

net : Reinventing the Internet

lightweight, scalable and userfriendly

Bernd Paysan

#wefixthenet, 31c3, Hamburg



Outline



Motivation

Topology

Low-Overhead Packet Format

Encryption

Key Exchange

Symmetric Crypto

Flow Control

Commands

Distributed Data

Applications

Apps in a Sandbox

API Basics

1.5 years after Snowden



What happened to change the world:

Politics More spying, more cyberwar, more terrorist panic — don't count on them

Users 700 million users changed their behavior (that's probably 700 million terrorists)

Software Lots of work, even WhatsApp got some crypto!

Protocols Most of the Internet still is a complete mess with security tugged in

1.5 years after Snowden



What happened to change the world:

Politics More spying, more cyberwar, more terrorist panic — don't count on them

Users 700 million users changed their behavior (that's probably 700 million terrorists)

Software Lots of work, even WhatsApp got some crypto!

Protocols Most of the Internet still is a complete mess with security tugged in

1.5 years after Snowden



What happened to change the world:

Politics More spying, more cyberwar, more terrorist panic — don't count on them

Users 700 million users changed their behavior (that's probably 700 million terrorists)

Software Lots of work, even WhatsApp got some crypto!

Protocols Most of the Internet still is a complete mess with security tugged in

1.5 years after Snowden



What happened to change the world:

Politics More spying, more cyberwar, more terrorist panic — don't count on them

Users 700 million users changed their behavior (that's probably 700 million terrorists)

Software Lots of work, even WhatsApp got some crypto!

Protocols Most of the Internet still is a complete mess with security tugged in

1.5 years after Snowden



What happened to change the world:

Politics More spying, more cyberwar, more terrorist panic — don't count on them

Users 700 million users changed their behavior (that's probably 700 million terrorists)

Software Lots of work, even WhatsApp got some crypto!

Protocols Most of the Internet still is a complete mess with security tugged in

The Enemies of the Internet



Criminals malware, DDoS attacks, spam, ...

Corporations walled gardens, censorship, big honeypots for dragnet surveillance, ...

Government dragnet surveillance, censorship, ...

Users careless, uninformed, annoying, ...

Software bloated, buggy, insecure, ...

The Enemies of the Internet



Criminals malware, DDoS attacks, spam, ...

Corporations walled gardens, censorship, big honeypots for dragnet surveillance, ...

Government dragnet surveillance, censorship, ...

Users careless, uninformed, annoying, ...

Software bloated, buggy, insecure, ...

The Enemies of the Internet



Criminals malware, DDoS attacks, spam, ...

Corporations walled gardens, censorship, big honeypots for dragnet surveillance, ...

Government dragnet surveillance, censorship, ...

Users careless, uninformed, annoying, ...

Software bloated, buggy, insecure, ...

The Enemies of the Internet



Criminals malware, DDoS attacks, spam, ...

Corporations walled gardens, censorship, big honeypots for dragnet surveillance, ...

Government dragnet surveillance, censorship, ...

Users careless, uninformed, annoying, ...

Software bloated, buggy, insecure, ...

The Enemies of the Internet



Criminals malware, DDoS attacks, spam, ...

Corporations walled gardens, censorship, big honeypots for dragnet surveillance, ...

Government dragnet surveillance, censorship, ...

Users careless, uninformed, annoying, ...

Software bloated, buggy, insecure, ...

How many defects?



- DAN GEER: buy all zero-days
- Condition: The number of bugs are finite. Are they?
- Bug density between $1/100\text{LoC}$ (CMM 1) to $<1/10\text{kLoC}$ (Correct by Design [3])
- Networked applications and protocol stacks in orders of $1\text{M}-100\text{MLoC}$
- Unless we stop bloating, we are doomed
- Therefore: **Keep it simple!**

How many defects?



- DAN GEER: buy all zero-days
- Condition: The number of bugs are finite. Are they?
- Bug density between $1/100\text{LoC}$ (CMM 1) to $<1/10\text{kLoC}$ (Correct by Design [3])
- Networked applications and protocol stacks in orders of $1\text{M}-100\text{MLoC}$
- Unless we stop bloating, we are doomed
- Therefore: **Keep it simple!**

How many defects?



- DAN GEER: buy all zero-days
- Condition: The number of bugs are finite. Are they?
- Bug density between $1/100\text{LoC}$ (CMM 1) to $<1/10\text{kLoC}$ (Correct by Design [3])
- Networked applications and protocol stacks in orders of $1\text{M}-100\text{MLoC}$
- Unless we stop bloating, we are doomed
- Therefore: **Keep it simple!**

How many defects?



- DAN GEER: buy all zero-days
- Condition: The number of bugs are finite. Are they?
- Bug density between $1/100\text{LoC}$ (CMM 1) to $<1/10\text{kLoC}$ (Correct by Design [3])
- Networked applications and protocol stacks in orders of $1\text{M}-100\text{MLoC}$
 - Unless we stop bloating, we are doomed
 - Therefore: **Keep it simple!**

How many defects?



- DAN GEER: buy all zero-days
- Condition: The number of bugs are finite. Are they?
- Bug density between $1/100\text{LoC}$ (CMM 1) to $<1/10\text{kLoC}$ (Correct by Design [3])
- Networked applications and protocol stacks in orders of $1\text{M}-100\text{MLoC}$
- Unless we stop bloating, we are doomed
- Therefore: **Keep it simple!**

How many defects?



- DAN GEER: buy all zero-days
- Condition: The number of bugs are finite. Are they?
- Bug density between $1/100\text{LoC}$ (CMM 1) to $<1/10\text{kLoC}$ (Correct by Design [3])
- Networked applications and protocol stacks in orders of $1\text{M}-100\text{MLoC}$
- Unless we stop bloating, we are doomed
- Therefore: **Keep it simple!**



Where are the defects?




Figure : Bugs by phase [2]

net2o in a nutshell



net2o consists of the following 6 layers (implemented bottom up):

2. Path switched packets with 2^n size writing into shared memory buffers
3. Ephemeral key exchange and signatures with Ed25519, symmetric authenticated encryption+hash+prng with Keccak, symmetric block encryption with Threefish
4. Timing driven delay minimizing flow control
5. Stack-oriented tokenized command language
6. Distributed data (files) and distributed metadata (prefix hash trie)
7. Apps in a sandboxed environment for displaying content

net2o in a nutshell



net2o consists of the following 6 layers (implemented bottom up):

2. Path switched packets with 2^n size writing into shared memory buffers
3. Ephemeral key exchange and signatures with Ed25519, symmetric authenticated encryption+hash+prng with Keccak, symmetric block encryption with Threefish
4. Timing driven delay minimizing flow control
5. Stack-oriented tokenized command language
6. Distributed data (files) and distributed metadata (prefix hash trie)
7. Apps in a sandboxed environment for displaying content

net2o in a nutshell



net2o consists of the following 6 layers (implemented bottom up):

2. Path switched packets with 2^n size writing into shared memory buffers
3. Ephemeral key exchange and signatures with Ed25519, symmetric authenticated encryption+hash+prng with Keccak, symmetric block encryption with Threefish
4. Timing driven delay minimizing flow control
5. Stack-oriented tokenized command language
6. Distributed data (files) and distributed metadata (prefix hash trie)
7. Apps in a sandboxed environment for displaying content

net2o in a nutshell



net2o consists of the following 6 layers (implemented bottom up):

2. Path switched packets with 2^n size writing into shared memory buffers
3. Ephemeral key exchange and signatures with Ed25519, symmetric authenticated encryption+hash+prng with Keccak, symmetric block encryption with Threefish
4. Timing driven delay minimizing flow control
5. Stack-oriented tokenized command language
6. Distributed data (files) and distributed metadata (prefix hash trie)
7. Apps in a sandboxed environment for displaying content

net2o in a nutshell



net2o consists of the following 6 layers (implemented bottom up):

2. Path switched packets with 2^n size writing into shared memory buffers
3. Ephemeral key exchange and signatures with Ed25519, symmetric authenticated encryption+hash+prng with Keccak, symmetric block encryption with Threefish
4. Timing driven delay minimizing flow control
5. Stack-oriented tokenized command language
6. Distributed data (files) and distributed metadata (prefix hash trie)
7. Apps in a sandboxed environment for displaying content

net2o in a nutshell



net2o consists of the following 6 layers (implemented bottom up):

2. Path switched packets with 2^n size writing into shared memory buffers
3. Ephemeral key exchange and signatures with Ed25519, symmetric authenticated encryption+hash+prng with Keccak, symmetric block encryption with Threefish
4. Timing driven delay minimizing flow control
5. Stack-oriented tokenized command language
6. Distributed data (files) and distributed metadata (prefix hash trie)
7. Apps in a sandboxed environment for displaying content

net2o in a nutshell



net2o consists of the following 6 layers (implemented bottom up):

2. Path switched packets with 2^n size writing into shared memory buffers
3. Ephemeral key exchange and signatures with Ed25519, symmetric authenticated encryption+hash+prng with Keccak, symmetric block encryption with Threefish
4. Timing driven delay minimizing flow control
5. Stack-oriented tokenized command language
6. Distributed data (files) and distributed metadata (prefix hash trie)
7. Apps in a sandboxed environment for displaying content

Objectives



net2o's design objectives are

- lightweight, fast, scalable
- easy to implement
- secure
- media capable
- works as overlay on current networks (UDP/IP), but can replace the entire stack

Objectives



net2o's design objectives are

- **lightweight, fast, scalable**
- easy to implement
- secure
- media capable
- works as overlay on current networks (UDP/IP), but can replace the entire stack

Objectives



net2o's design objectives are

- lightweight, fast, scalable
- easy to implement
- secure
- media capable
- works as overlay on current networks (UDP/IP), but can replace the entire stack



Objectives



net2o's design objectives are

- lightweight, fast, scalable
- easy to implement
- secure
- media capable
- works as overlay on current networks (UDP/IP), but can replace the entire stack

Objectives



net2o's design objectives are

- lightweight, fast, scalable
- easy to implement
- secure
- media capable
- works as overlay on current networks (UDP/IP), but can replace the entire stack



Objectives



net2o's design objectives are

- lightweight, fast, scalable
- easy to implement
- secure
- media capable
- works as overlay on current networks (UDP/IP), but can replace the entire stack



Switching Packets, Routing Connections



- Switches are faster and easier to implement than routers
- Routing then is a combination of destination resolution and routing calculation (destination path lookup)

Path Switching

- Take first n bits of path field and select destination
- Shift target address by n
- Insert bit-reversed source into the rear end of the path field to mark the way back
- The receiver bit-flips the path field, and gets the return address
- Easy handover possible



Switching Packets, Routing Connections



- Switches are faster and easier to implement than routers
- Routing then is a combination of destination resolution and routing calculation (destination path lookup)

Path Switching

- Take first n bits of path field and select destination
- Shift target address by n
- Insert bit-reversed source into the rear end of the path field to mark the way back
- The receiver bit-flips the path field, and gets the return address
- Easy handover possible



Switching Packets, Routing Connections



- Switches are faster and easier to implement than routers
- Routing then is a combination of destination resolution and routing calculation (destination path lookup)

Path Switching

- Take first n bits of path field and select destination
- Shift target address by n
- Insert bit-reversed source into the rear end of the path field to mark the way back
- The receiver bit-flips the path field, and gets the return address
- Easy handover possible



Switching Packets, Routing Connections



- Switches are faster and easier to implement than routers
- Routing then is a combination of destination resolution and routing calculation (destination path lookup)

Path Switching

- Take first n bits of path field and select destination
- Shift target address by n
- Insert bit-reversed source into the rear end of the path field to mark the way back
- The receiver bit-flips the path field, and gets the return address
- Easy handover possible



Switching Packets, Routing Connections



- Switches are faster and easier to implement than routers
- Routing then is a combination of destination resolution and routing calculation (destination path lookup)

Path Switching

- Take first n bits of path field and select destination
- Shift target address by n
- Insert bit-reversed source into the rear end of the path field to mark the way back
- The receiver bit-flips the path field, and gets the return address
- Easy handover possible



Switching Packets, Routing Connections



- Switches are faster and easier to implement than routers
- Routing then is a combination of destination resolution and routing calculation (destination path lookup)

Path Switching

- Take first n bits of path field and select destination
- Shift target address by n
- Insert bit-reversed source into the rear end of the path field to mark the way back

- The receiver bit-flips the path field, and gets the return address

- Easy handover possible



Switching Packets, Routing Connections



- Switches are faster and easier to implement than routers
- Routing then is a combination of destination resolution and routing calculation (destination path lookup)

Path Switching

- Take first n bits of path field and select destination
- Shift target address by n
- Insert bit-reversed source into the rear end of the path field to mark the way back

- The receiver bit-flips the path field, and gets the return address

- Easy handover possible



Routing Algorithm



- A node publishes ISP switch+label in the DHT
- The ISP publishes peering switch+label in the DHT
- Assumption is a hierarchical network, so a recursive lookup will give a good solution
- Splice the labels together, and you get a path



Routing Algorithm



- A node publishes ISP switch+label in the DHT
- The ISP publishes peering switch+label in the DHT
- Assumption is a hierarchical network, so a recursive lookup will give a good solution
- Splice the labels together, and you get a path



Routing Algorithm



- A node publishes ISP switch+label in the DHT
- The ISP publishes peering switch+label in the DHT
- Assumption is a hierarchical network, so a recursive lookup will give a good solution
- Splice the labels together, and you get a path



Routing Algorithm



- A node publishes ISP switch+label in the DHT
- The ISP publishes peering switch+label in the DHT
- Assumption is a hierarchical network, so a recursive lookup will give a good solution
- Splice the labels together, and you get a path



Why Source Routing



Three possible schemes

1. switched circuit (POTS, virtual: ATM, MPLS)
 2. unique identifier (IP)
 3. source routing
- Separation of network gear and computers: Fast, dumb, stateless equipment for routing/switching
 - The hierarchical topology is a derived “law of nature”: people cluster together and connect clusters
 - Attack vector is only bandwidth-based, and this can be mitigated (see “fair routing” below)
 - Routing slice is an implementation detail of each network segment (i.e. is a unique identifier within each subnet)



Why Source Routing



Three possible schemes

1. switched circuit (POTS, virtual: ATM, MPLS)
 2. unique identifier (IP)
 3. source routing
- Separation of network gear and computers: Fast, dumb, stateless equipment for routing/switching
 - The hierarchical topology is a derived “law of nature”: people cluster together and connect clusters
 - Attack vector is only bandwidth-based, and this can be mitigated (see “fair routing” below)
 - Routing slice is an implementation detail of each network segment (i.e. is a unique identifier within each subnet)

Why Source Routing



Three possible schemes

1. switched circuit (POTS, virtual: ATM, MPLS)
 2. unique identifier (IP)
 3. source routing
- Separation of network gear and computers: Fast, dumb, stateless equipment for routing/switching
 - The hierarchical topology is a derived “law of nature”: people cluster together and connect clusters
 - Attack vector is only bandwidth-based, and this can be mitigated (see “fair routing” below)
 - Routing slice is an implementation detail of each network segment (i.e. is a unique identifier within each subnet)



Why Source Routing



Three possible schemes

1. switched circuit (POTS, virtual: ATM, MPLS)
 2. unique identifier (IP)
 3. source routing
- Separation of network gear and computers: Fast, dumb, stateless equipment for routing/switching
 - The hierarchical topology is a derived “law of nature”: people cluster together and connect clusters
 - Attack vector is only bandwidth-based, and this can be mitigated (see “fair routing” below)
 - Routing slice is an implementation detail of each network segment (i.e. is a unique identifier within each subnet)



Why Source Routing



Three possible schemes


1. switched circuit (POTS, virtual: ATM, MPLS)
 2. unique identifier (IP)
 3. source routing
- Separation of network gear and computers: Fast, dumb, stateless equipment for routing/switching
 - The hierarchical topology is a derived “law of nature”: people cluster together and connect clusters
 - Attack vector is only bandwidth-based, and this can be mitigated (see “fair routing” below)
 - Routing slice is an implementation detail of each network segment (i.e. is a unique identifier within each subnet)



Packet Format

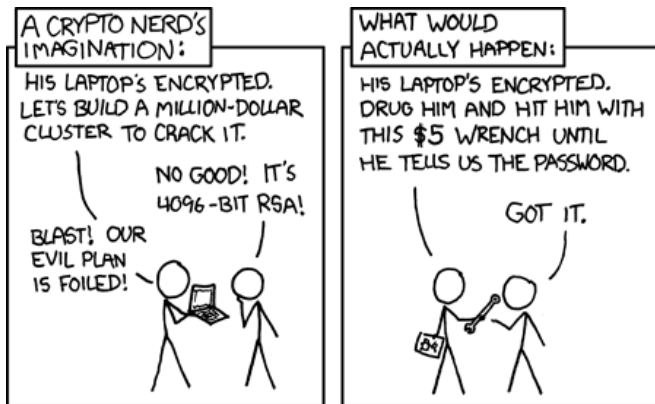


	<i>Bytes</i>	<i>Comment</i>
<i>Flags</i>	2	priority, length, flow control flags
<i>Path</i>	16	Internet 1.0 terminology: "address"
<i>Address</i>	8	address in memory, \approx port+sequence number
<i>Data</i>	$64 * 2^{0..15}$	up to 2MB packet size, enough for the next 40 years
<i>Chksum</i>	16	cryptographic checksum





Security: Indirect Attacks are Cheaper



Key Exchange



ECC Elliptic Curve Cryptography has still only a generic attack (i.e. can be considered “unscratched”, as the attack uses a fundamental property of the problem), and therefore 256 bit keys (32 bytes) have a strength of 128 bits

Therefore the choice now is Ed25519, a variant of Curve25519 from DAN BERNSTEIN that supports signatures, too. This is a curve where the parameters are of high quality.

I use Ed25519 both for Diffie–Hellman–Exchange and signatures with the same key; PETER SCHWABE warned me that this might be insecure in some circumstances and that they are working on some recommendations how to do this securely.

Key Exchange



ECC Elliptic Curve Cryptography has still only a generic attack (i.e. can be considered “unscratched”, as the attack uses a fundamental property of the problem), and therefore 256 bit keys (32 bytes) have a strength of 128 bits

Therefore the choice now is Ed25519, a variant of Curve25519 from DAN BERNSTEIN that supports signatures, too. This is a curve where the parameters are of high quality.

I use Ed25519 both for Diffie–Hellman–Exchange and signatures with the same key; PETER SCHWABE warned me that this might be insecure in some circumstances and that they are working on some recommendations how to do this securely.

Key Exchange



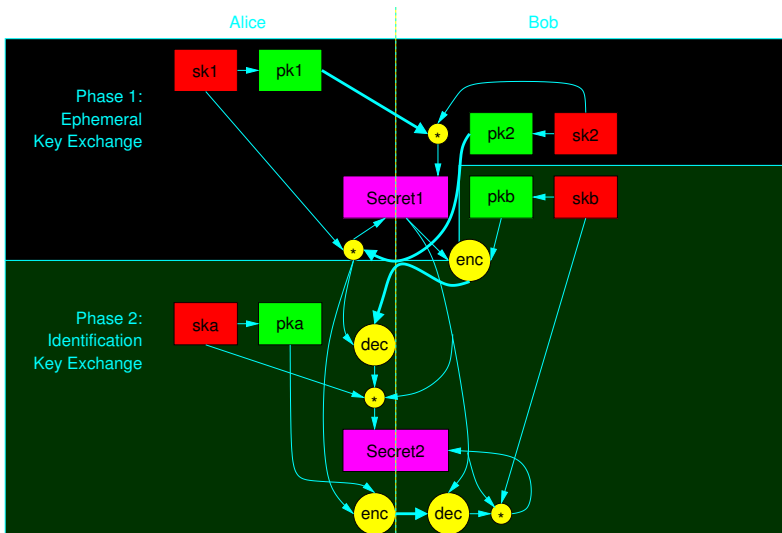
ECC Elliptic Curve Cryptography has still only a generic attack (i.e. can be considered “unscratched”, as the attack uses a fundamental property of the problem), and therefore 256 bit keys (32 bytes) have a strength of 128 bits

Therefore the choice now is Ed25519, a variant of Curve25519 from DAN BERNSTEIN that supports signatures, too. This is a curve where the parameters are of high quality.

I use Ed25519 both for Diffie–Hellman–Exchange and signatures with the same key; PETER SCHWABE warned me that this might be insecure in some circumstances and that they are working on some recommendations how to do this securely.



Ephemeral Key Exchange+Validation



Key Replacement



Problem with key revocation: you really want to *replace* your compromised key; signing the revocation with the secret key is a bad idea, because the secret key is compromised (the attacker might be the source of the key replacement)

- Only the creator of the secret key can revoke it
- A thief of the secret key can't (i.e. further information is necessary)
- Revocation must present a trustworthy replacement key
- Third parties must trust both the revocation and the replacement key without another trustworthy instance, i.e. trusting only their communication partner
- Solution: "proof of creation", i.e. you prove you made the key with a separately stored secret

Key Replacement



Problem with key revocation: you really want to *replace* your compromised key; signing the revocation with the secret key is a bad idea, because the secret key is compromised (the attacker might be the source of the key replacement)

- Only the creator of the secret key can revoke it
- A thief of the secret key can't (i.e. further information is necessary)
- Revocation must present a trustworthy replacement key
- Third parties must trust both the revocation and the replacement key without another trustworthy instance, i.e. trusting only their communication partner
- Solution: "proof of creation", i.e. you prove you made the key with a separately stored secret

Key Replacement



Problem with key revocation: you really want to *replace* your compromised key; signing the revocation with the secret key is a bad idea, because the secret key is compromised (the attacker might be the source of the key replacement)

- Only the creator of the secret key can revoke it
- A thief of the secret key can't (i.e. further information is necessary)
- Revocation must present a trustworthy replacement key
- Third parties must trust both the revocation and the replacement key without another trustworthy instance, i.e. trusting only their communication partner
- Solution: "proof of creation", i.e. you prove you made the key with a separately stored secret



Key Replacement



Problem with key revocation: you really want to *replace* your compromised key; signing the revocation with the secret key is a bad idea, because the secret key is compromised (the attacker might be the source of the key replacement)

- Only the creator of the secret key can revoke it
- A thief of the secret key can't (i.e. further information is necessary)
- Revocation must present a trustworthy replacement key
- Third parties must trust both the revocation and the replacement key without another trustworthy instance, i.e. trusting only their communication partner
- Solution: "proof of creation", i.e. you prove you made the key with a separately stored secret



Key Replacement



Problem with key revocation: you really want to *replace* your compromised key; signing the revocation with the secret key is a bad idea, because the secret key is compromised (the attacker might be the source of the key replacement)

- Only the creator of the secret key can revoke it
- A thief of the secret key can't (i.e. further information is necessary)
- Revocation must present a trustworthy replacement key
- Third parties must trust both the revocation and the replacement key without another trustworthy instance, i.e. trusting only their communication partner
- Solution: "proof of creation", i.e. you prove you made the key with a separately stored secret

Key Replacement



Problem with key revocation: you really want to *replace* your compromised key; signing the revocation with the secret key is a bad idea, because the secret key is compromised (the attacker might be the source of the key replacement)

- Only the creator of the secret key can revoke it
- A thief of the secret key can't (i.e. further information is necessary)
- Revocation must present a trustworthy replacement key
- Third parties must trust both the revocation and the replacement key without another trustworthy instance, i.e. trusting only their communication partner
- Solution: “proof of creation”, i.e. you prove you made the key with a separately stored secret

Proof of Creation



- Create two 256 bit random numbers s_1 and s_2
- Create pubkeys $p_1 = base * [s_1]$ and $p_2 = base * [s_2]$
- Compute $[s] = [s_1 * p_2]$ as "work secret" and $p = base * [s]$, the pubkey
- Publish p and p_1 , destroy s_1 (no longer needed), keep s_2 as offline copy (e.g. on paper)
- To revoke a key, publish p_2 , which the recipient can validate by $p_1 * [p_2] \equiv p$.
- To proof possession of all secrets, sign new key with s_2 , s , and s_{new}

Proof of Creation



- Create two 256 bit random numbers s_1 and s_2
- Create pubkeys $p_1 = base * [s_1]$ and $p_2 = base * [s_2]$
- Compute $[s] = [s_1 * p_2]$ as "work secret" and $p = base * [s]$, the pubkey
- Publish p and p_1 , destroy s_1 (no longer needed), keep s_2 as offline copy (e.g. on paper)
- To revoke a key, publish p_2 , which the recipient can validate by $p_1 * [p_2] \equiv p$.
- To proof possession of all secrets, sign new key with s_2 , s , and s_{new}

Proof of Creation



- Create two 256 bit random numbers s_1 and s_2
- Create pubkeys $p_1 = base * [s_1]$ and $p_2 = base * [s_2]$
- Compute $[s] = [s_1 * p_2]$ as "work secret" and $p = base * [s]$, the pubkey
- Publish p and p_1 , destroy s_1 (no longer needed), keep s_2 as offline copy (e.g. on paper)
- To revoke a key, publish p_2 , which the recipient can validate by $p_1 * [p_2] \equiv p$.
- To proof possession of all secrets, sign new key with s_2 , s_1 , and s_{new}

Proof of Creation



- Create two 256 bit random numbers s_1 and s_2
- Create pubkeys $p_1 = base * [s_1]$ and $p_2 = base * [s_2]$
- Compute $[s] = [s_1 * p_2]$ as "work secret" and $p = base * [s]$, the pubkey
- Publish p and p_1 , destroy s_1 (no longer needed), keep s_2 as offline copy (e.g. on paper)
- To revoke a key, publish p_2 , which the recipient can validate by $p_1 * [p_2] \equiv p$.
- To proof possession of all secrets, sign new key with s_2 , s_1 , and s_{new}

Proof of Creation



- Create two 256 bit random numbers s_1 and s_2
- Create pubkeys $p_1 = base * [s_1]$ and $p_2 = base * [s_2]$
- Compute $[s] = [s_1 * p_2]$ as "work secret" and $p = base * [s]$, the pubkey
- Publish p and p_1 , destroy s_1 (no longer needed), keep s_2 as offline copy (e.g. on paper)
- To revoke a key, publish p_2 , which the recipient can validate by $p_1 * [p_2] \equiv p$.
- To proof possession of all secrets, sign new key with s_2 , s_1 , and s_{new}

Proof of Creation



- Create two 256 bit random numbers s_1 and s_2
- Create pubkeys $p_1 = base * [s_1]$ and $p_2 = base * [s_2]$
- Compute $[s] = [s_1 * p_2]$ as "work secret" and $p = base * [s]$, the pubkey
- Publish p and p_1 , destroy s_1 (no longer needed), keep s_2 as offline copy (e.g. on paper)
- To revoke a key, publish p_2 , which the recipient can validate by $p_1 * [p_2] \equiv p$.
- To proof possession of all secrets, sign new key with s_2 , s , and s_{new}

Symmetric Crypto: Keccak



Keccak used for the following reasons:

- Good cryptanalysis
- Keccak in duplex mode provides perfect side-channel protected AEAD operation (no constant key to snoop)
- Strength >256 bits: very good security margin
- Keccak is a universal crypto primitive (hash+encrypt+authenticate)
- Keccak is both NIST-approved and (still) NSA-independent. I use Keccak with $r = 1024$ and capacity $c = 576$ as suggested by the Keccak authors.

Symmetric Crypto: Keccak



Keccak used for the following reasons:

- Good cryptanalysis
- Keccak in duplex mode provides perfect side-channel protected AEAD operation (no constant key to snoop)
- Strength >256 bits: very good security margin
- Keccak is a universal crypto primitive (hash+encrypt+authenticate)
- Keccak is both NIST-approved and (still) NSA-independent. I use Keccak with $r = 1024$ and capacity $c = 576$ as suggested by the Keccak authors.

Symmetric Crypto: Keccak



Keccak used for the following reasons:

- Good cryptanalysis
- Keccak in duplex mode provides perfect side-channel protected AEAD operation (no constant key to snoop)
- Strength >256 bits: very good security margin
- Keccak is a universal crypto primitive (hash+encrypt+authenticate)
- Keccak is both NIST-approved and (still) NSA-independent. I use Keccak with $r = 1024$ and capacity $c = 576$ as suggested by the Keccak authors.

Symmetric Crypto: Keccak



Keccak used for the following reasons:

- Good cryptanalysis
- Keccak in duplex mode provides perfect side-channel protected AEAD operation (no constant key to snoop)
- Strength >256 bits: very good security margin
- Keccak is a universal crypto primitive (hash+encrypt+authenticate)
- Keccak is both NIST-approved and (still) NSA-independent. I use Keccak with $r = 1024$ and capacity $c = 576$ as suggested by the Keccak authors.

Symmetric Crypto: Keccak



Keccak used for the following reasons:

- Good cryptanalysis
- Keccak in duplex mode provides perfect side-channel protected AEAD operation (no constant key to snoop)
- Strength >256 bits: very good security margin
- Keccak is a universal crypto primitive (hash+encrypt+authenticate)
- Keccak is both NIST-approved and (still) NSA-independent. I use Keccak with $r = 1024$ and capacity $c = 576$ as suggested by the Keccak authors.

Symmetric Crypto: Keccak

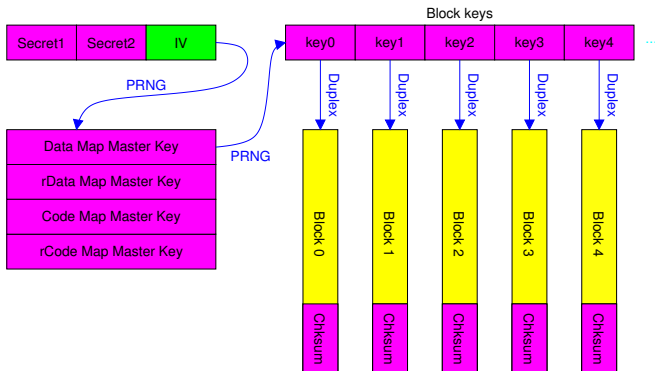


Keccak used for the following reasons:

- Good cryptanalysis
- Keccak in duplex mode provides perfect side-channel protected AEAD operation (no constant key to snoop)
- Strength >256 bits: very good security margin
- Keccak is a universal crypto primitive (hash+encrypt+authenticate)
- Keccak is both NIST-approved and (still) NSA-independent. I use Keccak with $r = 1024$ and capacity $c = 576$ as suggested by the Keccak authors.

Key Usage

All keys are one-time-use only!



Symmetric Crypto: Threefish



Keccak has one disadvantage: No ECB mode. Application for ECB mode:

- Encryption of hash values for the DHT — to store key/value pairs in a public DHT without revealing the content.
- For net2o-in-net2o tunnels (to be used for onion-routing), no authentication and no IV is desirable, so use an ECB mode algorithm.
- Strength >256 bits, tweakable to make ECB mode more secure (counter as tweak)
- SHA-3 finalist, so sufficiently good cryptanalysis

Symmetric Crypto: Threefish



Keccak has one disadvantage: No ECB mode. Application for ECB mode:

- Encryption of hash values for the DHT — to store key/value pairs in a public DHT without revealing the content.
- For net2o-in-net2o tunnels (to be used for onion-routing), no authentication and no IV is desirable, so use an ECB mode algorithm.
- Strength >256 bits, tweakable to make ECB mode more secure (counter as tweak)
- SHA-3 finalist, so sufficiently good cryptanalysis

Symmetric Crypto: Threefish



Keccak has one disadvantage: No ECB mode. Application for ECB mode:

- Encryption of hash values for the DHT — to store key/value pairs in a public DHT without revealing the content.
- For net2o-in-net2o tunnels (to be used for onion-routing), no authentication and no IV is desirable, so use an ECB mode algorithm.
- Strength >256 bits, tweakable to make ECB mode more secure (counter as tweak)
- SHA-3 finalist, so sufficiently good cryptanalysis

Symmetric Crypto: Threefish



Keccak has one disadvantage: No ECB mode. Application for ECB mode:

- Encryption of hash values for the DHT — to store key/value pairs in a public DHT without revealing the content.
- For net2o-in-net2o tunnels (to be used for onion-routing), no authentication and no IV is desirable, so use an ECB mode algorithm.
- Strength >256 bits, tweakable to make ECB mode more secure (counter as tweak)
- SHA-3 finalist, so sufficiently good cryptanalysis

Symmetric Crypto: Threefish



Keccak has one disadvantage: No ECB mode. Application for ECB mode:

- Encryption of hash values for the DHT — to store key/value pairs in a public DHT without revealing the content.
- For net2o-in-net2o tunnels (to be used for onion-routing), no authentication and no IV is desirable, so use an ECB mode algorithm.
- Strength >256 bits, tweakable to make ECB mode more secure (counter as tweak)
- SHA-3 finalist, so sufficiently good cryptanalysis



Flow Control (Broken)



- TCP fills the buffer, until a packet has to be dropped, instead of reducing rate before. Name of the symptom: "Buffer bloat". But buffering is essential for good network performance.




Figure : Buffer Bloat

Alternatives?



- LEDBAT tries to achieve a low, constant delay: Works, but not good on fairness
- CurveCP's flow control is still “a lot of research”
- Therefore, something new has to be done




Figure : That's how proper flow control should look like

Alternatives?



- LEDBAT tries to achieve a low, constant delay: Works, but not good on fairness
- CurveCP's flow control is still “a lot of research”
- Therefore, something new has to be done




Figure : That's how proper flow control should look like

Alternatives?



- LEDBAT tries to achieve a low, constant delay: Works, but not good on fairness
- CurveCP's flow control is still “a lot of research”
- Therefore, something new has to be done




Figure : That's how proper flow control should look like

net2o Flow Control




Figure : Measure the bottleneck using a burst of packets

Client Measures, Server Sets Rate



Client records the *time* of the first and last packet in a burst, and calculates the achieved rate for received packets, extrapolating to the achievable rate including the dropped packets. This results in the requested *rate*.

$$rate := \Delta t * \frac{burstlen}{packets}$$

Server would simply use this rate

Client Measures, Server Sets Rate



Client records the *time* of the first and last packet in a burst, and calculates the achieved rate for received packets, extrapolating to the achievable rate including the dropped packets. This results in the requested *rate*.

$$rate := \Delta t * \frac{burstlen}{packets}$$

Server would simply use this rate

Fairness



Fairness means that concurrent connections achieve about the same data rate, sharing the same line in a fair way.

- Ideally, a router/switch would schedule buffered packets round-robin, giving each connection a fair share of the bandwidth (fair queuing). That would change the calculated rate appropriately, and also be a big relief for current TCP buffer bloat symptoms, as each connection would have its private buffer to fill up.
- Unfortunately, routers use a single FIFO policy for all connections
- Finding a sufficiently stable algorithm to provide fairness
- We want to adopt to new situations as fast as possible, there's no point in anything slow. Especially on wireless connections, achievable rate changes are not only related to traffic.

Fairness



Fairness means that concurrent connections achieve about the same data rate, sharing the same line in a fair way.

- Ideally, a router/switch would schedule buffered packets round-robin, giving each connection a fair share of the bandwidth (fair queuing). That would change the calculated rate appropriately, and also be a big relief for current TCP buffer bloat symptoms, as each connection would have its private buffer to fill up.
- Unfortunately, routers use a single FIFO policy for all connections
- Finding a sufficiently stable algorithm to provide fairness
- We want to adopt to new situations as fast as possible, there's no point in anything slow. Especially on wireless connections, achievable rate changes are not only related to traffic.

Fairness



Fairness means that concurrent connections achieve about the same data rate, sharing the same line in a fair way.

- Ideally, a router/switch would schedule buffered packets round-robin, giving each connection a fair share of the bandwidth (fair queuing). That would change the calculated rate appropriately, and also be a big relief for current TCP buffer bloat symptoms, as each connection would have its private buffer to fill up.
- Unfortunately, routers use a single FIFO policy for all connections
- Finding a sufficiently stable algorithm to provide fairness
- We want to adopt to new situations as fast as possible, there's no point in anything slow. Especially on wireless connections, achievable rate changes are not only related to traffic.

Fairness



Fairness means that concurrent connections achieve about the same data rate, sharing the same line in a fair way.

- Ideally, a router/switch would schedule buffered packets round-robin, giving each connection a fair share of the bandwidth (fair queuing). That would change the calculated rate appropriately, and also be a big relief for current TCP buffer bloat symptoms, as each connection would have its private buffer to fill up.
- Unfortunately, routers use a single FIFO policy for all connections
- Finding a sufficiently stable algorithm to provide fairness
- We want to adopt to new situations as fast as possible, there's no point in anything slow. Especially on wireless connections, achievable rate changes are not only related to traffic.

Fairness



Fairness means that concurrent connections achieve about the same data rate, sharing the same line in a fair way.

- Ideally, a router/switch would schedule buffered packets round-robin, giving each connection a fair share of the bandwidth (fair queuing). That would change the calculated rate appropriately, and also be a big relief for current TCP buffer bloat symptoms, as each connection would have its private buffer to fill up.
- Unfortunately, routers use a single FIFO policy for all connections
- Finding a sufficiently stable algorithm to provide fairness
- We want to adopt to new situations as fast as possible, there's no point in anything slow. Especially on wireless connections, achievable rate changes are not only related to traffic.

net2o Flow Control — Fair Router




Figure : Fair queuing results in correct measurement of available bandwidth

net2o Flow Control — FIFO Router




Figure : Unfair FIFO queuing results in twice the available bandwidth calculated



Fairness I



- To improve stability of unfair queued packets, we need to improve that P regulator (proportional to measured rate) to a full PID regulator
- The integral part is the accumulated slack (in the buffer), which we want to keep low, and the D part is growing/reducing this slack from one measurement to the next
- We use both parts to decrease the sending rate, and thereby achieve better fairness
- The I part is used to exponentially lengthen the rate Δt with increasing slack up to a maximum factor of 16.

$$s_{exp} = 2^{\frac{slack}{T}} \quad \text{where } T = \max(10ms, \max(slacks))$$



Fairness I



- To improve stability of unfair queued packets, we need to improve that P regulator (proportional to measured rate) to a full PID regulator
- The integral part is the accumulated slack (in the buffer), which we want to keep low, and the D part is growing/reducing this slack from one measurement to the next
- We use both parts to decrease the sending rate, and thereby achieve better fairness
- The I part is used to exponentially lengthen the rate Δt with increasing slack up to a maximum factor of 16.

$$s_{exp} = 2^{\frac{slack}{T}} \quad \text{where } T = \max(10ms, \max(slacks))$$

Fairness I



- To improve stability of unfair queued packets, we need to improve that P regulator (proportional to measured rate) to a full PID regulator
- The integral part is the accumulated slack (in the buffer), which we want to keep low, and the D part is growing/reducing this slack from one measurement to the next
- We use both parts to decrease the sending rate, and thereby achieve better fairness
- The I part is used to exponentially lengthen the rate Δt with increasing slack up to a maximum factor of 16.

$$s_{exp} = 2^{\frac{slack}{T}} \quad \text{where } T = \max(10ms, \max(slacks))$$

Fairness I



- To improve stability of unfair queued packets, we need to improve that P regulator (proportional to measured rate) to a full PID regulator
- The integral part is the accumulated slack (in the buffer), which we want to keep low, and the D part is growing/reducing this slack from one measurement to the next
- We use both parts to decrease the sending rate, and thereby achieve better fairness
- The I part is used to exponentially lengthen the rate Δt with increasing slack up to a maximum factor of 16.

$$s_{exp} = 2^{\frac{slack}{T}} \quad \text{where } T = \max(10ms, \max(slacks))$$

Fairness I



- To improve stability of unfair queued packets, we need to improve that P regulator (proportional to measured rate) to a full PID regulator
- The integral part is the accumulated slack (in the buffer), which we want to keep low, and the D part is growing/reducing this slack from one measurement to the next
- We use both parts to decrease the sending rate, and thereby achieve better fairness
- The I part is used to exponentially lengthen the rate Δt with increasing slack up to a maximum factor of 16.

$$s_{exp} = 2^{\frac{slack}{T}} \quad \text{where } T = \max(10ms, \max(slacks))$$

Fairness D



- To measure the differential term, we measure how much the slack grows (a Δt value) from the first to the last burst we do for one measurement cycle (4 bursts by default, first packet to first packet of each burst)
- This is multiplied by the total packets in flight (head of the sender queue vs. acknowledged packet), divided by the packets within the measured interval
- A low-pass filter is applied to the obtained D to prevent from speeding up too fast, with one round trip delay as time constant
- $\max(\text{slacks})/10ms$ is used to determine how aggressive this algorithm is
- Add the obtained Δt both to the rate's Δt for one burst sequence and wait that time before starting the next burst sequence.

Fairness D



- To measure the differential term, we measure how much the slack grows (a Δt value) from the first to the last burst we do for one measurement cycle (4 bursts by default, first packet to first packet of each burst)
- This is multiplied by the total packets in flight (head of the sender queue vs. acknowledged packet), divided by the packets within the measured interval
- A low-pass filter is applied to the obtained D to prevent from speeding up too fast, with one round trip delay as time constant
- $\max(\text{slacks})/10ms$ is used to determine how aggressive this algorithm is
- Add the obtained Δt both to the rate's Δt for one burst sequence and wait that time before starting the next burst sequence.

Fairness D



- To measure the differential term, we measure how much the slack grows (a Δt value) from the first to the last burst we do for one measurement cycle (4 bursts by default, first packet to first packet of each burst)
- This is multiplied by the total packets in flight (head of the sender queue vs. acknowledged packet), divided by the packets within the measured interval
- A low-pass filter is applied to the obtained D to prevent from speeding up too fast, with one round trip delay as time constant
- $\max(\text{slacks})/10\text{ms}$ is used to determine how aggressive this algorithm is
- Add the obtained Δt both to the rate's Δt for one burst sequence and wait that time before starting the next burst sequence.

Fairness D



- To measure the differential term, we measure how much the slack grows (a Δt value) from the first to the last burst we do for one measurement cycle (4 bursts by default, first packet to first packet of each burst)
- This is multiplied by the total packets in flight (head of the sender queue vs. acknowledged packet), divided by the packets within the measured interval
- A low-pass filter is applied to the obtained D to prevent from speeding up too fast, with one round trip delay as time constant
- $\max(\text{slacks})/10ms$ is used to determine how aggressive this algorithm is
- Add the obtained Δt both to the rate's Δt for one burst sequence and wait that time before starting the next burst sequence.

Fairness D



- To measure the differential term, we measure how much the slack grows (a Δt value) from the first to the last burst we do for one measurement cycle (4 bursts by default, first packet to first packet of each burst)
- This is multiplied by the total packets in flight (head of the sender queue vs. acknowledged packet), divided by the packets within the measured interval
- A low-pass filter is applied to the obtained D to prevent from speeding up too fast, with one round trip delay as time constant
- $\max(\text{slacks})/10ms$ is used to determine how aggressive this algorithm is
- Add the obtained Δt both to the rate's Δt for one burst sequence and wait that time before starting the next burst sequence.

VDSL




Figure : One connection on a VDSL-50 line

VDSL, Congestion




Figure : One of four connections on a VDSL-50 line



Unreliable Air Cable (WLAN)

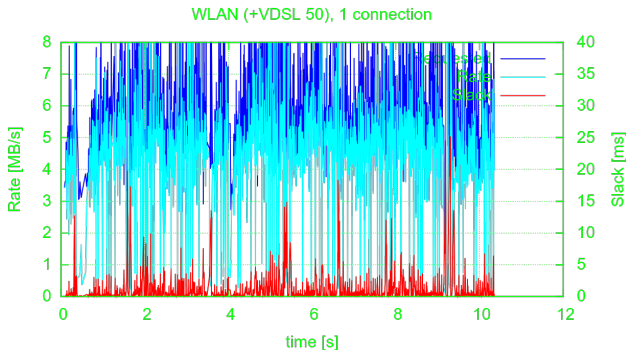


Figure : Single connection using WLAN



Unreliable Air Cable, Congestion




Figure : One of four connections using WLAN

LAN, 1GBE

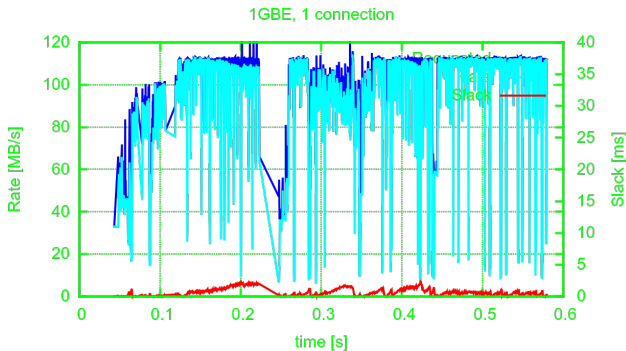


Figure : Single connection using 1GBE

LAN 1GBE, Congestion (4 servers)




Figure : One of four connections using 1GBE



LAN 1GBE, Congestion (1 server)




Figure : One of four connections using 1GBE, fair queuing

Data and Commands



- Data of several files/streams can be transferred interleaving, so a single connection can do multiple things in parallel
- Commands are send in command blocks, i.e. there is not just one command per block, but a sequence of commands!
- Commands are encoded like protobuf, i.e. 7 bits per byte, and if the MSB of the byte is 1, there's another byte to follow (allowing arbitrary many commands)
- The command “machine” is a stack architecture.
- The command VM is object oriented, i.e. commands are messages to objects
- The command interpreter itself is extremely simple

Data and Commands



- Data of several files/streams can be transferred interleaving, so a single connection can do multiple things in parallel
- Commands are send in command blocks, i.e. there is not just one command per block, but a sequence of commands!
- Commands are encoded like protobuf, i.e. 7 bits per byte, and if the MSB of the byte is 1, there's another byte to follow (allowing arbitrary many commands)
- The command "machine" is a stack architecture.
- The command VM is object oriented, i.e. commands are messages to objects
- The command interpreter itself is extremely simple

Data and Commands



- Data of several files/streams can be transferred interleaving, so a single connection can do multiple things in parallel
- Commands are send in command blocks, i.e. there is not just one command per block, but a sequence of commands!
- Commands are encoded like protobuf, i.e. 7 bits per byte, and if the MSB of the byte is 1, there's another byte to follow (allowing arbitrary many commands)
- The command “machine” is a stack architecture.
- The command VM is object oriented, i.e. commands are messages to objects
- The command interpreter itself is extremely simple

Data and Commands



- Data of several files/streams can be transferred interleaving, so a single connection can do multiple things in parallel
- Commands are send in command blocks, i.e. there is not just one command per block, but a sequence of commands!
- Commands are encoded like protobuf, i.e. 7 bits per byte, and if the MSB of the byte is 1, there's another byte to follow (allowing arbitrary many commands)
- The command “machine” is a stack architecture.
- The command VM is object oriented, i.e. commands are messages to objects
- The command interpreter itself is extremely simple

Data and Commands



- Data of several files/streams can be transferred interleaving, so a single connection can do multiple things in parallel
- Commands are send in command blocks, i.e. there is not just one command per block, but a sequence of commands!
- Commands are encoded like protobuf, i.e. 7 bits per byte, and if the MSB of the byte is 1, there's another byte to follow (allowing arbitrary many commands)
- The command “machine” is a stack architecture.
- The command VM is object oriented, i.e. commands are messages to objects
- The command interpreter itself is extremely simple

Data and Commands



- Data of several files/streams can be transferred interleaving, so a single connection can do multiple things in parallel
- Commands are send in command blocks, i.e. there is not just one command per block, but a sequence of commands!
- Commands are encoded like protobuf, i.e. 7 bits per byte, and if the MSB of the byte is 1, there's another byte to follow (allowing arbitrary many commands)
- The command “machine” is a stack architecture.
- The command VM is object oriented, i.e. commands are messages to objects
- The command interpreter itself is extremely simple

Example: Download three files



reading three files

```
0 file-id "net2o.fs" 0
  open-file get-size get-stat endwith
1 file-id "data/2011-05-13_11-26-57-small.jpg" 0
  open-file get-size get-stat endwith
2 file-id "data/2011-05-20_17-01-12-small.jpg" 0
  open-file get-size get-stat endwith
```

Reading Files: Reply



reading three files: replies

```
0 file-id 12B9A set-size
  138D607CB83D0F06 1A4 set-stat endwith
1 file-id 9C65C set-size
  13849CAE1F3B6EA8 1A4 set-stat endwith
2 file-id 9D240 set-size
  13849CAE2643FDCC 1A4 set-stat endwith
```

Basics



- Five data types: Integer (64 bits signed+unsigned), flag, string (generic byte array), IEEE double float, objects
- Instructions and data encoding derived from Protobuf (7 bits per byte, MSB=1 means “data continues”, most significant part first)
- Four stacks: integer, float, objects, strings
- `endwith` and `endcmd` for ending object message blocks and commands
- `oswap` to transfer the current object to the object stack, to be inserted in the outer object
- `words` for reflection (words are listed with token number, identifier and stack effect to make automatic bindings possible)

Basics



- Five data types: Integer (64 bits signed+unsigned), flag, string (generic byte array), IEEE double float, objects
- Instructions and data encoding derived from Protobuf (7 bits per byte, MSB=1 means “data continues”, most significant part first)
- Four stacks: integer, float, objects, strings
- `endwith` and `endcmd` for ending object message blocks and commands
- `oswap` to transfer the current object to the object stack, to be inserted in the outer object
- `words` for reflection (words are listed with token number, identifier and stack effect to make automatic bindings possible)

Basics



- Five data types: Integer (64 bits signed+unsigned), flag, string (generic byte array), IEEE double float, objects
- Instructions and data encoding derived from Protobuf (7 bits per byte, MSB=1 means “data continues”, most significant part first)
- Four stacks: integer, float, objects, strings
- `endwith` and `endcmd` for ending object message blocks and commands
- `oswap` to transfer the current object to the object stack, to be inserted in the outer object
- `words` for reflection (words are listed with token number, identifier and stack effect to make automatic bindings possible)

Basics



- Five data types: Integer (64 bits signed+unsigned), flag, string (generic byte array), IEEE double float, objects
- Instructions and data encoding derived from Protobuf (7 bits per byte, MSB=1 means “data continues”, most significant part first)
- Four stacks: integer, float, objects, strings
- `endwith` and `endcmd` for ending object message blocks and commands
- `oswap` to transfer the current object to the object stack, to be inserted in the outer object
- `words` for reflection (words are listed with token number, identifier and stack effect to make automatic bindings possible)

Basics



- Five data types: Integer (64 bits signed+unsigned), flag, string (generic byte array), IEEE double float, objects
- Instructions and data encoding derived from Protobuf (7 bits per byte, MSB=1 means “data continues”, most significant part first)
- Four stacks: integer, float, objects, strings
- `endwith` and `endcmd` for ending object message blocks and commands
- `oswap` to transfer the current object to the object stack, to be inserted in the outer object
- `words` for reflection (words are listed with token number, identifier and stack effect to make automatic bindings possible)

Basics



- Five data types: Integer (64 bits signed+unsigned), flag, string (generic byte array), IEEE double float, objects
- Instructions and data encoding derived from Protobuf (7 bits per byte, MSB=1 means “data continues”, most significant part first)
- Four stacks: integer, float, objects, strings
- `endwith` and `endcmd` for ending object message blocks and commands
- `oswap` to transfer the current object to the object stack, to be inserted in the outer object
- `words` for reflection (words are listed with token number, identifier and stack effect to make automatic bindings possible)

Why binary encoding?



- Faster and simpler to parse (simpler means smaller attack vector)
- Ability to enter commands on the fly in text form through a frontend interpreter still exists
- Debugging with a de-tokenizer is also very easy
- Object-oriented approach makes writing application-specific logic extremely simple

Why binary encoding?



- Faster and simpler to parse (simpler means smaller attack vector)
- Ability to enter commands on the fly in text form through a frontend interpreter still exists
- Debugging with a de-tokenizer is also very easy
- Object-oriented approach makes writing application-specific logic extremely simple

Why binary encoding?



- Faster and simpler to parse (simpler means smaller attack vector)
- Ability to enter commands on the fly in text form through a frontend interpreter still exists
- Debugging with a de-tokenizer is also very easy
- Object-oriented approach makes writing application-specific logic extremely simple

Why binary encoding?



- Faster and simpler to parse (simpler means smaller attack vector)
- Ability to enter commands on the fly in text form through a frontend interpreter still exists
- Debugging with a de-tokenizer is also very easy
- Object-oriented approach makes writing application-specific logic extremely simple

Why a programming language as data?

Lemma: every glue logic will become Turing complete



- Implement only the things you need — but you shouldn't have to implement more than *one* generic interpreter
- Typical idea of sending remote procedure calls: serialize the entire object (with subobjects), and call a function on that object
- Net2o idea (derived from ONF): Keep the entire object synchronized by sending only the changes to it — these changes are simple messages (setters)
- This allows multi-message passing, and reduces latency

Why a programming language as data?

Lemma: every glue logic will become Turing complete



- Implement only the things you need — but you shouldn't have to implement more than *one* generic interpreter
- Typical idea of sending remote procedure calls: serialize the entire object (with subobjects), and call a function on that object
- Net2o idea (derived from ONF): Keep the entire object synchronized by sending only the changes to it — these changes are simple messages (setters)
- This allows multi-message passing, and reduces latency

Why a programming language as data?

Lemma: every glue logic will become Turing complete



- Implement only the things you need — but you shouldn't have to implement more than *one* generic interpreter
- Typical idea of sending remote procedure calls: serialize the entire object (with subobjects), and call a function on that object
- Net2o idea (derived from ONF): Keep the entire object synchronized by sending only the changes to it — these changes are simple messages (setters)
- This allows multi-message passing, and reduces latency

Why a programming language as data?

Lemma: every glue logic will become Turing complete



- Implement only the things you need — but you shouldn't have to implement more than *one* generic interpreter
- Typical idea of sending remote procedure calls: serialize the entire object (with subobjects), and call a function on that object
- Net2o idea (derived from ONF): Keep the entire object synchronized by sending only the changes to it — these changes are simple messages (setters)
- This allows multi-message passing, and reduces latency

Security



Lemma: every sufficiently complex format can be exploited

Therefore stick to a very simple format, i.e.: simplify and factor the code

Interpreter (pseudocode)

```
get_cmd:  p -> p cmd
          cmd = *p++;
n2cmd:   n -> call
          call = o ? token_table[n] : setup_table[n];
cmd_dispatch:  p -> p
              p, cmd = get_cmd(p);
              invoke(n2cmd(n));
cmd-loop:  p -> void
          do { p = cmd_dispatch(p);
              } while(len(p) > 0)
```

Security



Lemma: every sufficiently complex format can be exploited

Therefore stick to a very simple format, i.e.: simplify and factor the code

Interpreter (pseudocode)

```
get_cmd:  p -> p cmd
          cmd = *p++;
n2cmd:   n -> call
          call = o ? token_table[n] : setup_table[n];
cmd_dispatch:  p -> p
              p, cmd = get_cmd(p);
              invoke(n2cmd(n));
cmd-loop:  p -> void
           do { p = cmd_dispatch(p);
             } while(len(p) > 0)
```

Distributed Data



- Following the “everything is a file” principle, every data object is a file
- Data objects are accessed by their hash. The associated metadata are “tags”
- Metadata is organized as a distributed prefix hash tree
- Efficient distribution of data is important!

Distributed Data



- Following the “everything is a file” principle, every data object is a file
- Data objects are accessed by their hash. The associated metadata are “tags”
- Metadata is organized as a distributed prefix hash tree
- Efficient distribution of data is important!

Distributed Data



- Following the “everything is a file” principle, every data object is a file
- Data objects are accessed by their hash. The associated metadata are “tags”
- Metadata is organized as a distributed prefix hash tree
- Efficient distribution of data is important!

Distributed Data



- Following the “everything is a file” principle, every data object is a file
- Data objects are accessed by their hash. The associated metadata are “tags”
- Metadata is organized as a distributed prefix hash tree
- Efficient distribution of data is important!

Tree Distribution Network




Figure : Avalanche distribution with quad-tree of depth 2

Possible Performance



- Trees with a bigger base reduce latency. Example: To transfer a Justin Bieber tweet to 50 million followers, a binary tree needs 25.5 hops on average, a quad-tree 12.8 hops, and an oct-tree 8.5 hops.
- A typical domestic (inside e.g. Germany) hop-to-hop time is just 20ms. International hops can be in the order of 250ms. Assuming there is only one international hop in the chain, the latency to distribute Justin Bieber's babbling is typically just 500ms in a quad-tree.
- Rule of thumb: $bandwidth = latency$, i.e. if it takes 20ms from hop to hop, each node should replicate data for 20ms — if we make the tree wider, the linear effort of replicating data will dominate transfer time, if we make the tree more narrow, the hop-to-hop time will dominate.
- The tree-like graph greatly reduces the number of nodes to know

Possible Performance



- Trees with a bigger base reduce latency. Example: To transfer a Justin Bieber tweet to 50 million followers, a binary tree needs 25.5 hops on average, a quad-tree 12.8 hops, and an oct-tree 8.5 hops.
- A typical domestic (inside e.g. Germany) hop-to-hop time is just 20ms. International hops can be in the order of 250ms. Assuming there is only one international hop in the chain, the latency to distribute Justin Bieber's babbling is typically just 500ms in a quad-tree.
- Rule of thumb: *bandwidth = latency*, i.e. if it takes 20ms from hop to hop, each node should replicate data for 20ms — if we make the tree wider, the linear effort of replicating data will dominate transfer time, if we make the tree more narrow, the hop-to-hop time will dominate.
- The tree-like graph greatly reduces the number of nodes to know

Possible Performance



- Trees with a bigger base reduce latency. Example: To transfer a Justin Bieber tweet to 50 million followers, a binary tree needs 25.5 hops on average, a quad-tree 12.8 hops, and an oct-tree 8.5 hops.
- A typical domestic (inside e.g. Germany) hop-to-hop time is just 20ms. International hops can be in the order of 250ms. Assuming there is only one international hop in the chain, the latency to distribute Justin Bieber's babbling is typically just 500ms in a quad-tree.
- Rule of thumb: $bandwidth = latency$, i.e. if it takes 20ms from hop to hop, each node should replicate data for 20ms — if we make the tree wider, the linear effort of replicating data will dominate transfer time, if we make the tree more narrow, the hop-to-hop time will dominate.
- The tree-like graph greatly reduces the number of nodes to know

Possible Performance



- Trees with a bigger base reduce latency. Example: To transfer a Justin Bieber tweet to 50 million followers, a binary tree needs 25.5 hops on average, a quad-tree 12.8 hops, and an oct-tree 8.5 hops.
- A typical domestic (inside e.g. Germany) hop-to-hop time is just 20ms. International hops can be in the order of 250ms. Assuming there is only one international hop in the chain, the latency to distribute Justin Bieber's babbling is typically just 500ms in a quad-tree.
- Rule of thumb: $bandwidth = latency$, i.e. if it takes 20ms from hop to hop, each node should replicate data for 20ms — if we make the tree wider, the linear effort of replicating data will dominate transfer time, if we make the tree more narrow, the hop-to-hop time will dominate.
- The tree-like graph greatly reduces the number of nodes to know

Distributed Hashes



- Most DHT approaches have poor performance
- Still working out what is both simple and fast
- Model: Directory servers know how stores which subset of the hashes
- Replicated servers send updates through a distribution tree (low latency mirroring)

Distributed Hashes



- Most DHT approaches have poor performance
- Still working out what is both simple and fast
- Model: Directory servers know how stores which subset of the hashes
- Replicated servers send updates through a distribution tree (low latency mirroring)

Distributed Hashes



- Most DHT approaches have poor performance
- Still working out what is both simple and fast
- Model: Directory servers know how stores which subset of the hashes
- Replicated servers send updates through a distribution tree (low latency mirroring)

Distributed Hashes



- Most DHT approaches have poor performance
- Still working out what is both simple and fast
- Model: Directory servers know how stores which subset of the hashes
- Replicated servers send updates through a distribution tree (low latency mirroring)



Content or Apps?



- The current web is defined by content — web apps (JavaScript) are an afterthought
- Therefore, the application logic is usually on the server side
- This doesn't work for a P2P network!
- Content is structured text, images, videos, music, etc.



Content or Apps?



- The current web is defined by content — web apps (JavaScript) are an afterthought
- Therefore, the application logic is usually on the server side
- This doesn't work for a P2P network!
- Content is structured text, images, videos, music, etc.



Content or Apps?



- The current web is defined by content — web apps (JavaScript) are an afterthought
- Therefore, the application logic is usually on the server side
- This doesn't work for a P2P network!
- Content is structured text, images, videos, music, etc.



Content or Apps?



- The current web is defined by content — web apps (JavaScript) are an afterthought
- Therefore, the application logic is usually on the server side
- This doesn't work for a P2P network!
- Content is structured text, images, videos, music, etc.

App-Centric World



- There's a phenomenon I call "Turing creep": Every sufficiently complex system contains a user-accessible Turing-complete language
- Corollary: Every efficient sufficiently complex system can execute native machine code
- The application logic is to present the data; data itself is as above: structured text, images, videos, music, etc.
- Executing (especially efficient) code from the net raises obvious questions about security



App-Centric World



- There's a phenomenon I call "Turing creep": Every sufficiently complex system contains a user-accessible Turing-complete language
- Corollary: Every efficient sufficiently complex system can execute native machine code
- The application logic is to present the data; data itself is as above: structured text, images, videos, music, etc.
- Executing (especially efficient) code from the net raises obvious questions about security



App-Centric World



- There's a phenomenon I call "Turing creep": Every sufficiently complex system contains a user-accessible Turing-complete language
- Corollary: Every efficient sufficiently complex system can execute native machine code
- The application logic is to present the data; data itself is as above: structured text, images, videos, music, etc.
- Executing (especially efficient) code from the net raises obvious questions about security

App-Centric World



- There's a phenomenon I call "Turing creep": Every sufficiently complex system contains a user-accessible Turing-complete language
- Corollary: Every efficient sufficiently complex system can execute native machine code
- The application logic is to present the data; data itself is as above: structured text, images, videos, music, etc.
- Executing (especially efficient) code from the net raises obvious questions about security

How to securely execute code?



There are several options tried; as usual, things are broken:

1. Execute code in a controlled secure VM, see for example Java. This is broken by design, as securing something from the inside doesn't work.
2. Execute code in a sandbox. This has shown as more robust, depending on how complex the outside of the sandbox is.
3. Public inspection of code. This is how the open source world works, but the underhanded C contest shows that inspection is tricky.
4. Scan for known evil code. This is the security industry's approach, and it is not working.
5. Code signing can work together with public inspection — but using it for accountability doesn't work

Therefore the choice is to sandbox public inspected code.

How to securely execute code?



There are several options tried; as usual, things are broken:

1. Execute code in a controlled secure VM, see for example Java. This is broken by design, as securing something from the inside doesn't work.
2. Execute code in a sandbox. This has shown as more robust, depending on how complex the outside of the sandbox is.
3. Public inspection of code. This is how the open source world works, but the underhanded C contest shows that inspection is tricky.
4. Scan for known evil code. This is the security industry's approach, and it is not working.
5. Code signing can work together with public inspection — but using it for accountability doesn't work

Therefore the choice is to sandbox public inspected code.

How to securely execute code?



There are several options tried; as usual, things are broken:

1. Execute code in a controlled secure VM, see for example Java. This is broken by design, as securing something from the inside doesn't work.
2. Execute code in a sandbox. This has shown as more robust, depending on how complex the outside of the sandbox is.
3. Public inspection of code. This is how the open source world works, but the underhanded C contest shows that inspection is tricky.
4. Scan for known evil code. This is the security industry's approach, and it is not working.
5. Code signing can work together with public inspection — but using it for accountability doesn't work

Therefore the choice is to sandbox public inspected code.

How to securely execute code?



There are several options tried; as usual, things are broken:

1. Execute code in a controlled secure VM, see for example Java. This is broken by design, as securing something from the inside doesn't work.
2. Execute code in a sandbox. This has shown as more robust, depending on how complex the outside of the sandbox is.
3. Public inspection of code. This is how the open source world works, but the underhanded C contest shows that inspection is tricky.
4. Scan for known evil code. This is the security industry's approach, and it is not working.
5. Code signing can work together with public inspection — but using it for accountability doesn't work

Therefore the choice is to sandbox public inspected code.

How to securely execute code?



There are several options tried; as usual, things are broken:

1. Execute code in a controlled secure VM, see for example Java. This is broken by design, as securing something from the inside doesn't work.
2. Execute code in a sandbox. This has shown as more robust, depending on how complex the outside of the sandbox is.
3. Public inspection of code. This is how the open source world works, but the underhanded C contest shows that inspection is tricky.
4. Scan for known evil code. This is the security industry's approach, and it is not working.
5. Code signing can work together with public inspection — but using it for accountability doesn't work

Therefore the choice is to sandbox public inspected code.

How to securely execute code?



There are several options tried; as usual, things are broken:

1. Execute code in a controlled secure VM, see for example Java. This is broken by design, as securing something from the inside doesn't work.
2. Execute code in a sandbox. This has shown as more robust, depending on how complex the outside of the sandbox is.
3. Public inspection of code. This is how the open source world works, but the underhanded C contest shows that inspection is tricky.
4. Scan for known evil code. This is the security industry's approach, and it is not working.
5. Code signing can work together with public inspection — but using it for accountability doesn't work

Therefore the choice is to sandbox public inspected code.

How to securely execute code?



There are several options tried; as usual, things are broken:

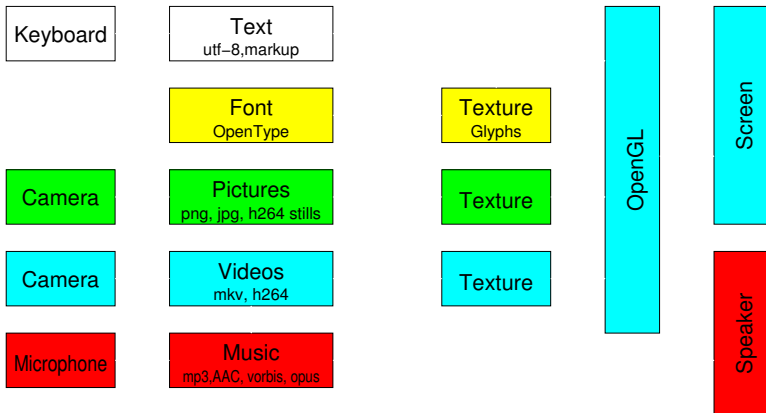
1. Execute code in a controlled secure VM, see for example Java. This is broken by design, as securing something from the inside doesn't work.
2. Execute code in a sandbox. This has shown as more robust, depending on how complex the outside of the sandbox is.
3. Public inspection of code. This is how the open source world works, but the underhanded C contest shows that inspection is tricky.
4. Scan for known evil code. This is the security industry's approach, and it is not working.
5. Code signing can work together with public inspection — but using it for accountability doesn't work

Therefore the choice is to sandbox public inspected code.



Formats&Requirements

How to display things



Why OpenGL?

OpenGL can do everything



OpenGL renders:

1. Triangles, lines, points — simple components
2. Textures and gradients
3. and uses shader programs — the most powerful thing in OpenGL from 2.0.

Real requirement: visualization of *any* data. OpenGL can do that.

Why OpenGL?

OpenGL can do everything



OpenGL renders:

1. Triangles, lines, points — simple components
2. Textures and gradients
3. and uses shader programs — the most powerful thing in OpenGL from 2.0.

Real requirement: visualization of *any* data. OpenGL can do that.

Why OpenGL?

OpenGL can do everything



OpenGL renders:

1. Triangles, lines, points — simple components
2. Textures and gradients
3. and uses shader programs — the most powerful thing in OpenGL from 2.0.

Real requirement: visualization of *any* data. OpenGL can do that.

Why OpenGL?

OpenGL can do everything



OpenGL renders:

1. Triangles, lines, points — simple components
2. Textures and gradients
3. and uses shader programs — the most powerful thing in OpenGL from 2.0.

Real requirement: visualization of *any* data. OpenGL can do that.

Why OpenGL?

OpenGL can do everything



OpenGL renders:

1. Triangles, lines, points — simple components
2. Textures and gradients
3. and uses shader programs — the most powerful thing in OpenGL from 2.0.

Real requirement: visualization of *any* data. OpenGL can do that.

How to connect the media?

Lemma: every glue logic will become Turing complete



- currently used glue: HTML+CSS+JavaScript
- containers with Flash, Java, ActiveX, PDF, Google's NaCl...
- conclusion: use a powerful tool right from start!
- browser: run-time and development tool for applications

How to connect the media?

Lemma: every glue logic will become Turing complete



- currently used glue: HTML+CSS+JavaScript
- containers with Flash, Java, ActiveX, PDF, Google's NaCl...
- conclusion: use a powerful tool right from start!
- browser: run-time and development tool for applications

How to connect the media?

Lemma: every glue logic will become Turing complete



- currently used glue: HTML+CSS+JavaScript
- containers with Flash, Java, ActiveX, PDF, Google's NaCl...
- conclusion: use a powerful tool right from start!
- browser: run-time and development tool for applications

How to connect the media?

Lemma: every glue logic will become Turing complete



- currently used glue: HTML+CSS+JavaScript
- containers with Flash, Java, ActiveX, PDF, Google's NaCl...
- conclusion: use a powerful tool right from start!
- browser: run-time and development tool for applications

Frameworks



- libsoil for images (PNG+JPEG loading into a texture)
- freetype-gl for fonts (TrueType/OpenType into a texture)
- OpenMAX on Android, gstreamer on Linux: videos into a texture
- MINOΣ2: Lightweight OpenGL-based widget library in Forth (still a lot of work in progress)

Frameworks



- libsoil for images (PNG+JPEG loading into a texture)
- freetype-gl for fonts (TrueType/OpenType into a texture)
- OpenMAX on Android, gstreamer on Linux: videos into a texture
- MINOΣ2: Lightweight OpenGL-based widget library in Forth (still a lot of work in progress)

Frameworks



- libsoil for images (PNG+JPEG loading into a texture)
- freetype-gl for fonts (TrueType/OpenType into a texture)
- OpenMAX on Android, gstreamer on Linux: videos into a texture
- MINOΣ2: Lightweight OpenGL-based widget library in Forth (still a lot of work in progress)

Frameworks



- libsoil for images (PNG+JPEG loading into a texture)
- freetype-gl for fonts (TrueType/OpenType into a texture)
- OpenMAX on Android, gstreamer on Linux: videos into a texture
- MINOΣ2: Lightweight OpenGL-based widget library in Forth (still a lot of work in progress)



For Further Reading I



BERND PAYSAN

net2o source repository and wiki

<http://fossil.net2o.de/net2o>



HEALTH & SAFETY EXECUTIVE HSE – UK

Out of control, 2nd edition 2003

<http://www.hse.gov.uk/pubns/priced/hsg238.pdf>



MARTIN CROXFORD and DR. RODERICK CHAPMAN

Correctness by Construction: A Manifesto for High-Integrity Software

[http://www.crosstalkonline.org/storage/
issue-archives/2005/200512/200512-Croxford.pdf](http://www.crosstalkonline.org/storage/issue-archives/2005/200512/200512-Croxford.pdf)