

1. From: Mike Hearn <mike@plan99.net>  
2. Date: Sun, Apr 12, 2009 at 12:46 PM  
3. To: satoshin@gmx.com  
4.  
5.  
6. Hi Satoshi,  
7.  
8. I read your paper on BitCoin with great interest. I found it a bit  
9. confusing though - I believe it may be easier to follow if you provide  
10. some examples.  
11.  
12. Specifically, it's not quite clear to me what blocks contain. If I  
13. understand correctly, there is only one (or maybe a few) global  
14. chain[s] into which all transactions are hashed. If there is only one  
15. chain recording "the story of the economy" so to speak, how does this  
16. scale? In an imaginary planet-wide deployment there would be millions  
17. of even billions of transactions per hour being hashed into the chain.  
18. I realize that each PoW can wrap many transactions in one block,  
19. nonetheless, that's a large amount of data to hash. If there are many  
20. chains, how are transactions assigned to each chain such that it is  
21. still difficult to overpower the network? Eg, if there are 10 global  
22. chains, the amount of cpu power you need to beat the system is only  
23. 10% of what it was previously.  
24.  
25. I also wonder if the assumption of 1 core = 1 vote is sound. If the  
26. majority of nodes are on standard computers, it seems likely that an  
27. attacker could use FPGA or custom ASICs to get significantly better  
28. performance. What are your thoughts on using custom hardware to beat  
29. the chain?  
30.  
31. I found the section on incentives hard to follow. In particular, I'm  
32. not clear on what triggers the transition from minting new coins as a  
33. reason to run a node, to charging transaction fees (isn't the point of  
34. BitCoin largely to zero transaction costs anyway?). Presumably there's  
35. some human in charge of the system - eg, you decided somehow that 24  
36. million coins was a good number to have, and would distribute some  
37. kind of rules file saying "coins minted after this timestamp must have  
38. an N+1 zero bits prefix", which honest nodes enforce.  
39.  
40. How did you decide on the inflation schedule for v1? Where did 24  
41. million coins come from? What denominations are these coins? You  
42. mention a way to combine and split value but I'm not clear on how this  
43. works. For instance are bitcoins always denominated by an integer or  
44. can you have fractional bitcoins?  
45.  
46. So many questions :) But it's rare that I encounter truly  
47. revolutionary ideas. The last time I was this excited about a new  
48. monetary scheme was when I discovered Ripple. If you have any thoughts  
49. on Ripple, I'd also love to hear them.  
50.  
51. thanks -mike  
52.  
53. -----  
54. From: Satoshi Nakamoto <satoshin@gmx.com>

55. Date: Sun, Apr 12, 2009 at 10:44 PM

56. To: Mike Hearn <mike@plan99.net>

57.

58.

59. Hi Mike,

60.

61. I'm glad to answer any questions you have. If I get time, I ought to write a FAQ to supplement the paper.

62.

63. There is only one global chain.

64.

65. The existing Visa credit card network processes about 15 million Internet purchases per day worldwide. Bitcoin can already scale much larger than that with existing hardware for a fraction of the cost. It never really hits a scale ceiling. If you're interested, I can go over the ways it would cope with extreme size.

66.

67. By Moore's Law, we can expect hardware speed to be 10 times faster in 5 years and 100 times faster in 10. Even if Bitcoin grows at crazy adoption rates, I think computer speeds will stay ahead of the number of transactions.

68.

69. I don't anticipate that fees will be needed anytime soon, but if it becomes too burdensome to run a node, it is possible to run a node that only processes transactions that include a transaction fee. The owner of the node would decide the minimum fee they'll accept. Right now, such a node would get nothing, because nobody includes a fee, but if enough nodes did that, then users would get faster acceptance if they include a fee, or slower if they don't. The fee the market would settle on should be minimal. If a node requires a higher fee, that node would be passing up all transactions with lower fees. It could do more volume and probably make more money by processing as many paying transactions as it can. The transition is not controlled by some human in charge of the system though, just individuals reacting on their own to market forces.

70.

71. Eventually, most nodes may be run by specialists with multiple GPU cards. For now, it's nice that anyone with a PC can play without worrying about what video card they have, and hopefully it'll stay that way for a while. More computers are shipping with fairly decent GPUs these days, so maybe later we'll transition to that.

72.

73. A key aspect of Bitcoin is that the security of the network grows as the size of the network and the amount of value that needs to be protected grows. The down side is that it's vulnerable at the beginning when it's small, although the value that could be stolen should always be smaller than the amount of effort required to steal it. If someone has other motives to prove a point, they'll just be proving a point I already concede.

74.

75. My choice for the number of coins and distribution schedule was an educated guess. It was a difficult choice, because once the network is going it's locked in and we're stuck with it. I wanted to pick something that would make prices similar to existing currencies, but without knowing the future, that's very hard. I ended up picking something in the middle. If Bitcoin remains a small niche, it'll be worth less per unit than existing currencies. If you imagine it being used for some fraction of world commerce, then there's only going to be 21 million coins for the whole world, so it would be worth much more per unit. Values are 64-bit integers with 8 decimal places, so 1 coin is represented internally as 100000000. There's plenty of granularity if typical prices become small. For example, if 0.001 is worth 1 Euro, then it might be easier to change where the decimal point is displayed, so if you had 1 Bitcoin it's now displayed as 1000, and 0.001 is displayed as 1.

76.

77. Ripple is interesting in that it's the only other system that does something with trust besides concentrate it into a central server.

78.  
79. Satoshi  
80.  
81. -----  
82. From: Mike Hearn <mike@plan99.net>  
83. Date: Mon, Apr 13, 2009 at 1:39 PM  
84. To: Satoshi Nakamoto <satoshin@gmx.com>  
85.  
86.  
87. Thanks Satoshi,  
88.  
89. I tried the app yesterday. It seems to work pretty well running on  
90. Wine (I tried it on MacOS but it should run on Linux too, and will try  
91. that next week when I am back at work).  
92.  
93. In the lower right hand corner it has a block count which increases  
94. rapidly and then stops. Is this the length of the global chain? It  
95. seems to advance far too fast for that. Or is this the number of  
96. genesis blocks that have been tried but did not result in a partial  
97. collision? I'm not sure if the way it stops and starts is expected, or  
98. some glitch caused by it running under emulation. My best guess - it  
99. is the length of the global chain, and the rapid advance at the start  
100. is as the software downloads and verifies the preceding blocks in the  
101. chain as being valid.  
102.  
103. With regards to the buyer/seller experience, I understand that the  
104. global chain advances at about 6-7 blocks per hour under the current  
105. settings. If we assume that 0.1% is a good risk rate, then  $z=5$  thus  
106. any transaction must wait a bit less than an hour before being  
107. solidified in the chain. As micropayments for things like web content  
108. or virtual goods are by definition something that requires low  
109. overhead, waiting an hour seems like quite a significant hurdle.  
110.  
111. I understand that nodes attempt to find a POW to advance the global  
112. chain in an uncoordinated fashion. This sentence however:  
113.  
114. "If a majority of CPU power is controlled by honest nodes, the  
115. honest chain will grow the fastest and outpace any competing chains."  
116.  
117. is confusing for me, because it appears the only way the honest chain  
118. can grow faster than a chain worked on by 1 attacking cpu is if the  
119. keyspace to scan looking for a partial collision is sharded evenly  
120. amongst the participating honest nodes. That way the speed at which  
121. collisions are found would be proportional to the number of nodes. Yet  
122. I don't see any discussion of such work sharding, which obviously adds  
123. complexity. Likewise:  
124.  
125. "To compensate for increasing hardware speed and varying interest  
126. in running nodes over time,  
127. the proof-of-work difficulty is determined by a moving average  
128. targeting an average number of  
129. blocks per hour. If they're generated too fast, the difficulty increases."  
130.  
131. How is the required difficulty of each block communicated through the

132. network and agreed upon?  
133.  
134. Thanks once again. I have yet more questions but this is enough for  
135. one email :) I will be happy to summarize these discussions into an  
136. FAQ-like document at some point. Apologies if the questions seem  
137. trivial.  
138.  
139. -mike  
140.  
141. -----  
142. From: Mike Hearn <mike@plan99.net>  
143. Date: Mon, Apr 13, 2009 at 10:51 PM  
144. To: Satoshi Nakamoto <satoshin@gmx.com>  
145.  
146.  
147. Something else that isn't clear to me - does the global chain only get  
148. extended when there is actual work to do? Currently it seems to grow  
149. all the time, although there are only a few people in the network. So  
150. presumably it gets extended with null blocks. Is this actually  
151. required? The timestamping doesn't have to be actually in parallel  
152. with real time does it ... it's merely establishing an ordering of  
153. events.  
154.  
155. -----  
156. From: Satoshi Nakamoto <satoshin@gmx.com>  
157. Date: Mon, Apr 13, 2009 at 11:00 PM  
158. To: Mike Hearn <mike@plan99.net>  
159.  
160.  
161. Mike Hearn wrote:  
162.  
163. My best guess - it  
164. is the length of the global chain, and the rapid advance at the start  
165. is as the software downloads and verifies the preceding blocks in the  
166. chain as being valid.  
167.  
168.  
169. Right. I'm trying to think of more clear wording for that, maybe "%d network blocks" or "%d block  
chain".  
170.  
171.  
172. If we assume that 0.1% is a good risk rate, then  $z=5$  thus  
173. any transaction must wait a bit less than an hour before being  
174. solidified in the chain. As micropayments for things like web content  
175. or virtual goods are by definition something that requires low  
176. overhead, waiting an hour seems like quite a significant hurdle.  
177.  
178.  
179. For the actual risk, multiply the 0.1% by the probability that the buyer is an attacker with a huge  
network of computers.  
180.  
181. For micropayments, you can safely accept the payment immediately. The size of the payment is too small  
for the effort to steal it. Micropayments are almost always for intellectual property, where there's no

physical loss to the merchant. Anyone trying to steal a micropayment would probably not be a paying customer anyway, and if they want to steal intellectual property they can use the file sharing networks.

Currently, businesses accept a certain chargeoff rate. I believe the risk with 1 or even 0 confirming blocks will be much less than the rate of chargebacks on verified credit card transactions.

The usual scam against a merchant that doesn't wait for confirming blocks would be to send a payment to a merchant, then quickly try to propagate a double-spend to the network before the merchant's copy. What the merchant can do is broadcast his transaction and then monitor the network for any double-spend copies. The thief would not be able to broadcast during the monitoring period or else the merchant's node would receive a copy. The merchant would only have to monitor for a minute or two until most of the network nodes have his version and it's too late for the thief's version to catch up and reach many nodes. With just a minute or two delay, the chance of getting away without paying could be made much too low to scam. A thief usually needs a high probability of getting an item for free to make it worthwhile. Using a lot of CPU power to do the brute force attack discussed in the paper in addition to the above scam would not increase the thief's chances very much.

Anything that grants access to something, like something that takes a while to download, access to a website, web hosting, a subscription or service, can be cancelled a few minutes later if the transaction is rejected.

is confusing for me, because it appears the only way the honest chain can grow faster than a chain worked on by 1 attacking cpu is if the keyspace to scan looking for a partial collision is sharded evenly amongst the participating honest nodes. That way the speed at which collisions are found would be proportional to the number of nodes. Yet I don't see any discussion of such work sharding, which obviously adds complexity.

The keyspace is huge,  $2^{256}$ . The thing being hashed includes the node's public key and a random nonce, so the chance of any two nodes duplicating work on the same space is negligible.

How is the required difficulty of each block communicated through the network and agreed upon?

It's not communicated. The formula is hardcoded in the program and every node does the same calculation to know what difficulty is required for the next block. If someone diverged from the formula, their block would not be accepted by the majority.

Thanks once again. I have yet more questions but this is enough for one email :) I will be happy to summarize these discussions into an FAQ-like document at some point. Apologies if the questions seem trivial.

No problem, thanks for testing it on Mac Wine.

Satoshi

219. -----  
220. From: Satoshi Nakamoto <satoshin@gmx.com>  
221. Date: Mon, Apr 13, 2009 at 11:11 PM  
222. To: Mike Hearn <mike@plan99.net>  
223.  
224.  
225. It keeps getting extended all the time. If it stopped, an attacker would have time to catch up. Don't  
worry, empty blocks aren't very big.  
226.  
227. As you say, it's the order of events that matters.  
228.  
229. -----  
230. From: Mike Hearn <mike@plan99.net>  
231. Date: Mon, Apr 13, 2009 at 11:18 PM  
232. To: Satoshi Nakamoto <satoshin@gmx.com>  
233.  
234.  
235. Oh yes, of course, that's fundamental. Silly me. Thanks for your  
236. answers. I'd recommend being over-explicit for early versions of the  
237. software, something like "Global chain is currently %d blocks long".  
238.  
239. I guess the key problem right now is that once you generate coins,  
240. there's nobody to test it with, even for dummy transactions. Is there  
241. a plan for a mailing list or some kind of trivial marketplace to give  
242. people something to do with their newly minted bitcoins?  
243.  
244. -----  
245. From: Satoshi Nakamoto <satoshin@gmx.com>  
246. Date: Tue, Apr 14, 2009 at 7:41 PM  
247. To: Mike Hearn <mike@plan99.net>  
248.  
249.  
250. I started implementing a marketplace feature earlier that facilitates offering things for sale and  
taking orders, it's only half done though. A bit like e-bay but without auctions, just "buy now". Among  
other things, it would make it easy for anyone to offer currency exchange.  
251.  
252. If you send to 1PhUXucRd8FzQved2KGK3g1eKfTHPGjgFu and e-mail me your bitcoin address, or IP if you can  
accept incoming connections, I'll send back the same amount +50.  
253.  
254. -----  
255. From: Mike Hearn <mike@plan99.net>  
256. Date: Sat, Apr 18, 2009 at 3:08 PM  
257. To: Satoshi Nakamoto <satoshin@gmx.com>  
258.  
259.  
260. Hi Satoshi,  
261.  
262. I sent you 32.51 coins, my bitcoin address is 1JuEjh9znXwqsy5RrnKqgzqY4Ldg7rnj5n  
263.  
264. My IP is currently 84.73.233.199, however, it's a laptop so may or may  
265. not be online at the time you act on this mail. I suggest using the  
266. bitcoin address instead. It'd be convenient if the same comment  
267. functionality was available via indirect transfer. Can the comment be  
268. encrypted using the public key of the receiver and placed into a

269. block?  
270.  
271. -----  
272. From: Satoshi Nakamoto <satoshin@gmx.com>  
273. Date: Sat, Apr 18, 2009 at 6:16 PM  
274. To: Mike Hearn <mike@plan99.net>  
275.  
276.  
277. I sent back 32.51 and 50.00.  
278.  
279. I badly wanted to find some way to include a comment with indirect transfers, but there just wasn't a way to do it. Bitcoin uses EC-DNA, which was essential for making the block chain compact enough to be practical with today's technology because its signatures are an order of magnitude smaller than RSA. But EC-DNA can't encrypt messages like RSA, it can only be used to verify signatures.  
280.  
281. -----  
282. From: Mike Hearn <mike@plan99.net>  
283. Date: Sat, Apr 18, 2009 at 9:25 PM  
284. To: Satoshi Nakamoto <satoshin@gmx.com>  
285.  
286.  
287. Thanks. I sent you back 50, so now we're even.  
288.  
289. For some reason your transfer to me shows up as "From: unknown" even  
290. though I added you to my address book.  
291.  
292. I have a "Generated (not accepted)" line in my transaction list, it  
293. seems like an attempt to generate a coin went wrong somehow. Not sure  
294. what happened here - presumably my node successfully solved a block  
295. but then I went offline before it was sent to the network?  
296.  
297. I suppose for sending metadata with a transaction some other mechanism  
298. will be needed, for instance, broadcast of encrypted messages  
299. associated with a transaction that persist for (say) a month, with  
300. some kind of budget on how much storage a node can use for messages.  
301. Alternatively, a payee could generate some reference number which is  
302. of some significance to themselves but otherwise opaque, and give it  
303. to the payer, thus it does not need to be encrypted and can be put  
304. into the block directly.  
305.  
306. -----  
307. From: Satoshi Nakamoto <satoshin@gmx.com>  
308. Date: Sat, Apr 18, 2009 at 10:52 PM  
309. To: Mike Hearn <mike@plan99.net>  
310.  
311.  
312. Got the 50.  
313.  
314. Transactions sent to a bitcoin address will always say "from: unknown". The transaction only tells who it's to. Sending by bitcoin address has a number of problems, but it's so nice having the fallback option to be able to send to anyone whether they're online or not. There are a number of ideas to try to improve things later. For now, if things work out like the real world where the vast majority of transactions are with merchants, they'll pretty much always make sure to set up to receive by IP. The

P2P file sharing networks seem fairly successful at getting a large percentage of their users to set up their firewalls to forward a port.

The "Generated (not accepted)" normally happens if two nodes find a block at close to the same time, one of them will not be accepted. It's normal and unavoidable. I plan in v0.1.6 to hide those, since they're just confusing and annoying and there's no reason for users to have to see them. While the network is still small like it is now, if you can't receive incoming connections you're at more of a disadvantage because you can't receive block announcements as directly.

-----  
From: Mike Hearn <mike@plan99.net>  
Date: Sat, Apr 18, 2009 at 11:23 PM  
To: Satoshi Nakamoto <satoshin@gmx.com>

Yes, I believe most P2P clients use the UPnP protocol to get routers to open up the port automatically. That would probably improve the listen rate significantly. I just discovered DMZ wasn't enabled on my router, though I thought it was. That's now fixed.

Is there a way to be told of new versions? Does the app auto update itself? Again, some kind of mailing list would be excellent.

I was thinking through how a practical micropayment implementation for the web might work in the last few days. One key issue is ensuring micropayments are fully automatic, yet can't be easily abused to drain the users account. I think the right approach would be to allow any website that presents an EV SSL cert to automatically request a micropayment, by default the browser always accepts as long as the charge is "low" and displays a small notification of what has occurred. Sites can then show that content requires payment in any way that suits their site design. Abusive sites that don't meet some simple guidelines (eg, showing unambiguously that clicking a link will trigger payment, or taking payment from direct search engine links) would simply have their SSL cert blacklisted, much like anti-phishing filters work today.

The protocol could be very straightforward and implemented by a Firefox extension or an IE BHO. Some static file (eg, a protocol buffer) is hosted on the site. It specifies the charge, a transaction description, the target IP and a URL for the browser to load after the transaction was accepted by the target node, to which the user identifier is sent in a URL parameter. The site can then give back a cookie and the paywalled content. The entire process is automatic and simply results in, say, a little coin animation in the URL bar. Thus it's as convenient as regular web browsing. The users software would have some limit on what payments are automatically accepted.

The main problem with this approach is that somebody has to decide what the user interface guidelines are, then enforce them via blacklisting, as well as decide what payment requirements are low enough to be automatic vs requiring a user prompt. This introduces a trusted authority back into the system. However, it's one that the user can choose in an open market.



363.  
364. By the way, if you're not already using protocol buffers for the  
365. node-to-node traffic, I recommend them. We use them here at Google for  
366. everything, they solve a lot of versioning problems simply and  
367. efficiently.  
368.  
369. -----  
370. From: Satoshi Nakamoto <satoshin@gmx.com>  
371. Date: Sun, Apr 19, 2009 at 2:14 AM  
372. To: Mike Hearn <mike@plan99.net>  
373.  
374.  
375. The list is:  
376. bitcoin-list@lists.sourceforge.net  
377. Subscribe/unsubscribe page:  
378. <http://lists.sourceforge.net/mailman/listinfo/bitcoin-list>  
379. Archives:  
380. [http://sourceforge.net/mailarchive/forum.php?forum\\_name=bitcoin-list](http://sourceforge.net/mailarchive/forum.php?forum_name=bitcoin-list)  
381.  
382. I'll always announce new versions there. Automatic update, or at least notification of new versions, is  
definitely on the list. There could potentially be necessary changes in the future where nobody will  
want to talk to you until you upgrade, and there needs to be code in the older version to convey that to  
the user. This is all the harder in the context of not trusting anyone.  
383.  
384. Your approach to micropayments sounds right. At first, it might be a good idea to default to asking  
permission until the user gets comfortable and is ready to set it to automatic. The end goal though  
should get to something like you describe, where it's similar to using your cell phone without really  
having to think about the per minute charges.  
385.  
386. I looked at Google protocol buffers when they were released last year, but I had already written  
everything by then. What I did was something similar to Boost Serialisation. For this application, where  
I was parsing messages from strangers who might have extreme incentive to hack the protocol, it was  
necessary to make it as basic as possible so I could crawl over every line of code to convince myself it  
was airtight. It became clear that any unnecessary degrees of freedom in the binary format multiplied  
the potential angles of attack. You guys are so right though to standardize across the company on  
protocol buffers. I think you've got the optimal solution in the general case.  
387.  
388. -----  
389. From: Mike Hearn <mike@plan99.net>  
390. Date: Thu, May 2, 2013 at 10:02 AM  
391. To: satoshiarchive@gmail.com  
392.  
393.  
394.  
395.  
396. Forwarded conversation  
397. Subject: Questions about BitCoin  
398. -----  
399.  
400. From: Mike Hearn <mike@plan99.net>  
401. Date: Sun, Apr 12, 2009 at 12:46 PM  
402. To: satoshin@gmx.com  
403.  
404. -----

405. From: Satoshi Nakamoto <satoshin@gmx.com>  
406. Date: Sun, Apr 12, 2009 at 10:44 PM  
407. To: Mike Hearn <mike@plan99.net>  
408.  
409.  
410. Hi Mike,  
411.  
412. I'm glad to answer any questions you have. If I get time, I ought to write a FAQ to supplement the  
paper.  
413.  
414. There is only one global chain.  
415.  
416. The existing Visa credit card network processes about 15 million Internet purchases per day worldwide.  
Bitcoin can already scale much larger than that with existing hardware for a fraction of the cost. It  
never really hits a scale ceiling. If you're interested, I can go over the ways it would cope with  
extreme size.  
417.  
418. By Moore's Law, we can expect hardware speed to be 10 times faster in 5 years and 100 times faster in  
10. Even if Bitcoin grows at crazy adoption rates, I think computer speeds will stay ahead of the number  
of transactions.  
419.  
420. I don't anticipate that fees will be needed anytime soon, but if it becomes too burdensome to run a  
node, it is possible to run a node that only processes transactions that include a transaction fee. The  
owner of the node would decide the minimum fee they'll accept. Right now, such a node would get nothing,  
because nobody includes a fee, but if enough nodes did that, then users would get faster acceptance if  
they include a fee, or slower if they don't. The fee the market would settle on should be minimal. If a  
node requires a higher fee, that node would be passing up all transactions with lower fees. It could do  
more volume and probably make more money by processing as many paying transactions as it can. The  
transition is not controlled by some human in charge of the system though, just individuals reacting on  
their own to market forces.  
421.  
422. Eventually, most nodes may be run by specialists with multiple GPU cards. For now, it's nice that anyone  
with a PC can play without worrying about what video card they have, and hopefully it'll stay that way  
for a while. More computers are shipping with fairly decent GPUs these days, so maybe later we'll  
transition to that.  
423.  
424. A key aspect of Bitcoin is that the security of the network grows as the size of the network and the  
amount of value that needs to be protected grows. The down side is that it's vulnerable at the beginning  
when it's small, although the value that could be stolen should always be smaller than the amount of  
effort required to steal it. If someone has other motives to prove a point, they'll just be proving a  
point I already concede.  
425.  
426. My choice for the number of coins and distribution schedule was an educated guess. It was a difficult  
choice, because once the network is going it's locked in and we're stuck with it. I wanted to pick  
something that would make prices similar to existing currencies, but without knowing the future, that's  
very hard. I ended up picking something in the middle. If Bitcoin remains a small niche, it'll be worth  
less per unit than existing currencies. If you imagine it being used for some fraction of world  
commerce, then there's only going to be 21 million coins for the whole world, so it would be worth much  
more per unit. Values are 64-bit integers with 8 decimal places, so 1 coin is represented internally as  
100000000. There's plenty of granularity if typical prices become small. For example, if 0.001 is worth  
1 Euro, then it might be easier to change where the decimal point is displayed, so if you had 1 Bitcoin  
it's now displayed as 1000, and 0.001 is displayed as 1.  
427.

428. Ripple is interesting in that it's the only other system that does something with trust besides  
concentrate it into a central server.

429.

430. Satoshi

431.

432. -----

433. From: Mike Hearn <mike@plan99.net>

434. Date: Mon, Apr 13, 2009 at 1:39 PM

435. To: Satoshi Nakamoto <satoshin@gmx.com>

436.

437. -----

438. From: Mike Hearn <mike@plan99.net>

439. Date: Mon, Apr 13, 2009 at 10:51 PM

440. To: Satoshi Nakamoto <satoshin@gmx.com>

441.

442.

443. Something else that isn't clear to me - does the global chain only get  
444. extended when there is actual work to do? Currently it seems to grow  
445. all the time, although there are only a few people in the network. So  
446. presumably it gets extended with null blocks. Is this actually  
447. required? The timestamping doesn't have to be actually in parallel  
448. with real time does it ... it's merely establishing an ordering of  
449. events.

450.

451. -----

452. From: Satoshi Nakamoto <satoshin@gmx.com>

453. Date: Mon, Apr 13, 2009 at 11:00 PM

454. To: Mike Hearn <mike@plan99.net>

455.

456.

457. Mike Hearn wrote:

458.

459. My best guess - it  
460. is the length of the global chain, and the rapid advance at the start  
461. is as the software downloads and verifies the preceding blocks in the  
462. chain as being valid.

463.

464.

465. Right. I'm trying to think of more clear wording for that, maybe "%d network blocks" or "%d block  
chain".

466.

467.

468.

469. If we assume that 0.1% is a good risk rate, then  $z=5$  thus  
470. any transaction must wait a bit less than an hour before being  
471. solidified in the chain. As micropayments for things like web content  
472. or virtual goods are by definition something that requires low  
473. overhead, waiting an hour seems like quite a significant hurdle.

474.

475.

476. For the actual risk, multiply the 0.1% by the probability that the buyer is an attacker with a huge  
network of computers.

477.

478. For micropayments, you can safely accept the payment immediately. The size of the payment is too small for the effort to steal it. Micropayments are almost always for intellectual property, where there's no physical loss to the merchant. Anyone trying to steal a micropayment would probably not be a paying customer anyway, and if they want to steal intellectual property they can use the file sharing networks.

479.

480. Currently, businesses accept a certain chargeoff rate. I believe the risk with 1 or even 0 confirming blocks will be much less than the rate of chargebacks on verified credit card transactions.

481.

482. The usual scam against a merchant that doesn't wait for confirming blocks would be to send a payment to a merchant, then quickly try to propagate a double-spend to the network before the merchant's copy. What the merchant can do is broadcast his transaction and then monitor the network for any double-spend copies. The thief would not be able to broadcast during the monitoring period or else the merchant's node would receive a copy. The merchant would only have to monitor for a minute or two until most of the network nodes have his version and it's too late for the thief's version to catch up and reach many nodes. With just a minute or two delay, the chance of getting away without paying could be made much too low to scam. A thief usually needs a high probability of getting an item for free to make it worthwhile. Using a lot of CPU power to do the brute force attack discussed in the paper in addition to the above scam would not increase the thief's chances very much.

483.

484. Anything that grants access to something, like something that takes a while to download, access to a website, web hosting, a subscription or service, can be cancelled a few minutes later if the transaction is rejected.

485.

486.

487.

488. is confusing for me, because it appears the only way the honest chain

489. can grow faster than a chain worked on by 1 attacking cpu is if the

490. keyspace to scan looking for a partial collision is sharded evenly

491. amongst the participating honest nodes. That way the speed at which

492. collisions are found would be proportional to the number of nodes. Yet

493. I don't see any discussion of such work sharding, which obviously adds

494. complexity.

495.

496.

497. The keyspace is huge,  $2^{256}$ . The thing being hashed includes the node's public key and a random nonce, so the chance of any two nodes duplicating work on the same space is negligible.

498.

499.

500.

501. How is the required difficulty of each block communicated through the

502. network and agreed upon?

503.

504.

505. It's not communicated. The formula is hardcoded in the program and every node does the same calculation to know what difficulty is required for the next block. If someone diverged from the formula, their block would not be accepted by the majority.

506.

507.

508.

509. Thanks once again. I have yet more questions but this is enough for

510. one email :) I will be happy to summarize these discussions into an

511. FAQ-like document at some point. Apologies if the questions seem

512. trivial.

513.

514.  
515. No problem, thanks for testing it on Mac Wine.  
516.  
517. Satoshi  
518.  
519. -----  
520. From: Satoshi Nakamoto <satoshin@gmx.com>  
521. Date: Mon, Apr 13, 2009 at 11:11 PM  
522. To: Mike Hearn <mike@plan99.net>  
523.  
524.  
525. It keeps getting extended all the time. If it stopped, an attacker would have time to catch up. Don't worry, empty blocks aren't very big.  
526.  
527. As you say, it's the order of events that matters.  
528.  
529. -----  
530. From: Mike Hearn <mike@plan99.net>  
531. Date: Mon, Apr 13, 2009 at 11:18 PM  
532. To: Satoshi Nakamoto <satoshin@gmx.com>  
533.  
534.  
535. Oh yes, of course, that's fundamental. Silly me. Thanks for your  
536. answers. I'd recommend being over-explicit for early versions of the  
537. software, something like "Global chain is currently %d blocks long".  
538.  
539. I guess the key problem right now is that once you generate coins,  
540. there's nobody to test it with, even for dummy transactions. Is there  
541. a plan for a mailing list or some kind of trivial marketplace to give  
542. people something to do with their newly minted bitcoins?  
543.  
544. -----  
545. From: Satoshi Nakamoto <satoshin@gmx.com>  
546. Date: Tue, Apr 14, 2009 at 7:41 PM  
547. To: Mike Hearn <mike@plan99.net>  
548.  
549.  
550. I started implementing a marketplace feature earlier that facilitates offering things for sale and taking orders, it's only half done though. A bit like e-bay but without auctions, just "buy now". Among other things, it would make it easy for anyone to offer currency exchange.  
551.  
552. If you send to 1PhUXucRd8FzQved2KGK3g1eKfTHPGjgFu and e-mail me your bitcoin address, or IP if you can accept incoming connections, I'll send back the same amount +50.  
553.  
554. -----  
555. From: Mike Hearn <mike@plan99.net>  
556. Date: Sat, Apr 18, 2009 at 3:08 PM  
557. To: Satoshi Nakamoto <satoshin@gmx.com>  
558.  
559.  
560. Hi Satoshi,  
561.  
562. I sent you 32.51 coins, my bitcoin address is 1JuEjh9znXwqsy5RrnKqgzqY4Ldg7rnj5n  
563.

564. My IP is currently 84.73.233.199, however, it's a laptop so may or may  
565. not be online at the time you act on this mail. I suggest using the  
566. bitcoin address instead. It'd be convenient if the same comment  
567. functionality was available via indirect transfer. Can the comment be  
568. encrypted using the public key of the receiver and placed into a  
569. block?  
570.  
571. -----  
572. From: Satoshi Nakamoto <satoshin@gmx.com>  
573. Date: Sat, Apr 18, 2009 at 6:16 PM  
574. To: Mike Hearn <mike@plan99.net>  
575.  
576.  
577. I sent back 32.51 and 50.00.  
578.  
579. I badly wanted to find some way to include a comment with indirect transfers, but there just wasn't a  
way to do it. Bitcoin uses EC-DSA, which was essential for making the block chain compact enough to be  
practical with today's technology because its signatures are an order of magnitude smaller than RSA. But  
EC-DSA can't encrypt messages like RSA, it can only be used to verify signatures.  
580.  
581. -----  
582. From: Mike Hearn <mike@plan99.net>  
583. Date: Sat, Apr 18, 2009 at 9:25 PM  
584. To: Satoshi Nakamoto <satoshin@gmx.com>  
585.  
586.  
587. Thanks. I sent you back 50, so now we're even.  
588.  
589. For some reason your transfer to me shows up as "From: unknown" even  
590. though I added you to my address book.  
591.  
592. I have a "Generated (not accepted)" line in my transaction list, it  
593. seems like an attempt to generate a coin went wrong somehow. Not sure  
594. what happened here - presumably my node successfully solved a block  
595. but then I went offline before it was sent to the network?  
596.  
597. I suppose for sending metadata with a transaction some other mechanism  
598. will be needed, for instance, broadcast of encrypted messages  
599. associated with a transaction that persist for (say) a month, with  
600. some kind of budget on how much storage a node can use for messages.  
601. Alternatively, a payee could generate some reference number which is  
602. of some significance to themselves but otherwise opaque, and give it  
603. to the payer, thus it does not need to be encrypted and can be put  
604. into the block directly.  
605.  
606. -----  
607. From: Satoshi Nakamoto <satoshin@gmx.com>  
608. Date: Sat, Apr 18, 2009 at 10:52 PM  
609. To: Mike Hearn <mike@plan99.net>  
610.  
611.  
612. Got the 50.  
613.

614. Transactions sent to a bitcoin address will always say "from: unknown". The transaction only tells who it's to. Sending by bitcoin address has a number of problems, but it's so nice having the fallback option to be able to send to anyone whether they're online or not. There are a number of ideas to try to improve things later. For now, if things work out like the real world where the vast majority of transactions are with merchants, they'll pretty much always make sure to set up to receive by IP. The P2P file sharing networks seem fairly successful at getting a large percentage of their users to set up their firewalls to forward a port.

615.

616. The "Generated (not accepted)" normally happens if two nodes find a block at close to the same time, one of them will not be accepted. It's normal and unavoidable. I plan in v0.1.6 to hide those, since they're just confusing and annoying and there's no reason for users to have to see them. While the network is still small like it is now, if you can't receive incoming connections you're at more of a disadvantage because you can't receive block announcements as directly.

617.

618. -----

619. From: Mike Hearn <mike@plan99.net>

620. Date: Sat, Apr 18, 2009 at 11:23 PM

621. To: Satoshi Nakamoto <satoshin@gmx.com>

622.

623. -----

624. From: Satoshi Nakamoto <satoshin@gmx.com>

625. Date: Sun, Apr 19, 2009 at 2:14 AM

626. To: Mike Hearn <mike@plan99.net>

627.

628.

629. The list is:

630. bitcoin-list@lists.sourceforge.net

631. Subscribe/unsubscribe page:

632. <http://lists.sourceforge.net/mailman/listinfo/bitcoin-list>

633. Archives:

634. [http://sourceforge.net/mailarchive/forum.php?forum\\_name=bitcoin-list](http://sourceforge.net/mailarchive/forum.php?forum_name=bitcoin-list)

635.

636. I'll always announce new versions there. Automatic update, or at least notification of new versions, is definitely on the list. There could potentially be necessary changes in the future where nobody will want to talk to you until you upgrade, and there needs to be code in the older version to convey that to the user. This is all the harder in the context of not trusting anyone.

637.

638. Your approach to micropayments sounds right. At first, it might be a good idea to default to asking permission until the user gets comfortable and is ready to set it to automatic. The end goal though should get to something like you describe, where it's similar to using your cell phone without really having to think about the per minute charges.

639.

640. I looked at Google protocol buffers when they were released last year, but I had already written everything by then. What I did was something similar to Boost Serialisation. For this application, where I was parsing messages from strangers who might have extreme incentive to hack the protocol, it was necessary to make it as basic as possible so I could crawl over every line of code to convince myself it was airtight. It became clear that any unnecessary degrees of freedom in the binary format multiplied the potential angles of attack. You guys are so right though to standardize across the company on protocol buffers. I think you've got the optimal solution in the general case.