

# Hashcash - Amortizable Publicly Auditable Cost-Functions

Adam Back  
e-mail: adam@cypherspace.org

1st August 2002

## Abstract

We present a distributed efficiently amortizable CPU cost-function with no trap-door. The absence of a trap-door allows us to avoid needing to trust any party.

Applications for such cost-functions are in distributed document popularity estimation, and metering of web advertising. None of the servers involved have any advantage over users in computing the cost function.

The amortized token has a small fixed sized representation independent of the amortized value. The valuation function is efficient.

Limits can be placed on the resources which can be expended in computing the cost-function to prevent the user inflating the value of his contribution by using more CPU time than the expected token value.

The amortized part of the cost result can be blinded so that clients can not obtain service but avoid contributing to the amortization by expending more resources than the expected token value.

Interactive and non-interactive variants of the cost-function can be constructed which can be used in situations where the server can issue a challenge (connection oriented interactive protocol), and where it can't (where the communication is store-and-forward, or packet oriented) respectively.

**Key Words:** hashcash, cost-functions, client puzzles, distributed metering, distributed popularity estimation

## 1 Introduction

CPU cost-functions have been proposed by Dwork and Naor in [1], and by Back in [2] to limit systematic abuse of unmetered internet resources such as Internet email, and anonymous remailers. Franklin and Malkhi in [3] propose a (non-amortizable) distributed web usage metering scheme based on incremental cost-functions. Naor and Pinkas in [4] explore the use of secret-sharing schemes to perform secure and accurate web metering.

## 2 Cost-Functions

A *cost-function* should be computationally cheap to verify, but parameterisably expensive to compute. We use the following notation to define a cost-function.

In the context of cost-functions we use *client* to refer to the user who must compute a *token* (usually denoted  $\mathcal{T}$ ) using a cost-function  $\text{MINT}()$  which is used to create tokens to participate in a protocol with a *server*. We use the term *mint* for the cost-function because of the analogy between creating cost tokens and minting physical money.

The server will check the value of the token using an evaluation function  $\text{VALUE}()$ , and only proceed with the protocol if the token has the expected value.

The functions are parameterised by the amount of work  $w$  that the user will have to expend on average to mint a token.

With *interactive cost-functions*, the server issues a challenge  $c$  to the client – the server uses the  $\text{CHAL}()$  function to compute the challenge. (The challenge function is also parameterised by the work factor.)

$$\begin{cases} \mathcal{C} = \text{CHAL}(w) & \text{server challenge function} \\ \mathcal{T} = \text{MINT}(\mathcal{C}) & \text{mint token based on challenge} \\ \mathcal{V} = \text{VALUE}(\mathcal{T}) & \text{token evaluation function} \end{cases}$$

With *non-interactive cost-functions* the client choses it's own challenge or random start value in the MINT() function, and there is no CHAL() function.

$$\begin{cases} \mathcal{T} = \text{MINT}(w) & \text{mint token} \\ \mathcal{V} = \text{VALUE}(\mathcal{T}) & \text{token evaluation function} \end{cases}$$

Clearly a *non-interactive cost-function* can be used in an interactive setting, whereas the converse is not possible.

## 2.1 Known solution

- A **known solution** cost-function is a cost-function where the server knows or can cheaply compute the result of the challenge it issues.

Where we can rely on an interactive protocol, the challenger can issue the client with a challenge with a single known solution. A simple way to create known solution cost-functions is to base them on a trap-door problem such as the factoring problem.

$$\left\{ \begin{array}{ll} \text{PUBLIC:} & n = pq \\ \text{PRIVATE:} & \text{primes } p \text{ and } q, \phi(n) = (p-1)(q-1) \\ \\ \mathcal{C} = \text{CHAL}(w) & \text{choose } r \in_R [0, n) \\ & \text{return } (r, w) \\ \\ \mathcal{T} = \text{MINT}(\mathcal{C}) & \text{compute } x = r^{r^w} \pmod{n} \\ & \text{return } (x, w) \\ \\ \mathcal{V} = \text{VALUE}(\mathcal{T}) & \text{if } r^{r^w} = x \pmod{n} \text{ return } w \\ & \text{else return } 0 \end{array} \right.$$

The client does not know  $\phi(n)$ , and so the most efficient method for the client to calculate MINT() is repeated exponentiation, which requires  $w$  exponentiations. The challenger knows  $\phi(n)$  which allows a more efficient computation of  $\mathcal{T} = r^{r^w} \pmod{n}$ , namely reducing  $r^w \pmod{\phi(n)}$ , so the challenger can execute VALUE() with 2 modular exponentiations. The challenger as a side-effect has a trapdoor in computing the cost-function as he can compute MINT() efficiently using the same algorithm. (This cost-function is related to a construct Rivest et al used for *time-lock puzzles* in [5].)

## 2.2 Publicly Auditable, Probablistic Cost

- A **publicly auditable** cost-function can be *efficiently* verified by any third party without access to any trapdoor or secret information. (When we say *publicly auditable* we mean implicitly that the cost-function is *efficiently* publicly auditable, all cost-functions are publicly auditable by definition, as the auditor can just repeat the work done by the client.)
- A **fixed cost** cost-function takes a fixed amount of resources to compute. The fastest algorithm to mint a fixed cost token is a deterministic algorithm. The factoring based cost-function given above is *fixed cost*.
- A **probablistic cost** cost-function is one where the cost to the client of minting a token has a predictable expected time, but a random actual time as the client can most efficiently compute the cost-function by starting at a random start value. Sometimes the client will get lucky and start very close to the solution.

There are two types of probablistic cost *bounded probablistic cost* and *unbounded probablistic cost*.

- An **unbounded probablistic cost** cost-function, can in theory take forever to compute, though the probability of taking significantly longer than expected tends to 0. (An example would be the cost-function of being required to throw a head with a fair coin; in theory the user could be unlucky and end up throwing thousands of tails, but in practice the probability of not throwing a head for  $k$  throws tends towards 0 rapidly as  $\lim_{k \rightarrow \infty} \frac{1}{2}^k = 0$ .)

- With a *bounded probabilistic cost* cost-function there is a limit to how unlucky the client can be in it's search for the solution; for example where the client is expected to search some key space for a known solution; the size of the key space imposes an upper bound on the cost of finding the solution.

Juels and Brainard in [6] give a cost function they call a *client puzzle*.

Client puzzles are interactive, known solution, trapdoor, cost functions with bounded probabilistic cost. Client puzzles are *not* publicly auditable, because private information is required to verify them. (Note the secret key material used in the client puzzle Juels and Brainard propose is introduced only as an optimization to avoid needing to keep state. If this optimization were removed the resulting cost-function would be *publicly auditable*.)

Client puzzles retain the trapdoor property, in that the server can cheaply mint tokens, although the ability to cheaply mint tokens does not derive from private key material. The trapdoor property derives from the fact that client puzzles are interactive, known solution cost functions: the server knows the solution to the puzzle at the time it creates the puzzle.

First we introduce some notation: consider bitstring  $s = \{0, 1\}^*$ , we define  $[s]_i$  to mean the bit at offset  $i$ , where  $[s]_1$  is the left-most bit, and  $[s]_{|s|}$  is the right-most bit.  $[s]_{i\dots j}$  means the bit-wise substring between and including bits  $i$  and  $j$ , so that  $[s]_{i\dots j} = [s]_i \parallel \dots \parallel [s]_j$  and  $s = [s]_{1\dots |s|}$ .

The client puzzle cost-function is:

PUBLIC:	hash function $\mathcal{H}(\cdot)$ with output size $k$ bits
PRIVATE:	server seed key $s$
$\mathcal{C} = \text{CHAL}(w)$	<b>choose</b> $r \in_R \{0, 1\}^k$ <b>compute</b> $x = \mathcal{H}(s \parallel r \parallel w)$ <b>compute</b> $y = \mathcal{H}(x)$ <b>set</b> $p = [x]_{w+1\dots k}$ <b>return</b> $(p, y, r, w)$
$\mathcal{T} = \text{MINT}(\mathcal{C})$	<b>find</b> $z \in_R \{0, 1\}^w$ <b>st</b> $y = \mathcal{H}(z \parallel p)$ <b>set</b> $x' = z \parallel p$ <b>return</b> $(x', y, r, w)$
$\mathcal{V} = \text{VALUE}(\mathcal{T})$	<b>if</b> $y = \mathcal{H}(x')$ <b>and</b> $x' = \mathcal{H}(s \parallel r \parallel w)$ <b>return</b> $w$ <b>else return</b> $0$

Juels and Brainard also give a method for reducing the *variance* of the bounded probabilistic cost of their client puzzles. The technique is to compose each puzzle of a number of sub puzzles such that the expected cost remains the same. This property can be useful in interactive settings where you don't want the user to experience large variance in latency.

## 2.3 Trapdoor-free

A disadvantage of known solution cost-functions is that the challenger can cheaply create tokens of arbitrary value. This precludes public audit where the server may have a conflict of interests, for example in web hit metering, where the server may have an interest to inflate the number of hits on it's page where it is being paid per hit by an advertiser.

- A *trapdoor-free* cost-function is one where the server has no advantage in minting tokens.

An example of a trapdoor-free cost-function is the Hashcash [2] cost-function.

Hashcash is a non-interactive, publicly auditable, trapdoor-free cost function with unbounded probabilistic cost.

We define a pair of binary infix comparison operators  $\stackrel{\text{left}}{=} b$  where  $b$  is the length of the common left-substring from the two bit-strings, and  $\stackrel{\text{right}}{=} b$  is the length of the common right-substring.

$$\begin{aligned}
 x \stackrel{\text{left}}{=}_0 y & \quad [x]_1 \neq [y]_1 \\
 x \stackrel{\text{left}}{=} b y & \quad \forall_{i=1\dots b} [x]_i = [y]_i \\
 \\ 
 x \stackrel{\text{right}}{=}_0 y & \quad [x]_{|x|} \neq [y]_{|y|} \\
 x \stackrel{\text{right}}{=} b y & \quad \forall_{i=1\dots b} [x]_{|x|+1-i} = [y]_{|y|+1-i}
 \end{aligned}$$

Hashcash is computed relative to a service-name  $s$ , to prevent tokens minted for one server being used on another (servers only accept tokens minted using their own service-name). The service-name can be any bit-string which uniquely identifies the service (eg. host name, or email address).

The hashcash function is defined as:

$$\left\{ \begin{array}{ll} \text{PUBLIC:} & \text{hash function } \mathcal{H}(\cdot) \text{ with output size } k \text{ bits} \\ \mathcal{T} \leftarrow \text{MINT}(s, w) & \mathbf{find } x \in_R \{0, 1\}^* \text{ st } \mathcal{H}(s\|x) \stackrel{\text{left}}{=}_w 0^k \\ & \mathbf{return } (s, x) \\ \mathcal{V} \leftarrow \text{VALUE}(\mathcal{T}) & \mathcal{H}(s\|x) \stackrel{\text{left}}{=}_v 0^k \\ & \mathbf{return } v \end{array} \right.$$

The hashcash cost-function is based on finding *partial hash collisions* with the all-zero bit-string  $0^k$ . The fastest algorithm for computing partial collisions is brute force. There is no challenge as the client can safely choose his own random challenge, and so the hashcash cost-function is a *trapdoor-free* and *non-interactive* cost-function.

In addition the Hashcash cost-function is *publicly auditable*, because anyone can efficiently verify any published tokens.

(In practice  $|x|$  should be chosen to be large enough to make the probability that clients collectively reuse a previously used start value negligible;  $|x| = 128$  bits should be enough even for a busy server.)

The server needs to keep a double spending database of spent tokens, to detect and reject attempts to spend the same token again. To prevent the database growing indefinitely, the service string can include the time at which it was minted. This allows the server to discard entries from the spent database after they have expired. Some reasonable expiry period should be chosen to take account of clock inaccuracy, computation time, and transmission delays.

Hashcash was originally proposed as a counter-measure against spam in email, and against systematic abuse of anonymous remailers. It is necessary to use *non-interactive cost-functions* for these scenarios as there is no channel for the server to send a challenge over. However one advantage of *interactive cost-functions* is that it is possible to prevent pre-computation attacks. For example, there is a cost associated with sending each email, which tends to discourage email abuse; however a determined adversary may spend a year pre-computing tokens to all be valid on the same day, and on that day be able to overload the system.

It would be possible to mitigate the pre-computation attack by using a slowly changing *beacon* (unpredictable broadcast authenticated value changing over time) such as say this weeks winning lottery numbers. In this event the public value is included in the start string, limiting pre-computation attacks to being conducted within the time period between beacon value changes.

## 2.4 Amortizable

- A *probabilisticly amortizable* token is one where there exists an efficient algorithm to combine tokens into a single token with workably compact representation with value which probabilisticly approximates the combined value. Repeatedly combining the same tokens again or combining the token with itself should not increase it's value.
- A *publicly amortizable* token is one where the amortization function does not require any private information.
- A *publicly auditable amortizable* token is one where the amortized token's value is efficiently verifiable by anyone without any private information.

*Non-interactive Amortizable Hashcash* is a non-interactive, publicly auditable, trapdoor-free, unbounded probabilistic cost, probabilisticly amortizable cost-function.

The amortizable hashcash function is defined:

$$\left\{ \begin{array}{ll} \text{PUBLIC:} & \text{hash function } \mathcal{H}(\cdot) \text{ with output size } k \text{ bits} \\ \mathcal{T} = \text{MINT}(s, w) & \text{find } x \in_R \{0, 1\}^* \text{ st } \mathcal{H}(s\|x) \stackrel{\text{left}}{\equiv}_w 0^k \\ & \text{return } (s, x) \\ \mathcal{V} = \text{VALUE}(\mathcal{T}) & \mathcal{H}(s\|x) \stackrel{\text{left}}{\equiv}_v 0^k \\ & \text{return } (v) \\ \mathcal{T} = \mathcal{T}_1 + \mathcal{T}_2 & \text{return MAX}(\mathcal{T}_1, \mathcal{T}_2) \end{array} \right.$$

For applications such as web metering, the Non-interactive Amortizable Hashcash protocol has a weakness in that the protocol allows forms of cheating called *under-contributing* and *over-contributing* by clients. With *under-contributing* the client spends some extra computation to ensure that he minimally complies with the token value required, to avoid contributing to the amortized value.

The converse client action of *over-contributing* may sometimes be considered an attack, depending on the use. For web metering, where the amortized token is intended to measure hits served by a web hosting service, over-contributing could be a problem, as users have the ability to inflate the measurement of web-hits. Over-contributing as such can not be prevented, as the user can simply request the document multiple times. But there remains some value to defending against over-contribution, as the user has to download the whole document which is an additional cost for the user (and the hosting service), and the server can to some extent detect and refuse multiple downloads by the same user.

For distributed document popularity where the estimate of hits served is not relevant, over-contributing may be encouraged as a more bandwidth efficient way for users to rate documents.

- A *fair amortizable* cost-function is one where the client can not *under-contribute*. *Over-contributing* may optionally be allowed.

To simultaneously achieve both document popularity metric collection and amortizable document hit statistics we would use a pair of tokens: a non-interactive amortizable token for document popularity, and a fair amortizable token for hit statistics. Both tokens would be independently amortized.

*Interactive Fair Amortizable Hashcash* is an interactive, publicly auditable, trapdoor-free, unbounded probabilistic cost, fair probabilistically amortizable cost-function.

The Interactive Fair Amortizable Hashcash function is defined below.

$$\left\{ \begin{array}{ll} \text{PUBLIC:} & \text{hash function } \mathcal{H}(\cdot) \text{ with output size } k \text{ bits} \\ \mathcal{C} = \text{CHAL}(s, w) & \text{choose } b \in_R \{0, 1\}^* \\ & \text{compute } b' = \mathcal{H}(s\|b) \\ & \text{return } (s, b', w) \\ \mathcal{T}' = \text{MINT}(\mathcal{C}) & \text{find } x \in_R \{0, 1\}^* \text{ st } b' \stackrel{\text{left}}{\equiv}_w \mathcal{H}(s\|b'\|x) \\ & \text{return } (s, b', w, x) \\ \mathcal{T} = \text{UNBLIND}(\mathcal{T}') & \text{return } (s, b, w, x) \\ \mathcal{V} = \text{VALUE}(\mathcal{T}) & \text{compute } b' = \mathcal{H}(s\|b) \\ & \text{if } b' \stackrel{\text{left}}{\neq}_w \mathcal{H}(s\|b'\|x) \text{ return } 0 \\ & \mathcal{H}(s\|b'\|x) \stackrel{\text{right}}{\equiv}_u b \\ & \text{return } (w + u) \\ \mathcal{T} = \mathcal{T}_1 + \mathcal{T}_2 & \text{return MAX}(\mathcal{T}_1, \mathcal{T}_2) \end{array} \right.$$

(To admit over-contribution, the VALUE() function should be modified to return  $(w' + u)$  where  $w'$  is computed as the actual left-most bits of collision between  $b'$  and  $\mathcal{H}(s\|b'\|x)$ , namely  $w'$  is computed as:  $b' \stackrel{\text{left}}{\equiv}_{w'} \mathcal{H}(s\|b'\|x)$ .)

Some notes on the design criteria for Fair Amortizable Hashcash:

- The cost of the verification function is minimised. The verification function costs  $2 \mathcal{H}()$  operations.

- The server should not be able to trick users into computing tokens on other service-names.  $s$  is included in  $\mathcal{H}(s||b'||x)$  for this reason even though the service-name  $s$  will later become apparent when  $b'$  is unblinded as  $b' = \mathcal{H}(s||b)$ .
- There should be no computational advantage for the server, so to obtain a token of value  $v = w + u$  the server should be forced to do an average of  $2^{w+u}$  trial computations. The  $'b = \mathcal{H}(s||b), a = \mathcal{H}(s||b'||x), b' \stackrel{\text{left}}{\equiv}_w a \stackrel{\text{right}}{\equiv}_u b$  construct ensures this. Changing any of the values  $s, b$  or  $x$  ( $b'$  is changed by changing  $b$ ) in any of the corresponding expressions for calculating  $b', a$  or  $b$  has side effects on at minimum one of the other expressions. This property prevents the server finding a pair of collisions and then varying an independent parameter in the final expression calculation to match, which would admit a  $2^w + 2^u < 2^{w+u}$  work factor attack.
- The parameters  $s, b$  and  $b'$  are all hashed before or during being used to find a collision. This ensures that similar service-names (service-names which are pre-fixes of other service-names), and similar blinding values don't offer the attacker any reduction in work factor by combining his attempts to create tokens for multiple services.
- Note the hash function should not admit any significant advantage in computing the hash operation by choosing a stride to match the hash operation internal structure (ie rather than choosing  $x_0 \in_R \{0, 1\}^*$  and  $x_{i+1} = x_i + 1$ , use recurrence relation  $x_{i+1} = f(x_i)$  with  $f()$  chosen to reuse internal computation steps. Similarly no significant advantage should be achieved by computing a batch of candidate  $x_i$  values in parallel.

### 3 Improving Accuracy

Clearly the metering estimate afforded by using the hashcash amortizable cost-function is probabilistic and can be quite inaccurate. This problem can be combatted by instead of keeping a single highest valued token for each document, keeping the  $n$  highest valued tokens.

## References

- [1] Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In *Proceedings of Crypto*, 1992. Also available as [http://www.wisdom.weizmann.ac.il:81/Dienst/UI/2.0/Describe/ncstrl.weizmann\\_il/CS95-20](http://www.wisdom.weizmann.ac.il:81/Dienst/UI/2.0/Describe/ncstrl.weizmann_il/CS95-20).
- [2] Adam Back. Hashcash, May 1997. Published at <http://www.cypherspace.org/hashcash/>.
- [3] Matt Franklin and Dalia Malkhi. Auditable metering with lightweight security. In *Financial Cryptography*, pages 151–160, 1997. Also available as <http://www.parc.xerox.com/csl/members/franklin/>.
- [4] Moni Naor and Benny Pinkas. Secure and efficient metering. *Lecture Notes in Computer Science*, 1403, 1998. Also available as <http://citeseer.nj.nec.com/naor98secure.html>.
- [5] Ronald L Rivest, Adi Shamir, and David A Wagner. Time-lock puzzles and timed-release crypto. Technical Report MIT/LCS/TR-684, 1996. Also available as <http://theory.lcs.mit.edu/~rivest/publications.html>.
- [6] Ari Juels and John Brainard. Client puzzles: A cryptographic countermeasure against connection depletion attacks. In *Network and Distributed System Security Symposium*, 1999. Also available as <http://www.rsasecurity.com/rsalabs/staff/bios/ajuels/publications/client-puzzles/>.