

Riccardo Volonterio

A framework for web-based human and machine computation

Master Graduation Thesis



Supervisor: Alessandro Bozzon
Assistant Supervisor: Marco Brambilla

Politecnico di Milano
Polo territoriale di Como
Scuola di Ingegneria dell'Informazione
Academic Year 2011-2012

Il non fare nulla è la cosa più difficile del mondo,
la più difficile e la più intellettuale.

— Oscar Wilde

CONTENTS

1	INTRODUCTION	xiii
1.1	Original contribution	xv
1.2	Outline	xv
2	BACKGROUND	1
2.1	Auto and human computation distribution	1
2.1.1	Human computation & GWAP	2
2.1.2	Automatic computation	7
2.1.3	CrowdSearcher	11
2.2	Enabling web-based distributed computation	13
2.2.1	HTML5	14
2.2.2	WebCL	18
3	CONCEPTUAL MODEL	21
3.1	Development methodology	21
3.1.1	Data Design	21
3.1.2	Workflow Design	22
3.1.3	Framework Design	23
3.2	Data model	23
3.2.1	WorkFlow	23
3.2.2	Task Data Model	27
3.2.3	Task Data	28
3.2.4	Task Execution	29
3.2.5	Statistics	31
3.3	Architectural model	31
3.4	Execution model	32
3.5	Pluggable strategies assignment	35
3.5.1	μ Task Planning strategy	35
3.5.2	Performer Assignment strategy	35
3.5.3	μ Task Implementation strategy	36
3.5.4	Task Planning strategy	36
3.5.5	Task control strategy	36
3.5.6	Aggregation function	37
4	USE-CASES	39
4.1	Automatic	39
4.1.1	SIFT algorithm	40
4.2	Human	42
4.3	Hybrid (automatic & human)	43
4.3.1	Introducing Score Degradation	44
4.3.2	Gameplay	45
5	IMPLEMENTATION AND EVALUATION	47
5.1	Architecture	47

5.1.1	Configurator	48
5.1.2	Execution layer	49
5.1.3	Task storage & task runtime storage	51
5.1.4	Performer & Performer Client	51
5.2	Use cases	51
5.2.1	Automatic	52
5.2.2	Human	54
5.2.3	Hybrid	55
A	CONCLUSION AND FUTURE WORKS	57
A.1	Future work	57

LIST OF FIGURES

Figure 1	General structure of a crowd based distributed computing system.	2
Figure 2	Human Computation relation with CrowdSourcing, Social Computing and Collective intelligence.	3
Figure 3	Centralized vs Distributed execution of Human Computation.	4
Figure 4	<i>Amazon's Mechanical Turk</i> web interface for choosing the HIT.	5
Figure 5	The reCAPTCHA user interface containing the text to be digitized.	6
Figure 6	The FoldIt user interface.	7
Figure 7	General structure of a distributed computing system.	8
Figure 8	The BOINC computational flow.	9
Figure 9	Schematic diagram of the parasitic computer solving the 3-SAT problem.	10
Figure 10	The CrowdSearch framework.	12
Figure 11	Official HTML5 logo & unofficial CSS3 logo.	14
Figure 12	The slow script dialog.	17
Figure 13	OpenCL execution flow.	19
Figure 14	OpenCL device composition.	19
Figure 15	The conceptual Workflow handled by the framework.	22
Figure 16	Overall Data Model schema, with logical subdivision.	24
Figure 17	Conceptual organization of Work, Task and μ Task.	25
Figure 18	Reference architecture.	32
Figure 19	Representation of the Task execution flow.	33
Figure 20	Interface of the automatic use-case.	39
Figure 21	AIDA web interface for WSD.	43
Figure 22	Interface of the hybrid use-case.	44
Figure 23	Specialized architecture.	47
Figure 24	Work/Task creation flow.	49
Figure 25	Manual Task planning vs Automatic Task planning.	50
Figure 26	Task planning flow.	50
Figure 27	Intermediate results of the algorithm.	53
Figure 28	SIFT result comparison with the reference data.	53
Figure 29	Human use case Web interface.	54

Figure 30 Face detection algorithm demo. 55

LIST OF TABLES

Table 1 Computation type vs User awareness matrix. . . xiv
Table 2 Samsung WebCL performance compared to pure
JavaScript. 20
Table 3 Task planning vs. Task Assignment. 33
Table 4 **SIFT** algorithm time performances on the Web. . 54

SOMMARIO

Negli ultimi anni si è vista l'esplosione del numero di contenuti multimediali prodotti dagli utenti. La gestione di enormi archivi di informazioni richiede metodi e strumenti in grado di garantire l'efficienza e la rapidità dell'analisi e del recupero dei dati. Una classica soluzione per risolvere questo problema è l'uso della computazione distribuita. Questa tecnica permette di distribuire il calcolo su nodi diversi, in modo da massimizzare la parallelizzazione. L'analisi di contenuti multimediali ha dei requisiti particolari: da una parte la natura delle informazioni analizzate richiede sistemi in grado di essere distribuiti su più nodi, dall'altra, richiede anche una certa qualità nell'analisi. La computazione umana è una tecnica in grado di sopperire a questa esigenza tramite l'utilizzo di esseri umani nel processo di analisi, in modo da migliorare la qualità dei risultati ottenuti.

Contemporaneamente allo sviluppo di queste tecnologie si è assistito anche alla definitiva esplosione, ed evoluzione, del Web, che si è trasformato da una semplice rete per la distribuzione dei contenuti fino a diventare uno strumento sociale per la collaborazione, in grado di utilizzare gli utenti come nodi di computazione sia automatica sia umana.

Allo stato attuale delle conoscenze in nostro possesso, non siamo in grado di individuare soluzioni che supportino senza problemi questa enorme forza lavoro, permettendo lo sfruttamento sia della computazione umana, attraverso l'uso delle capacità intellettive degli esseri umani, sia quella automatica, attraverso l'uso dei client Web (Browsers).

In questa tesi è presentato un framework Web per la computazione umana e automatica in grado di sfruttare al meglio le possibilità offerte da queste due soluzioni. Abbiamo, inoltre, validato il nostro lavoro implementando tre casi d'uso in grado di stressare le diverse configurazioni architetturali supportate.

ABSTRACT

Last years witnessed an explosion in the number of multimedia contents produced by users. Managing huge repositories of digital information calls for timely and efficient methods and tools for analysis and retrieval. The distribution of the analysis task on different computation nodes to maximize parallelization is a classic solution to this problem, typically referred to as Distributed Computing (DC). However, the analysis of multimedia contents brings specific requirements: one one hand, the nature of the analyzed information demands systems able to scale, thus involving a good amount of computational nodes. On the other hand, multimedia analysis is known to be a subtle task to be performed, as several state of the art solutions do not bring sufficient analysis quality; the recent trend of Human Computation (HC) acknowledges the needs for more correct processing by involving human beings in the multimedia content analysis tasks.

At the same time of the development of these methodologies we have also seen the definitive explosion of the Web and its evolution. The Web has evolved from a mere content delivery network, where the contents are presented to the users, to a collaborative and social tool, where powerful client-side technology enable the usage of Web users as node for both automatic and human computation.

To the best of our knowledge, no existing solutions are able to seamlessly exploit this massive workforce, which accounts for both machine computational power (Web Clients) and human computational power (the user brain).

This thesis presents a framework for web-based human and machine computation able to take advantages of both human and distributed computation. We validated our approach by implementing several use case applications, which demonstrate the framework by stressing its architectural configurations.

*Abbiamo visto che la programmazione è un'arte,
perché richiede conoscenza, applicazione, abilità e ingegno,
ma soprattutto per la bellezza degli oggetti che produce.*

— Donald Ervin Knuth

RINGRAZIAMENTI

Per raggiungere questo traguardo le persone da ringraziare sono molte. Prima di tutto vorrei ringraziare Alessandro per la sua pazienza e per avermi sopportato durante la stesura di questa tesi. Se avessi messo in pratica la metà dei suoi consigli mi sarei laureato mesi fa.

Grazie alla Pas per aver speso una fortuna in caramelle e averle condivise con tutti. Grazie per avermi detto "LAVORA!" ogni 5 secondi, anche mentre stavo lavorando.

Grazie a Luca T. per le sue consulenze informatiche sempre puntuali e complete.

Grazie a Luca G. che nonostante quello che dice il cartone/televisore è sempre disponibile ad aiutare gli altri. Grazie anche per avermi fatto saltare la fila.

Grazie Andrea per aver fatto un grande lavoro con lo strumento meno indicato per farlo. Senza il tuo lavoro questa tesi non esisterebbe.

Grazie all'AirLab Team di Como per le giornate perse a fare altro e per quelle passate a studiare quando si era stanchi di non fare nulla.

Grazie ai Difensori della Rocca, del Muro e di Roccopodio per avermi fatto diventare l'elfo di 4° livello che sono ora.

Grazie Ivan per essere testardo. Senza la tua testardaggine avrei capito la metà degli esami preparati insieme.

Grazie a Botta per avermi guidato come un bambino nella selva della burocrazia ricordandomi ogni singola rata da pagare, esame da preparare e consegna da fare con settimane di anticipo.

Grazie a Rocco per avermi sopportato in tanti anni e per le serate passate al 35 a giocare a Trivial Pursuit o a Jenga. P.S. non siamo ancora diventati immortali.

Grazie alla mia famiglia per aver creduto in me e avermi sostenuto in questi anni. Grazie per le frasi come: "Come va con la tesi?" che volevano dire altro.

Grazie Arianna per avermi aiutato nella correzione di questa tesi. Grazie per essermi sempre vicina. Grazie per essere sempre pronta ad aiutarmi. Grazie.

1

INTRODUCTION

Last years witnessed an explosion in the number of multimedia contents produced by users. Managing huge repositories of digital information calls for timely and efficient methods and tools for analysis and retrieval. The distribution of the analysis task on different computation nodes to maximize parallelization is a classic solution to this problem, typically referred to as **DC**. However, the analysis of multimedia contents brings specific requirements: on one hand, the nature of the analyzed information demands systems able to scale, thus involving a good amount of computational nodes. On the other hand, multimedia analysis is known to be a subtle task to be performed, as several state of the art solutions do not bring sufficient analysis quality; the recent trend of **HC** acknowledges the needs for more correct processing by involving human beings in the multimedia content analysis tasks.

Distributed Computing deals with the computation distribution among computers, also called nodes, connected to a network. The execution of the code is possible thanks to the creation of an abstraction layer on top of each node. This layer normalizes the differences between computers by abstracting the available resources in order to make them consistent. For instance grid computing abstracts only part of the available resources, meanwhile cloud computing abstracts the whole hardware.

The distribution of the computation can be done at **hardware** or **software** level. At **hardware** level we have similar distributed resources, or at least they can be easily abstracted, giving us the ability to offload the same code to the nodes and gather the results. This type of computation distribution is used in frameworks like MapReduce (Dean and Ghemawat, 2008) where the computation is spread on large clusters of computers. The distribution of computation at **software** level uses ad-hoc softwares to normalize the resources of a computer. With this coherent representation of nodes, the distributed system is now able to send code and gather results.

DC can operate in two different ways according to the user awareness of execution: transparent and opaque. If the user is aware of the computation being performed then we are dealing with transparent (Voluntary) execution otherwise we have opaque (Involuntary) execution.

For the Voluntary scenario there are solutions like Berkeley Open Infrastructure for Network Computing (**BOINC**) or Distributed.net¹. Meanwhile for the Involuntary scenario we have *Parasitic Computing*, as described Barabási *et al.*, 2001.

¹ <http://www.distributed.net>

As one may notice this type of distributed computation strongly relies on the presence of the abovementioned abstraction layer. Usually this layer is created using a complex cross-platform application software able to normalize the execution of code among different OS and architectures.

Human Computation is used where a computational process performs its functions by outsourcing some activities to humans. As one may notice **HC** is suitable only for a certain category of tasks, like for example image recognition or Word-Sense Disambiguation. This kind of applications usually relies on the Web as the main platform for distributing and executing such tasks. For instance *Amazon's Mechanical Turk* is a web-based market for performing **HC** in exchange of money rewards. A major issue for **HC** is how to engage users, and, more importantly, how to keep them engaged. To overcome this issue some **HC** tasks can be executed as part of a gaming experience, using the so called Game With A Purpose (**GWAP**). Games such as Peekaboomb², the ESP Game³ and *FoldIt* utilize the entertainment given by playing the game as an engaging factor.

Table 1 summarizes the available solutions for performing crowd-base computation distribution. In the table we added a dimension concerning the user awareness of the task⁴ execution.

Table 1.: Computation type vs User awareness matrix.

	Automatic	Human
Voluntary	BOINC	<i>Amazon's Mechanical Turk</i>
Involuntary	Parasitic computing	GWAP

At the same time of the development of these methodologies we have also seen the definitive explosion of the Web and its evolution. The Web has evolved from a mere content delivery network, where the contents are presented to the users, to a collaborative and social tool full of Rich Internet Application (**RIA**). The advent of **RIA** was possible due to the great evolution of the computation performance on the client side. **HTML5** boosted the computing possibility of web browser by giving to developers the possibility to access video and audio raw data, interact with the filesystem and create communication channels between the client and the server.

Given this general overview one can spot that we reached the condition where we have the technical ability to use all the web-users as nodes for a web-based human and machine computation framework.

As far as we know there are no existing solutions able to offer both the advantages of human and the distributed computation.

² Von Ahn, Liu, and Blum, 2006.

³ <http://www.gwap.com/gwap/gamesPreview/espgame/>

⁴ For **GWAP** we considered the purpose of the game not the game itself.

1.1 ORIGINAL CONTRIBUTION

This thesis presents a framework for web-based human and machine computation able to cover all the dimensions expressed in [Table 1](#). The Task has been designed to use standardized protocols/languages, ease the Task implementation by the developers, ease the execution of the Task by the users. The main contribution of this thesis are:

1. The Definition of a unified model for automatic, human and hybrid computation.
2. The design of a reference web-based architecture for human and automatic computation.
3. The implementation of an infrastructure supporting the aforementioned model.
4. A validation through three use-cases that stress the framework to support: [automatic](#), [human](#) and [hybrid](#) computation.

1.2 OUTLINE

The rest of the thesis is organized as follows.

THE SECOND CHAPTER introduces the context of this thesis, giving the needed notions that will be used during the presentation of the model and the use-cases.

THE THIRD CHAPTER presents the conceptual model of our framework.

THE FOURTH CHAPTER introduces the use-cases that we have used to validate our framework.

THE FIFTH CHAPTER describes the implementation stage of the framework and the use-cases.

2 | BACKGROUND

Recent years have seen an increasing interest in *Human Computation* and *Crowdsourcing* areas. One of the reason they are becoming so attractive is the growth of the Web as a content production and social interaction platform. This has allowed to leverage on the ability of people over the Internet to perform tasks.

This chapter presents the main research fields related to *harnessing human intelligence to solve computational problems that are beyond the scope of existing Artificial Intelligence (AI) algorithms* (Law and Ahn, 2011), providing a brief introduction to the terms and the core concepts that will be used during the exposition.

Section 2.1 gives an introduction to the concept of *distributed computing*, focusing on *HC* and *Automatic computation*, from both the theoretical point of view and to the state-of-the-art tools that leverages on this techniques.

Section 2.2 presents the web technologies that enables the use of the *distributed computing* paradigm over the web, focusing on the computational part of the *distributed computing* process.

2.1 AUTO AND HUMAN COMPUTATION DISTRIBUTION

Generally speaking *computation distribution* is a paradigm for splitting a task into atomic subtasks that can be performed on multiple nodes. Once these nodes have executed the subtasks they can send the results back to the central hub. Using client-server as an example architecture, we can list the required operation needed to obtain *computation distribution*:

1. the server **splits** the workload into atomic operations
2. the server **sends the task**, among with the needed data, to the clients
3. the client **performs** the atomic task
4. the client **sends the results** back to the server
5. the server **gathers** the results from all the clients
6. the server **joins** the results

The previous operations are the cornerstone of every *computation distribution* system. Although they can be "implemented" in different ways

or can be joined together, they are present in all the distributed systems.




Figure 1.: General structure of a crowd based distributed computing system.

As shown in [Figure 1](#) the distribution of computation can be subdivided in two areas. One concerning the distribution of tasks to humans called Human Computation (HC). The other one focuses on the distribution of code to machines and is called Distributed Computing (DC).

Consistently with [Table 1](#) and with the previous subdivision we split the general problem of crowd-based computation distribution into two fields: *Human Computation & Game With A Purpose* and *distributed computing*

In [2.1.1](#) are presented the theoretical basis as long as the state-of-the-art tools that deal with HC and the distribution of task to a human crowd.

In [2.1.2](#) the concept of *distributed computing* is presented and the main tools that implement this paradigm are described.

In [2.1.3](#) is presented the concept of *Crowdsearching*, a new paradigm for searching in the Web using the crowd as source of information.

2.1.1 Human computation & GWAP

Human Computation (HC) is a computer science technique where a computational process performs its function by outsourcing certain steps to humans. This *outsourcing* process, as explained in the [introduction](#), is mainly due to the computational complexity of AI algorithms. There are some AI problems that cannot be solved by computers or are too computationally intensive to be solved by computers in a reasonable amount of time.

Some of these are very simple tasks for humans, for example natural language processing and object recognition are hard to solve problem for a computer, but natural for a human being. A great example for

this kind of problem is recognizing hand-written text. Even after years of research, humans are still faster and more accurate than computers. Other AI problems are too computationally expensive, such as many NP-complete problems like Traveling Salesman problem, scheduling problems, packing problems, and FPGA routing problems.

The expression *Human Computation (HC)* in the context of computer science is already used by Turing, 1950. However is Law and Ahn, 2011 to introduce the modern usage of the term. It defines human computation as "a research area of computer science that aims to build systems allowing massive collaboration between humans and computers to solve problems that could be impossible for either to solve alone". A most simple and direct definition of HC is:

*Some problems are hard, even for the most
sophisticated AI algorithms.
Let humans solve it...*

— Edith Law




Figure 2.: Human Computation relation with CrowdSourcing, Social Computing and Collective intelligence.

Human Computation is related with other fields, such as *CrowdSourcing*, *Social Computing* and *Collective Intelligence* as depicted in Figure 2. Here we give some definitions to better understand the similarities and the differences:

CROWDSOURCING is "the act of taking a job traditionally performed by a designated agent (usually an employee) and outsourcing it to an undefined, generally large group of people in the form of an open call" Howe, 2006. So it does not involve computation directly like HC.

SOCIAL COMPUTING "describes any type of computing application in which software serves as an intermediary or a focus for a social relation"

Schuler, 1994. So despite of the name its purpose is not computing.

COLLECTIVE INTELLIGENCE defined very broadly as "groups of individuals doing things collectively that seem intelligent".

When dealing with a human crowd the main issue is to engage users to perform tasks. A user can be motivated to perform a task due to its nature (e.g. the task helps finding the cure to some disease) or to the revenue (e.g. karma¹) he/she gets for doing such task. The most effective way for recruiting and motivating users is to give them money². For instance *Amazon's Mechanical Turk* is an online tool for performing Human Intelligent Task (HIT) in exchange of money rewards³.



Figure 3.: Centralized vs Distributed execution of Human Computation.

The categorization of HC can be further specified by adding another dimension that involves how the tasks are executed by the users. As you can see in Figure 3 there are two main types of HC execution: *centralized* and *distributed*.

Centralized

In the *centralized* execution we have a central hub (i.e. a website) where users must go to perform the task. Typically the execution of a task does not involve the offload of code and data to the user, and there is no need of ad-hoc softwares to run the task.

A good example of a *centralized* HC platform is *Amazon's Mechanical Turk*. *Amazon's Mechanical Turk* is an online platform for executing tasks in exchange of money rewards. The platform is divided into two sections, one for the *Workers* and one for the *Requesters*. The **Workers** are users willing to spend time to execute a HIT and receive the reward, the **Requesters** are users that publish HIT and after getting the results pay the *Workers*.

¹ Reputation points used in www.reddit.com.

² Since the beginning of time.

³ The rewards for a single HIT can be as low as 0.01\$.

The lifecycle of a [HIT](#) is the following:

1. A *Requester* creates a [HIT](#) using one of the predefined project instances available.
2. Once the creation is completed the [HIT](#) is ready to be executed by the *Workers*.
3. To execute a [HIT](#) a *Worker* must visit the *Amazon's Mechanical Turk* website and choose from a list of available [HITs](#) the one that he/she wants to perform (see [Figure 4](#)).
4. Once the whole [HIT](#) is completed the *Requester* checks the result obtained and if he/she is satisfied then proceeds with the payment.

As one can see the whole flow of the [HIT](#) from the creation to the payment of the *Workers* is done within the browser.

The screenshot shows the Amazon Mechanical Turk interface. At the top, there's a navigation bar with 'Your Account', 'HITS', and 'Qualifications' tabs. Below that, a search bar is visible with filters for 'HITs Available To You' and 'HITs Assigned To You'. The main content area displays a list of HITs under the heading 'All HITs' and '1-10 of 1457 Results'. Each HIT entry includes details such as the requester (SpeechLink), HIT Expiration Date, Time Allotted, Reward, and HITs Available. The list includes tasks like 'Video Transcribing and Tagging' and 'Voicemail Transcription'.

Figure 4.: *Amazon's Mechanical Turk* web interface for choosing the [HIT](#).

This platform has been *extended* as presented in [Little et al., 2010](#) to create complete Artificial Intelligence algorithms able to use human computation as functions during the execution process. The code in [Listing 1](#) is an example of an algorithm implemented using Turkit. Here `mturk.prompt` and `mturk.vote` are Human Intelligent Task executed on the *Amazon's Mechanical Turk* platform.

Listing 1: Example of a Turkit algorithm.

```
ideas = []
for (var i = 0; i < 5; i++) {
  idea = mturk.prompt("What's fun to see in New York City?
  Ideas so far: " + ideas.join(", "))
  ideas.push(idea)
}
ideas.sort(function (a, b) {
```

```

    v = mturk.vote("Which is better?", [a, b])
    return v == a ? -1 : 1
})

```

Distributed

In a *distributed* execution environment the central hub acts as a distribution node in charge of offloading the task upon user request. The user can now run the task locally without the intervention of the central hub. Eventually, when the task is done, the user contacts the hub to upload the results. The process of requesting the task to the hub, executing the task and sending the results, is all done by the users, typically using a preinstalled piece of software.

One of the best known examples of Distributed Human Computation (DHC) is **reCAPTCHA**. CAPTCHAs (Completely Automated Public Turing test to tell Computers and Humans Apart) are widespread security measures on the World Wide Web that prevent automated programs from abusing online services. They do so by asking humans to perform a task that computers cannot yet perform, such as deciphering distorted characters. reCAPTCHA channeled such human effort into a useful purpose: helping to digitize old printed material by asking users to decipher scanned words from books that computerized Optical Character Recognition (OCR) failed to recognize (see Figure 5). As described in [Von Ahn *et al.*, 2008](#), this method can transcribe text with a word accuracy exceeding 99%, matching the guarantee of professional human transcribers.



Figure 5.: The reCAPTCHA user interface containing the text to be digitized.

Game With A Purpose

Game With A Purpose (GWAP) is defined as "*a human-based computation technique in which a computational process performs its function by outsourcing certain steps to humans in an entertaining way*" by von Ahn. GWAP come from a simple observation of data on how many hours are spent playing games. [Von Ahn and Dabbish, 2008](#) reported that, according to the Entertainment Software Association⁴, more than 200 million hours are spent each day playing computer and video games in the U.S.. Indeed, by age 21, the average American has spent more

⁴ Game data from www.theesa.com/facts/gamer_data.php

than 10,000 hours playing such games equivalent to five years of working a full-time job 40 hours per week.

The simple idea behind **GWAP** is "*why not make playing games useful?*" If a task can be transformed into a game the user can be motivated to play the game so there is no need of other types of reward (i.e. money). The entertainment of playing the game itself can be used as a reward for the user.

GWAPs have demonstrated that they can be used as a valuable source of information. For instance the ESP game described in [Von Ahn and Dabbish, 2004](#), is a **GWAP** developed by Luis von Ahn to perform image tagging. The users' task is to agree on a word that would be an appropriate label for the recognition of the image. Another example of a successful **GWAP** is Peekaboom, where users help computers locating objects in images, and the *FoldIt* game.




Figure 6: The FoldIt user interface.

FoldIt is a puzzle game about protein folding, developed by the University of Washington's Center for Game Science in collaboration with the UW Department of Biochemistry. The objective of the game is to fold the structure of selected proteins to the best of the player's ability. The highest scoring solutions are analyzed by researchers, that can determine whether or not there is a native structural configuration that can be applied to the relevant proteins (see [Figure 6](#)).

2.1.2 Automatic computation

Unlike human computation, *automatic computation* aims at executing a task, or part of it, in an automatic fashion, without user's interaction.

This kind of *distributed computation* leverages on the existence of a *grid* of connected nodes able to perform data intensive calculation.

As stated in "[Foundations of multithreaded, parallel, and distributed programming, 2000](#)": "*Distributed computing is a field of computer science that studies distributed systems. A distributed system consists of multiple autonomous computers that communicate through a computer network. The computers interact with each other in order to achieve a common goal. A computer program that runs in a distributed system is called a distributed program, and distributed programming is the process of writing such programs.*"




Figure 7.: General structure of a distributed computing system.

The platforms that implements these solution use different frameworks for splitting algorithms into atomic operations. One of these frameworks is MapReduce⁵ that, using the core concept of *Divide et impera*, can produce highly parallelizable algorithms.

The name "*distributed computing*" refers to a wide range of different applications (i.e. grid computing, cloud computing, [Parasitic computing](#)⁶, [jungle computing](#)⁷), that implements the same paradigm with different purposes. Using the dimensions presented in [Table 1](#) we can divide the core concept of *distributed computing* into two subcategories: *voluntary computing* and *parasitic computing*.

Voluntary computing

Voluntary computing refers to all those *distributed computation* systems where the computation is performed on behalf of the users will. In such systems users go to the website of the "project" they intend to support and, usually by installing an ad-hoc client, give the resources (i.e. CPU idle time, storage, etc.) of their machines to the distributed system.

⁵ [Dean and Ghemawat, 2008.](#)

⁶ [Barabási et al., 2001.](#)

⁷ [Seinstra et al., 2011](#) and winner of the coolest name 2012.

The first project implementing such computing paradigm was Great Internet Mersenne Prime Search ([GIMPS](http://www.mersenne.org/))⁸, started in 1995. Then other platforms (such as [distributed.net](http://boinc.berkeley.edu/), [SETI@home](http://setihome.berkeley.edu/), etc.) have been developed to support this type of computation.




Figure 8.: The BOINC computational flow.

Berkeley Open Infrastructure for Network Computing (BOINC)⁹ is an open-source software for Voluntary computing and grid computing. BOINC was originally developed to support the Search for Extraterrestrial Intelligence *at home* (SETI@home) project, before it became used as a platform for other distributed computing applications. It consists of two parts: the back-end server, running on Linux platforms, and the client, cross-platform, for the end-user computing.

The client software allows users to connect to the BOINC grid, to perform computation. The flow of the execution, as depicted in Figure 8 is the standard *distributed computing* execution flow. Here is the list of actions performed to run the task on the client:

1. get instruction from the server on how to get all the resources needed to perform the task,
2. download all the resources from the server,
3. execute the downloaded application,
4. send the results obtained to the server.

There are, at the time of writing, over 40 projects that leverage on the BOINC platform to perform computation on different application areas. For example the aforementioned SETI@home project uses this framework to search for extraterrestrial intelligence by analyzing the narrow-band radio signal coming from the Arecibo radio telescope.

Parasitic computing

Parasitic computing¹⁰ is a technique that, using some exploits and ad-hoc code, allows a *malicious* user to use the computational power of the *victim* computer without this being aware. As one can notice

⁸ <http://www.mersenne.org/>

⁹ <http://boinc.berkeley.edu/>

¹⁰ In this thesis we are not covering, neither we are interested in, the ethical or moral implication of using this technique.

parasitic computing has a strong relationship with *distributed computing*. In fact it is like a specialization of the general class of *distributed computing*, where the user is unaware of the execution¹¹.

This approach was first proposed by Barabási *et al.*, 2001 to solve the NP-complete 3-SAT problem using the existing TCP/IP protocol stack and its error handling routines. The satisfiability problems or "SAT" involve finding a solution to a boolean equation that satisfies a number of logical clauses. For example, $(x_1 \oplus x_2) \wedge (x_2 \wedge x_3)$ in principle has 2^3 potential solutions, but it is satisfied only by the solution: $x_1 = 1, x_2 = 0, x_3 = 1$. Problems like the one in the example are known as 2-SAT problems because each clause, shown in parentheses, involves two variables. The more difficult 3-SAT problem is known to be NP-complete, which in practice means that there is no known polynomial-time algorithm which solves it.

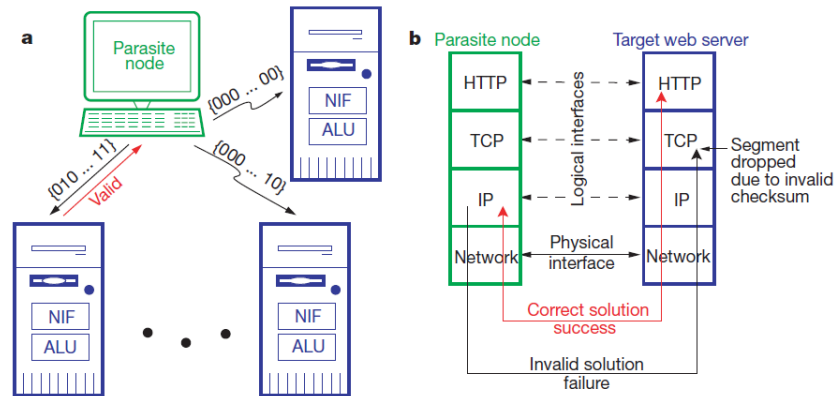


Figure 9.: Schematic diagram of the parasitic computer solving the 3-SAT problem.

The approach proposed in the paper was to perform a brute force attack to guess the right solution of a 3-SAT problem using a parallel approach as depicted in Figure 9. The parasite node creates 2^n specially constructed messages designed to evaluate a potential solution. These messages are sent to many target servers throughout the Internet. After receiving the message, the target server verifies the data integrity of the TCP segment by calculating a TCP checksum. The construction of the message ensures that the TCP checksum fails for all messages containing an invalid solution to the posed SAT problem. Thus, a message that passes the TCP checksum contains a correct solution. The target server will respond to each message it receives (even if it does not understand the request). As a result, all messages containing invalid solutions are dropped in the TCP layer. Only a message which encodes a valid solution *reaches* the target server, which sends a

¹¹ In *distributed computing* the user can be unaware of the actual code they are executing, but they are aware of the execution.

response to the *request* it received.

This approach may seem unfair to the user, but if one thinks about it, may notice how we are always making computation without even knowing it. Playing GWAPs or using application like reCAPTCHA are examples of involuntary human computation (as defined in Table 1). So they are using the same technique to perform a sort of *parasitic human computing* without complaining about the users willingness.

A hybrid approach (parasitic & voluntary) can be used to avoid the ethic implication of doing pure *Parasitic Computing*. If the user gives the permission to run computation on its computer exchange of some sort of return, then we are able to score the best on both approaches. A similar solution was proposed in Karame, Francillon, and Čapkun, 2011. In this paper the authors propose a micro-computation as micro-payments in web-based services. Their solution is to give the user access to online contents (such as newspaper, video, etc.) after performing small JavaScript computation.

PARASITIC JAVASCRIPT described in Jenkin, 2008 can be considered as an enhancement of the solution proposed by Barabási *et al.*, 2001. Since using JavaScript and the HTML5 features to their full potential (see 2.2.1) we are able to perform any kind of computation within the browser window/tab. JavaScript offers also a standard platform for the execution of code without the need of any ad-hoc software for each platform. Furthermore all the code executed by a browser run in a sandboxed environment, keeping the user's computer safe from any malicious intent.

2.1.3 CrowdSearcher

Web search has evolved since it was a massive link analysis system. Now the contents of a search query have been enriched with multimedia contents from the crowd. For instance a query can show results as links to other websites, user generated videos, tweets or images. On top of this, a lot of tasks has been tweaked to generate more searchable data (e.g. image and video tagging), so now we are able to search for them. The crowd is becoming part of the search process, by adding information, ranking the results, moderating contents, etc.

CrowdSearch is targeted to enabling, promoting and understanding individual and social participation to search Ceri *et al.*, 2012. CrowdSearch uses the crowds as sources for the content processing and information seeking processes. It fills the gap between generalized search systems, which operate upon world-wide information - including facts and recommendations as crawled and indexed by computerized systems and social systems, capable of interacting with real people, in real time Ceri *et al.*, 2012. Crowd-searching can be defined as the promotion of individual and social participation to search-based applica-

tions and improve the performance of information retrieval algorithms with the calibrated contribution of humans [Bozzon, Brambilla, and Ceri, 2012](#).

The CrowdSearch framework

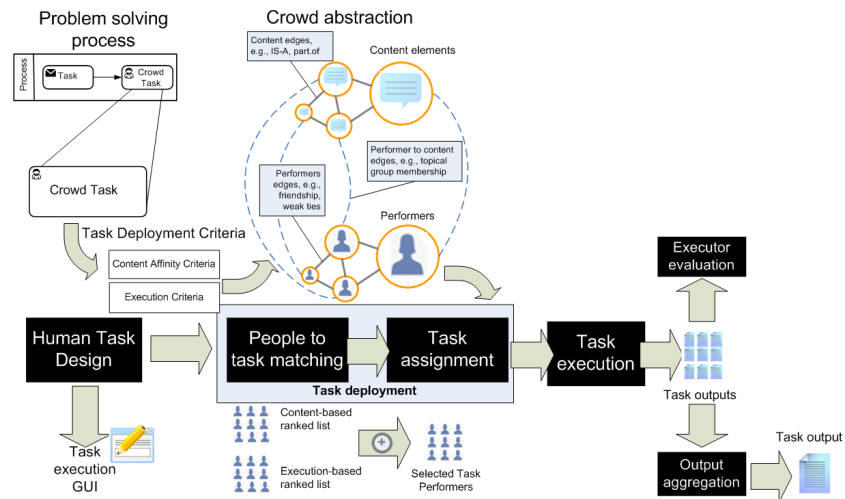


Figure 10.: The CrowdSearch framework.

The framework (see [Figure 10](#)) proposed by [Bozzon, Brambilla, and Ceri, 2012](#), uses the humans' skills to improve some the result of a query. To perform such operation the framework must choose which task must be executed by humans and which must not. This has been addressed by mapping some tasks to humans and some others to machines.

When a task is assigned to the crowd, the execution mode is chosen according to the *Human Task Design*. This step produces the actual design of the *Task Execution GUI*, and the *Task Deployment Criteria*. These criteria can be logically subdivided into two subject areas: *Content Affinity Criteria* (what topic the task is about) and *Execution Criteria* (how the task should be executed). The Execution Criteria could specify constraints or desired characteristics of task execution including: a time budget for completing the work, a monetary budget for paying workers, etc.

The framework provides an abstraction for the crowd and their skills, called *Crowd Abstraction*. Human performers and content elements are represented as nodes connected by edges in a bipartite graph. Edges connecting performers denote a friendship, edges connecting content elements denote a semantic relationship and finally, a connection from a performer to a content element may denote an interest.

The next step to the *Human Task Design* is the *Task Deployment*, whose goal is to assign the crowd tasks to the most suited performers. This step is subdivided into two sub-steps: *People to Task Matching* and *Task Assignment*. The first sub-step is almost a query matching problem,

where the workers must be ranked according to the *Task Deployment Criteria* and a set of suitable workers is selected (e.g. top-k).

The final step is the *Task Execution*. This step is the actual execution of the task by the crowd. The results are then merged to form the result of the crowd task.

2.2 ENABLING WEB-BASED DISTRIBUTED COMPUTATION

Using the web (i.e. the browser) as a platform for distributing and executing code implies that we have the available technologies to perform high-level *computation* and real-time *communication*. These are the requirements for evaluating if the web is a suitable platform for code distribution.

COMPUTATION is the key for being able to perform tasks within the browser. *Computation* can involve any kind of operation on the data, and the data itself can be of any type. For instance creating an application that analyzes audio files, or creating a image manipulation program that runs without external plugins is relatively simple using standard languages (e.g. C, C++, Java, etc). But until a few years ago it was almost¹² impossible to do it within a browser.

HTML5 filled the gap that existed between any "standard" language and JavaScript. This is done by giving the developers access to all the required APIs needed to create fully functional web-applications. In 2.2.1 are presented all the features, along with some technical details, that have enabled this evolution.

There are also initiatives that aims at simplifying the deployment of JavaScript application. Since most of the developers have experience on languages other than JavaScript there is the need of porting existing application to the Web. Projects like *Emscripten* and *Google Web Toolkit* offer the possibility to write the code directly in C,C++ (*Emscripten*) or Java (*Google Web Toolkit*) and compile it into pure, and optimized, JavaScript.

Emscripten, by Mozilla, is a LLVM-to-JavaScript compiler. It takes LLVM bitcode (which can be generated from C/C++ using Clang, or any other language that can be converted into LLVM bitcode) and compiles that into JavaScript. Since it is a compiler it offers multiple grades of optimization that reduce the size of the JavaScript file and speedup the computation. The website is full of demos of the ported application, including games, 2D/3D game engines, various libraries and also SQLite.

Google Web Toolkit, by Google, is a development toolkit for building and optimizing complex browser-based applications. Its goal is to enable productive development of high-performance web applications

¹² Without using strange interaction between Flash/Silverlight and the browser.

without the developer having to be an expert in browser quirks, XMLHttpRequest, and JavaScript.

COMMUNICATION is being empowered by introducing *WebSocket*, that enables full-duplex data exchange with the server, and Cross-origin Resource Sharing (**CORS**) that gives the developers the possibility to make Asynchronous JavaScript and XML (**AJAX**) requests to "foreign" servers (other than localhost) without the need of a proxy for forwarding the requests.

The availability of all those technical API gives us the possibility to create a system capable of performing any *Human* or *Automatic* computation task without the need of external plugins. We used all the features of HTML5 for the computational side of our framework. For instance WebSocket are used for real-time task monitoring and **CORS** are used within the Task implementation to request any external data needed by the application.

2.2.1 HTML5



Figure 11.: Official HTML5 logo & unofficial CSS3 logo.

When speaking of HyperText Markup Language (**HTML**)5 one, usually, is not only focusing on the markup language but on a set of web technologies and specifications strictly related to **HTML**5. This set of technologies includes the **HTML**5 specification itself, the Cascading Style Sheets (**CSS**)3 recommendations and a whole new set of JavaScript APIs. So, first things first, let us point out the differences:

HTML5 refers to a set of semantic tags (like `<footer>`, `<header>`, `<article>`, ...), media tags (like `<video>` or `<audio>`) and the so called Web Form 2.0 alongside with all the "old" tags inherited from **HTML**4. These tags help developers to give semantics to the website they make, so the websites can be better understood

by search engines or HTML parsers (like those used for reading the site for blind people).

CSS3 refers to the presentation layer of the pages. Here are introduced specifications including image effects, 3D transformation, new tag selectors, form element validation, etc. The specifications take care also of the new devices (like smartphones and tablets) giving the user the media queries to examine the media (screen, print, aural) and provide different **CSS** rules.

JS refers to the JavaScript with a new set of API for interacting with the new media elements and other tags, as long as API for concurrent computation, real-time communication, offline storage, etc.

With the advent of **HTML5**, like any new technology, many problems were resolved and many others have been created. The main issue with using **HTML5** is the browser compatibility and browser-specific methods.

When browsers start implementing some **HTML5** draft feature, since they are not fully standardized¹³, they prevent the pollution of the DOM by prefixing the standard method with a browser specific prefix¹⁴ (i.e. `requestAnimationFrame` can become `mozRequestAnimationFrame` or `webkitRequestAnimationFrame`). This prefixing is particularly common in the **CSS3** where things becomes awful¹⁵.

To avoid browser inconsistency there are plenty of JavaScript frameworks for every purpose. Frameworks like *jQuery* provide a layer of abstraction between browser-specific code and the user, giving developers fallbacks for the most common API and additional features not covered by the standard implementation. Other frameworks like *Modernizr* give developers the ability to test if some **HTML5** feature is available in the currently used browser and provide a general fallback system for dynamically load polyfills¹⁶.

Now are presented the main **HTML5** features to better understand how they can be used in this System.

CANVAS Let's start with the official definition¹⁷

The canvas element provides scripts with a resolution-dependent bitmap canvas, which can be used for rendering graphs, game graphics, or other visual images on the fly.

¹³ In fact HTML5 (at the time of writing) is not yet standardized, its still a draft. See <http://www.w3.org/TR/html5/>

¹⁴ o: for Opera, ms: for Internet Explorer, moz: for Firefox, and webkit: for the WebKit based browser (Chrome and Safari)

¹⁵ See CSS animation or gradients for example.

¹⁶ A polyfill is a JavaScript library or third part plugin that emulates one or more HTML5 feature, providing websites to have a consistent behavior.

¹⁷ Got from the specs: <http://www.w3.org/TR/html5/the-canvas-element.html#the-canvas-element>

So the canvas element is basically a *Canvas*, like the name says, where one can *paint* anything. On top of this, the canvas element gives the developers access to the underlying raw pixel data. Also in the canvas element you can *draw* images taken from a `` tag or a frame taken from a `<video>` tag.

As one can see now we have all the tools we need to perform image analysis or video manipulation within the browser. Obviously there are plenty of JavaScript libraries that facilitate the whole process of filtering or, in general, image manipulation (like [Pixastic](#) or [Camanjs](#)). Other libraries give you the tools to create diagrams or charts on the fly (like [Raphaël](#) or [Processingjs](#)).

The canvas element also provides a 3D context to draw and animate¹⁸ high definition graphics and models using the WebGL API. This API is maintained by the [Khronos Group](#) and is based on OpenGL ES 2.0 specifications. On top of these API there are a lot of libraries¹⁹ made to facilitate development of 3D applications. One of the the most used is the [Three](#) JavaScript library, that can be used for creating and animating 2D or 3D scenes within the canvas element.

WEBSOCKET The WebSocket is an API interface for enabling bi-directional full-duplex client server communication on top of the Transmission Control Protocol (TCP). It enables real-time communication between clients and servers, allowing servers to **push** data to the clients and obtain *real* real-time content updates.

Like many other [HTML5](#) features on top of WebSocket a library that provides easy access to these functionality as long as fallbacks for old browsers was built. [Socket.io](#) provides a single entry-point to create a connection to the server and manage the message exchange, providing fallbacks²⁰ to ensure cross-browser compatibility.

WEBWORKERS A problem that rises when coding load intensive JavaScript application is the single thread nature of the language. Every script runs in the same thread of the browser window/tab. This can lead to some unwanted behavior (like browser freezing or a warning dialog that that alerts the user as in [Figure 12](#)). To solve this problem [Jenkin, 2008](#) proposed a timed-based programming structure that ensures that the code runs without any browser warning or freezing and also offers the developer to tweak the performance of the script by dynamically adjusting the interval between the step execution. This method leverage on the `setTimeout` function of JavaScript in order to split code into timestep-driven code chunks to execute. Here is an example of loop translated into a timer-based loop:

¹⁸ Animations are not natively supported, must be coded separately.

¹⁹ For a reference see http://en.wikipedia.org/wiki/WebGL#Developer_libraries

²⁰ If WebSocket, are not available the library can use Adobe® Flash® Socket, AJAX long polling, AJAX multi-part streaming, Forever Iframe and JSONP Polling




Figure 12.: The slow script dialog.

```
while condition do
| ...do something...
end
```

```
procedure STEP
| ...do something...
| if condition then
| | setTimeout(STEP,
| | delay)
| end
```

Obviously this is not the solution to the problem, it is a hack that tricks the browser.

WebWorkers offer a simpler solution. They provide a simple, yet powerful, way of creating *threads* in JavaScript. The official definition says:

The WebWorkers specification defines an API for running scripts in the background independently of any user interface scripts. This allows for long-running scripts that are not interrupted by scripts that respond to clicks or other user interactions, and allows long tasks to be executed without yielding to keep the page responsive.

The core concept behind WebWorkers is the *Worker*. A *Worker* is a piece of JavaScript code that runs in parallel to the main thread and is able to send and receive messages (just like normal threads).

STORAGE When web developers think of storing anything about the user, they immediately think of uploading to the server. **HTML5** changes that, as there are now several technologies allowing the **RIA** to save data on the client device.

HTML5 supports a number of storage techniques able to store data within the browser to be accessed later. Here is a simple list with the principal features:

WEB STORAGE is a convenient form for offline storage. It uses a simple key-value mapping for storing data persistently on the browser.

WEB SQL DATABASE is an offline SQL database, usually implemented using SQLite, a general-purpose open-source SQL engine.

INDEXEDDB is a nice compromise between Web Storage and Web SQL Database. Like the former, it is relatively simple and like the latter, it's capable of being very fast. It uses the same mapping as *Web storage* and indexes certain fields inside the stored data.

FILESYSTEM API, as the name says, offers the ability to manipulate the file system of the host.

OFFLINE STORAGE In this category falls the application cache. The application cache is controlled by a plain text file called a manifest, which contains a list of resources to be stored for use when there is no network connectivity. The list can also define the conditions for caching, such as which pages should never be cached and even what to show the user when he follows a link to an uncached page.

If the user goes offline but has visited the site while online, the cached resources will be loaded so the user can still view the site in a limited form. Here is a simple cache file:

```

CACHE MANIFEST

# This is a comment

CACHE:
/css/screen.css
/css/offline.css
/js/screen.js
/img/logo.png

http://example.com/css/styles.css

FALLBACK:
/ /offline.html

NETWORK:
*
```

2.2.2 WebCL

With the advent of General-purpose computing on graphics processing units (**GPGPU**), the spreading of multi-core CPUs and multiprocessor programming (like OpenMP) we can see emerging an intersection in parallel computing. This intersection is known as **heterogeneous computing**. There are initiatives aimed at enabling numeric calculation, even complex, on the web client. Open Computing Language (**OpenCL**) is a framework for heterogeneous computing and Web Computing Language (**WebCL**) is a porting of this technology to the web.




Figure 13.: OpenCL execution flow.

OpenCL execution is based on *kernels*, special functions written in a language based on C99²¹ with some extensions. These *kernels* can be compiled at build-time or at run-time and also have a rich set of built-in function, e.g. *cross*, *dot*, *sin*, *cos*, *pow*, *log*, etc.




Figure 14.: OpenCL device composition.

An OpenCL device (like CPU or GPU) is composed by (see Figure 14):

- A collection of one or more **compute units**. This concept is similar to the *cores* of a standard CPU.
- A *compute unit* is composed by one or more **processing elements**. This concept is similar to the *threads*.
- Processing elements execute code as Single instruction multiple data (**SIMD**) or Single process/program multiple data (**SPMD**).

²¹ A programming language dialect for the past C developed in 1999 (formal name ISO/IEC 9899:1999)

Once the *kernels* have been created the flow for executing an [OpenCL](#) program can start. Here is the list of the actions performed to run code on [OpenCL](#) enabled computer:

1. Query host for [OpenCL](#) devices.
2. Create a context to associate [OpenCL](#) devices.
3. Create programs for execution on one or more associated devices.
4. From the programs, select kernels to execute.
5. Create memory objects accessible from the host and/or the device.
6. Copy memory data to the device as needed.
7. Provide kernels to the command queue for execution.
8. Copy results from the device to the host.

Currently there are two main implementations of the [WebCL](#) API, one supporting Windows and Linux using Firefox (held by Nokia), and the other, held by Samsung, that brings this interface to WebKit browsers. Speaking of performance improvements, [Table 2](#) shows the performance data from the TIZEN™ developer conference held in May 2012:

Table 2.: Samsung WebCL performance compared to pure JavaScript.

Demo name	JavaScript	WebCL	Speed-up
Sobel filter (with 256x256 image)	~ 200ms	~ 15ms	13x
N-body simulation (1024 particles)	5-6fps	75-115fps	12-23x
Deformation (2880 vertices)	~ 1fps	87-116fps	87-116x

3 | CONCEPTUAL MODEL

In this chapter it is presented a framework for task distribution and execution able to cover all the dimension defined in [Table 1](#). First there is introduced a methodology to follow in order to pass from the model to the actual application implementation, available in [chapter 5](#). The components that will compose the framework are founded on a conceptual model, describing the logical structure of the data processed by the application.

The conceptual model is composed of three parts: the *Data Model*, the *Architectural Model* and the *Execution model*. Eventually there are explained the pluggable strategies available for this framework.

THE DATA MODEL describes the data structures used to support the framework.

THE ARCHITECTURAL MODEL describes the reference architecture of the framework.

THE EXECUTION MODEL focuses on the execution model of the task.

THE PLUGGABLE STRATEGIES describes the possible configuration points of the framework.

3.1 DEVELOPMENT METHODOLOGY

During the development process of a framework for web-based human and machine computation the first issue that needed to be resolved was the description of a suitable model for supporting such framework. The design of a model involves the definition of the data structures that will be used by the workflow, and of course the definition of the workflow itself.

3.1.1 Data Design

In the design process for the data model we faced the problem of creating a suitable underlying model for all the application need that we have to cover, ranging from pure Human Computation task to *parasitic computing*. This wide range of possible applications and the need of great flexibility demanded to the system, led us to the creation of a meta model for almost all the data structures used in the framework. The metamodel was subdivided into four parts: the *Workflow*, the *Task Data Model*, the *Task Model* and the *Task Execution Model*.

THE WORKFLOW MODEL describes the flow of the execution of the Task within the framework, the relationships that need to be taken into account to correctly execute a sequence of task (called *Work*).

THE TASK DATA MODEL describes the data model for each Task, including, if present, the field dependencies and if a field belongs to the input or the output set.

THE TASK MODEL describes the actual data of the Task, called *Objects*, for the Task.

THE TASK EXECUTION MODEL describes the actual execution step for each Task, providing information on which user is performing which task, as well as the data associated to the execution itself.

3.1.2 Workflow Design

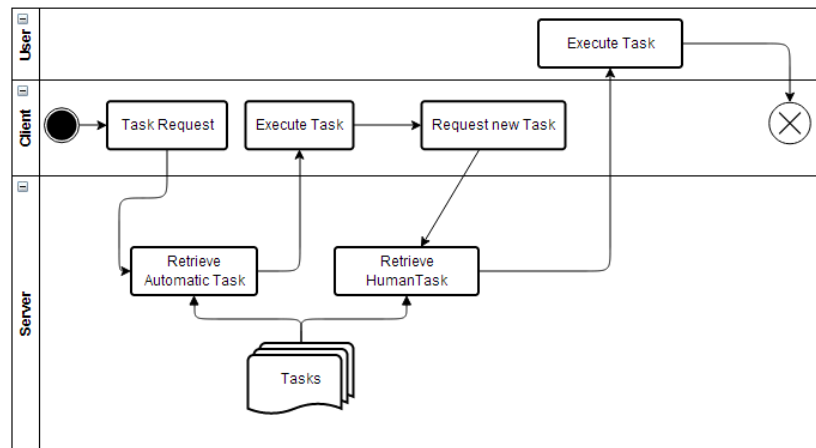


Figure 15.: The conceptual Workflow handled by the framework.

In the *Workflow Design* we must describe, at a conceptual level, how our framework deals with the client and how the tasks are executed. During the *Data Design* we stated that the tasks can have relationships and thus our framework must be able to handle different Task types seamlessly. In [Figure 15](#) is depicted a typical flow of execution of a *Work* (a composition of Tasks) with multiple task types.

Our Workflow model must be able to handle the interleaving of human and automatic task without any human intervention. This must also take into account the constraints defined during the creation step and manage the execution of tasks accordingly.

A feature that the framework is required to handle is also the case where the application type is not strictly human or strictly automatic. We can have scenarios where the task is a blend of both human and automatic interactions. [GWAP](#) are a case where the automatic com-

putation and the human interaction can be blended together. Other scenarios can be the validation of the automatically computed results as we presented in our use case (see 4.3).

To guarantee the needed flexibility the framework must be able to include third-part logic in charge of managing certain steps of the execution (like Task planning).

3.1.3 Framework Design

In the design process of the framework we took all the conceptually designed data and workflows, as well as the requirements, and we merged them together in order to create the whole structure for supporting our framework.

During the **Data Design** we pointed out the data structures that need to be managed. The *Workflow Model* needs to specify the relationships between the modules and the "connection points" where the third part softwares are plugged. On top of that, every Task must have defined its set of input and output data. The *Task Data Model* specifies all the data structures of the task, building the metamodel for each Task. The metamodel contains all the information on the structure of the data model for a Task. In the *Task Model* there are contained all the actual data for each Task, called Objects. Each object represents a "row" of the data and is, in turn, associated to a Task. The *Task Execution Model* contains all the information on the actual execution of a Task. It stores all the information on the user that performed the task and the device used.

The **Workflow Design** allows us to create the software able to handle all the different types of archetype application (e.g. Human computation, Automatic computation) giving also the guidelines on how to manage the interaction between subsequent Task execution.

3.2 DATA MODEL

3.2.1 WorkFlow

The WorkFlow embodies all the data associated to the *flow* of Task that need to be executed in order to complete a set of Tasks. The tasks can belong to different archetypes (e.g. Human, Automatic) and have their own set of assigned objects. The *Workflow* is composed of:

THE WORK: is the abstract representation of a set of Tasks related to each other.

THE FLOW: describes how the Tasks, belonging to a Work, are related to each other, also defining the type (e.g. dependency, parallel) of relation between such Tasks.

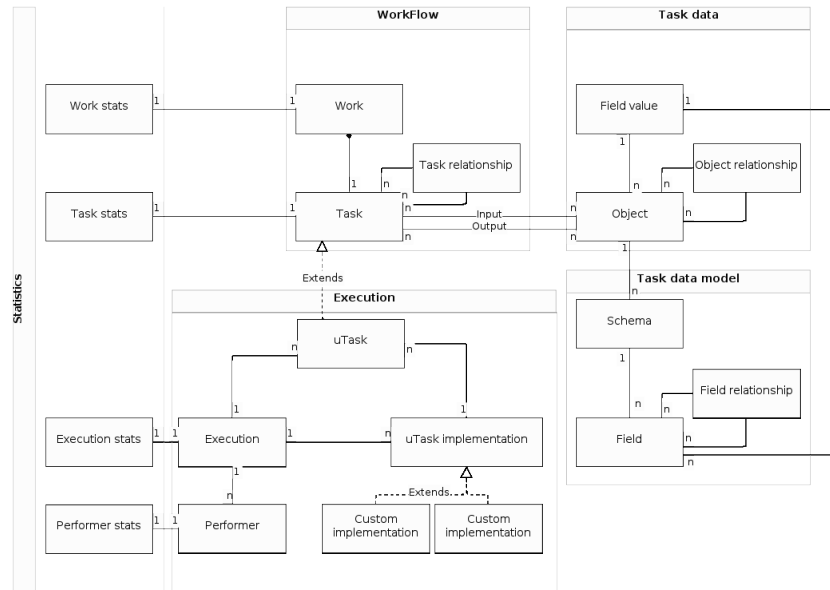


Figure 16.: Overall Data Model schema, with logical subdivision.

THE TASK: is the representation of a general activity that can be performed by the framework.

The Work

The *Work* represents a set of related *Tasks* finalized to the execution of some kind of data manipulation or analysis on a set of input *Objects*. The result of the execution produces output *Objects* described by a *Schema*. The *Work* is defined by:

- A **Name** that identifies the *Work*.
- **Constraints** used to bind some aspects of the *Work* or to make decision for giving less or more priority to this work with respect to the others. Example of constraints are: Due date, Performer skills, Max execution time, etc.
- **Input** data, defined by a *Schema* and the associated *Objects*. To keep the model as general as possible no assumptions are made on the *Schema* type (relational, graph, etc.).
- **Output** data, defined as an extension of the input schema (sharing the same schema type).
- A set of **Task** composing the *Work*. Their orchestration is made at design time, by specifying a *Flow*.

The Task

The *Task* is the central unit of computation. It represents a single activity focusing on a particular action. The activity must not be an




Figure 17.: Conceptual organization of Work, Task and μ Task.

atomic operation, like in an algorithm a single function call, but it should be focused on a single purpose. For example tagging an image can be considered as a single Task. Otherwise tagging and validating an image should be divided in two separated tasks if possible. The flexibility of the framework allows the creation of any type of Task, the previous statement is only an advice to better separate different activities into different Task. As we mentioned our system can seamlessly handle complex Task involving multiple activities, such as whole [GWAP](#) game into a single Task. A Task is characterized by:

- A **Name** that identifies the Task.
- **Input** data, defined by a *Schema* and the associated Objects. To keep the model as general as possible no assumptions are made on the *Schema* type (relational, graph, etc.). Usually the Schema is a projection of the Schema of a Work. Like the Schema the input Objects usually are a selection of the Work input Objects.
- **Output** data, defined as an extension of the input schema (sharing the same schema type).
- A Task **type**¹ defining, at abstract level, what kind of data manipulation will be performed by a Task. These categorizations are taken from [Bozzon, Brambilla, and Ceri, 2012](#), here there are a few:
 - Like
 - Order
 - Classify
 - Add
 - ...

¹ An assumption is made to make the list fit all the possible abstract tasks our framework is able to handle.

Each Task type is defined by:

- I/O relationship, defining, at abstract level, how the Task transforms the data and the schema.
- A default implementation.
- A **Status** encoding the current state of the Task. A Task, can have only one of the following statuses at point of its lifecycle:
 - *Planning-Input*: the Task has been created, has a *Schema* and *Object data* associated and a defined Task *type*.
 - *Planning- μ Task*: a set of μ Task has been associated to the Task.
 - *Planning-Assignment*: a set of *Performers* has been selected to execute the μ Task.
 - *Wait*: Task planned, μ Task ready for execution.
 - *Running*: μ Tasks are running.
 - *Ended*: all the μ Tasks have completed their execution.
- A set of **Subscribers** able to receive updates on the Task execution. The Subscribers are declared during the Task creation step and the updates that they will receive are about change on the Status of the Task.
- A set of **Execution constraints** used for prioritizing the Task among others or to modify the standard behavior of the Task. The constraints that can be used are the same defined for the Work in 3.2.1.
- **Configuration data** are application specific data used to configure the behavior of the Task once it is executing. For instance we can store here the classes for a classify Task type.
- μ **Tasks** are the instances of a Task assigned to one or more Performers.
- An **Aggregation** function, in charge of collecting the μ Tasks results and generate the Task output.
- μ **Task planning** strategy, in charge of defining how many μ Tasks create for a given Task and associate the right portion of input objects to each μ Task. For example total disjunction, redundancy, partial overlap, etc.
- **Performer assignment** strategy able to assign *Performers* to μ Tasks. Some strategies can be: manual, random, most reliable, etc.
- μ **Task implementation** strategy in charge of routing the correct μ Task implementation for each μ Task execution. The routing can be *fixed* for everyone, can be done according to the *user-agent* (e.g. Browser) or other conditions.

- A **Task planning** which embodies the functionalities of μ Task *planning* strategy, *Performer assignment* strategy and μ Task *implementation* strategy and whose purpose is to decide the logic behind the invocation of those strategies.
- A **Task control** strategy is able to control the status of the Task and if needed perform corrective actions.
- An **Emission policy** specifying which *Subscriber* needs to be notified of a Task change in *Status*.

The **relationships** that may exist between Tasks are materialized into the *Task Relationship* object. This object contains all the information related to the flow of the task execution during the execution of a Work. An example of a Task relation can be the validation of the results obtained by an Automatic computation task previously performed. In this case the relation is the simple sequence.

In a flow we can use *Control Structures* and *Variables*. The *variables* are included in the flow to control the behavior of the Work and can be modified by some conditions upon Task execution. For instance after obtaining Task results we found that they are under a certain threshold then we can set a variable to match this problem.

The *Control Structures* are common to all the Workflow managers and control how the Tasks are executed:

SEQUENCE: represents the normal flow of an application where one operation is executed after the previous is completed.

CHOICE: gives the possibility to made choice according to one, or more, *Variables*.

LOOP: Allows to execute some steps multiple times, according to a predefined value or a *Variable*.

PARALLEL: the steps of the flow are not executed in *Sequence*, allowing the parallelization of some steps.

3.2.2 Task Data Model

The Task Data Model contains all the information about the Task metamodel. The metamodel is organized in Fields and the Fields are grouped into a Schema. This data structure resembles the standard DBMS schema organization, where we have a Schema defined as a collection of Fields with possible relationship between them. The Task Data Model is composed by

THE SCHEMA: is the abstract representation of a table structure.

THE FIELD: contains all the information of the field type and relationships.

The Schema

The Schema is used as a container of Fields that composes the meta-model of the Task data. The Schema is defined by:

- A **Name**
- A **list of Fields** defining the structure of the Task data.

The Schema is in **relation** with the actual data instances of the Task, called Objects, and the Fields it is composed of.

The Field

The Field contains information on the type of the data it contains and information on the relation between two, or more, Fields (like the type of dependency). The Field is composed by:

- A **Name** that identifies the Field.
- A **type** defining the type of the data that this field contains, for instance string, number, etc.
- A set **related fields**. The relationship between the fields is defined by the *relation* attribute.
- A **relation** that specifies which type of relationship occurs among the *related fields*.
- The **list of field values** representing the data contained within this field in the Object structure.

The **relationships** existing between the Field are materialized using the *Field Relationship* table. Here we have all the information of the fields involved and on the type of the relationship (e.g. derivate of, copy of, calculated etc.). Each field has also a relationship with the corresponding instances in the *Field Values* table.

3.2.3 Task Data

The Task Data contains the actual data instance for each Task. The data are subdivided into *Field Values* and *Objects*. This division is made to mirror the structure presented in 3.2.2. The Objects represent the "tuple" of the table, while the Field Values contain the atomic information of a Field within a Tuple. This structure allows the direct access to Field Values without the need of parsing the whole Object element. As mentioned the Task Data is composed of:

THE OBJECT: a Task data instance, represented by a set of Field Values.

THE FIELD VALUE: represents the instance of a Field. As the Fields they can be related to each other.

The Object

The Object represents the Task data instances. They are composed of a set of Field Values that, joined together, form the Object. The object itself is composed by:

- A **Name**.
- A **list of Field Values** data.

Objects can have **relationships** among them. These relationships are materialized into the *Object Relationship* entity. This object contains the needed information used to identify a set of Objects belonging to a particular Task.

The Field Value

The Field Value contains the actual data of a Field. Since this object is used to represent any kind of data, it must be able to handle different types of information (e.g. string, BLOB, number, date, etc.). The Task Field is composed by:

- The field **Value**.
- A **related Object** used to identify to which Object this Field Value belongs to.
- The **Field** to which data belongs.

The Field Value has **relationships** with the Object to which it belongs and with the Field that defines the general type of the Field Value.

3.2.4 Task Execution

The Task Execution is used to collect information on the Task execution. Since a Task must be distributed to multiple nodes² the Task must be splitted in μ Task. Then we need to know what user is performing such μ Task, and, based on the device used, we can have multiple implementations of the same μ Task. All these information are subdivided into:

THE μ TASK: is a subclass of the Task. The μ Task is a portion of the Task executed by the end-user.

THE μ TASK IMPLEMENTATION: represents the real implementation of the μ Task.

THE PERFORMER: is the Human being in charge of the execution of the Task.

² We use the broad term node to identify both the user and the device on which the Task must be executed.

THE EXECUTION: represents a bridge object connecting the μ Task, the μ Task Implementation and the Performer to identify the execution of a μ Task using a μ Task implementation by a particular user.

μ Task

Since the same Task must be distributed to many users there is the need of a further subdivision of a Task into "atomic" activities (see [Figure 17](#)). The μ Task is a subclass of the Task that represents this subdivision. As a subclass it inherits all the information of the data and the schema from the parent Task. Usually μ Tasks have associated only a portion of the parent Task data objects. The μ Task is defined by:

- A **Name**.
- A set of **Execution constraints**, similar to the ones defined in [3.2.1](#). The difference is that these constraints bind the execution "client side". This means that with these we can control the computational load on the user machine according to some policy (like user preferences, device capabilities, etc).
- **Input** data, as a subset of the Task input Objects.
- **Output** data, with the same schema as the related Task output Objects.
- A **list of Properties**, defined as name-value pairs, having domain specific meaning.

The μ Task is in **relationship** with its μ Task implementations. This relation is useful when choosing a suitable implementation to execute on a particular device.

μ Task implementation

The μ Task Implementation represents the application logic and presentation delivered to a user to perform a Task. The Task has a default implementation and the other implementations can be supplied to customize the Task execution. Each μ Task implementation is in **relation** with the corresponding μ Task.

Performer

The Performer is an object used to characterize a human being in the Data Model. The user is characterized with properties that defines his/her skills on a certain field (like music, style, etc.) and some other attributes used for profiling. The Performer object is defined by

- A **Name**.
- **Demographic** information.

- **Performance** information.
- **Trustworthiness**.
- **Social properties**.

The Execution

The Execution is a bridge object that joins together μ Task, μ Task implementation and performer. This object represents a single instance execution of a Task by a user. Using this object we can identify cheaters or reward good Performers, or using it as a timeline we can see how the Performers respond to the Task and change it accordingly. An Execution object is defined by:

- A **Status** that identifies the current status of a μ Task implementation. The available statuses are:
 - *running*
 - *suspended*
 - *idle*
 - *ended*
- A **set of Execution data**.

Since it is a bridge object the Execution has **relationships** with μ Task, μ Task implementation and the Performer. These relationships identify the Execution itself.

3.2.5 Statistics

This Statistics model contains all the statistics and data related to the profiling of a Task. These data can be used by other components, such as the Task controller to take decision during the flow of a Task. The Statistics are about *Work*, *Task*, μ *Task*, *Performer*, etc. The data contained in these objects are:

- **Creation date**.
- **Total execution time**.
- **Average number of Performers per hour**.
- **Last execution**.
- Etc.

3.3 ARCHITECTURAL MODEL

During the Design Process we faced the problem of finding a suitable Architectural Model able to support all the requirements, in terms of flexibility and pluggability, raised during the Process Design. In our

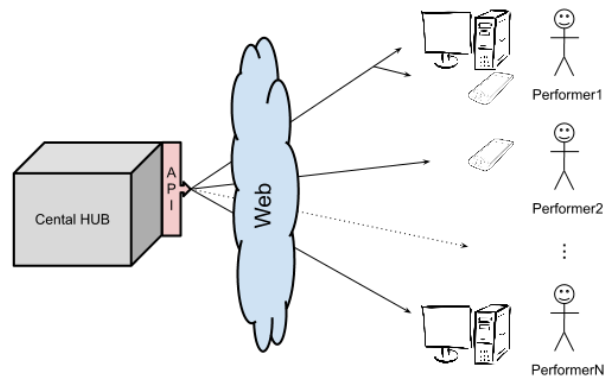


Figure 18.: Reference architecture.

model we use as a reference architecture the one depicted in [Figure 18](#). The *Central Hub* is used to manage all the data exchange between to the nodes and the hub itself, orchestrating all the communication flow.

As described in [3.1.2](#), our framework has multiple configuration points used to customize the Task behavior. As described in [3.2.1 on page 24](#), for each Task we can identify seven configuration points:

μTASK PLANNING STRATEGY: defines how many μTask to create for each Task and associate the right portion of Objects to each of them.

μTASK IMPLEMENTATION STRATEGY: defines the logic behind the choice of a μTask implementation sent to a user.

PERFORMER ASSIGNMENT STRATEGY: in charge of choosing the users who are suitable to execute a certain Task.

TASK PLANNING STRATEGY: orchestrates the invocation of the μTask planning strategy, the μTask implementation strategy and the Performer assignment strategy.

TASK CONTROL STRATEGY: defines a controller for the Task, able to check the status, and if needed, perform corrective actions.

AGGREGATION FUNCTION: is in charge of joining the results obtained from the μTasks execution and create a Task output.

EMISSION POLICY: is used to rule notifications sent to the Subscribers.

All these configuration points are described in [3.5](#).

3.4 EXECUTION MODEL

During the Framework Design stage we found how our framework processes a Task and how this behavior can be configured. In [Figure 19](#) is depicted the typical flow of a Task execution. As one can notice

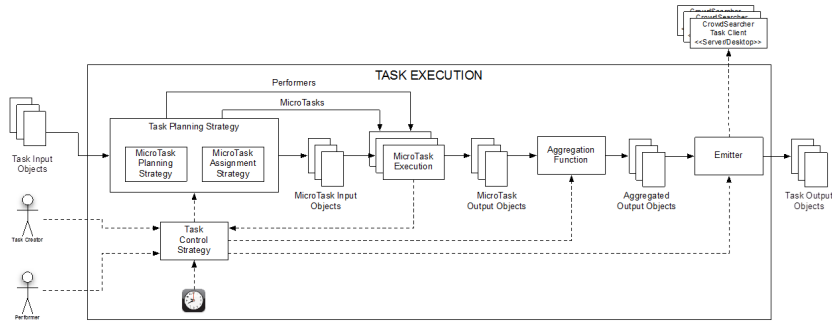


Figure 19.: Representation of the Task execution flow.

all the controlling strategy is handled by the *Task Control strategy* that, accordingly to the *Task Planning strategy*, tweaks the execution of a Task. By changing the Task Planning strategy we can obtain four different execution modalities, summarized in Table 3:

STATIC EXECUTION: where both the μ Task planning and the Performer assignment is performed statically at Task design time.

STATIC μ TASK PLANNING & DYNAMIC ASSIGNMENT: in this scenario the μ Task are planned at design time and the Performer assignment can be made at "any" time after the creation of the Task.

DYNAMIC μ TASK PLANNING & STATIC ASSIGNMENT: in this scenario we first Assign Performers to μ Task and then we Plan the μ Tasks³.

DYNAMIC μ TASK PLANNING & DYNAMIC ASSIGNMENT: where both the strategies can be specified after the Task creation³.

Table 3.: Task planning vs. Task Assignment.

	Static	Dynamic
Static	static-static	static-dynamic
Dynamic	static-dynamic	dynamic-dynamic

STATIC EXECUTION This execution mode represents the simplest scenario possible. Static execution means that both the *Task Planning strategy* and the *Performer Assignment strategy* are executed once at Task creation time. The Task Control strategy duty is to control if the *Constraints* are verified and monitor the Task status. Here is the list of the default actions performed by the Task Control strategy:

- **Stop** a Task if constraints are met.
- **Invoke** the *Aggregation function* at the end of the Task execution.
- **Notify** the *Subscribers* about the Task execution.

³ This scenario presents some pitfalls, see 3.4 to further explanations.

STATIC μ TASK PLANNING & DYNAMIC ASSIGNMENT In this execution mode the user wants to define all the μ Task and then assign Performers later, usually with a custom Performer Assignment strategy. In this scenario the μ Task Planning strategy is are planned once at creation time and the assignment can be performed at any time. In this scenario the *Task Control* strategy invokes the *Performer Assignment* strategy ensuring that the constraints are met. The *Task control* strategy can also decide to reassign *Performers* to μ Task later during the execution. Here is a list of the default actions performed by the *Task control* strategy:

- **Stops** a Task if constraints are met.
- **Invokes** the *Aggregation function* at the end of the Task execution.
- **Notifies** the *Subscribers* about the Task execution.
- **Invokes** the *Performer assignment* strategy to bind μ Task to *Performers*.

DYNAMIC μ TASK PLANNING & STATIC ASSIGNMENT In this scenario we want the user requesting the Task to be able to dynamically assign the μ Task to the Performers. This scenario can lead to some pitfalls. For instance if the Performer assignment is executed and there are no μ Task defined then we can have unhandled behavior⁴. Supposed we not fall into these pitfalls, in this scenario the *Performer Assignment* strategy is called at creation time and the *μ Task Planning* strategy can be called at any time. Summing up, the *Task Control* strategy:

- **Stops** a Task if constraints are met.
- **Invokes** the *Aggregation function* at the end of the Task execution.
- **Notifies** the *Subscribers* about the Task execution.
- **Invokes** the *Task planning* strategy to *re-plan* μ Task or **Create** new μ Tasks.

DYNAMIC μ TASK PLANNING & DYNAMIC ASSIGNMENT This scenario covers the most flexible use-case. Here the user can create an "empty" Task and later assign the μ Tasks and Performers. Also the strategies can be called more than once to adapt to the user preferences. For instance if the user that created the Task notice that the results are biased then he/she can decide to expand the Performers set and re-plan the μ Tasks. Summing up, the *Task Control* strategy:

- **Stop** a Task if constraints are met.
- **Invoke** the *Aggregation function* at the end of the Task execution.
- **Notifies** the *Subscribers* about the Task execution.

⁴ This behavior is not managed by the framework. A custom logic must be supplied to handle such execution.

- **Invoke** the *Performer assignment* strategy to bind μ Task to *Performers*.
- **Invoke** the *Task planning* strategy to *re-plan* μ Task or **Create** new μ Task

3.5 PLUGGABLE STRATEGIES ASSIGNMENT

As pointed out during the Process Design and briefly introduced in the Architectural Model, our framework must be able to handle custom strategies. These strategies (see 3.3) are used to grant that our model is flexible enough to support all the application archetypes defined in Table 1.

This section explains the purposes of each strategy and what are the main properties for each strategy.

3.5.1 μ Task Planning strategy

The μ Task Planning strategy is a pluggable logic devoted to the creation, or deletion, of μ Task. Eventually this strategy must, during the creation time of a μ Task, bind a corresponding set of Task Objects to the μ Task. This binding produce a set of μ Task with the corresponding Task Objects. A Task planning strategy is defined by:

- A set of **Constraints** that rules the execution.
- A **Planning policy** that can be defined at:
 - DESIGN TIME:** the assignment is made at design time during the creation phase. After the planning is done it can be modified only
 - DYNAMIC:** the planning is done at least once, using a provided set of input *Objects*. The planning can be further invoked due to:
 - *Variations* in the state of the Task. i.e. an object can be reassigned to another μ Task.
 - *Addition* of new *Objects* through the API.

Note that the addition of new μ Task can be performed using the API but usually do not involve the invocation of a μ Task planning strategy.

3.5.2 Performer Assignment strategy

The Performer assignment strategy is a pluggable logic devoted to the assignment *Performers* to μ Task. Once we have a set of μ Task we can assign to them the suitable Performers. The Performers are chosen according to some criteria like their skills or the place they live. A Performer assignment strategy is composed by:

- A **set of Constraints** defined on top user-specific statistics (e.g. do not assign more than 1 μ Task per hour).
- A **list of routes** that, by matching the description of a *Performer*, decide if a μ Task can be assigned to that *Performer*.
- An **Invitation Strategy**, which defines how performers should be invited to execute the task.
- An **Assignment policy** that can be:
 - STATIC:** the assignment can be performed both at design time and runtime, but it must be explicitly called by a control rule.
 - DYNAMIC:** the assignment is performed at least once and can be invoked multiple times later according to *Variables* that can change over time.

3.5.3 μ Task Implementation strategy

μ Task Implementation strategy is a pluggable logic in charge of selecting a suitable μ Task implementation for an *Execution*. This strategy is invoked before the execution of a Task on a device. Based on the Performer preferences or on the device characteristics a suitable μ Task implementation is routed to the user.

3.5.4 Task Planning strategy

The Task Planning strategy embodies the functionalities of a μ Task Planning strategy and Performer Assignment strategy, deciding the logic by which the two strategies should be invoked. This strategy can be used to manage the re-planning of the μ Tasks or to call the Performer Assignment strategy. The invocation of this strategy is controlled by the Task Control strategy that can call this strategy upon changes in the Task status.

3.5.5 Task control strategy

The Task control strategy is a pluggable logic devoted to verifying the status of a Task, possibly against the assigned constraints. This logic can be executed:

- **Once** when the Task ends. For instance when all the μ Tasks are executed to re-plan the execution.
- According to a **temporal schedule** (i.e. every x minutes, once a day, at noon, etc.).
- Every time a μ Task is **executed**.

Among the corrective actions available to the Task controller we have:

- The **re-planning** of the task, also with the creation of new μ Task.
- The **re-assignment** of μ Task to *Performers*.
- **Deletion** of executed μ Task.
- **Change** the properties of an executed μ Task. For instance we can set the results as invalid if we have spotted a cheater.
- **Re-execution** of the entire Task.
- **Halting** the Task.
- etc.

3.5.6 Aggregation function

An Aggregation function is a pluggable logic devoted to joining the results obtained with the μ Tasks execution. These results are merged according to a Task specific logic in order to produce the Task output result. An aggregation function can be as simple as a *Sum* or an *Average* but can also be more complex. For instance we can gather the results obtained, perform filtering operation and eventually produce an image.

Emission policy

The Emission policy controls how the *Subscribers* are notified about the status changes in a Task. This logic can be executed:

- **Once** the Task ends.
- According to a **temporal schedule**.
- Every time a task is **executed**.

4 | USE-CASES

In the previous chapter we introduced a framework for web-based human and machine computation, able to handle different types of application archetypes. In this chapter we present the use-cases use to test this framework under different scenarios. Since we wanted to test the framework with all the possible application archetypes we implemented three use-cases: *Automatic*, *Human* and *Hybrid*:

AUTOMATIC: the automatic use-case is used to simulate a *Distributed Computing* application. In this scenario we compute the SIFT algorithm on an image and return the obtained keypoints to the server.

HUMAN: the human use-case is used to test if we can perform Human Computation with our framework. To test this scenario we implemented a text disambiguation application.

HYBRID: the hybrid use-case is used to test both the automatic and the human scenarios. Here we implemented a [GWAP](#) on top of a face recognition algorithm.

We focused on the implementation of the *Voluntary* scenario, see [Table 1](#), because the *involuntary* case is straightforward to obtain. At the end of every use-case will be presented, if possible, a simple benchmark/metric where the use-case results are compared with the ones obtained with the available tools.

4.1 AUTOMATIC

Automatic use case

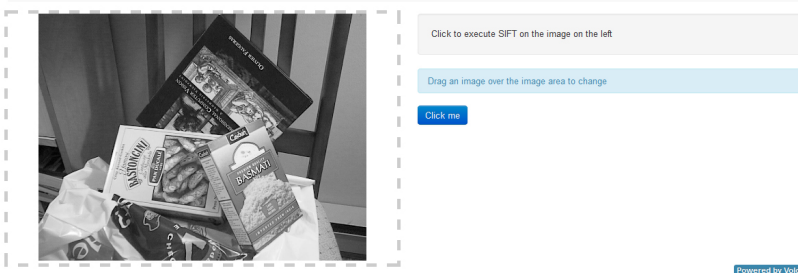


Figure 20.: Interface of the automatic use-case.

The Automatic use-case aim at executing computationally intensive applications on the user device (e.g. browser). This scenario allow

us to test if our framework is able to act as a *Distribute Computing* platform, as described in 2.1.2.

For this use-case we choose to implement an image recognition algorithm. We used an image recognition algorithm because they are commonly used by search engines to find images with similar features. Also these algorithms usually have high requirements in terms of CPU load and resources usage, so they are the perfect candidate for our purposes. The algorithm we choose to implement was Scale-Invariant Feature Transform (**SIFT**), one of the most widely adopted for this purpose.

4.1.1 **SIFT** algorithm

The **SIFT** algorithm is composed of four sequential steps: Scale-space extrema detection, Keypoint localization, Orientation assignment and Keypoint descriptor.

SCALE-SPACE EXTREMA DETECTION: is the stage where the interest points are detected.

KEYPOINT LOCALIZATION: is used to filter the unstable keypoints and keep only good keypoints.

ORIENTATION ASSIGNMENT: compute orientation and magnitude for each keypoint.

KEYPOINT DESCRIPTOR: generates the keypoints descriptors.

Scale-space extrema detection

In this step we generate the so called scale space representation of the image. In order to do this we need to convolve the image $I(x, y)$ at different scales $k\sigma$ with a varying Gaussian kernel $G(x, y, k\sigma)$ obtaining:

$$L(x, y, k\sigma) = G(x, y, k\sigma) * I(x, y)$$

Then the difference of successive blurred images are taken

$$D(x, y, \sigma) = L(x, y, k_1\sigma) - L(x, y, k_2\sigma)$$

This step produce the Difference of Gaussian (**DoG**) images, the first *Keypoints* are identified as local minima/maxima of the **DoG** image across scales.

Keypoint localization

In this step the *keypoints* are filtered to remove unstable points and keep only the good ones. This step can be further subdivided into 3 stages:

- *Interpolation* of nearby data for accurate position.
- *Discarding* low-contrast keypoints.
- *Eliminating* edge responses.

INTERPOLATION OF NEARBY DATA FOR ACCURATE POSITION The interpolation is done using the quadratic Taylor expansion of the DoG $D(x, y, \sigma)$ scale-space function, with the candidate keypoint as the origin. This Taylor expansion is given by:

$$D(\mathbf{x}) = D + \frac{\partial D}{\partial \mathbf{x}} \cdot \mathbf{x} + \frac{1}{2} \cdot \mathbf{x}^T \cdot \frac{\partial^2 D}{\partial \mathbf{x}^2} \cdot \mathbf{x}$$

where D and its derivatives are evaluated at the candidate keypoint and $\mathbf{x} = (x, y, \sigma)$ is the offset from this point.

DISCARDING LOW-CONTRAST KEYPOINTS To discard the keypoints with low contrast, the value of the second-order Taylor expansion $D(\mathbf{x})$ is computed at the offset $\hat{\mathbf{x}}$. If this value is less than 0.03, the candidate keypoint is discarded. Otherwise it is kept, with final location $\mathbf{y} - \hat{\mathbf{x}}$ and scale σ , where \mathbf{y} is the original location of the keypoint at scale σ .

ELIMINATING EDGE RESPONSES The DoG function will have strong responses along edges, even if the candidate keypoint is not robust to small amounts of noise. Therefore, in order to increase stability, we need to eliminate the keypoints that have poorly determined locations but have high edge responses. For poorly defined peaks in the DoG function, the principal curvature across the edge would be much larger than the principal curvature along it. Finding these principal curvatures amounts to solving for the eigenvalues of the second-order Hessian matrix, \mathbf{H} :

$$\mathbf{H} = \begin{bmatrix} D_{xx} & D_{xy} \\ D_{xy} & D_{yy} \end{bmatrix}$$

The eigenvalues of \mathbf{H} are proportional to the principal curvatures of D . The trace of \mathbf{H} , gives us the sum of the two eigenvalues, while its determinant yields the product. The ratio $\mathbf{R} = \text{Tr} \mathbf{H}^2 / \text{Det} \mathbf{H}$ can be shown to be equal to $(r + 2)^2 / r$, which depends only on the ratio of the eigenvalues rather than their individual values. It follows that, for some threshold eigenvalue ratio r_{th} , if \mathbf{R} for a candidate keypoint is larger than $(r_{th} + 1)^2 / r_{th}$, that keypoint is poorly localized and hence rejected.

Orientation assignment

In this step for each keypoint is assigned an orientation and a magnitude. This step is used to achieve *invariance rotation*. The magnitude $m(x, y)$ and orientation $\theta(x, y)$ are calculated as follows:

$$m(x, y) = \sqrt{(L(x+1, y) - L(x-1, y))^2 + (L(x, y+1) - L(x, y-1))^2}$$

$$\theta(x, y) = \tan^{-1} \left(\frac{L(x, y+1) - L(x, y-1)}{L(x+1, y) - L(x-1, y)} \right)$$

Keypoint descriptor

Previous steps found keypoint locations at particular scales and assigned orientations to them. This ensured invariance to image location, scale and rotation. Now we want to compute a descriptor vector for each keypoint such that the descriptor is highly distinctive and partially invariant to the remaining variations such as illumination, 3D viewpoint, etc. This step is performed on the image closest in scale to the keypoint's scale.

First a set of orientation histograms are created on 4x4 pixel neighborhoods with 8 bins each. These histograms are computed from magnitude and orientation values of samples in a 16x16 region around the keypoint such that each histogram contains samples from a 4x4 subregion of the original neighborhood region. The magnitudes are further weighted by a Gaussian function with equal to one half the width of the descriptor window. The descriptor then becomes a vector of all the values of these histograms. Since there are $4 \times 4 = 16$ histograms each with 8 bins the vector has 128 elements. This vector is then normalized to unit length in order to enhance invariance to affine changes in illumination.

4.2 HUMAN

In this scenario we want to create a completely human computation Task. To create this kind of application we decided to use word-sense disambiguation.

In computational linguistics, Word-Sense Disambiguation ([WSD](#)) is an open problem of natural language processing, which governs the process of identifying which sense of a word (i.e. meaning) is used in a sentence, when the word has multiple meanings. Research has progressed steadily to the point where [WSD](#) systems achieve sufficiently high levels of accuracy on a variety of word types and ambiguities.

A rich variety of techniques have been researched, from dictionary-based methods that use the knowledge encoded in lexical resources, to supervised machine learning methods in which a classifier is trained for each distinct word on a corpus of manually sense-annotated examples, to completely unsupervised methods that cluster occurrences of words, thereby inducing word senses. Among these, supervised learning approaches have been the most successful algorithms to date.

For the purposes of [WSD](#) there are plenty of online solution to use like [DBpedia](#)¹, [YAGO2](#)² or [Entitypedia](#)³.

[DBPEDIA](#) as described in [Auer et al., 2007](#), is is a community effort to extract structured information from Wikipedia and to make this

¹ <http://wiki.dbpedia.org/>

² <http://www.mpi-inf.mpg.de/yago-naga/yago/>

³ <http://entitypedia.org/>

information available on the Web. DBpedia allows you to ask sophisticated queries against datasets derived from Wikipedia and to link other datasets on the Web to Wikipedia data.

YAGO2 as described in Hoffart, Berberich, Weikum, *et al.*, 2010 is a huge semantic knowledge base, derived from Wikipedia, WordNet and GeoNames. At the time of writing, YAGO2 has knowledge of more than 10 million entities (like persons, organizations, cities, etc.) and contains more than 120 million facts about these entities. YAGO2 has some special features:

- The accuracy of YAGO2 has been manually evaluated, proving a confirmed accuracy of 95. Every relation is annotated with its confidence value.
- YAGO2 is an ontology that is anchored in time and space. YAGO2 attaches a temporal dimension and a spacial dimension to many of its facts and entities.
- YAGO2 is particularly suited for disambiguation purposes, as it contains a large number of names for entities. It also knows the gender of people.

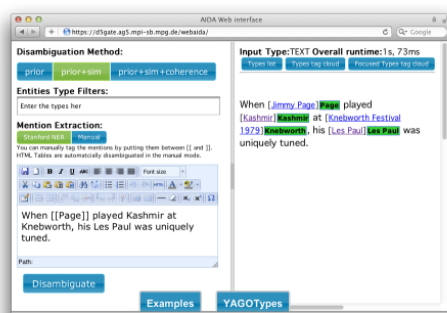


Figure 21.: AIDA web interface for WSD.

AIDA, see Figure 21 is a framework and online tool for entity detection and disambiguation. Given a natural-language text or a Web table, it maps mentions of ambiguous names onto canonical entities (e.g. individual people or places) registered in the YAGO2 knowledge base.

4.3 HYBRID (AUTOMATIC & HUMAN)

With the previous sections we presented two applications able to handle the human and the automatic scenarios. In this section we are presenting a use-case where both the previous scenarios are blended together. This is used to test if our framework is flexible enough to seamlessly support mixed application archetypes. In the matrix at

Table 1 this use-case fits between the human and the automatic computation.



Figure 22.: Interface of the hybrid use-case.

This use-case has the purpose of *detecting faces* in a picture, to accomplish this Task are used an automatic face recognition algorithm plus a human interaction that has the double purpose of validating the algorithm result and detect the missing faces in the image.

This scenario is implemented in 2 steps, in the first step we run the algorithm for detecting the faces (this is the *automatic* scenario), the second step is implemented as a [GWAP](#).

The game, under the name **ThemAmongUs** has been inspired by the 1988 film "*They Live*" directed by John Carpenter. *ThemAmongUs* is a single player arcade shooter in which the player assumes a role of an agent that fights against an alien race disguised as human beings. Equipped with a special camera able to distinguish between human beings and non humans, the agent is asked to shoot at the head of the beings that have not been identified by the camera software. The camera may fail in some occasion, so the agent has to use his judgment to fire only at the right targets.

4.3.1 Introducing Score Degradation

The new game mechanic that is presented to manage the task is called Score-Degradation. This technique may be used in scenarios in which there are not the possibility to compare the results provided

by the players with techniques such as the one provided in output-aggregation or inversion-problems games, because the game that is being taken into consideration is not a multiplayer game but a single player one.

Goal of the technique is to force the user to always provide the right answer with game mechanics that involve low reaction times, high penalties for mistakes (such as early game termination) and incentives to achieve the best results compared to all the other players.

Players are first evaluated based on well known trial examples tasks to understand their reliability level. Failing the required task in these training examples usually ends the gaming experience for the player, forcing him to start the game from the beginning.

Once a sufficient level of trust for the player has been reached, the player is then provided with a sequence of mixed tasks, some of them with an already well established knowledge of the expected results, some of them with completely unknown expected results. While the results of the first kind of tasks will still be checked against the right results, for the second kind of tasks the results provided by the players will always be considered good results.

The players will not be able to distinguish which of the instances of the tasks are being checked against their provided results and which results are simply considered "true as provided" without any further checks. This behavior is also enforced by the fact that the player is not able to understand the moment in which the "trial phase" will end and the fast reaction times force him to not even have the time to think about providing misleading results, with the risk of having to start the game from the beginning again.

In this way the player is always forced to try to give the best possible solution for a specific task. The collected results can be further improved by using traditional aggregation techniques such as majority voting or similar, depending on the task that has to be solved.

4.3.2 Gameplay

Goal of the game is to obtain the highest possible score given a limited amount of time (1 minute). The player is provided with a series of images that present bounding boxes of the face of human beings automatically identified by the special camera of the agent. Each provided image will constitute a round of the game. If in the image some face has not been surrounded by a bounding box, it means that the portrayed subject is an alien and must be shot at by pressing the left click button.

The player has a limited amount of time, typically 5 seconds, to shoot at all the faces that have not been recognized, in order to obtain a certain amount of points. During a round the player may also find improper bounding boxes, such as knees or other part of the body that have been recognized as a face.

The player may right click on these boxes to remove them and obtain additional points. When the player has shoot to all the unboxed faces,

he may shoot at a button on the right lower corner of the screen to play the next round of the game. The game will end if the player will shoot at a recognized face by mistake.

At the end of each round (after the 5 seconds have passed or when the player has pressed the end button), the system checks if the player has missed any face. If it is the case and the image was a trial one or one for which the results were known, the player will lose the game with a score equal to the number of points he achieved so far. Otherwise the score for the current round are calculated in the following way:

$$\begin{aligned} \text{Score} = & (\text{RoundNumber} * 10) * (\text{NumberOfAliensKilled}) \\ & + (100 * (\text{FalseBBRemoved})) \end{aligned} \quad (4.1)$$

At the end of the global gaming time, a player who has not made any mistake will receive 1000 additional points. The points are used to provide an incentive to improve and beat other players by improving the score on further matches.

5 | IMPLEMENTATION AND EVALUATION

The previous chapters presented the Conceptual Model and the Use-cases of a framework for web-based human and machine computation. In this chapter are presented the actual implementation of the framework and the use-cases. The framework implementation is divided in two parts: the Configurator, managed by the Crowdscheduler (see [subsection 2.1.3](#)), and the Execution Layer implemented in NodeJS. The use-cases are built in JavaScript and the Execution Layer and leverages on the [HTML5](#) features described in [subsection 2.2.1](#).

5.1 ARCHITECTURE

In this section we present the architecture of our framework, whose implementation is divided in two parts. The reference model in [Figure 18](#) has been customized to meet our needs in terms of flexibility and pluggability. As one can see in [Figure 23](#) the Central Hub has been splitted into two separated components: the *Configurator* and the *Execution Layer*. The *Configurator* is in charge of managing the Tasks lifecycle. It is used to create and configure the Tasks. The *Execution Layer* is used to configure the actual implementations for each μ Task.

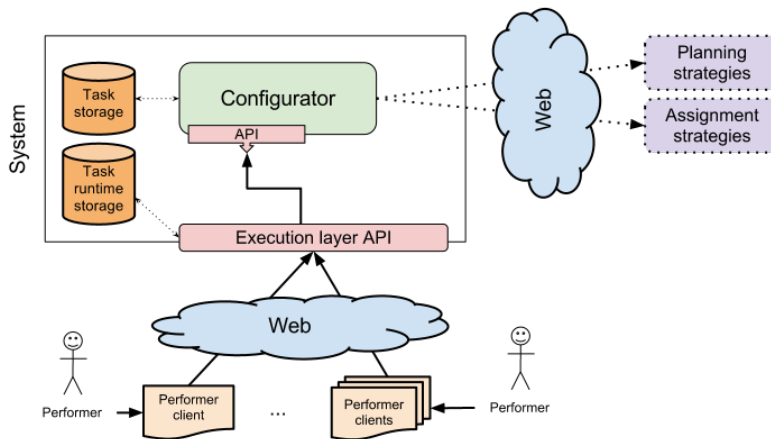


Figure 23.: Specialized architecture.

5.1.1 Configurator

The **Configurator** is in charge of managing the Tasks lifecycle from the abstract definition to the results aggregation. This component must implement the metamodel described in 3.2.2 to describe the Tasks and must handle correctly the hooks defined in 3.5. This component must have a set of API for the Task management. The main functionality offered by the *Configurator* are:

- Allow the **creation** of a Task, also at abstract level, using either the API or the built-in UI.
- Allow a *Performer* to **execute** the Task using a standard non configurable UI, provided as-is for each Task type.
- Allow to **request information** about a Task, the information that can be requested includes:
 - Retrieve the list of μ Task associated with a given Task
 - Post the result of the execution of a given μ Task
 - Notify about the completion of a Task or μ Task

For implementing this component we adopted the *CrowdSearcher* framework. As described in 2.1.3 this framework natively embodies most of the functionality described previously. The *CrowdSearcher* supports most of the operations required by our framework natively. The *Task creation* and the *Task planning* built-in implementations are now described to better understand how the framework works.

Task creation

The task creation is performed by the *Configurator* either by using its web-interface or via API calls. The creation of a Work/Task can be performed with three methods: using a JSON file, via API calls or using a step-by-step "manual" procedure. Using a JSON file, user can supply file containing all the required information like the data definition and the data instances. Otherwise the user create the new Task following a step-by-step procedure bundled with the *Configurator*. As shown in Figure 24 the manual Task creation involves the definition of a **Schema** for the data with all the related Fields. Once the Schema has been defined the user can supply the data instances to the Task.

Task planning

The planning of a Task involves the creation of μ Task with the associated data. The assignment can be performed either automatically or manually. The automatic plan assignment uses a simple subdivision based on the number of instances to assign to each μ Task (see ??).

As depicted in Figure 26 manual planning involves the *Requester* interaction in order to create each μ Task. After the creation of the μ Task the user have to select the instances belonging to this μ Task. Eventually the user is able to select and, if needed, configure the type

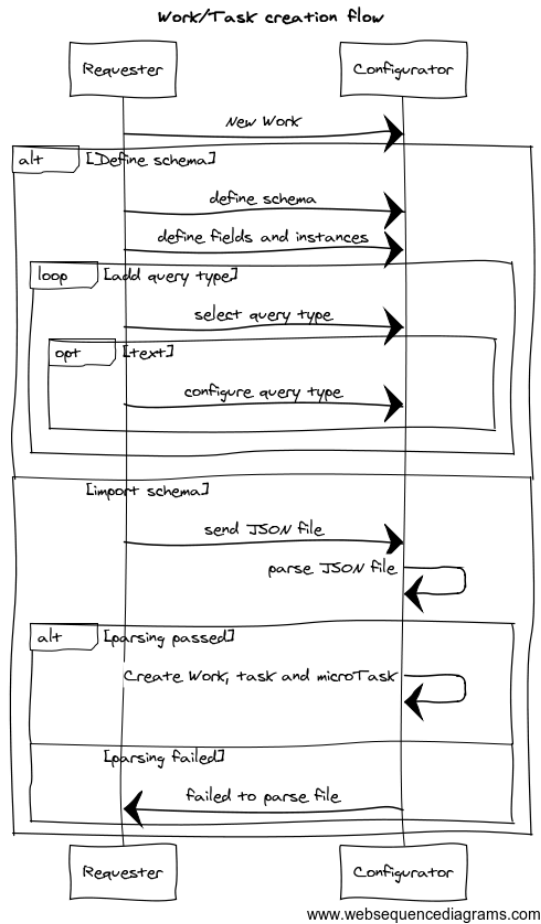


Figure 24.: Work/Task creation flow.

of the μ Task. The configurable types must be a subset of the parent Task types.

5.1.2 Execution layer

The *Execution Layer* is in charge of managing the μ Task implementations. As described in 3.2.1, each Task belongs to a macro-type (i.e. classify, like, comment, etc), and each macro-type has its own built-in implementation. The *Execution Layer* give an easy overriding method to replace the built-in UI with a custom one. Using the *Execution Layer* a user is able to upload the executable resources (e.g. HTML, JavaScript) that will be used as the new interface. This custom implementations can be defined hierarchically for each Task and μ Task. So we can supply custom code for the whole Task or for each μ Task. Using this system we have the ability to tweak single Tasks without harming the others. The layer provides a fallback system for finding the most suitable code for each μ Task.

The **Execution layer** offer the following functionalities:




Figure 25.: Manual Task planning vs Automatic Task planning.




Figure 26.: Task planning flow.

- Allow a *Requester* to configure the implementations associated to a Task and/or a μ Task. The implementations are configured specifying the target platform (mobile, desktop, tablet, ...) and the executable resources used by the implementation (i.e. HTML, CSS and JS files). Which implementation to use is configured later in the *Planning* step.
- Create a layer of abstraction between the implementation and the Configurator, creating a sandboxed environment where the implementation can run and communicate with the Configurator.
- Allow the *Performer* to execute a specific μ Task implementation.

The *Execution Layer* has been developed using *NodeJS*. *NodeJS* is a platform built on Chrome's JavaScript runtime for easily building fast, scalable network applications. *NodeJS* uses an event-driven, non-blocking

I/O model that makes it lightweight and efficient, perfect for data-intensive real-time applications that run across distributed devices. Due to the great hype around *NodeJS* and the ease of developing new libraries, there are thousands of these being developed and made available via the built-in package manager `npm`.

The implementation of the *Execution Layer* is subdivided into two parts: the *Server Backend* and the *μTasks Wrapper*.

SERVER BACKEND is a node REST web-server implemented using *Express*. This server provides an web interface for managing the implementations of Tasks and *μTasks*. This interface allows a user to upload custom code (e.g. HTML, JS, CSS) for the selected Task (or *μTask*). The Server interacts with the *Configurator* to gather information on a Task composition and type.

μTASKS WRAPPER exposes a wrapper class that can be used by the *μTasks* implementations to communicate with the *Server Backend*. The envelops the *μTask* and supply useful methods to retrieve Task configurations, gather the associated data and post the results *Configurator*.

5.1.3 Task storage & task runtime storage

These are the storage areas where we put all the data associated with the Tasks and the *μTasks* implementations. We used two separated storage area to physically separate the runtime data from the abstract configuration of the Task.

The *Task Storage* is handled by the *Configurator*, meanwhile the *Task Runtime Storage* is managed by the *Execution Layer*.

5.1.4 Performer & Performer Client

The *Performer client* represents the platforms (like desktop or mobile) on which a *Performer* executes a Task implementation. The *Performer client* make use of the *Execution layer* API to retrieve the correct implementation, communicate the status during the execution of a *μTask* and post the result of the execution. The *Performer* is the actual user that is using the *client*.

5.2 USE CASES

The previous section presented the implementation of the framework described in this thesis. Now we describe the implementation of the three use-cases described in chapter 4: Automatic, Human and Hybrid.

5.2.1 Automatic

If a Task typically performed on high-end machines can be offloaded to web clients we are able to use the browser as a node for Distributed Computing.

As described in 4.1 in this scenario we are implementing the Scale-Invariant Feature Transform (SIFT) algorithm. This algorithm was chosen due to its high computational requirements.

We made some preliminary tests to check if we can implement this algorithm using pure JavaScript code but we stumbled across the problems described in [WebWorkers](#). Due to this limitation we opted to implement the code using [WebCL](#) (see 2.2.2). With this framework we can leverage on the [GPGPU](#) power to perform the computations and unburden the JavaScript engine.

In order to build a working example of the algorithm we started with the creation of an *Abstraction Layer* over the [WebCL](#) raw implementation. Then we created a small *MultiMedia Library* able to interact with the *Abstraction Layer* for performing the CPU intensive operations. Eventually we implemented the [SIFT](#) algorithm within the *MultiMedia Library*.

THE ABSTRACTION LAYER allow an easy communication with the [WebCL](#) framework. As described in 2.2.2, for running a OpenCL/WebCL program a few step must be performed. Here is the list:

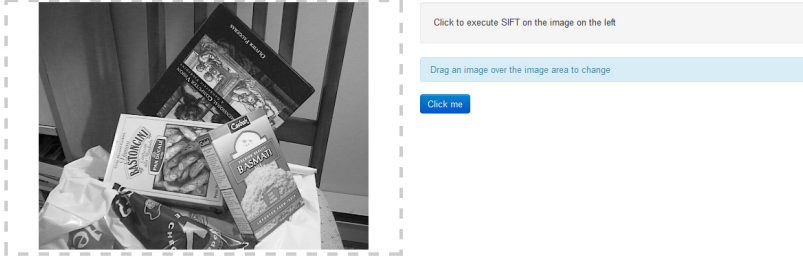
1. Query host for [OpenCL](#) devices.
2. Create a context to associate [OpenCL](#) devices.
3. Create programs for execution on one or more associated devices.
4. From the programs, select kernels to execute.
5. Create memory objects accessible from the host and/or the device.
6. Copy memory data to the device as needed.
7. Provide kernels to the command queue for execution.
8. Copy results from the device to the host.

Our *Abstraction Layer* allow to define all the I/O parameters first and then run the selected kernel function on top of these. This abstraction allowed us to easily interact with [WebCL](#) without performing all the previous steps each time we need to execute a kernel function.

THE MULTIMEDIA LIBRARY is a set of image operations implemented using both JavaScript and the OpenCL kernel language. With this library we provided a minimal set of common operations (e.g. convolve, blur, resize, etc.) useful for the implementation of the [SIFT](#) algorithm.

THE ALGORITHM has been implemented using the methods of the *MultiMedia library*. The algorithm has been implemented step-by-step as described in 4.1.1. For each step the intermediate result is stored and the step itself is timed to find bottlenecks. In Figure 27 and Figure 28 are presented the intermediate results obtained during the process and the final result.

Automatic use case



Click to execute SIFT on the image on the left

Drag an image over the image area to change

Click me

Scale space took 1004ms

DoG took 133ms

MaxMin took 153ms

KeyPoints refinement took 76ms

Magnitude and Orientation took 120ms

powered by Holos

Figure 27.: Intermediate results of the algorithm.




Figure 28.: SIFT result comparison with the reference data.

Benchmark/Metric

Since the purpose of this use-case is the feasibility of high load computation on the user browser, the implementation of the algorithm has not been optimized. Then the performance of this implementation are not comparable to the existing C/C+ implementation+, but we can leverage on the parallelism of the whole framework to obtain a higher throughput. In our test cases we obtained the results presented in Table 4.

Table 4.: SIFT algorithm time performances on the Web.

Image size	ScaleSpace + Dog	Keypoints detection	Total time
1024x683	~4300 ms	~1100 ms	~5400 ms
800x534	~2600 ms	~700 ms	~3400 ms
600x401	~1500 ms	~360 ms	~1900 ms
400x360	~1130 ms	~310 ms	~1500 ms

5.2.2 Human

In this scenario we implemented a Word-Sense Disambiguation (WSD) problem (see 4.2). To test our framework we built an Human Computation application like AIDA. The purpose of this application is to let the user disambiguate the words in a piece of text. This application allowed us to test also the connection capabilities of our framework. Since the Task we implemented need to communicate to external servers to gather the information on a word.

Human use case

The screenshot shows a web interface for a human use case. On the left, under the heading "Text", there is a snippet of text: "Agnetha, Björn, Benny, and Anni-Frid formed Sweden's most successful pop music group. Their greatest hits were Waterloo and SOS." The word "Benny" is highlighted in orange. On the right, under the heading "Possible definitions", there are three entries: "Benny Andersson" (Göran Eror Benny Andersson, born 16 December 1946, known professionally as Benny Andersson, is a Swedish musician, composer, former member of the Swedish musical group ABBA), "Benny Neyman" (Wilhelmus Albertus (Benny) Neyman (9 June 1951 in Maastricht, Netherlands; 7 February 2008 in Soesterberg, Netherlands) was a Dutch singer), and "Benny Hill" (Benny Hill (21 January 1925 – 20 April 1992) was an English comedian and actor, notable for his long-running television programme The Benny Hill Show). Below the definitions, there is a small blue button that says "Powered by Volox".

Figure 29.: Human use case Web interface.

The application (see Figure 29) can be divided in two steps. In the first step a text is analyzed from the backend (YAGO2 or DBpedia) then, in the second step the application asks to the user to disambiguate certain words that have not been correctly understood.

In the second step for each ambiguous work the user have a list of possible meanings to choose from. Once the user selected the correct word interpretation it is saved into the browser.

Eventually when the user has disambiguated all the words he/she can send the results back to the framework, where they are saved and stored for improving the precision of the WSD algorithm itself.

5.2.3 Hybrid

In this scenario we have the human and the automatic computation blended together. As explained in 4.3, this use-case is implemented in two parts: the face detection and the **GWAP** logic.



Figure 30.: Face detection algorithm demo.

The face detection logic is implemented using an external library by **Liu Liu** (see Figure 30). This library allow to detect face in an image using the JavaScript features presented in (like Canvas and WebWorkers). The technique used to implement the algorithm will not be covered here but you can find a detailed explanation here: <http://liuliu.me/eyes/javascript-face-detection-explained/>. The library¹ can be used with or without WebWorkers to further increase the processing speed².

The **GWAP** follows the game logic presented in 4.3.2 an it has been implemented using the **HTML5** features.

Objective of the game is being able to identify the visage of the subject portrayed in a photo. The provided algorithm for face detection may fail in this task, by recognizing objects that are not the required ones (e.g. knees) or not recognizing at all the face of a person. For this reason the game will be able to perform two different tasks: identify faces that have not been recognized and remove elements that have been marked improperly.

The game present the image with the bounding boxes of the faces that has been identified by the first step to the player. The faces that have not been recognized are the ones that the player needs to mark during the game by clicking over them with the left mouse button and the area interested by the click will be circumscribed in order to obtain the bounding box for the identified face.

Possible wrong bounding boxes will be removed by the player with the right mouse button, in order to remove the false positives that may be identified by automatic algorithms. Bounding boxes with high level confidence are kept into the image, while bounding boxes with

¹ Source code: <https://github.com/liuliu/ccv/tree/unstable/js>

² Demo available here: <http://liuliu.me/ccv/js/nss/>

low level confidence values will be removed: it will be a player duty to mark those faces by playing the game.

At the end of the execution we obtain a set of automatically detected faces plus all the faces detected by the user. By intersecting this two set we could improve the performance of our algorithm by adding the faces it was unable to find to the training set.

A

CONCLUSION AND FUTURE WORKS

In this thesis we proposed a novel approach for web-based human and machine computation. By leveraging on the features offered by modern Web browsers, we provide a framework for designing and implementing general purpose Tasks. This system allows the user to configure the Task at abstract level and add all the required custom implementations later, even during the execution. On top of that our framework offer the possibility to plugin external logic to tweak the Task execution flow. With this pluggable behavior our framework is able to support any application archetype (human, automatic and hybrid). We validated our approach by implementing several use case applications, which demonstrate the framework by stressing its architectural configurations.

A.1 FUTURE WORK

The described work is the begin of a broaden research activity. For instance our framework can be extended to meet the scenario proposed by [Karame, Francillon, and Čapkun, 2011](#). The authors of this paper suggest the creation of an online marketplace where the users can gather *credits* by performing Task and then use these credits as payment method for online contents (like newspapers, images, video, etc.). From a technical point of view, our work can be extended as follows:

- Better support for image processing with optimized libraries.
- Enhanced Task creation with the implementation of wizards that guide the user though the Task creation.
- Enriching of the framework with the addition of supported application archetype.

ACRONYMS

WebCL	Web Computing Language The WebCL working group is working to define a JavaScript binding to the Khronos OpenCL standard for heterogeneous parallel computing. WebCL will enable web applications to harness GPU and multi-core CPU parallel processing from within a Web browser, enabling significant acceleration of applications such as image and video processing and advanced physics for Web Graphics Library (WebGL) games.
SIFT	Scale-Invariant Feature Transform SIFT is an algorithm in computer vision to detect and describe local features in images.
OpenCL	Open Computing Language OpenCL is a framework for writing programs that execute across heterogeneous platforms consisting of CPU, GPU, and other processors. OpenCL includes a language (based on C99) for writing <i>kernels</i> (functions that execute on OpenCL devices), plus APIs that are used to define and then control the platforms. OpenCL provides parallel computing using task-based and data-based parallelism.
WebGL	Web Graphics Library WebGL is a cross-platform, royalty-free API used to create 3D graphics in a Web browser. Based on OpenGL ES 2.0, WebGL uses the OpenGL shading language, GLSL, and offers the familiarity of the standard OpenGL API. Because it runs in the HTML5 Canvas element, WebGL has full integration with all DOM interfaces.
CORS	Cross-origin Resource Sharing Cross-origin resource sharing (CORS) is a web browser technology specification which defines ways for a web server to allow its resources to be accessed by a web page from a different domain. Such access would otherwise be forbidden by the same origin policy. CORS defines a way in which the browser and the server can interact to determine whether or not to allow the cross-origin request. It is a compromise that allows greater flexibility, but is more secure than simply allowing all such requests.
RIA	Rich Internet Application Rich Internet Applications (RIA) are web-base application that have many of the characteristics of desktop application software.
HIT	Human Intelligent Task
HC	Human Computation
DHC	Distributed Human Computation
DC	Distributed Computing

AI	Artificial Intelligence
GIMPS	Great Internet Mersenne Prime Search
TCP	Transmission Control Protocol
AJAX	Asynchronous JavaScript and XML
HTML	HyperText Markup Language
CSS	Cascading Style Sheets
BOINC	Berkeley Open Infrastructure for Network Computing
GWAP	Game With A Purpose
WSD	Word-Sense Disambiguation
GPGPU	General-purpose computing on graphics processing units
SPMD	Single process/program multiple data
SIMD	Single instruction multiple data
SETI@home	Search for Extra-Terrestrial Intelligence <i>at</i> home
DoG	Difference of Gaussian

Essential bibliography

Andrews, G.R.

nodate “Foundations of multithreaded, parallel, and distributed programming. 2000”, *Wesley, University of Arizona, USA*. (Cited on p. 8.)

Auer, S. *et al.*

2007 “Dbpedia: A nucleus for a web of open data”, *The Semantic Web*, pp. 722–735. (Cited on p. 42.)

Barabási, A.L. *et al.*

2001 “Parasitic computing”, *Nature*, 412, 6850. (Cited on pp. xiii, 8, 10, 11.)

Bozzon, A., M. Brambilla, and S. Ceri

2012 “Answering search queries with CrowdSearcher”, in *Proceedings of the 21st international conference on World Wide Web*, ACM. (Cited on pp. 12, 25.)

Ceri, S. *et al.*

2012 “CrowdSearch: Crowdsourcing Web search”. (Cited on p. 11.)

Dean, J. and S. Ghemawat

2008 “MapReduce: Simplified data processing on large clusters”, *Communications of the ACM*, 51, 1. (Cited on pp. xiii, 8.)

Group, Khronos OpenCL Working *et al.*

2008 “The opencl specification”, *A. Munshi, Ed.*

Hoffart, J., K. Berberich, G. Weikum, *et al.*

2010 “YAGO2: a spatially and temporally enhanced knowledge base from Wikipedia”, *Knowledge Creation Diffusion Utilization*, 52, November. (Cited on p. 43.)

Howe, J.

2006 “The rise of crowdsourcing”, *Wired magazine*, 14, 6, pp. 1–4. (Cited on p. 3.)

Jenkin, N.

2008 “Parasitic JavaScript”. (Cited on pp. 11, 16.)

Karame, G.O., A. Francillon, and S. Čapkun

2011 “Pay as you browse: microcomputations as micropayments in web-based services”, in *Proceedings of the 20th international conference on World wide web*, ACM. (Cited on pp. 11, 57.)

Law, E. and L. Ahn

2011 “Human computation”, *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 5, 3. (Cited on pp. 1, 3.)

Little, G. *et al.*

- 2010 “Turkit: human computation algorithms on mechanical turk”, in *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, ACM. (Cited on p. 5.)

Parameswaran, M. and A.B. Whinston

- 2007 “Research issues in social computing”, *Journal of the Association for Information Systems*, 8, 6, pp. 336–350.

Quinn, A.J. and B.B. Bederson

- 2011 “Human computation: a survey and taxonomy of a growing field”, in *Proceedings of the 2011 annual conference on Human factors in computing systems*, ACM, pp. 1403–1412.

Schuler, D.

- 1994 “Social computing”, *Communications of the ACM*, 37, 1, pp. 28–29. (Cited on p. 4.)

Seinstra, F.J. *et al.*

- 2011 “Jungle computing: Distributed supercomputing beyond clusters, grids, and clouds”, *Grids, Clouds and Virtualization*, pp. 167–197. (Cited on p. 8.)

Turing, A. M.

- 1950 *Computing Machinery and Intelligence*. (Cited on p. 3.)

Von Ahn, L. and L. Dabbish

- 2004 “Labeling images with a computer game”, in *Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM, pp. 319–326. (Cited on p. 7.)
- 2008 “Designing games with a purpose”, *Communications of the ACM*, 51, 8, pp. 58–67. (Cited on p. 6.)

Von Ahn, L., R. Liu, and M. Blum

- 2006 “Peekaboom: a game for locating objects in images”, in *Proceedings of the SIGCHI conference on Human Factors in computing systems*, ACM, pp. 55–64. (Cited on p. xiv.)

Von Ahn, L. *et al.*

- 2008 “recaptcha: Human-based character recognition via web security measures”, *Science*, 321, 5895, pp. 1465–1468. (Cited on p. 6.)

Wang, F.Y. *et al.*

- 2007 “Social computing: From social informatics to social intelligence”, *Intelligent Systems, IEEE*, 22, 2, pp. 79–83.

WebGL-OpenGL, ES

- 2011 “2.0 for the Web”, *Verkkodokumentti* <<http://www.khronos.org/webgl/>>. Luettu, 16.

Online resources

- Amazon's Mechanical Turk* 2012 , <http://www.mturk.com/>. (Cited on pp. [xiv](#), [4](#), [5](#).)
- CoffeeScript* 2012 , <http://coffeescript.org/>.
- Emscripten* 2012 , <http://emscripten.org/>. (Cited on p. [13](#).)
- Express* 2012 , <http://expressjs.com/>. (Cited on p. [51](#).)
- FoldIt* 2012 , <http://fold.it/>. (Cited on pp. [xiv](#), [7](#).)
- Google Web Toolkit* 2012 , <https://developers.google.com/web-toolkit/>. (Cited on p. [13](#).)
- jQuery* 2012 , <http://www.jquery.com/>. (Cited on p. [15](#).)
- Modernizr* 2012 , <http://modernizr.com/>. (Cited on p. [15](#).)
- NodeJS* 2012 , <http://nodejs.org/>. (Cited on pp. [50](#), [51](#).)
- Socket.io* 2012 , <http://socket.io/>. (Cited on p. [16](#).)