FOUNDATIONS OF COMPUTING AND DECISION SCIENCES
Vol. 42 (2017) No. 1
DE GRUYTER OPEN
DOI: 10.1515/fcds-2017-0001
ISSN 0867-6356
e-ISSN 2300-3405

# BROWSER-BASED HARNESSING OF VOLUNTARY COMPUTATIONAL POWER

Tomasz FABISIAK, Arkadiusz DANILECKI*

**Abstract.** Computers connected to internet represent an immense computing power, mostly unused by their owners. One way to utilize this public resource is via world wide web, where users can share their resources using nothing more except their browsers. We survey the techniques employing the idea of browser-based voluntary computing (BBVC), discuss their commonalities, recognize recurring problems and their solutions and finally we describe a prototype implementation aiming at efficient mining of voluntary-contributed computing power.

**Keywords:** voluntary computing, desktop grids, browser-based computing, internet computing, survey, crowdsourced supercomputer, peer-to-peer computing

## 1. Introduction

With regard to a computing power, modern personal computers are comparable to supercomputers from two decades ago. In 1996, the last supercomputers in the list achieved between 4.6 to 6.5 Gflops (depending on the measure used)[1] – a performance easily matched by laptops such as Lenovo Y50, available for less than 1000$. By even the most conservative estimates there are more than one billion personal computers currently in use[2],

* Institute of Computing Science, Poznań University of Technology

1 http://www.top500.org/list/1996/11/?page=5

2 http://www.worldometers.info/computers/. Some estimates give a number of more than two billions personal computers in use worldwide.

and their solutions in Section 4. Our secondary contribution is contained in Section 5 where we describe our work-in-progress: a new BBVC platform PovoCop, its proof-of-concept prototype implementation and some initial experimental results. We discuss non-browser-based voluntary computing in Related Work (Section 6). Section 7 concludes the paper.

## 2. Characteristic features of browser-based voluntary computing

The idea of utilizing idle cycles of volunteering users is associated with many labels, denoting same or at least heavily overlapping concepts. We encountered names such as peer-to-peer computing [28][52], global computing [28][62], public resource computing [3], internet computing [75][52], crowdsourced supercomputer [79], desktop grid [75] and cycle stealing [28]. The name most commonly used seems to be "volunteer computing" or "voluntary computing" [13][75][4][26]. As mentioned earlier, in this paper we concentrate only on browser-based solutions, i.e. on a subset of voluntary computing, Henceforth we will use the name "browser-based voluntary computing" (BBVC).

Browser-based voluntary computing is a subset of distributed computing, differentiated by several characteristics of the environment.

First, there are two kind of the nodes in the system: the generally long-living servers and the machines controlled by volunteers – providing resources we want to use. Obviously, in BBVC we assume that every volunteer has a browser installed. Browsers provide execution environment common for the applications running in the system.

The second characteristics is that, in contrast to e.g. grids and clusters where work is "pushed" from server to slave nodes, in voluntary computing the work must be actively pulled from server by workers [75]. The workers must find the server, inquire about the existing tasks, choose some of them for download and execution, and have to voluntarily upload the results to the server at time suitable for the worker. The fact that the worker volunteered for some task does not guarantee that it will actually complete the task.

The third distinguishing characteristic is that resources used in voluntary computing are ephemeral and unreliable in nature, and are not managed nor authenticated [26][75]. Especially in BBVC, users may connect and disconnect at will, and may (or might not) reappear later. The interaction between users and the web sites may last as short as ten seconds and is rarely longer than few minutes [14]. Resource characteristics often are unknown (i.e. it is not known in advance what computing power will be available for use) and, even if they were known, they could quickly become obsolete (as computers are in use and may occasionally experience bursts of activitity).

The distinguishing characteristics of computing environment pose a set of challenges inherent in voluntary computing in general [28][69], and those are applicable also for BBVC: accessibility, programmability, fault tolerance, security, scalability, usability, availability, heterogenity, adaptability/dynamicity.

*Accessibility* means that the proposed solution must be easily accessible to users irrespective of their machines' architectures and capabilities. Users should require no technical knowledge to share their resources. Projects should be attractive to the users irrespective of their demographic characteristics.

Therefore, the ideal jobs for BBVC should be easily divisible into small, computationally independent tasks – or, at least, if the computation runs in phases, tasks should be computationally autonomous within each phase. The general problem categories considered by researchers working on BBVC were called piecework computations [16], master-worker [26], coarse-grained [22], malleable applications [19] or bags of tasks [23].

During our survey we encountered the following problems identified as well-suited for BBVC: evolutionary/genetic algorithms [23][25][44][58][64], simulated annealing [23][64], colony optimization [23], particle swarm optimization [23], artificial bee colonies [23], brute force search [64], cryptography [9][64][80], game-tree evaluation [64], map-reduce problems with computation-intensive map phase [22][47][50][57][67][79], monte-carlo simulations [19][54][55] branch-and-bound [16] and its limited subset, bulk-synchronous-parallel (BSP) [68], and, finally, the algorithms based on directed acyclic graphs [19].

BBVC could be used in social media to facilitate image tagging for improving users' browsing experience or reducing load on servers. Hence, the following problems were suggested: face recognition [80][83], logo recognition [83], image scaling and sentiment analysis [80]. Similarly, it was suggested that near similarity detection can be used for detection of copyrighted media [83], and images can be classified by neural networks [57].

During experimental evaluation of the surveyed proposals few problems were regularly reappearing. Raytracing was very popular, starting with one of the first applet-based approach (distributed applet-based massive parallel processing, DAMPP) from 1997 [77]. Other work which suggested raytracing as example problem are [16][23][71], though the author of [71] admitted his test results were not encouraging.

Other problems which in reviewed works were used as benchmarks ranged from the very simple, like searching for prime numbers [22] and large internet number 1234567890123456789 factorization [68], to genome analysis [79] and protein sequence alignments using Needleman-Wunsch algorithm. [19]. Other problems used in experiments were: DES code breaking [10], matrix multiplication [16][68], Traveling Salesman Problem [16][24][67], Mandelbrot generation [69][70], searching for Mersenne primes using Lucas-Lehmer primality test [56][29], Collatz hypothesis verification [56], the factorization of large integer into two primes and Pearson correlation evaluation on set of genetic samples [9].

For those problems linear or even superlinear speedup [16][57] was often achieved, with experiments employing from 80 [67] through 200 [47] up to 400 clients [22]. A number of researchers [23][47][54][55][57][71] noticed the problem of “speedup decay”: adding more workers caused speedup to level off or even drop. With more workers, each of them received less work, meaning communication costs imposed by work distribution overshadowed the gains. In solutions using centralized architecture, the central server could become overloaded with the increased number of workers [57][67]. Finally, the bandwidth depletion was suggested as another explanation of the phenomenon [47].

After presenting the problems well-suited for BBVC and discussing the characteristics of the environment and the challenges associated with voluntary computing in general, we are now ready to review the browser-based proposals which appeared during the last two decades.

### 3.1. The first generation: Dead applets society

In the first-generation of browser-based voluntary computing proposals, jobs were implemented as Java applets (called “distriblets” in [29] and “computelets” in [64]). Typically, the applications were listed at a website. User had to choose which computation to run, which usually required him or her to click on the link and download an applet (e.g. [6][7][29]). The applets either contained the necessary input data or connected to a server requesting work (possibly using UDP instead of HTTP as in DAMPP [77]). After finishing the work, the applet would send back the results and ask for another task to execute. Some solutions allowed offline computation: in [29] and [16] a user could download applet, disconnect, and when later he or she would reconnect; the applet would send back the computed results.

To the best of our knowledge, the first implementation of the idea was Charlotte [7], appearing in 1996. The applications had to be written especially for Charlotte and had limited scalability (as its intended use was limited to the authors' home university). Charlotte provided Distributed Shared Memory abstraction, via wrapper classes of common types and explicit “set” and “get” function calls. Programs had to be registered at the project's website. User visiting the site could click one of the listed URLs, loading and executing “worker” applets. In addition, a standalone Java programs, *manager*, resided on a web server's site, one for each distributed application. The web was used only for coordination between clients and job creators. Charlotte programs could indirectly communicate via *managers*. In their followup work [6] they described a KnittingFactory directory service, implemented as JavaScript-enriched HTML page. JavaScript constructed new URL, with search state passed as a part of the created URL, and instructed the browser to follow the URL. Searching took place solely on users' browsers.

We will now describe three first-generation solutions in more detail. Two of them were chosen because they are the most known in the literature and may be considered the early paragons of browser-based solution: Javelin [16] and Bayanihan [68]. In addition, we decided to discuss POPCORN [64] because of its interesting economy-based interaction between job owners and worker nodes.

*Javelin* is a prototype implementation of the SuperWeb architecture [2]. The architecture consists from brokers, hosts and clients and, in general, does not depend upon Java/browser combination. *Client* create Java programs, applets, which they then register at *broker*. Applets belonging to the same application could communicate via broker. Programs are not sent to broker; only their URL addresses are sent. During registration clients estimate how much resources will be used by their tasks.

*Hosts* (i.e. worker nodes) were the machines willing to provide their resources (CPU power, storage, memory) to clients. Hosts had to register at brokers. Hosts could specify when their resources may be used and what percentage of a given resource would be available for use. Their SuperWeb architecture demanded the hosts to provide a sandboxed environment in which tasks would be safely executed. This was their primary motivation for choosing applets and browsers, as they obviously provide the required safe environment.

Brokers coordinated between clients and hosts. They ran accounting, giving credits to users. Those credits were envisaged to be later exchanged for “cyber money, access to software licenses, cheap internet access or free cpu time”. When host registered at broker, it received URL of jobs submitted by clients. Host downloaded an applet from a specified

Clients could be applets or standalone Java applications. Each client was a chassis, containing an engine. Watcher clients used watch engine, which contacted a watch manager at the server side in order to investigate problem's statistics, computation progress etc. Worker clients would use work engine instead.

Workers had to invoke a remote method of a worker advocate, which then communicated with worker manager and sent back a work object. Worker client then called *doWork* method of the received work object. When the computation was finished, it called *sendDone* remote method of the work advocate. Result was then put into a result pool. After work pool was empty, worker could request another bunch of work objects; if a programmer wanted worker to to wait and synchronize with other workers, all he or she had to do was to reimplement the appropriate interface methods.

The most interesting feature of *POPCORN* [64] was its "capitalist" approach, visible even in the chosen nomenclature. The parties seeking for the resources were called *buyers*, machines voluntary contributing their resources were named *sellers*, and the server allowing their cooperation was called *market*.

Applets in POPCORN contained both the code and the data and were called *computelets*. Computelets were sent by buyers into the market. Sellers registered at the market, received computelets, executed them, and sent back the results to the market, which forwarded them to the buyer.

The jobs – *computation packets* – contained also information about their pricing and information how to handle different events, such as receiving results from sellers or premature termination. POPCORN provided rudimentary fault tolerance: it guaranteed that for every job, eventually either its *completed* or *failed* method would be called. Computelets losses were detected via timeouts. Programmers were responsible for result verification. There were no direct communication between the computelets.

In POPCORN, sellers were selling "java operations". Computelets invoke periodically high-priority thread – basically a microbenchmark – which computed few operations and reported the performance back to the marker. Sellers were paid by "popcoin". Each user had its account, and could spent popcoins on resources, or earn them by renting his or hers resources to other users. In theory those popcoins could be converted to real money.

Buyers offered their price for resources when advertising their computelets. Whenever a user wanted to sell his CPU time, he visited the page with an applet and clicked the appropriate "execute" button. There might be many potential buyers for his CPU time. The highest bidders were chosen, selected by vickrey auction (where buyers could specify their lowest/highest prices for the resources they want).

Additionally, the authors suggested that some users may become "publishers": they could create web pages with animated POPCORN logo. Visitors to this page inadvertently would execute applets executing computelets. The publishers would receive popcoins, hopefully trading interesting web content for visitors' CPU power.

The idea of brokers earning a profit from mediating between hosts selling their CPU power and clients wishing to buy the power was also discussed in *Unicorn* [66], the last BBVC solution based on Java applets we are aware of.

visit a web page, unwittingly running a client-side code: fetching the fitness tests from server, running them within PushScript engine and sending back results to server, and then – after a small delay – asking for new problem to solve. They implemented server-side code mainly in PHP, except Push3 engine which was in C++.

The third solution, *RABC* (Riken AJAX Browser Computing) came from Japan [45]. RABC aimed at creating simple environment SETI@Home-like projects. The distinguishing feature of this proposal was the fact that the kernel agent – JS code embedded within a HTML page – was trying to adapt its behavior (requesting tasks, returning the results) to network bandwidth.

The second proposal we know to use the genetic algorithms appeared in 2008 [58]. They used JSON instead of XML, hence the name *AGAJAJ* (Asynchronous Genetic Algorithm with JavaScript and JSON). AGAJAJ employed one central server. Volunteers visiting the server's page downloaded a code (written JavaScript) which fetched the parameters from the servers, ran several preset number of generation and then sent back the results (best individual according to fitness function) to the server. Server chose the best individual from all workers and sent it back for another round of generation. Eventually clients were shown the best results and were given chance to restart old experiment or start a new one.

The apparent performance problems with JavaScript led to create at least one Flash-based solution, Online Community Grid (OGC) [61]. Adobe Flash speed was comparable to Java applets. OGC envisaged webmasters putting a small flash widget on their sites. Users visiting the site for the first time would be able to choose whether to participate in one of available projects. Because work would be terminated whenever user would leave the site (by closing a tab, surfing to another page or closing the browser) the running processes were checkpointed every few seconds, in order to allow the work to be resumed when user would later return to the page.

We consider the second-generation to end at 2009, with two proof-of-concept implementations, of distributed hashing in [8] and of map-reduce by a blogger and web developer Ilya Grigorik [34]. In Grigorik's prototype worker node first makes a GET request to job-server, and receives an allocated unit of work with redirection to a prepared site containing mostly JavaScript code to execute. This code would contain typical map-reduce functions and then would emit the results back to the server by creating invisible form, filling it with the results, and submitting it.

Second-generation suffered from the JavaScript's lack of support for threads. For example, in [9] authors discussed at length how to re-engineer the programming loops in order for web pages to being responsive. Their solution was a creation of special function, which run only few iterations of the loop and then rescheduled itself for further future execution.

Another problem was poor performance of browser's JavaScript engines. In 2007, JavaScript was between 9.8 to 23.2 times slower (depending on the platform ) than Java code for a simple mathematical task of random numbers generation [45]. Compared to C, the results were even worse: the browser-based JavaScript engine was between 20 to 200 times slower (depending on platform and browser) [44].

Developers also had to fight with JavaScript same-origin policy (where script from some address domain, e.g. put.poznan.pl, for security reason could not visit URLs from different domain, say github.pl – see Section 4.1). Finally, there had not existed standardized, mature and popular solutions for direct browser-to-browser communication, preventing any direct

Restful, accepting HTTP requests formatted in JSON. Using those requests one can GET job status, GET job or POST a new job to a server. It is implemented as a module of Node.js, with support for five different opensource databases: redis, mongoDB, sqlite, postgresSQL, couchDB. The web server implements presentation layer and serves web pages. Worker nodes have the usual role of executing the tasks.

To create jobs for QMachine, application developers have to reimplement the QM class, with *submit*, *map* and *mapreduce* functions. *Submit* allows a remote execution of single function on worker, while *map* is dedicated method describing invocation of many functions on set of data. Finaly, *mapreduce* function describes map-reduce [21] pattern. The *map* function may have set of input data URLs (in the example experiment, with twenty URLs describing genomes of different streptococcus pneumoniae).

QMachine uses CORS [76] (cross-origin sharing) in order to allow the code to import libraries and the data from third-party servers. The example job described in the paper get jmat library from Google Code, usmlibrary from GitHub and input data (genome for the computational experiment) from NCBI. Volunteers pull the task descriptions and start executing them after downloading the required code and data. Individual workers may download the data from different servers, using potentially different communication links – not just computation is partitioned, but also data transfer bandwidth. Results are sent back to the QMachine server. The job owner polls for the updates in the task status and eventually retrieves the results.

To assist programmers solving concurrency problems with asynchronous data transfer QMachine provides special library (quanah). In addition of JavaScript, they support also CoffeeScript (essentially, a pretty-ish language wrapper for JS).

The authors described the impressive experiment carrying DNA sequence analysis of bacteria genomes. It lasted 12 months, and involved 2100 different IPs from 87 countries issuing 2.2 million API calls to QMachine.

Another typical third-generation solution is MRJS [67], one of the first to use WebWorkers. MRJS is based on map-reduce paradigm. Computation is divided into two phases, requiring modicum synchronization between the workers.

Job owners submit jobs via simple web interface. Input data is divided into uniformly sized parts, called chunks. It is the job owner's responsibility to specify the number of chunk into which input data will be divided.

The workers first receive the task code and map chunks from JobServer with HTTP GET request. Results are sent back via HTTP POST. One chunk may be assigned to many workers, and majority-voting is used for result verification. JobServer aggregates the results from the workers and sends them reduce chunks in the second phase. The second phase may not proceed for a given job until server finishes the aggregation phase.

The input data (map chunks) may be stored on an external server. In such a case, address of input data is contained within a chunk. Because of JavaScript same-origin policy, JobServer must function as a proxy for external data server, forwarding data to workers when necessary. Only after finishing the tests authors discovered CORS standard, which would remove the need for JobServer functioning as a proxy.

*JSMapReduce* created by Langhans et al [50] (based on earlier Langhans bachelor thesis from 2012) shares many similarities to MRJS. In contrast to [67], the chunks are variable-sized, with more capable workers getting more data. Moreover, chunks are compressed

Model training is divided into the iterations. Each worker is told how long it should run, based on estimation of the network latency to the master. New worker nodes must wait for end of the iteration before they may join the computation. Unfortunately, in each iteration master waits for the slowest worker. Master, uploads and allocates new data initializes new workers, handles the reduce step, monitors latency, detects lost worker nodes via timeout and reallocates the data when necessary, and finally broadcasts new parameters.

Krupa et al [47] and their followup work Turek et al [74] proposed to create a search engine consisting of ensemble of browsers and two layers of servers. Browsers are responsible for lexical analysis and parsing, detection of URLs links in texts, processing the content and index building (e.g. of words in parsed pages). Servers distribute work. Each second layer server manages a subset of first layer servers, but in order for the system to be more fault tolerant, the second layer servers also know some randomly chosen servers managed by other second layer servers. That allows failover in case of failure of some second layer server.

In their proposal, search scripts are written first by experts in new formal language and then submitted to second layer servers, which rewrite them into JavaScript, containing definitions of search rules and implementation of the text processing algorithm. JS scripts are sent to the first layer servers. The work is further distributed based on addresses of pages to be searched: second layer servers manage their own range of overlapping URLs, those ranges are then partitioned amongst servers from the first layer, which further partitions them amongst the workers. Work distribution uses the data about workers' performance gathered by the first layer servers. Second layer servers then collect the results: pages meeting search criteria plus newly detected URLs and return them to job owner.

*Comcute* [12] and a followup work by another team, *Comcute.js* [22] are another proof-of-concept implementations. The architecture consists from comcute cores, responsible for starting and stopping the jobs, partitioning the data and aggregating the results, the dispatchers (called also proxies in some papers) controlled by one of the cores, responsible for assigning the jobs to the workers (browsers) and separating the core servers from direct interaction with the workers [10]. Based on a job owner description, cores may automatically divide input data and create an appropriate number of tasks [13][48]. A separate layer provides an entry point for job owners wishing to submit, delete or monitor their jobs. Somewhat surprisingly, in experimental evaluation of *Comcute.js* [22] the centralized architecture with one dispatcher worked the best.

The jobs submitted to comcute consist of data and JS code, and are divided into tasks (having the same code, but different input data). Dispatchers download JSON description of tasks from the core, dispatch the tasks to workers, and then send the results back to the core. Core aggregates the results and monitors the computation, though it's not clear to us whether the monitoring signifies anything more besides tracking the progress of the computation.

*Capataz* [54][55] is a library created to support voluntary computing. The applications must import capataz library and run within Node.js. Application can post a job to the server, which returns them a *Future* object. When workers return the results, the *Future* object is resolved allowing application to continue.

They suggested job bundling, where server sends the same code with many different sets of input data, speculating that browser's just-in-time compiler reuses the JavaScript code in case of job bundling, while it is not reusing the code when the same script is fetched several times in a row.

in Japanese). They used a rarely used extensions, NaCL (Native Client) and PNaCL (Portable Native Client), The former allows native codes to be run by browsers and the latter first requires compilation into LLVM bytecodes, which then can be run at near-native speed. Unfortunately, the future of both extensions is doubtful, as Mozilla has no plans to support them and Google, the original proponent, seems to stop further development[4]. The authors have compiled Gnu Mathematic Library into format compatible with PNaCL, potentially enabling the use of existing scientific code base without the rewriting it into JavaScript. Unfortunately, they've shown that the performance is still up to 9 times slower than the original native codes.

*BrowserCloud.js* [23] is a browser-based computation grid. It uses emscripten to generate very fast JavaScript code, indexedDB for storing data at the browser's side, and WebRTC [27] for direct browser-to-browser communication.

The most distinguishing feature of browserCloud is its decentralized nature. The main servers are used only as a randez-vous points between the browsers. After new clients learns about other clients, it uses WebRTC for direct communication with those clients. New clients register at signaling server, receive unique identifiers. Other peers are then notified about then existence of new peer.

Signaling server holds the network state in its memory, which allows it to generate new unique identifiers. Messages between the peers are routed via Chord routing. Distributed Hash Table (DHT) holds which peer is responsible for routing message destined to a particular peer. In addition, peers have so called “fingers” i.e. direct connections to randomly chosen other peers.

We conclude our survey by briefly mentioning an interesting BBVC project, called crowdprocess.com; it was referenced by [54]. It was supposedly a company which embedded JS code in webmasters websites; unfortunately, the project seems to be inactive, and its main website redirects to james.finance, the site definitely not related to BBVC.

## 4. Issues in browser-based voluntary computing

Browser-based voluntary computing is still relatively unpopular. In this section we identify commonly recurring problems in browser-based voluntary computing which may contribute to its lack of wider acceptance, and how those problems were or could be solved. Since many problems are common to all voluntary computing projects, when discussing the possible solutions we will also sometimes mention non-browser based approaches. In addition, we discuss the fault tolerance issues in Section 4.6 and in Section 4.7 we present the scheduling policies employed in BBVC.

### 4.1. Limitations imposed by browser's security model

The browser provides a sandboxed execution environment for the tasks distributed by the BBVC platform, greatly simplifying security and safety management. Faulty code usually (at least, in absence of errors in browser's code) should cause just unresponsiveness of a

4 http://developer.chrome.com/native-client/nacl-and-pnacl

communication. Applet first passes its reference to a broker (called “initiator”) which passes the RMI reference to all other applets participating in the same session. As described earlier, browserCloud.js [23] uses WebRTC [27] to achieve the same goal. Central server is used only to register browser handles, and message routing is done by the browsers themselves. There exist already libraries such as webGC [11] (based on WebRTC) which could be used to implement direct task-to-task communication. Unfortunately, WebRTC is still not fully supported by all browsers[5].

Lack of widely accepted and universally available mechanism enabling browsers to communicate constrains the design choices for algorithms developed for BBVC, and is one of the factors, even if minor, limiting the wider acceptance of BBVC. Support for direct task-to-task communication is limited in the papers we reviewed.

### 4.3. Recruiting and keeping the volunteers

Voluntary computing, including browser-based voluntary computing, ultimately cannot exist without recruiting users willing to contribute their machines' resources for a public cause. One have to both attract the users to a project and to keep them interested. Since those problems are shared by both BBVC and voluntary computing in general, the offered solutions should be shared too. In this section we do not concentrate on BBVC only and we discuss the papers analyzing the problem from the perspective of voluntary computing in general.

The great advantage of browser-based voluntary computing is its potential to recruit users from a broader audience than is typical for traditional projects. For example, BOINC seem to attract only particular type of users – 96.5% of BOINC users were identified as male, 83% of them were labeled as advanced users. Almost no children, women or elders participated in BOINC projects, which stands in stark contrast to a typical internet user profile [18].

The most basic way to get computing power is to pay users or to force them to participate [68], effectively meaning that the voluntary computing becomes voluntary in name only. New volunteers could be recruited by sending email invitations and blog posts [45][58], social media [19][83] or by being reported upon in popular new sites [79]. To attract new users, a flash widget used by OGC [61] allowed to inspect the news fetched from Yahoo Top Stories feed, and in theory could be used to allow social interaction or playing games. Users could run parasitic games, paying with their CPU time by allowing some computation job running in background [2]. This concept was tested in [50] where users played simple snake game, while in the background workers were doing useful work. The experiment was quite successful: an average game lasted about 2 minutes and impact of background computation on users' experience was not perceptible. None of the users rejected the idea of contributing their CPU cycles to the computation.

Some problems may be more interesting to users than others. Fox et al [30] identified environmental protection, weather forecasting and economic growth projection as kind of problems which potentially may attract many volunteers. Whatever the problem is, its

5 http://iswebrtcreadyyet.com/ Accessed on 9.01.2017

above). Most of the proposals we reviewed either no not report the number of attracted volunteers, or admits that the number is in the range of less than few hundreds. Even QMachine's results seem less impressive when remembering that those were collected over a period of almost a year, for many different projects. Therefore, the problem of recruiting the users is the most important challenge faced by browser-based platforms. We think choosing the right audience and combining them with methods mentioned in this section may be a key to a success of future BBVC platforms.

### 4.4. Trust and security issues

The issues of trust in voluntary computing are twofold. Firstly, BBVC solution must prevent job owners from abusing the trust of the volunteering users, otherwise it could become a platform for distribution of malicious software [51]. A malicious job owner could try to submit trojan horses, viruses, install backdoors, or try to use volunteers' computers to create denial of service attacks on internet services [26]. Volunteers must trust the applications that it does what it claims, that it won't attack their privacy and harm their system [26]. Secondly, users must be prevented from sabotaging the computation. Moreover, if the solution rewards the users for contributing their power, users could falsely claim the credit for the job they had not completed.

We stress that same problems would appear even without malicious intentions from volunteers and job owners. Software bugs may inadvertently harm users' machines, and faulty hardware or software errors on users' machines may be indistinguishable from intentionally submitted wrong results.

To deal with the first issue, the job owners' software can be checked manually and/or tested before being accepted for distribution on a BBVC platform. For example, XtremWeb allows submissions only from designated repository of trusted jobs. The jobs are accepted only from trusted institutions, tested first on set of dedicated users, and finally uses pair of public/private keys is used to ensure code integrity [28]. Unfortunately, sometimes data in distributed computation could be sensitive, and would have to be encrypted while residing on hosts' side [2]. Similarly, application code could be obfuscated [68][8]. We think, however, that if researchers want to use free resources provided by the volunteers, they should reciprocate with sharing their code and results, especially since it may building in build users' trust [59].

Several solutions to the problem of fake results submitted by users were proposed in the literature. Redundant computation is most often used. The canonically correct results may be then chosen by a simple majority voting [48][69][80]. The discrepancies between the results sometimes may sometimes be expected, and may depend on used compiler, architecture, or operating system. Because of that, scheduler in BOINC used homogenous redundancy, sending tasks of a given job only to machines with the same configuration [3]. Voting may be improved by credibility-based weighting, where trusted users' results are given more consideration than novice users [26][49]. Application-specific function may compare the results and choose the canonical, correct version [3][53]. To make sabotageur's work harder, the scheduler may limit number of tasks sent to a single host [3].

In [68] a “spot-checking” technique was proposed, where server would occasionally sent a job to client knowing a correct result in advance. Computation of some software

**Table 2. Accuracy of multiplication in JavaScript [84]**

| Multiplication | Result |
|---|---|
| 10000000000000000.77777777777777777*1 | 10000000000000000.0 |
| 1000000000000000.77777777777777777*1 | 1000000000000000.8 |
| 100000000000000.77777777777777777*1 | 100000000000000.78 |
| ... | ... |
| 10.77777777777777777*1 | 10.777777777777779 |
| 1.77777777777777777*1 | 1.777777777777777 |
| 1.77777777777777777*1 | 1.777777777777777 |
| 1.77777777777777777*1 | 1.777777777777777 |

Problems with JS performance may limit the scalability of the applications dedicated to the BBVC. JavaScript codes may be substantially sped up using asm.js [35][6]. Since asm.js is a subset of JavaScript, code using it should run on most if not all browsers, though older versions may encounter minor problems. An emscripten compiler [81] is able to convert LLVM bytecodes into very fast JavaScript using asm.js, rivaling even a native code in speed [23][80][41]. Therefore, any language which can be compiled into LLVM bytecodes can be converted into fast JavaScript code using this tool. While several works mention possibility of using the emscripten, we are aware of only one actually using it [23]. Scientific code may in future benefit enormously from technologies such as simd.js [37] allowing JavaScript to use accelerated vector operations. Finally, the recently proposed WebAssembly[7] standard may solve the performance problem.

All of the mentioned limitations, lack of ready-to-use components, JS limitations and performance issues limit the acceptance of the BBVC.

## 4.6. Fault tolerance

The are several kinds of faults which can affect distributed computation in BBVC (a *fault tolerance* challenge from Section 3). Worker nodes may fail, users may close the system, the browser, or just the tab running the scripts. All those events mean the task sent to the worker node is lost, and they all can be eventually detected by a job scheduler, for example by using timeouts or by subscribing to *disconnect* event with WebSocket technology. The lost tasks can be then reallocated by the job scheduler to another worker node [9][57][83]. Assigning tasks in round-robin fashion may result in natural redundancy and achieving natural resistance to worker node failures [50][55][63][67][70]. Task re-allocation has positive side-effect of ensuring the computation is not blocked by slowest worker nodes [7][68]. Dealing with faults might be left to the developer, who may choose to ignore lost results [64].

6 http://www.2ality.com/2013/02/asm-js.html
7 http://webassembly.org/

the development for new applications. Finally, increasing the efficiency of the platform should be another priority.

In this section we present our plans for a platform which, in a completed form, hopefully will solve all of the problems and challenges mentioned above, allowing users to volunteer their computing power to solve the scientific problems. PovoCop will consist of set of servers maintained by our team and set of recruiting sites. Any web site could become recruiting site by injecting small code into the pages they serve. Researchers would upload their codes into our platform. Whenever users (volunteers) will visit one of the recruiting sites, task will be automatically fetched from our infrastructure. Computed results will be then aggregated, verified and presented to the researchers submitting the jobs (and, possibly, any other interested party). Currently PovoCop is in an early stage of work, though we have a working prototype.

## 5.1. The motivation

Why BBVC is still relatively unpopular? In Section 4.5 we have pointed to the performance issues and lack of ready-to-use JS components as one factor. We conjecture, however, that there is another reason the scientific community seemed to be hesitant to embrace the possibilities offered by BBVC. Simply, the community has access to clouds, clusters or established non-browser based voluntary computing platform such as BOINC. Therefore, we motivate our work by choosing to target an another audience. Nowadays it is not uncommon for amateurs attempting to solve problems which not so long ago were thought to require professional teams working on supercomputers or grids. Such amateurs usually have modest technical knowledge, sufficient to maintain and configure a blog. Sometimes they are PhD students or scientists, who would want to quickly check some hunch or experiment with data, while lacking access for larger computing infrastructure. While our completed solution should be universal enough to be used by everyone, at least initially we will target this class of users. That will allow the platform to mature and then, we hope, it will attract the more “serious” users.

We envisage that the potential users would insert simple JavaScript code to their blog software (possibly just a single line <*script src*=...>). Whenever their blog would be visited, a script will run, showing some basic information on the sidebar and connecting to infrastructure we intend to establish. As blog authors usually form cliques, some of them could be recruited to mobilize the computing power and to proselytize the solution.

Software will be also provided for so called “power users” who would be able to run their own infrastructure based on Povocop. Therefore, after PovoCop would gain a popularity amongst amateurs, professional users would have easy way to adopt the platform for their own uses. Eventually we would want all those infrastructures to be able to share the harvested computing power.

we can either try a master-slave replication with fail-over in case of primary replica failure, or we could exploit the natural load balancing offered by DNS, by offering several servers operating under the same DNS name.

*Scheduler* (d) is a server that gives tasks (input data for the code) to workers, and

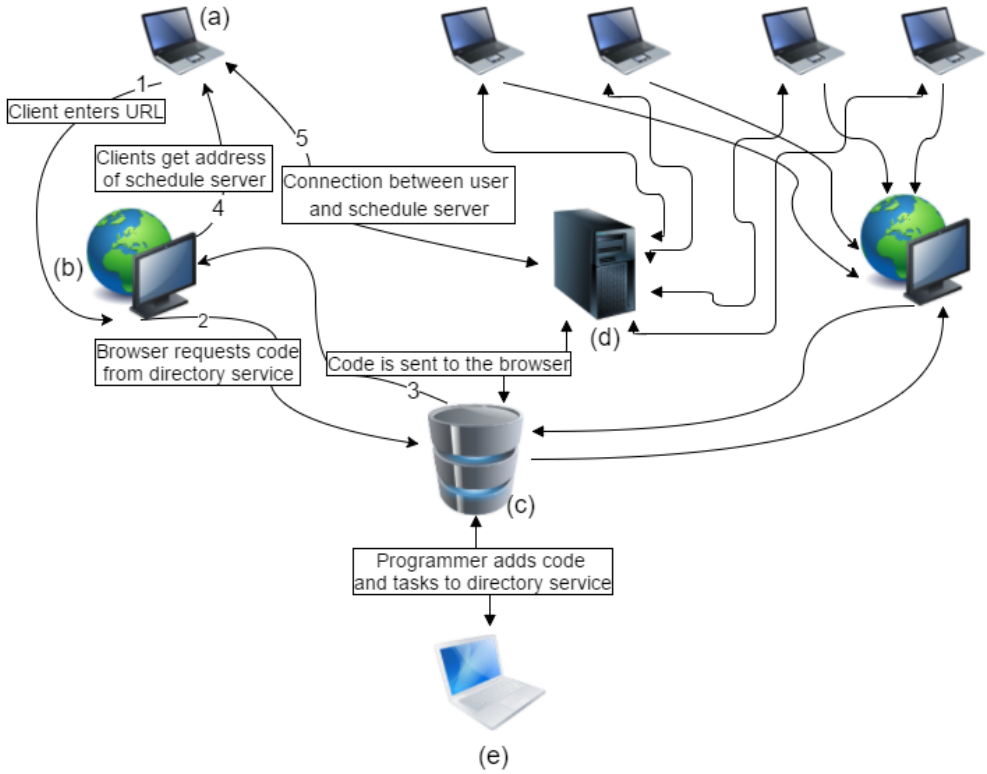


receives the computing results from them. Along with workers it is the main actor in our system. In addition, it monitors if the workers are still computing tasks. In future it will verify the correctness of the results sent by users. It stores the results in directory service. One scheduler can maintain many jobs; ideally, jobs maintained by one scheduler should share the code and differ only in parameters (i.e. input data). In addition, jobs sharing the code should not be maintained by different schedulers. In the current implementation we have only one Scheduler, but we are currently working on adding more. The required functionality should be finished in the next version of the prototype. We want to develop new scheduling policies, taking into account the fact that some Plantations will be more reliable and predictable than other.

**Figure 1. Architecture of PovoCop**

header on both Scheduler and Directory, we can enable browser to access the resources on either server.

The initial script uses a quick benchmark to estimate the number of CPU cores of client's browser and the single core's speed. The Scheduler can use this knowledge to send more tasks to multicore computers or prioritize clients with faster CPU's when deciding on tasks distribution. Benchmark computation will be used in the future as part of spot-checking technique, to ensure user actually runs the tasks and returned correct results.

The browser receives code from Directory Service (server written in Node.js) and tasks from Scheduler (also written in Node.js) and creates a Web Worker with the received code and arguments for each task. When the task is done, the Worker emits result to the main browser scope via postMessage/onMessage JavaScript mechanism. Result is then re-emitted to Scheduler via WebSockets. Scheduler saves the result into the database. Scheduler periodically sends PUT requests to Directory Service with the number of currently connected clients and number of tasks left to do 100% of the Job – this data will be used in the next version of the prototype.

The tasks to execute and the results of computing should be automatically sent to/from server without page reload. For the required communication between workers and servers we decided to use WebSockets. Compared to basic HTTP requests (with AJAX), WebSockets do not need headers for each message. Another benefit is the server's ability to send data to any client at any time. Before that technology, a browser had to send a request to server in order to get some data (using techniques such as reverse AJAX/Comet). Many structures can be sent over Web Sockets – String, JSON, Number, Boolean. Reliability of the data is ensured by underlying TCP protocol. WebSocket allows two way communication – server can send message to browser at any moment. Instead of standard WebSockets browser implementation, we use Socket.io extension. Our choice was dictated by additional features provided by Socket.io, such as long pooling or broadcasting.

In order to execute many threads we use WebWorkers. It allows us to run it in a separate runtime. It is safe against XSS attacks (injecting malicious JavaScript code to the WebSites without having permission, i.e. in blog comments) or CSRF (sending requests to other websites using the cookie for logged in users on these websites).

We need a database to store code, results, scope of tasks arguments, and job arguments. We choose noSQL database MongoDB because of it's scalability, ease of use and integration with Node.js with frameworks like Mongoose.

Every Job must have the code, the IP address of scheduler and the scope of tasks to do inserted into the Data Dictionary's database. Job owner must provide an input data for tasks, the interval of arguments to be sent to clients and the code to be executed by them. There has to be a Scheduler working on public IP address. Job owner has to provide an optional function that verifies results, and a non-optional function that divides input arguments to clients.

Jobs posted by job owners may contain unintended errors, e.g. the code may end up in never ending loop. The working version of the prototype has no built-in defenses against such behavior. In the next version, however, we will use Node.js' ability to run code as a child process in the background. Code will be first run as a small task with parameters on server side. If the code won't end after a certain time, child process will be killed and the code will not be inserted into the database.

As the second algorithm we evaluated the matrix multiplication. We chose this problem to check the abilities of JavaScript in simple math calculations. We also wanted to see the impact of heavy message passing on overall time. In experiments every node received two columns as arrays, multiplied them, sending the results afterward. In this case sequential version was faster than the distributed. The second experiment (Figure 3) revealed that it's not the computing speed, but rather the network bandwidth is crucial in running algorithms with heavy input or output messages – the time the browser spent downloading the matrices was higher than time actually spent on computation. We saw similar results in the third experiment (results not shown) with distributed file sort, when distributed algorithm performance was worse in order of magnitude, mostly because of high communication costs.

We conclude our prototype is better suited to the algorithms downloading input data only once, doing basic calculations and sending results back to the server using lightweight messages.

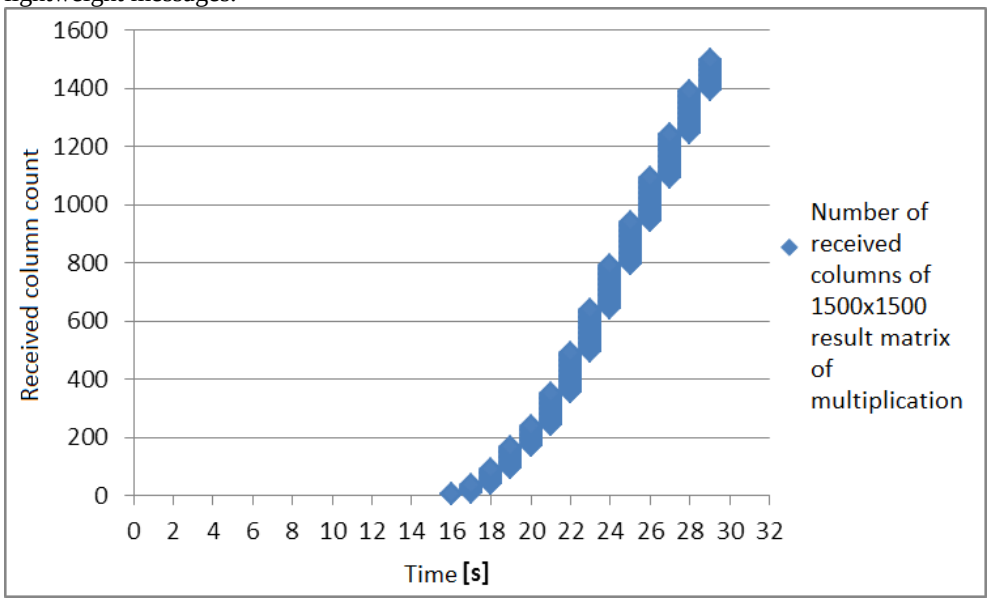


**Figure 3. Graph showing moments of receiving succeeding columns of matrix multiplication of size 1500 x 1500 – for 13 computers using our prototype[84]**

## 6. Related work

One of the first successful large-scale attempts at voluntary computing was GIMPS (Great Internet Mersenne Prime Search) [82], started in January 1996. Until January 2016, fifteen previously unknown prime numbers were found[9], and total average computing power of

9 "GIMPS Project Discovers Largest Known Prime Number: $2^{74,207,281}$-1", http://www.mersenne.org/primes/?press=M74207281

GIMPS project amounted to almost 300TFlops[10] (as of June 2016). Another early precursor was distributed.net [51], created in 1997 to answer the RSA Secret-key challenge[11]. The third, and probably the most known, was Seti@Home project, a search for extraterrestial radio signals of artificial origin, oficially started in 1999.

Based on experiences gathered during Seti@Home project, a team from Berkeley university created BOINC (Berkeley Open Infrastructure for Network Computing) [3]. BOINC is a common infrastructure for distributed voluntary computing. Currently there are at least 57 projects based on BOINC, with close to half million volunteers and more than nine hundred thousands computers, contributing on average 10.5 PetaFLOPS. The largest BOINC project is still Seti@Home, with 1.6 million users, while the smallest project, BMG@Home attracted just 170 volunteers[12]. Other popular projects are Predictor@Home (study of protein behavior), Folding@Home (protein folding), climateprediction.net, climate@home (study of climate changes) and many more. Almost half of compute power is dedicated to top five BOINC projects [75].

BOINC applications can be screensavers, window services or standalone application. To decrease the amount of effort in developing new BOINC applications and their maintenance, a virtual machine approach was implemented [36]. First, the application is compiled for a particular virtual machine. Later, a virtual machine image is created (virtualization-friendly linux plus the instantiation data). Finally, users must install a manager responsible for downloading, unwrapping and running virtual machine images. A lot of effort was put into assuring the last step would be as easy as possible.

Another project with similar goals is XtremWeb [28]. Anyone can set up a project using XtremWeb, starting his or hers own web page. The only requirement is that the "collaborators" must contribute some of their leftover computing power to original XtremWeb.

In XtremWeb, worker initiates the connection with the main server and receives list of servers providing the jobs, as well as how to contact to those servers (i.e what protocol they use and on which port they listen). Clients request work matching their specification. This specification contains description of client's operating system and architecture, and specifies what binaries already are in client's possession. Server answers by sending back the job and address to which client should upload the results. Client periodically pings the server to assure the server its still alive. Clients which do no respond for a long time are considered dead. Job assigned to dead clients is rescheduled. Clients are written in Java, using an interface allowing to call system calls (in C). XtremWeb contains dispatcher which is responsible to choosing the jobs (for example, based on their priority) which are then send to scheduler responsible for scheduling their execution.

The EDGeS project [75] aimed to join computing resources of institutional grids (in this case, from EGEE – Enabling Grids for E-sciencE) and desktop grids, i.e. XtremWeb and BOINC. The ultimate goal was to create special "bridges" enabling grids to act as clients for BOINC and XtremWeb-based projects, and also to allow using BOINC and XtremWeb infrastructure for executing jobs submitted to EGEE grids.

10 "PrimeNet Activity Summary 2016-06-24 15:00 UTC", http://www.mersenne.org/primenet/

11 "distributed.net History & Timeline", http://www.distributed.net/History

12 Retrieved at June the 23 2016 from http://boinc.berkeley.edu/

The existing surveys on voluntary computing in general [26][78], but the surveys we are familiar with dedicate a very limited space for browser-based solutions and describe only the most popular proposals.

A concept related to voluntary computing is crowd computing or "citizen science" (leveraging human abilities for scientific purposes) where humans participate in solving scientific problems. While they contribute their computing resources (by running applications provided by the project teams), the key difference is that they contribute also their skills and knowledge.

There are still problems where a human can relatively quickly find the solution, or good approximation of the solution, while there are no known algorithms or they are slow. One example is pattern recognition. Zooniverse [72] exploits unique human abilities in pattern recognition by encouraging volunteers to participate in projects involving classification of the galaxies, finding exoplanets or classifying wild bees captured by motion cameras in jungle near Serengeti. In 2014 900.000 volunteers participated in 20 projects, and 59 scholarly articles were published with Zooniverse findings.

A related concept are Games With A Purpose (GWAP) or Human Computing Games. In [17] it was shown that when NP-complete graph-related problems are presented as task of determining whether some puzzle is solveable (or not), human players can to find a solution within minutes (when algorithmic approach required hours or even days). In one case novice player devised a novel strategy which eluded an expert researcher (and his three students) studying this problem for three years. Using games has advantage of appealing to larger user base: about 71% of people playing puzzle games are women, and 37% are people in their 40s [18].

FoldIt was descendant of Rosetta@Home – the participants in that project were often frustrated that they were seeing an obvious solution, but were not able to refine the results found by algorithms [43]. In FoldIt players manipulate protein structures using interactive visual tools. Players can create micro-programs, called recipes (a name most likely chosen to sound as non-technical as possible), which can be edited and shared, optimized collectively by society of FoldIt players (though recipes can be private to their creator and/or his team). Some of the recipes are comparable to algorithms created by experts. Within 3 and half month, 721 players run 5488 unique recipes. 568 players wrote 5202 recipes. In January 2013, 2200 players were active.

The input of volunteers to solving scientific problems is much appreciated. "FoldIt players" are credited as coauthors or at least two scientific papers [42][43]. For example, within ten days players were able to find a workable 3D structure of Mason-Pfizer monkey virus M-PMV retroviral protease – a problem for which there was no solution for previous 15 years. As retroviral proteases have critical role in virus proliferation and maturation, this potentially can help design anti-retroviral drugs, including AIDS drugs.

Browser-based approaches may also be used to create distributed file system [49][46], "stealing" disk space from users.

## 7. Conclusions

We have surveyed three generations of browser-based voluntary computing, discussing their commonalities and differences. We have presented the issues shared by all BBVC (and

[8] Berry K., Distributed and Grid Computing via the Browser, In *Proceedings of the 3rd Villanova University Undergraduate Computer Science Research Symposium (CSRS 2009)*, Villanova University, 2009, Retrieved: http://www.csc.villanova.edu/~tway/courses/csc3990/f2009/csrs2009/Kevin_Berry_Grid_Computing_CSRS_2009.pdf

[9] Boldrin F., Taddia C., Mazzini G., Distributed Computing Through Web Browser. *2007 IEEE 66th Vehicular Technology Conference*, Baltimore, MD, 2007, 2020-2024

[10] Brudło P., Foundations of Grid Processing for the Comcute System, in: Balicki J., Krawczyk H., Nawarecki E. (eds.), *Grid and Volunteer Computing*, Gdańsk University of Technology Press, Gdańsk, 2012, 33-39

[11] Carvajal-Gómez R., Frey D., Simonin M., Kermarrec A. M., WebGC Gossiping on Browsers Without a Server. In *Web Information Systems Engineering–WISE 2015, Springer International Publishing*, 332-336

[12] Czarnul P., Kuchta J., Matuszak M., Parallel Computations in the Volunteer–Based Comcute System, in: Goos G., Hartmanis J., van Leeuwen *Parallel processing and Applied Mathematics, Lecture Notes in Computer Science*, **8384**, Berlin, Heidelberg, Springer-Verlag, 2014, 261–271

[13] Czarnul P., On configurability of Distributed Volunteer-based Computing in the Comcute system, in: Balicki J., Krawczyk H., Nawarecki E. (eds.), *Grid and Volunteer Computing*, Gdańsk University of Technology Press, Gdańsk, 2012, 53-69

[14] Liu Ch., White R.W., Dumais S., Understanding web browsing behaviors through Weibull analysis of dwell time. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval* (SIGIR '10). ACM, New York, NY, USA, 2010, 379-386. DOI=http://dx.doi.org/10.1145/1835449.1835513

[15] Charalampidis I., Berzano D., Blomer J., Buncic P., Ganis G., Meusel R., Segal, B. CernVM WebAPI-Controlling Virtual Machines from the Web. In *Journal of Physics: Conference Series*, **664**, 2, 2015, 022010, IOP Publishing

[16] Christiansen B. O., Cappello P., Ionescu M. F., Neary, M. O., Schauser, K. E., Wu, D., Javelin: Internet-based parallel computing using Java. *Concurrency: Pract. Exper.*, **9**, 1997, 1139–1160

[17] Cusack C., Largent J., Alfuth R., Klask K.. Online games as social-computational systems for solving np-complete problems. In: *Meaningful play*, 2010

[18] Cusack C., Martens C., Mutreja P., Volunteer computing using casual games. In: *Proceedings of Future Play 2006 International Conference on the Future of Game Design and Technology*, October 2006, 1-8

[19] Cushing R., Putra G.H.H., Koulouzis S., Belloum A., Bubak M., de Laat C., Distributed Computing on an Ensemble of Browsers. *IEEE Internet Computing*, **17**, 5, Sept.-Oct., 2013, 54-61

[20] Danilecki A., Fabisiak T., Kaszubowski M., Job Description Language for a Browser-based Computing Plaftorm: A Preliminary Report, Accepted for the *9th Asian Conference on Intelligent Information and Database Systems*, April 2017, Kanazawa, Japan

[21] Dean J., Ghemawat S., MapReduce: simplified data processing on large clusters. *Communications ACM* , **51**, 1, January 2008, 107-113

[39] Kajitani S., Nogami Y., Fukushi M., Amano N., A performance evaluation of Web-based volunteer computing using applications with GMP, In: *2015 IEEE International Conference on Consumer Electronics - Taiwan (ICCE-TW),*, Taipei, 2015, 41-42

[40] Kajitani S., Nogami Y., Miyoshi S., Austin T., Volunteer Computing for Solving an Elliptic Curve Discrete Logarithm Problem, In: *2015 Third International Symposium on Computing and Networking (CANDAR)*, Sapporo, pp. 122-126

[41] Khan F., Foley-Bourgon V., Kathrotia S., Lavoie E., Using JavaScript and WebCL for numerical computations: a comparative study of native and web technologies. In *Proceedings of the 10th ACM Symposium on Dynamic Languages (DLS'14)*, ACM, New York, NY, USA, 2014, 91-102

[42] Khatib F., Cooper S., Tyka M. D., Xu K., Makedon I., Popović Z., FoldIt Players, Algorithm discovery by protein folding game players. *Proceedings of the National Academy of Sciences*, *108*(47), 2011, 18949-18953

[43] Khatib F., DiMaio F., Foldit Contenders Group, Foldit Void Crushers Group, Cooper S., Kazmierczyk M., Baker, D., Crystal structure of a monomeric retroviral protease solved by protein folding game players. *Nature Structural & Molecular Biology*, **18**, 10, 2011, 1175–1177

[44] Klein J., Spector L., Unwitting distributed genetic programming via asynchronous JavaScript and XML. In: *Proceedings of the 9th annual conference on Genetic and evolutionary computation* (GECCO '07). ACM, New York, NY, USA, 2007, 1628-1635

[45] Konishi F., Ohki S., Konagaya A., Umestu R., Ishii M., RABC: A conceptual design of pervasive infrastructure for browser computing based on AJAX technologies. In *Seventh IEEE International Symposium on Cluster Computing and the Grid, 2007, CCGRID 2007*, IEEE, 2007, 661-672

[46] Kruliš M., Falt Z., Zavoral F., Exploiting HTML5 Technologies for Distributed Parasitic Web Storage. *Databases, Texts*, 2014, 71.

[47] Krupa T., Majewski P., Kowalczyk B., Turek W., On-Demand Web Search Using Browser-Based Volunteer Computing. In: *Proceedings of 6th International Conference on Complex, Intelligent and Software Intensive Systems (CISIS)*, Palermo, 2012, 184-190

[48] Kuchta P., in: Balicki J., Data partitioning and Task Management in the Clustered Server Layer of the Volunteer-based Computation System, in: Krawczyk H., Nawarecki E. (eds.), *Grid and Volunteer Computing*, Gdańsk University of Technology Press, Gdańsk, 2012, 40-52

[49] Kuhara M., Amano N., Watanabe K., Nogami Y., Fukushi M., A peer-to-peer communication function among Web browsers for Web-based Volunteer Computing, *Communications and Information Technologies (ISCIT), 2014 14th International Symposium on*, Incheon, 2014, 383-387

[50] Langhans, P., Wieser, C., Bry, F. Crowdsourcing MapReduce: JSMapReduce. In: *Proceedings of the 22nd international conference on World Wide Web companion*, International World Wide Web Conferences Steering Committee, May 2013, 253-256

[51] Lawton G., Distributed net applications create virtual supercomputers. *Computer*, 2000, **6**, 16-20

[67] Ryza S., Wall T., *MRJS:* A JavaScript MapReduce Framework for Web Browsers. Retrieved: *http://www.cs.brown.edu/courses/csci2950-u/f11/papers/mrjs.pdf* (2010).
[68] Sarmenta L.F.G., Bayanihan: Web-based volunteer computing using Java. In: *Worldwide Computing and Its Applications—WWCA'98*. Springer Berlin Heidelberg, 1998, 444-461
[69] Sarmenta L.F.G., Hirano S., Bayanihan: building and studying web-based volunteer computing systems using Java. *Future Generation Computer Systems*, **15**, 5–6, October 1999, 675-686
[70] Sarmenta L.F.G., An adaptive, fault-tolerant implementation of BSP for Java-based volunteer computing systems. In: *Parallel and Distributed Processing*. Springer Berlin Heidelberg, 1999, 763-780
[71] Simonarson S., Browser Based Distributed Computing, Retrieved: https://www.tjhsst.edu/~rlatimer/techlab10/SimonarsonPaperQ1-09.pdf, 2010
[72] Simpson R., Page K.R., De Roure D., Zooniverse: observing the world's largest citizen science platform. In *Proceedings of the 23rd International Conference on World Wide Web* (WWW '14 Companion). ACM, New York, NY, USA, 2014, 1049-1054
[73] Tilkov S., Vinoski S.,Node.js: Using JavaScript to Build High-Performance Network Programs. *IEEE Internet Computing*, **14**, 6, November 2010, 80-83
[74] Turek, W., Nawarecki, E., Dobrowolski G., Krupa T., Majewski P., Web Pages Content Analysis Using Browser-based Volunteer Computing. *Computer Science*, **14,**2, 2013, 215
[75] Urbah E., Kacsuk P., Farkas Z., Gilles F., Kecskemeti G., Lodygensky O., Marosi A., Balaton Z., Caillat G., Gombas G., Kornafeld A., Kovacs J., He H., Lovas R., EDGeS: Bridging EGEE to BOINC and XtremWeb. *Journal of Grid Computing*, **7**, 3, 2009, 335-354
[76] Van Kesteren A., Cross-origin resource sharing. *World Wide Web Consortium (W3C) Recommendation*, Retrieved: https://www.w3.org/TR/cors/, January 2014
[77] Vanhelsuw L., Create Your Own Supercomputer with Java, *JavaWorld Online Magazine*, Jan 1997, Retrieved: http://www.javaworld.com/jw-01-1997/jw-01-dampp.html
[78] Venkataraman, N., A Survey on Desktop Grid Systems-Research Gap. In *Proceedings of the 3rd International Symposium on Big Data and Cloud Computing Challenges (ISBCC–16')* , Springer International Publishing, 2016,183-212
[79] Wilkinson S. R., Almeida J. S., QMachine: commodity supercomputing in web browsers. *BMC bioinformatics*, **15**, 1, 2014, 1
[80] Yao P., White J., Sun Y., Gray J.. Gray computing: an analysis of computing with background JavaScript tasks. In: *Proceedings of the 37th International Conference on Software Engineering (ICSE '15)*, **1**, IEEE Press, Piscataway, NJ, USA, 2015, 167-177
[81] Zakai A., Emscripten: an LLVM-to-JavaScript compiler. In *Proceedings of the ACM international conference companion on Object oriented programming systems languages and applications companion* (OOPSLA '11). ACM, New York, NY, USA, 2011, 301-312. DOI=http://dx.doi.org/10.1145/2048147.2048224
[82] Ziegler G. M., The great prime number record races. *Notices of the AMS*, **51**, 4, 2004, 414-416