模

mó

https://scale.qihardware.org

2019 . Week 10 . Mar10 - Mar 16

This page left intentionally blank
to power your imagination
of what interesting art, ads,
sponsorship, standard
frontmatter or blank space
should be included in
future editions of scale.

模 mó

Mó means that which is left over after division. In algebra this is called the module or modulo operator.

https://en.wikipedia.org/wiki/Module_(mathematics)

ITNEXT BY LINKIT

[Part 1] How to build a Telegram Bot using Node.js (Building a Bookmark Manager Bot series)

Using Node.js and Vue.js

Mike Tasset

Jun 26, 2018 · 8 min read

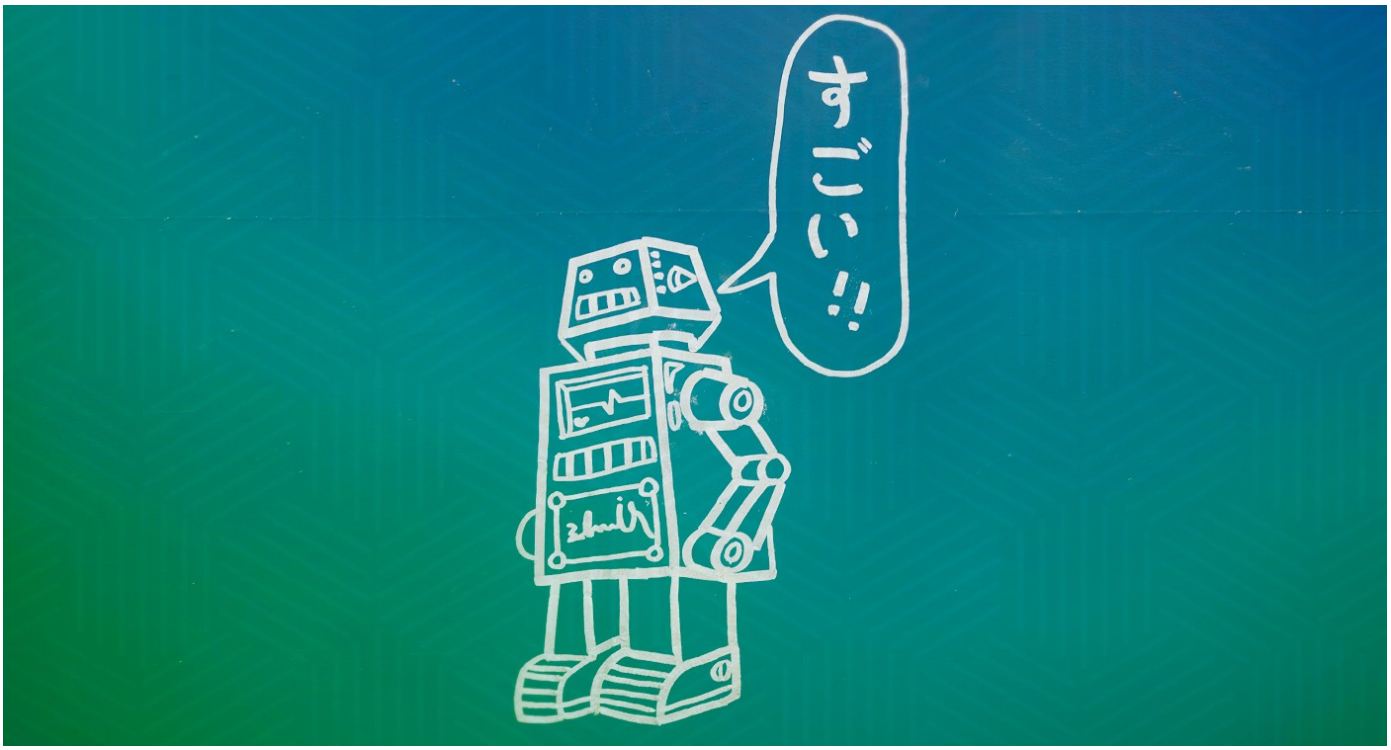

Original photo by

This article is part of a series on how to build a Telegram bot and save our bookmarks, deploy the bot and create a simple interface in Vue.js to manage our bookmarks. In this part we will take a look at how to build the bot.

Even though browsers are implementing shared bookmarks between devices and there are a bunch of services helping you keep track of your bookmarks, I always found something that I was not satisfied with. As you will see, it is actually very easy

found something that I was not satisfied with. As you will see, it is actually very easy to create your own bookmark manager which you can access from anywhere by creating a Telegram bot!

Before we start

I'm assuming that you have some sort of package manager installed, I will be using Yarn but NPM will work just as fine (however the commands for installing packages will be different). Because we will be using Node.js, you will, obviously, need to have it installed.

You can find the full code for the bot here on GitHub.

1: Let's get busy!

The first thing we need to do is creating our project folder, all we need for now is a project folder and a `index.js` file. In the Terminal go to your project folder, create a `package.json` using `init` and add/install the modules we will be using, `node-telegram-bot-api` , `open-graph-scraper` and `firebase` :

```
cd your-project-folder
yarn init
yarn add node-telegram-bot-api open-graph-scraper firebase
```

`node-telegram-bot-ap` is used to interact with the Telegram Api and we will use `open-graph-scraper` to easily scrap the OG data from our bookmarked websites. Because we will be using **Firebase** for storing the data you will also need the `firebase` module.

The BotFather

Luckily Telegram makes it very easy to create and manage bots. Add the BotFather to Telegram and start creating a new bot by sending the `/newbot` command. BotFather will guide you through the process. He will give you a *token* which you will need to use the **HTTP API**.

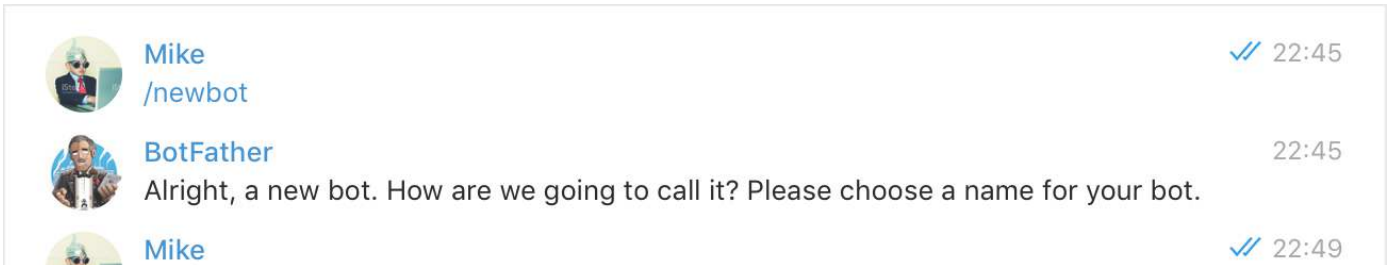

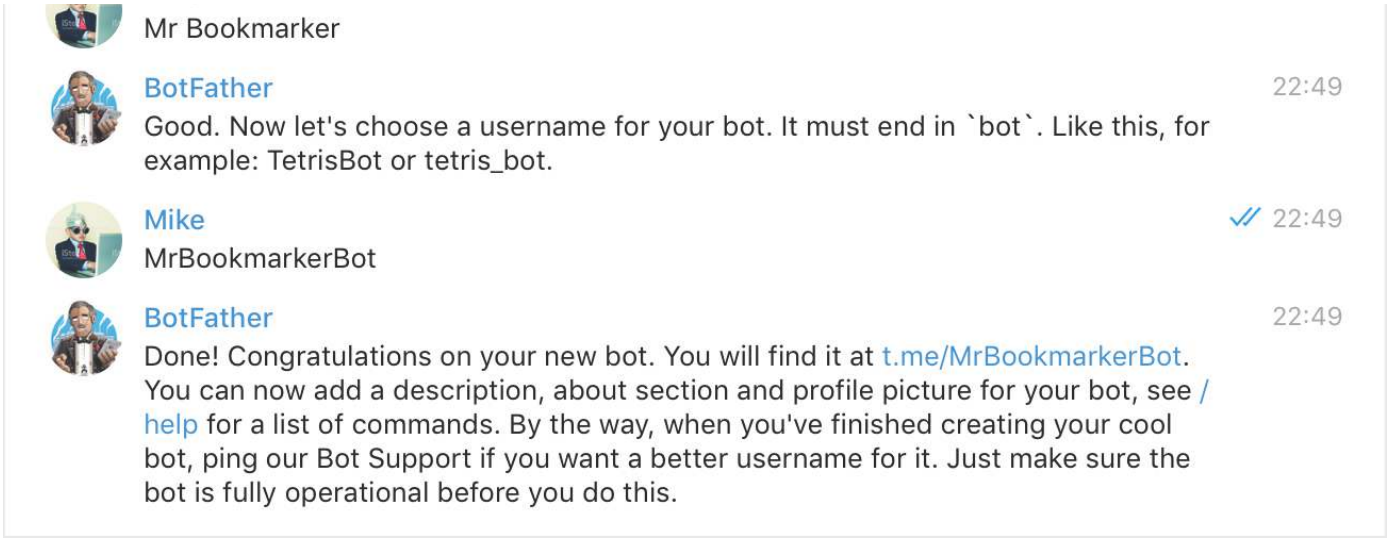

2: Let's finally build something

We are all set up now and ready to start programming. Your project should have a `node_modules` folder containing the required modules. First we need make sure to load our modules using Node.js's `require()` in the `index.js` file:

```
const TelegramBot = require('node-telegram-bot-api');
const ogs = require('open-graph-scraper');
const firebase = require('firebase');
```

Next we need to create some constants to hold some of our bot's information. Constants are similar to variables, except their value cannot change (see the MDN page for more information):

```
const token = 'your-token';
const bot = new TelegramBot(token, {polling: true});
```

Make sure to replace `your-token` with your actual token, else your bot won't connect! The constant `bot` creates a new bot object from the *Node.js Telegram Bot API* using the token constant for the token.

We now have our basics set up! But it is a bit boring when it does not do anything. Let's make it talk back to us! To make our bot do something we need to capture an event so it can trigger the bot to do whatever we instruct it to do.

```
bot.on('message', (msg) => {
  bot.sendMessage(msg.chat.id, 'Ill have the tuna. No crust.');
});
```

Now if the bot receives a message, it will reply back to us! In my example it is ordering a sandwich, but you can use any message you want of course. Let's check if it is working!

3: Testing the bot

Because we are using Node.js it is very easy to test our bot. Just go into the terminal and run the following commands:

```
cd your-project-folder
node index.js
```

Node will now run our script for us and we can go test our bot. Add your bot as a friend on Telegram and try sending it a message, if everything is working correctly you will receive a message back from your bot! Now before we continue let's stop Node running our script by pressing *Control + C* while in the Terminal.

4: Connecting to Firebase

Our bot can receive our messages and reply back, but if we want it to bookmark our links we will need to connect to a database. With Firebase we can create a cloud based database and sync/read our data real-time. This makes it ideal for our bot (actually Firebase can do more than just that, this article gives a good overview). Sign into your Google account and go to Firebase console. Click on **Add project,** choose a name for your project and walk trough the steps to step up your database.

When your project is ready you will be automatically redirected to the dashboard. In the left side menu click on **Database** under **Develop**. Create a new **Realtime Database** and start in **Test Mode.** Great, we now have our database ready as well. To get the information needed to connect to the database go to **Project Overview** from the left side menu and click on **Add Firebase to your web app**. You can find your credentials in the popup (in the object `config`).

You can remove the *message* event we created for testing as we won't need it

anymore. To connect to your Firebase add the following code to your script and add the credential we just got from the Project settings:

```
const app = firebase.initializeApp({
  apiKey: "your-api-key",
  authDomain: "your-auth-domain",
  databaseURL: "your-database-url",
  projectId: "your-project-id",
  storageBucket: "your-storage-bucket",
  messagingSenderId: "your-sender-id"
});
```

Next we create two constants, `ref` to hold the reference to our database, and `sitesRef` which holds the reference to the *child* “sites” in the database which we created earlier in the Firebase console:

```
const ref = firebase.database().ref();
const sitesRef = ref.child("sites");
```

Because we want to add more functions to our bot in the future we will create a command to make our bot bookmark a link, let’s call it `/bookmark` .

```
let siteUrl;

bot.onText(/\/bookmark (.+)/, (msg, match) => {
  siteUrl = match[1];

  bot.sendMessage(msg.chat.id,'Got it, in which category?', {
    reply_markup: {
      inline_keyboard: [[
        {
          text: 'Development',
          callback_data: 'development'
        },{
          text: 'Music',
          callback_data: 'music'
        },{
          text: 'Cute monkeys',
          callback_data: 'cute-monkeys'
        }
      ]]
    }
```

```
  });
});
```

We created an empty let which we will use to store the received URL. Then we use the Telegram API to catch the `onText` event and use a Regular Expressions to separate the URL from the command. This will create an array from which we will select the second value (which is `[1]` because arrays start from zero) and set the variable `siteUrl` to this value.

The bot will reply with “Got it, in which category?” (you can of course use any message you want here) and uses the Telegram API to pull up the inline keyboard (because typing is too much work, we don’t want too tired). The inline keyboard is a build in function in Telegram which allows us to reply by pressing a button, rather then having to type the answers! The keyboard contains three buttons which are our three categories. `text:` is the value that will be shown in the button, and `callback_data:` is the value we can use in our script. Let’s save and run our script to test if our bot is replying by running `node index.js` from the Terminal and sending it an URL.

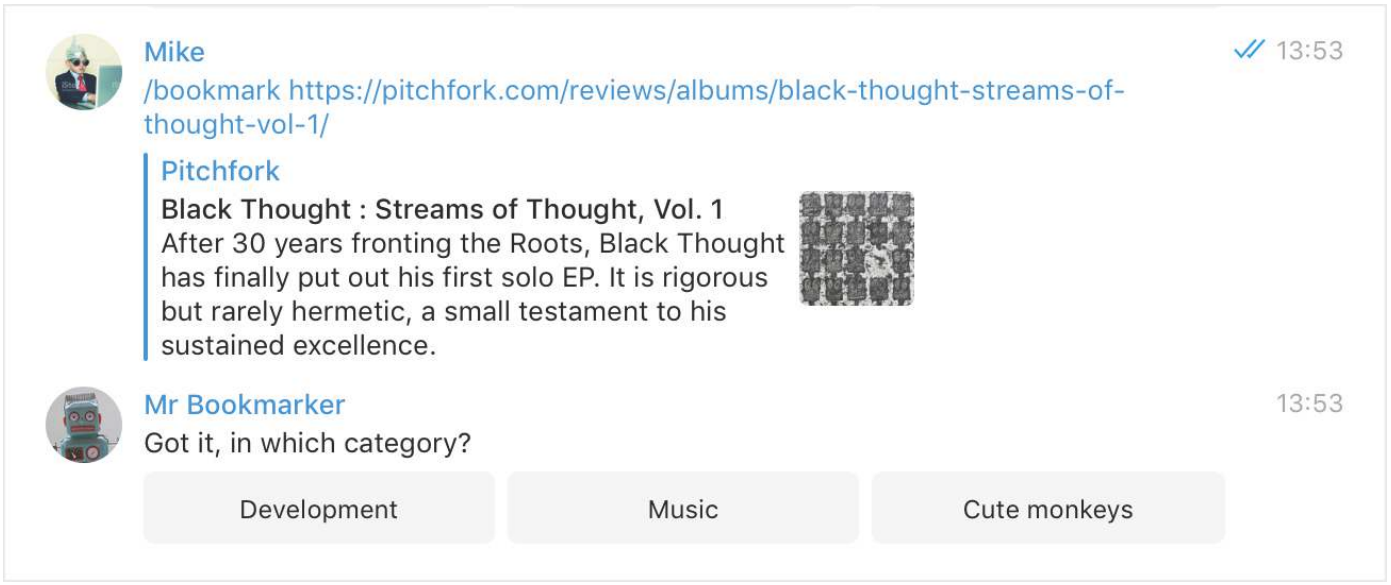

Our bot is working! We can now use Telegram’s inline keyboard to reply to our bot. Now we need the bot to scrape the OG data and save it:

```
bot.on("callback_query", (callbackQuery) => {
  const message = callbackQuery.message;

  ogs({'url': siteUrl}, function (error, results) {
    if(results.success) {
      sitesRef.push().set({
```

```
      sitesRef.push().set({
        name: results.data.ogSiteName,
        title: results.data.ogTitle,
        description: results.data.ogDescription,
        url: siteUrl,
        thumbnail: results.data.ogImage.url,
        category: callbackQuery.data
      });

      bot.sendMessage(message.chat.id,'Added \"' + results.data.ogTitle +'\" to category \"' +
callbackQuery.data + '\"!')
} else {
      sitesRef.push().set({
        url: siteUrl
      });
      bot.sendMessage(message.chat.id,'Added new website, but there was no OG data!');
    }
  });
});
```

If the bot receives a callback query (meaning that we used the inline keyboard to sent something to him in this case) it will receive this information and process it. The `callbackQuery` contains a bunch of information, but we only care about the message so we save that in a constant. `ogs` (Open Graph Scraper) will scrap the URL we passed it into a JSON file. If the scraping is successful (we check by using the `success` flag) it pushes the OG information, along with the site URL, to Firebase and sends us a reply letting use know it added the bookmark. Then it will trigger the bot to sent back a message, we select the `chat.id` from the `message` constant to get the chat ID so the bot knows to which chat it needs to send the reply.

If it was unable to scrape any OG data it will still push the site to our database and send us a different message letting us know he could not scrape any OG data, but still bookmarked the website. In my experience I would say about 95% of the sites I bookmarked worked fine, so it's not very common to get no information.

You did it!

Go into the terminal and start our script again to start the bot. Let's try it out by sending him a URL:

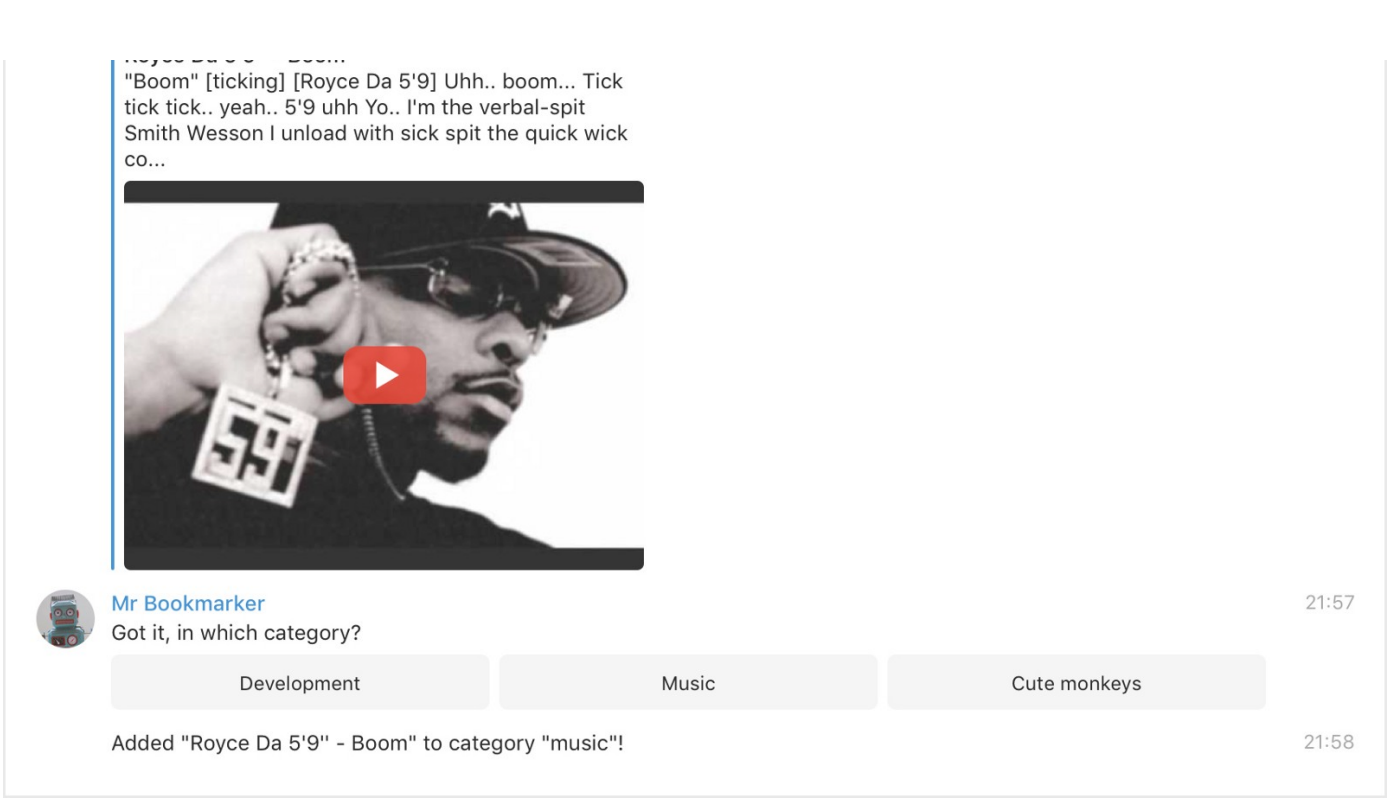

It's working!

It seems like everything working fine! The final thing to check is our database. Go to your Firebase Project and select **Database** from the left side menu. If everything is working correctly you will see your bookmark appear in the database (you might need to expand the childs using the little '+'-sign next to it).

Congratulation, you have now build your own bookmark assistant!

Deploying

Click here to go to the next part of the series, where we take a look at how to deploy our bot. In the future parts we will create a simple interface for managing our bookmarks using Vue.js!

Series index

Series Index

[Part 1] How to build a Telegram Bot using Node.js

[Part 2] How to deploy a Telegram Bot using Now

Next parts coming soon!

JavaScript Nodejs Telegram Web Development Chatbots

200 claps



7


Mike Tasset

Frontend Engineer at Birdman
Tokyo

Follow

ITNEXT

ITNEXT is a platform for IT developers & software engineers to share knowledge, connect, collaborate, learn and experience next-gen technologies.

Follow

Never miss a story from **ITNEXT**

GET UPDATES

Worlding Raga: 2 – What is a World?

March 5, 2019 By Ian Cheng

This entry is part 2 of 3 in the series Worlding Raga

Hi, I'm Ian Cheng. I'm an artist. Over the last six years, I've been creating a series of simulations that explore an agent's capacity to deal with an ever-changing environment. These works culminated in the Emissaries trilogy, which introduced a narrative agent — the emissary — whose motivation to enact a story was set into conflict with the open-ended chaos of a simulation. In the process, I began to see the edges of a new layer of artistic activity. One that could organize my base ingredients — deterministic stories and open-ended simulations — into something more than the sum of its parts. Something meaningful yet alive, bounded yet transforming. I've been calling this activity Worlding.

At Venkat's invitation, I'm contributing to Worlding Raga with the hope of further developing a literacy around Worlding. As a tourist of this part of the blogosphere,

Blogchains

Domestic Cozy
Elderblog Sutra
Infinite Machine
Mediocratopia
Predictable Identities
Refactorings Roundups
Weirding Diary
Worlding Raga

Crash Early, Crash Often

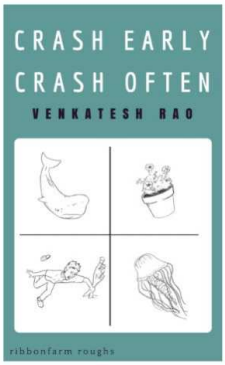

Be Slightly Evil

I've been drawn to the spiritual dimension of Ribbonfarm again and again. It is the side of Ribbonfarm that is hungry to identify phenomena in the wild that don't seem to die, and to name them so that they are enduring tools for others to see and act anew. This voluntary desire to surf chaos, metabolize it into new order, and then do it all over again, is sometimes called "walking with god." Maybe it's more like slouching with god around here. Either way, it is a spirit I resonate with and one that I believe is highly suited to Worlding.

First things first. What is a World?

We could say a World is something like a gated garden. A World has borders. A World has laws. A World has values. A World has dysfunction. A World can grow up. A World has members who live in it. A World gives its members permission to act differently than outside of it. A World incentivizes its members to keep it alive, often with the pleasures of its dysfunction. A World counts certain actions inside it as relevant and meaningful. A World undergoes reformations and disruptions. A World has mythic figures. A World is a container for all the possible stories of itself. A World manifests evidence of itself in its members, emissaries, symbols, tangible artifacts, and media, yet it is always something more.

All these qualities describe a World from the perspective of living inside one, already deep within its ongoing history. But this view takes for granted how a World begins, how it comes to be, and how it could ever originate from an individual author. As an artist, this is at the heart of my desire to understand what a World is. Because the dream is to be able to possess the agency to create Worlds – now more than ever – not just inherit and live within existing ones.

To think about beginnings, we have to go back to the moment before a World is born, to the moment of a curious creator looking at Reality — chaotic, meaningless, scary, but latent with potential — and wondering what to do with it. Philip K Dick said, "Reality is that which, when you stop believing in it, doesn't go away." A World is conceived when a creator decides to pick some part of Reality and start believing in it again. Think of JK Rowling stuck on a broken train for four hours, her mom dying, and the life arc of Harry Potter coming to her fully formed. Think of Elon, after the Russians rejected his bid to buy a used rocket, firing up his spreadsheet and costing out how to build one from scratch. Think of George Lucas, unable to secure the rights to Flash Gordon serials, thinking I'll just roll my own episodic space opera. Think of your mom looking at your dad for the first time and imagining, I could build something with that. The belief is fragile at first, but immediately suggests both a stabilizing structure and interesting generative dysfunction. The creator sets about trying to structure this belief and channel its potential. And at the same time, the creator begins to imagine another satisfaction: putting aside the role of creator and being a person living inside the belief, the beneficiary of its potentiality, a believer.

So what is a World from the perspective of a creator? Here is a working definition:

Gervais Principle

Tempo

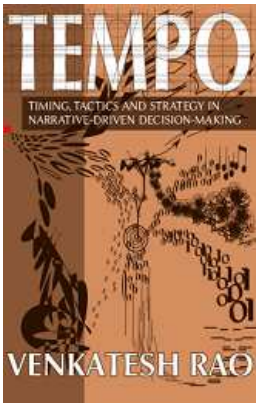

Every Cradle is a Grave

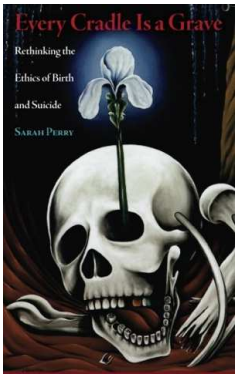

> A World is a reality you can believe in: one that promises to bring about habitable structure from the potential of chaos, and aim toward a future transformative enough to metabolize the pain and pleasure of its dysfunction.

Such a notion of a World — with its built-in dynamic between structure and dysfunction — is the sufficiently complex and enduring response to the chaos of bald Reality. One that a creator of a World must account for and encompass within its creation. It invites both sides of your brain to the party: the open/spiritual part of yourself who seeks change and transformation, alongside the conservative/dogmatic part of yourself who seeks stabilizing structure and evolutionary enduring procedures. It ensures a World that is perpetually interesting in its dysfunction in ways that a perfect utopia can never be; nurturing and habitable in ways that a no-rules dystopia can never be; and spiritually transformative and meaningful in ways that routine everyday life lived within aging institutions can no longer be. In short, it's a World that can keep going because it is a World worth living in and contributing to.

Let's wrap here with a pragmatic application of this notion of a World. There are three minimum archetypal forces that need expression in a World to keep a World alive, interesting, and structured enough to want to keep Worlding it. They are:

1. Chaos: an encroaching, unknowable, nature force.
2. Order: a dogmatic, conservative, structuring force.
3. Transformer: A spiritual, open, metabolizing force.

For example, in the World of *Game Of Thrones*, George RR Martin establishes:

1. Chaos: expressed via the white walkers and the dragons.
2. Order: expressed via the seven ruling kingdoms and its various representative characters vying for control of the age-old hierarchical ruling throne structure.
3. Transformer: expressed via individuals like John Snow and Daenerys Targaryen who transcend their given roles and attempt to metabolize the encroaching chaos and the classic kingdom structure into something more. In doing so, they irreversibly update Westeros.

With these forces at play, Martin's stories have a sufficiently dramatic domain to exercise in, there's enough dysfunction to surprise him as he writes, and there is enough complexity and artistic aim to grow beyond Fire and Ice in the minds of its members (showrunners, actors, fans, etc). The proof of its success as a World: HBO doesn't need Martin, the creator, to write the rest of the show. This is a World that can perpetuate itself.

So what is Worlding?

> Worlding is the art of devising a World: by choosing its dysfunctional present, maintaining its habitable past, aiming at its transformative future, and ultimately, letting it outlive your authorial control.

Series Navigation
<< Worlding Raga: 1 Worlding Raga: 3 — Slouching with God >>

2 Comments

About Ian Cheng

Ian Cheng is an artist based in New York. Follow his work on Instagram and at iancheng.com

Comments

says
March 5, 2019 at 2:25 pm

Wow, LOVE what u are saying. I totally "get it". Thank u!!!

Reply

http://Liz says
March 6, 2019 at 8:35 pm

Love it! Reminds me of franchise structures. Colin Burnett is working on a book about the James Bond franchise, an example of an industry structure particularly conducive to a most inclusive and expansive Worlding.
https://artsci.wustl.edu/ampersand/serial-bonds-many-lives-007

Reply

Leave a Comment

Name *

Email *

Website

POST COMMENT

This site uses Akismet to reduce spam. Learn how your comment data is processed.

Meta

- Log in
- Entries RSS
- Comments RSS
- WordPress.org

Useful Links

- Blogging Residencies
- For New Readers
- Now Reading
- You Are Here
- Consulting
- Refactor Camp

Recent Comments

- Peli Grietzer on Domestic Cozy: 2
- Senthil on Domestic Cozy: 2
- Martin on Domestic Cozy: 2
- Paul on Domestic Cozy: 2
- Nathaniel Peter Eliot on Worlding Raga: 3 — Slouching with God
- Machado on Worlding Raga: 3 — Slouching with God
- Romeo Stevens on Domestic Cozy: 2

Return to top of page

The State of Ethereum 2.0

Executive Summary

Research and development of Ethereum 2.0 continues at a rapid pace, with a testnet of the beacon chain scheduled for release in March of this year. However, a number of coordination problems are slowing implementation and the public's view of what will be delivered and when differs from reality. In this paper, we outline these problems and propose solutions for how Kyokan and the Moloch DAO can help fix them.

Overview

"Ethereum 2.0" refers to a set of specifications that will dramatically improve the performance characteristics of the Ethereum blockchain. As of this writing, it does so by merging and improving upon research from two older specifications: "Casper," which introduces a proof-of-stake consensus mechanism, and "sharding," which introduces the splitting of

transactions across a number of “shards” secured by the main chain. These specs confer the following benefits to Ethereum users:

1. Proof-of-stake removes the need to invest in equipment and burn electricity to secure the chain. Furthermore, it improves Ethereum’s finality characteristics by making certain types of 51% attacks dramatically more expensive and reducing reliance upon mining cartels to secure the chain.
2. Sharding improves the Ethereum network’s maximum transactions per second by orders of magnitude.

To determine the current state of the project, we interviewed the researchers and implementers working on the protocol.

Interview Methodology

We interviewed the following implementation teams via video call:

- Nimbus (Status)
- Lodestar (ChainSafe Systems)
- Artemis (PegaSys)
- Lighthouse (Sigma Prime)
- Prysm (Prysmatic)

Each implementation team was asked questions encompassing the following functional areas:

1. Team status;
2. Development status;
3. Roadmap;
4. Launch considerations;
5. Dependencies;
6. Comparisons to other implementation teams; and
7. Recommendations.

We also interviewed Danny Ryan, one of the core Ethereum Foundation researchers working on the project, via phone.

Observations and Potential Consequences

We now present our findings from the interviews described above. Quotes from individual interviewees are placed in quotation marks, and are reproduced verbatim.

Implementation Teams are Committed, But Funding Is A Concern

We asked each team their likelihood of and under what conditions they would give up development. All of the implementation teams we spoke to were committed to seeing Ethereum 2.0 through to completion as long as funding exists to continue development. This is an important point to underscore, as it implies two things. First, the implementation teams care deeply about shipping Ethereum 2.0, and are willing to weather the frustrations and problems that crop up along the way. Specifically, we received answers such as “We’d be dead before giving up.” and “This is going to happen no matter what.” in response to our questions around what it’d take for them to give up. However, the implementation teams are not immune to market realities. If EF funding dries up, or the larger entities funding individual implementation efforts (e.g., ConsenSys or Status) turn off funding, then there is a possibility that teams will be forced to find other work.

The Spec Continues To Churn, But It’s Getting Better

The Ethereum 2.0 spec has experienced a high level of churn over the past year. According to one person we interviewed, the spec has “entirely changed since the middle of last year” and continues to undergo regular “surgery” as issues are found and ironed out by the research team. Every aspect of the spec is subject to change. For example, until recently the names of important data structures were changing “emergently” - that is, a researcher would update their mental model of what a particular data structure represented, and change the spec accordingly without consideration of the effects that such a change would have on implementation efforts. For example Issue #358, 35 individual fields were renamed, yet the GitHub discussion received no participation from implementers. This has forced implementation teams to redo large swathes of work as the spec shifts under them - leading to frustration, wasted time, and in some cases a reduction of resources allocated to Ethereum 2.0 projects until the spec stabilizes.

There have been several promising developments over the past few weeks to reduce churn. First, according to the research team there is an ongoing effort to start versioning specific areas of the spec in order to make clear which areas are stable enough for implementation and which are still being actively researched (this first release was published the day this report was published for internal review). Second, the research team believes that “changes are slowing down” and that “deep reorgs” of the spec itself should be rarer now. A culture shift is occurring as well: the impact a particular change will have on implementation teams is now considered as part of new spec proposals (see this issue as an example). As a result of these developments, implementation teams unanimously agree that the spec in its current state is implementable.

Implementation Teams Do Not Push Back Against Researchers

Most implementation teams do not push back against the research team. They stated two reasons for this: implementers either feel unqualified to push back, or they feel like the chances

of successfully pushing back are too low to warrant doing so. These feelings are reinforced by how the research team describes the changes they have inserted into the spec: changes are usually described “clearly better” and “hard to push back against” given the qualifications of the people proposing the changes. While it is true that some areas of the spec can only be critiqued by a select few individuals, that feeling of “research exclusivity” currently extends to the entire spec as well as overall Ethereum 2.0 plan of execution.

The Definition Of Done Varies From Team To Team

All of the implementers we spoke to are working towards the testnet launch in March. What that launch looks like - and what happens afterwards - varies significantly from team to team. For example, it is unlikely that day one of the testnet will support inter-node peering because the peer protocol specification has not been fully accepted yet. Some teams have inter-node operability as a specific goal of the testnet launch, others do not. As a result, it is difficult to say with clarity what the actual deliverable in March will consist of.

Things become increasingly hazy past the testnet. None of the teams were able to estimate when Phase 2 - that is, the complete Ethereum 2.0 specification including cross-shard operations and EVM - would be mainnet-ready. Since some teams received grant money for the beacon chain specifically, it is likely that implementation teams will need additional funding to complete the spec.

Finally, only one of the teams we interviewed had user adoption - specifically, “100 validators staking using [their] software” - as one of their stated goals. The others were focused on completion of their respectively committed parts of the spec.

What Comes After Ethereum 2.0 Is Unclear For Implementers

Many teams expressed concerns around what comes next for their respective businesses after successfully delivering Ethereum 2.0. None of the teams we spoke to had a clear answer for how they would monetize post-implementation, so securing funding for continued maintenance (particularly if ETH continues to drop in price) is of major concern to implementation teams.

There Is No Ethereum 2.0 Lead

From an organizational perspective, no single person is responsible or accountable for making sure that Ethereum 2.0 lands and lands in a way that matters to the Ethereum community at large. Danny Ryan fills part of this role. He took it upon himself to be the liaison between the implementation and research teams, and his efforts are highly appreciated. Access to Danny isn’t consistent across implementation teams, however - some expressed that they wished that they could have more access to him.

Ethereum 2.0's Narrative Is Controlled By People Outside The R&D Process

Consider James Prestwich's popular post "What To Expect When ETH's Expecting." It includes claims like the following:

- "The tools and contracts we've written for ETH1.X will likely need to be completely redesigned and rewritten for ETH2.0."
- "Phase 1 doesn't have anything particularly interesting in it. Fundamentally it's a bootstrapping phase for crosslinking, and the symmetric mechanism by which shards reference the beacon chain. The designers seem confident that these mechanisms will work."
- "Interestingly, Phase 0 implementation has happened concurrently with specification. Even today, less than three months from testnet, the Phase 0 specification changes regularly. This implies that future ETH2.0 phases will have extremely high variance in development time. While optimists have told me six months, it is easy to see Phase 1 taking 12–18 months of development after Phase 0 enters testing."
- "[Beyond eWASM, EVM2, and storage rent] we don't know what to expect from Phase 2. It's still in very early stages of research and includes several major unsolved problems. Given the informal specification and development process, as well as Phase 2's expanded scope over Phase 1, it doesn't seem reasonable to suggest that Phase 2 could launch before 2020. Which is to say, while ETH2.0 may launch this year, don't expect ETH2.0 to support asset transfer or smart contracts until at least 2020."
- "We have very little information about ETH2.0's communication model. We know that it can't provide cross-shard contract calls without sacrificing almost all scaling benefits. I won't blame you if you stop reading here, as Phase 4 only has a mind map and a few vague links. A non-obvious consequence of this is that ETH2.0 will not provide significant scaling benefits to complex smart contract systems until Phase 4. Until then, contracts wishing to interact with other contracts must cohabitate a shard and are limited to the speed and scale of that shard."

These are concrete details about how developers can expect to make use of Ethereum 2.0. The post also includes plenty of technical detail about the spec, but as an Ethereum developer the most relevant information contained therein is about how the spec will impact my work and when to expect scalability improvements. Media created by the implementation teams and the EF, in contrast contrast, tends to focus instead on new research or the completion of specific pieces of the spec. Consider the following outline of Prysmatic's Development Update #20[1]:

[1] This isn't limited to Prysmatic specifically - the majority of public-facing media from implementers and researchers alike focuses on the tech. A Prysmatic development update was chosen because Prysmatic has the highest volume of public media.

- Latest Research
    - Phase 0 Validator Client Responsibilities
- Merged Code, Pull Requests, and Issues
    - State Transition Block Processing E2E Testing Complete
    - State Transition Epoch Processing Integration Complete
    - Implementing Deposit Listener For The Validator Deposit Contract
    - Implementer Validator Deposits Trie
- Upcoming Work
    - GHOST Fork-Choice Rule for the Ethereum Beacon Chain
    - Full End-to-End Testing of Beacon Chain With Validator Deposits
    - Deprecating Our Solidity Contract to Vyper
    - Creating the Beacon Chain's Transactions Pool
    - Refactor Validator Client
    - Validator Private Key Management and Other Secrets
- Misc
    - Ethereum 2.0 Implementers Call Jan 17, 2019
    - Blog Post on Ethereum 2.0 (which includes a link to Prestwich's post above)

The things that drive Ethereum 2.0's "narrative" - i.e., when it'll ship, what it'll be useful for, and how developers can use it - is driven more by posts like Prestwich's than Prysmatic's since the information is more directly relevant to the day-to-day work of Ethereum's users. We commend the Ethereum 2.0 teams for their commitment to transparency, and obviously want technically-focused updates such as the one above to continue. However, if nobody from the research or implementation teams provides additional context around when Ethereum 2.0 will be ready and what it will look like when it is, others will continue to step in and do it for them. As a result, it will be difficult for correct expectations to be set and lived up to.

Addendum: A conversation of Ethereum community members pointed us towards this high level outline of Ethereum's roadmap prepared by EthHub. The link to the official sharding roadmap, while instructive, isn't as helpful for setting the expectations of platformer developers who want to know what each phase of the spec means for them.

Open Questions

What Does The Community Expect?

Currently, the narrative around when Ethereum 2.0 will be delivered and what it will look like is somewhat like this:

- Ethereum 2.0 is coming, and will be ready for public use soon.
- Ethereum 2.0 will be on testnet starting in March.
- Ethereum 2.0 will solve the majority of Ethereum's scalability concerns.

Given what we know after talking to both the research team and the implementation teams, it's clear that delivery of an Ethereum 2.0 implementation with real utility to dApp developers is not feasible for at least another year and a half. From our understanding, the deliverables included in each phase of the Ethereum 2.0 roadmap are as follows[2]:

- **Phase 0:** The beacon chain.
- **Phase 1:** Shards without the EVM.
- **Phase 2:** EVM on shards and cross-shard communication.

For developers to derive the same level of utility from Ethereum 2.0 as they did on Ethereum 1.0, phase 2 must be delivered.

Furthermore, for the later phases of Ethereum 2.0's rollout it is possible that new research will invalidate or otherwise reshape the roadmap profoundly. It is unclear to us whether or not the Ethereum community at large is aware of this. A gap between what the community expects and what is actually delivered could hurt implementation efforts substantially by reinforcing the narrative that Ethereum will not be able to scale and cause new developers to begin exploring other blockchains.

Were Implementers Consulted When Designing The Rollout?

It is unclear to us how much, if any, implementer input was included in the decision to roll out Ethereum 2.0 in phases and what goes into each phase. While we understand the value of a phased deployment process - i.e., that it gives new technology like proof-of-stake time to "burn in" in a quasi-production environment - it stands to reason that the individuals responsible for implementing each phase are some of the most qualified people to design them in the first place. This includes which technology is introduced when, as well as when each phase is set to be delivered. If implementers were not consulted, would the launch of the beacon chain be a good time to reset and bring implementers into the process?

Would Making Danny Ryan Official "Ethereum 2.0 Lead" Help?

Many of the frustrations around Ethereum 2.0 stem from a lack of coordination between the research and implementation teams. The protocol consists of numerous disparate components that must be integrated as part of a plan spanning several years. Until Danny Ryan assumed the role of coordinator, there was no single person to whom both researchers and implementers alike could turn to to oversee that integration. Danny has already demonstrated his value as a lead. His name came up repeatedly during our interviews as someone that implementers would

[2] Note that the phases described on the Sharding Roadmap document are different. The development phases described to us during our interviews indicated that those phases are no longer canonical.

like to see more of, and his efforts on the early versions of the specification show that he is knowledgeable enough as a researcher to oversee the project.

It's important to be clear about what 'lead' means in this context. We are using 'lead' to mean the single person who is:

- Accountable for making sure that Ethereum 2.0 lands.
- Accessible to all implementation teams and researchers.
- Empowered with veto authority to act as a tiebreaker for critical decisions.
- Empowered with delegation authority to make sure that the right people are solving the right problems.

This is most definitely a centralization of control. However, given the role he is already filling empowering Danny with official leadership seems fitting, and could make sure that the entire project is integrated smoothly.

Recommendations

Include "Product Context" In Public-Facing Media

Given the importance of Ethereum 2.0 to the success of the network, clearly communicating what will be delivered and when as well as how to prepare for the release of Ethereum 2.0 is paramount. To make media produced by Ethereum 2.0 teams and researchers more relevant to the community and regain control of the Ethereum 2.0 narrative, we recommend clearly articulating the following things with new public updates:

1. How the latest update will impact developers.
2. How roadmap changes might impact Ethereum 2.0's timeline or roadmap.
3. Areas where research or product churn is expected.

Including the above things will go a long way towards returning control of the Ethereum 2.0 narrative to those working on it.

Provide Clear Avenues For Continued Funding

We believe that incentivizing long-term, continuous development on Ethereum 2.0 clients is critical to a successful Ethereum 2.0 launch. The source of continued funding for implementation is ambiguous and a source of worry. If the Ethereum Foundation or an alliance of other interested parties pooled money and provided clear funding amounts with associated timelines, it would remove many worries around how client projects will continue to fund maintenance and new features once Ethereum 2.0 is launched.

Rigorously Define And Enforce A Formal Standards Process

The coordination problems inherent to implementing a spec while it is being defined are exacerbated by the lack of a formal standards process. By developing and publishing specifications, the Ethereum Foundation is acting as the de-facto standards body for Ethereum 2.0. As such, defining a formal process through which research can go from proposal to implementable can further reduce the amount of churn around the spec. There are numerous standards bodies from whom the EF can draw examples, however we recommend a variant of ECMA'S TC39 Standards Process. Our reasoning is as follows:

1. The TC39 process is open, and embraces modern development practices such as pull requests on GitHub that contributors are already familiar with.
2. The TC39 process bakes acceptance tests and reference implementations into the process itself.
3. The TC39 process is more digestible and has less overhead than other standards processes.
4. The TC39 process favors incremental releases of new standards on an set cadence.
5. The TC39 process has a track record of success. As a result of the TC39 process, the JavaScript ecosystem successfully recovered from 10 years of language stagnation.
6. Many of Ethereum's developers come from a JavaScript background, and as such are already familiar with the TC39 process (specifically, proposal "stages") via Babel.

We recommend introducing, at the very least, the concept of "stages" from the TC39 process. For the unfamiliar, TC39 proposals go from stage 0 ("strawman") to stage 4 ("complete"), after which they are ratified as new standards at an annual meeting of TC39 members. The objective of having proposal stages is to make it extremely clear how ready an individual proposal is for implementation. Furthermore, since test vectors and reference implementations are requirements to transition from one phase to another, dialogue between researchers and implementers is encouraged. While an implementer may not be qualified to comment on the specifics of an individual algorithm, for example, they *are* qualified to comment on how that algorithm gets implemented. Under TC39, both the research and the implementation would be required to transition from stage 3 to 4.

How Kyokan/Moloch Can Help

Kyokan is a blockchain-native software consultancy based in the bay area. In the past, we have worked with MetaMask, SpankChain, Cosmos, Dfinity, and Uniswap. In addition, we received a grant from the Ethereum Foundation to build an implementation of Plasma MVP, which is currently preparing for mainnet launch. Our team has significant experience shipping production software at prominent consumer and enterprise tech companies.

Moloch is a grant-making DAO / Guild and a radical experiment in voluntary incentive alignment to overcome the "tragedy of the commons". Our objective is to accelerate the development of public Ethereum infrastructure that many teams need but don't want to pay for on their own. By pooling our ETH and ERC20 tokens, ETH investors and teams building on Ethereum can collectively fund open-source work we decide is in our common interest.

Kyokan, with funding from Moloch and others, is positioned to provide support for the ETH 2.0 effort in the following ways:

- Prepare reports and analysis, like this one, to inform the community of ETH 2.0 progress
- Evaluate the internal processes of the ETH 2.0 R&D effort
- Help develop organizational structure as needed
- Assist in coordinating standards across teams
- Help plan the development roadmap
- Provide launch coordination and prepare clients for production release
- Help with developer recruitment

Moloch is positioned to provide support for the ETH 2.0 effort as follows:

- Provide additional funding for ETH 2.0 teams
- Funding key hires to provide cross-team support
- Funding open-source tools (e.g. testing) to help 2.0 development

Conclusion

The ETH 2.0 effort may be unique, but the challenges it faces are generally what one might expect to observe in an emergently organized R&D effort of 8+ researchers and 50+ developers. With stronger coordination through leadership, a standards process, a well-communicated shared roadmap, and funding security, the ETH 2.0 effort is set to accelerate and meet the expectations of the Ethereum community.

We're already seeing the beginnings of this acceleration taking place. In December, Vitalik (non-giver of ETH) said "YOLO" and gave a 1000 ETH grants to each of the Prysmatic, Lighthouse, and Lodestar teams. Another prominent ETH investor followed Vitalik and contributed 2,800 ETH to Prysmatic, helping Preston Van Loon leave Google to join the ETH 2.0 effort full-time. As more community members step up to share responsibility for ETH 2.0 delivery, we anticipate even greater results. We're all in this together.

Disclaimer

Funding for Kyokan's efforts in preparing this report is anticipated to come from the Moloch DAO, but has also been guaranteed personally by Ameen Soleimani, should there be complications. The report was authored primarily by Matt Slipper and Dan Tsui of Kyokan.

Patterns for Decentralised Organising

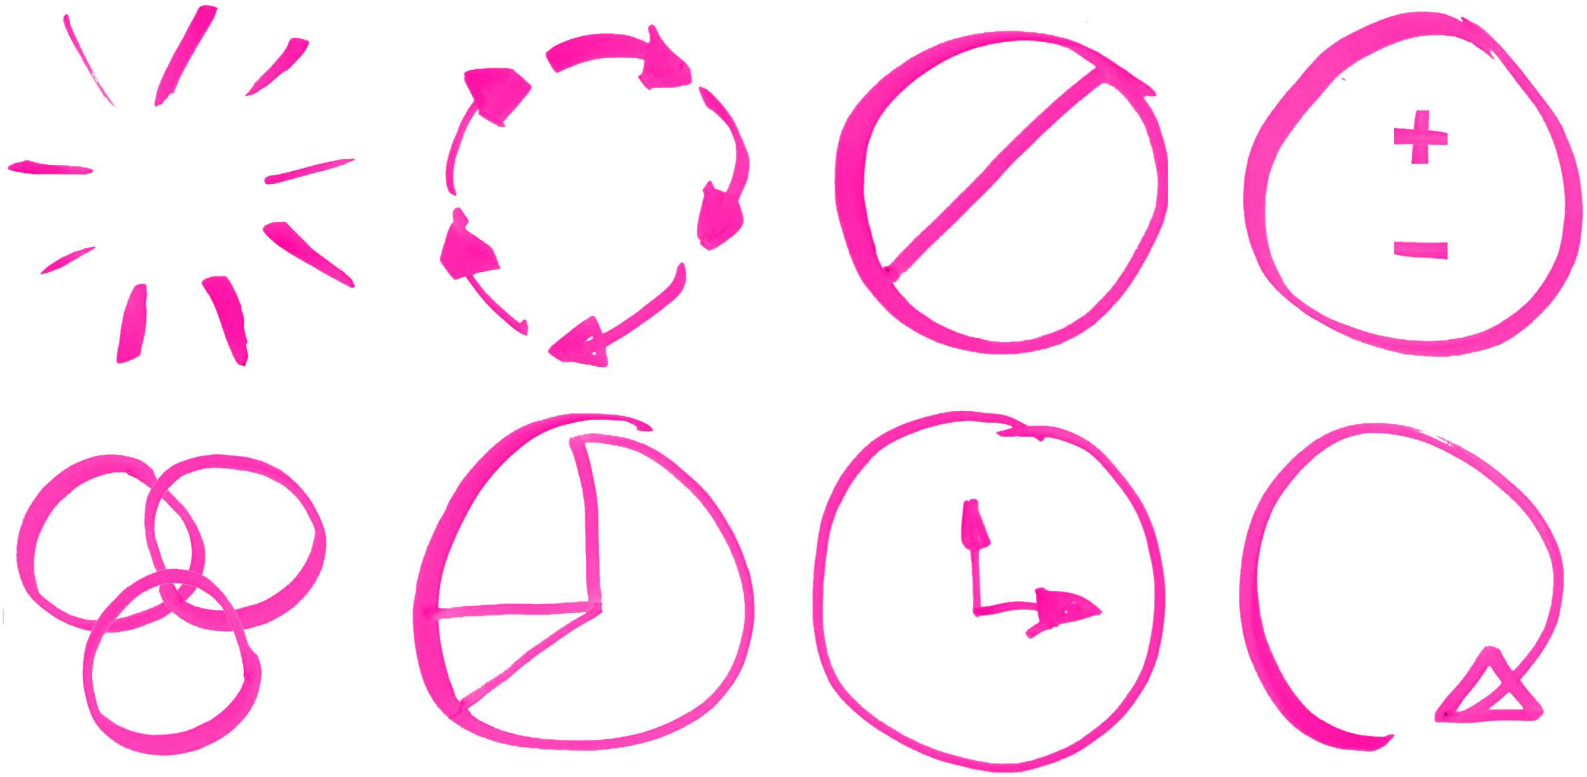

Practical guidance for humming teams

from TheHum.org

A note from the author

As a cofounder and coordinator at Loomio and Enspiral, I worked with dozens of teams experimenting with decentralised organising. These folks are prototyping organisational structures for distributed leadership, high autonomy, and shared ownership. I noticed we all faced similar, painful challenges:

> *How can we be inclusive without spending all our time in meetings? How do we deal with power imbalances? How do we prioritise what to work on? How do we undo our programming and develop an open, collaborative culture? And where does accountability come from if there is no boss!?*

I spent 2017 travelling the world with my partner, working with hundreds of groups to understand how to translate my experiences from Loomio and Enspiral into lessons that can help any team that wants to work with more collaboration and less hierarchy.

The result is these "**patterns**", naming the most common *challenges* of working without a management hierarchy, and practical *responses* you can adapt to your local context and apply immediately.

— *Richard D. Bartlett*
TheHum.org

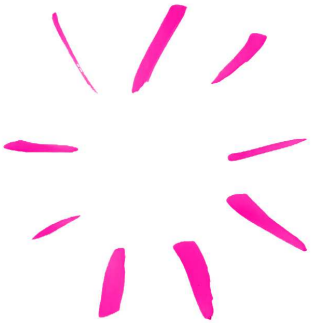

1. Intentionally produce counter-culture

Challenge:
You want to be non-hierarchical but you have **hierarchical habits**, e.g. telling people what to do, or looking to others for answers. We are **conditioned by culture**: if there is sexism and racism in your environment, it can be imprinted into your habits.

Response:
We can **unlearn hierarchies** together. We can **co-design a culture** that encourages each of us to develop our best qualities, making us all more generous, respectful, trusting, courageous, etc.

How do you produce culture? **Fermentation!** To make sourdough bread, you have a **starter dough**, mixed with **fresh ingredients,** and put it somewhere dark and safe for some time. To ferment a new group culture, your "starter dough" is a person or people who **embody some of the qualities** you want to develop. The “fresh ingredients” are **new people** who have a **desire** to grow in a specific way. We combine these ingredients in a **retreat**: safe, quiet, isolated from the outside world for a few days.

Results:
We learn about each other’s **dreams** and **fears**, building deep **relationships of trust** and **belonging**: the most important resource for all your upcoming challenges.

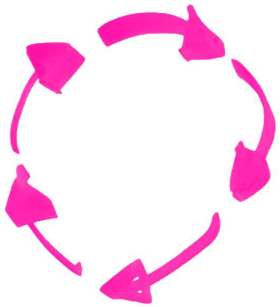

2. Systematically distribute care labour

Care includes the practical stuff of **hospitality**: preparing a comfortable room with food, lighting, decoration, refreshments, collaboration tools, and tidying up after. It also includes **emotional work**, like noticing tension between colleagues and supporting them to resolve it.

Challenge:
Hierarchical culture trains us to not **share the care labour fairly**. Most groups have one or two people, usually women, doing most of the care work. If they get overwhelmed or frustrated, they'll stop, and the group loses its gravity.

Response:
Make all work visible, so you can share it fairly. E.g. the Loomio team uses "**stewardship**", a peer-to-peer support system. Everyone supports one person, and is supported by someone else. Each pair meets once per month, the steward asks "how can I support you?" and they figure out the answer together. More info: loomio.coop/stewarding.html

Results:
Builds deep **trusting relationships**; dissolves **conflicts**; continuously **improving emotional intelligence** of everyone in the group; more distribution = **more resilience**.

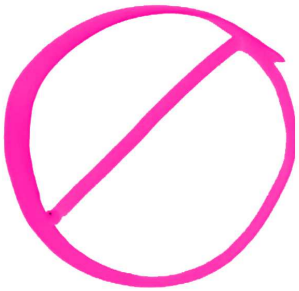

3. Make explicit norms and boundaries

Norms = **how we do things** around here. Boundaries = **what we don't do** around here. Many groups leave these things unsaid, relying on "common sense".

Challenge:
Conflicts grow when people have different unspoken assumptions (everyone has different common sense). When you cross an invisible boundary, it takes huge energy to make the boundary explicit, before you can get to the behaviour.

Response:
Talk about your norms: **how do we want to be together?** e.g. open, honest, authentic, gentle, inquisitive...

Talk about your boundaries: **what behaviour do we want to exclude?** e.g. no mean feedback, no sexist jokes.

Results:
Buy-in — clarity helps people evaluate whether or not they want to be here. Expectations are clear. There is a process for challenging destructive behaviour, and a process for updating our agreements.

E.g. see roles + responsibilities described in Enspiral's People Agreement: handbook.enspiral.com/agreements/people.html

4. Keep talking about power

Challenge:
Power, influence, status, rank, social capital, volume, access... whatever you call it, I've never met a group where it was equally distributed between all members. Equality is a compass point to navigate towards, not a destination I've ever arrived at.

Response:
Groups thrive when anyone can safely talk about power differentials. Imbalance can be bad, e.g. inherited privilege, coercion, manipulation, the "old boys club". Some imbalance can be good: earned trust, reputation, eldership.

Transparency reduces toxicity. Discuss together: "**How's the power?** Who has it? How do you earn it?"

Some roles attract power (e.g. manager, facilitator, spokesperson, coordinator, director). Rotation increases access: take turns, step out, encourage others to step in.

E.g. Loomio team coordinators are elected by the team for a limited term; we intentionally support less experienced people to try the role. See loomio.coop/coordination.html

The best 'elders' use their status to praise, acknowledge, and encourage people with less.

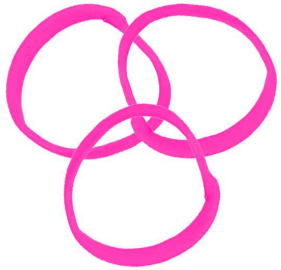

5. Agree how you're using your tech

Challenge
Many groups are dissatisfied with their communication technology. Information overload: too much data but can never find the thing you want. Half the team uses this tool, the other half uses another one. Too many tools, don't know how to get everyone's attention, can never find the document I need.

Response
Agree together what tools are for what job. E.g. the 'trinity of digital comms':

1. **Realtime**, like chat, messenger, or Slack. Informal, quick, organised around **time**: it's about right now.
2. **Asynchronous**, like email, forum or Loomio. More formal, organised around **topic**. Has a subject + context + invitation. Can take days or weeks. Makes a useful archive, considered comments rather than random messy chatter.
3. **Static**, like a wiki, Google Docs, handbook, or FAQ. Very formal, usually with an explicit process for updating content.

Depending on your work, you will need different tools. The important thing is that you have an agreement together about what tools are for what job. With a shared understanding of the tools, they all fit together beautifully. When people have different ideas, it gets messy.

A note about introducing new tools

Introducing a new communications tool usually makes the problem worse. Most groups don't know how to introduce new tech well.

This method makes it less likely to go badly:

1. **Agree the problem.** What issue do you want to solve? Do other team members agree it's a problem?
2. **Volunteer(s) test prototypes.** One or a few people research options and come back to the wider group with a recommendation.
3. **Support people to learn.** Once you've chosen a new tool to evaluate, make space for people to learn together how to use it.
4. **Reminders build a habit.** It can take weeks to develop a new communication habit. Remind each other gently, "hey we said we'd try using Loomio for these kinds of conversations..."
5. **Evaluate + repeat.** Most importantly, set a time-limit, e.g. "We'll try this tool for 2 months and then evaluate together. Is the problem solved? Or do we need more training, or a different tool?"

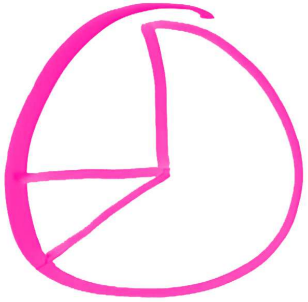

6. Make decisions asynchronously

Challenge:
Most collaborative groups make decisions in meetings or conference calls. Meetigs are a kind of **synchronised** or **realtime** communication: you have to synchronise people's calendars to find a time that works, then when they arrive, everyone has to pay attention to everything at the same time. It's very **expensive**, **excluding** the input of people who can't attend, and often results in **hurried** decisions.

Response:
With a little effort, you can develop a habit of **asynchronous decision-making**. People can participate in their own time, contributing only to the issues relevant to them. This is what **Loomio** is for: more inclusion and collective intelligence, less time in meetings.

E.g. I'm on a Board of Directors. We meet monthly. We co-create the agenda in a Loomio thread ahead of time. A few days before the meeting, a Loomio poll confirms everyone is happy with the agenda and we've all read the reports. We all arrive at the meeting prepared and focussed. We'll make some decisions face-to-face. For decisions that require input from more people, or more time to consider options, one of the Directors will take the decision to Loomio. We also use the software to sign off the minutes, and find another meeting time.

Results:
Over time you learn the unique qualities of realtime and async communication. Meetings are good for bonding, brainstorming, and dealing with complex or sensitive topics. Loomio creates more space for deliberation: you can take more time, consider more options, hear from more people, and keep a record.

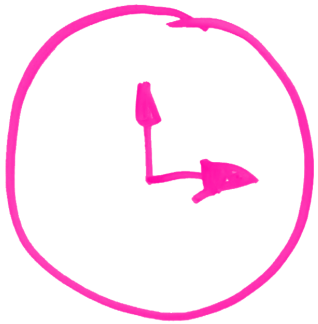

7. Use rhythm to balance flexibility and focus

Challenge
Hierarchies are designed to manage flows of communication and decision-making. When you remove the hierarchy, you need to replace it with something. If there is no agreed structure, your group can suffer from information overload (everyone asked about everything all the time) and exclusion (decisions made without appropriate input).

Response
Rhythm helps balance **speed** with **participation.** People can trust each other to seek input at the right time, so they don't need to be involved in every decision.

We create distinct communication spaces for different timeframes, e.g. today's work is discussed **every morning**; if you want to discuss the long term strategic direction, we have a dedicated space for that **every month**.

Example
Here's a set of working rhythms we use in the Loomio team. You can adapt to your context, e.g. maybe it makes sense to align with seasons or moon cycles:

- **Daily** "standup" meeting. Everyone answers, "What did you do yesterday? What are you doing today? Are there any obstacles we can help you with?" Quick info exchange, accountability and support.
- **Weekly** "sprints", a regular working period. E.g. on Monday we agree what work we're going to do this week. On Friday we share progress and have a "retrospective" looking for

improvements to try next week.

- **Quarterly** objectives. Every 3 months we have a planning day, looking for agreement on 3 or 4 measurable targets to align all of the work in the cooperative. After we finalise the decision on Loomio, everyone has freedom to do whatever work they feel is most relevant to achieve those outcomes.
- **Bi-annual** retreats. Every 6 months we go away together for 3 or 4 days. This deepens our relationships, and creates a space for the kind of conversations that can't happen in the office, e.g. dreaming together about our shared vision, or dealing with a complex tension.

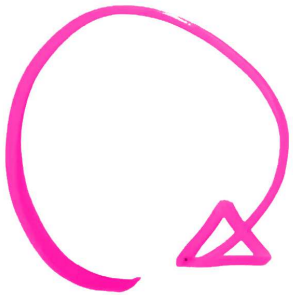

8. Generate new patterns together

Challenge
There is no such thing as an organisational structure that suits every team. Processes that worked for you last year are made obsolete by changing environmental conditions and team makeup. You need a reliable way to notice what's not working and make improvements without losing people along the way.

Response
Retrospectives turn frustrations into improvements.

You can choose a frequency that suits you, but let's say weekly. At the end of each week, stop working. Have a **retrospective** meeting. Review the week just been. **What was good?** Notice it and do more. **What was bad?** Discuss. **Agree a change** that you're going to try next week to make it more good and less bad.

I've shared a bunch of collaboration patterns. These are just my way of describing things I've seen. We discovered them by a lot of invention, remixing ideas, making them our own, adjusting them to local conditions. This booklet is not a recipe for you to copy and paste, it’s a guide to a way of thinking. The retrospective is where you learn to notice your own collaboration patterns, and co-design new ones.

See the Retrospective Wiki for ideas of how to host a structured reflection process: retrospectivewiki.org

Going deeper

This booklet is an extremely condensed summary of what we've learned about decentralised organising. If you're looking for more stories, conversations, training or consulting, check TheHum.org.

Credits

Contributor	Contribution
Richard D. Bartlett rich@thehum.org	Writing first draft in English
Natalia Lombardo nati@thehum.org	Co-designed first draft with Rich.
Lucas Tauil de Freitas lucas@enspiral.com	Translation for Portuguese (Brasil)
Sandra Chemin sandra.chemin@enspiral.com	Revision & editing in English + Portuguese

License: CC0 1.0 Universal

Trust

First published Mon Feb 20, 2006; substantive revision Mon Aug 3, 2015

Trust is important, but it is also dangerous. It is important because it allows us to form relationships with people and to depend on them—for love, for advice, for help with our plumbing, or what have you—especially when we know that no outside force compels them to give us such things. But trust also involves the risk that people we trust will not pull through for us; for, if there were some guarantee that they *would* pull through, then we would have no need to trust them.[1] Thus, trust is also dangerous. What we risk while trusting is the loss of the things that we entrust to others, including our self-respect, perhaps, which can be shattered by the betrayal of our trust.

Because trust is risky, the question of when it is warranted is of particular importance. In this context, "warranted" means justified or well-grounded (where well-grounded trust successfully targets a trustworthy person). If trust is warranted in these senses, then the danger of it is either minimized, as with justified trust, or eliminated altogether, as with well-grounded trust. Leaving the danger of trust aside, one could also ask whether trust is warranted in the sense of being plausible. Trust may not be warranted in a particular situation because it is simply not plausible: the conditions necessary for it do not exist, as is the case when people feel only pessimism toward one another. This entry on trust is framed as a response to the general question of when or whether trust is warranted, where "warranted" is broadly construed to include "justified," "well-grounded" and "plausible."

A complete philosophical answer to this question must explore the various philosophical dimensions of trust, including the conceptual nature of trust and trustworthiness, the epistemology of trust, the value of trust, and the kind of mental attitude trust is. To illustrate how each of these concerns is relevant, note that trust is warranted, that is,

i. Plausible, again, only if the conditions required for trust exist (e.g. optimism about one another's ability). Knowing what these conditions are requires understanding the nature of trust.

ii. Well-grounded, only if the trustee is trustworthy, which makes the nature of trustworthiness important in determining when trust is warranted.

iii. Justified, sometimes when the trustee is not in fact trustworthy, which suggests that the epistemology of trust is relevant.

iv. Justified, often because some value will emerge from the trust or because it is valuable in and of itself. Thus, the value of trust is important.

v. Plausible, only when it is possible for one to develop trust, given one's circumstances and the sort of mental attitude trust is. For example, trust may not be the sort of attitude that one can will oneself to have without any evidence of a person's trustworthiness.

This piece explores these different philosophical issues about trust. It also deals predominantly with interpersonal trust, which I take to be the dominant paradigm of trust. Although some philosophers write about trust that is not interpersonal, including "institutional trust" (i.e. trust in institutions; see e.g. Potter 2002, Govier 1997, Townley and Garfield 2013), trust in government (Hardin 2002) or in artificial intelligence (which may only be person-like; see e.g. Coeckelbergh 2012, Taddeo and Floridi 2011), and "self-trust" (Govier 1993, Lehrer 1997, Foley 2001, McLeod 2002, Goering 2009, Jones 2012b, Potter 2013), most would agree that these forms of "trust" are coherent only if they share important features of (i.e. can be modeled on) interpersonal trust. Hence, I assume that the dominant paradigm is interpersonal.

---

1. The Nature of Trust and Trustworthiness

Trust is an attitude that we have towards people whom we hope will be trustworthy, where trustworthiness is a property, not an attitude. Trust and trustworthiness are therefore distinct although, ideally, those whom we trust will be trustworthy, and those who are trustworthy will be trusted. For trust to be warranted (i.e. plausible) in a relationship, the parties to that relationship must have attitudes toward one another that permit trust. Moreover, for trust to be warranted (i.e. well-grounded), both parties must be trustworthy.

Trusting requires that we can, 1) be vulnerable to others (vulnerable to betrayal in particular); 2) think well of others, at least in certain domains; and 3) be optimistic that they are, or at least will be, competent in certain respects. Each of these conditions for trust is relatively uncontroversial.[2] There is a further condition which is controversial, however: that the trustor is optimistic that the trustee will have a certain kind of motive for acting. Controversy surrounds this last criterion, because it is unclear what, if any,

sort of motive we expect from people we trust.

Likewise, it is unclear what, if any, sort of motive a *trustworthy* person must have. Clear conditions for trustworthiness are that the trustworthy person is competent and committed to do what s/he is trusted to do. But this person may also have to be committed in a certain way or for a certain reason (e.g. s/he cares about the trustor). This section explains these various conditions for trust and trustworthiness, and focuses in particular on the controversy that surrounds the condition about motive.

One important criterion for trust is that the trustor can accept some level of risk or vulnerability (Becker 1996). Minimally, what this person risks, or is vulnerable to, is the failure by the trustee to do what s/he depends on that person to do. The truster might try to reduce this risk by monitoring or imposing certain constraints on the behavior of the trustee; yet after a certain threshold perhaps, the more monitoring and constraining s/he does, the less s/he *trusts* that person. Trust is relevant "before one can monitor the actions of … others" (Dasgupta 1988, 51) or when out of respect for others one refuses to monitor them. One must be content with them having some discretionary power or freedom (Baier 1986; Dasgupta 1988). Hence, one cannot reject being vulnerable.

A related condition for trust is the potential for betrayal (and, as noted below, the corresponding condition for trustworthiness is the power to betray). Annette Baier writes that "trusting can be betrayed, or at least let down, and not just disappointed" (1986, 235). In her view, disappointment is the appropriate response when one merely *relied* on someone to do something, but did not trust him or her to do it. Although people who monitor and constrain other people's behavior and do not allow them to prove their own trustworthiness may rely on others, they do not trust them. For, while their reliance could be disappointed, it could not be betrayed. Consider that one can rely on inanimate objects, such as alarm clocks; but when they break, one is not betrayed, although one may be disappointed. Reliance without the possibility of betrayal is not trust. Thus, people who rely on one another in a way that makes betrayal impossible do not trust one another. (For some resistance to this view, see O'Neil 2012, 307.)

People also do not, or cannot, trust one another if they are easily suspicious of one another (Govier 1997, 6). If one assumes the worst about someone—"she is late because she has no regard for my feelings," or "I bet he is talking about me behind my back"—then one distrusts, rather than trusts the person. Paradigmatically, trust involves being optimistic, rather than pessimistic, that the trustee will do something for us (or for others perhaps), which is in part what makes us vulnerable by trusting. As Karen Jones writes, such optimism "restricts the inferences we will make about the likely actions of another. Trusting thus opens one up to harm, for it gives rise to selective interpretation, which means that one can be fooled, that the truth might lie, as it were, outside one's gaze" (Jones 1996, 12).

Some—including Jones in her later work on trust—argue that optimism about the trusted person is present in typical instances of trust, but not in all instances of it (Jones 2004, McGeer 2008, Walker 2006 cited in McGeer). Such optimism is absent, for example, in cases of "therapeutic trust" (Horsburgh 1960). To illustrate this type of trust, consider parents who "trust their teenagers with the house or the family car, believing that their offspring may well abuse their trust, but *hoping* by such trust to elicit, in the fullness of time, more responsible and responsive trustworthy behaviour"

(McGeer 241, her emphasis; see also Pettit 1995). The claim of Jones and others is that such trust involves the normative attitude that the trustee *ought* to do what one trusts him or her to do, rather than optimism that s/he *will* do it. Therapeutic trust is unusual in this respect and in others (which will become evident later on in this entry). The rest of this section deals with usual rather than unusual forms of trust and trustworthiness.

Failing to be optimistic about people's competence also makes trust impossible. Without being confident that people will display some competence, we cannot trust them. We usually trust people *to do* certain things—for example, to look after our children, to give us advice, or to be honest with us—but we would not do so if we thought they had none of the relevant skills (including moral skills of knowing what it means to be honest or caring; Jones 1996, 7; McLeod 2002, 19). Rarely, if ever, do we trust people completely (i.e. A simply trusts B). Instead, "trust is generally a three-part relation: A trusts B to do X" (Hardin 2002, 9).[3] To have trust in a relationship, therefore, we do not need to assume that the other person will be competent in *every* way. Optimism about the person's competence in at least one area, however, is essential.

When we trust people, we are optimistic not only that they are competent to do what we trust them to do, but also that they are committed to doing it. One can talk about this commitment either in terms of what the trustor expects of the trustee or in terms of what the trustee possesses: that is, as a condition for trust or as a condition for trustworthiness (and the same is true, of course, of the competence condition). For simplicity's sake and to focus some of this section on trustworthiness rather than on trust, I will refer to the commitment mostly as a condition for trustworthiness.

Although both the competence and motivational elements of trustworthiness are crucial, the exact nature of the latter is unclear. For some philosophers, it matters only *that* the trustee is committed. The central problem of trustworthiness in their view concerns the ongoing commitment of the trustee, and in particular, under what circumstances, if any, one could expect such a commitment from another person (see, e.g., Hardin 2002, 28). By contrast, for other philosophers, an ongoing commitment from another person is not sufficient for trustworthiness; according to them, the origins of the commitment matter, not just its existence or duration. The central problem of trustworthiness for these philosophers is not just *whether* but also *how* the trustee is motivated to act. Some motives are simply incompatible with trustworthiness, on this view. To determine which view is correct, we need to consider different possible motives that could underlie trustworthy behavior.

Some philosophers believe that trustworthiness can be "compelled by the force of norms" or, more generally, by the force of social constraints (Hardin 2002, 53; see also O'Neill 2002, Dasgupta 1988). In an effort to be trustworthy, people can subject themselves to social constraints, as someone does when she publicly declares her intention to lose weight, putting herself at risk of public censure if she fails. Alternatively, the *trustor* in a relationship can introduce the constraints by requiring that the trustee sign a contract, for example. The constraint imposed could be the primary motivation for being trustworthy. It would compel on ongoing commitment grounded in self-interest. Call this view of trustworthiness "the social contract view."

Many philosophers would agree that the social contract view is only a partial account

of what could motivate trustworthiness. While social constraints can shore up trustworthiness, they cannot account for trustworthiness altogether. For, if they could, then the following sort of person would be trustworthy: a sexist employer who treats female employees well only because he believes that he would face legal sanctions if he did not (Potter 2002, 5). Many would argue that while this person's behavior is predictable or reliable, it is not *trustworthy* in any genuine sense. These theorists may distinguish mere reliability from trustworthiness on the grounds that people known or considered to be trustworthy have the power to betray us, whereas people known or considered to be merely reliable can only disappoint us (Holton 1994). The female employees might know that their employer treats them well only because he fears social sanctioning. In that case, he could not betray them, although he could disappoint them. And if that were true, he would not be trustworthy for them.

An alternative to the social contract view is a view according to which trustworthy people are motivated by their own interest to maintain the relationship they have with the trustor, which in turn encourages them to encapsulate the interests of that person in their own interests. Russell Hardin defends this "encapsulated interest" view (2002). However, it too is problematic. To see why, consider how it applies to the sexist employer. He is *not* motivated by an interest to sustain his relationships with female employees: if he could easily fire them, or even avoid hiring them altogether, then he would do that. He is therefore not trustworthy. Imagine, however, that he did have an interest in maintaining these relationships and as a result he treated the women well, yet his interest stemmed from a desire to keep them around mainly so that he could daydream about having sex with them. (Hence, he remains a sexist employer.) To satisfy this interest, he would have to encapsulate enough of their interests in his own to keep the relationships going. And this would make him trustworthy on Hardin's account. But *is* he trustworthy? The answer is surely "no," if the women have an interest in being treated well in more than just a surface manner. My point here is that being motivated by a desire to maintain a relationship (the central motivation of a trustworthy person on the encapsulated interests view) may not require one to adopt all of the interests of the trustor that would actually make one trustworthy to that person. In the end, like the social contract view, the encapsulated interests view might describe only reliability, not trustworthiness.

The social contract view and the encapsulated interests view are both instances of what Karen Jones calls "risk-assessment views" of trust or trustworthiness (1999, 68). According to them, people trust other people whenever they assume that the risk of relying on other people to act a certain way is low—because it is in the self-interest of these people to act that way—and so they rely on them. Self-interest determines trustworthiness on these accounts. Risk-assessment theories in general are popular among rational decision theorists and social contract theorists who presume that people are naturally self-interested.

A different type of view is what Jones calls a "will-based" account of trustworthiness, which finds trustworthiness only where the trustee is motivated by goodwill (Jones 1999, 68). This view originates in the work of Annette Baier and it is influential, even outside of moral philosophy (e.g. in bioethics and law, especially fiduciary law; see e.g. Pellegrino and Thomasma 1993, O'Neill 2002, and Fox-Decent 2005). According to it, a trustee who is actually trustworthy will act out of goodwill toward the trustor, to what or to whom the trustee is entrusted with, or both. While proponents of risk-assessment

views would likely find the goodwill view too narrow—surely we can trust people without presuming their goodwill!—it does seem immune to criticisms that apply to those views. To summarize those criticisms: these accounts, unlike the goodwill account, fail to demand that the trustworthy person *care* at all about the trustor, or care about what he or she cares about. As we have seen, such caring appears to be central to a complete account of trustworthiness.

The particular reason why care is central is that it allows us to distinguish between trust and mere reliance. I have said the two differ, putatively, because only the former can be betrayed. But why is that true? Why can trust be betrayed, while mere reliance can only be disappointed? The answer Baier gives is that betrayal is the appropriate response to someone on whom one relied to act out of goodwill, as opposed to ill will, selfishness, or habit bred out of indifference (1986, 234–5). Those who say trusting could involve relying on people to act on *any* of these motives cannot distinguish trust from reliance. Examples are risk-assessments theorists, who, again, make trustworthiness a matter of self-interest. Although self-interest as a motive is compatible with goodwill toward others, it is also compatible with ill will and selfishness. The question is: how would trustworthiness be different from mere reliability if trust could target any of these attitudes?

While useful in distinguishing trustworthiness from reliability, or trust from reliance, Baier's will-based account is not perfect, however. Criticisms have been made of it that suggest goodwill is neither necessary nor sufficient for trustworthiness. It is not necessary, some say, because we can trust other people without presuming that they have goodwill. In fact, "[w]e are often content to trust without knowing much about the psychology of the one-trusted, supposing merely that they have psychological traits sufficient to get the job done" (Jones 2004, 4; citing Blackburn 1998). This objection implies that trust is compatible with relying on different kinds of motives, which is the case on a risk-assessment view of trust, but not on Baier's will-based view. Although they are problematic, risk-assessment accounts do have their merits. For example, they seem to make better sense of trusting that strangers will be morally decent toward us, which presumably can occur without assuming that strangers feel goodwill toward us.

One could make sense of trust in strangers without adopting a risk-assessment view, however, or in other words, without assuming that self-interest would be the motive of the trustworthy stranger. Motives that could underlie trustworthiness in the absence of goodwill include the motive to stand by one's moral commitments, to fulfill a moral obligation, or to adhere to a social norm. For example, I could trust a stranger to be decent simply by presuming that he is committed to common decency. Ultimately, what I am presuming about the stranger is moral integrity, which some say is the relevant motive for trust relations (McLeod 2002, 21–27). Others similarly identify this motive as moral obligation, and say it is ascribed to the trustee by the very act of trusting him or her (Nickel 2007; and for a similar account, see Cohen and Dienhart 2013). Out of concern that such views "moralize" trust inappropriately, Amy Mullin argues that trustworthy people are actually motivated by a commitment that may not be moral: a commitment to a particular social norm (2005, 316).

As well as being unnecessary, goodwill may not be sufficient for trustworthiness for three reasons. First, someone trying to manipulate "you"—a "confidence trickster" (Baier 1986)—could "rely on your goodwill without trusting you" (Holton 1994, 65).

Hence, Zac Cogley claims that trust involves the belief not simply that trustees will display goodwill but that they owe it to us (2012). Second, basing trustworthiness on goodwill alone cannot explain unwelcome trust. When people do not welcome your trust, they object not to your optimism about their goodwill (who would object to that?), but only to the fact that you are counting on them. Thus, optimism about goodwill is insufficient and, according to Karen Jones, needs to be coupled with the expectation that the trustee is "favorably moved by the thought that [you are] counting on her" (1996, 9). Third, you can expect people to be reliably benevolent toward you without trusting them (Jones 1996, 10). You can think that their benevolence is not shaped by the sorts of values that, for you, are essential to trustworthiness.[4] It follows that some expectation about shared values or norms may be an important element of trust (Lahno 2001, McLeod 2002, Mullin 2005, Smith 2008).

Further criticism of will-based accounts concerns how "goodwill" ought to be interpreted. In much of the discussion above, it is narrowly conceived so that it must involve friendly feeling or personal liking. Yet Jones urges us in her earlier work to understand goodwill so that it could be grounded in benevolence, conscientiousness, or the like, *or* in friendly feeling (1996, 7). In her later work, however, she worries that by defining goodwill so broadly we "turn it into a meaningless catchall that merely reports the presence of some positive motive, and one that may or may not even be directed toward the truster" (2012, 67). If this claim is correct, and the narrow understanding of goodwill is indeed too narrow, then will-based theories are in serious trouble.

To recapitulate on the topic of motive, although will-based theories are influential, they are still open to substantial criticisms. Philosophers who are less concerned with distinguishing trustworthiness from reliability reject such accounts on the grounds that they are too narrow. By contrast, those who think this distinction is important follow Baier in identifying the relevant motive as goodwill (e.g. Potter 2002); or combine reliance on goodwill with certain expectations (Jones 1996), including in one case a normative expectation of goodwill (Cogley 2012); or abandon the goodwill requirement altogether and replace it with something else, such as moral integrity or moral obligation. In addition, the literature includes some debate about just how to understand "goodwill".

Notably, the majority of views discussed so far describe trustworthiness as a relation between the trustee and some attitude, commitment, or structuring condition (e.g. a social constraint). According to a different position, found in the work of Richard Holton (1994), the conditions giving rise to trustworthiness reside not in the trustee's relationship to certain of his or her attitudes, commitments, etc., but rather in the stance that the trustor takes toward the trustee. Holton argues that this stance (a "participant stance"; Strawson 1974) involves a readiness on the part of the trustor to feel betrayal. Although this view has garnered positive attention (e.g. by Hieronymi 2008, McGeer 2008), some find it dissatisfying because it does not obviously explain what would justify a reaction of betrayal, rather than mere disappointment, when someone fails to honor one's trust (see Nickel 2007, 318). By contrast, others say that the participant stance does not do "any distinctive work" in some cases, such as when a son trusts his mother simply because he knows she loves him; and hence, a participant-stance theory of trust is subject to counter examples (Simpson 2012, 553).

However, some have expanded on Holton's theory in a way that deflects at least some

of the criticism of it. Margaret Urban Walker explains that in taking a participant stance toward others, we hold them responsible (2006, 79). Thus, we expect them to act not simply as we assume they *will*, but as they *should*. In other words, we have normative, rather than merely predictive, expectations of them. Walker believes that all trust, at bottom--not just therapeutic trust--involves having normative expectations (and Jones, in her later work, agrees; 2012a). If that is right, then to be trustworthy is to live up to these expectations, and a failure to do so can result in betrayal. This is how a participant-stance theory can explain how trust can be betrayed.

One last view about trustworthiness is that it is a virtue. Consider first why will-based accounts do not capture this view. Someone can show goodwill towards another and be trustworthy within the scope of their relationship (think of a convicted felon with his mother) without being someone who we would describe as *trustworthy* (Potter 2002, 8). Sometimes, we think of trustworthiness as a character trait that virtuous people possess. Nancy Nyquist Potter refers to the trait as “full trustworthiness,” and distinguishes it from “specific trustworthiness”: trustworthiness that is specific to certain relationships (25). To be fully trustworthy, one must have a disposition to be trustworthy toward everyone. Call this the “virtue” account.

It may sound odd to insist that trustworthiness is a virtue or, in other words, a moral disposition to be trustworthy (Potter 2002, 25; Hardin 2002, 32). What disposition exactly is it meant to be? Is it a disposition normally to honor people’s trust? That would be strange, since trust can be unwanted if the trust is immoral (e.g., being trusted to hide a murder) or if it misinterprets the nature of one’s relationship with the trusted person (e.g., being trusted to be friends with a mere acquaintance). Perhaps trustworthiness is instead a disposition to respond to trust in appropriate ways, given “who one is in relation” to the trustor and given other virtues that one possesses or ought to possess (e.g., justice, compassion) (Potter 25). This is essentially Potter’s view of trustworthiness. Modeling trustworthiness on an Aristotelian conception of virtue, she defines a trustworthy person as “*one who can be counted on, as a matter of the sort of person he or she is, to take care of those things that others entrust to one and* (following the Doctrine of the Mean) *whose ways of caring are neither excessive nor deficient”* (her emphasis; 16).[5]

Recent criticism of a virtue account comes from Karen Jones (2012a). As she explains, if being trustworthy were a virtue, then being untrustworthy would be a vice, but that cannot be right because we can never be required to exhibit a vice, yet we can be required to be untrustworthy (84). An example occurs when we are counted on by two different people to do two incompatible things and being trustworthy to the one requires that we are untrustworthy to the other (83). To defend her virtue theory, Potter would have to insist that in such a situation, either one is required simply to disappoint someone’s trust rather than to be untrustworthy, or that the kind of trustworthiness at issue is specific not full trustworthiness.

But rather than cling to a virtue theory, why not just endorse the thin conception of trustworthiness (i.e. “specific trustworthiness”), according to which X is trustworthy for me just in case I can trust X? Two things can be said. First, the thick conception—that is, of trustworthiness as a virtue—is not meant to displace the thin conception. We can and do refer to some people as trustworthy in the specific or thin sense and others as trustworthy in the full or thick sense. Second, one could argue that the thick conception explains better than the thin one why fully trustworthy people are as dependable as

they are. It is ingrained in their character. Hence, they must have an ongoing commitment, and better still, their commitment must come from a source that is compatible with trustworthiness (i.e. virtue as opposed to mere self-interest).

An account of trustworthiness that includes the idea that trustworthiness is a virtue will seem ideal only if one thinks that the genesis of the trustworthy person's commitment is important. If one thinks that it matters only whether—not how—the trustor will be motivated to act, then one could assume that social constraints and self-interest will do the job as well as a moral disposition. Such controversy explains why philosophical accounts of trustworthiness diverge from one another on the question of what could motivate a trustworthy person (e.g. social constraints, interests, goodwill, a moral disposition?).[6] Not surprisingly, a similar divergence occurs in accounts of trust, as those accounts bear on the issue of what sort of motivation we seek from people we trust.

Recent work on trust has taken one or more of the following paths in dealing with the controversy about motive. 1) It has given "genealogical" accounts of trust or trustworthiness, which involve asking why we have these concepts, given the kinds of beings we are (e.g. social, reflective; see Jones 2012a, Simpson 2012). For example, Jones argues that the concepts exist essentially because of the need we have to be able to count on one another. Hence, in her view, trustworthiness involves a willingness to factor into our practical deliberation the likelihood that others are counting on us; and it can even require that we signal this willingness to others. 2) Some philosophers have adopted a pluralist view, according to which there are different forms of trust or trustworthiness, rather than a single variant or concept of them (e.g. Simpson 2012). And 3), some have investigated the nature of betrayal closely as a way to get clearer on what trust is like, in particular its normative dimensions (O'Neil 2012, Cogley 2012, Harding 2011). This last development is particularly welcome, given how little philosophical work there has been on betrayal.

The controversy about motive still persists, however. Thankfully, in spite of it, there are things we can say for certain about trust, which are relevant to deciding when it is warranted. The trustor must be able to accept that by trusting, s/he is vulnerable, usually to betrayal. The trustee must be competent and committed to do what the trustor expects of him or her, and may have to be committed in a particular way. Last, in cases of paradigmatic trust at least, the trustor must be optimistic that the trustee is indeed competent and committed.

2. The Epistemology of Trust

Writings on the epistemology of trust obviously bear on the issue of when trust is warranted (i.e. justified). The central epistemological question about trust is, "Ought I to trust or not?" That is, given the way things seem to me, is it reasonable for me to trust? People tend to ask this sort of question only in situations where they cannot take trustworthiness for granted—that is, where they are conscious of the fact that trusting could get them into trouble. Examples are situations similar to those in which they have been betrayed in the past or unlike any they have ever been in before. Thus, the question, "Ought I to trust?" is particularly pertinent (though not restricted) to a somewhat odd mix of people that includes victims of infidelity, abuse, or the like, as well as foreign immigrants and travelers.

What guidance, if any, could one offer to these people who struggle with when to trust? Philosophical work on this question appears either under the general heading of the epistemology or rationality of trust (e.g. Baker 1987, Webb 1992, Wanderer and Townsend 2013) or under the specific heading of testimony—that is, of putting one's trust in the testimony of others (e.g. Fricker 1995, Hardwig 1991, Coady 1992, Jones 1999, Goldman 2001, Foley 2001, Daukas 2006, Faulkner 2007 and 2011, Koenig and Harris 2007, McMyler 2011, Zagzebski 2012). This section focuses on the epistemology of trust generally, rather than on trust in testimony specifically. For discussion of the latter, see related entry on testimony.

Philosophers sometimes ask whether it could ever be rational to trust other people. This question arises for two main reasons. First, it appears that trust and rational reflection (e.g. on whether one should be trusting) are in tension with one another. Since trust inherently involves risk, any attempt to eliminate the risk through rational reflection could eliminate the trust at the same time. Second, trust tends to give us blinkered vision: it makes us resistant to evidence that may contradict our optimism about the trustee (Baker 1987; Jones 1996). For example, if I trust my brother not to harm anyone, I will resist the truth of any evidence to the contrary. Here, trust and rationality seem to come apart.

But even if some of our trust could be rational, one might insist that not all of it could be rational for a number of reasons. First, if we trust people in a myriad of ways every single day, as some claim that we do (e.g. Baier 1986, 234), then we could not possibly subject *all* of our trust to rational reflection. We certainly could not reflect on every bit of knowledge we have acquired through the testimony of others, such as that the earth is round or that Afghanistan exists (Webb 1993, Fricker 1995, Coady 1992). Second, bioethicists point out that some trust is unavoidable and occurs in the absence of rational reflection (e.g. trust in emergency room nurses and physicians; see Zaner 1991). Lastly, some trust—namely therapeutic trust—purposefully leaps beyond any evidence of trustworthiness in an effort to engender trustworthiness in the trustee. Is this sort of trust rational? One might say "no," given that there is not sufficient evidence for it.

Philosophers have responded to such skepticism about the rationality of trust by saying that rationality, when applied to trust, needs to be understood differently than it is in each of the skeptical points above. There, 'rationality' means something like this: it is rational to believe in something only if one has verified that it will happen or done as much as possible to verify it. For example, it is rational for me to believe that my brother has not harmed anyone only if the evidence points in that direction and I have discovered that to be the case. However, problems exist with applying this view of rationality to trust. And the rationality of trust need not be understood according to it. While this view is both "truth-directed" and "internalist," for instance, the rationality of trust could—as discussed below—be "end-directed" or "externalist."

In discussing the rationality of trust, some authors distinguish between truth-directed (or epistemic) rationality and end-directed (or strategic) rationality (e.g. Baker 1987). One could say that one is rational in trusting emergency room physicians, for example, not necessarily because one has good reason to believe that they are trustworthy, but because by trusting them, one can remain calm in a situation over which one has little control. Similarly, it may be rational for me to trust my brother not because I have good evidence of his trustworthiness but rather because trusting him is essential to our

having a loving relationship.[7]

Thus, trust can be rational depending on whether one conceives of rationality as truth-directed or end-directed. But notice that it also matters how one conceives of trust. For example, if trust is a belief in someone's trustworthiness (see Trust and the Will), then whether the rationality of trust can be end-directed will depend on whether the rationality of a belief can be end-directed. To put the point more generally, how trust is rationally justified will depend on how beliefs are rationally justified (Jones 1996).

Some work on trust and rationality concerns whether the rationality of trust can indeed be end-directed and also what could make therapeutic trust and the like rational. Pamela Hieronymi argues that the ends for which we trust cannot provide reasons for us to trust in the first place (2008). Considerations about how useful or valuable trust is do not bear on the truth of a trusting belief, that is, a belief in someone's trustworthiness. But Hieronymi claims that trust, in a pure sense at least, always involves a trusting belief. How then does she account for trust that is motivated by how therapeutic (i.e. useful) the trust will be? She believes that trust of this sort is not pure or full-fledged trust. As she explains, people can legitimately complain about not being trusted fully when they are trusted in this way, which occurs only when other people lack confidence in them but trust them nonetheless (230; see also Lahno 2001, 184–185).

By contrast, Victoria McGeer believes that trust is *more* substantial or pure when the available evidence does not support it (2008). She describes how trust of this sort—what she calls "substantial trust"—could be rational, and does so without appealing to how important it might be or, in other words, to the ends it might serve, but instead to whether the trustee will be trustworthy.[8] According to McGeer, what makes "substantial trust" rational is that it involves hope that the trustees will do what they are trusted to do, which "can have a galvanizing effect on how [they] see themselves, as trustors avowedly do, in the fullness of their potential" (252). Rather than complain (as Hieronymi would assume that trustees might) about trustors being merely hopeful about rather than confident in their trustworthiness, they could respond well to the trustors' attitude toward them. Moreover, if it is likely that they will respond well, then the trust in them must be rational.

Philosophers who agree that trust can be rational—regardless, perhaps, of whether it is "substantial" (to use McGeer's language)—tend to disagree about the extent to which reasons that confer rationality must be accessible to the trustor. Some say that all of these reasons must be available to this person in order for the trust to be rational; in that case, the person is or could be *internally* justified in trusting as s/he does. Others say that the reasons need not be internal, but can instead be *external* to the trustor and can lie in what caused the trust, or, more specifically, in the epistemic reliability of what caused it. Furthermore, the trustor need not have access to, or be aware of the reliability of, such reasons. The latter's epistemology of trust is therefore externalist, while the former's is internalist.

Some epistemologists of trust write as though trust is only rational if the trustor him or herself has rationally estimated the likely trustworthiness of the trustee. For example, Russell Hardin implies that if my trust in you is rational, then "I make a rough estimate of the truth of [the] claim … that you will be trustworthy under certain conditions … and

then I correct my estimate, or 'update,' as I obtain new evidence on you" (2002, 112). On this view, I must have reasons for my estimate or for my updates (Hardin 130), which could come from inductive generalizations I have made about my past experience, from my knowledge that social constraints exist that will encourage your trustworthiness or what have you. Such an internalist epistemology of trust is valuable because it coheres with the common sense idea that one ought to have good or at least decent reasons for trusting other people, especially when something important is at stake (Fricker 1995).

But such an epistemology is also open to criticisms. For example, it suggests that rational trust will always be partial rather than complete, given that the rational trustor is open to evidence that contradicts his or her trust on this theory, while someone who trusts completely in someone else lacks such openness. The theory also implies that the reasons for trusting well (i.e. in a justified way) are accessible to the trustor, at some point or another, which may be false. Some reasons for trust may be too "cunning" for this to be the case. Here, I have in mind a reason for trusting others that Philip Pettit discusses (1995): trust signals to people that they are being held in esteem, which is something that they will want to maintain; they will honor our trust because they are naturally "esteem-seeking." However, consciously having this reason for trust is incompatible with actually trusting (Wanderer and Townsend 2013, 9), that is, if trust involves being optimistic that people will act out of motive other than self-interest.

Others suggest that reasons for trust are often too numerous and varied to be open to the conscious consideration of the trustor (e.g. Baier 1986). There can be very subtle reasons to trust or distrust someone—for example, reasons that have to do with body language, with systematic yet veiled forms of oppression, or with a complicated history of trusting others about which one could not easily generalize. Such factors may influence the trustor without him or her knowing it.

The above worries explain why some philosophers defend externalist epistemologies of trust. Some do so explicitly (e.g. McLeod 2002). They argue for reliabilist theories that make trust rationally justified if and only if it is formed and sustained by reliable processes (i.e. "processes that tend to produce accurate representations of the world"; Goldman 1992, 113). Others gesture towards externalism (Webb 1993; Baier 1986), as Baier does with what she calls "a moral test for trust." The test is that "knowledge of what the other party is relying on for the continuance of the trust relationship would … itself destabilize the relation" (1986, 255). The other party might be relying on concealment of untrustworthiness or on a threat advantage, in which case the trust would probably fail the test. Because Baier's test focuses on the causal basis for trust, or for what maintains the trust relation, it is externalist. Moreover, because the trustor often cannot gather the information needed for the test without ceasing to trust the other person (Baier 1986, 260), the test cannot be internalist.

Although an externalist theory of trust deals well with some of the worries one might have with an internalist theory, it has problems of its own. The most serious problem, perhaps, is the absence of any requirement that trustors themselves have good reasons for trusting, especially when their trust makes them seriously vulnerable. Again, it appears that common sense dictates the opposite: that sometimes we, as trustors, ought to be able to back up our decisions about when to trust.

Presumably to avoid having to defend either internalism or externalism, some

philosophers simply provide a list of common “justifiers” for trust (i.e. “facts or states of affairs that determine the justification status of [trust]”; Goldman 1999, 274), which a trusting agent could take into account in deciding when to trust (Govier 1998, Jones 1996). The lists include such factors as the social role of the trustee, the domain in which the trust occurs, an “agent-specific” factor that concerns how good a trustor the agent tends to be (Jones 1996, 21), and the social or political climate in which the trust occurs. Philosophers have tended to emphasize that last factor as a justification condition for trust, and so I will describe it in some detail.

While, paradigmatically, trust is a relation that holds between two individuals, forces larger than those individuals inevitably shape their trust in one another. Social or political climate contributes to how trustworthy people tend to be and therefore to whether trust is justified. A climate of virtue, for example, is one in which trustworthiness tends to be pervasive, because the virtues (i.e. those other than trustworthiness) tend to enhance trustworthiness (Baier 2004). Similarly, a democratic society is one in which people can trust one another more often than in other sorts of political societies, such as totalitarian ones (Uslaner 1999). Alternatively, societies that are oppressive make it irrational in general for the people who are oppressed to trust those who oppress them (Baier 1986, 259; Potter 2002, 24). “Social trust,” as some call it, is low in these circumstances (Govier 1997, Welch 2013).

Social or political climate has a significant influence on the default stance that one ought to take toward people’s trustworthiness (see e.g. Walker 2006). One needs such a stance because one cannot always stop to reflect carefully on when trust is appropriate. For example, some say that the correct stance toward the testimony of others is always trust: that people always have a *prima facie* right to trust what other people say (e.g. Coady 1992). Others disagree that the correct stance could be so universal and claim instead that it is relative to climate, as well as to other factors listed above, such as domain (Jones 1999).

Our trust or distrust may be *prima facie* justified if we have the correct default stance, although most philosophers assume that it could only be fully justified either by reasons that are internal to us or by the causal processes that created the attitude in the first place. (Lately, however, an alternative to pure internalism or externalism about trust has been developed, though only with respect to trust in people’s testimony. The theory is an “inheritance model” of testimony, according to which “by trusting a speaker, audiences allow their beliefs to inherit the evidential support enjoyed by the speaker’s belief” (Keren 2014, 2612, citing McMyler 2011). The trustor/hearer does not gather this support him or herself, and thus is not aware of the reasons it provides in favour of his or her beliefs. Nevertheless, s/he must have reasons for thinking that the speaker is speaking knowledgeably and sincerely. In this way, some, but not all, of the reasons that justify her trust are internal to him or her.)

But whichever epistemology of trust we choose, it ought to be sensitive to the tension that exists between trusting somebody and rationally reflecting on the grounds for that trust. It would be odd, to say the least, if what made an attitude justified destroyed that attitude. At the same time, however, if possible, our epistemology ought to cohere with common sense, insofar as common sense requires that we inspect rather than have pure faith in whatever makes us seriously vulnerable to other people, which trust can most definitely do.

3. The Value of Trust

The one who asks, “When is trust warranted?” might be interested in knowing what the point of trust is. In other words, what value does it have? Although the value it has for a particular person will depend on his or her circumstances, the value it *could* have for any particular person will depend on why trust is valuable, generally speaking. Trust can have enormous instrumental value and may also have some intrinsic value. In discussing its instrumental value, I will refer to the “goods of trust,” which include opportunities for cooperative activity, meaningful relationships, knowledge, autonomy, self-respect, and overall moral maturity. Because these goods may benefit the trustor, the trustee, or society in general, they are therefore social as well as individual goods, where the most relevant individuals tend to be parties to the trust relationship. A further point about these goods is that they accompany *justified* trust, rather than any old trust.[9]

However, if trust produced no goods independent of it, then would there be no point in trusting? One might say “no,” on the grounds that trust is (or at least can be; O’Neil 2012, 311) a sign of respect for others which makes it intrinsically worthwhile. This idea is closely linked to the view that trustworthiness is a virtue, which makes it intrinsically worthwhile (see The Nature of Trust and Trustworthiness). Trust would be a sign of respect for others if it were an attitude of optimism about the trustee’s character: that is, if it assumed that virtue resided within this person’s character. Moreover, trust that has intrinsic value of this sort presumably must be justified. If optimism about the person’s character was inappropriate, then the respect would be misplaced and the intrinsic value would be lost. But I am speculating here. Philosophers have said comparatively little about trust being worthwhile in itself as opposed to worthwhile because of what it produces, or because of what accompanies it. Thus, I will say no more about the former and focus on the latter, in particular, on the goods of trust.

Turning first to the instrumental value of trust to the trustor, some argue that trusting vastly increases our opportunities for cooperating with others and for benefiting from that cooperation, although of course we would only benefit if people we trusted cooperated as well (Gambetta 1988; Hardin 2002). Trust enhances cooperation (but may not be necessary for it; Cook *et al* 2005, Skyrms 2008). Because trust removes the incentive to check up on other people, it makes cooperation with trust less complicated than cooperation without it (Luhmann 1979). This point applies to *justified* trust only if trust is justified in external rather than internal ways, however (see The Epistemology of Trust). If trust were only justified when the trustor could produce legitimate reasons for why the trustee is trustworthy, then cooperation with justified trust would not be a whole lot easier than cooperation without it. That is true, at least, if having proper reasons required frequently checking up on the trustee.

Trust makes cooperation possible, rather than simply easier, if trust is essential to promising. Daniel Friedrich and Nicholas Southwood defend what they call the “Trust View” of promissory obligation (2011), according to which “making a promise involves inviting another individual to trust one to do something” (Friedrich and Southwood 277). If this view is correct, then cooperation through promising is impossible without

trust. Moreover, cooperation of this sort will not be fruitful unless the trust is justified.

Trusting provides us with goods beyond those that come with cooperation, although again, for these goods to materialize, the trust must be justified. Sometimes, trust involves little or no cooperation, so that the trustor is completely dependent on the trustee, although the reverse is not true. Examples are the trust of young children in their parents and the trust of severely ill or disabled people in their care providers. Trust is particularly important for these people, because they tend to be powerless to exercise their rights or to enforce any kind of contract (which is not to say they could not be parties to a contract, including the social contract; see Silvers and Francis 2005). Moreover, since the trust that the ill or disabled place in their care providers contributes to them being vulnerable, it is essential that they *can* trust these people; that is, it is important that their trust be justified. The goods at stake for them are all the goods involved in having a good or decent life.

Among the specific goods that philosophers associate with trusting are meaningful relationships with others (rather than simply cooperative relationships that further individual self-interest; Harding 2011), knowledge, and autonomy.[10] Let me expand on the points about knowledge and autonomy. Philosophers writing on testimony argue that scientific knowledge (Hardwig 1991), moral knowledge (Jones 1999), or almost all knowledge in fact (Webb 1993) depends for its acquisition on trust in the testimony of others. The basic argument for the need to trust what others say is that no one person has the time, intellect, and experience necessary to learn, independently, facts about the world that many of us do know. Examples include the scientific fact that the earth is round, the moral fact that the oppression of people from social groups different from our own can be severe (Jones 1999), and the mundane fact that we were born on such-in-such a day (Webb 1993, 261). Of course, trusting the people who testify to these facts could only generate knowledge if that trust were justified. If we were told our date of birth by people who were determined, oddly, to deceive us about when we were born, then we would not know when we were born.

Autonomy is another good that flows from trust, at least insofar as being autonomous is a skill that we acquire and exercise only in social environments where we can trust people to support it. Feminists in particular tend to conceive of autonomy this way—that is, as a relational, or socially-constituted, property (Mackenzie and Stoljar 2000). Many feminists note that sexist or oppressive social environments can inhibit autonomy, and some say explicitly that conditions necessary for autonomy (e.g. adequate options, knowledge relevant to one's decisions) can exist only with the help of people or institutions that are trustworthy (McLeod 2002, Oshana 2014). Justified trust in others to ensure that these conditions exist is essential for autonomy, if in fact autonomy is relational.

The kind of trust that is necessary for both knowledge and autonomy is arguably self-directed as well as other-directed. Keith Lehrer claims that to be able to do the epistemic work that his internalist epistemology requires of us (i.e. evaluating the truth of our beliefs and the worth of our desires), we need to trust ourselves to do this work (1997). Richard Foley similarly argues that to have knowledge, including knowledge acquired through others' opinions and through our own past opinions, we need ultimately to be able to trust ourselves (2005). (For a related view, one that links self-trust with the "diachronic exercise of practical reason," see Hinchman 2003). Trudy Govier (1993) and Carolyn McLeod (2002) both assert that to be motivated to choose

and act in accordance with our own values—that is, to choose and act autonomously—we need to trust ourselves to do so. Although some self-trust might be better than none at all as far as knowledge and autonomy go (for no self-trust would leave us incapacitated and with no opportunity to learn from our mistakes), justified self-trust is best overall. Without being justified in trusting ourselves to be good epistemic or autonomous agents, we cannot *be* either.

Specific goods of trust that are instrumental to the well-being of the trust*ee* also materialize only if the trust is justified. Trust can improve the self-respect and moral maturity of this person. In particular, if trust involves optimism about a person's moral character, it can engender self-respect in that person. (For on such an account of trust, to trust is to show respect, which can then be internalized.) To be trusted can allow us to be more respectful not only toward ourselves, but also toward others. It can therefore enhance our overall moral maturity. The explicit goal of therapeutic trust is precisely to bring about such maturity (see The Epistemology of Trust). Therapeutic trust leaps ahead of the evidence, which means that it is hard to justify epistemically. But as suggested above (The Epistemology), perhaps it can be justified in a truth-directed way over time, provided that the trust has its intended effect of making the trustee more trustworthy (McGeer 2008; Baker 1987, 12). Clearly, for therapeutic trust to benefit trusted people, it would have to be justified in this way (i.e. the therapy would normally have to work).

The social goods of trust are linked with the individual goods that concern moral maturity and cooperation. These social goods include the practice of morality, the very existence of society perhaps, as well as strong social networks. Morality itself is a cooperative activity, which can only get off the ground if people can trust one another to try, at least, to be moral. Yet to be able to make meaningful attempts in this regard, people have to be somewhat morally mature, which can only come from a moral education grounded in trust. For such reasons, Baier claims that trust is "the very basis of morality" (2004, 180). It could also be the very basis of society, insofar as trust in fellow citizens to honor social contracts makes those contracts possible. Even if trust did not make society possible, however, it certainly makes it better or more livable. Some argue that trust is a form of "social capital," meaning roughly that it enables "people to work together for common purposes in groups and organizations" (Fukuyama 1995, 10; quoted in Hardin 2002, 83). Hence, "high-trust" societies have stronger economies and stronger social networks in general than "low-trust" societies (Fukuyama; Inglehart 1999). Of course, these facts could only be true of these societies if, on the whole, the trust within them was justified (that is, if trustees tended not to "defect" and destroy chances for cooperating in the future).

If without trusting or being trusted in justified ways, we could not have morality or society and could not be morally mature, autonomous, knowledgeable, or invested with opportunities for collaborating with others, then the value of justified trust is hard to over-estimate. The same is not true, however, of trust in general (Baier 1986). People can trust too much or too little; and either way, their trust can be harmful since it can deprive them of the goods that go along with justified trust. Too much trust in particular leaves people open to betrayal, abuse, terror, and deception. And since there are people who tend to elicit too much trust from others, the question, "Why be *dis*trusting?" is as legitimate as "Why be trusting?" (For responses to the former question, see Hardin 2004.)

4. Trust and the Will

Trust may not be warranted (i.e. plausible) because the agent has lost the ability to trust. People lose this ability often as a result of trauma (Herman 1991). The trauma caused by rape, for example, can profoundly reduce one's sense that the world is a safe place, with caring people in it. How can people ever recover this trusting sensibility? A similar but broader question, "How can trust be restored once it has been lost?", is relevant to people who lose trust not in everyone or everything, but rather in particular people or particular institutions. For example, in some parts of the world people tend to trust the medical profession much less than they did in the past (O'Neill 2002; Pellegrino 1991b). How could their trust in this profession be restored? Although often destroying trust is quick and dirty, creating trust is slow and painful (Uslaner 1999; Baier 1986). The reasons have to do with what kind of mental attitude trust is. It is not the sort of attitude that we can simply will ourselves to have, although we can cultivate it.

While the cultivation of any trust depends on what sort of mental attitude trust is, the cultivation of *justified* trust depends on how trust is justified (see The Epistemology of Trust). Some philosophers, most notably Baier, deny that useful rules exist for when to trust so that one's trust will be justified. The process of trusting well is too complicated for that to be the case. Even so, giving some guidance on how to trust well is possible: for example, Baier and others list factors that will at least improve our chances of trusting well, if we take them into account (see The Epistemology). This section focuses on how to trust at all, rather than on how to trust well. To receive the goods of trust, one has to do both: trust and trust well.

Under the heading "how to trust," consider why trust cannot be willed but can be cultivated. Baier questions whether people are able "to trust *simply* because of encouragement to trust" (1986, 244; my emphasis). She writes, "'Trust me!' is for most of us an invitation which we cannot accept at will—either we do already trust the one who says it, in which case it serves at best as reassurance, or it is properly responded to with, 'Why should and how can I, until I have *cause* to?'" (my emphasis; 1986, 244). Baier is not a voluntarist about trust, just as most people are not voluntarists about belief. In other words, she thinks that we cannot simply decide to trust for purely motivational rather than epistemic reasons (i.e., merely because we want to, rather than because we have reason to think that the other person is actually trustworthy; Mills 1998). That many people feel compelled to say, "I wish I could trust you," suggests that Baier's view is correct: wishing or wanting is not enough. But Holton interprets Baier's view differently. According to him, Baier's point is that we can never decide to trust, not that we can never decide to trust for motivational purposes (1994). This interpretation ignores, however, the attention that Baier gives to situations in which all we have is encouragement (trusting "*simply* because of encouragement"). The "cause" she refers to ('Why should and how can I, until I have *cause* to [trust]?') is an epistemic cause. Once we have one of those, we can presumably decide whether to trust on the basis of it. But we cannot decide to trust simply because we want to, according to Baier.

Thus, trust resembles belief in being non-voluntary in the above sense. *Is* trust a belief

then (i.e., a belief in someone's trustworthiness)? While many philosophers believe that it is (e.g., Hieronymi 2008), others insist that it is not (e.g. Faulkner 2007). The latter say it is possible to trust without believing that the trustee is trustworthy. Holton gives the nice example of trusting a friend to be sincere without believing that the friend will be sincere (75). Arguably, if one already believed that to be the case, one would have no need to trust the friend. It is also possible to believe that someone is trustworthy without trusting that person, which suggests that trust could not *merely* be a belief in someone's trustworthiness (McLeod 2002, 85). I might think that a particular person is trustworthy without trusting him or her because I have no cause to do so.

One reason for thinking that trust is not a belief, but is instead an emotion, is that trust resembles an emotion in having characteristics that are unique to emotions, at least according to an influential account of them (de Sousa 1987, Calhoun 1984, Rorty 1980, Lahno 2001). These characteristics concern how emotions narrow our perception to certain "fields of evidence": those fields that support the emotion (Jones 1996, 11). Thus, when we are in the grip of an emotion, we tend to see those facts that affirm its existence and are resistant to facts that negate it. To illustrate, if I am really angry at my mother, then I tend to think about the things that justify my anger while ignoring or refusing to see the things that make it unjustified. I can only see those other things once my anger subsides. Similarly with trust: if I truly trust my mother in certain domains, my attention falls on those aspects of her that justify my trust and is averted from evidence that suggests she is untrustworthy in these domains (Baker 1987). (For a view about how trust could be a belief yet still possess such a characteristic—that is, an insensitivity to counter evidence—see Keren 2014.)

The characteristics of trust that indicate it might be an emotion are ones that we can try to mimic in our attitude toward other people, in an effort to be more trusting. In other words, we could purposefully try to focus our attention on what makes other people trustworthy, and in doing so cultivate trust in them. Our goal could simply be self-improvement: that is, becoming more trusting, hopefully in a good way, so that we reap the benefits of justified trust. Alternatively, we might strive for the improvement of others: that is, making them more trustworthy. (See the above discussions of therapeutic trust.)

Institutions or groups trying to restore public trust could also use the above method of cultivating trust. For example, the medical profession could try to hammer home to the public the message that most physicians and other medical professionals help people and care about them. For many, such a campaign could make it impossible for them *not* to think that these professionals are trustworthy. (However, for others such as Black people in the United States who have social histories of serious harm, including exploitation, by the medical profession, the result may be very different. See e.g. Almassi 2014.)

Whether such a campaign is even morally appropriate would depend on whether the resulting trust would be justified. In general, cultivating trust is only morally wise if trusting would be wise in the circumstances, which in turn would depend on whether factors are present that roughly indicate that trust would be justified. Is, for example, the social or political climate of one's society conducive to one trusting well? It may not be. Consider a woman who has been raped in a society that effectively condones such acts toward women. For her to cultivate the trust that was lost as a result of the rape may be imprudent, especially if she continues to live in such a sexist place.

5. Conclusion

This entry has centered on an important practical question about trust: “When is trust warranted?” Different answers to this question give rise to different philosophical puzzles. For example, in responding to the question, one might appeal to the nature of trust and trustworthiness, and consider whether the attitude of the proposed trustor could support trust, and whether the qualities of the proposed trustee indicate that this person is trustworthy. But to decide such matters, one must first settle the difficult issue of whether, to be trustworthy, a person must have a particular kind of motive for acting.

Alternatively, in deciding whether trust is warranted, one might consider whether the trust would either be rationally justified or valuable. The two can overlap, and do overlap when rational justification is understood in an end-directed way—that is, where trust is rationally justified because it is instrumentally valuable, or because it serves one’s ends. Regardless of whether the justification of trust is end-or truth-directed, however, the exact nature of its justification is puzzling. Should the rationality of trust or distrust be interpreted using an internalist epistemology or an externalist one (or an epistemology that combines aspects of the two)? Because good arguments exist on different sides, it is not clear how trust is rationally justified. Neither is it entirely clear what sort of value trust can have, given the nature of trust. For example, trust may or may not have intrinsic moral value depending on whether it signals respect for others.

Lastly, one might focus on the fact that trust cannot be warranted when it is impossible, which is the case when the agent does not already trust and cannot simply will him or herself to do so. While it appears that trust is not the sort of attitude that one can will oneself to have, trust can still be cultivated. The manner in which it is cultivated, however, would depend again on what sort of attitude it is.

Since one can respond to the question, “When is trust warranted?” by referring to each of the above dimensions of trust, a complete philosophical answer to this question is inevitably complex. But such an answer would also be philosophically intriguing and socially important. It would be intriguing both because of its complexity and because it would draw on a number of philosophical areas, including epistemology, philosophy of mind, and value theory. It would be important because trust that is warranted contributes to the foundation of a good society. It helps people to thrive through healthy cooperation with others and to be morally mature human beings.

Bibliography

Almassi, B., 2014. “Trust and the Duty of Organ Donation,” *Bioethics*, 28(6): 275–283.

Baier, A. C., 1986. “Trust and Antitrust,” *Ethics*, 96: 231–260.

——, 1991. “Trust and Its Vulnerabilities” and “Sustaining Trust,” *Tanner Lectures on Human Values*, Volume 13, Salt Lake City: University of Utah Press.

——, 1995. *Moral Prejudices: Essays on Ethics*, Cambridge, MA: Harvard University Press.

——, 2004. “Demoralization, Trust, and the Virtues,” in Calhoun (ed.) 2004.

Baker, J., 1987. “Trust and Rationality,” *Pacific Philosophical Quarterly*, 68: 1–13.

Becker. L. C., 1996. “Trust as Noncognitive Security about Motives,” *Ethics*, 107(1): 43–61.

Blackburn, S., 1998. *Ruling Passion: A Theory of Practical Reasoning*, Oxford: Clarendon Press.

Calhoun, C. (ed.), 1984. “Cognitive Emotions?” in *What is an Emotion?*, C. Calhoun and R. C. Solomon (eds.), New York: Oxford University Press.

——, 2004. *Setting the Moral Compass: Essays by Women Philosophers* , New York, NY: Oxford University Press.

Coady, C.A.J., 1992. *Testimony: A Philosophical Study*, Oxford: Clarendon Press.

Coeckelbergh, M., 2012. “Can we trust robots?” *Ethics and Information Technology*, 14: 53–60.

Cogley, Z., 2012. “Trust and the Trickster Problem,” *Analytic Philosophy*, 53(1): 30–47.

Cohen, M.A., and J. Dienhart, 2013. “Moral and Amoral Conceptions of Trust, with an Application to Organizational Ethics,” *Journal of Business Ethics*, 112: 1–13.

Cook, K. R., Hardin, R. and M. Levi, 2005. *Cooperation Without Trust?*, New York: Russell Sage Foundation.

Dasgupta, P., 1988. “Trust as a Commodity,” in Gambetta (ed.) 1988.

Daukas, N., 2006. “Epistemic Trust and Social Location,” *Episteme*, 3(1–2): 109–124.

de Sousa, R., 1987. *The Rationality of Emotion*, Cambridge: MIT Press.

Faulkner, P., 2007. “A Genealogy of Trust,” *Episteme*, 4(3): 305–321.

——, 2011. *Knowledge on Trust*, Oxford: Oxford University Press.

Foley, R., 2001. *Intellectual Trust in Oneself and Others*, Cambridge: Cambridge University Press.

——, 2005. “Universal Intellectual Trust,” *Episteme*, 2(1): 5–11.

Fox-Decent, E., 2005. “The Fiduciary Nature of State Legal Authority,” *Queen’s Law Journal*, 31: 259–310.

Fricker, E., 1995. “Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony,” *Mind*, 104 (414): 393–411.

Friedman, M., 2004. “Diversity, Trust, and Moral Understanding,” in Calhoun (ed.) 2004.

Friedrich, D. and N. Southwood, 2011. “Promises and Trust,” in *Promises and Agreements: Philosophical Essays*, H. Sheinman (ed.), Oxford: Oxford University Press.

Fukuyama, F., 1995. *Trust: The Social Virtues and the Creation of Prosperity*, New York, NY: The Free Press.

Gambetta, D. (ed.), 1988. *Trust: Making and Breaking Cooperative Relations*, New York: Basil Blackwell.

——, 1988b. “Can We Trust Trust?” in Gambetta (ed.) 1988.

Goering, S., 2009. “Postnatal Reproductive Autonomy: Promoting Relational Autonomy and Self-Trust in New Parents,” *Bioethics*, 23(1): 9–19.

Goldman, A. I., 1992. *Liaisons: Philosophy Meets the Cognitive and Social Sciences*, Cambridge, MA: MIT Press.

——, 1999. "Internalism Exposed," *Journal of Philosophy*, 96(6): 271–293.

——, 2001. "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research*, LXIII (1): 85–110.

Govier, T., 1993. "Self-Trust, Autonomy, and Self-Esteem," *Hypatia*, 8(1): 99-120.

——, 1997. *Social Trust and Human Communities*, Montreal and Kingston: McGill-Queen's University Press.

——, 1998. *Dilemmas of Trust*, Montreal and Kingston: McGill-Queen's University Press.

Hardin, R., 1996. "Trustworthiness," *Ethics*, 107: 26–42.

——, 2002. *Trust and Trustworthiness*, New York, NY: Russell Sage Foundation.

—— (ed.), 2004, *Distrust*, New York: Russell Sage Foundation.

Harding, M., 2011. "Responding to Trust," *Ratio Juris*, 24(1): 75–87.

Hardwig, J., 1991. "The Role of Trust in Knowledge," *The Journal of Philosophy*, 88(12): 693–708.

Herman, J. L., 1991. *Trauma and Recovery*, New York: Basic Books.

Hertzberg, L., 1988. "On the Attitude of Trust," *Inquiry*, 31: 307–322.

Hieronymi, P., 2008. "The Reasons of Trust," *Australasian Journal of Philosophy*, 86(2): 213–236.

Hinchman, E. S., 2003. "Trust and Diachronic Agency," *Noûs*, 37(1): 25–51.

Hollis, M., 1998. *Trust Within Reason*, Cambridge: Cambridge University Press.

Holton, R., 1994. "Deciding to Trust, Coming to Believe," *Australasian Journal of Philosophy*, 72(1): 63–76.

Horsburgh, H.J.N., 1960. "The Ethics of Trust," *Philosophical Quarterly*, 10: 343–354.

Inglehart, R., 1999. "Trust, Well-being and Democracy." In Warren ed. 1999.

Jones, K., 1996. "Trust as an Affective Attitude," *Ethics*, 107: 4–25.

——, 1999. "Second-Hand Moral Knowledge," *Journal of Philosophy*, 96(2): 55–78.

——, 2004. "Trust and Terror," in *Moral Psychology: Feminist Ethics and Social Theory*, P. DesAutels and M. U. Walker (eds.), Lanham, Maryland: Rowman & Littlefield.

——, 2012a. "Trustworthiness," *Ethics*, 123(1): 61–85.

——, 2012b. "The Politics of Intellectual Self-Trust," *Social Epistemology*, 26(2): 237–251.

Keren, A., 2014. "Trust and Belief: A Preemptive Reasons Account," *Synthese*, 191: 2593–2615.

Koenig, M. A. and P. L. Harris, 2007. "The Basis of Epistemic Trust: Reliable Testimony or Reliable Sources?" *Episteme*, 4(3): 264–284.

Lagerspetz, O., 1998. *Trust: The Tacit Demand*, Dordrecht: Kluwer.

Lahno, B., 2001. "On the Emotional Character of Trust," *Ethical Theory and Moral Practice*, 4: 171–189.

Lehrer, K., 1997. *Self-Trust: A Study of Reason, Knowledge, and Autonomy*, New York: Oxford University Press.

Luhmann, N., 1979. *Trust and Power*, Toronto: Wiley.

Mackenzie, C. and N. Stoljar (eds.), 2000. *Relational Autonomy: Feminist Perspectives on Autonomy, Agency, and the Social Self*, New York: Oxford University Press.

McGeer, V., 2008. "Trust, Hope, and Empowerment," *Australasian Journal of Philosophy*, 86(2): 237–254.

McLeod, C., 2002. *Self-Trust and Reproductive Autonomy*, Cambridge, MA: MIT Press.

McMyler, B., 2011. *Testimony, Trust, and Authority*, Oxford: Oxford University Press.

Mills, E., 1998. "The Unity of Justification," *Philosophy and Phenomenological Research*, 58(1): 27–51.

Mullin, A., 2005. "Trust, Social Norms, and Motherhood," *Journal of Social Philosophy*, 36(3): 316–330.

Nickel, P. J., 2007. "Trust and Obligation-Ascription," *Ethical Theory and Moral Practice*, 10: 309–319.

O'Neil, C., 2012. "Lying, Trust, and Gratitude," *Philosophy and Public Affairs*, 40(4): 301–333.

O'Neill, O., 2002. *Autonomy and Trust in Bioethics*, Cambridge: Cambridge University Press.

Oshana, M., 2014. "Trust and Autonomous Agency," *Res Philosophica*, 91(3): 431–447.

Pellegrino, E. D., 1991. "Trust and Distrust in Professional Ethics," in Pellegrino *et al*. (eds.) 1991.

Pellegrino, E. D., R. M. Veatch and J. P. Langan (eds.), 1991, *Ethics, Trust and the Professions: Philosophical and Cultural Aspects*, Washington, DC: Georgetown University Press.

Pellegrino, E. D., and D. C. Thomasma, 1993. *The Virtues in Medical Practice*, New York: Oxford University Press.

Pettit, P., 1995. "The Cunning of Trust," *Philosophy and Public Affairs*, 24: 202–225.

Potter, N. N., 2002. *How Can I be Trusted? A Virtue Theory of Trustworthiness*, Lanham, Maryland: Rowman & Littlefield.

——, 2013. "Narrative Selves, Relations of Trust, and Bipolar Disorder," *Philosophy, Psychiatry, & Psychology*, 20(1): 57–65.

Rorty, A. O., 1980. "Explaining Emotions," in *Explaining Emotions*, A. O. Rorty (ed.), Berkeley: University of California Press.

Skyrms, B., 2008. "Trust, Risk, and the Social Contract," *Synthese*, 160: 21–25.

Silvers, A. and L. Pickering Francis, 2005. "Justice through Trust: Disability and the 'Outlier Problem' in Social Contract Theory," *Ethics*, 116: 40–76.

Simpson, T. W., 2012. "What is Trust?" *Pacific Philosophical Quarterly*, 93: 550–569.

Smith, M. N., 2008. "Terrorism, Shared Rules and Trust," *The Journal of Political Philosophy*, 16(2): 201–219.

Strawson, P. F., 1974. “Freedom and Resentment,” in his *Freedom and Resentment*, London: Methuen.

Taddeo, M., and L. Floridi, 2011. “The case for e-trust,” *Ethics and Information Technology*, 13: 1–3.

Townley, C., and J. L. Garfield, 2013. “Public Trust,” in *Trust: Analytic and Applied Perspectives*, P. Makela and C. Townley (eds.), Amsterdam: Rodopi Press.

Uslaner, E. M., 1999. “Democracy and Social Capital,” in Warren (ed.) 1999.

Walker, M. U., 2006. *Moral Repair: Reconstructing Moral Relations After Wrongdoing*, Cambridge: Cambridge University Press.

Wanderer, J., and L. Townsend, 2013. “Is it Rational to Trust?” *Philosophy Compass*, 8(1): 1–14.

Warren, M. (ed.), 1999. *Democracy and Trust*, Cambridge: Cambridge University Press.

Webb, M. O., 1992. “The Epistemology of Trust and the Politics of Suspicion,” *Pacific Philosophical Quarterly*, 73: 390–400.

——, 1993. “Why I Know About as Much as You: A Reply to Hardwig,” *The Journal of Philosophy*, 90(5): 260–270.

Welch, S., 2013. “Transparent trust and oppression,” *Critical Review of International Social and Political Philosophy*, 16(1): 45–64.

Zagzebski, L. T., 2012. *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*, New York: Oxford University Press.

Zaner, R., 1991. “Trust and the Patient-Physician Relationship,” in Pellegrino *et al*. (eds.) 1991.

Academic Tools

How to cite this entry.

Preview the PDF version of this entry at the Friends of the SEP Society.

Look up this entry topic at the Indiana Philosophy Ontology Project (InPhO).

Enhanced bibliography for this entry at PhilPapers, with links to its database.

Other Internet Resources

- ‘A Question of Trust,’ BBC Reith Lectures.
- Russell Sage Foundation Program on Trust.
- Building Trust, *OECD Observer*.
- ‘Trust and Mistrust,’ Philosophy Talk.

Related Entries

Aristotle, General Topics: ethics | autonomy: in moral and political philosophy | autonomy: personal | criminal law, theories of | emotion | epistemology: social | faith | feminist philosophy, interventions: social epistemology | feminist philosophy, topics: perspectives on the self | friendship | justification, epistemic: internalist vs. externalist conceptions of | rights: of children | scientific knowledge: social dimensions of | testimony: epistemological problems of

Acknowledgments

Thanks to Julie Ponesse and Ken Chung for their research assistance, to Andrew Botterell for his helpful comments, and to the Lupina Foundation for funding.

Browse

Table of Contents
What's New
Random Entry
Chronological
Archives

About

Editorial Information
About the SEP
Editorial Board
How to Cite the SEP
Special Characters
Advanced Tools
Contact

Support SEP

Support the SEP
PDFs for SEP Friends
Make a Donation
SEPIA for Libraries

Mirror Sites

View this site from another server:

USA (Main Site)
CSLI, Stanford University

Library of Congress Catalog Data: ISSN 1095-5054

Open Call for E-Waste in Exchange for Hope Coins Cryptocurrency

Written on March 12th, 2019 by Rain Wong

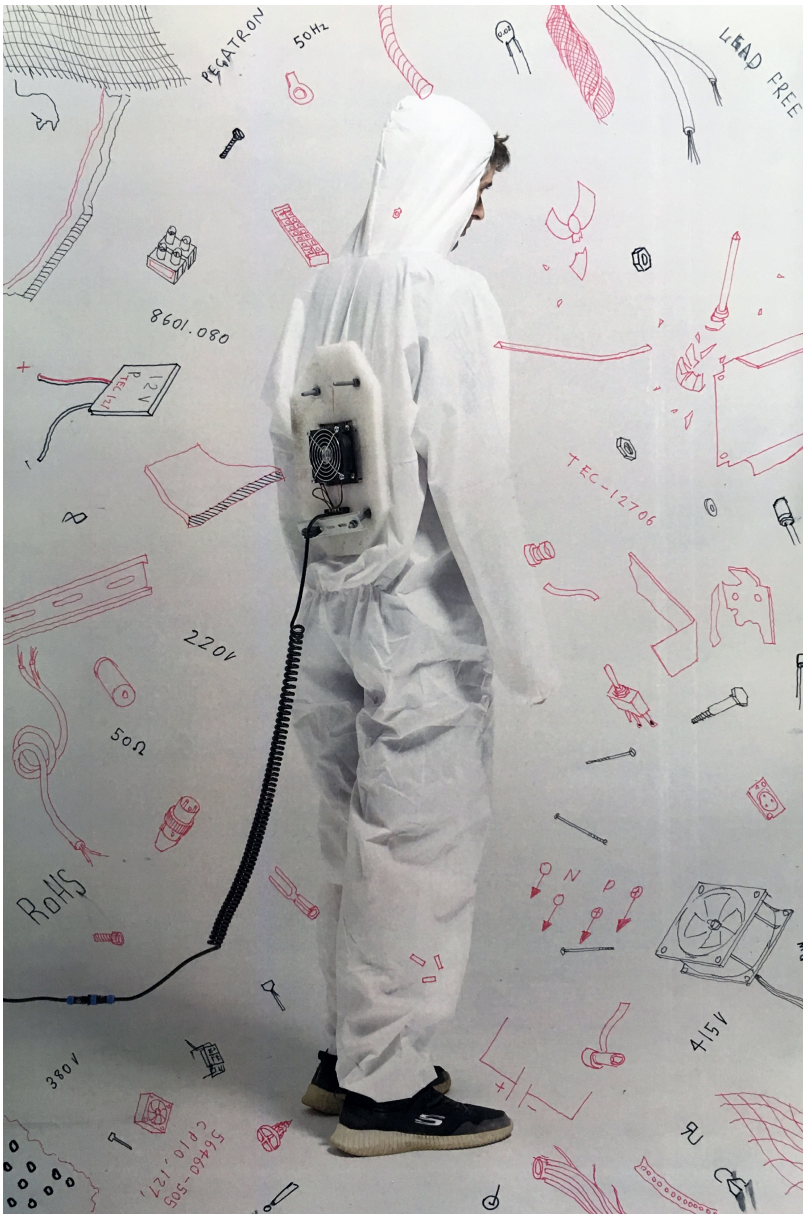

PRESS RELEASE - FOR IMMEDIATE RELEASE.

Hong Kong — March 13, 2019 — Today an Open Call for E-Waste was made by artist Matt Hope, Oui.Gallery Hong Kong, Qi Hardware and Fabricatorz Foundation. Beginning March 13, an E-waste Collection Point in Central Hong Kong is open from 11 am - 7 pm at Oui.Gallery Hong Kong to make the currency exchange to Hope Coins as part of Matt Hope's "Entropic Bodies" project.

The Hope Coin is a cryptocurrency based upon Ethereum and is transferrable based upon the number of minutes one spends collecting e-waste between today March 13 and March 31, 2019. During the last week of March, in sync with Hong Kong Art Week, artist Matt Hope will lead a community of volunteers in transforming the e-waste into artwork, valued in Hope Coins.

The total supply of Hope Coins is 25,920 "Hopes," the number of minutes between the initiation and end of the project. The value of the coin is governed based upon volunteerism in collecting e-waste and the works that artist Matt Hope creates along with volunteers, valued in Hope coins. On Friday, March 29, 2019, "Entropic Bodies" artworks produced by Matt Hope will be available in Hope Coins at Oui.Gallery Hong

Kong.

“The fight for zero waste artwork has begun,” said Beijing-based British artist Matt Hope. “My previous work creating the Breathing Bike from e-waste, 10+ years living in Beijing, and deep thinking while making ink drawings has brought me to this critical point. We must be conscious as artists and citizens of earth to know and understand how much energy we use, often in excess as waste. We must be willing to become more efficient in not only the production of artwork, but also in how we live our lives on planet Earth.”

The Open Call for E-Waste is specifically looking for electronic waste and other materials including: packing foam, electric wire, foam board, electronics, electric connectors, insulation, insulation foil, pharmaceutical and other plastic casings, cables, small appliances, and rechargeable batteries. The e-waste may be brought to the newly launched Oui.Gallery Hong Kong from 11 am - 7 pm, Tuesday - Sunday.

About Matt Matt Hope

Matt Hope (born 1976, London, U.K.) lives and works in Beijing, China. Hope received his M.F.A. from the University of California, San Diego in 2004. Selected recent solo exhibitions include Art Lights up Life: People’s Power Station – Lighting Up Project, Guangzhou, China, 2016; Sun Dragon Hardware, Ace Gallery, Los Angeles, 2015; and Spectrum Divide, Saamlung Gallery, Hong Kong, China, 2012. Selected recent group exhibitions include Desert Island - Epicenter Projects, Coachella Valley Art Center Indio, California, 2017; Shenzhen Biennale of Contemporary Art, Shenzhen, China, 2017; BI-City Biennale of Urbanism\Architecture, Shenzhen, China, 2016; BOOSTER: Art Sound Machine, MARTA Herford Contemporary Art Museum, Herford, Germany, 2016; The Solutions, International Design Exhibition, Chengdu Biennale, Chengdu, China, 2011; and What if, Beijing International Design Triennial, China National Museum, Beijing, China, 2011. https://matthope.org

About Qi Hardware

Qi Hardware is a non-profit bringing together: Cryptocurrency Blockchain Software & Hardware Innovators, Open Source Software and Hardware Developers, and Enthusiasts of Cryptocurrency, China and Shenzhen. https://qihardware.org

About Fabricatorz Foundation

Fabricatorz Foundation supports creative technologists with community-focused initiatives that drive innovation through art and technology. The Foundation is made up of a team who have lead projects for Fortune 500 companies, A-list international artists, top museums, galleries, and educational institutions around the world. Notable sponsored projects include Bassel Khartabil Fellowship, New Palmyra, and Qi Hardware. http://fabricatorz.org

About Oui.Gallery

Oui.Gallery is an international contemporary art gallery featuring emerging artists and making innovative shows in Hong Kong and Saint Louis. https://oui.gallery

Oui.Gallery Hong Kong
1009 Yu Yuet Lai Building
43-55 Wyndham St
Central, Hong Kong

Tuesday - Sunday, 11 am - 7 pm

For more information

- [https://oui.gallery/entropic-bodies]
- Project Posters: JPG PDF SVG
- Project Graphic: JPG

E-waste Graphics: (Coming soon…)

Press Contact

Jon Phillips +1-314-484-8042 ai@qihardware.org

Rain Wong +852-6911-0860 press@oui.gallery

###

Feel free to share!

You may also enjoy...

Oui.Talk: GONG HEI Artist's Maryam Mobasseri & Bao Ho Artist Talk

恭喜

"GONG HEI" Group Show Launches Oui.Gallery Hong Kong in Central with Maryam Mobasseri, Bao Ho & Rachel Youn

EARTH OVERSHOOT DAY USA

MARCH 14, 2019, UNIVERSITY OF CALIFORNIA, IRVINE

JOIN THE CELEBRATION OF “EARTH OVERSHOOT DAY USA”

The Earth Overshoot Day is the calculated calendar date on which humanity’s resource consumption for the year exceeds Earth’s capacity to regenerate those resources that year. A country’s overshoot day is the date on which Earth Overshoot Day would fall if all of humanity consumed like the people in this country. This year “Earth Overshoot Day USA” falls on March 14th, which is right before a worldwide youth Climate Strike and the spring equinox, a critical day for Earth celebrations since time immemorial.

TOOLKIT FOR POST-GROWTH FUTURES

Contribute to a crowdsourced intellectual reinvention toolkit, a collection of radical strategies, concepts and stories for the stimulation of new perspectives, practices, value systems and socio-political alternatives to the ongoing environmental breakdown. This open-source initiative aims to collect and popularize concepts and narratives for awareness, preparedness and reskilling for a near future defined by fossil-fuel decline, decreasing

near future defined by fossil fuel decline, decreasing economic growth and climate change. It puts particular emphasis on collective intellectual preparation, regeneration and seeding from the complementary fields of art, critical theory, speculative fiction and indigenous knowledges.

LOCAL CONTRIBUTIONS
↳ March 14, 11AM-4PM, UC Irvine - [outdoor meeting in front of the Student Center]

REMOTE CONTRIBUTIONS
↳ Use #OvershootDayUSA or ping[at]earth-overshoot-day.org

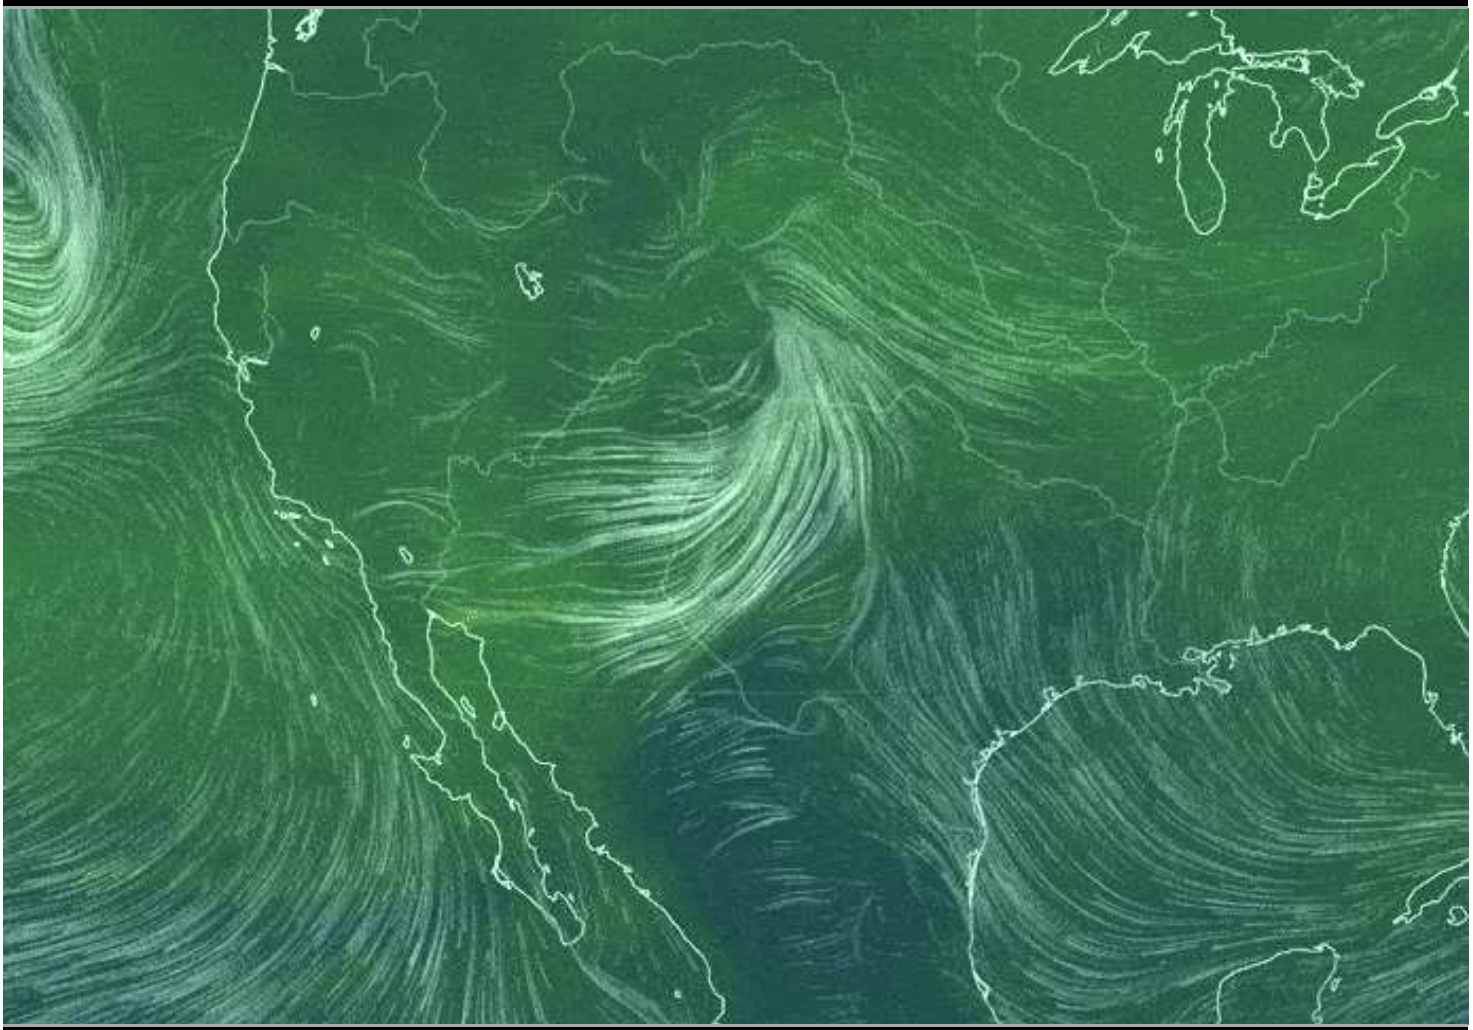

OUR PEOPLEEarth Overshoot Day USA, is a public event org
art collective Disnovation.org together with Geo
Bowker, Seanna Rose O’Leary and the Departm
Informatics, University of California, Irvine. Vide
documentation by Yubo Dong. Numerous, scho
activists and artists have informed this project, a
will progressively be made accessible on Earth-
Day.org

(pt. 2) ECIP-1049: Why Ethereum Classic should Adopt Keccak256 for its Proof of Work Algorithm.

Alexander
Tsankov

Mar 10 · 5 min read

Since publishing ECIP-1049 (here) as a response to the 51% attack I have received quite a bit of input and questions. I have also opened up a pull request to multi-geth to create Astor (here), an Ethereum testnet that will use Keccak256 as its proof of work method. If you agree with this article, or just want to learn more about how consensus works in Ethereum, I would appreciate any help. My goal is to have the testnet running by January 2020. Read first article in series (here)

Question (courtesy of Phyro): “Why does the amount of hashes matter to secure the network, shouldn’t it be energy expended?”

Answer: This is a great question, and really goes to the heart of the ECIP-1049 proposal, and why moving away from Ethash to Keccak256 will help secure Ethereum Classic. First of all, the expenditure of energy in a proof-of-work system is the most important factor for decentralization. Because energy is distributed all throughout the world, and energy is necessary to mine a proof of work coin, it follows that mining is distributed all throughout the world. This is especially true in Bitcoin mining, where miners set up in a variety of locations—the biggest being hydro-electric rich areas in the Northwestern United States and Southern China. Other popular areas to mine are particularly cool climate with cheap power like Russia, Canada, Norway, and Iceland.

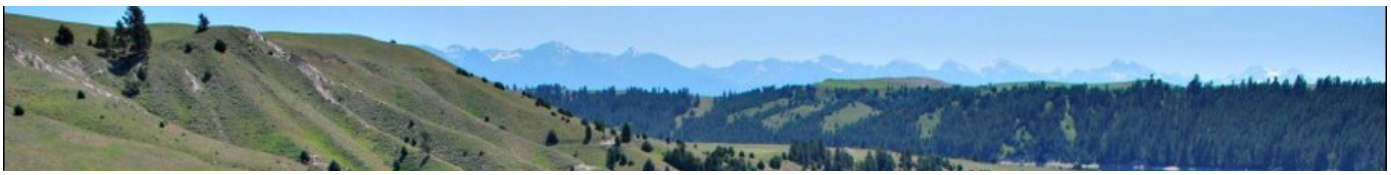

Valatie Falls, upstate New York. Most power from the dam pictured is used to mine Bitcoin.

It's fairly clear that if power is the limiting factor for mining, the system will be distributed globally. However, an important topic of debate is whether or not ASICs cause the system to be fragile because centralized chip makers can control most of the hashpower. This may have been true in the early days of Bitcoin, but it is false today. With a number of companies ranging from Innosilicon and Bitmain (China), to GMO (Japan), Bitfury (Europe and USA) all vying to create the most efficient miners possible. Keep in mind this is for an entirely new chip industry that didn't exist even 10 years ago. Compare this to the fact that the **only** major graphics card manufacturers are AMD and Nvidia (publicly traded companies both based in USA)

This competitive liquidity exists in Bitcoin, because Bitcoin uses SHA256, a very standardized and efficient algorithm that is easy to produce ASICs for. Because it is so easy to produce ASICs, there are many companies that can do this. It decentralizes manufacturing. This is good, **we want the deciding factor of mining profitability to be access to cheap power only.**

Miners measure the effectiveness of their miners using a simple metric: W/Gh (watts of energy / gigahashes per second)—this is the only number that matters to miners, like MPG (miles per gallon) matters with fuel efficiency. We can see the increase in efficiency starting from the first SHA256 miners in 2013 (source)

Butterfly Labs Jalapeno (from 2013) does 4 W/GH

Antminer S3 (from 2014) does 1.3 W/GH (3.1x better)

Antminer S7 (from 2015) does .25 W/GH (5.2x better)

Antminer S9 (from 2016) does 0.098 W/GH (2.5x better)

Innosilicon T2 Terminator (from 2018) does 0.083 W/GH (1.18x better)

Innosilicon T3 Terminator (from 2019) does 0.049 W/Gh (1.69x better)

The lower the W/Gh number, the better because it means you can do more hashes with expending less energy. As you can see, every new year devices get more efficient. While Bitcoin ASICs are so new (less than 7 years) we can already start to see the same pattern that emerged in integrated circuits—Moore's Law which says the number of transistors on chips doubles every 2 years. In crypto, we will see a "Nakamoto's Law"—Bitcoin miner efficiency doubles every two years. The fact we have maintained Moore's law since the 1970s is an amazing achievement of modern electrical engineering and economics.

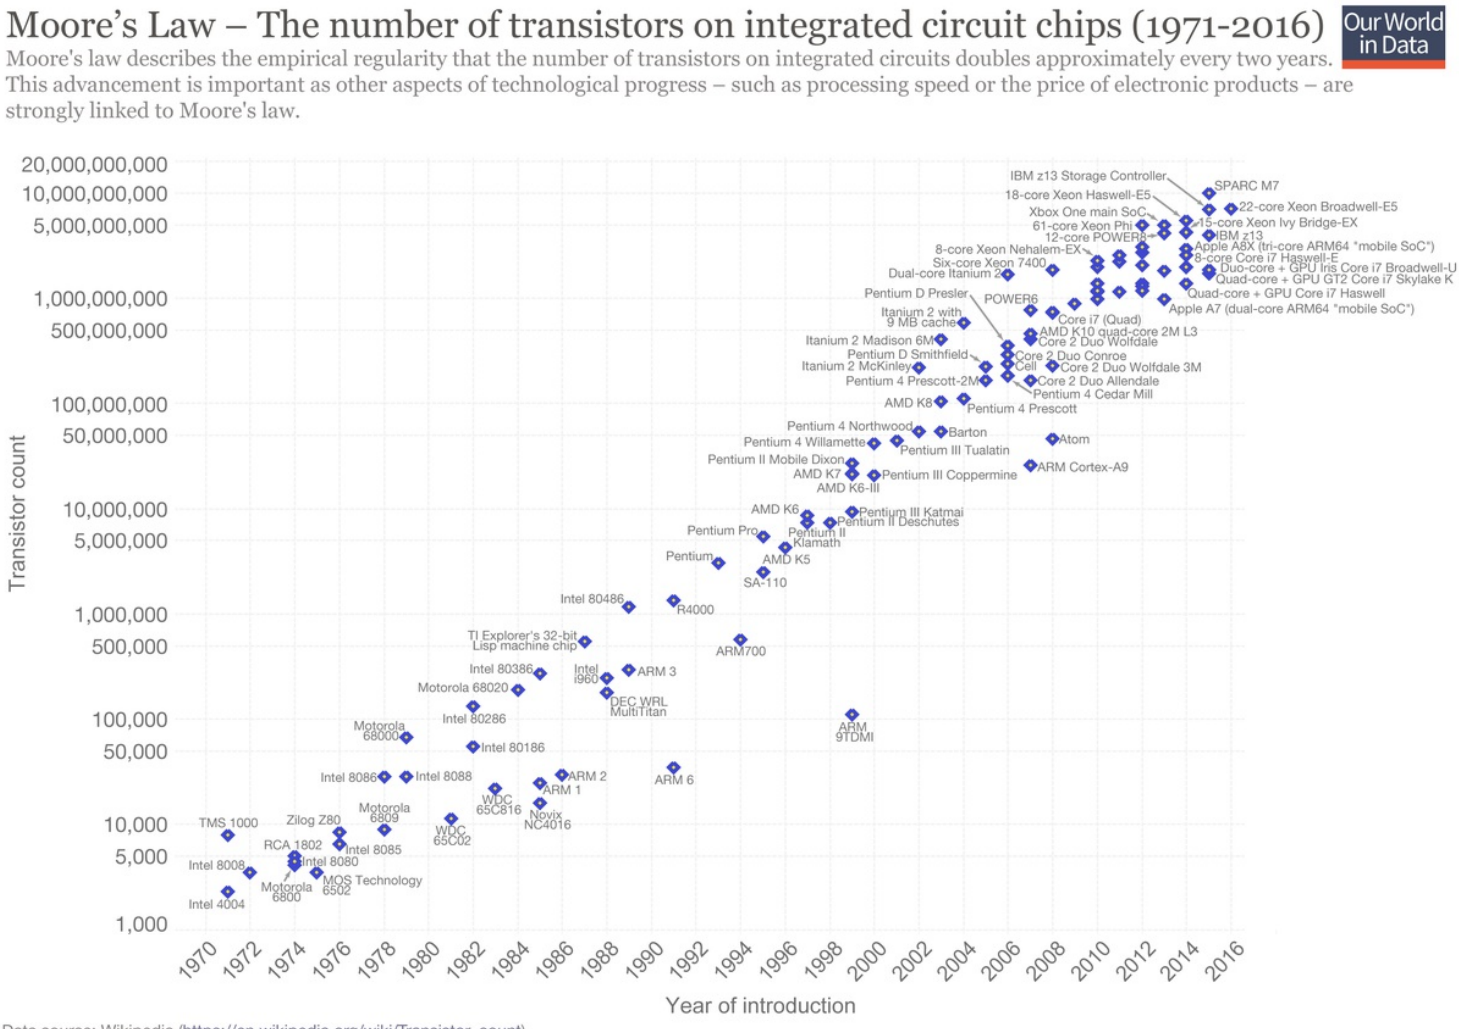

The issue with “ASIC-resistant” algorithms is that they introduce a new complication to the proof of work equation—which is access to ASIC miners. If one company or country successfully creates ASICs for an ASIC-resistant algorithm before anybody else can, it means they can control the network with very minimal power expenditure. In fact, a government could easily overpower a network with very minimal energy spent, if they can manufacture a few ASICs when nobody else has them. **I am worried about algorithms like Ethash and Progpow because if there is a flaw in them, it can make it easy for an attacker to silently control the network for years with very minimal power expenditure.** Keccak256, the algorithm I am proposing in ECIP-1049 on the other hand is very similar to Bitcoin’s SHA256, and amazingly, with exponentially higher performance. In fact, the only reason Satoshi didn’t use Keccak256 AKA SHA3 was that it was only SHA certified as safe in 2015, almost 6 years after the Bitcoin blockchain began, and around when Ethereum launched. Satoshi was a gentleman, and gentlemen know not to create a new hashing algorithm for a software project. Make no mistake: Keccak256, the product of decades of research, is state of the art crytpographic technology that we can leverage to protect our young network.

Fig. 6. Combined Throughput vs. Area graph for multiple hardware architectures of the 256-bit variants of BLAKE, Groestl, JH, Keccak, Skein, and SHA-2 implemented in Xilinx Virtex 5 FPGAs.

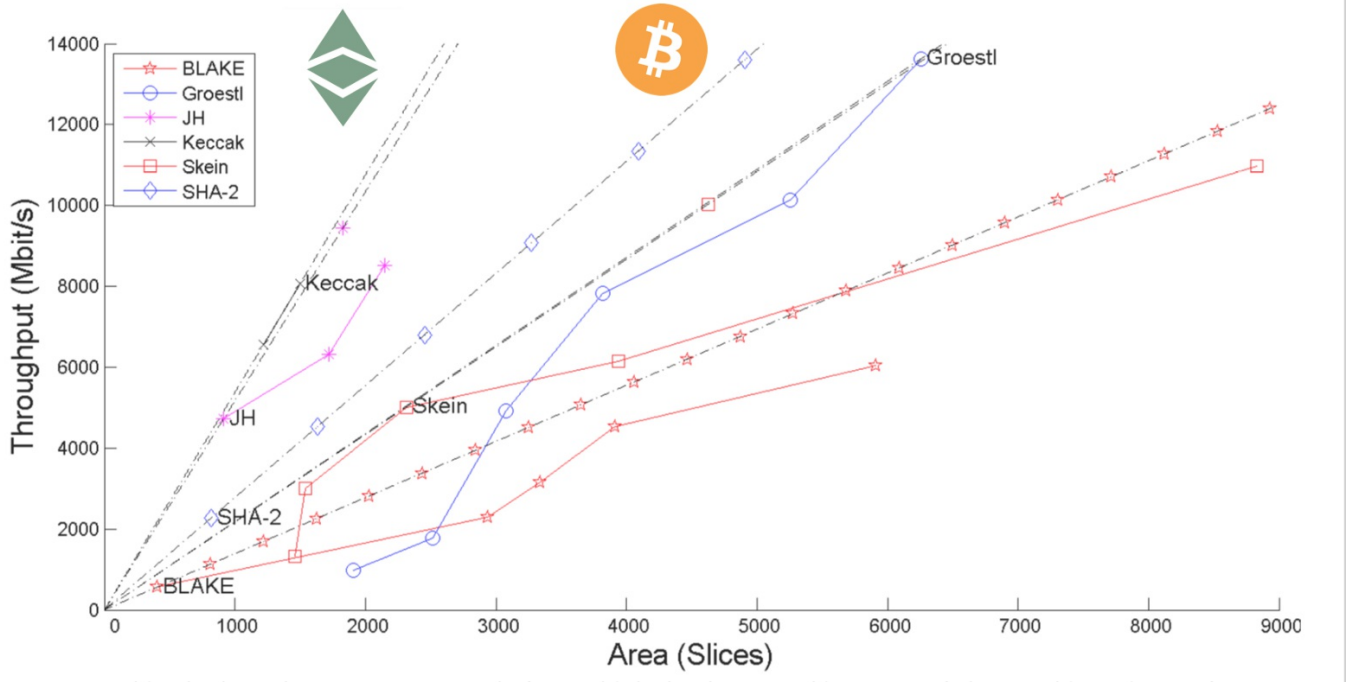

From a 2011 paper comparing FPGA implementations of hashing algorithms. Notice SHA-2 and Keccak

Now why does this steady increase in efficiency matter for mining? Well it's important because it means it's very difficult for a new ASIC manufacturer to jump ahead and make an ASIC that is 10x, or 100x, or 1000x faster than anybody else. Imagine how dangerous it would be if an attacker that can mine with 10,000x the efficiency of everybody else! With very few computers and little energy, the attacker can easily overpower the rest of the network. So these ASIC-resistant algorithms are so dangerous because—if they fail (see XMR, LTC, ZEC, ETH, etc.) and an ASIC is produced, it can be so much stronger than the rest of the network. This is why Keccak256 is such a good algorithm because it is very efficient and will have a highly liquid mining market. In the Keccak implementation document they show you the exact architecture to produce an ASIC yourself, and we know that the market for mining equipment will be able to follow this open standard and manufacturers can compete against each other to produce the best product. This means miners and pools have access to better mining equipment. A transparent market with clear standards is always better than an opaque market with asymmetric information. It's clear to me that what worked for Bitcoin, SHA256, will work for Ethereum Classic: Let's switch to Keccak256!

Sources:

Keccak implementation document: https://keccak.team/files/Keccak-implementation-3.2.pdf

Paper profiling different hashing algorithms on an FPGA: http://www.ecrypt.eu.org/hash2011/proceedings/hash2011_07.pdf

Bitcoin | Ethereum Classic | Proof Of Work | Ethereum | Blockchain Development

Alexander Tsankov
Pushing boundaries

Navigation

- Mixtape Torrents
- Promo Tracks
- Mixtape News
- Search
- Guide: How to Download
- Contact Us

User login

Username: *

Password: *

Log in

- Request new password

Browse By DJ

Ali Vegas
Beatsnblends
Beltway 8 Boyz
Biggy Jiggy
Big Mike
Big Stress
Black Bill Gates
Care Bears
Chinese Assassin
Clinton Sparks
Color Changin Click
Cutmaster C
Dirty Harry
Dirty Money Records
DJ 007
DJ 151
DJ 2Evil
DJ 2Mello
DJ 31 Degreez
DJ Absolut
DJ Arson
DJ Artillary
DJ Bape
DJ Bedz
DJ Benzi
DJ Berocke
DJ Barry Bee
DJ BeyondReset
DJ Big Tobacco
DJ Biz
DJ Black Jesus
DJ Block
DJ Bobby Black
DJ Break
DJ Burn One
DJ Butta
DJ Cali
DJ Capcom
DJ Capone
DJ Cease Fire
DJ Chuck T
DJ Cinema
DJ Clue
DJ Cobra
DJ Coolbreeze
DJ Crowd
DJ DBF
DJ Deals
DJ Decko
DJ Delz
DJ Diggz
DJ Dollar Bill
DJ Domination
DJ Drama
DJ Dutty Laundry
DJ E.Nyce
DJ E Stacks
DJ Egg Nice
DJ Envy
DJ Exclusive
DJ Explicit
DJ EZ Cutt

Lil Wayne - Young Moula Baby

Submitted by hardb0dy on Mon, 09/22/2008 - 6:14am. Dirty South | Mixtape Torrents

1. Lil Wayne Young Money Roster Intro
2. Lil Wayne, Drake, Kid Kid I Want Dis Forever
3. Drake, Lil Wayne, Kid Kid Stunt Hard
4. Lil Wayne, Kid Kid Bullshit
5. Lil Wayne, Kid Kid Gorilla
6. Kid Kid, Lil Wayne Real Shit
7. Lil Chuckee, Lil Wayne Too Clean
8. Lil Chuckee Ya Heard Me
9. Lil Twist, Lil Wayne Thatll Be Cool
10. Drake, Lil Wayne Ransom
11. Lil Wayne, Kid Kid Fuck A Nigga Thoughts
12. Tyga, Lil Wayne & Rich Boy Coconut Juice (Remix)
13. Lil Wayne, Mack Maine & T-Pain Got Money (Remix)
14. Jae Millz, Lil Wayne Holla At A Playa
15. Kid Kid, Mack Maine, Lil Wayne Talk To The Pillow
16. Bobby Valentino, Lil Wayne Smile
17. Lil Wayne, Kid Kid Dats My Nigga
18. Drake, Lil Wayne, Kid Kid I Can Take Your Girl
19. Drake, Lil Wayne Brand New
20. Lil Wayne Carter 4 Outro

Attachment	Size
youngmoulababy-MIXFIEND.torrent	19.71 KB

How To DOWNLOAD

» login to post comments

hot

Submitted by Jomel (not verified) on Mon, 09/22/2008 - 6:32am.

young moula baby

» login to post comments

hatin

Submitted by yizzo (not verified) on Mon, 09/22/2008 - 6:46am.

it still amazes how this dude got rich. He is fuckn wack as hell. He got good producers and a lucky ass break. Wish I could rob him

» login to post comments

Another bangin Ass Request

Submitted by SkoolNU on Mon, 09/22/2008 - 8:37am.

Granted big ups mixfiend/torrent give credit where its do. so lil wayne nice he kind like fell off a lil bit but his shit be ji going when hes not on that T Payne shit still one of the best

» login to post comments

Popular content

Today's:

- Wu-Tang Clan - The Sword Chamber
- Prodigy - Can't Take Back
- La Chat - Drama Queen (Who Want Smoke)
- Lil Gotit - It's Crazy But It's True
- Chief Keef & Zaytoven - Glotoven
- DJ Scream Presents Lotto Savage - Georgia Lottery
- Benny Foxmore - More Drip 7
- Crooked I - The Weeklys Vol. 1
- CMG Presents 42 Dugg - Young And Turnt
- Lil Flip - La Clover Nostra: Clover Gang
- Marty Baller - International Baller 2
- DJ Triple Exe - Street Kings 73
- Guide: How To Download
- DJ Ben Frank - Trap Excellence 5
- Vado - V-Day 3

Last viewed:

- DJ Kenny Mac - NBA 2K20
- Source Founders Dave Mays & Ray "Benzino" TERMINATED From The Source
- YFN Kay - Death Date

Browse By Mixtape Genre

Acapellas
Blends
Chopped & Screwed
Classic/Old
Dance Hall
Dirty South
East Coast
General
Instrumentals
International
Reggae
Remixes
RnB
Unsigned Hype
West Coast

DJ Fade
DJ Famous
DJ Fatal
DJ Felli Fel
DJ Finesse
DJ Flipcyide
DJ Gera
DJ Got Now
DJ G-Spot
DJ Green Lantern
DJ Greg Nasty
DJ Grip
DJ Gutta
DJ Haze
DJ Heat
DJ Hitz
DJ Hpnotiq
DJ Hype
DJ Ideal
DJ Ill Will
DJ Jaycee
DJ J-Boogie
DJ J. Period
DJ Jamad
DJ Jelly
DJ Joey Fingaz
DJ JS-1
DJ Kay Slay
DJ Kep
DJ Keyz
DJ Khaled
DJ Killa K
DJ King Assassin
DJ Kool Kid
DJ Kronik
DJ Kurupt
DJ L
DJ L-Gee
DJ LA Luv
DJ Laz
DJ Lennox
DJ Lex
DJ Love Dinero
DJ LRM
DJ Lu
DJ Lust
DJ LV
DJ Maku
DJ Mark S
DJ Mello
DJ Mystery
DJ Neil Armstrong
DJ Nice
DJ Nik Bean
DJ Noize
DJ Noodles
DJ Obscene
DJ On Point
DJ Papa Smurf
DJ P Exclusivez
DJ Pimp
DJ Premier
DJ Purfiya
DJ Quess
DJ Quote
DJ Radio
DJ Rah2K
DJ Ramadan
DJ RC
DJ Red Devil
DJ Reeg
DJ Reflex
DJ Rell
DJ Rhude
DJ Rondevu
DJ RPM
DJ Rukiz
DJ Scarface
DJ Scoob Doo
DJ Scope
DJ Scream
DJ Screw
DJ Semi
DJ Set It Off
DJ Seye Daddy
DJ Shef
DJ Skee
DJ Smallz
DJ Smarts
DJ Smooth Denali
DJ Spinatik
DJ Spinz
DJ Spyda
DJ Stashman

yizzo

Submitted by Guest95 (not verified) on Mon, 09/22/2008 - 1:20pm.

hahah hatin ass faggot..im not even really a weezy fan..just a fan of success and you're a haterrr haha you faggot

» **login** to post comments

its almost like people go

Submitted by imrickyduh on Mon, 09/22/2008 - 1:58pm.

its almost like people go out of their way to say how bad lil wayne is lol. like that nigga care what u think, and like i care either. if u hatin the game take the walk of shame cuz we'll never see ur picture on the wall of fame.

» **login** to post comments

Weezy Baaaaby

Submitted by Dezil (not verified) on Mon, 09/22/2008 - 5:55pm.

This nigga fire, his lyrics are lyrical homicide, he murders the track! Big ups Wayne!

» **login** to post comments

man..

Submitted by ddddddddd (not verified) on Mon, 09/22/2008 - 6:05pm.

all this shit b old... i'm ready for some new wayne.. come on dude...

» **login** to post comments

its funny how everyone

Submitted by kjlkj (not verified) on Tue, 09/23/2008 - 4:21am.

its funny how everyone thinks its wayne puttin out these mixtapes when its djs just doing different beats to his songs....yall are stupid as fuck for thinkin wayne puts all these out

» **login** to post comments

I promise, if you a hater,

Submitted by Hardbody7 (not verified) on Tue, 09/23/2008 - 2:36pm.

I promise, if you a hater, you have a serious mental disorder and I suggest you get some counseling. I lot of times you guys aren't even aware of the fact that you're hating. I mean seriously, the way YOU feel about YOUR favorite artists is exactly the way WE feel about Wayne since we're Wayne fans and you like whoever you like. But the difference between us and you is we don't hate on your favorite artists. Real individuals respect your choice. Bitches and busters think otherwise. That being said, this is a nice mixtape. Kid Kid is a muthafuckin monster. I'm glad he came back.

» **login** to post comments

I dunno but every since the

Submitted by polo_229 (not verified) on Tue, 09/23/2008 - 6:33pm.

I dunno but every since the Dedication 2 came out all of Waynes songs sound th same to me. I guess he found his bread and butter in what most people will accept as a dscent sone which isn't very much. Don't get me wrong because I use to be a Wayne fan to the heart but lately I'm like whatever about his music. If you're a fan of his thats all good..to each his own. The thing that erks me is just because someone isn't a fan of Wayne or not really feeling a song of his people wanna call them a hater. The last time I checked people have a right to their own opinion. I'd rather have a mind of my own listen to a wider variety of music than just jump on somebodies bang wagon just to fit in with everybody else. Does Wayne has some hot songs? Yeah he does, but lets not be in denial and say that he hasn't mad some wack shit. I mean waht artist hasn't. And for you to say that Wayne fans don't hate on other artist is straight up bullshit. I could go to any of these mixtapes on this site and find a Weezy fan hating on an artists mixtape that Lil Wayne isn't even on.

» **login** to post comments

The realest non-lil wayne

Submitted by Guest502 (not verified) on Wed, 09/24/2008 - 2:08am.

The realest non-lil wayne fan i've heard on the site ever..... dig, brah

» **login** to post comments

My comment applies to all

Submitted by Hardbody7 (not verified) on Wed, 09/24/2008 - 10:22am.

My comment applies to all haters. Even Wayne lovers who HATE on others. And the reason the Dedications and Droughts went so hard is because Wayne was superhungery. We probably won't see him do anything else like that again. I been listening to dude since the first Big Tymers album around 1998. It's time for him to

DJ Steelz
DJ Storm
DJ Strong
DJ Suss One & DJ L
DJ Teknikz Mixtapes
DJ Thoro
DJ Trigga
DJ Unexpected
DJ Vlad
DJ Warrior
DJ White Owl
DJ Whoo Kid
DJ WizKid
DJ Woogie
DJ Wreck
DNA
Don Bishop Agallah
Don Cannon
Dow Jones
Dub Floyd
Evil Empire
Exclusive J
Father Bentley
Focus Entertainment
Funkmaster Flex
Grind City Mob
Hevehitta
Hustle Squad DJs
Idol
Iron Fist Records
J.A.
Jay Classik
J. Armz
Jail Break Recordz
J-Love
J. Period
King Smij
Kochece
La Profecy
Mark Ronson
Mastertapes
MC Assault
Mick Boogie
MIDIMarc
Mixtape Assassin
MVP Sound Crew
Nu Jerzey Devil
NY C.E.O.
O.G.B.
OG Ron C
P Cutta
Professor Green
Rapid Ric
Rob E. Rob
Screwed Up Click
Screwston
Selector Rondon
September 7th
Shiest Bubz
Sinister Shan
Soundwave
Southern Style DJs
Statik Selektah
Street Pharmacy
Suge White
Supa Mario
Superstar Jay
Swishahouse
Tapemasters Inc
Team Invasion
Team Invasion Midwest
Terry Urban
The Aphilliates
The Architect
The Cartel
The Empire
The Firemen DJs
The Grit Boys
The Hitmen
The League Crew
Tony Touch
Trap-A-Holics
Ty Boogie
Wally Sparks
Wit-E Beats
World Wide Legacy

Syndicate

more

rap/sing about something other than crack, bling bling, and violence. I download a lot of music and I just don't see this much hate thrown at anybody else. But then if it wasn't for haters, Wayne probably wouldn't be where he's at right now. Anybody notice dude ain't re-reciting on this one? I guess he found something else to prove. And that's why I'm a fan. Regardless of whether you all think a song is garbage or not, at least he's willing to put it out. All of the greats of the past have something in common...they put out an enormous amount of music and not all of it was great either (2pac,stevie wonder, james brown, mary j blige). I guess it's the law of average kicking in. Bottom line.

» **login** to post comments

You're not a stupid hater.

Submitted by Hardbody7 (not verified) on Wed, 09/24/2008 - 10:33am.

You're not a stupid hater. I'm talking about them stupid haters. They don't really hate Wayne. They probably just hate the fact that their favorite rapper ain't in Wayne's spot. Or maybe they were picked on by the popular kids at school and decided to root for the underdog. I don't know, but to HATE on ANYBODY is just stupid. You ain't gotta like em. You ain't gotta listen to em. And these real haters probably got more Wayne on their computers than me. How else would they know if he's garbage or if he's falling off?

» **login** to post comments

lil chuckie is the shit ta

Submitted by young juve aka mr freestlye (not verified) on Wed, 09/24/2008 - 7:56pm.

lil chuckie is the shit ta be 12...he betta that most rappers in there 30s lol

» **login** to post comments

2 bad lil wayne writes his

Submitted by Lol (not verified) on Thu, 09/25/2008 - 3:06am.

2 bad lil wayne writes his lines

» **login** to post comments

SIMILAR MIXTAPES:

- **Lil Wayne- Young Moula Baby (Chopd n Slowd by DJ Rico)**
- **Young Money Ent Presents Lil Wayne & Drake - Young Money Year 2K9 Vol 1**
- **Young Money All-Stars**
- **DJ Spinatik & Lil Wayne - Got What You Need 3**
- **The Cartel Presents Lil Wayne And Drake - The Carter Meets The Cartel 2**
- **DJ Hitz & Young Money Ent - Young Money Massacre**
- **DJ 31 Degreez - Young Money Family 2**
- **DJ Fletch - Young Money Millionaires**
- **DJ Hitz And Young Money Ent - Young Money Massacre 2**
- **Lil Wayne - Return Of The Leak**
- **Lil Wayne - Public Enemy**
- **Lil Wayne - The Leak 4**

Affiliates

MixFiend
Hip-Hop Posters
Hip Hop Ave
GNX Music
Nah Right
Balls & My Word
Got Beats Online
RockNRollIsDead
Queens Bridge
IllRoots
Thug Online - Hip Hop News
HH Live News
HoodFever.com
The Hype Factor
Shadowville
TheHoodUp
imHipHop
MusicVideoCast
Tastemaker Magazine
Rap Beats
Escape Entertainment
Pardon My Fresh
Green Hitz
UrbanFreshNation
Drop-Bomb
Ususpects
HipHopGiant
BFochs Beats Blog
Jordan Release Dates
Air Force Ones
Urban Music Blog
BestOfBothOffices
Air Jordans
CentralStationLovesYou
IndustryFinest
CredibleMusicReviews
DaStreetBuzz
ThatFireBF
New Era Hats
MadRapVideos
SouthernStyleDJs
Free Acapella Downloads
DopeHood
RapVet
HHHead
Jordans
XclusivesZone
Down-South
WeLiveThis
2Fresh2Cool
urbanmusicdaily
DJ Smallz
FreshlyServedHipHop
NewBloodMusic
ForbezDVD
TruBeats
SoulfullVibes
KarenCivil
RespectOrganization
SamHoody
iDJBlast
DJ 5150's Official Website
HitHipHop
2024 Productions
WeeklyDrop
LilWayneHQ
MiamiStreetWear
Rapgra Forum
NewLilWayne
DefPenRadio
DirtyGloveBastard
MuzikFene
thadoubleo
PiffAvenue
TrapsnTrunks
DopeTV

Hit Us Up To Become An Affiliate

This repo is a collection of vm statistics, gathered from a geth instance.

Every data point represents `10k` blocks. Every time an opcode executed, a timer was stopped, the time since start was noted, and the new opcode timer was started. After `10K` blocks, the data was dumped into a `json` file. The files are available in the `/m5.2xlarge/`-folder.

The benchmarks are from a `m5.2xlarge` aws instance that did a full-sync. Such a machine has

- `8 vcpu`,
- `32 Gb RAM`
- `300 GB NVMe SSD`.

Time spent

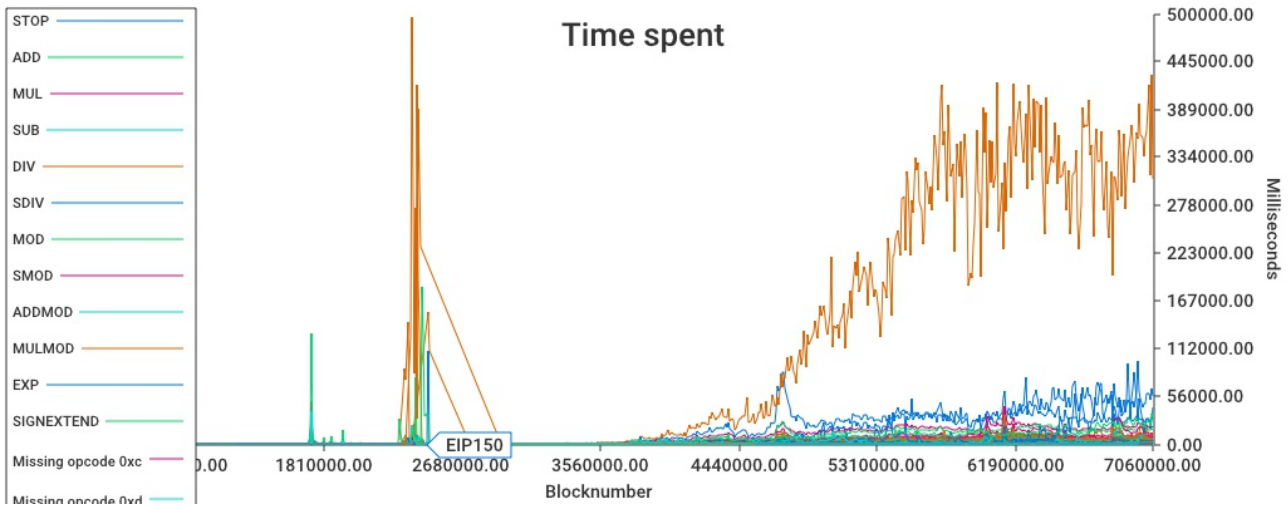

The large thing there is `SLOAD`.

If we place a cap on it, and filter so we only see the top four things that take up the execution time, we can see that the top ones are

- `SLOAD`, obviously,
- `SSTORE` - hardly surprising. Even though `SSTORE` is expensive, it's also one that touches disk.
- `EXP`
- `JUMPI` - This is a *very* common opcode. Also, Geth nowadays defer the jump analysis until it's actually needed - the first time a `JUMP` / `JUMPI` is called.

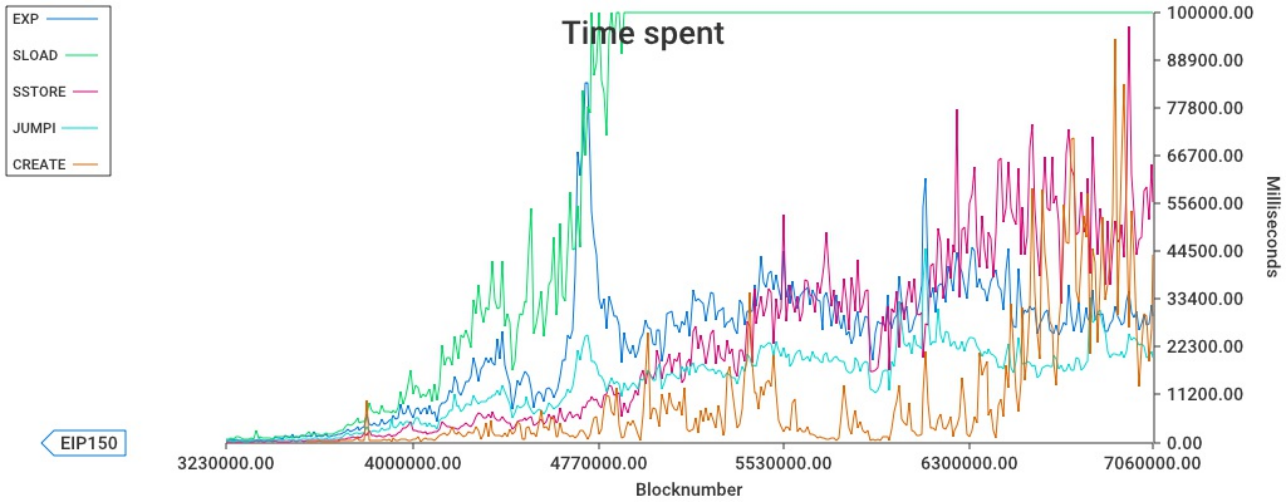

It seems that aside from `SLOAD`, there's no other single operation that dominates the remaining time that is spent on block processing.

Time Pies!

Let's see how the time spent has progressed through time.

Blocks 0 to 1000000 - Time spent

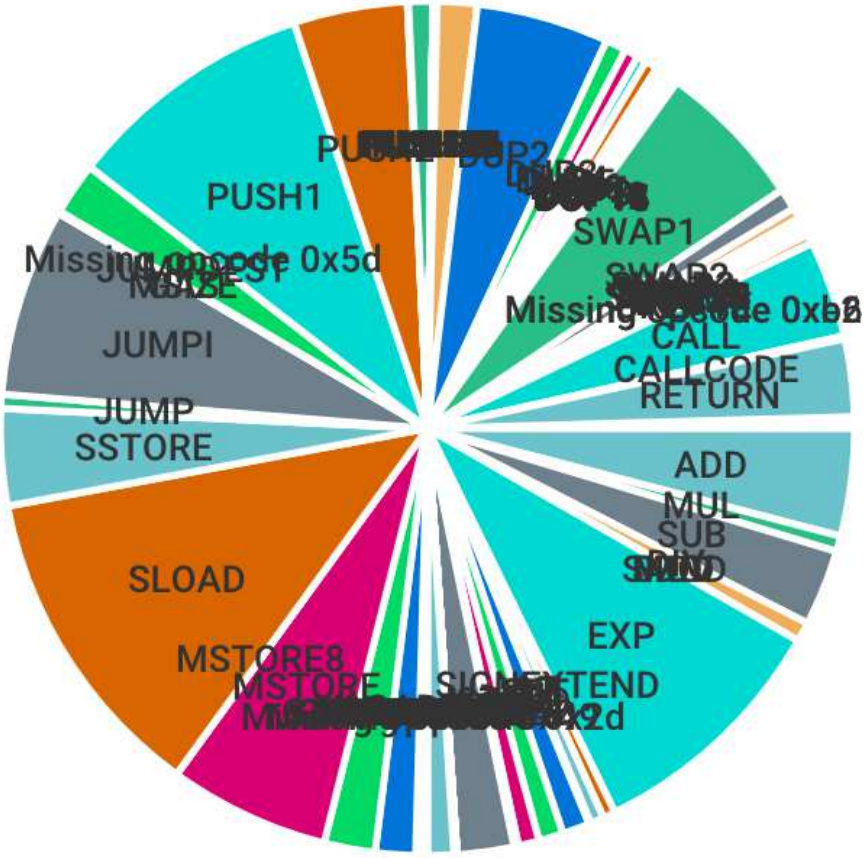

Blocks 1000000 to 2000000 - Time spent

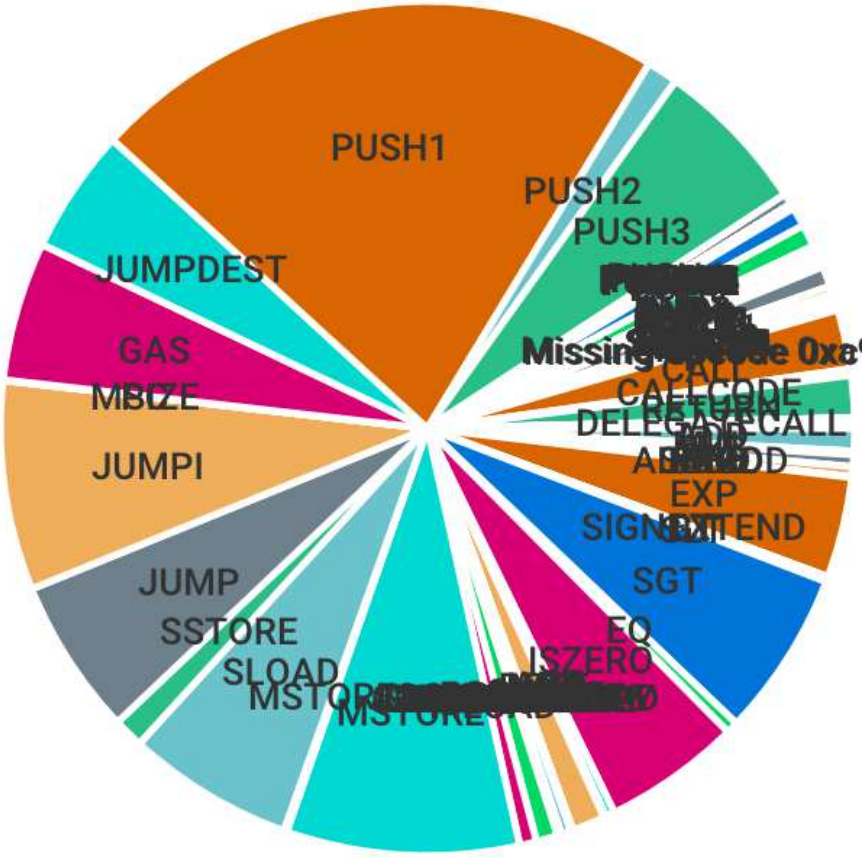

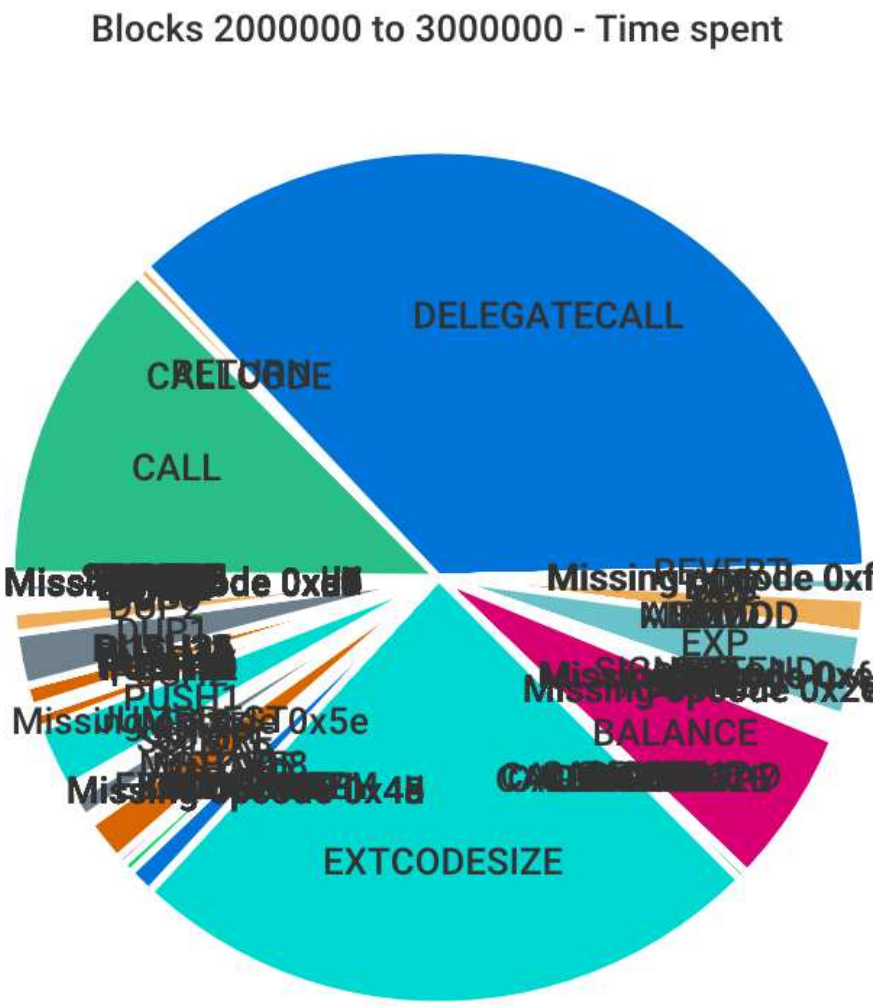

The `2-3M` range contains the shanghai attacks.

Blocks 3000000 to 4000000 - Time spent

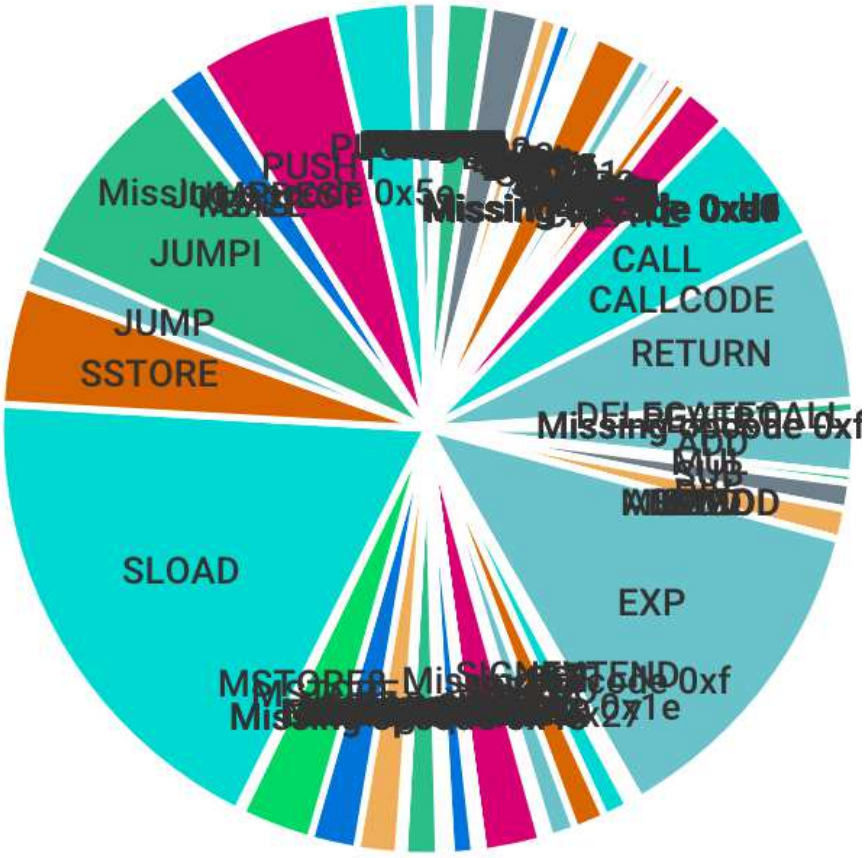

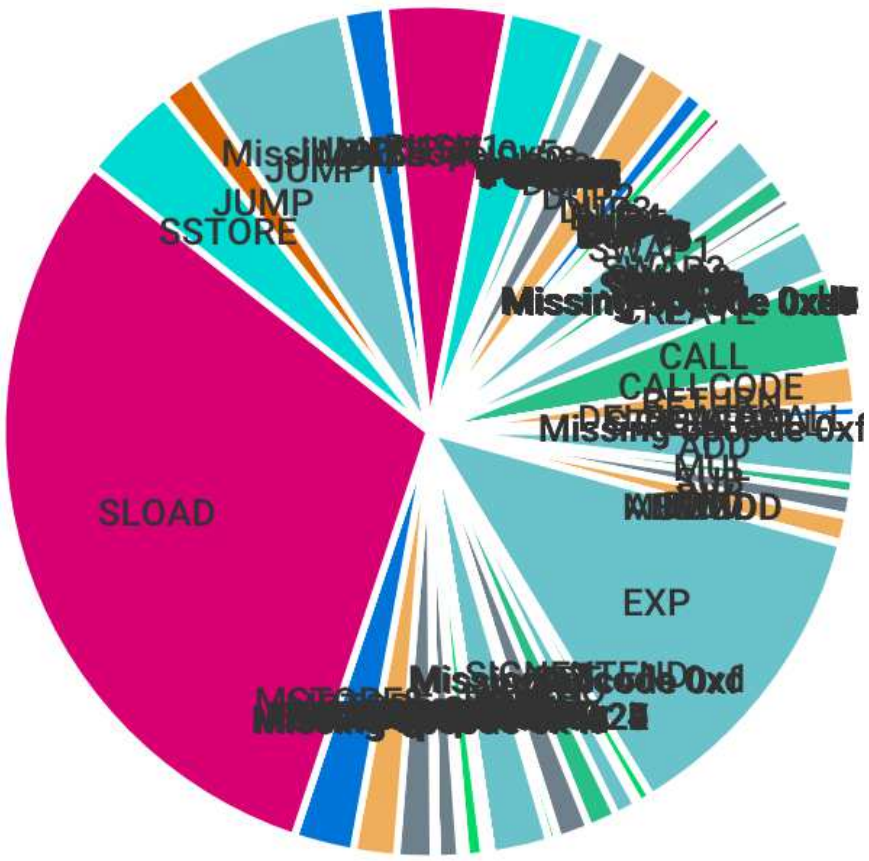

In the `3M-4M` range, we start seeing `SLOAD` take a larger piece of the pie.

Blocks 5000000 to 6000000 - Time spent

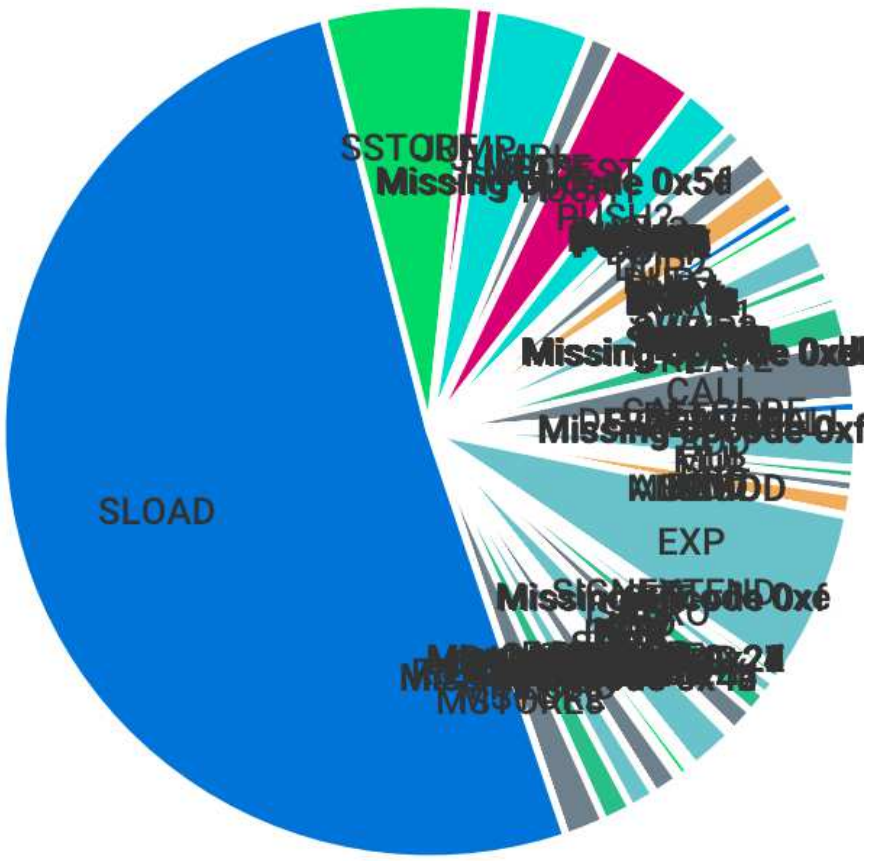

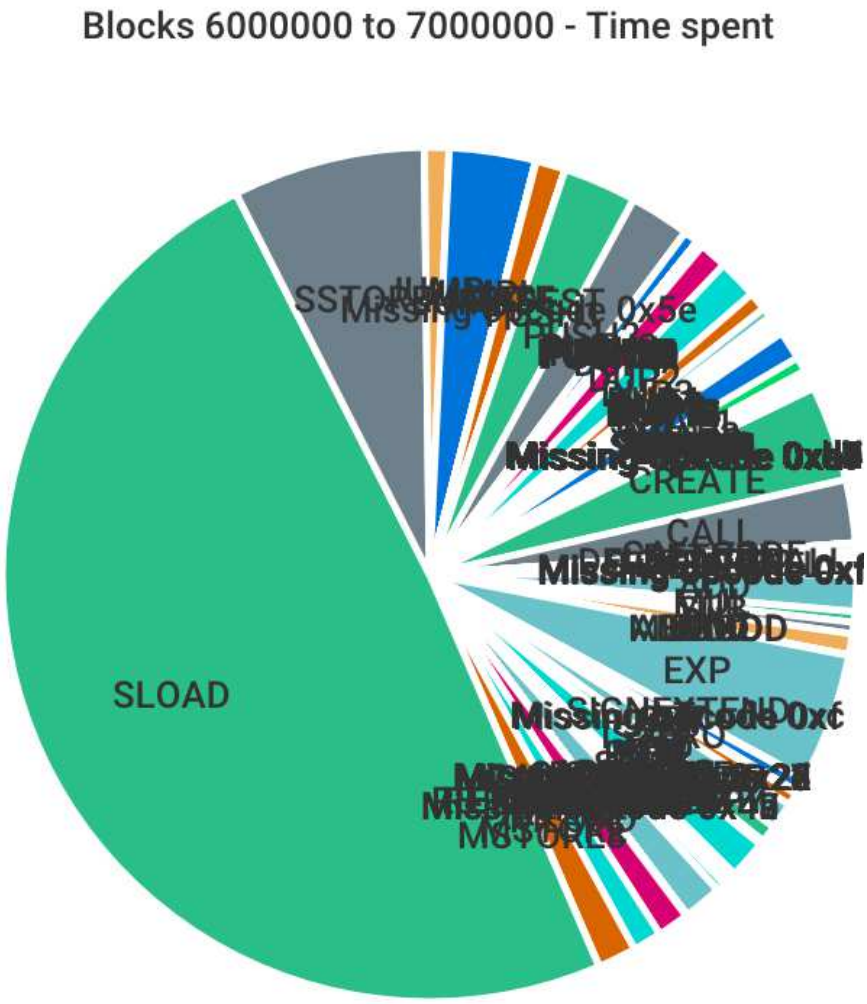

... and accounting for ~`50%` between `5M` and `6M`.

Count Pies!

We can also look at the*prevalence* of opcodes -- that is, how*common* is an opcode, and how has that varied over time?

Blocks 0 to 1000000 - Total count

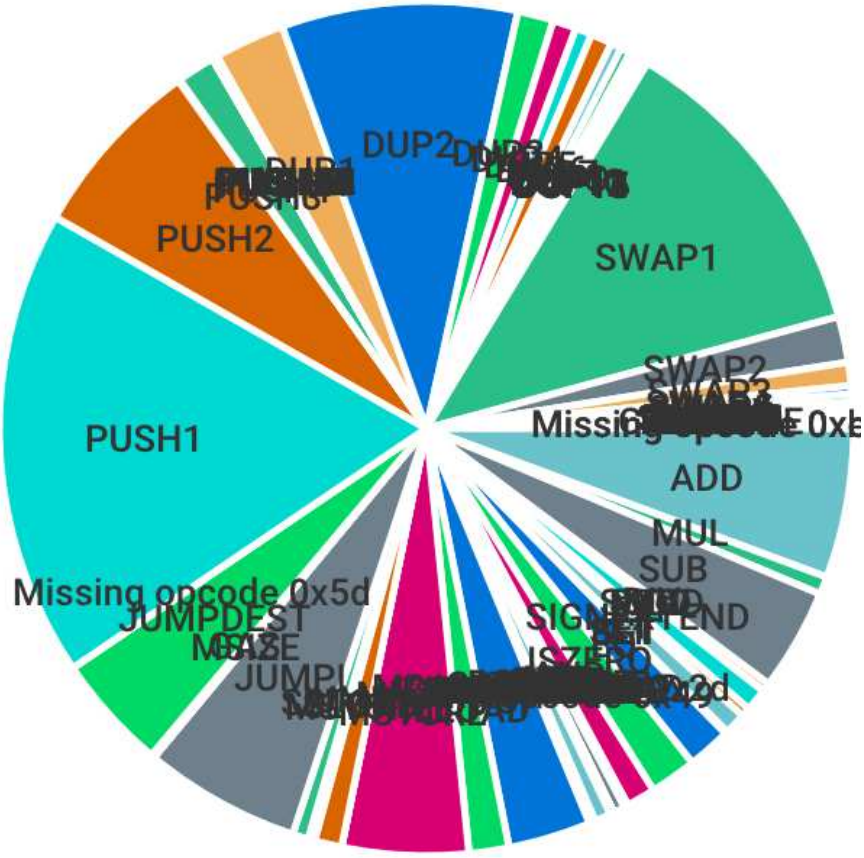

Blocks 1000000 to 2000000 - Total count

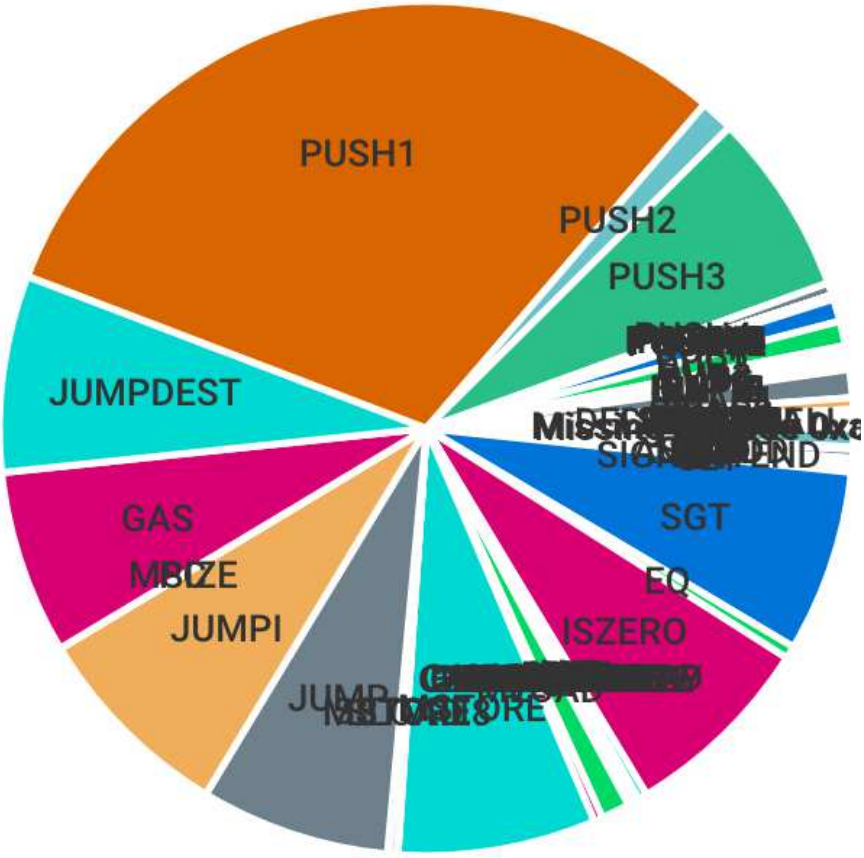

Blocks 2000000 to 3000000 - Total count

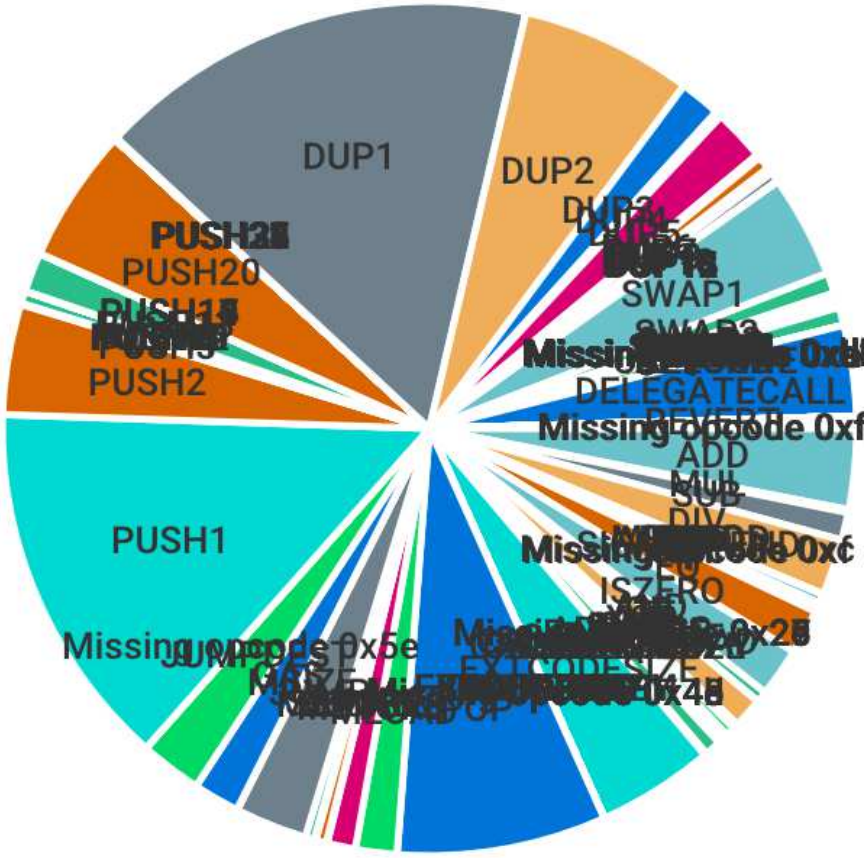

Blocks 3000000 to 4000000 - Total count

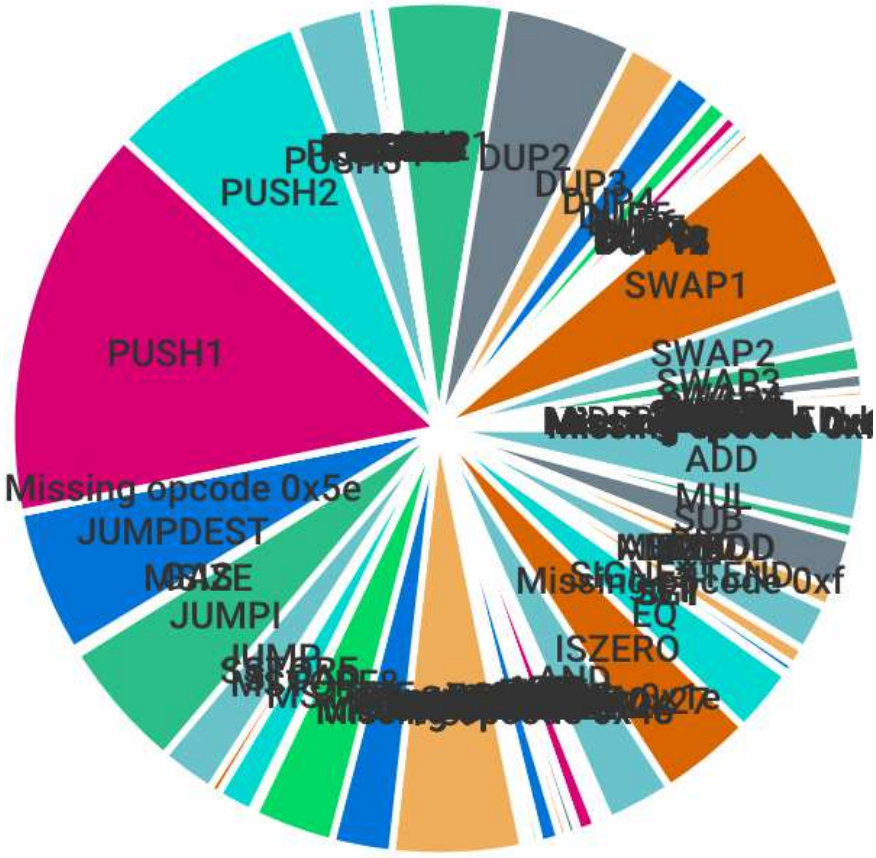

Blocks 4000000 to 5000000 - Total count

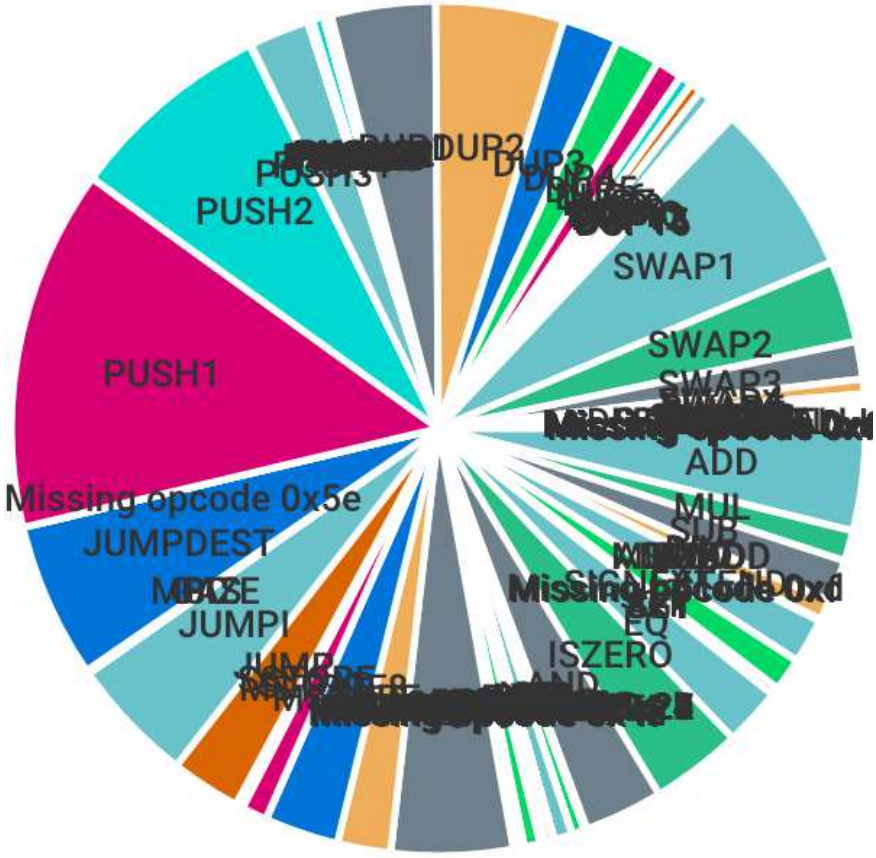

Blocks 5000000 to 6000000 - Total count

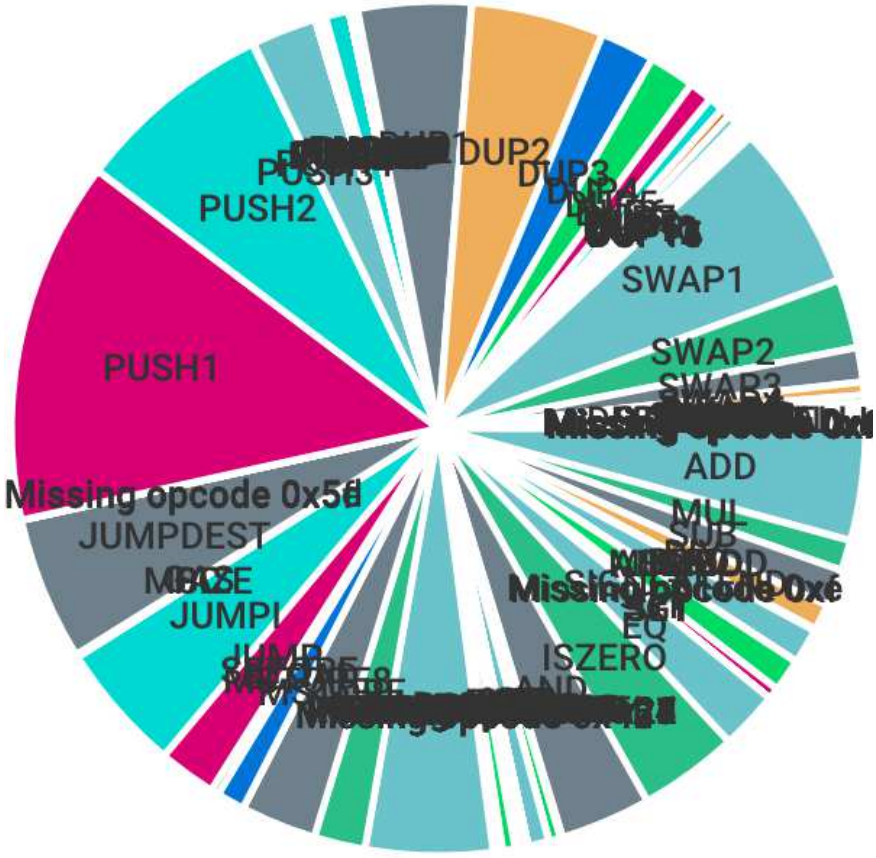

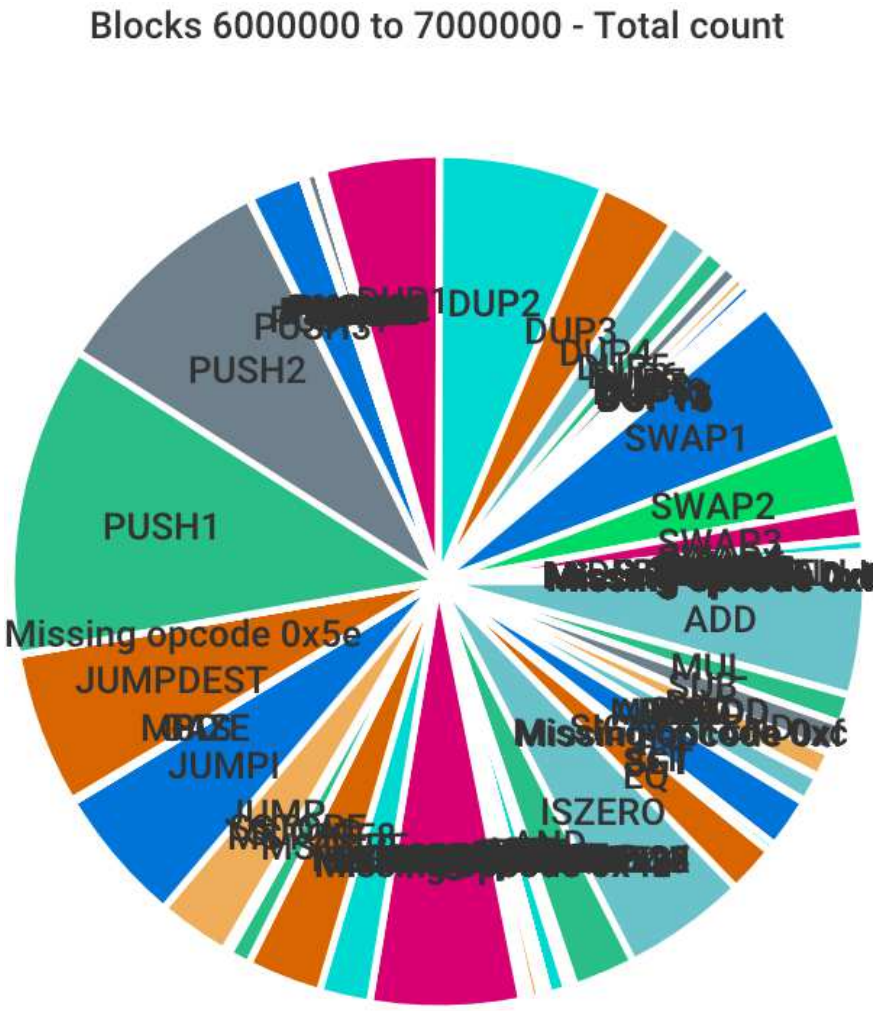

Cost of ops

Are operations well-balanced, gas-wise?

Arithmetic ops (`0x00` -range)

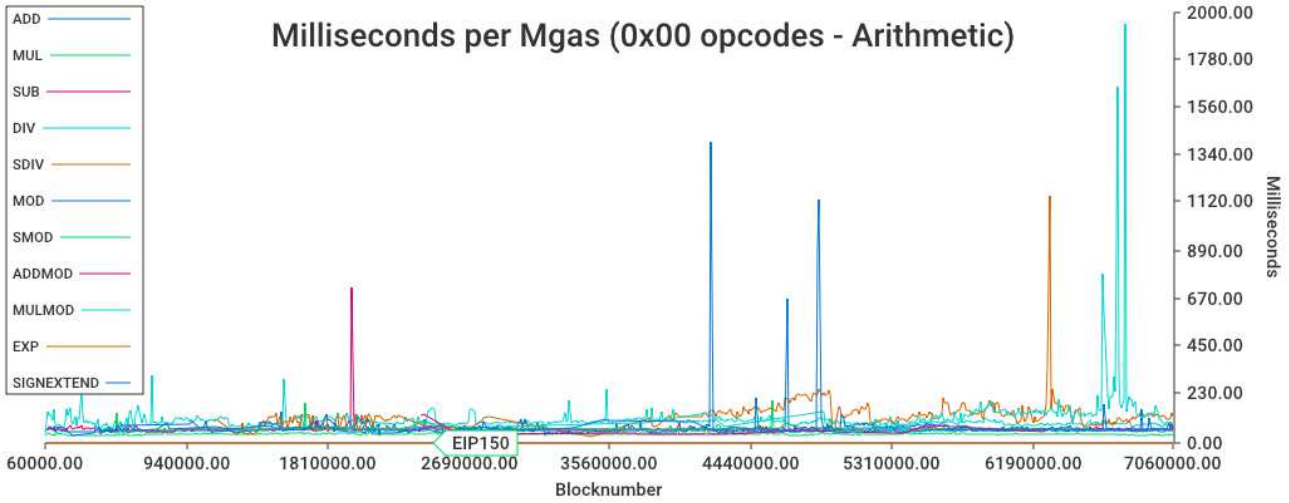

The arithmetic ops seems to vary within one order of magnitude -- some peaks which is likely due to noisy data. Here's a graph where it's been capped at `250`

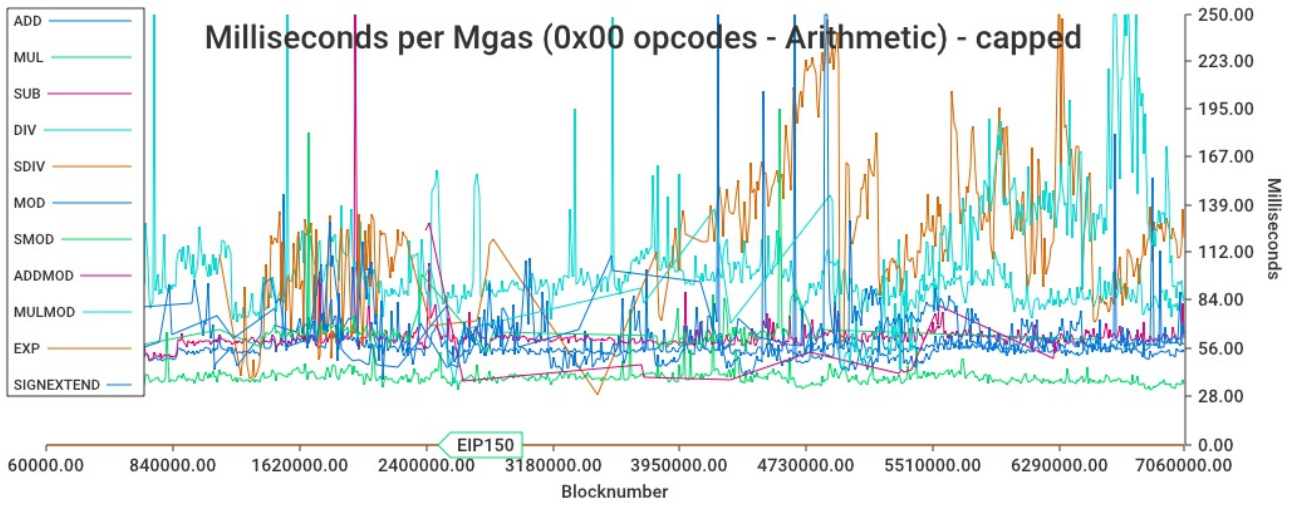

Comparison ops (`0x10` -range)

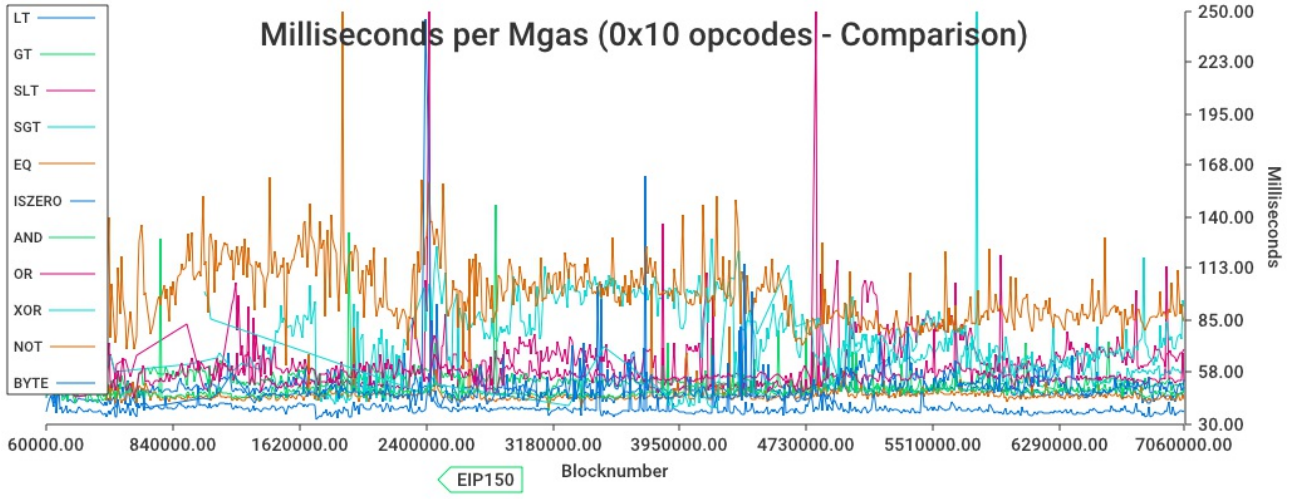

The comparisons ops (also capped at `250`) are fairly aligned within `35` and `130`. The `NOT` looks pretty over-represented.

SHA3 ops (`0x20` -range)

The `SHA3` operation has it's own range. Since it'a dynamically priced operation, the gas has not been part of the data collection. What we can do, however, is plot the time spent during `SHA3`, and the number of invocations.

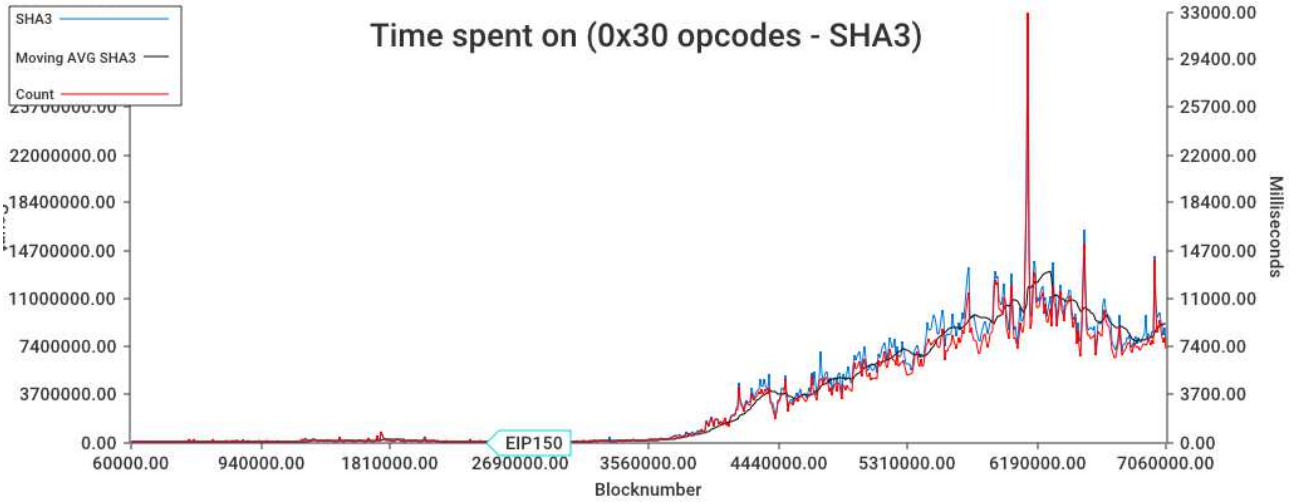

The two line up pretty well, so the op doesn't 'degrade' over time. As expected.

Context ops (`0x30` -range)

The following ops have dynamic gas, and are not charted:

```
CALLDATALOAD,
CALLDATACOPY,
CODECOPY,
EXTCODECOPY,
RETURNDATACOPY,
```

Here are the remainders (split up into two charts, both capped at `500`):

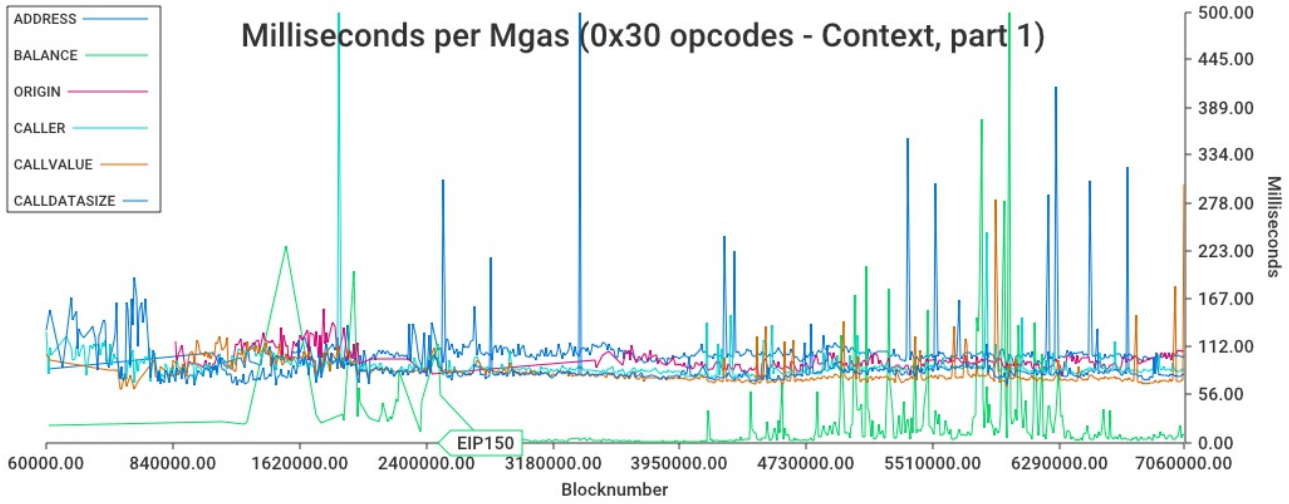

The big spike around `2.08M` is, surprisingly, `CALLER`. Aside from a few spikes, they're fairly well aligned -- except for `BALANCE`, which is starting to fluctuate and reaching high peaks. It was repriced in `EIP150`, anad was low for a few million blocks.

Note: `BALANCE` should really be called `EXTBALANCE`, since it fetches balance for (potentially) external accounts.

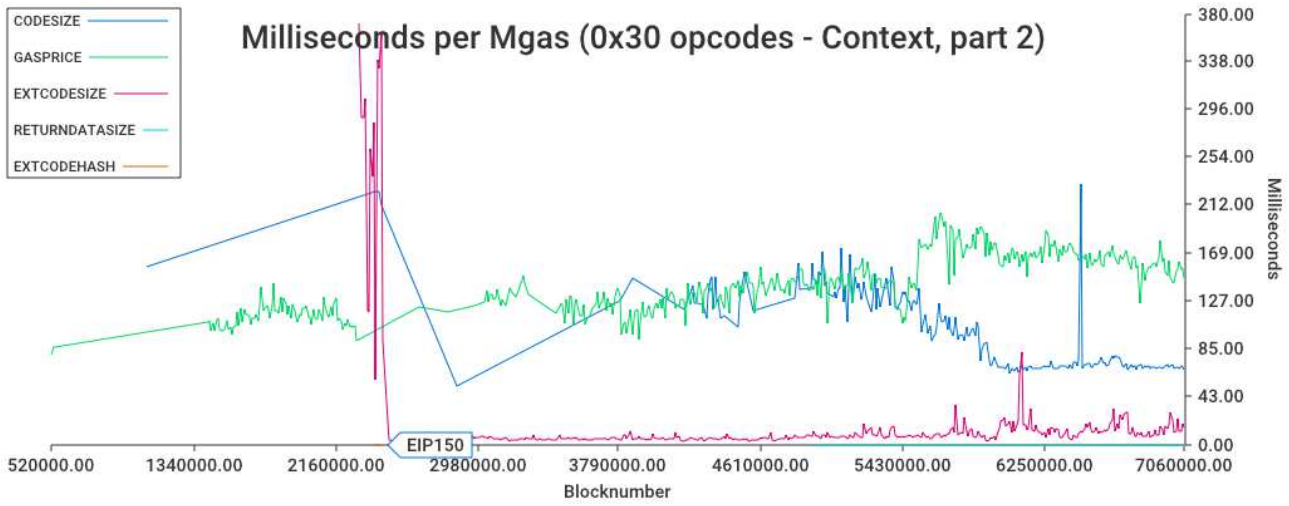

The `EXTCODESIZE` looks *extremely* cheap -- which makes sense, since it was repriced at block `2463000` from `20` to `700`.

The Tangerine Whistle HF also contained these changes:

- Increase the gas cost of `EXTCODESIZE` to `700` (from `20`).
- Increase the base gas cost of `EXTCODECOPY` to `700` (from `20`).
- Increase the gas cost of `BALANCE` to `400` (from `20`).
- Increase the gas cost of `SLOAD` to `200` (from `50`).
- Increase the gas cost of `CALL`, `DELEGATECALL`, `CALLCODE` to `700` (from `40`).
- Increase the gas cost of `SELFDESTRUCT` to 5000 (from `0`).

Block operations (`0x40` range)

This chart capped at `600`.

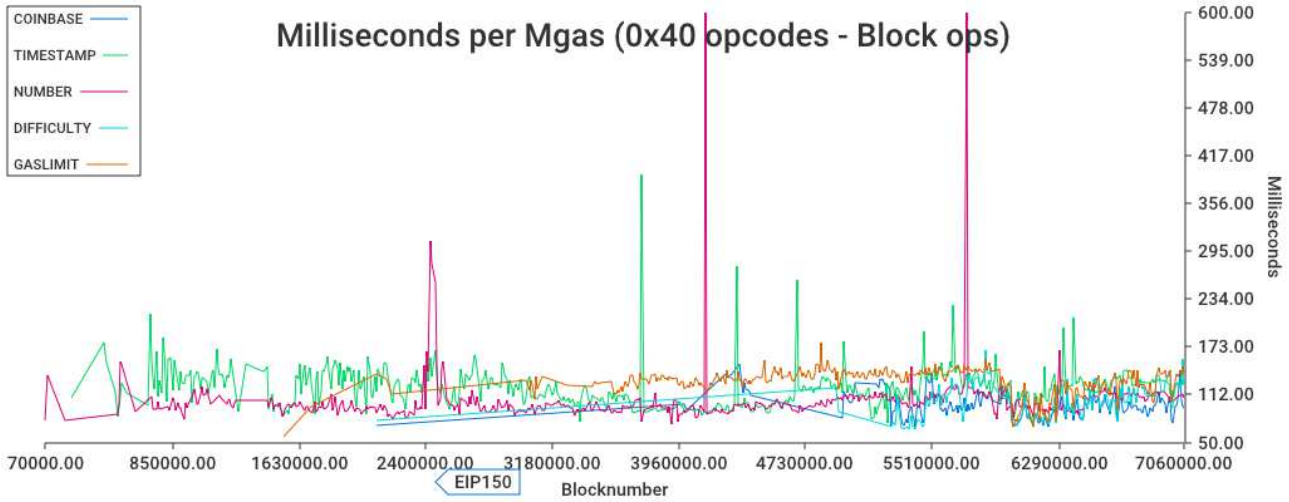

Missing from this chart is `BLOCKHASH`, which deserves it's own chart. It's plotted along with it's opcode count -- this operation is a bit quirky, and is more expensive when it's executed in isolation. So the more this op is used, the cheaper the operations become (to a degree). This chart capped at `3000`.

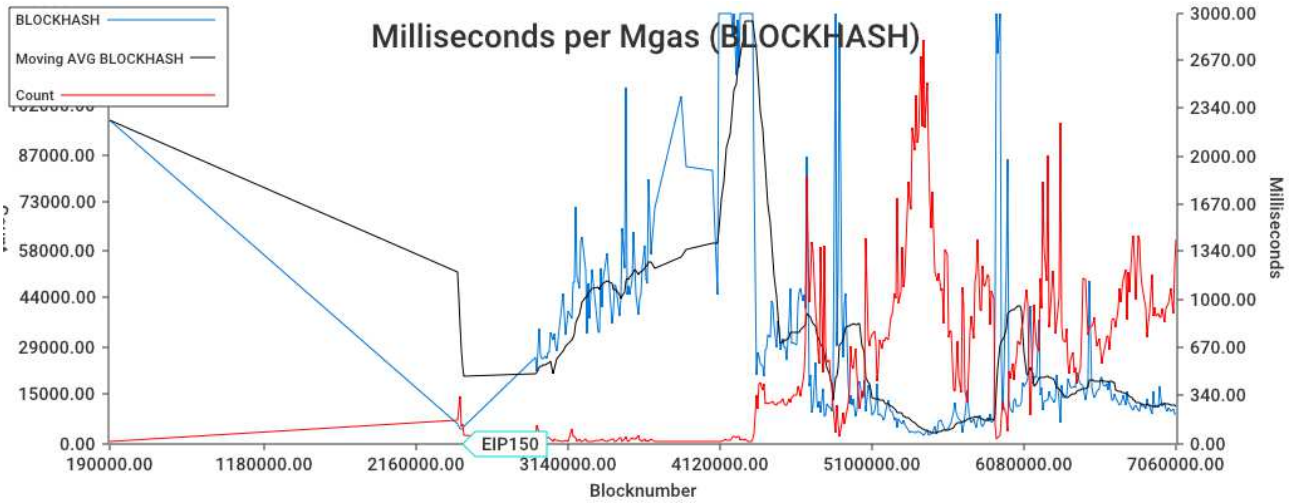

There are some further optimizations that can be done on this op, at least for geth.

Storage and execution (`0x50` range)

These are the ops

```
POP
MLOAD
MSTORE
MSTORE8
SLOAD
SSTORE
JUMP
JUMPI
PC
MSIZE
GAS
JUMPDEST
```

However,

- `MSTORE` and `MSTORE8` have an additional cost for expanding memory

- `SLOAD` varies depending on previous value
- `JUMP` / `JUMPI` have a hidden cost: they require jumpdest analysis.

So the first graph here shows `POP`, `MLOAD`, `SLOAD`, `PC`, `MSIZE` and `GAS`,

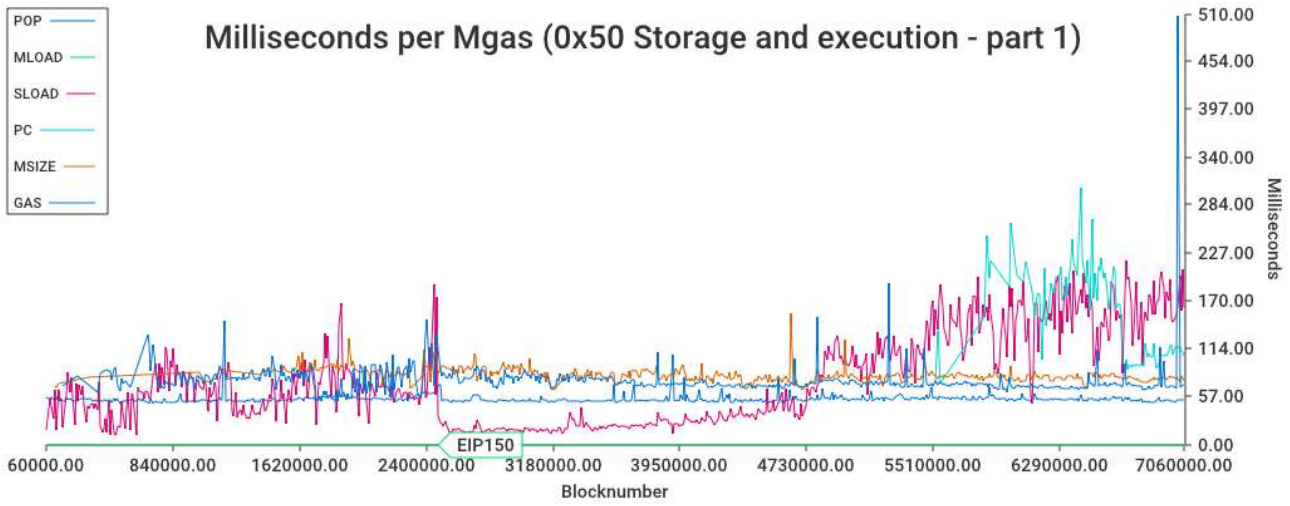

It's clear that `SLOAD` went down at `EIP150`, but has started to rise significantly since then. At around `5M`, it was back to the same levels as before `EIP150`.

`0x60` range

These are `DUPX`, `SWAPX` and `PUSHX`. Here capped for spikes, at `600`:

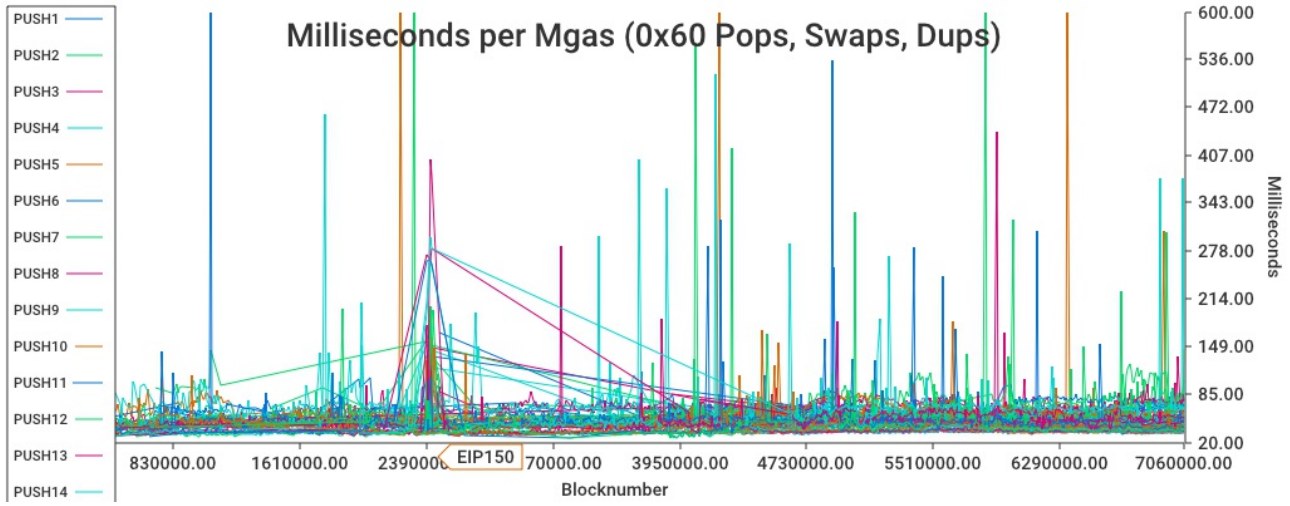

Here capped at `100`:

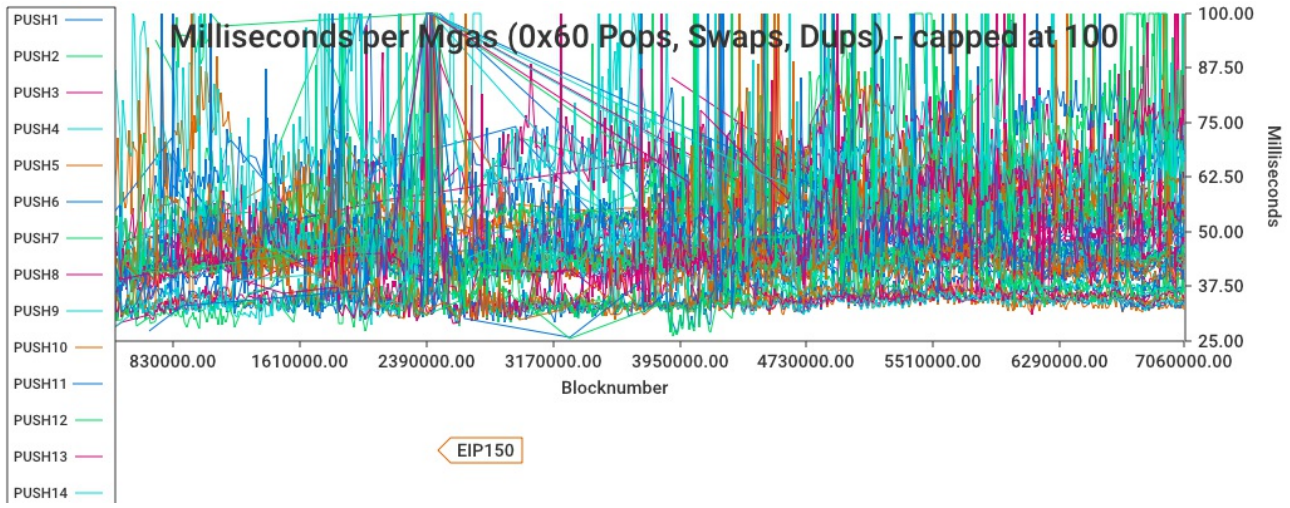

`0xa0` - Logging

The `LOG` opcodes are dynamically priced, depending on the memory size, so we don't have gas/time charts, but here's a time spent-chart:

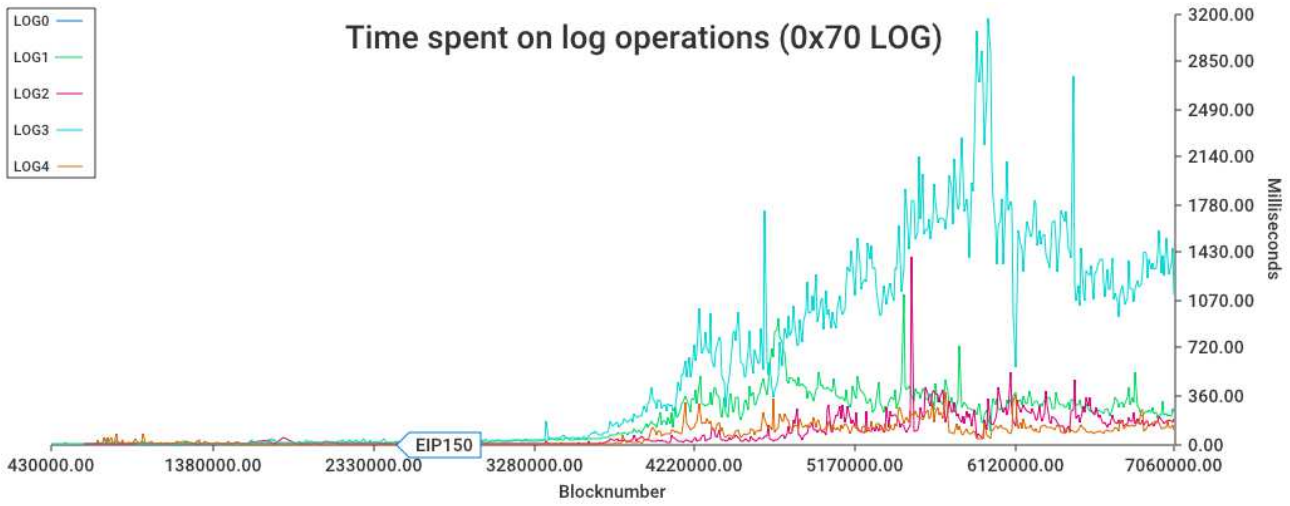

`0xf0` - calling

Remaining ops are 'special' - both dynamic costs and non-trivial effects, such as starting new call contexts or exiting from call contexts.

```
CREATE
CALL
CALLCODE
RETURN
DELEGATECALL
CREATE2
STATICCALL
REVERT
SELFDESTRUCT
```

Some individual OPS

SLOAD

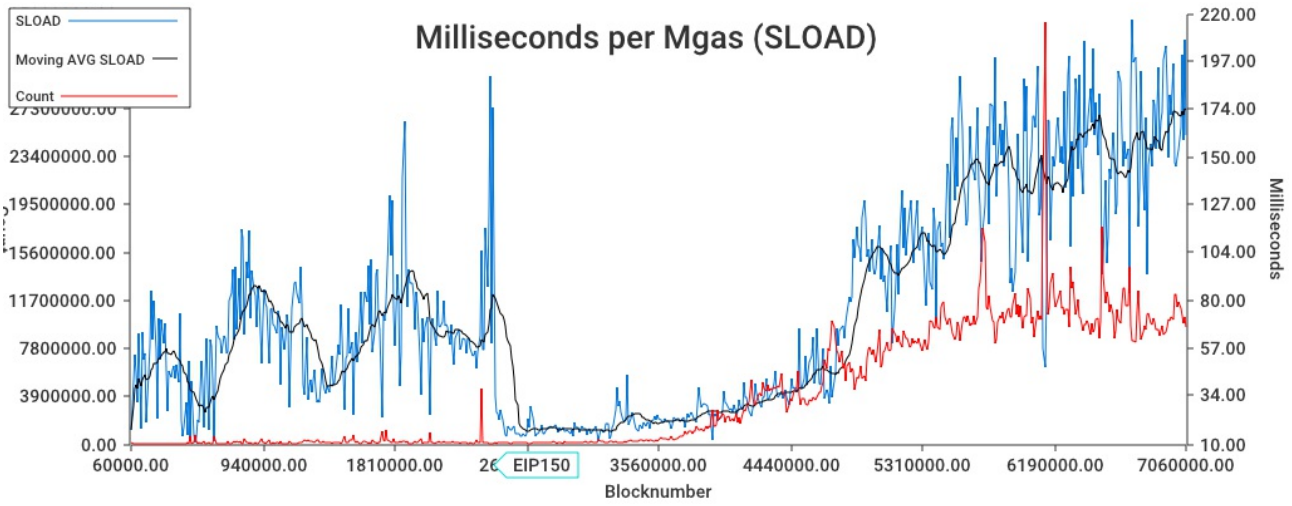

BALANCE

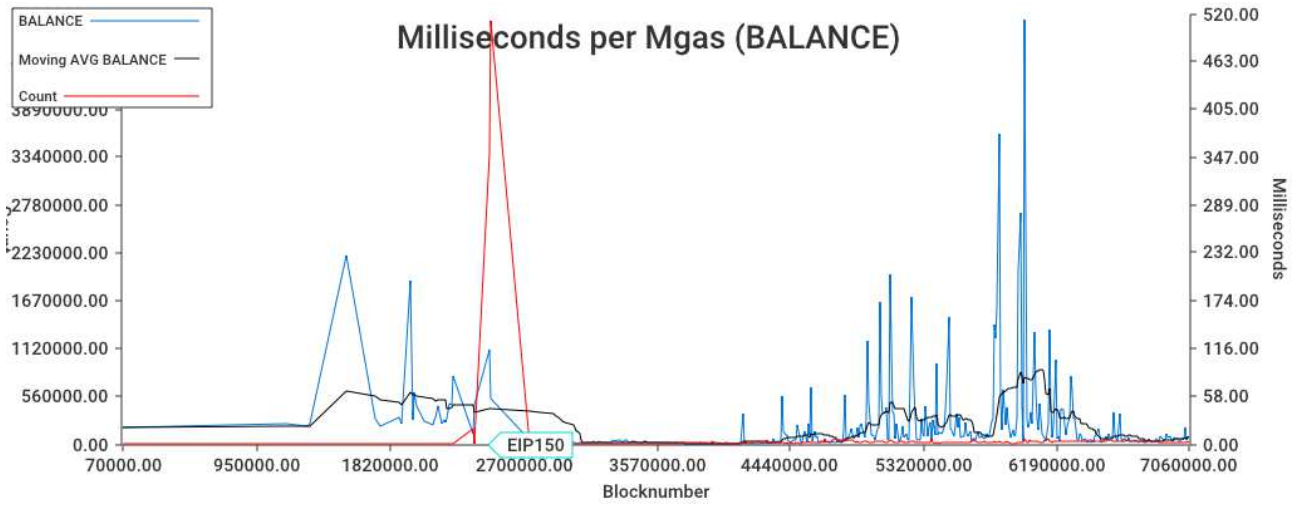

So what's the most 'heavy'

If 'heavy' means large time per gas unit. Here are some charts, where you can also see the actual cost. For the really cheap opcodes, like `PC`, those are probably over-represented, since they are so brief that the actual execution is in roughly on the same order of magnitude as actually performing the measurement.

Anyway, some graphs:

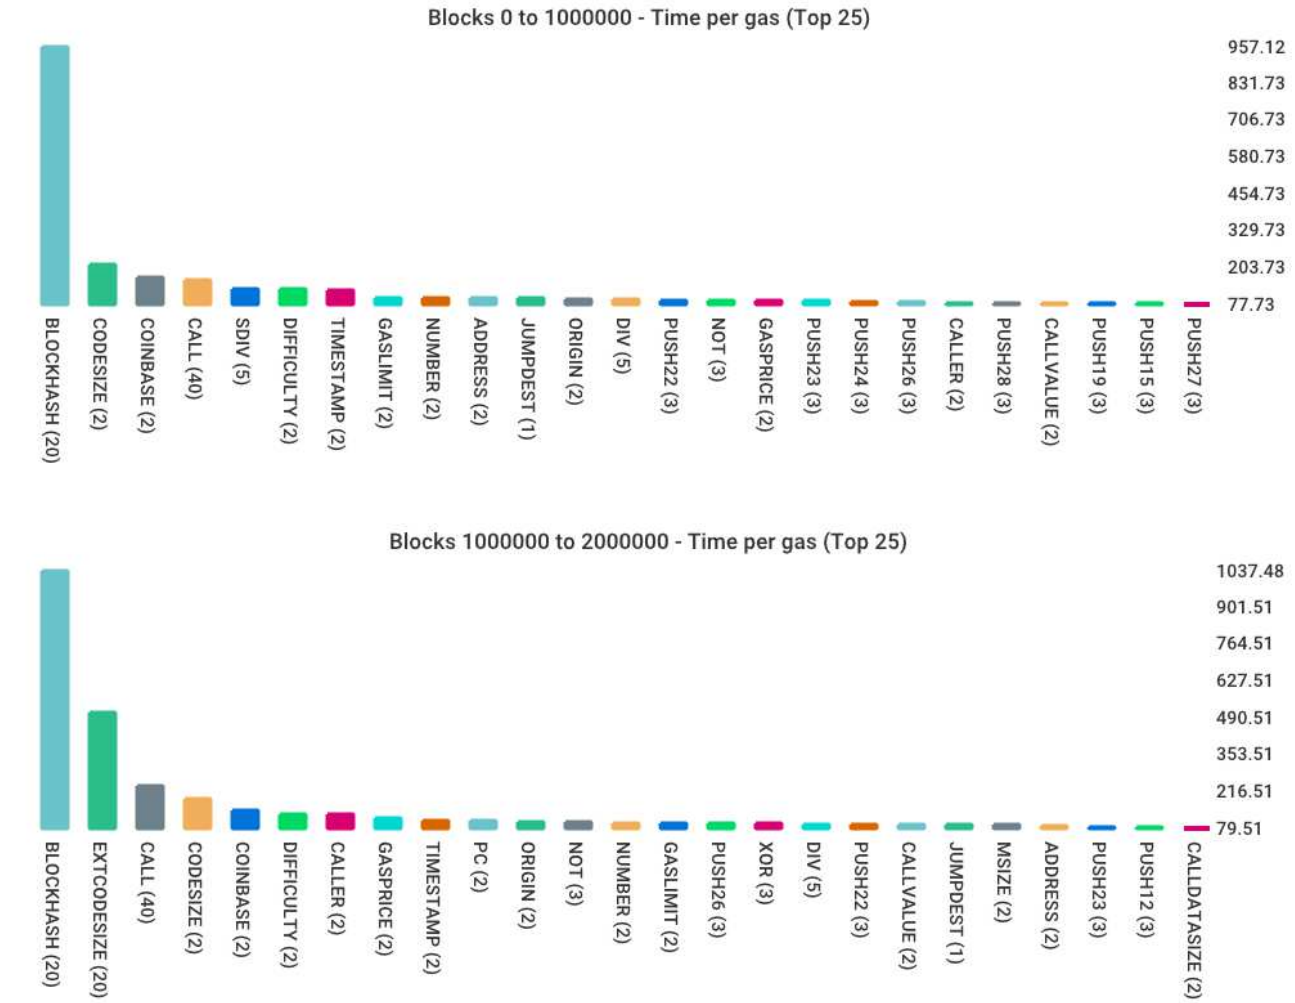

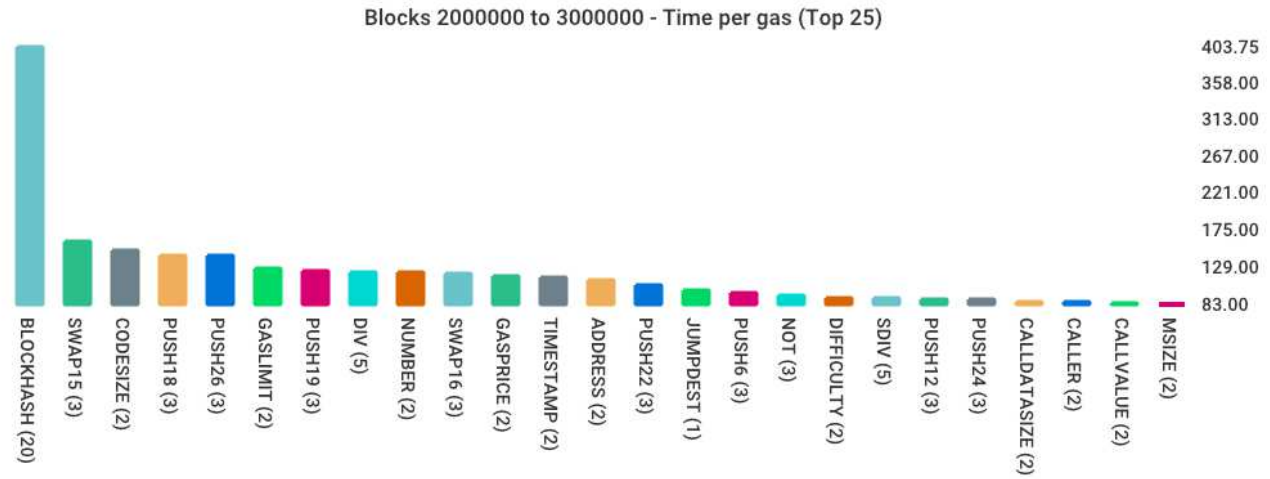
Blocks 2000000 to 3000000 - Time per gas (Top 25)
403.75
358.00
313.00
267.00
221.00
175.00
129.00
83.00
BLOCKHASH (20)
SWAP15 (3)
CODESIZE (2)
PUSH18 (3)
PUSH26 (3)
GASLIMIT (2)
PUSH19 (3)
DIV (5)
NUMBER (2)
SWAP16 (3)
GASPRICE (2)
TIMESTAMP (2)
ADDRESS (2)
PUSH22 (3)
JUMPDEST (1)
PUSH6 (3)
NOT (3)
DIFFICULTY (2)
SDIV (5)
PUSH12 (3)
PUSH24 (3)
CALLDATASIZE (2)
CALLER (2)
CALLVALUE (2)
MSIZE (2)

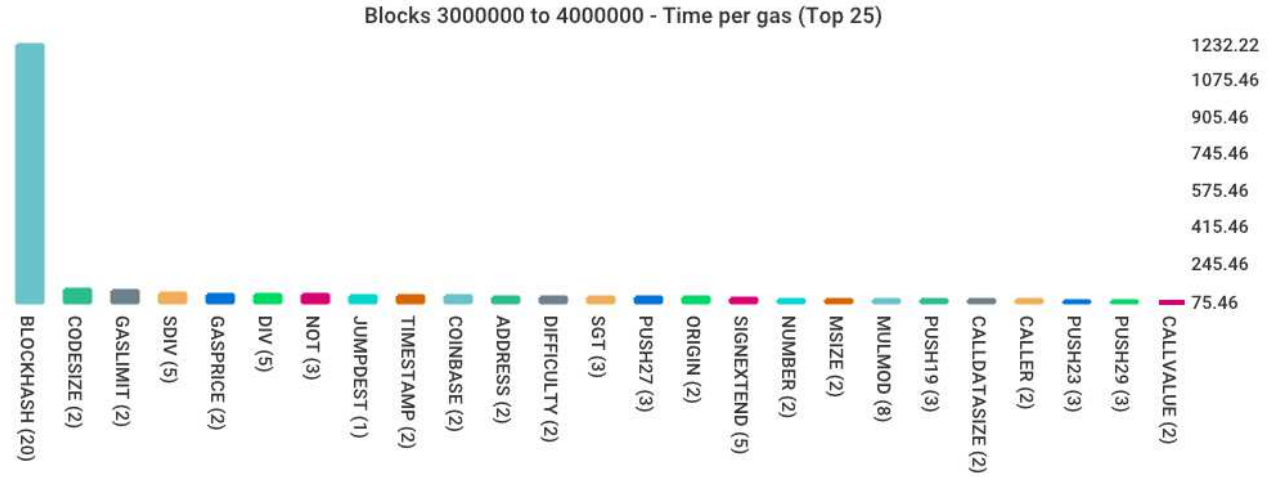
Blocks 3000000 to 4000000 - Time per gas (Top 25)
1232.22
1075.46
905.46
745.46
575.46
415.46
245.46
75.46
BLOCKHASH (20)
CODESIZE (2)
GASLIMIT (2)
SDIV (5)
GASPRICE (2)
DIV (5)
NOT (3)
JUMPDEST (1)
TIMESTAMP (2)
COINBASE (2)
ADDRESS (2)
DIFFICULTY (2)
SGT (3)
PUSH27 (3)
ORIGIN (2)
SIGNEXTEND (5)
NUMBER (2)
MSIZE (2)
MULMOD (8)
PUSH19 (3)
CALLDATASIZE (2)
CALLER (2)
PUSH23 (3)
PUSH29 (3)
CALLVALUE (2)

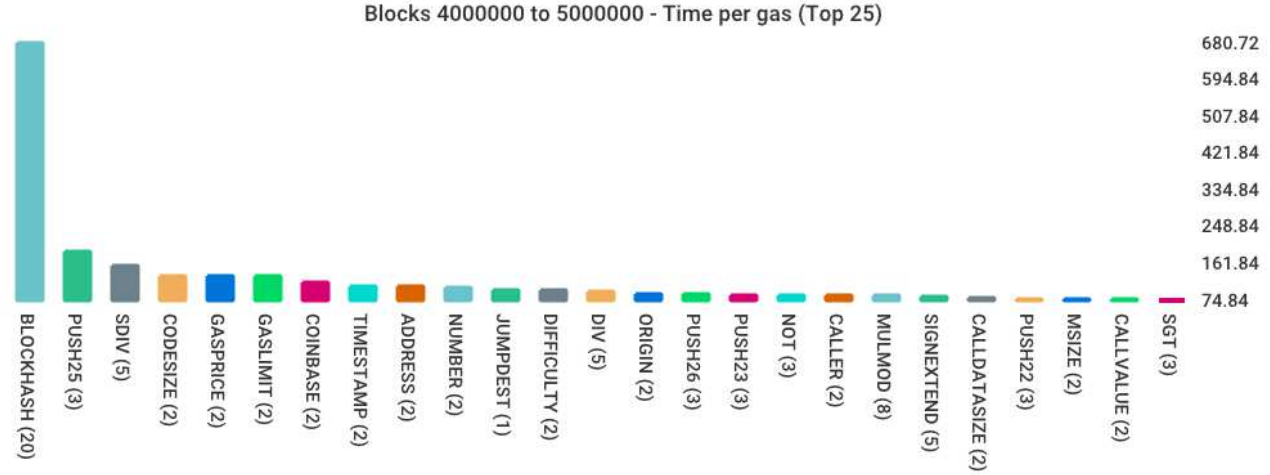
Blocks 4000000 to 5000000 - Time per gas (Top 25)
680.72
594.84
507.84
421.84
334.84
248.84
161.84
74.84
BLOCKHASH (20)
PUSH25 (3)
SDIV (5)
CODESIZE (2)
GASPRICE (2)
GASLIMIT (2)
COINBASE (2)
TIMESTAMP (2)
ADDRESS (2)
NUMBER (2)
JUMPDEST (1)
DIFFICULTY (2)
DIV (5)
ORIGIN (2)
PUSH26 (3)
PUSH23 (3)
NOT (3)
CALLER (2)
MULMOD (8)
SIGNEXTEND (5)
CALLDATASIZE (2)
PUSH22 (3)
MSIZE (2)
CALLVALUE (2)
SGT (3)

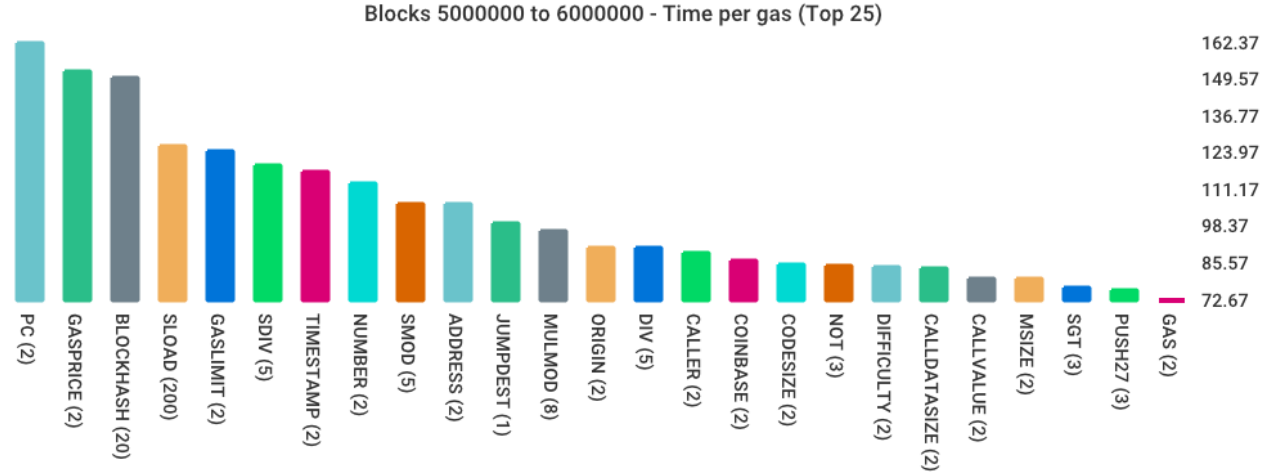
Blocks 5000000 to 6000000 - Time per gas (Top 25)
162.37
149.57
136.77
123.97
111.17
98.37
85.57
72.67
PC (2)
GASPRICE (2)
BLOCKHASH (20)
SLOAD (200)
GASLIMIT (2)
SDIV (5)
TIMESTAMP (2)
NUMBER (2)
SMOD (5)
ADDRESS (2)
JUMPDEST (1)
MULMOD (8)
ORIGIN (2)
DIV (5)
CALLER (2)
COINBASE (2)
CODESIZE (2)
NOT (3)
DIFFICULTY (2)
CALLDATASIZE (2)
CALLVALUE (2)
MSIZE (2)
SGT (3)
PUSH27 (3)
GAS (2)

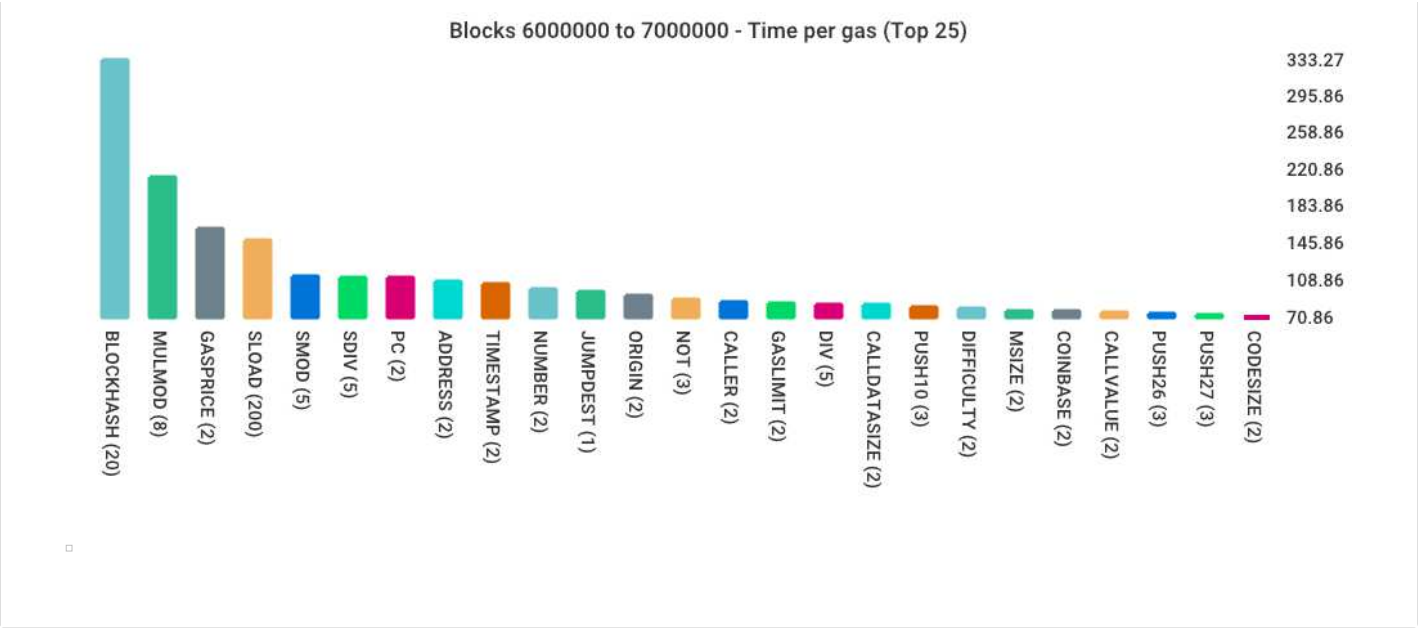

Computer Science > Artificial Intelligence

Unethical Research: How to Create a Malevolent Artificial Intelligence

Federico Pistono, Roman V. Yampolskiy

(Submitted on 10 May 2016 (v1), last revised 1 Sep 2016 (this version, v2))

Cybersecurity research involves publishing papers about malicious exploits as much as publishing information on how to design tools to protect cyber-infrastructure. It is this information exchange between ethical hackers and security experts, which results in a well-balanced cyber-ecosystem. In the blooming domain of AI Safety Engineering, hundreds of papers have been published on different proposals geared at the creation of a safe machine, yet nothing, to our knowledge, has been published on how to design a malevolent machine. Availability of such information would be of great value particularly to computer scientists, mathematicians, and others who have an interest in AI safety, and who are attempting to avoid the spontaneous emergence or the deliberate creation of a dangerous AI, which can negatively affect human activities and in the worst case cause the complete obliteration of the human species. This paper provides some general guidelines for the creation of a Malevolent Artificial Intelligence (MAI).

Subjects:	**Artificial Intelligence (cs.AI)**
Journal reference:	In proceedings of Ethics for Artificial Intelligence Workshop (AI-Ethics-2016). Pages 1-7. New York, NY. July 9 -- 15, 2016
Cite as:	**arXiv:1605.02817 [cs.AI]**
	(or **arXiv:1605.02817v2 [cs.AI]** for this version)

Submission history

From: Roman Yampolskiy [view email]
[v1] Tue, 10 May 2016 01:39:38 UTC (305 KB)
[v2] Thu, 1 Sep 2016 18:29:13 UTC (337 KB)

Which authors of this paper are endorsers? | *Disable MathJax* (*What is MathJax?*)

ISSN 1551-3483

https://scale.qihardware.org