# The Bitcoin protocol 

Balázs Kőműves

Falkstenen $A B$

Budapest, 30 May 2014

## Overview of the talk

I will try to explain how Bitcoin works under the hood, going into details as far as time and my limited knowledge allows. I will cover the Bitcoin protocol as of mid-2013 (so newer developments like SPV, HD wallets or the payment protocol will be excluded).

Short overview:

- overview of Bitcoin
- the blockchain
- structure of a block
- mining
- bitcoin adresses
- transactions
- scripts
- elliptic curves
- ECDSA


## High-level overview of Bitcoin

Bitcoin is a set of protocols and software which allows individuals and organizations:

- to send digital currency tokens to each other over the internet
- almost instantaneously (depending on confidence level, somewhere between 1 minute and 1 hour)
- cheaply (approximately 1-10 HUF per transaction at the moment, independently of the value of the tokens)
- pseudo-anonymously (more on this later)
- without the ability to cheat (uses digital signatures for proof of ownership)
- without any centralized control (distributed protocol)
- with a prefixed supply curve of said tokens


## Low-level overview of Bitcoin

Bitcoin is a decentralized, peer-to-peer digital currency protocol.

- bitcoin addresses correspond to private/public key pairs (more precisely, addresses are hashes of public keys)
- addresses can hold amounts of coins (varying in time)
- transactions spend the output(s) of previous transaction(s)
- a transaction specifies how its output(s) can be spend in the future
- typical transcation: send the coins to one or more address(es)
- the spending condition: the next spender(s) must prove that they own the private key(s) corresponding to these addreses
- transactions are stored in the public ledger (which is a decentralized database)
- it's the the set of transactions which is fundamental - the "coins" are just invariants.


## Public key cryptography

Each party has a pair $(d, Q)$ of corresponding keys: $Q$ is public (say, on published on their homepage) and $d$ is private (typically only a single person knows it).

Main applications:

- Estabilishing a shared secret without meeting (key exchange)
- Send messages which only the intended recepient can decrypt (encryption)
- Prove that a message was really written by the person who claims it (signature)

From these basic building blocks, a huge set of really interesting applications can be built. Bitcoin only uses the signature algorithm.

Public-key crypto is widely used on the internet today: HTTPS, SSH, TLS, PGP, Bitcoin, OTR messaging, etc...

## Bitcoin transaction example, part I.

I expect income from two sources:

- I sold a book to a Alice for 0.3 BTC
- I made a bet with Bob, and won 0.2 BTC

I generate two key-pairs $\left(d_{1}, Q_{1}\right)$ and $\left(d_{2}, Q_{2}\right)$ and derive the corresponding addresses $A_{1}$ and $A_{2}$. Give the address $A_{1}$ to Alice and the address $A_{2}$ to Bob. They send me the agreed amounts. Two transactions $T_{1}$ and $T_{2}$ will appear in the ledger:

- $T_{1}$ says that 0.3 BTC can be spent by anybody who can prove that they know $d_{1}$
- $T_{2}$ says that 0.2 BTC can be spent by anybody who can prove that they know $d_{2}$

This means that I "own" 0.5 BTC alltogether (since nobody else knows $d_{1}$ or $d_{2}$ ).

Next, let's say I want to buy a bottle of whisky for 0.4 BTC!

## Bitcoin transaction example, part II.

To send 0.4 BTC for the address of the whisky shop $A_{w}$, I need to craft a transaction $T$ which "spends" the outputs of the previous transactions $T_{1}$ and $T_{2}$, sends 0.4 BTC of it to the whisky shop, and sends the remaining 0.1 BTC back to me (usually to a freshly generated address $A_{3}$ )


The transaction $T$ is signed with the private keys corresponding to the addresses $A_{1}$ and $A_{2}$, to prove that they belong to me.

## The blockchain

The public ledger, called the blockchain in Bitcoin-slang, is a distributed, (mostly) append-only database.

- the blockchain is built up from blocks; on average 1 new block is produced and appended every 10 minutes
- each block contains a small header, and a set of transactions
- the blocks are organized in a linked list (chain), starting from the "genesis block"
- each block is produced by a single participant "node"; to decide which node will produce the next block, the protocol uses a hash based lottery, with winning chance proportional to CPU-power (this process is usually called "mining", which is rather confusing)
- people wanting to transact simply broadcast their (signed) transactions to the nodes
- to incentivise nodes to participate in this process, a reward is given to the winner, and they also get all the transaction costs for the set of transactions they include in that block
- typically, each participant node has a full copy of the database (currently around 18 gigabytes)


## A block

A block consists of a block header (80 bytes), plus a set of transactions.

The header consists of:

- block version number (4 bytes; currently 0x00000002, little-endian)
- hash of the previous block's header (32 bytes)
- the Merkle root of all transaction in this block (32 bytes)
- a unix timestamp (4 bytes)
- the current difficulty target (4 bytes)
- a nonce (4 bytes)
- (number of transactions (varint) - not actually part of the header)

A block is considered valid, if:

- all the previous blocks in the chain are valid
- the hash of the block header is smaller than the current target
- all the individual transactions are valid, and the Merkle root matches


## Merkle trees

A Merkle tree is a hash tree. Bitcoin uses double-SHA256 as its primary hash function, that is, hash=SHA256(SHA256(msg)).

- for each transaction, a transaction hash is computed
- a binary tree is built, such that the leaves are labeled with the hashes of the individual transactions
- rows with odd number of elements are extended to be even by duplicating the last hash
- for each node, the hashes of its two children are concatenated, and the node is labelled by the hash of this string
- the Merkle root is the hash at the root node

If the set of transactions changes any way, the Merkle root also changes (and it changes unpredictably).

## The lottery (or "mining", or proof-of-work)

To be able to decide who produces the next block, a lottery process is used (confusingly called "mining"):

- the Merkle tree of the set of current transactions is computed
- the first transaction must be always the so-called coinbase transaction (or generation), which gives a reward to the miner
- thus the Merkle root will be different for all participants
- these days, an extra nonce is also contained in this transaction
- the nonce (and the extra-nonce) is set to whatever (typically increasing numbers for simplicity)
- the hash of the block header is computed
- if the hash is smaller than the current target, we won, and broadcast the new block to the network
- othewise, new nonces are set and the hash is computed again

If there are conflicting blocks (winners), the chain with the most "work" is chosen. Abandoned blocks are called orphan blocks.
( $\rightarrow$ standard practice of waiting for 6 confirmations to be safe)

## The difficulty (or target)

The difficulty is a number (approx. $10^{10}$ now), which encodes a 32 byte long target. The block hash (which can be considered random for all practical purposes) must be below this target for a block to be valid.

Thus the difficulty controls that on average how often there is a winner of the lottery. To have a balanced stream of blocks (in time), this difficulty is adjusted every two weeks (more precisely, every 2016 blocks), based on the number of the blocks in the previous two weeks, so that on average there are 10 blocks (winners) per minute.

While in principle a miner could use a smaller difficulty (higher target) than the others, the protocol says that in case of competing chains (branches), one should choose the one with "more proof-of-work"; more-or-less, this means the one which contains blocks with higher summed difficulty. There is also a standard formula for adjusting the difficulty, but it is done individually by each client.

Note: The difficulty is encoded in the block headers in 4 bytes using a custom exponential encoding.

## The blocks on wire (or disk)

If you want to write a blockchain parser, you need to download the blockchain in some format. The official client simply stores the blocks on the disk (in the files blk00 NNN. dat) more-or-less as they come on the wire, which is:

- magic bytes 0xf9beb4d9 (4 bytes, big-endian)
- length of the block (4 bytes, little-endian)
- the block data as shown before:
- header (80 bytes)
- number of transactions (varint)
- the transactions

Unfortunately, sometimes there is a random number of zero bytes between the blocks. I couldn't find any pattern of where, why, and how many zero bytes are there...

## Bitcoin addresses

Bitcoin addresses correspond to private/public keypairs. "Owning" an address is equivalent to knowing the private key. A user can have any number of addresses, and an address can hold any number of "coins".

Recall that an ECC private key is a random number $d \in[1, n-1]$, and the corresponding public key is the point $Q=d G \in \mathcal{G}$ on the elliptic curve $\mathcal{G}$, where $G \in \mathcal{G}$ is the fixed generator. This point can be represented by its affine coordinates $d G=(x, y) \in \mathcal{G} \subset \mathbb{A}^{2} \subset \mathbb{P}^{2}$.

Because of the curve equation $\mathcal{G}=\left\{(x, y): y^{2}=x^{3}+7\right\}$, the $y$ coordinate can be computed (up to parity) from the $x$ coordinate, giving another representation consisting of $x$ and the parity of $y$.

Unfortunately, because of historical reasons, both versions are used on the Bitcoin network, thus there are two addresses corresponding to a single private key.

## Construction of Bitcoin addresses

- a representation of the public key is chosen, and encoded as follows
- the coordinates are encoded as 32 bytes big-endian integers
- full public key: 0x04 ( $x$ coord) ( $y$ coord)
- compr. pub. key, $y$ even: $0 \times 02$ ( $x$ coord)
- compr. pub. key, $y$ odd: 0x03 ( $x$ coord)
- compute hash=RIPEMD160(SHA256(pubkey))
- add a version byte in the front, which is $0 \times 00$ for the main Bitcoin network ( $0 \times 6 \mathrm{f}$ for testnet3; 0 x 05 for pay-to-script-hash addresses)
- compute SHA256(SHA256(0x00||hash)), and discard everything except the first 4 bytes; this will be the checksum.
- create the string ( $0 \times 00| | h a s h \|$ checksum)
- convert this to an ASCII string using Base58 encoding.

Note: This process can be easily inverted to recover the hash of the public key, but not the public key itself.

## Transactions

A bitcoin transaction consists of several inputs and several outputs. Inputs refer back to a selected output of an earlier transaction, which this transaction spends; and outputs are conditions on how to spend them. The creator of a transaction must prove, for each individual input, that he is allowed to spend them.

The most often used condition, called pay-to-address, is that the spender must prove that he owns a given address (that is, he knows the corresponding private key).

Another useful transaction is the generation or coinbase transaction, which is always the very first transaction in a block, and serves dual purposes:

- it is a reward for the miner for providing the service of maintaining the network
- but more importantly, it is a way to "mint" coins. So this is how the supply of bitcoins is created and distributed.

There are a few more standard transaction types, and in fact a full scripting language to specify spending conditions, though most of it is disabled (because of security fears).

## A randomly chosen transaction (source: http://blockexplorer.com)

Hash? ${ }^{?}$ 4elab2affa72a7393515383a82e60342abb3211c8bbaec8b1c7a6b5572aafb93
Appeared in block 301742 (2014-05-20 13:01:22)
Number of inputs ${ }^{2}: 3$ (Jump to inputs)
Total BTC in ${ }^{2}: 0.15690866$
Number of outputs: 2 (Jump to outputs)
Total BTC out ${ }^{2}: 0.15680866$
Size ${ }^{?}: 618$ bytes
Fee ${ }^{?}: 0.0001$
Raw transaction?

## Inputs?

| Previous output (index)? | Amount ${ }^{\text {? }}$ | From address ${ }^{\text {? }}$ | Type ${ }^{\text {? }}$ | ScriptSig ${ }^{\text {? }}$ |
| :---: | :---: | :---: | :---: | :---: |
| 8f97f1623e76 . 1 | 0.00311054 | 1M177PRYquDJ2Ar6anL 1qKhhRyabwUi9GV | Address | 3046022100d8e3975d9a397284f86faaaa2b6d1 046d8f773f4702272aa7629b7b8cec4949348a6 $\square$ $+$ |
| eea00af26d1b... 1 | 0.04453239 | 1M177PRYquDJ2Ar6anL 1qKhhRyabwUi9GV | Address | 3044022047elbc2c70e34971613db59af6f06ad 046d8f773f4702272aa7629b7b8cec4949348a6 $\square$ $\stackrel{\rightharpoonup}{2}$ |
| e50edce98c3c... 0 | 0.10926573 | 1M177PRYquDJ2Ar6anL 1qKhhRyabwUi9GV | Address | 3045022100cd998cc94904fbf73a01cbbff5e037, 046d8f773f4702272aa7629b7b8cec4949348a6 $\square$ III |

## Outputs ${ }^{\text {? }}$

| Index $^{2}$ | Redeemed at input $^{?}$ | Amount $^{?}$ | To address ${ }^{?}$ ? | Type $^{?}$ | ScriptPubKey? |
| :--- | :--- | :--- | :--- | :--- | :--- |

## A randomly chosen transaction (source: http://blockexplorer.com)

```
    "hash":"4e1ab2affa72a7393515383a82e60342abb3211c8bbaec8b1c7a6b5572aafb93",
    "ver":1,
    "vin_sz":3,
    "vout_sz":2,
    "lock_time":0,
    "size":618,
    "in":[
            {
        "prev_out":{
            "hash":"8f97f1623e768deff46829344fa8358c63074bffa3d1f848cb146ce0f1a44e8f",
            "n":1
            },
            "scriptSig":"3046022100d8e3975d9a397284f86faaaa2b6d1eaee7cde4d4c9690a53b4aefbc94677a956022100e92ac9ba2df27d9e4c9ce9;
        },
        "prev_out":{
            "hash":"eea00af26d1b6db69dcf29dd3c8b0aae7e4b5627d44b9a48d9ca04de65143d0a",
            "n":1
            },
            "scriptSig":"3044022047e1bc2c70e34971613db59af6f06ad6913d0b4590e1e83767f2a59f85bb4af002205e72ab9ed4cc9a7aca0b0ba83f:
        },
            "prev_out":{
                    "hash": "e50edce98c3cfa12eb45a16a351cf3924395a1788ab0ec7306eeea59fa15cf06",
                    "n":0
            },
            "scriptSig":"3045022100cd998cc94904fbf73a01cbbff5e037e796d0954cde6007e702b6e48345139de40220243ca524a8ad6933980fbac6:
        }
    ],
    "out":[
    {
        "value":"0.10654490",
            "scriptPubKey":"OP_DUP OP_HASH160 b3f290c116ad4d76dfbadb91713438b14d527567 OP_EQUALVERIFY OP_CHECKSIG"
        },
            "value":"0.05026376",
            "scriptPubKey":"OP_DUP OP_HASH160 db679dbb781de77866cd2113153544fd78d59f8c OP_EQUALVERIFY OP_CHECKSIG"
        }
    ]
}
```


## The structure of a transaction (Tx)

A transaction is basically defined by a list of inputs and outputs. It is encoded as follows:

- transaction version number (4 bytes; currently 0x00000001)
- number of $t x$ inputs (varint ${ }^{1}$ )
- list of $t x$ inputs
- number of tx outputs (varint)
- list of $t \times$ outputs
- the lock-time ${ }^{2}$ (4 bytes)

[^0]
## Transaction inputs (TxIn)

A transaction input refers back to a single output of a previous transaction (possibly in the same block), and proves that one is allowed to spend that output. It is encoded as follows:

- the hash of the previous tx we refer back to ( 32 bytes)
- the index of a specific output from that transaction (4 bytes, indexing starts from 0)
- the length of the signature script (varint)
- the signature script (often called "sigScript")
- the sequence number ( 4 bytes; typically 0 )

The signature script contains the proof that we are allowed to spend the given output. In the case of a pay-to-address previous tx , it will contain:

- the public key corresponding to the address we own
- the signature with the corresponding private key of a modifiction of this transaction (it must be modified to avoid self-signing, which is impossible).


## Transaction outputs (TxOut)

A transaction output simply specifies an amount of bitcoins, and a condition of how it can be spent. The most often used condition, called pay-to-address, specifies that the spender must own the address given here.

It is encoded as follows:

- transaction value (8 bytes, encoded as an 8 decimal digits fixed point number)
- length of the pubKeyScript
- the pubKeyScript (or pkScript) (containing the condition)

The spending condition is encoded as an arbitrary ${ }^{3}$ script, which, combined with the future signature script, must produce true for the spending tx to be vaid.

The sum of the input coins must be equal or more than the sum of the output coins; the difference is the transaction fee, which will go the miner who produces the block which includes the given transaction.

[^1]
## The coinbase (or generation) transaction

The first (and only the first) transaction in each block must the generation transaction. This is the only kind of transaction which has no (real) inputs. It serves dual purposes:

- to reward the miner who produced this block
- a constant reward, currently 25BTC, halving every 4 years
- plus all the transaction fees of the transactions included in the block
- to create new bitcoins

The reward is set so that this "minting" of new bitcoins diminishes over the years, and converges to $21,000,000$ bitcoins. Thus there will be never more 21 M bitcoins (unless this part of the protocol changes, which is really unlikely because none of the holders want that).

There is a field (called coinbase) containing arbitrary data; this is used for the extra-nonce, for political messages, the block height, and sometimes to vote whether miners accept a proposed protocol change.

## Script execution

The script language is a simple ${ }^{4}$, stack-based, non-Turing-complete language.

## Execution:

- first execute the scriptSig (= proof) of this transaction. This is not allowed to do anything else than push data to the stack (said data will contain the proof);
- then delete the secondary stack (which is typically not used anyway);
- finally execute the pkScript (= condition) of the previous transaction (whose output we spend).
If at the end the top of the stack contains true, we accept the proof, otherwise not.

Remark: Since each (non-coinbase) tx contains at least one input, which must contain a signature of this transaction ${ }^{5}$, third parties cannot modify transactions to redirect the outputs.

[^2]
## Pay-to-address script example

| step | stack after | opcode | description |
| :---: | :---: | :---: | :---: |
| 0 | (empty) | - | initialization |
| 1 | sig | PUSH <sig> | signature |
| 2 | sig \| pk | PUSH <pubkey> | pubkey |
| 3 | sig \| $\mathrm{pk} \mid \mathrm{pk}$ | DUP | duplicate pubkey |
| 4 | sig \| $\mathrm{pk} \mid \mathrm{addr}$ | HASH160 | compute address ( = hash of pubkey) |
| 5 | sig \| pk | addr | pkhash | PUSH <pkhash> | push the recepient address |
| 6 | sig \\| pk | EQUALVERIFY | check agreement |
| 7 | true | CHECKSIG | check signature |

## OP_CHECKSIG and signing transactions

The CHECKSIG opcode executes a rather involved procedure.

1. the public key and the signature is popped from the stack, and decoded (we fail if the encoding is invalid).
2. the so-called subscript is formed. Usually this is just pkScript ${ }^{6}$
3. extract the so-called sigHash byte from the signature (usually this is SIGHASH_ALL). Depending on the value of this, we may proceed diferently
4. replace all input scripts with empty strings, except the current input which is replaced with the subscript ( $\rightarrow \mathrm{tx}$ malleability)
5. append the sigHash byte, extended to a 4-byte little endian word, to the serialized, modified transaction
6. verify the signature against the (double-SHA256 hash of) this final bytestring.

The signing procedure is implied by the above checking procedure: The above is done for all inputs do derive all the signatures. For the rest of sigHash flag possibilities, the process is a bit different.

[^3]
## The SigHash byte

The sigHash byte controls how the outputs are signed. The above process is slightly modified based on this (for example, the outputs are also modified). This feature can be used for "smart contracts", which need interaction between different parties to create a transaction. The valid values of the sigHash byte are:

- SIGHASH_ALL $=0 \times 01^{7}$
- SIGHASH_NONE $=0 x 02$
- SIGHASH_SINGLE $=0 x 03$
- SIGHASH_ANYONECANPAY $=0 \mathrm{x} 80$ (flag)

ALL (the default) means that all outputs must be signed.
NONE means that none of the outputs should be signed - we don't care where the outputs go.
SINGLE means that only the output corresponding to the current input must be signed. ANYONECANPAY means that other people can extend this transaction with more inputs.

[^4]
## Key and signature encodings

Pubkeys are encoded as with the address: 33 or 65 byte, starting with $0 \times 02$, 0 x 03 or 0 x 04 (except that $0 \times 06$ also appears at some point instead of $0 \times 04 \ldots$ )

Signatures are $(R, S) \in \mathbb{Z}_{n} \times Z_{n}$ pairs encoded in ASN. 1 DER:


Except that

- there is an extra sigHash byte at the end
- DER encoding should be unique, but OpenSSL accepts some non-conforming encodings too, so you have to do that too ( $\rightarrow$ signature malleability)
- multisig sometimes looks a bit different

Unfortunately, a client must reproduce all past bugs of the official client (which it has a lot...), as people tested all the bugs by putting transactions in the blockchain. And when they sometimes fix one, you must branch on the block height when they fixed it.

Wallet Import Format: Base58 encoding form of the big-endian 32 byte private key, prepended with a version byte $0 \times 80$; except when the corresponding public key is in compressed form, then also add an extra $0 \times 01$ byte at the end

## Signing messages

The official client (was: Bitcoin-QT, now: Bitcoin-Core) allows signing of custom ASCII messages. This works a bit differently.

First, the message is encoded as $\left[n_{\text {magic }} \|\right.$ magic $\left.\left\|n_{\text {msg }}\right\| \mathrm{msg}\right]$, where the lengths are encoded with the Varlnt encoding, and magic is the string "Bitcoin Signed Message: $\backslash \mathrm{n}{ }^{\prime \prime} .^{8}$

Then the signature $(R, S)$, the pubkey format (compressed or uncompressed), the parity of $y$, and one more bit of extra information (whether $R=x$ ) is encoded in yet another custom encoding, which is then Base64 encoded (not Base58). The extra bits of information are necessary to recover the public key from the signature, so the other party needs to know only your address to check the signature.

The encoding is 1 byte containing the 3 bits (between $0 \times 1 \mathrm{~b}$ and $0 \times 22$ ) followed by 32 byte $R$ and 32 byte $S$ as big-endian integers. ${ }^{9}$

[^5]
## Multisig transactions

There can be transactions which need $N$ out of $K$ signatures to be valid ( $N \leq K$ ), called "multi-sig". This can be useful in various situations:

- two-factor authentication: with a 2-of-2 multisig transaction, you need both signatures to spend. If the two signing devices are physically separate, this adds extra security (it is not enough to hack one device).
- signing rights: similarly, a company may say that it needs 3 out of 5 executive's digital signature to spend some funds
- escrow services: If there is a seller and buyer on the internet who don't fully trust each other, they can use an escrow. However they both needs to trust the escrow. With a 2-of-3 multisig transactions, if there is no dispute, the two parties can sign the transaction; if there is dispute, the escrow can decide and sign it together with the party who is right.

This uses the CHECKMULTISIG opcode, which is even more complex (and contains some new bugs). Usually to funds are sent to a multisig address ${ }^{10}$ first, which can be spent only with the necessary number of signatures.
${ }^{10}$ this is pay-to-script-hash (P2SH) address

## Pay-to-script-hash (P2SH)

This is an ugly hack to enable reversion of control. It introduces a

- a new address type (starts with ' 3 '; version byte is $0 \times 05$ ) (BIP 13)
- a new "standard transaction type" and a special behaviour for transactions matching that template (BIP 16)
In the address, instead of the 20 -byte hash of the public key, we use the 20 -byte hash of the script which will specify the condition to spend.

The new standard transaction script looks like

$$
\begin{aligned}
\text { pkScript } & =\text { OP_HASH160 [20-byte-hash-value] OP_EQUAL } \\
\text { sigScript } & =\text { [signatures }]\{\text { serialized script }\}
\end{aligned}
$$

If a transaction matches the above pattern, the following special behaviour is invoked:

- sigScript is executed (fails if there is anything else than OP_PUSH)
- the serialized script is popped from the stack
- the hash of this script is checked against the one in pkScript
- the script is executed


## Networking

Bitcoin uses a more-or-less straightforward peer-to-peer networking scheme (based on broadcasting to connected peers). This includes

- peer discovery
- bootstrapping used to be via IRC, these days there are a few wired-in server addresses
- then periodically broadcasting our own IP, and asking other nodes for a list of peers
- selecting a few peers randomly
- relaying transactions
- relaying blocks
- initial download of older blocks
- etc

Newer developments:

- simplified payment verification (SPV) $\rightarrow$ thin clients
- some DOS protection is necessary


## Elliptic curves

What is an elliptic curve?

- $y^{2}=x^{3}+a x+b$ over a field $\mathbb{F}$ called the "short Weierstrass form"; $\left(\operatorname{char}(\mathbb{F}) \neq 2,3\right.$ and $\left.4 a^{3}+27 b^{2} \neq 0\right)$
- smooth cubic plane curve (with a base point)
- (...other mathematical definitions...)

Etimology / history:

- arc length of an ellipse
- elliptic integrals
- the inverse problem: elliptic functions
- elliptic curves


## Pictures of elliptic curves over $\mathbb{R}$

Table [Contourplot $\left.\left[y^{2}=x^{3}+a \star x+b,\{x,-10,10\},\{y,-10,10\}\right],\{a,\{-15,-5,+5\}\},\{b,\{-15,+5,+25\}\}\right]$










## The group law on elliptic curves

It's a kind of surprising fact, that elliptic curves has a group structure, which furthermore has a rather nice geometric interpretation.

Definitions (for the Weierstrass form):

- identity element: The point at infinity (denoted by $O$ )
- inverse: mirroring wrt. the X axis
- addition: if $P, Q$ and $R$ are on a straight line, we declare $P+Q+R=O$

Group laws:

- the identity satisfies what it should (trivial)
- addition is commutative (trivial)
- addition is associative (nontrival!)

Addition on elliptic curves


## Elliptic curve cryptography

Elliptic curve cryptography is based on the elliptic curve discrete logarithm problem. Consider an elliptic curve $\mathcal{G}$ over a (very large) finite field, usually either $\mathbb{F}_{2^{n}}$ or $\mathbb{F}_{p}$ for a prime $p$.

We fix the curve $\mathcal{G}$ and a generator $G \in \mathcal{G}$ of this group; that is, $\mathcal{G}=\{k G \in \mathcal{G}: k \in[0, n-1]\}$ where $n=|\mathcal{G}|$ is the size of this group.

- private key: a random number $d \in[1, n-1] \subset \mathbb{N}$
- public key: the group element $Q=d G \in \mathcal{G}$

It is easy to compute $Q$ from $d$ and $G$ in such a group, but very hard (at least for appropriate choice of the field and $\mathcal{G}$ ) to determine $d$ from $Q$.

## Key length comparison

The primary advantage of ECC is that we can have shorter keys for equivalent security:

| Security level | RSA key length | ECC key length | ratio |
| :---: | :---: | :---: | :---: |
| 80 | 1024 | $160-223$ | $5-6$ |
| 112 | 2048 | $224-255$ | $8-9$ |
| 128 | 3072 | $256-283$ | $11-12$ |
| 192 | 7680 | $384-511$ | $15-20$ |
| 256 | 15360 | $512-571$ | $27-30$ |

The shorter key length can be actually important in practice:

- embedded devices with limited resources (smartcards, etc)
- large number of signatures: both bandwidth and storage can be important (bitcoin)

A picture of $y^{2}=x^{3}+7$ over $\mathbb{F}_{59}$


Remark: $p \approx 2^{256} \approx 10^{77} \approx$ no. of elementary particles in the universe

## The secp256k1 curve

The safety of ECC depends on the hardness of the elliptic curve discrete logarithm problem. Not all curves are created equal!

Bitcoin uses the curve called secp256k1:

- the field is $\mathbb{F}_{p}$ with $p=2^{256}-2^{32}-2^{9}-2^{8}-2^{7}-2^{6}-2^{4}-1$
- the curve equation is $y^{2}=x^{3}+7$ (that is, $a=0$ and $b=7$ )
- the number of points on the curve $n$ is also a prime, and "close" to $p$ (to be more precise, $p-2^{129}<n<p-2^{128}$ )
- since $n$ is prime, the group is cyclic; thus any element (except the infinity) will do as the generator $G$
- but there is a concrete, randomly-looking $G$ in the standard

```
p = 0x FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFFF FFFFFFFE FFFFFC2F
n = 0x FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFE BAAEDCE6 AF48A03B BFD25E8C D0364141
G}\mp@subsup{G}{x}{}=0x 79BE667E F9DCBBAC 55A06295 CE870B07 029BFCDB 2DCE28D9 59F2815B 16F81798
G}=0\textrm{m}483ADA77 26A3C465 5DA4FBFC 0E1108A8 FD17B448 A6855419 9C47D08F FB10D4B8
```


## secp256k1 parameter choices

Arguably, apart from the choice generator, the parameters look sincere:

- the prime field $\mathbb{F}_{p}$ is very close to $2^{256}$, and differs from it by only a few bits
- the curve equation: if you look at $y^{2}=x^{3}+b$ with $b \in\{1,2,3, \ldots\}$, it seems that $b=7$ is the very first choice where you can have a somewhat reasonable curve
- the size of the curve $n$ is also a prime (the co-factor is $h=1$ ), which makes the code simpler
- the generator $G$, we have no idea how they chose that. On the other hand, the group is cyclic, so in principle any $G$ would do (though it may happen that for special choices of $G$, the discrete logarithm problem is simpler; I don't know anything about that)


## Bernstein's SafeCurves analysis (http://safecurves.cr.yp.to)

|  |  | Parameters: |  |  | ECDLP security: |  |  |  | ECC security: |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Curve | Safe? | field | equation | base | rho | transfer | disc | rigid | ladder | twist | complete | ind |
| Anomalous | False | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | False | False | True $\sqrt{\text { V }}$ | False | False | False | False |
| M-221 | True $\sqrt{ }$ | ve $\sqrt{ }$ | rue $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True ${ }^{\text {d }}$ | True $\sqrt{ }$ | True $\sqrt{ }$ |
| \|NIST P-256 | False | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ } \mid$ | True $\sqrt{ }$ | True $\sqrt{ }$ | \|False | False | True $\sqrt{ }$ | False | False |
| secp256k1 | False | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | False | True $\sqrt{ }$ | False | True $\sqrt{ }$ | False | False |
| E-382 | True $\sqrt{\text { V }}$ | $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True | ue | True | True $\sqrt{ }$ | True $\sqrt{ }$ |
| M-511 | $\mid$ True $\sqrt{\text { l }}$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ } \mid$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ } \mid$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ |
| E-521 | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ | True $\sqrt{ }$ |

What does it mean?

- disc: is the discriminant large enough? (small discriminant is not an explicit problem, but the story becomes more complicated, better be safe)
- ladder: is there a simple, fast and constant-time scalar multiplication algo (similar to Montgomery ladder)?
- complete: is there a complete addition law (that is, one without special cases like doubling)?
- ind: does an elliptic curve point representation looks random?

Conclusion: secp256k1 is probably OK for what Bitcoin uses it (if you worry about timing attacks, maybe just add random delays?).

## Example: Diffie-Hellmann key exchange ${ }^{11}$

Recall that a key pair $(d, Q)$ consists of

- a private key, which is a random number $d \in[1, n-1] \subset \mathbb{N}$
- a public key, which is the group element $Q=d G \in \mathcal{G}$

Let there be two parties: Alice and Bob, with key pairs $\left(d_{A}, Q_{A}\right)$ and $\left(d_{B}, Q_{B}\right)$. They can compute a shared secret $S \in \mathcal{G}$ as follows:

$$
d_{A} Q_{B}=d_{A}\left(d_{B} G\right)=\underbrace{\left(d_{A} d_{B}\right) G}_{S}=d_{B}\left(d_{A} G\right)=d_{B} Q_{A}
$$

Alice can knows $d_{A}$, so she can compute the leftmost version; Bob knows $d_{B}$, so he can compute the rightmost version. But nobody else knows neither $d_{A}$ or $d_{B}$, thus $S$ is their secret.

They can then proceed and use $S$ for any purpose, for example as the key of a symmetric encryption scheme.

[^6]
## Elliptic Curve Digital Signature Algorithm

Alice writes a message $m$, and wants to prove that she wrote it. She already has key pair $\left(d_{A}, Q_{A}\right)$, and people accept that the public key $Q_{A}$ in fact belongs to her.

Construction of the signature:

1. compute a hash $z=\operatorname{HASH}(m) \in[1, n-1]$ of the message $m$
2. generate an ephemeral key pair: $k$ and $Q_{k}=k G=(x, y)$
3. let $r=(x \bmod n) \in \mathbb{Z}_{n}$
4. let $s=k^{-1}\left(z+r d_{A}\right) \in \mathbb{Z}_{n}$
5. the signature is $(r, s) \in \mathbb{Z}_{n} \times \mathbb{Z}_{n}$

Verification of the signature:

1. compute $z=\operatorname{HASH}(m) \in[1, n-1]$ as before
2. compute $u=s^{-1} z \in \mathbb{Z}_{n}$ and $v=s^{-1} r \in \mathbb{Z}_{n}$
3. compute the curve point $(x, y)=Q=u G+v Q_{A} \in \mathcal{G}$
4. the signature is valid iff $x=r$.

## Representations of elliptic curve points

There are many different representations of elliptic curves:

- Weierstrass affine coordinates $\mathbb{A}^{2}$
- Weierstrass projective coordinates $\mathbb{P}^{2}$
- Weierstrass weighted projective coordinates $\mathbb{P}(2,3,1)$
- Montgomery form
- Hessian form
- Jacobian form
- Edwards form
- etc...

Representation matters because of efficiency! And possibly also different security properties.

## Computations with elliptic curves, I.

How to compute $d Q \in \mathcal{G}$ efficiently, with $d \in \mathbb{Z}_{n}$ and $Q \in \mathcal{G}$ ?

Answer: "fast exponentation"! Write $d$ in binary form:
$d=\sum_{i=0}^{m-1} d_{i} 2^{i}$, where $d_{i} \in\{0,1\}$.

$$
d Q=\left(\sum_{i=0}^{m-1} d_{i} 2^{i}\right) \cdot Q=\sum_{i=0}^{m-1} d_{i}\left(2^{i} Q\right)=\sum_{i=0}^{m-1} d_{i} Q_{i}
$$

where $Q_{i}=2^{i} Q$ can be computated by repeated doubling:
$Q_{0}=Q, Q_{1}=2 Q_{0}, Q_{2}=2 Q_{1}, Q_{3}=2 Q_{2}$, etc $\ldots$
Thus we need addition and doubling (which is a special case of addition, but needs to be handled separately anyway).

## Elliptic curve addition and doubling in pictures




## Elliptic curve addition and doubling in Weierstrass form

In any field $\mathbb{F}(\operatorname{char}(\mathbb{F}) \neq 2,3)$, for two points $P \neq Q \neq O$ on the elliptic curve $y^{2}=x^{3}+a x+b$, with coordinates $P=\left(x_{p}, y_{p}\right)$ and $Q=\left(x_{q}, y_{q}\right)$, it is straightforward (?) to calculate the coordinates of $P+Q=R=\left(x_{r}, y_{r}\right)$ and $2 P=U=\left(x_{u}, y_{u}\right)$ :

$$
\begin{aligned}
s & =\frac{y_{q}-y_{p}}{x_{q}-x_{p}} & t & =\frac{3 x_{p}^{2}+a}{2 y} \\
x_{r} & =s^{2}-\left(x_{p}+x_{q}\right) & x_{u} & =t^{2}-2 x_{u} \\
y_{r} & =-y_{p}-s\left(x_{r}-x_{p}\right) & y_{u} & =-y_{p}-s\left(x_{u}-x_{p}\right)
\end{aligned}
$$

Here $s$ resp. $t$ are the slopes of the secant resp. tangent lines.


[^0]:    ${ }^{1}$ varint $=$ a customly encoded variable-length integer
    ${ }^{2}$ The lock-time is a rarely used feature, which specifies a time in the future (measured either in block height or in unix timestamp), before which the transaction cannot be included in any block.

[^1]:    ${ }^{3}$ in practice, the reference client only relays the so-called "standard" scripts

[^2]:    ${ }^{4}$ there are a lot of corner cases, though...
    ${ }^{5}$ at least with typical spending condition - so called puzzle transaction may not contain such a signature, and thus can be modified by third parties $\bar{\equiv}$

[^3]:    ${ }^{6}$ but there is a very complicated and confusing prodecure of how exactly do this in general, which, as far as I knew, was never executed in history..

[^4]:    ${ }^{7}$ but because of some old bug, $0 \times 00$ (or a missing byte) is also accepted

[^5]:    ${ }^{8}$ of course this isn't documented anywhere...
    ${ }^{9}$... and neither is this

[^6]:    ${ }^{11}$ not actually used by Bitcoin

