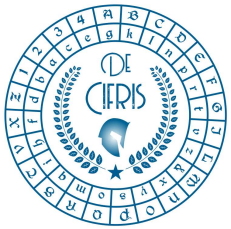

IOTA and the TANGLE

Guglielmo Morgari

Telsy SpA and Politecnico di Torino

De Cifris Augustae Taurinorum, 18 June 2021

Telsy

Talk Overview

Introduction to IOTA

The tangle

- validation
- confirmation
- double-spending
- tip selection

Attacks

Cryptography

The coordinator

What next (IOTA 1.5 and IOTA 2.0)

Open issues

IoT and Blockchains

IoT networks are typically made of a **huge number** of **low-power** devices **frequently** issuing **low-value** micro-transactions

Parameter	Typical blockchain	IoT needs
Throughput	Low (mining bottleneck)	High
Finalization Time	High (wait for N blocks)	Low
Cost (per-transaction)	High (fees)	Low
Decentralization	Questionable (too powerful mining pools)	Desirable
Sustainability	Unacceptable energy consumption (PoW)	Desirable

Current blockchain technology is not well suited for IoT world

IOTA aims to fill the gap

IOTA

- Created in 2015
  - by David Sønstebø, Dominik Schiener, Sergey Ivancheglo, and Serguei Popov
- Focused on Internet of Things (IOT)
  - physical devices able to collect, process and exchange data
- Supported by IOTA Foundation
  - established 2017, Berlin. No profit organization

- Partnerships and collaborations with Industry and Academia
  - Jaguar Land Rover, STM, Dell,Ubuntu/Canonical, Innogy, Microsoft, Cisco, Foxconn, Bosch, …

Some Numbers

- Fixed circulating supply
- 2.779.530.283.277.761 iota tokens
- Completely generated in the first special transactions (Genesis)
- Often exchanged as MIOTA (1 MIOTA=10^6 iota)

- MIOTA current value = $1.08

The Tangle

- IOTA is based on a Direct Acyclic Graph (DAG) structure (tangle)
- Transactions as nodes, validations as edges

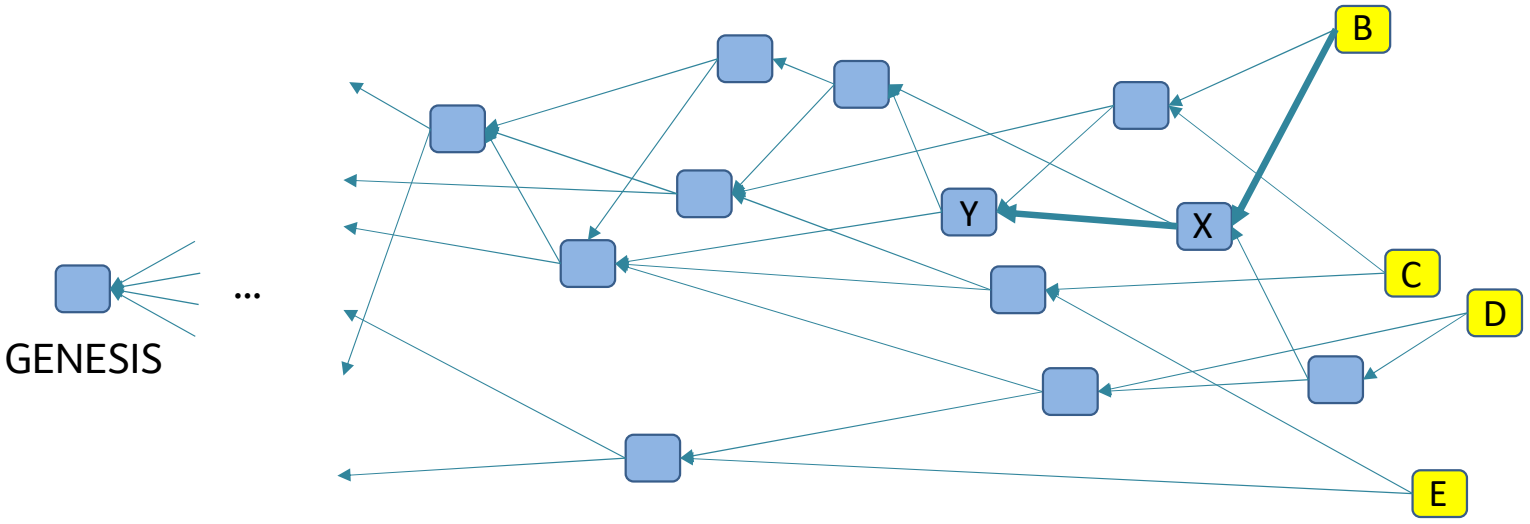

- Validation between transactions can be direct (B→X, X→Y) or indirect (B→Y)
- A not-yet validated transaction is named tip (B, C, D, E)

Validation

- Any new transaction must validate two old transactions (hopefully tips)
- To validate a transaction means to check its correctness
- A correct transaction must
  - be well formed
  - be correctly signed with the sender's private key
  - include a (relatively simple) proof of work
  - be consistent with its cone of past transactions
    - i.e. all the transactions directly or indirectly validated back to the Genesis
    - basically checks balances are always non negative and consistent
  - be consistent with the other selected transaction and its cone of past transactions
- Validation effort can be significant

Validation

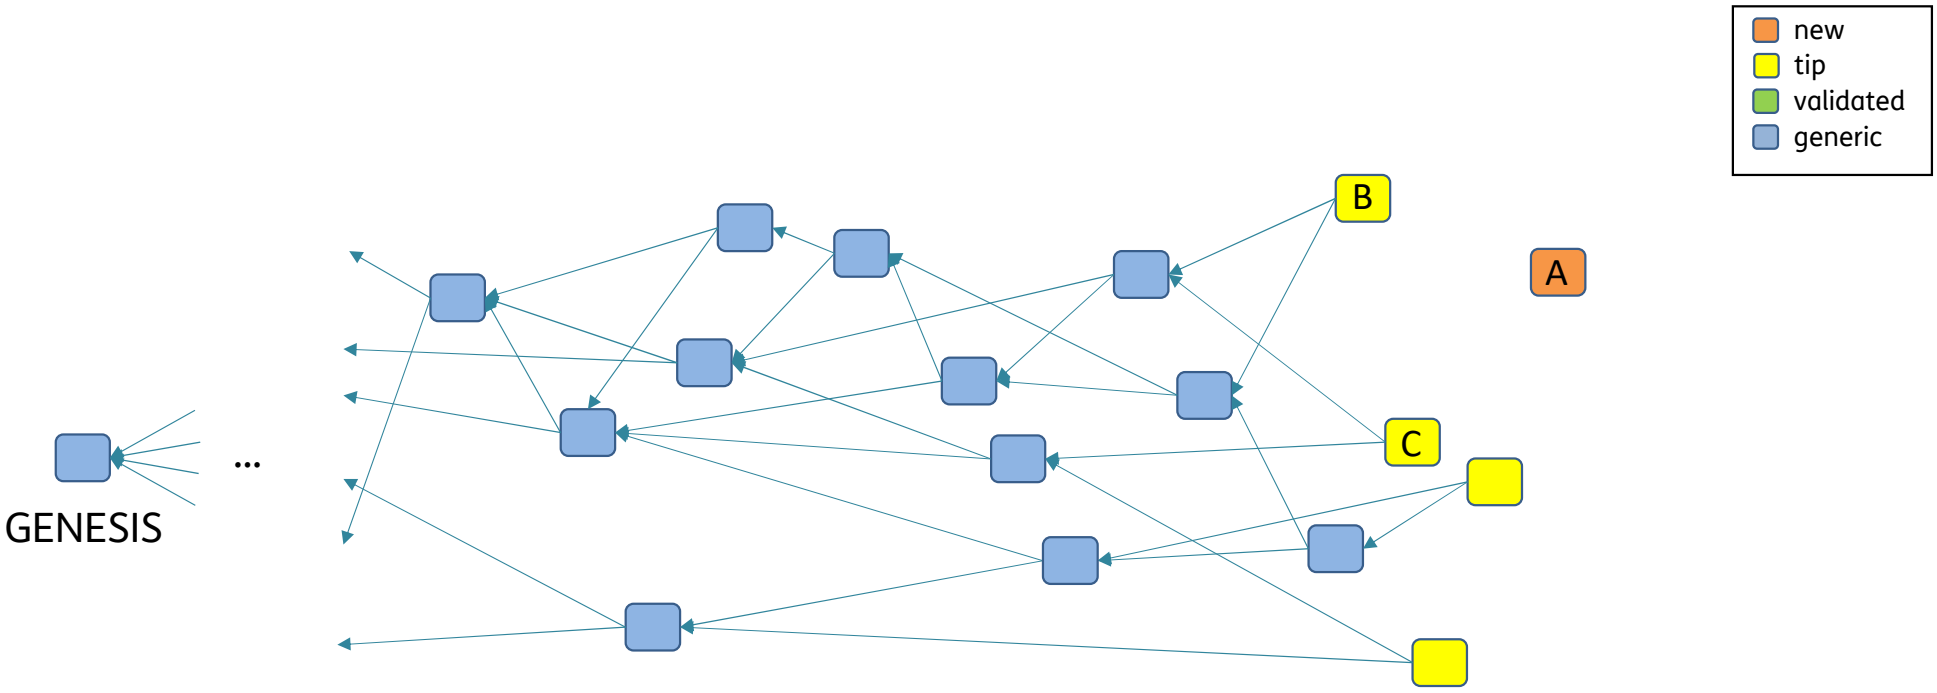

Validation

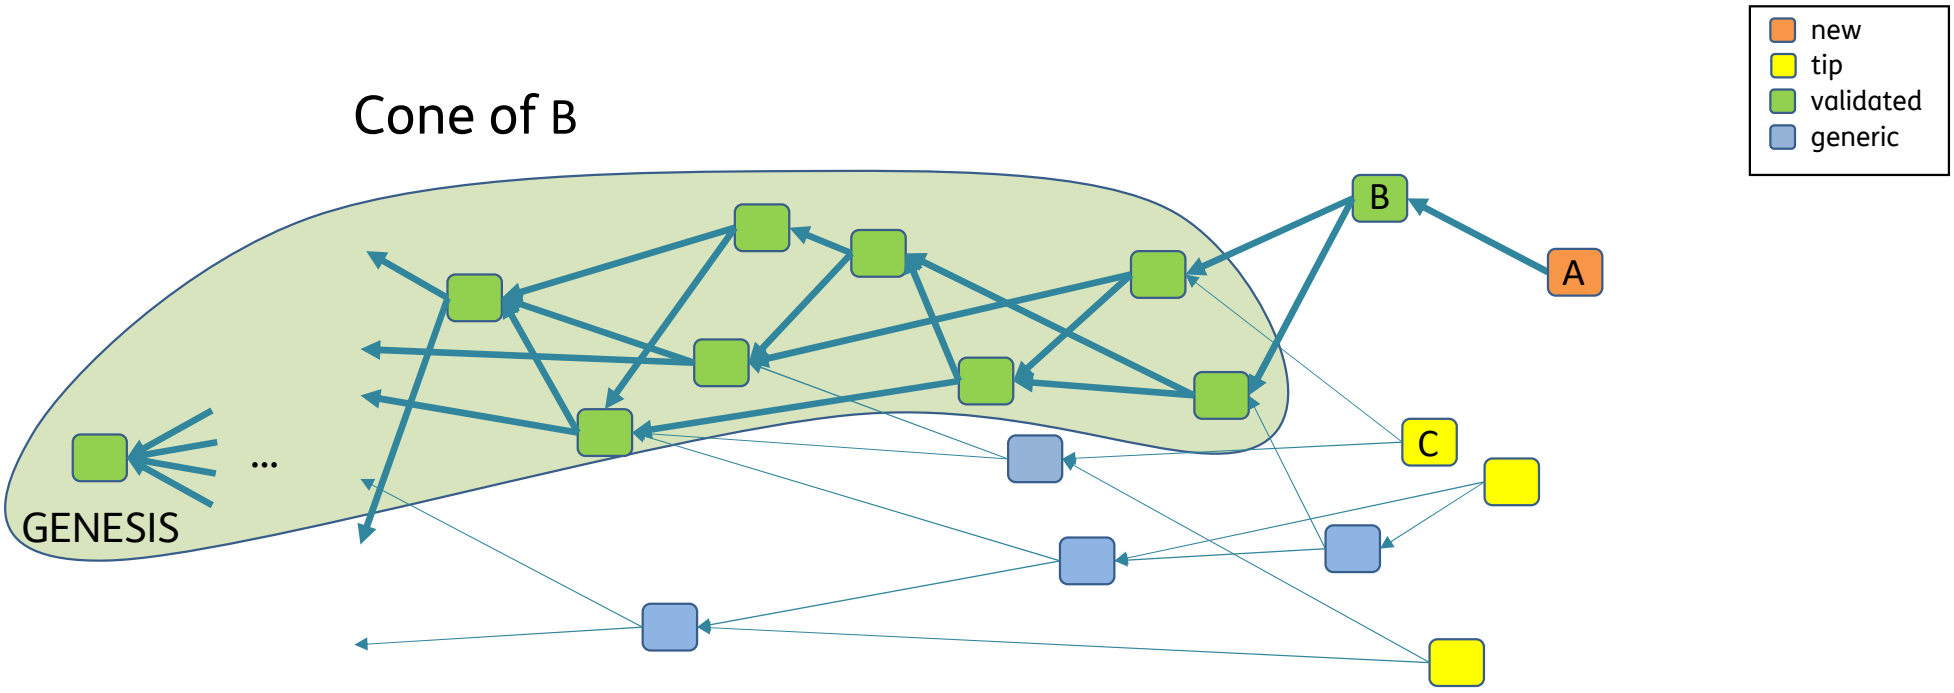

Validation

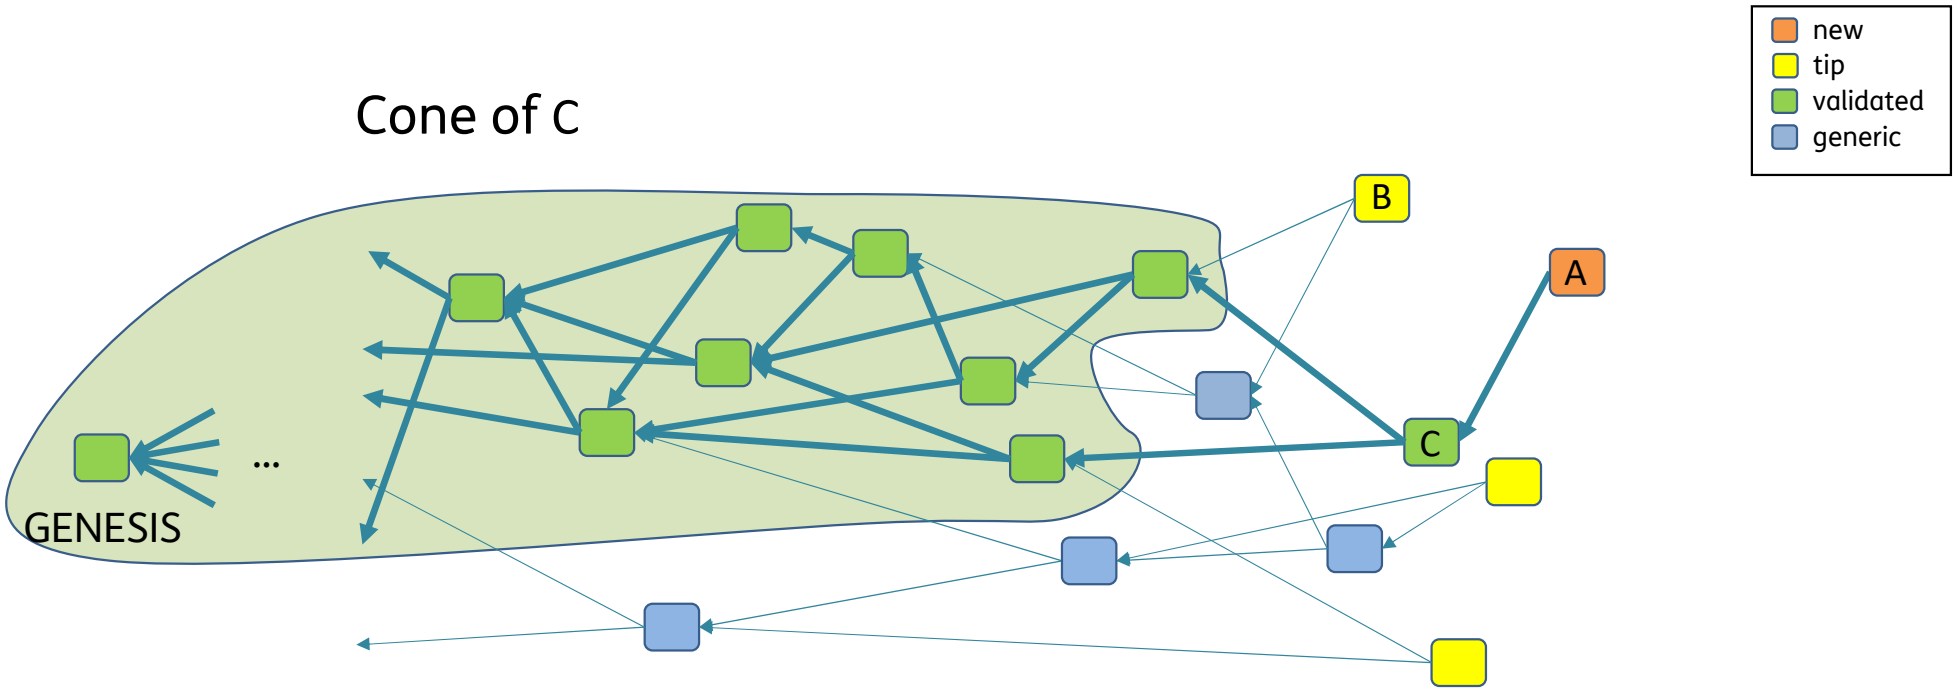

Validation

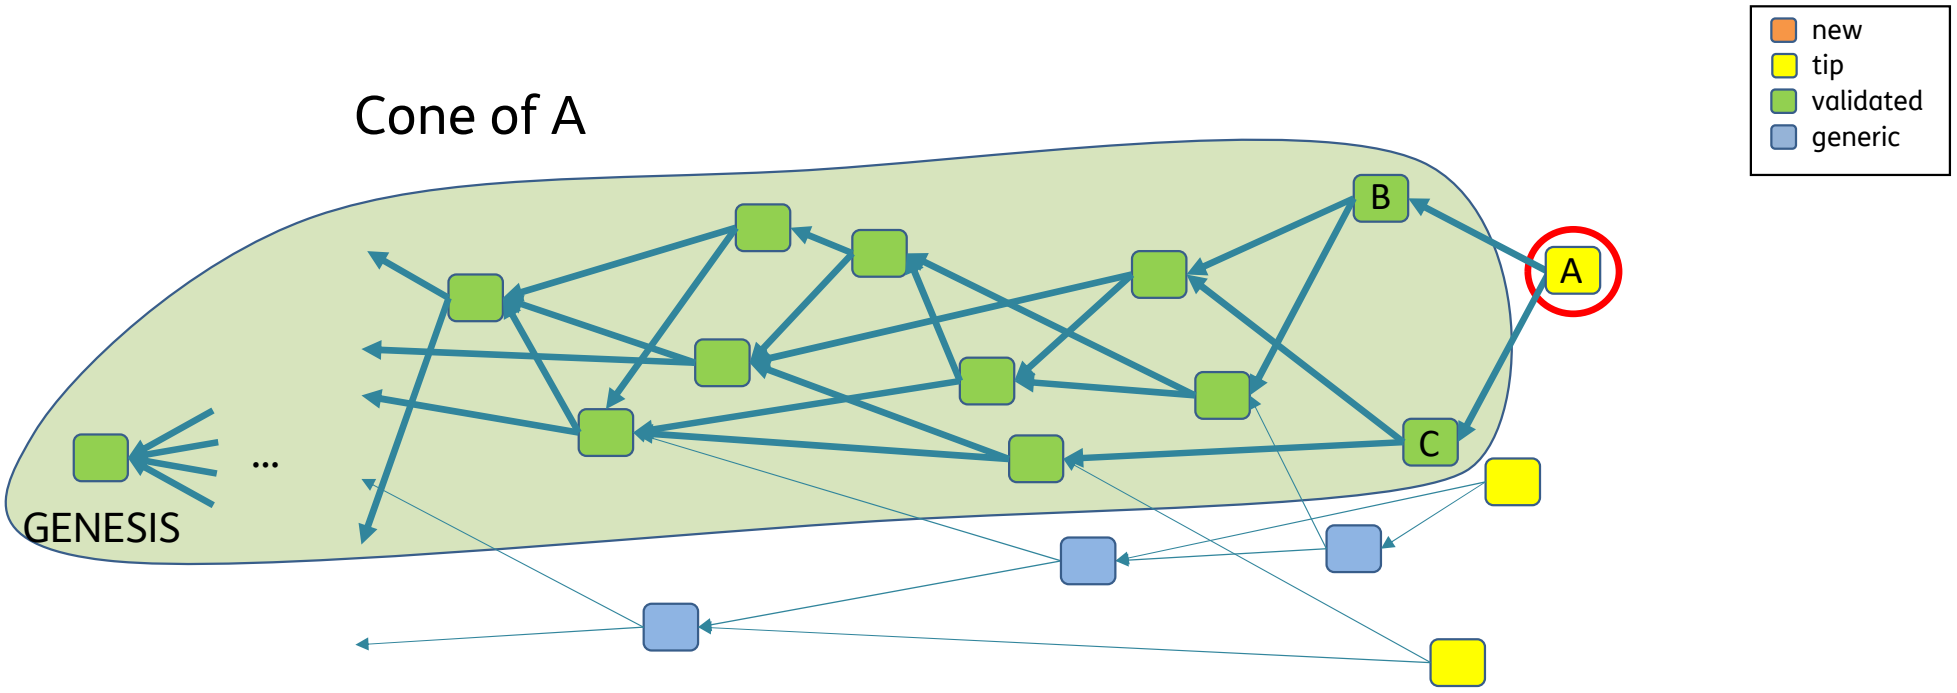

A validates B and C and becomes a new tip

Tangle vs Blockchain

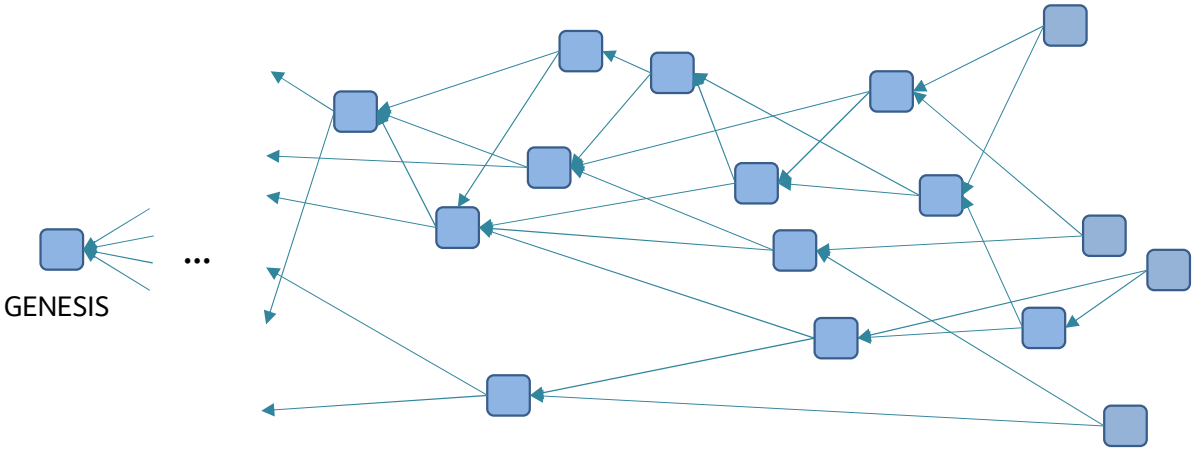

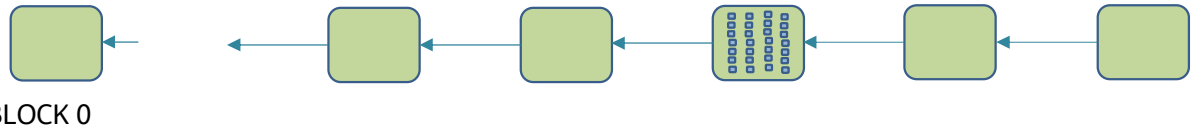

(claimed) Tangle advantages

Throughput
- No bottleneck (no limited size blocks)
- Full scalability (more transactions, more validation power)

Finalization
- Almost instantaneous
- More transactions, faster validation

Cost per user
- No miners, no fees
- Only a small PoW (to avoid spam)

Decentralization
- Fully decentralized
- No few miner pools with overwhelming power

Sustainability
- Only small PoW for each transaction
- Negligible overall cost

Double Spending (I)

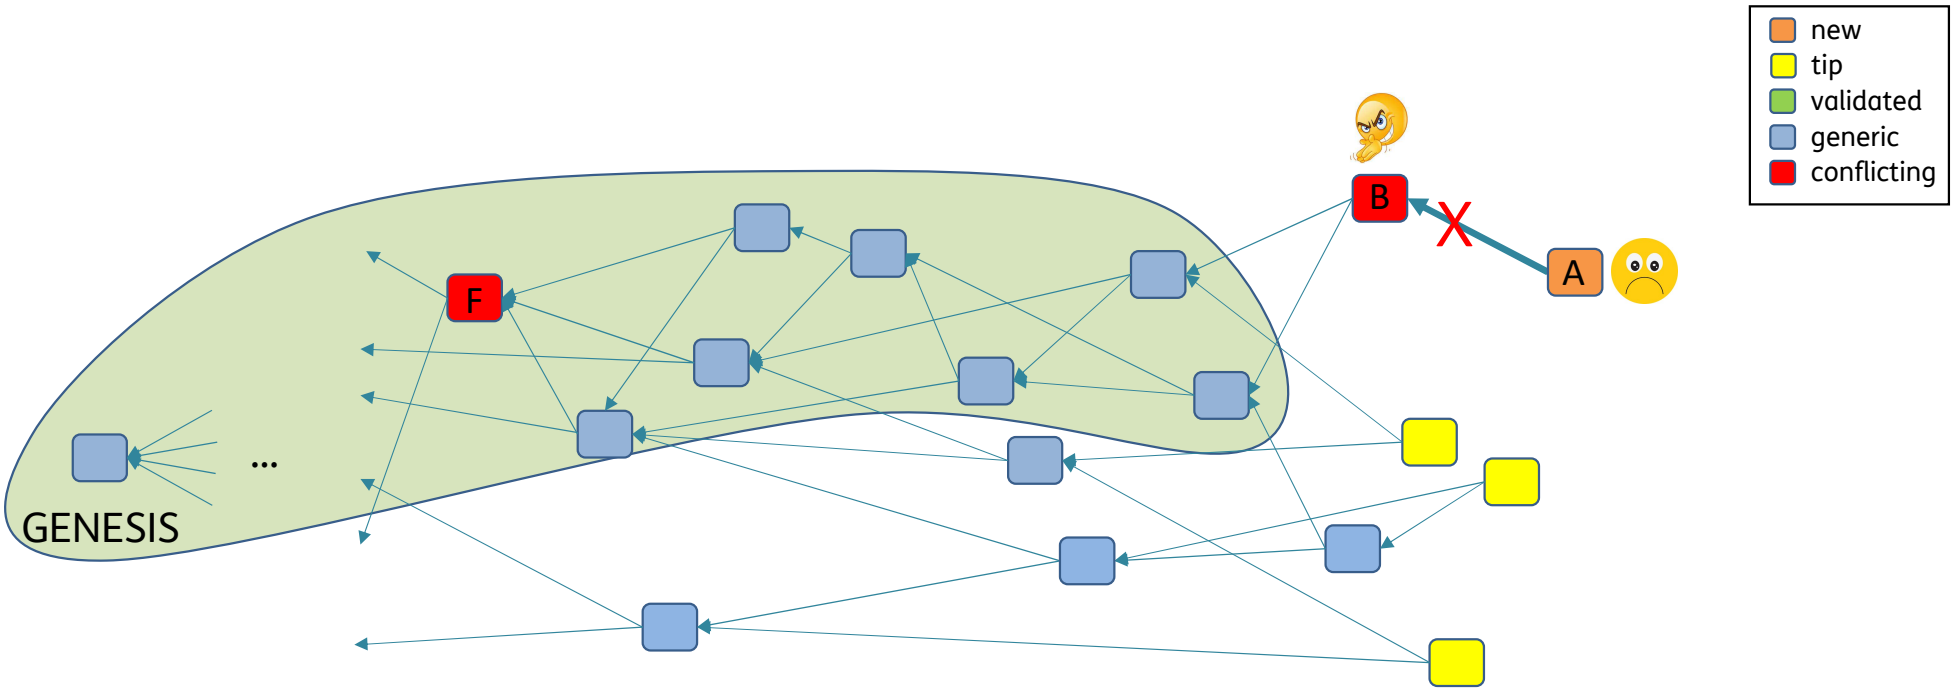

B and F conflicting, B (indirectly) validating F
B tip cannot be chosen for validation by A

Double Spending (II)

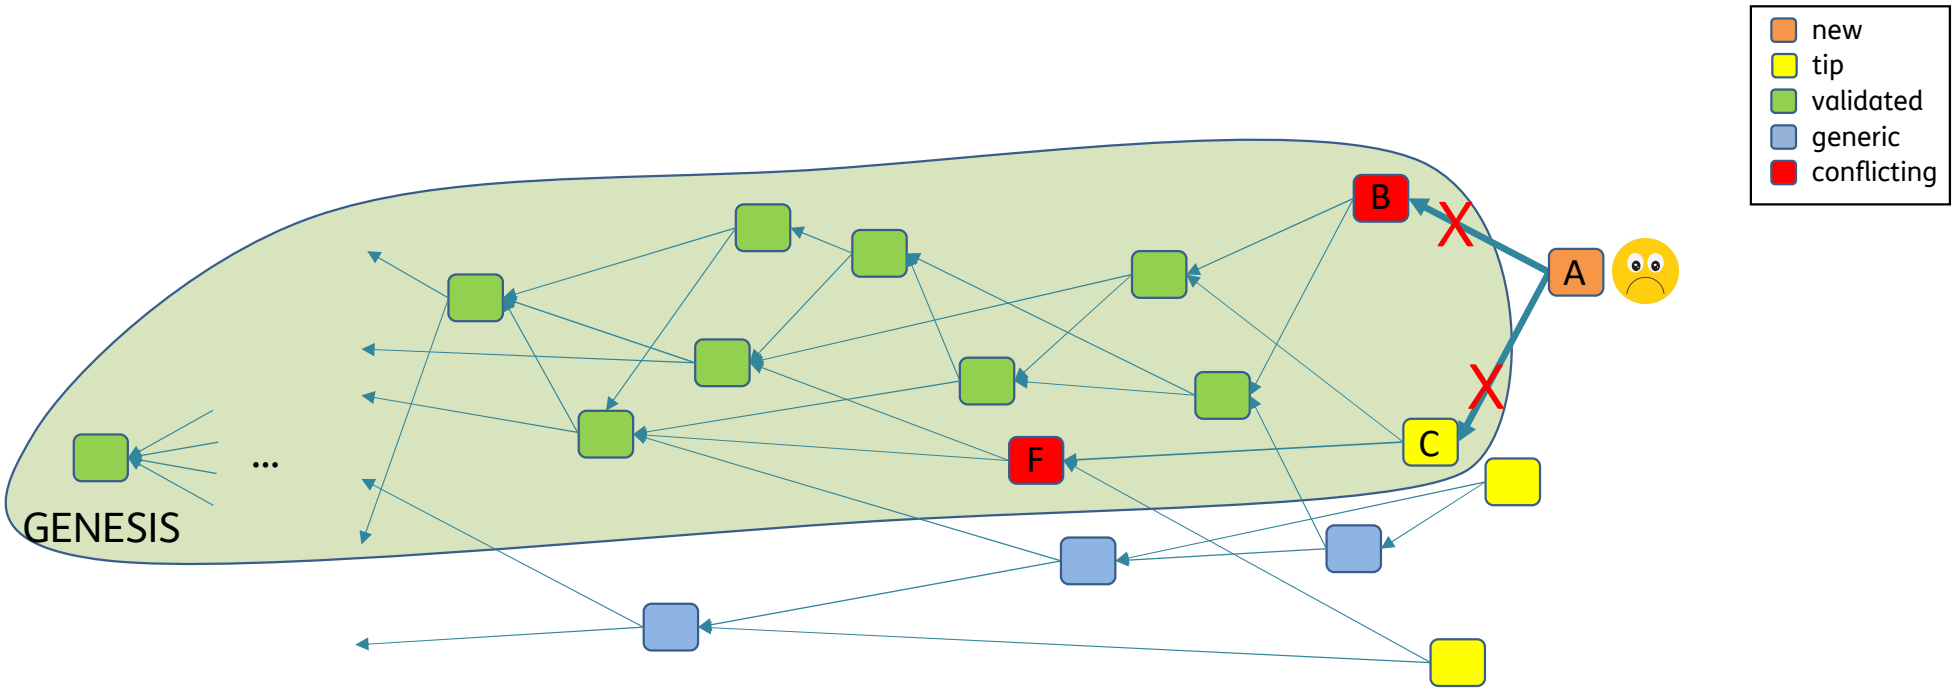

B and F conflicting
The (B, C) tips pair cannot be chosen for validation by A

Double Spending (III)

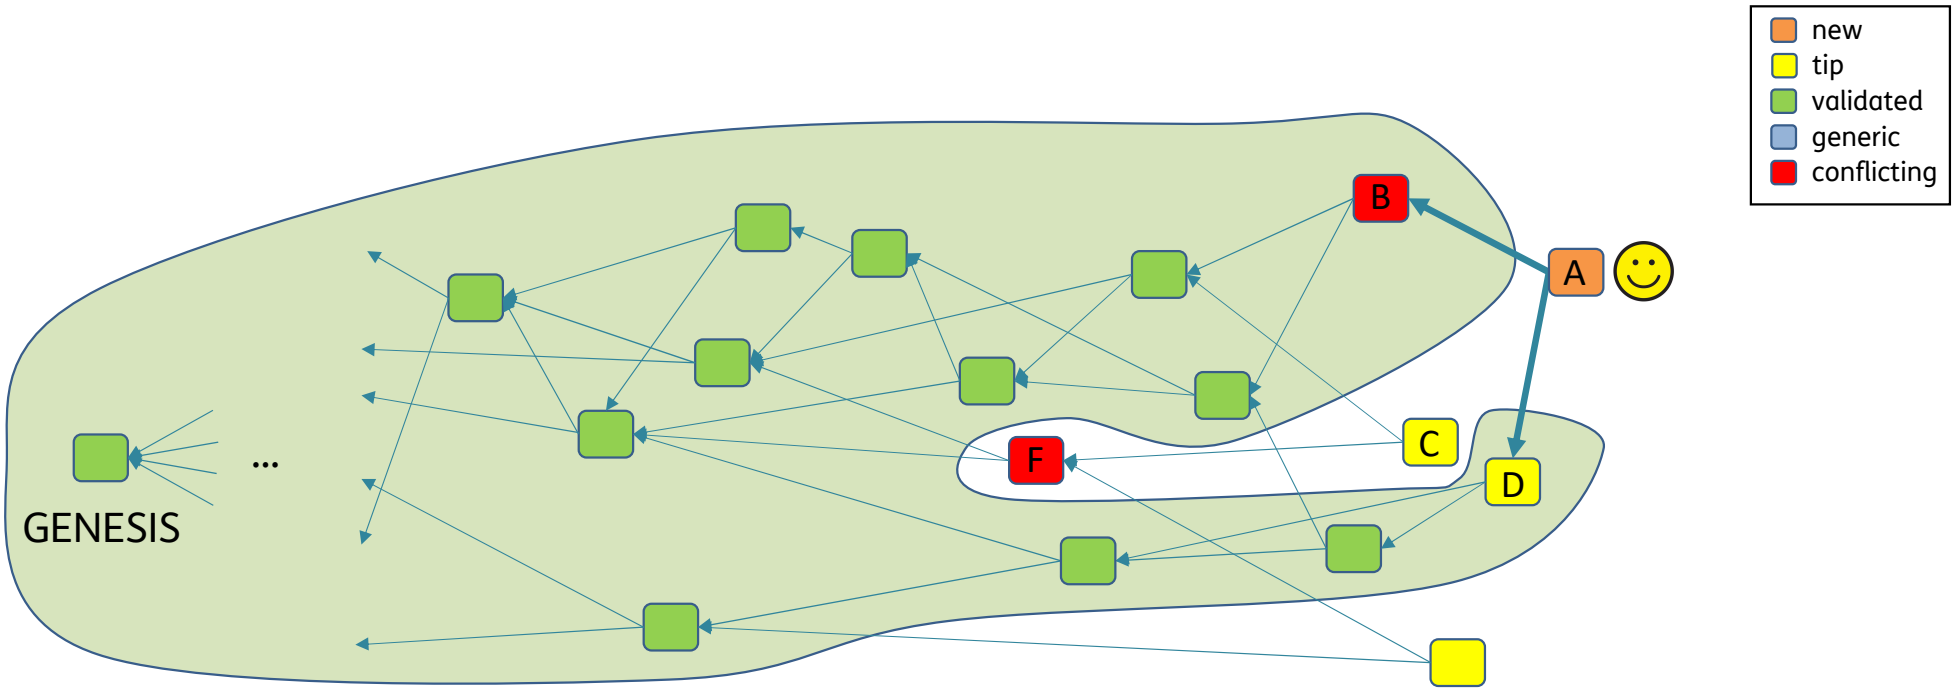

Conflicting transactions (double spending) can coexist in the tangle
How to know which one is “good” (consensus problem)?

Confirmation

- When can a transaction in the tangle be considered “safe”?
- How to define a “confirmation level”?

- Suggested strategy:
  - count how many tips directly or indirectly validate the transaction
  - check if a given threshold (according to the context) is reached, e.g.
    - 70% confirmation confidence: ok for small valued transactions
    - 100% confirmation confidence : required important transactions

- Rationale:
  - The more tips confirm a transaction, the more it is deep in the tangle and unlikely to be later discarded

Confirmation

Confirmation level of a transaction: number of tips (directly or indirectly) validating the transaction

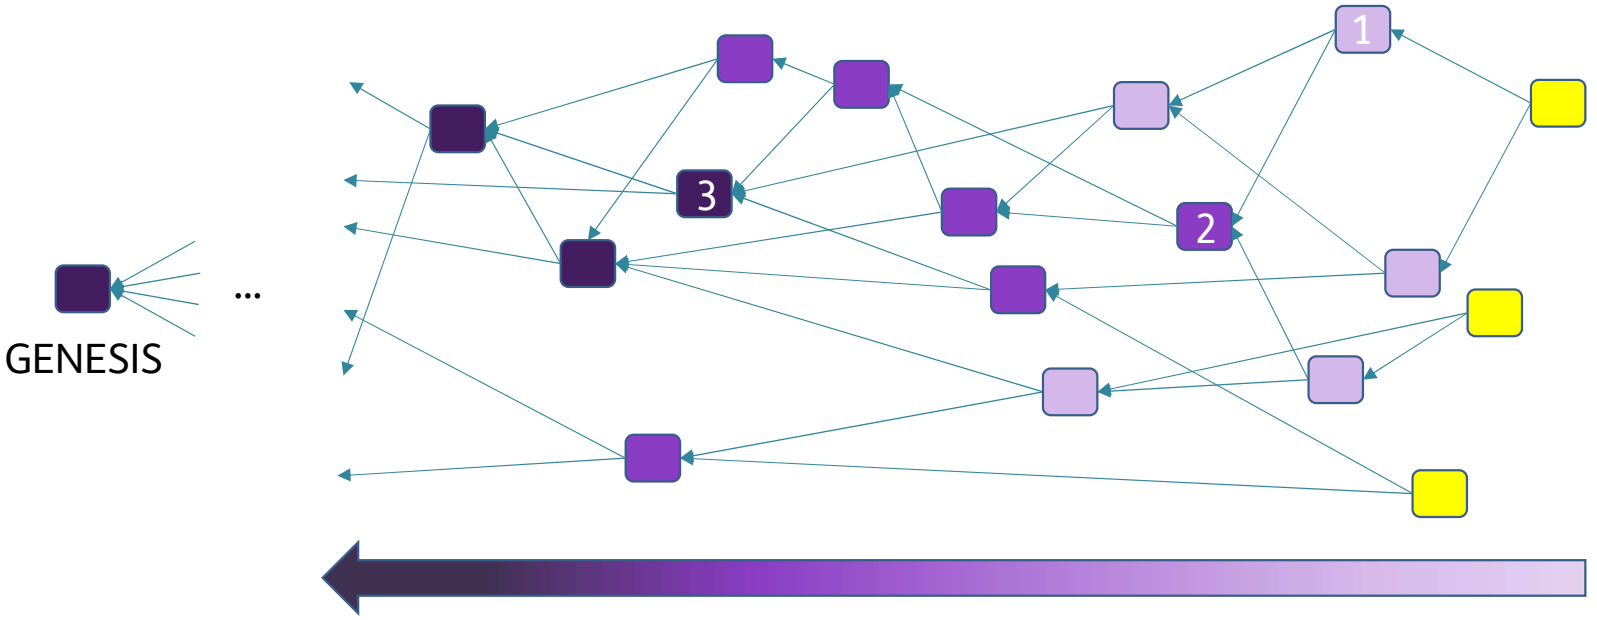

The deeper in the tangle, the higher the confirmation level

Tip Selection

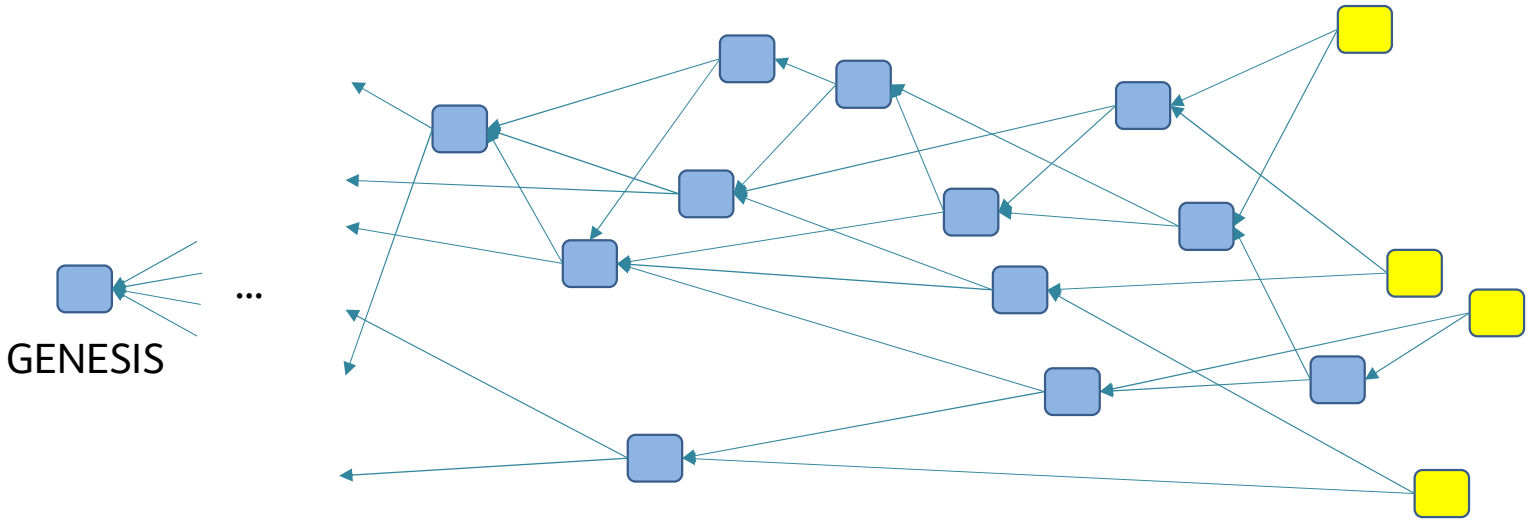

- Each new transaction must validate two existing transactions
- It is suggested that these transaction are tips, but it is not mandatory

Tip Selection

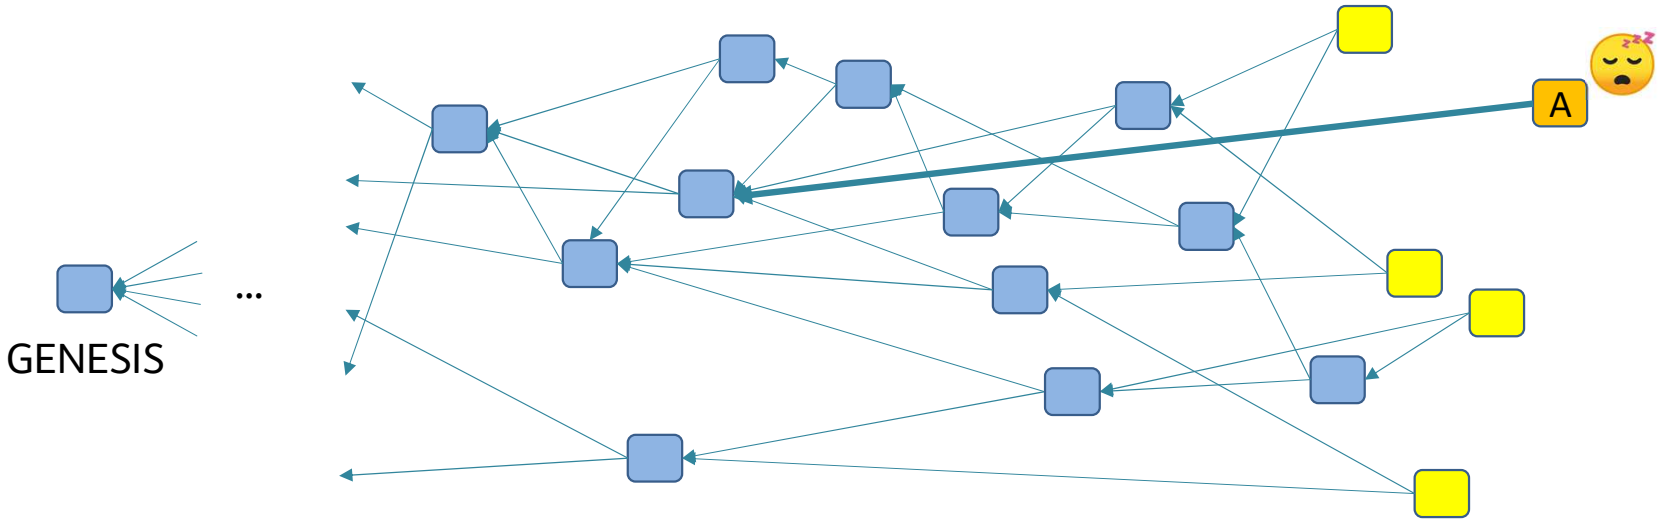

- Each new transaction must validate two existing transactions
- It is suggested that these transaction are tips, but it is not mandatory

- Selfish users may choose to validate old transactions ("lazy tips" issue, see A)
- No verification effort required, useless to help the tangle to grow (no new transactions are validated)
- Lazy tips must be discouraged!

Tip Selection

- How to discourage lazy tips?
  - Define a Tip Selection Strategy such that lazy tips are unlikely to be later validated
  - This way such tips will have low confirmation confidence and are thus penalized
- No guarantee that users follow any specific strategy
- But the one implemented in the reference code is likely to be dominant
- It is expected that the vast majority adopt it
- Two strategies proposed in IOTA white paper
  - Uniform Selection
  - (Unweighted / Weighted) Random Walk

Uniform Random Tip Selection

- Tips are uniformly randomly chosen
- Lazy tips are not penalized (A is as likely to be chosen as B, C, D and E)
- Users are encouraged to adopt selfish behaviour

→ Bad strategy choice!

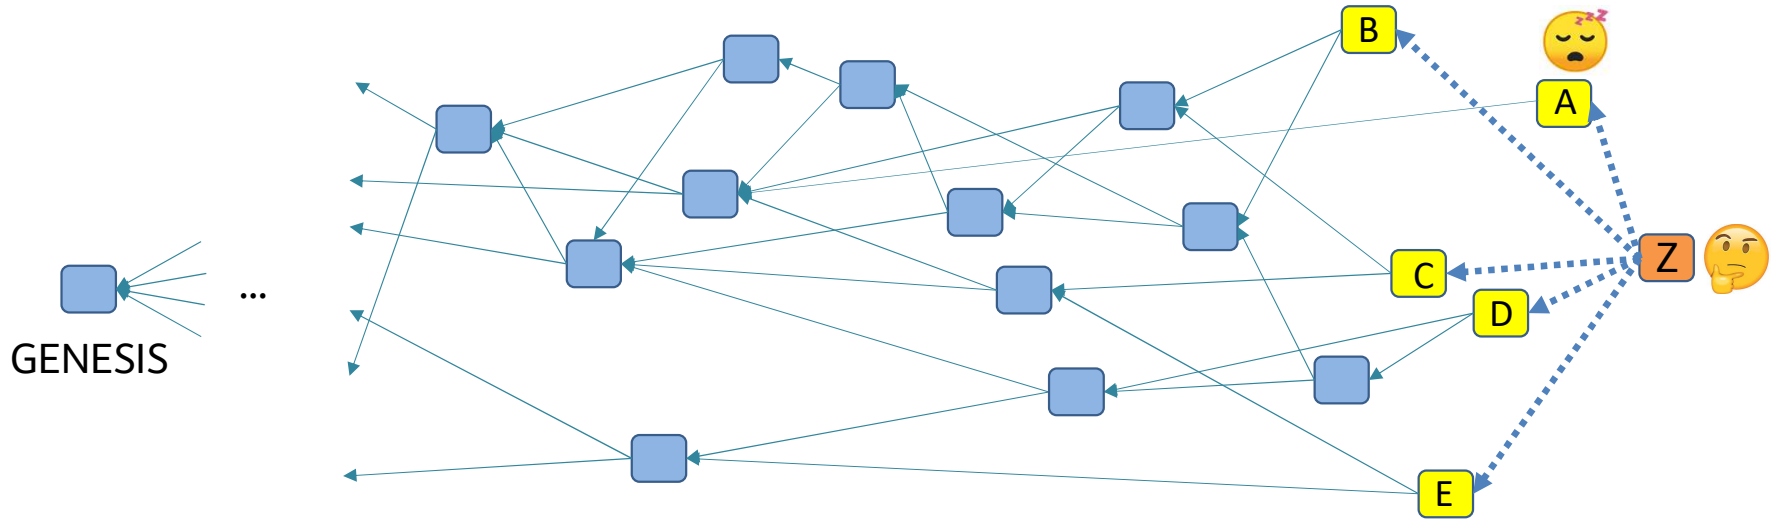

Random Walk

- Start from the Genesis
- Follow a random walk over the tangle through transactions

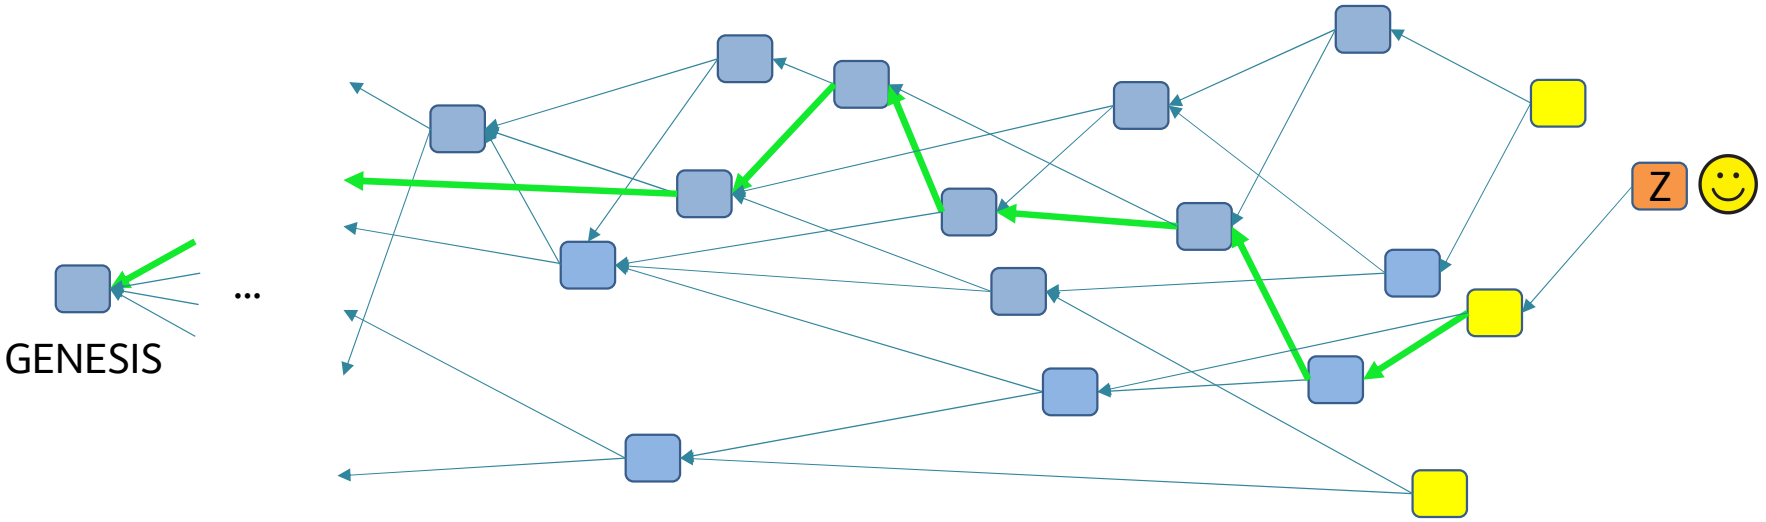

- Transition from X to Y is possible if and only if X is validated by Y
- When a tip is reached, select the tip and stop

Cumulative Weight

Given a transaction X, a (somehow defined) **weight (X)** is associated to X

$$\textbf{cumuluative weight}\ (\mathrm{X}) = \mathrm{weight}(\mathrm{X}) + \sum_{z} \mathrm{weight}(\mathrm{z})$$

$$\mathrm{z} \in Z = \{\text{set of transactions directly or indirectly validating x}\}$$

$$\text{if we set } \mathrm{weight}(\mathrm{X}) \stackrel{\text{def}}{=} 1$$

cumulative weight (X)
=
1 +
#{transactions directly or indirectly validating X}

Cumulative Weight

Cumulative weight of X = 1 + #{transactions directly or indirectly validating X}

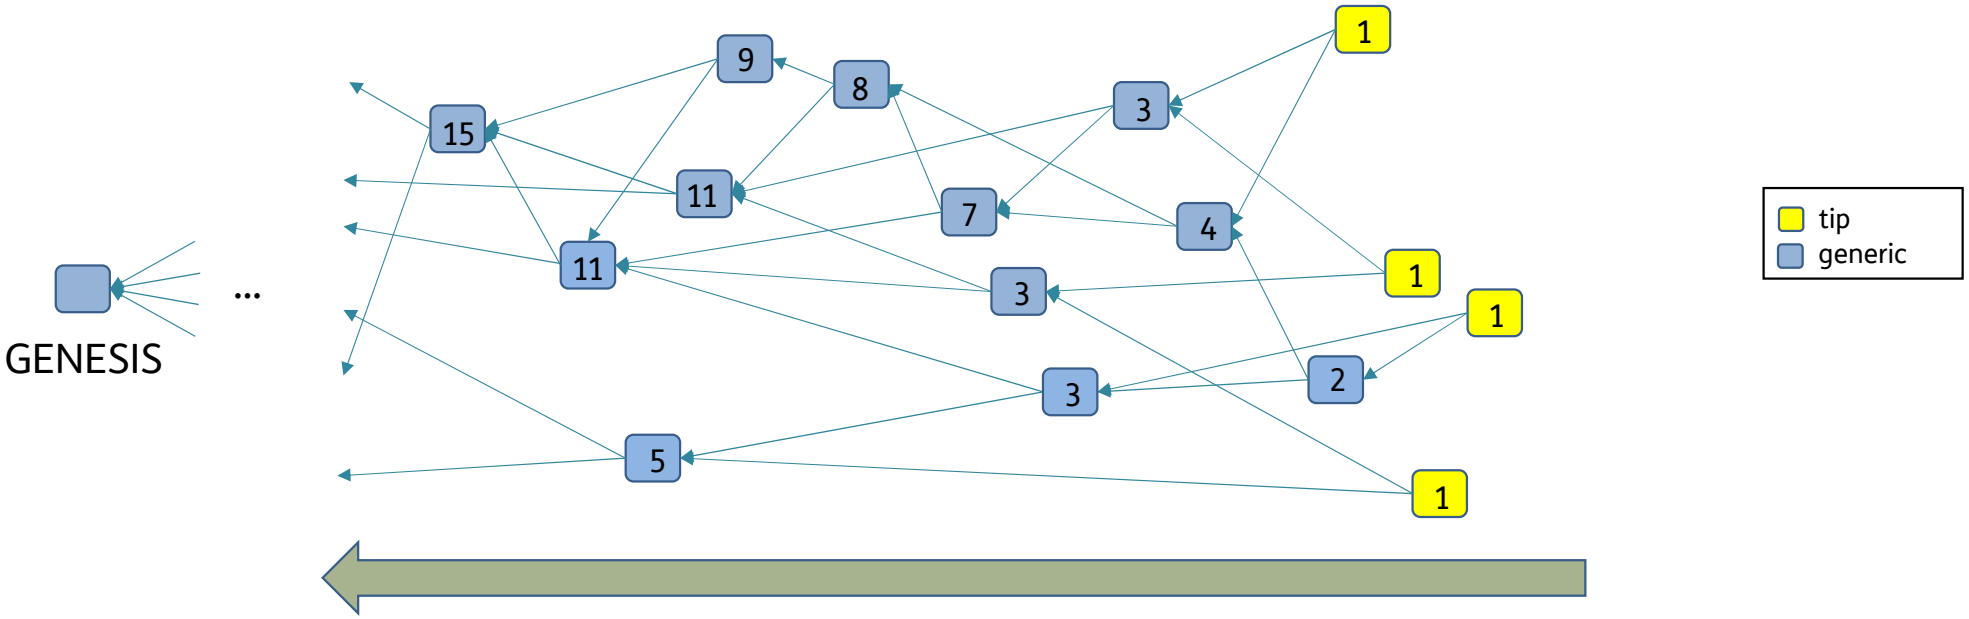

The deeper in the tangle, the higher the cumulative weight

Weighted Random Walk

Transition probability from X to Y

$$P_{XY} = \frac{e^{-\alpha(H_X - H_Y)}}{\sum_z e^{-\alpha(H_X - H_Z)}}$$

H_x= cumulative weight of x
z ϵ Z = {set of transactions directly validating x}
α = Random Walk parameter

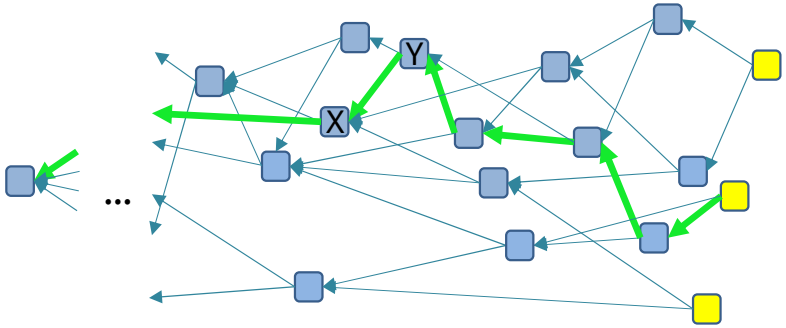

Weighted Random Walk

Transition probability from X to Y

$$P_{XY} = \frac{e^{-\alpha(H_X - H_Y)}}{\sum_z e^{-\alpha(H_X - H_Z)}}$$

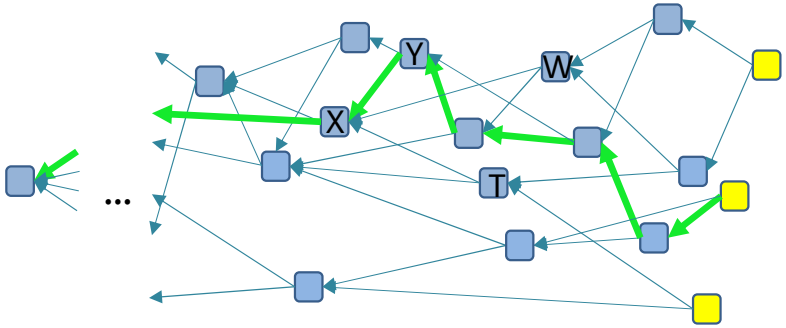

H_x= cumulative weight of x
z ϵ Z = {set of transactions directly validating x}
α = Random Walk parameter

	α=0	α=0.5	α=1	α=2
P_{XY}	0.333	0.821	0.976	0.9996
P_{XW}	0.333	0.111	0.018	0.0003
P_{XZ}	0.333	0.067	0.006	0.0001

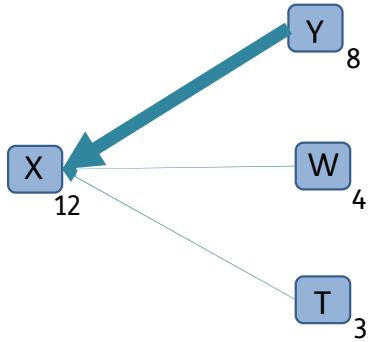

If α=0, all the transitions are equally likely
as α grows, transitions towards transactions with highest cumulative weight tend to probability one

(Unweighted) Random Walk: α=0

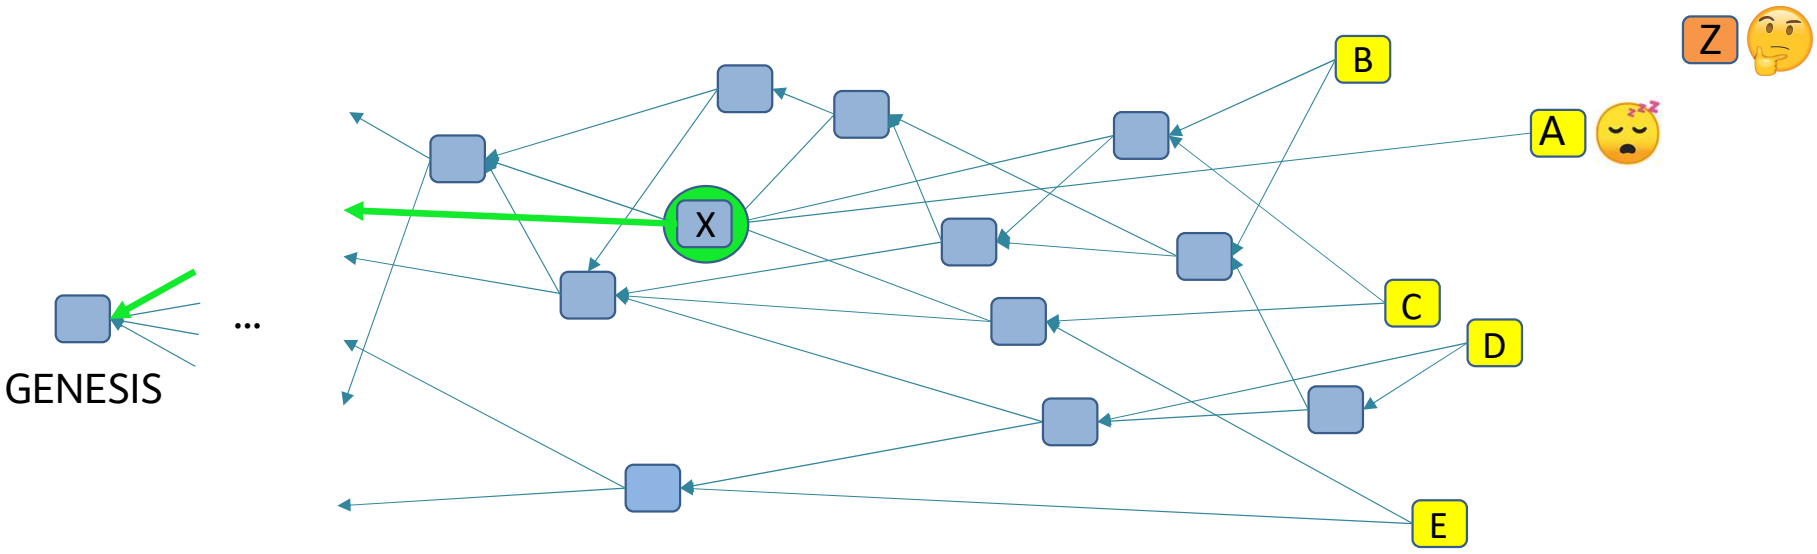

A new transaction Z needs to select tips to validate
A is a lazy tip

(Unweighted) Random Walk: α=0

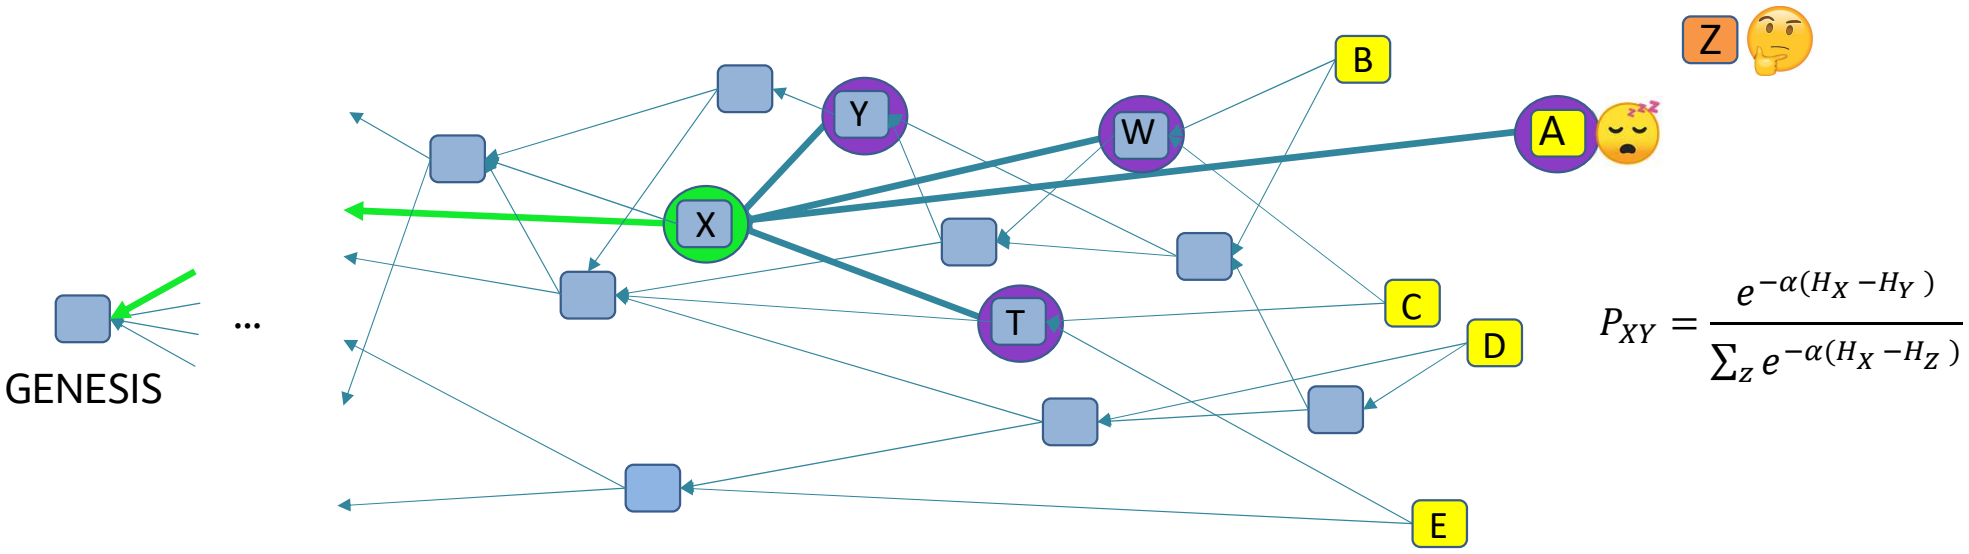

$$P_{XY} = \frac{e^{-\alpha(H_X - H_Y)}}{\sum_z e^{-\alpha(H_X - H_Z)}}$$

(Unweighted) Random Walk: α=0

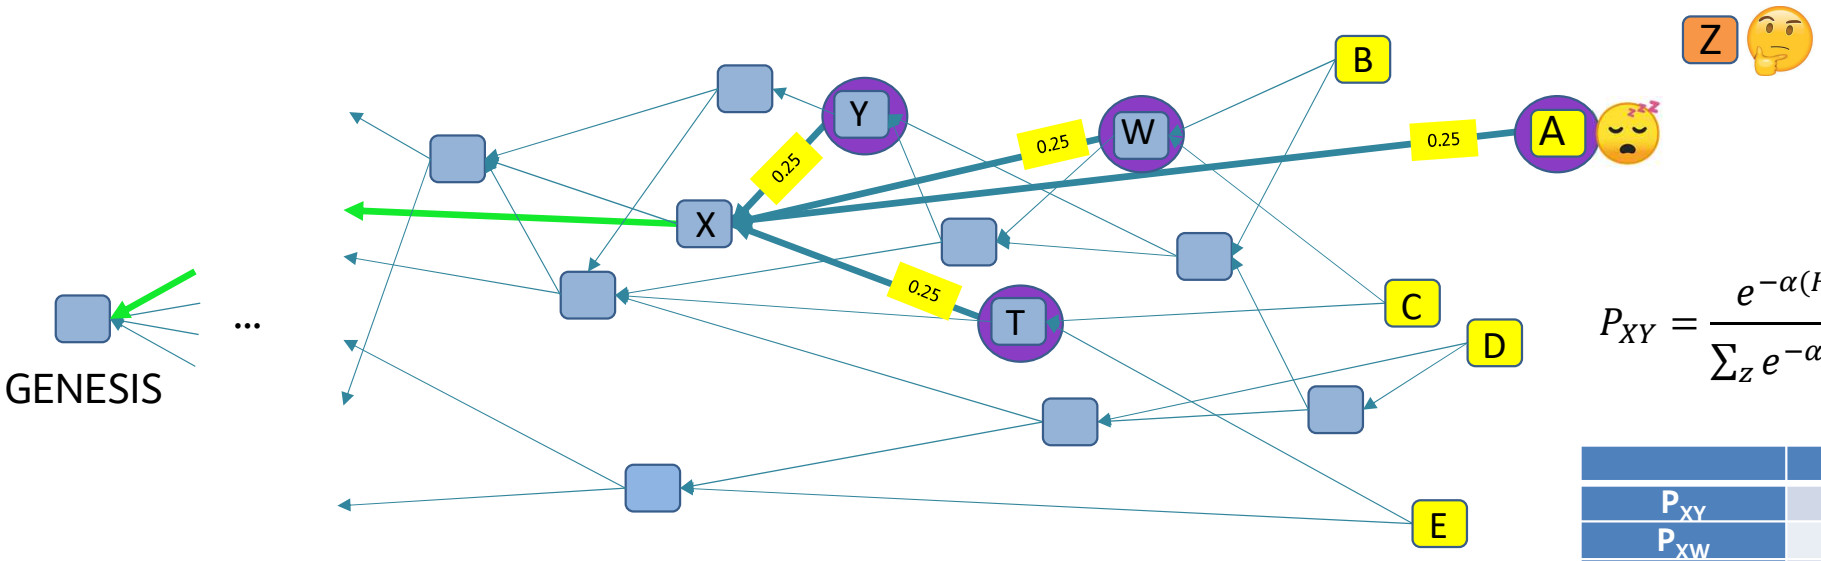

$$P_{XY} = \frac{e^{-\alpha(H_X - H_Y)}}{\sum_z e^{-\alpha(H_X - H_Z)}}$$

	α=0
P_{XY}	0.25
P_{XW}	0.25
P_{XT}	0.25
P_{XA}	**0.25**

(Unweighted) Random Walk: α=0

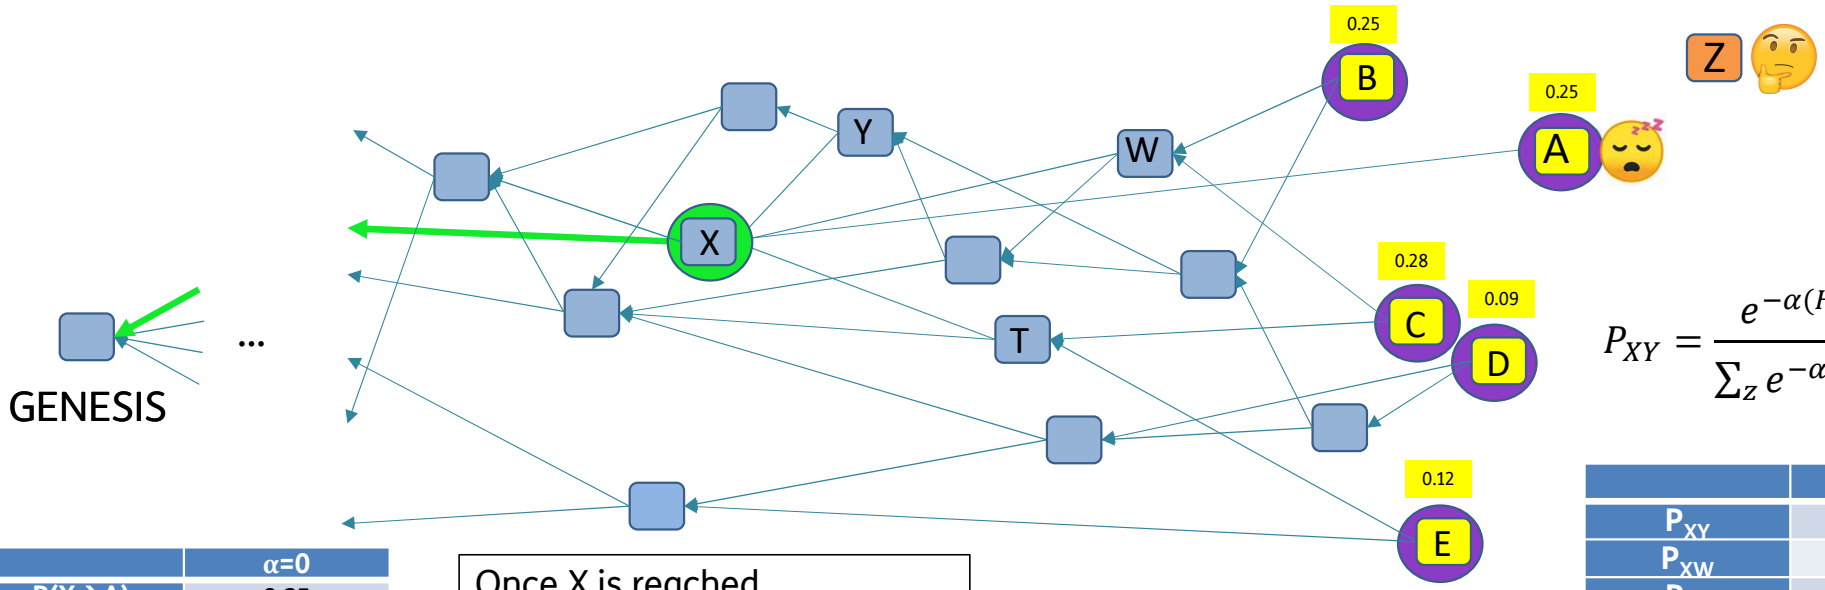

$$P_{XY} = \frac{e^{-\alpha(H_X - H_Y)}}{\sum_z e^{-\alpha(H_X - H_Z)}}$$

	α=0
P(X→A)	**0.25**
P(X→B)	0.25
P(X→C)	0.28
P(X→D)	0.09
P(X→E)	0.12

Once X is reached
average P(X→tip)=0.20
P(X→A)=0.25
A is not penalized by α=0

Note: of course B, C, D and E can also be reached through paths not including X. The deeper is X, the more likely is A to be selected as a tip

	α=0
P_{XY}	0.25
P_{XW}	0.25
P_{XT}	0.25
P_{XA}	**0.25**

Weighted Random Walk: α>0

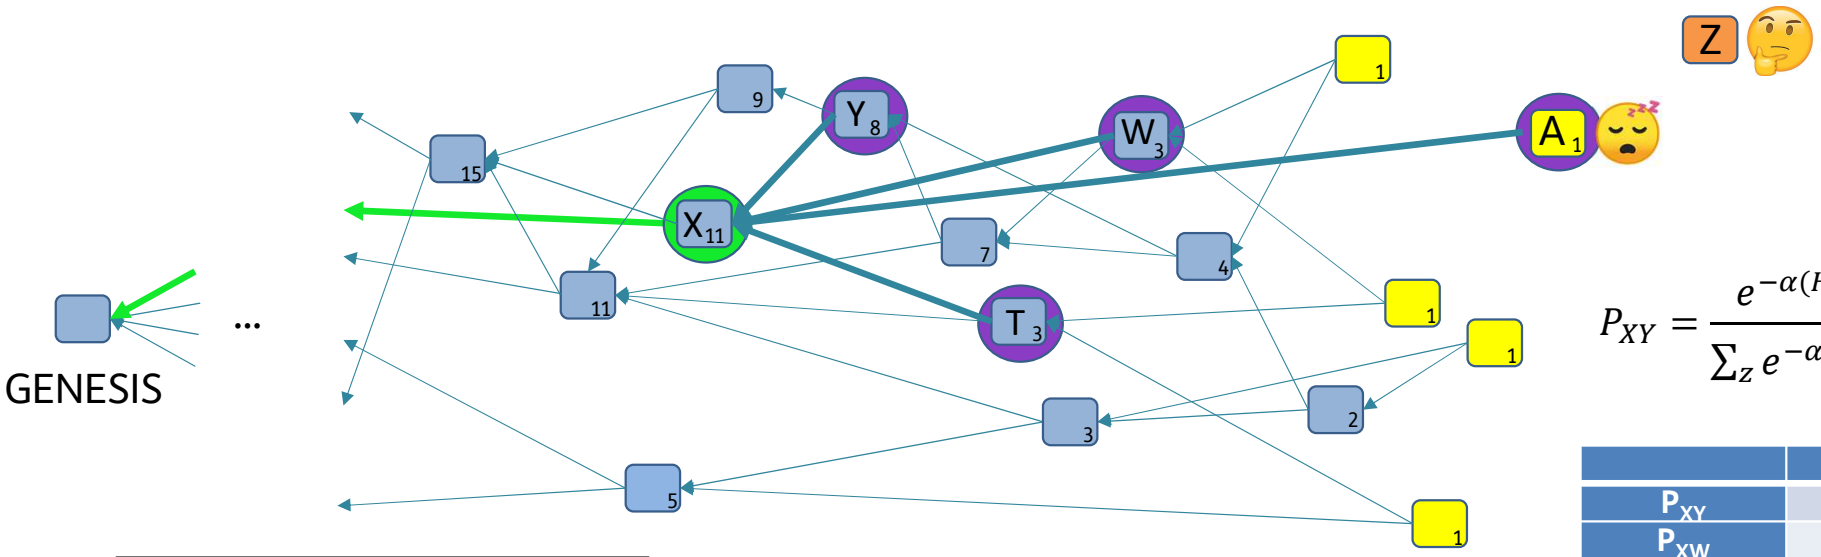

$$P_{XY} = \frac{e^{-\alpha(H_X - H_Y)}}{\sum_z e^{-\alpha(H_X - H_Z)}}$$

	α=1
P_{XY}	0.9858
P_{XW}	0.0066
P_{XT}	0.0066
P_{XA}	**0.0009**

Priority is given to tips on paths with high cumulative weight transactions

→

Lazy Tips have small cumulative weight
→ unlikely to be selected
→ **strongly disincentivized**

Weighted Random Walk: $\alpha>0$

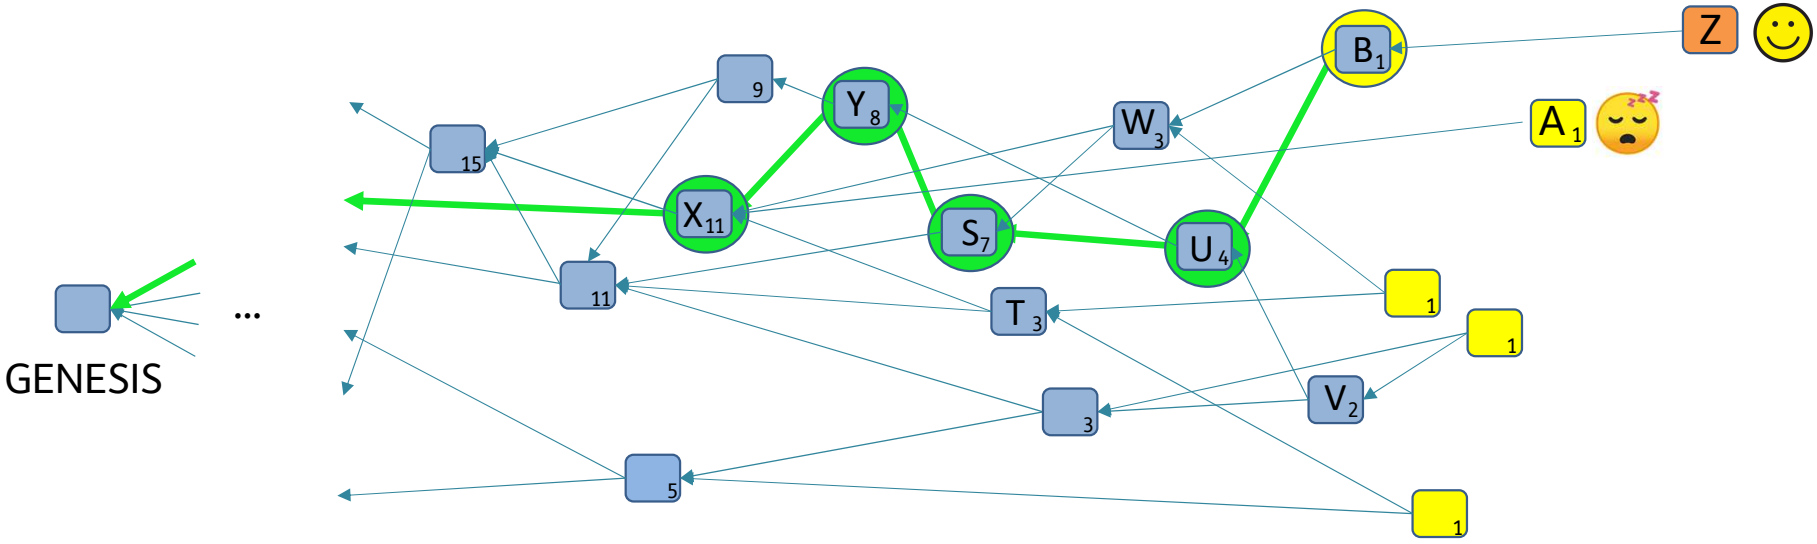

Weighted Random Walk: α>0

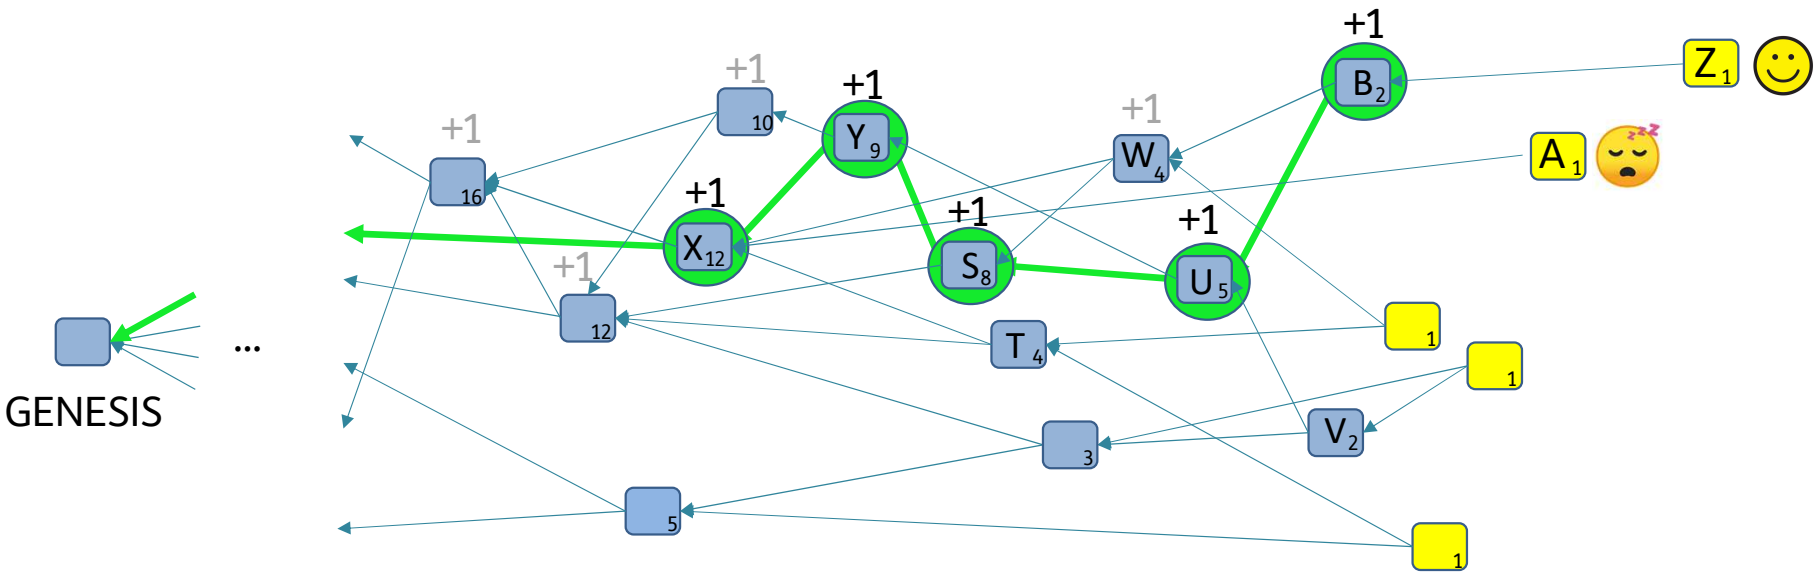

The weights of all the transactions in the Z cone increase by 1
More likely (higher weight) paths tend to become even more likely
Less likely paths (lower weight) tend to become even less likely

Weighted Random Walk: α>0

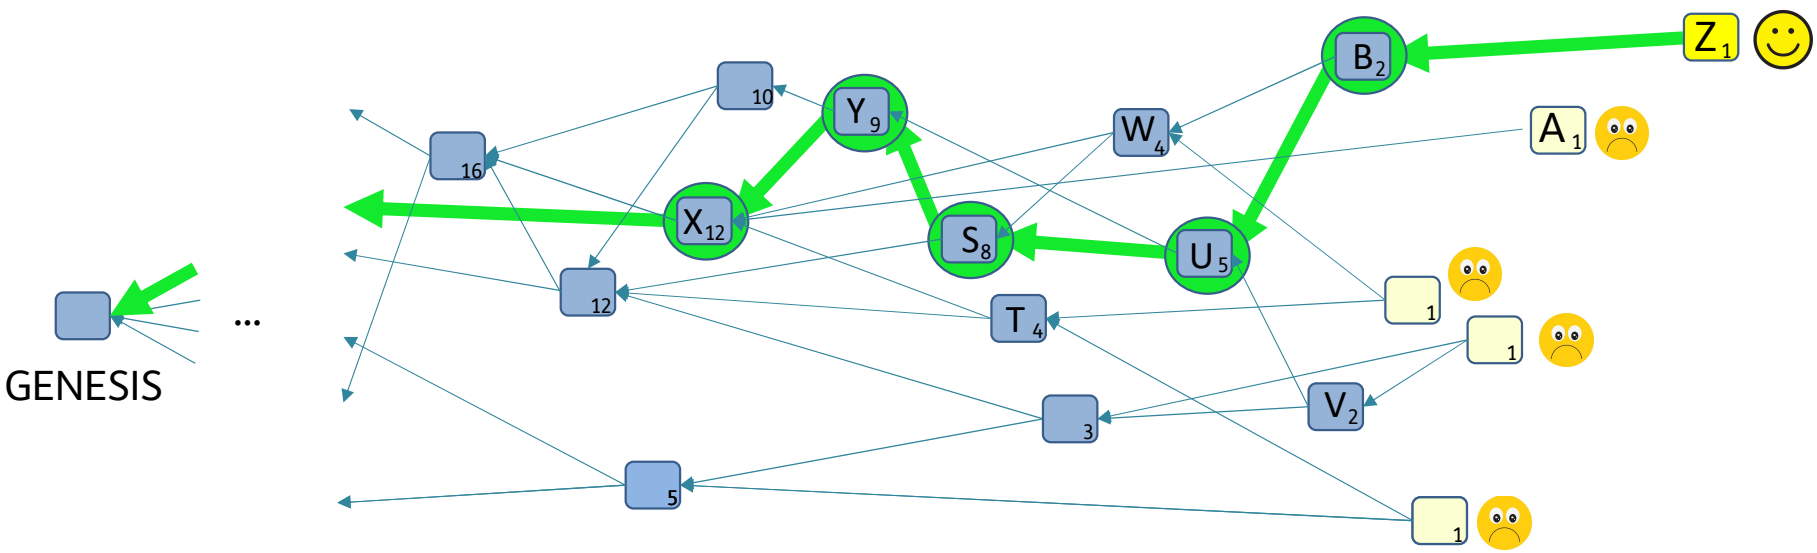

The weights of all the transactions in the Z cone increase by 1
More likely (higher weight) paths tend to become even more likely
Less likely paths (lower weight) tend to become even less likely

Some tips are left not validated

Weighted Random Walk: α>0

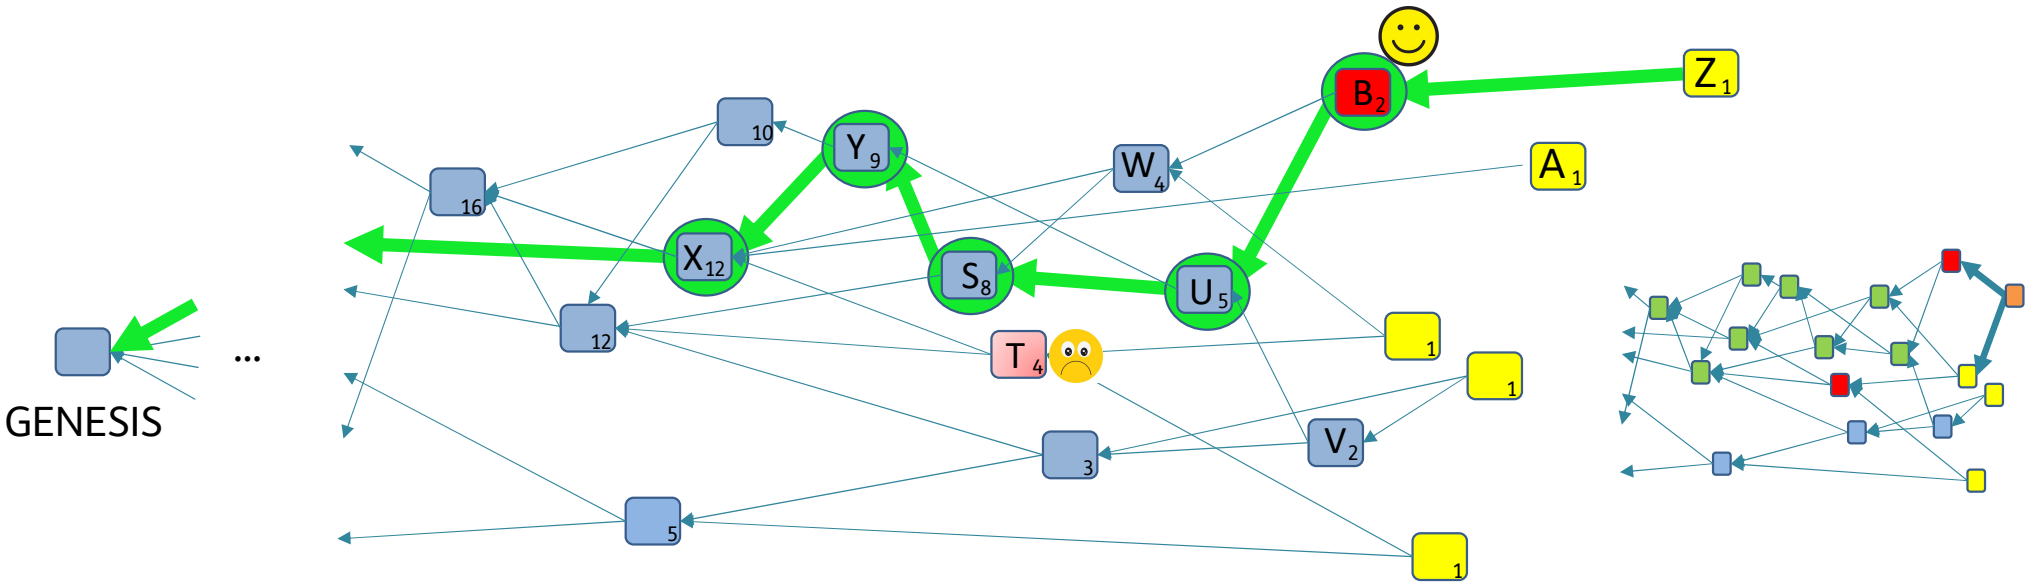

The weights of all the transactions in the Z cone increase by 1
More likely (higher weight) paths tend to become even more likely
Less likely paths (lower weight) tend to become even less likely

Some tips are left not validated
But consensus is helped
B (or T) will quickly prevail

Weighted Random Walk Comparison

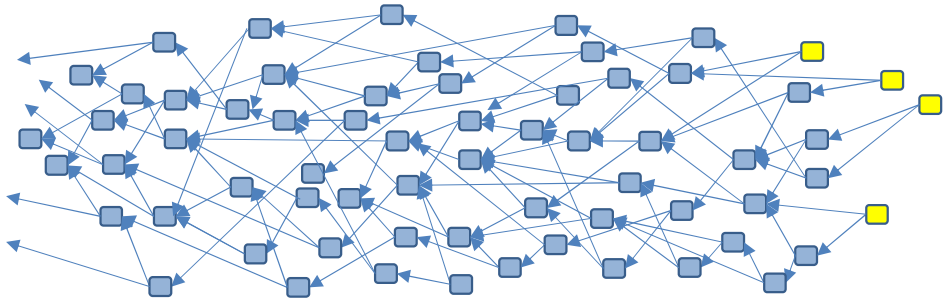

Low α:

☹ Lazy tips not discouraged

☺ New tips easily inserted

Finding the optimal α is a research problem

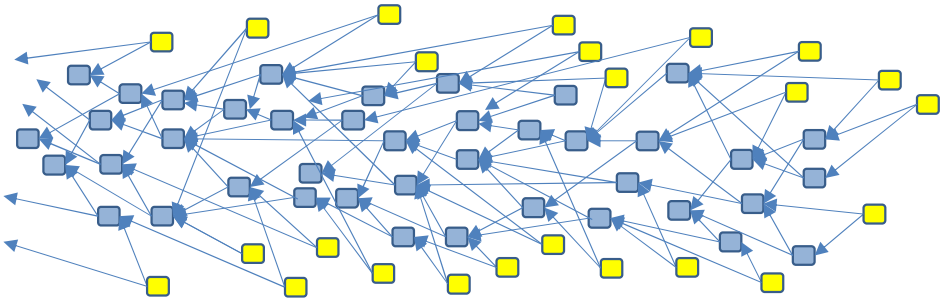

High α:

☺ Lazy tips discouraged

☹ Many tips lost

Attacks

- Many different (similar) attack models
  - Large weight attack
  - Parasite chain
  - Splitting attack
  - ...
- All of them eventually aim to double spend
- A correct mix between network access and computation power is required for the attacker
- In general «51% attacks» cannot be avoided

Basic Attack

The attacker

- creates a transaction T1 and waits until it is accepted (i.e. “spent” in the real world)
- creates a new double-spending transaction T2, causing two branches in the tangle
- issues many transactions validating **only T2**

If the attacker is poweful enough

- T2 branch will have cumulative weight higher than T1 branch
- T2 branch will become the main branch in the tangle and T1 branch will be discarded

Otherwise Random Walk Tip Selection will prevent the attack

- the attacker is not able to make T2 cumulative weight grows quicker that T1’s one
- when a bifurcation between T1 and T2 is found, T1 is thus likely to prevail

Cryptography

- **Private and public keys** (addresses) derived from seeds (~384 bits long)
- **Signature algorithm**
  - Winternitz One-Time Signature Scheme
  - Hash based → Quantum Resistant
  - Reveals a key portion each time a signature is published
  - Address reuse compromises private key
    - Each address can be used only once to withdraw iotas!
- **Hashing algorithm**
  - Curl (original scheme): vulnerable to a differential cryptanalysis attack
    - Used for PoW
  - Kerl (more conservative scheme, based on Keccak)
    - Used for address generation and signature creation/verification

The Coordinator

current IOTA traffic quite low

a single attacker may gain enough power to subvert the tangle

A coordinator is needed to make the tangle secure

- The coordinator is a closed-source special node run by IOTA Foundation
- The coordinator regularly issues special transactions known as **milestones**
- Milestones are assumed as 100% confirmed by all nodes

Coordicide in roadmap!

Trits and trytes

- Number of iotas

  2.779.530.283.277.761 = (3^33-1)/2

- Information basic unit is the trit {-1, 0, 1}
- Trits are encoded in trytes
- 1 tryte = 3 trits
- 1 tryte in {A, B, C, ... Z} U {9}: 27 possible values

A	(1, 0, 0)	B	(-1, 1, 0)	C	(0, 1 0)
D	(1, 1, 0)	E	(-1, -1, 1)	F	(0, -1, 1)
Y	(1, -1, 0)	Z	(-1, 0, 0)	9	(0, 0, 0)

Transaction in the Tangle

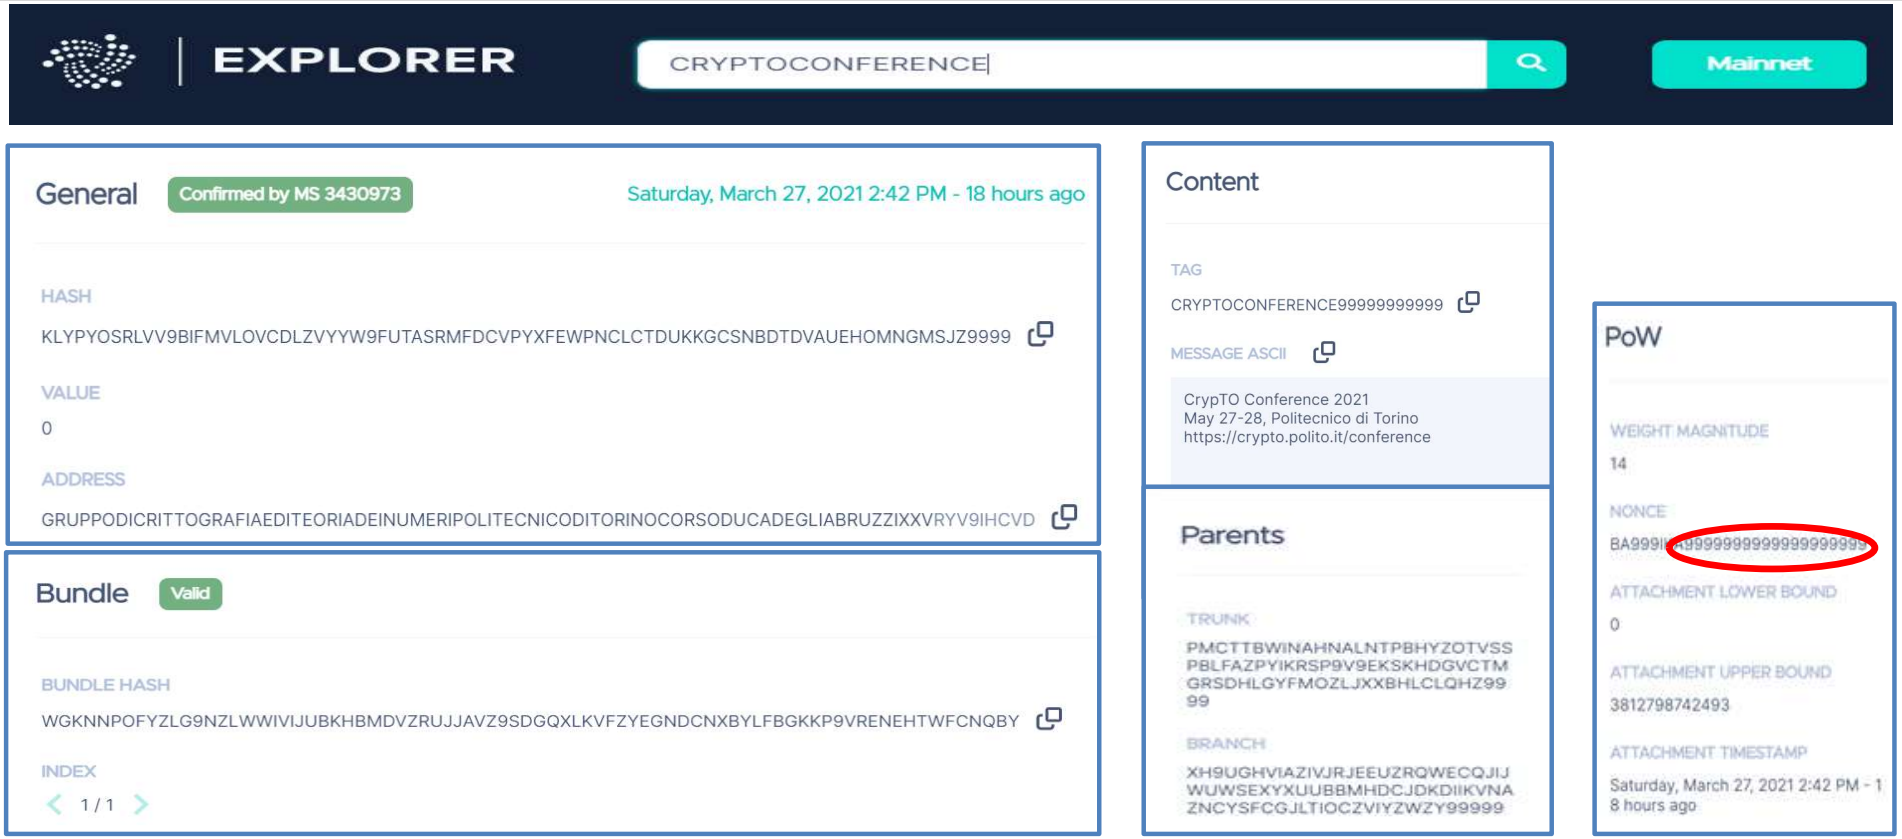

Roadmap

IOTA 1.0

(as described)

29th April 2021

IOTA 1.5

(Chrisalis)

Binary representation

EdDSA signature

Atomic transactions

UTXO model

…

?

2nd June 2021
DevNet
Nectar

IOTA 2.0

(Coordicide)

Coordinator removal

Open Issues

- **Scalability and decentralization**
  - Low power devices may be unable even to make signatures and PoW. Need of a *trusted proxy?*
  - Most IoT devices are *light nodes*, unable to store the full tangle
  - Tangle management (tips selection and validation) is thus left to *full nodes*
  - Who will maintain full nodes (core to the system)?
  - If traffic (hopefully) grows, only most powerful nodes will be able to manage it
  - [today] coordinator-based (should be coordicided soon)

- **PoW**
  - May be an issue for low-power devices
  - Is it a hidden fee? If weight is allowed to be > 1 and is related to PoW, more powerful nodes will pay a higher fee to make their transactions more likely to be confirmed

Open Issues

- **Questionable theoretical model and assumptions**
  - Stationary Poisson distribution for incoming transactions
  - All the devices have the same computing power
  - Big number of less powerful IoT devices out-compute attacker's dedicated high-power machines

- **Scenario evolution**
  - Only time will tell if IOTA's underlying assumptions fit the actual scenario in the next years (decades?)
  - Hard to predict network behaviour as traffic grows
  - Current situation say little about the future (conditions are going to change a lot)

The End

Thank you!
Questions?

guglielmo.morgari@telsy.it
guglielmo.morgari@polito.it