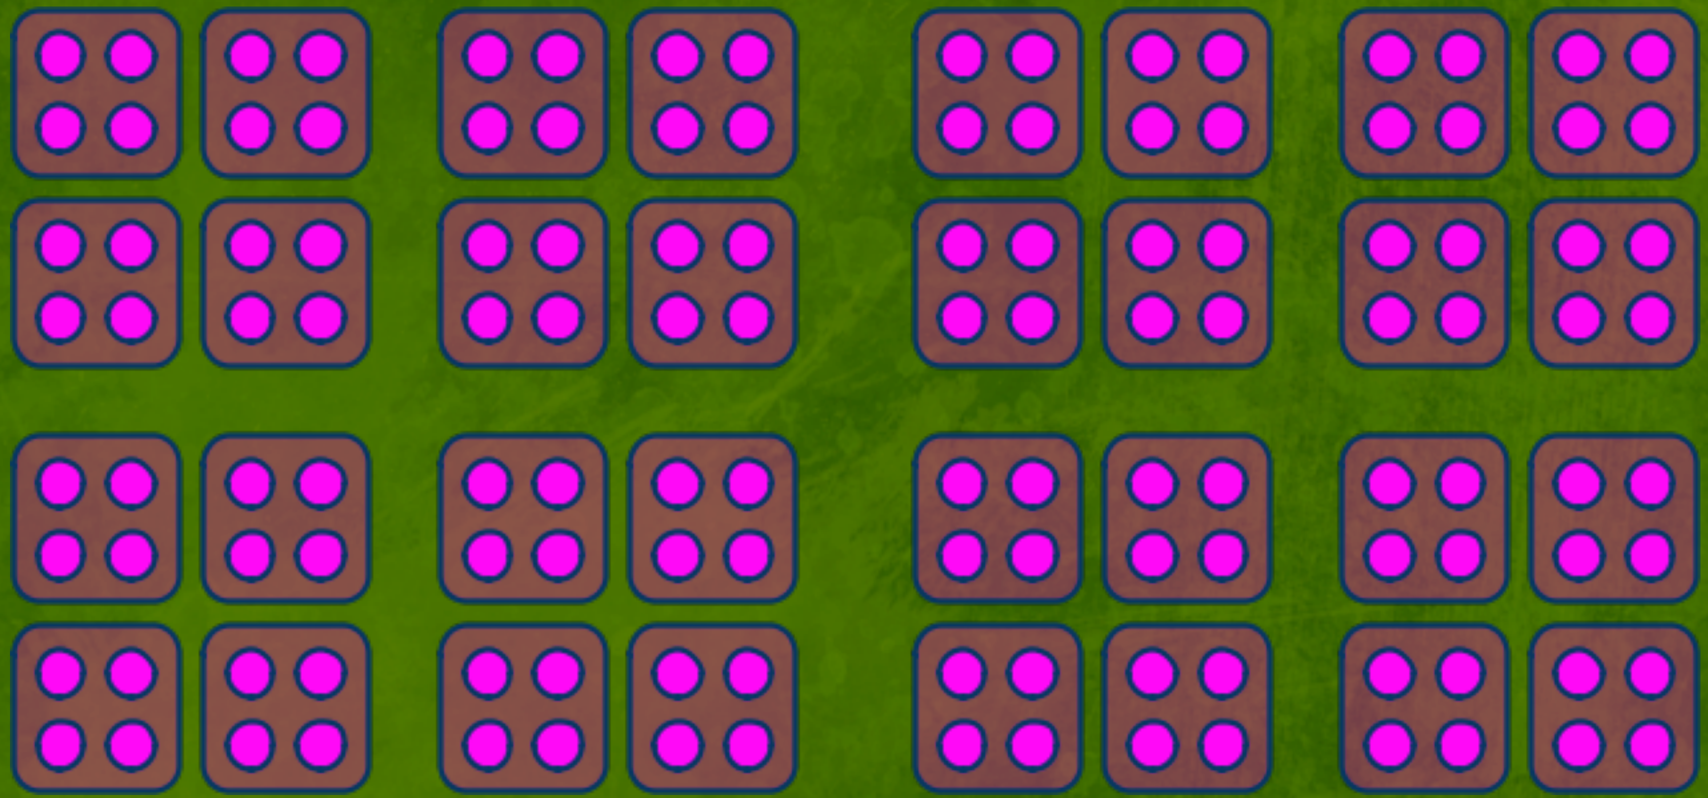

Hierarchical Structure

Groups of nodes at n levels

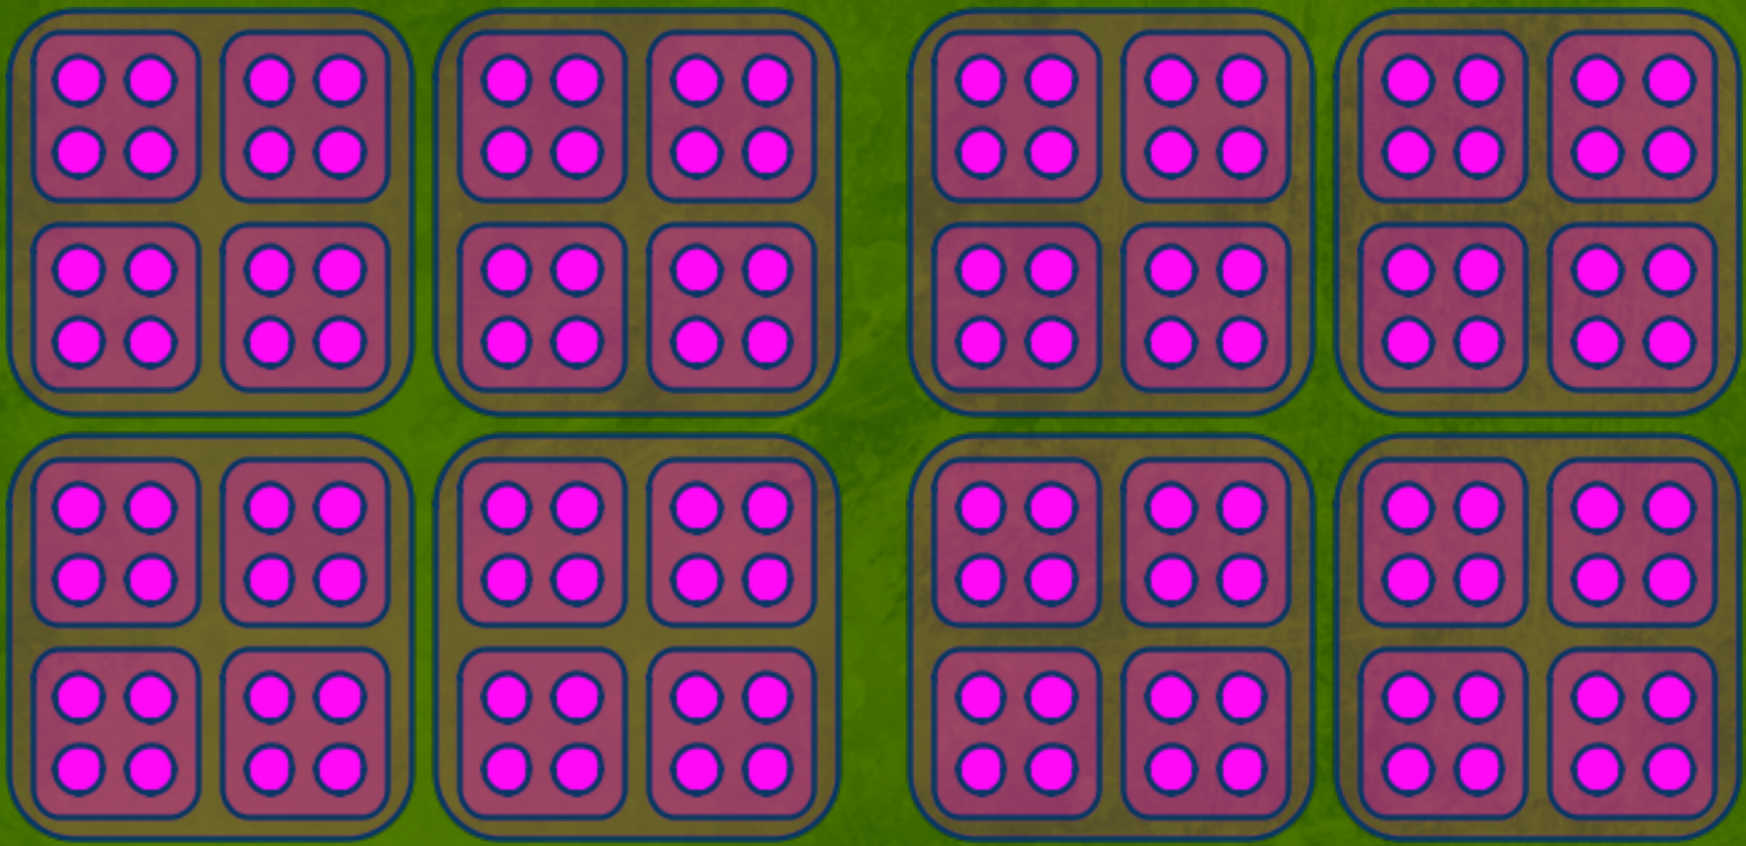

Hierarchical Structure

Groups of nodes at n levels

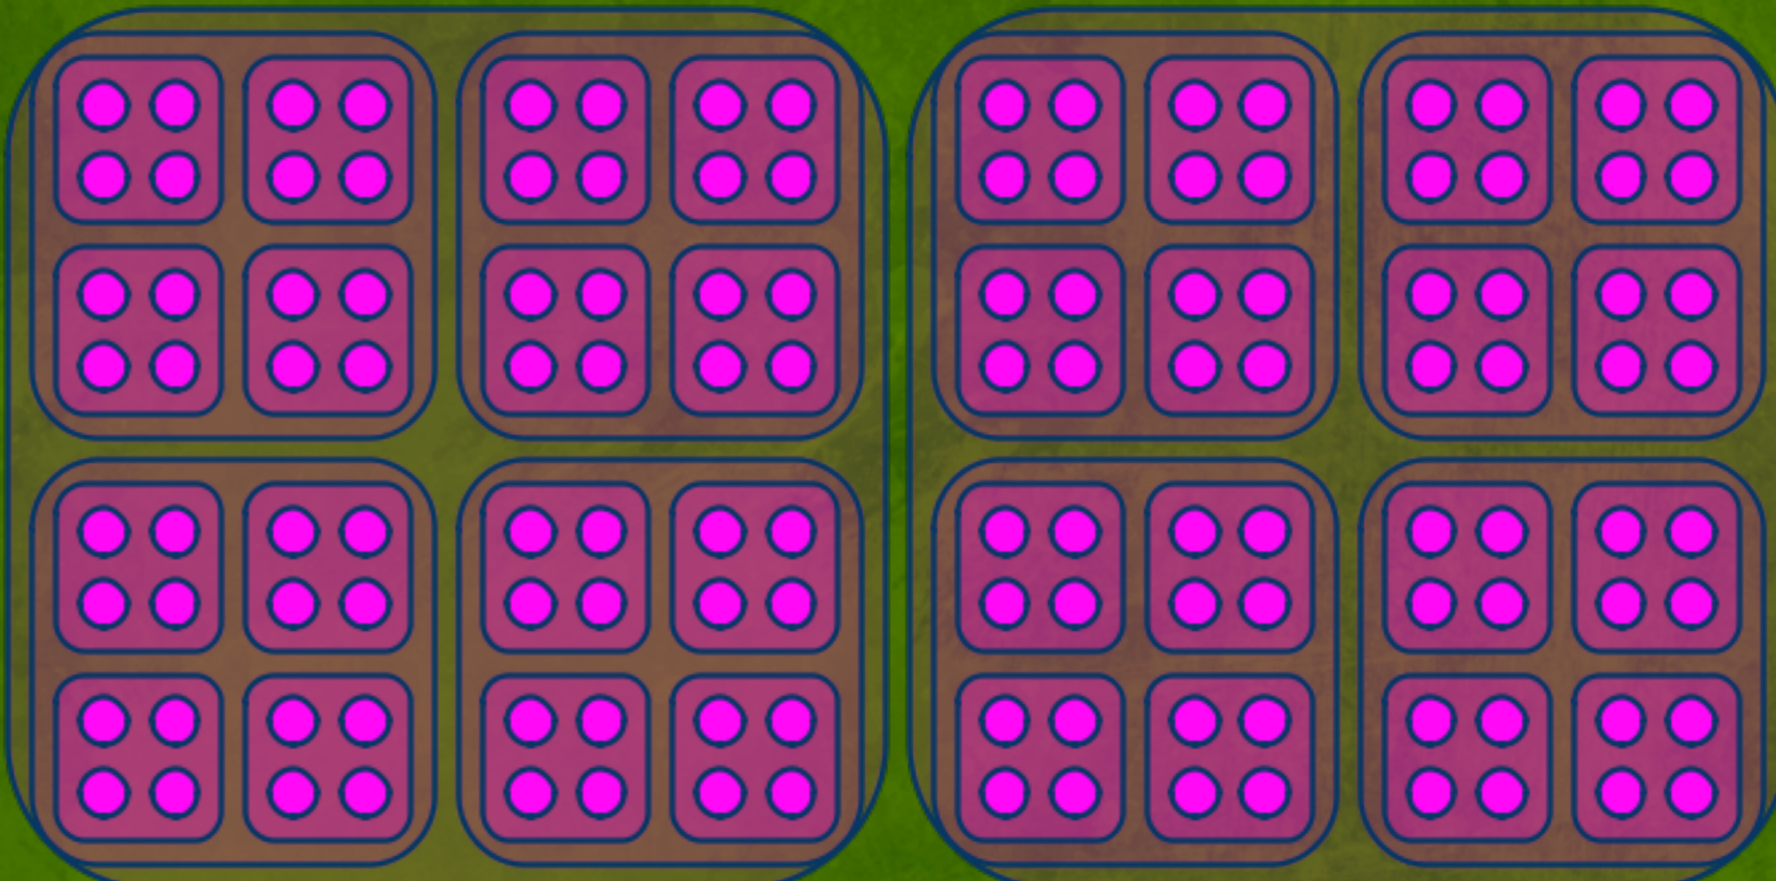

Hierarchical Structure

Groups of nodes at n levels

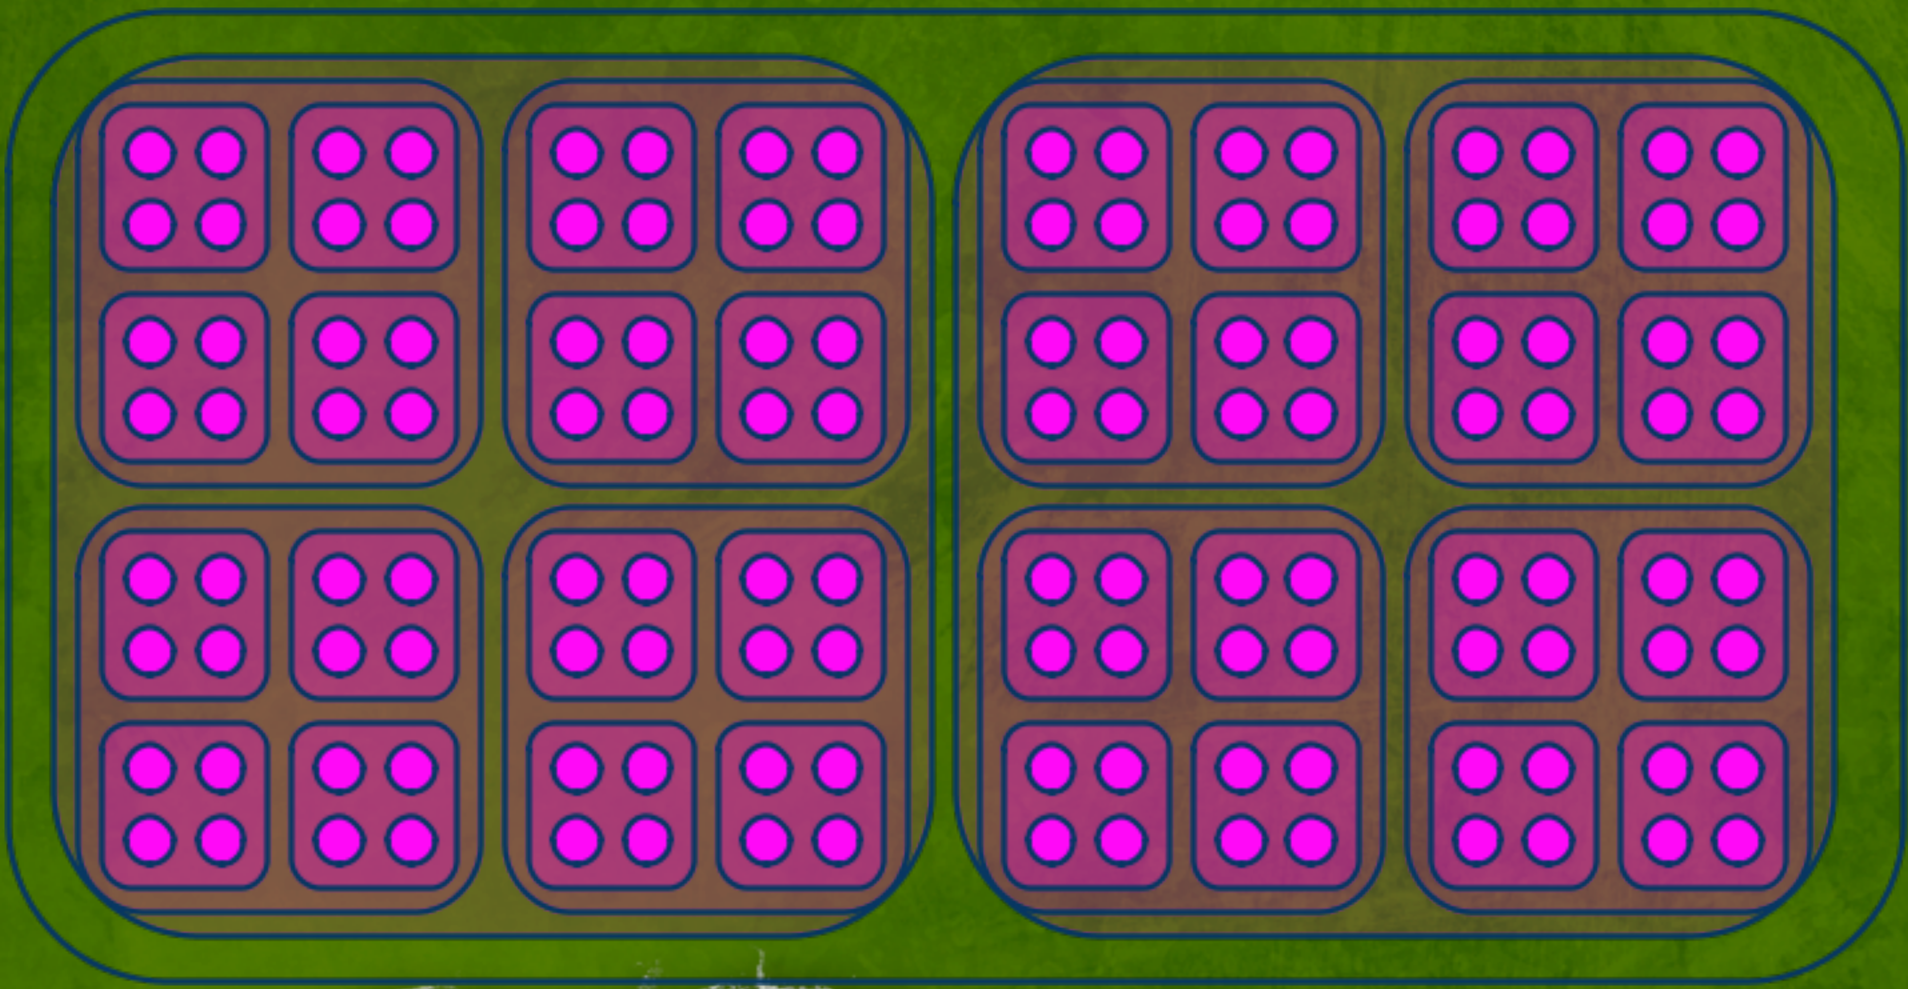

Hierarchical Structure

Groups of nodes at n levels

- Groups of m nodes. Up to level n. Max m^n nodes.

Hierarchical Structure

Groups of nodes at n levels

- Groups of m nodes. Up to level n. Max m^n nodes.
- m=256, n=4 : IPv4
- There is one gnode of level 4

Hierarchical Structure

Groups of nodes at n levels

- Groups of m nodes. Up to level n . Max m^n nodes.
- m=256, n=4 : IPv4
- There is one gnode of level 4
- It has max 256 gnodes of level 3
- Each gnode of level 3 has a unique ID, say 11.

Hierarchical Structure

Groups of nodes at n levels

- Groups of m nodes. Up to level n. Max m^n nodes.
- m=256, n=4 : IPv4
- There is one gnode of level 4
- It has max 256 gnodes of level 3
- Each gnode of level 3 has a unique ID, say 11.
- Gnode 11 has max 256 gnodes of level 2
- Each gnode inside it has a unique ID, say 11.22.

Hierarchical Structure

Groups of nodes at n levels

- Groups of m nodes. Up to level n. Max m^n nodes.
- m=256, n=4 : IPv4
- There is one gnode of level 4
- It has max 256 gnodes of level 3
- Each gnode of level 3 has a unique ID, say 11.
- Gnode 11 has max 256 gnodes of level 2
- Each gnode inside it has a unique ID, say 11.22.
- Gnode 11.22 has max 256 gnodes of level 1
- Each gnode inside it has a unique ID, say 11.22.33.

Hierarchical Structure

Groups of nodes at n levels

- Groups of m nodes. Up to level n. Max m^n nodes.
- m=256, n=4 : IPv4
- There is one gnode of level 4
- It has max 256 gnodes of level 3
- Each gnode of level 3 has a unique ID, say 11.
- Gnode 11 has max 256 gnodes of level 2
- Each gnode inside it has a unique ID, say 11.22.
- Gnode 11.22 has max 256 gnodes of level 1
- Each gnode inside it has a unique ID, say 11.22.33.
- Gnode 11.22.33 has max 256 nodes
- Each node inside it has a unique ID, say 11.22.33.44.

Hierarchical Structure

How this saves memory

- The node **g3.g2.g1.g0** will memorize info only for the direct children of the gnodes it belongs to.
  - 256 gnodes of level 3
  - 256 gnodes of level 2 inside gnode **g3** of level 3
  - 256 gnodes of level 1 inside gnode **g3.g2** of level 2
  - 256 nodes inside gnode **g3.g2.g1** of level 1

Hierarchical Structure

How this saves memory

- The node **g3.g2.g1.g0** will memorize info only for the direct children of the gnodes it belongs to.
  - 256 gnodes of level 3
  - 256 gnodes of level 2 inside gnode **g3** of level 3
  - 256 gnodes of level 1 inside gnode **g3.g2** of level 2
  - 256 nodes inside gnode **g3.g2.g1** of level 1
- Max num. of nodes is 256^4 = 4G.
- Each node has to memorize info for 256*4 = 1k gnodes.

Hierarchical Structure

How this saves memory

- The node **g3.g2.g1.g0** will memorize info only for the direct children of the gnodes it belongs to.
  - 256 gnodes of level 3
  - 256 gnodes of level 2 inside gnode **g3** of level 3
  - 256 gnodes of level 1 inside gnode **g3.g2** of level 2
  - 256 nodes inside gnode **g3.g2.g1** of level 1
- Max num. of nodes is 256^4 = 4G.
- Each node has to memorize info for 256*4 = 1k gnodes.
- If m=16, n=8 we have: 16^8 = 4G. 16*8 = 128

Hierarchical Structure

How this saves memory

- The same saving for the paths.
- For each destination we memorize all the hops.
- We will memorize only direct children of the gnodes we belong to.
- Worst case is not 256^4 hops, but just 256*4 hops.

Hierarchical Structure

How this saves CPU and net usage

- The same saving for the flooding of the network.
- We obtain a scalable discovery algorithm because each exploration of a given gnode is confined in the gnode itself.

- The goals of Netsukuku
  - Sustainable Infrastructure
  - Accessibility, freedom of expression and privacy
- How it works
  - Why a new distributed protocol
  - Operational Overview
  - Hierarchical Structure
  - Architectural weaknesses and solutions
- What it achieves

Architectural weaknesses and solutions

due to the hierarchical structure

- Allocation of free IPs isn't a trivial matter.
  - When the gnode level 4 is saturated a node must be in the same gnode of level 3 of one of its neighbours. Say we have to use gnode 11.
  - When gnode 11 (level 3) is saturated a node must be in the same gnode of level 2 of one of its neighbours. Say we have to use gnode 11.22.
  - ...
- A solution: Communicating Vessels
  - All the gnodes of the net will have approximately the same number of nodes, at any time.
- A downside: more network updates are required.

Architectual weaknesses and solutions

due to the hierarchical structure

- Routes are approximated.
  - A node does not know the paths outside of the gnode it belongs to.
  - Say we are in gnode 11.22.33, node X and node Y are inside 11.22.44.
  - The best route to X could be not the same of the best route to Y.
  - But we memorize just the routes to the gnode 11.22.44. So we have only a best route to the gnode 11.22.44.
- The problem is mitigated by the Communicating Vessels
- Another improvement: use Branching Multipath.

Architectural weaknesses and solutions

due to the independence from Internet

- Wi-Fi is not always feasible.
    - Wireless links are cheap to deploy, but not always performant enough or possible at all.
- A solution: Internet tunneling
    - Viphilama permits to connect, with Internet tunnels, nodes which aren't physically linked. Then whenever a physical link is available, we can remove the virtual one without interfering with the stability of the network.
    - a practice that makes a global Netsukuku network more feasible in the immediate future.
    - that should be discouraged in the long term.

Architectural weaknesses and solutions

due to the independence from Internet

- The IP address of a node in Netsukuku can clash with the IP of a node in Internet
  - This means that a node cannot be part of Netsukuku and Internet at the same time.
  - Netsukuku is completely independent from the Internet. However, a good level of integration between the two networks must be preserved.
- A solution: Use of restricted IP range
- Another solution: Net Split
  - a node is connected to the Internet and to a Netsukuku network.
  - an hack using netfilter in linux makes it possible to communicate with nodes from both the networks at the same time, without restriction of IP range.

Advantages

due to the use of a P2P system

- Distributed services can be easily built as overlays.
  - ANDNA hostnames management
    - Distributed and decentralised system
  - Cryptographic Layer, transparent for the users
    - Encrypted communications
    - Fully anonymous connections (hidden client)
    - Fully anonymous registration of hostnames (hidden server)

- The goals of Netsukuku
  - Sustainable Infrastructure
  - Accessibility, freedom of expression and privacy
- How it works
  - Why a new distributed protocol
  - Operational Overview
  - Hierarchical Structure
  - Architectural weaknesses and solutions
- What it achieves

What it achieves

anonymity

- A node can connect to the network from any point at any time, needing to retain no inherent identifying characteristics.
- The IP address identifying a computer is chosen randomly.
- This makes Netsukuku achieve a high level of anonymity, even without the cryptographic layer.

What it achieves

no censorship

- No need for a central organizing authority, since the network is dynamically configured.
- Who wants to connect does not rely on the service of a single entity.
- So, it's not possible to take control of the network. The only way to manipulate or demolish it would be to knock physically down each single node.

What it achieves

sustainable network. accessibility

- A node may have many links towards many neighbours.
- Full exploitation of all the links.
- This is an incentive for any user of the network to add more interconnections of greater quality to more nearby nodes, substantially improving the network for al.

Demo

Where we are

Contacts

Luca Dionisi: luca.dionisi@gmail.com
Website: www.netsukuku.org