USC
UNIVERSIDADE DE SANTIAGO DE COMPOSTELA
CENTRO INTERNACIONAL DE ESTUDOS DE DOUTORAMENTO E AVANZADOS DA USC (CIEDUS)

TESIS DE DOCTORADO

GRID, CLOUD AND HIGH PERFORMANCE COMPUTING: SCIENTIFIC APPLICATIONS' USE CASES

Mª del Carmen Cotelo Queijo

ESCUELA DE DOCTORADO INTERNACIONAL

PROGRAMA DE DOCTORADO EN INVESTIGACIÓN EN TECNOLOGÍAS DE LA INFORMACIÓN

SANTIAGO DE COMPOSTELA

2018

UNIVERSIDADE DE SANTIAGO DE COMPOSTELA

DECLARACIÓN DEL AUTOR DE LA TESIS

Grid, Cloud and High Performance Computing: scientific applications' use cases

Dña. Mª del Carmen Cotelo Queijo

Presento mi tesis, siguiendo el procedimiento adecuado al Reglamento, y declaro que:

1) *La tesis abarca los resultados de la elaboración de mi trabajo.*
2) *En su caso, en la tesis se hace referencia a las colaboraciones que tuvo este trabajo.*
3) *La tesis es la versión definitiva presentada para su defensa y coincide con la versión enviada en formato electrónico.*
4) *Confirmo que la tesis no incurre en ningún tipo de plagio de otros autores ni de trabajos presentados por mí para la obtención de otros títulos.*

En Santiago de Compostela, 10 de Julio de 2018.

Fdo. Mª del Carmen Cotelo Queijo

AUTORIZACIÓN DE LOS DIRECTORES DE LA TESIS

Grid, Cloud and High Performance Computing: scientific applications' use cases

Dr. D. J. Manuel Cotos Yañez
Dr. D. Andrés Gómez Tato

INFORMAN:

Que la presente tesis, corresponde con el trabajo realizado por Dña. Mª del Carmen Cotelo Queijo, bajo nuestra dirección, y autorizamos su presentación, considerando que reúne los requisitos exigidos en el Reglamento de Estudios de Doctorado de la USC, y que como directores de ésta no incurre en las causas de abstención establecidas en Ley 40/2015.

De acuerdo con el artículo 41 del Reglamento de Estudios de Doctorado, declaramos también que la presente tesis doctoral es idónea para ser defendida en base a la modalidad de COMPENDIO DE PUBLICACIONES, en los que la participación de la doctoranda fue decisiva para su elaboración.

La utilización de estos artículos en esta memoria, está en conocimiento de los coautores, tanto doctores como no doctores. Además, estos últimos tienen conocimiento de que ninguno de los trabajos aquí reunidos podrá ser presentado en ninguna otra tesis doctoral.

En Santiago de Compostela, 10 de Julio de 2018

Fdo. Dr. D. J. Manuel Cotos Yañez   Fdo. Dr. D. Andrés Gómez Tato

AGRADECIMIENTOS

En estas líneas quiero expresar mi agradecimiento a todas aquellas personas que han contribuido a que esta tesis fuera posible:

En primer lugar quiero agradecer a mis directores, el Dr. Andrés Gómez Tato y el Dr. José Manuel Cotos Yáñez, la confianza que depositaron en mí. Además de su buena dirección y tutorización, imprescindibles para llevar a cabo este trabajo, quiero agradecerles de forma especial su apoyo y comprensión, que me han permitido desarrollar esta tesis y compaginarla con los demás aspectos de mi vida.

A Andrés, responsable además de coordinar mi labor profesional durante estos años, quiero agradecerle también su implicación y disponibilidad, así como todo el ánimo que me ha transmitido.

A mis compañeros del CESGA, excepcionales tanto por su calidad profesional como humana, muchas gracias por estar siempre dispuestos a echarme una mano, por adoptar mis retos como vuestros y brindarme vuestro tiempo desinteresadamente. Más allá de lo profesional o lo académico, quiero agradeceros también toda la energía positiva que me habéis regalado en tantas charlas de café y tardes de terrazas. Es una gran suerte poder contar con vosotros.

Al CESGA como institución, mi agradecimiento por su apoyo, que me ha permitido completar los estudios de doctorado y esta tesis.

Gracias a todos los compañeros de los proyectos y a los coautores de artículos con los que he tenido oportunidad de colaborar y de los que sé que he aprendido muchísimo.

A mis amigas y amigos, por estar siempre ahí, esforzándose por hacer coincidir su tiempo con el mío, escuchándome y haciéndome reír.

A mis padres, por su gran esfuerzo y empeño en proporcionarme todos los medios a su alcance para ayudarme en la vida.

A mi hermana Laura y a Patricia, por estar siempre pendientes y ayudarme a cuidar de los que más quiero.

A Elsa y Álex, los motores de mi vida, gracias por tratar de entender y por ceder una parte de vuestro tiempo conmigo para que pudiese finalizar este trabajo.

A Fran, gracias por hacer de ti y de mí tantas veces, por caminar a mi lado y apoyarme siempre incondicionalmente

Santiago de Compostela, Julio 2018.

COMPENDIO DE ARTÍCULOS

Esta tesis doctoral se presenta como compendio de artículos publicados en revistas indexadas, de acuerdo con el artículo 41 del reglamento de estudios de doctorado de la Universidad de Santiago de Compostela (USC), aprobado en pleno ordinario del Consejo de Gobierno el 12 de junio de 2017.

Las publicaciones que se aportan como parte de esta tesis doctoral son los siguientes artículos:

1. ***Retelab: A geospatial grid web laboratory for the oceanographic research community.*** Carmen Cotelo, Andrés Gómez, J. Ignacio López, David Mera, José M. Cotos, J. Pérez Marrero, Constantino Vázquez. Future Generation Computer Systems 26 (2010) 1157-1164.
   DOI: 10.1016/j.future.2010.05.018.
   Factor de impacto JCR (2010): 2,371, Q1.
   Número de citas (Junio 2017): 9.

2. ***Fault-tolerant virtual cluster experiments on federated sites using BonFIRE.*** A. Gómez, L. M. Carril, R. Valin, J. C. Mouriño, C. Cotelo. Future Generation Computer Systems 34 (2014) 17-25.
   DOI: 10.1016/j.future.2013.12.027.
   Factor de impacto JCR (2014): 2,639, Q1.
   Número de citas (Junio 2017): 9.

3. ***On the successful coexistence of oceanographic operational services with other computational workloads.*** Carmen Cotelo, María Aránzazu Amo Baladrón, Roland Aznar, Pablo Lorente, Pablo Rey, Aurelio Rodríguez. International Journal of High Performance Computing Applications. Vol 32, Issue 2, (2018) pp. 302-313.
   DOI: 10.1177/1094342017692045.
   Factor de impacto JCR (2016): 2,097, Q2.

Conforme a la normativa de la USC vigente para esta modalidad, esta tesis se presenta en inglés, comenzando con un resumen en castellano al que sigue un primer capítulo de introducción al marco de trabajo, descripción general sobre la línea de investigación de los trabajos que se aportan y objetivos a alcanzar. El núcleo central de la tesis está constituido por una copia de los tres artículos que la componen. Finalmente se presenta un capítulo de conclusiones, seguido de las referencias bibliográficas y de un índice de proyectos de investigación en los que ha participado la autora.

RESUMEN

Los avances en computación y la reducción de costes experimentada en los componentes de equipamiento informático, a lo largo de las últimas décadas, han hecho posible el rápido desarrollo y la mejora continua de los sistemas de computación. En esta evolución destaca la capacidad de cálculo que se ha logrado alcanzar gracias a los sistemas distribuidos, que han sido y continúan siendo herramientas fundamentales para dar solución a problemas complejos, tanto en el ámbito científico como en entornos comerciales o industriales con importantes requerimientos computacionales.

Los sistemas de computación y almacenamiento distribuido, las redes de interconexión y los protocolos de comunicaciones han pasado a ser imprescindibles en el mundo empresarial, el entorno académico, de investigación, e incluso actualmente ya en los servicios y tecnologías de consumo más cotidianas. Sin embargo, la computación distribuida ha pasado por diferentes etapas en su desarrollo y ha evolucionado a través de varias líneas, dando lugar a diferentes modelos de implementación y uso. Según la naturaleza y las condiciones particulares de cada problema a resolver estos modelos de computación distribuida podrán resultar más o menos adecuados y complejos.

Para los investigadores, no expertos en computación, puede no resultar trivial decidir cuál es la mejor tecnología para resolver un problema o ejecutar una aplicación científica determinada. Cada aplicación científica dispone de diferentes restricciones relacionadas con el tipo y cantidad de recursos necesarios para su ejecución (memoria, almacenamiento, comunicaciones, etc.), pero además existe un factor externo a menudo determinante, el tiempo de obtención de resultados. La selección y configuración de las características más adecuadas para los recursos de computación a usar, el número de equipos necesarios o el tipo de interconexión entre ellos son cuestiones que pueden resultar complejas y que pueden entorpecer, e incluso comprometer, el proceso de desarrollo de un proyecto de investigación o de un servicio. Cabe la posibilidad de que los

investigadores dispongan de un sistema de cálculo propio y dedicado, aunque en la mayoría de casos estos sistemas resultarán en algún momento inadecuados o insuficientes y son, por lo general, difíciles de mantener, tanto por el coste económico que supone su adquisición como su mantenimiento. Además, aquellos recursos que son óptimos para la ejecución un tipo de aplicación científica pueden no serlo para resolver otro tipo de problemas, lo que supone una desventaja añadida que hace difícil la posibilidad de mantener recursos propios multipropósito o de que éstos ofrezcan las prestaciones necesarias para poder obtener resultados en un tiempo determinado. Por ello, gran parte de los investigadores se ven en algún momento en la situación de tener que usar alguna de las diferentes opciones de computación distribuida disponibles para poder ejecutar sus aplicaciones, cumpliendo con los requerimientos propios de sus proyectos o servicios. En la selección de la tecnología más adecuada a emplear pueden, según el caso, influir multitud de factores determinantes como la facilidad de uso, el número de recursos necesarios, la trasferencia de datos, los tiempos de espera antes de iniciar la ejecución, la tolerancia a fallos o la calidad de servicio ofrecida, etc.

El presente trabajo describe algunos de los principales paradigmas, asociados a la arquitectura de computación distribuida, más usados a lo largo de la última década, a través de tres casos de uso reales desarrollados e implementados en el Centro de Supercomputación de Galicia (CESGA) en el marco de diferentes proyectos de investigación europeos y nacionales.

Cada contribución explora un paradigma de computación distribuida: Grid, Cloud y supercomputadores basados en arquitecturas de tipo clúster. En cada caso se han analizado diferentes aspectos de estos modelos de computación a través de su aplicación en casos prácticos de aplicaciones científicas concretas. Durante el análisis de la aplicación de dichos casos prácticos se han detectado diversas mejoras y aportaciones a realizar a las tecnologías existentes, para poder aplicarlas con éxito estos casos, y se han diseñado e implementado soluciones concretas cumpliendo con los requisitos establecidos por sus usuarios finales.

Como resultado del trabajo realizado ha sido posible mejorar diversos servicios y procesos basados en el uso de computación distribuida. Aunque los casos de uso estudiados se centran en tres aplicaciones científicas concretas, el trabajo desarrollado no es de aplicación exclusiva a dichos casos, sino que puede ser aplicable también en otros de características similares, tanto para el entorno de investigación como de servicios operacionales.

El principal objetivo de esta tesis ha sido la simplificación de la ejecución de diferentes aplicaciones científicas en entornos distribuidos y la obtención de sus resultados en tiempo, según los requisitos específicos de cada caso. Para cada uno se diseñaron e implementaron diversos desarrollos para solucionar problemas, mejorar las funcionalidades del sistema y el rendimiento de la aplicación o del flujo de trabajo, y por tanto mejorar los tiempos de obtención de resultados.

En el primer caso de estudio se diseñó e implementó una solución basada en Grid, en un momento de gran auge para este tipo de paradigma de computación distribuida y de las tecnologías relacionadas. Bajo el marco del proyecto nacional RETELAB (Ministerio de Educación y Ciencia, ESP2006-13778-C04), se diseñó y desplegó un entorno de trabajo colaborativo y distribuido, configurado como como un laboratorio virtual para el desarrollo de proyectos de investigación interdisciplinares y relacionados con la teledetección oceanográfica. Este entorno estaba destinado a los miembros de la Red Nacional de Teledetección Oceanográfica (RETEMAR). Los miembros de esta Red demandaban una herramienta capaz de proporcionar la potencia de cálculo y almacenamiento requerida por las aplicaciones y datos científicos de sus proyectos. Sin embargo, además de la capacidad de cálculo requerida, la sencillez de uso del laboratorio virtual fue un factor fundamental en su diseño, ya que estaba destinado a usuarios finales no expertos en computación.

El Grid es típicamente una infraestructura distribuida geográficamente, formada por recursos heterogéneos proporcionados por diferentes organizaciones y que no están sujetos a un control

centralizado. Su funcionamiento se basa en la interacción de un gran número de servicios con diferentes cometidos, además de los diferentes sistemas de gestión propios de los recursos locales correspondientes a cada organización que aporta recursos al Grid.

Inicialmente los servicios proporcionados en un Grid computacional se exponían a los usuarios mediante interfaz de línea de comandos (en inglés CLI). Las CLI en Grid proporcionan herramientas con un modo de uso casi directo sobre los recursos de cómputo, pero debido a su complejidad, al tiempo necesario para aprender a utilizarlas adecuadamente y a que los principales usuarios Grid (físicos, químicos, ingenieros, etc.) no son necesariamente expertos en computación, sino en las aplicaciones que utilizan para su trabajo, se buscaron mecanismos alternativos para proveer entornos de acceso y uso más amigables, sin perder de vista otros importantes aspectos como la seguridad.

El desarrollo del laboratorio virtual RETELAB centró sus esfuerzos en ofrecer recursos distribuidos basados en tecnologías Grid mediante la implementación de métodos de acceso, uso y gestión lo más sencillos posible, principalmente desde el punto de vista del usuario pero también del administrador de los recursos involucrados.

Según las especificaciones iniciales de los investigadores de la Red RETEMAR, los requerimientos a cumplir por el laboratorio virtual debían incluir: la gestión de recursos compartidos, soporte para el almacenamiento, administración e intercambio de datos distribuidos, potencia de cálculo para realizar las simulaciones oceanográficas, seguridad en el acceso y compartición de los recursos, un interfaz amigable y herramientas para la visualización de datos y resultados de las simulaciones.

Para la creación de este laboratorio virtual se aunaron diferentes recursos computacionales proporcionados por la Universidad de Santiago de Compostela (a través del CITIUS) y por el Centro de Supercomputación de Galicia (CESGA). La solución final implementada se basó en el uso de tecnologías web, para el desarrollo de un portal, lo que permitió ofrecer a sus usuarios finales la sencillez de uso que habían requerido. La integración de diferentes módulos de software en las diferentes capas propias de un sistema Grid permitió construir la arquitectura fundamental del laboratorio y, mediante la adopción de estándares abiertos, garantizar una buena interoperabilidad entre los diferentes componentes. El uso de técnicas de virtualización hizo posible además la optimización del uso de los recursos físicos de computación y de almacenamiento disponibles, así como una buena gestión del servidor web del portal RETELAB.

El trabajo desarrollado y la solución implementada cumplieron con los requisitos indicados por los investigadores, demostrando la posibilidad de proporcionar un entorno sencillo para el acceso y ejecución de trabajos en una infraestructura computacional remota basada en integraciones complejas de recursos distribuidos de diferentes proveedores. Los investigadores de la Red RETEMAR verificaron y validaron el laboratorio virtual a través de la ejecución de diversos códigos propios, como los dedicados al estudio de la producción primaria, o la ejecución del modelo oceanográfico ROMS (Regional Ocean Model System).

El resultado de esta investigación, y de los desarrollos llevados a cabo en ella, permitieron mejorar la capacidad de cálculo de los investigadores de RETEMAR y por tanto también los tiempos de obtención de resultados, además de ofrecerles la posibilidad de acceder de forma sencilla a datos distribuidos y de llevar a cabo simulaciones más complejas de las que eran capaces empleando sus recursos locales.

Aunque la solución implementada en el laboratorio virtual de RETELAB fue desarrollada y particularizada para la comunidad de investigadores de la Red RETEMAR, y para aplicaciones de su área, su diseño e implementación se realizaron mediante herramientas basadas en estándares abiertos, proporcionando una solución de base adecuada para dar respuesta a diferentes tipos de aplicaciones científicas y problemas de investigadores de diferentes disciplinas. Así se refleja en los trabajos de otros autores que años más tarde siguieron el mismo esquema implementado por este laboratorio virtual para dar solución a otros casos de uso.

Sin embargo, demandar recursos en una infraestructura Grid no implica que puedan ser obtenidos de forma inmediata, ni garantiza al usuario poder disponer nuevamente de los recursos empleados en simulaciones o cálculos anteriores, lo que hace difícil que puedan hacer una estimación de tiempos basándose en simulaciones previas. Son precisamente el dinamismo y la heterogeneidad, propios de los recursos de un Grid, lo que supone una dificultad, especialmente para usuarios con restricciones temporales importantes en la obtención de resultados. La dificultad para estimar el tiempo necesario para obtener recursos en Grid y, dependiendo del tipo de recursos asignados y de su estado, la duración de la ejecución de los trabajos en dichos recursos compartidos provocaron que los usuarios con restricciones temporales importantes no encontrasen en el Grid una solución adecuada para el desarrollo de sus proyectos. Otros paradigmas como la computación Cloud proporcionaron más tarde una alternativa para dar solución a este tipo de aplicaciones que, siendo adecuados para su ejecución en este modelo de computación, disponen de restricciones de tiempo importantes para su ejecución.

Computación Grid y Cloud son dos paradigmas que comparten una idea común, la de ofrecer capacidad de cálculo mediante recursos remotos que pueden ser usados de forma conjunta para resolver un problema. Sin embargo, en Cloud los recursos se proveen como un servicio bajo demanda, lo que ofrece mayor rapidez en el aprovisionamiento, mayor interactividad y generalmente mayor homogeneidad. De esta manera, en Cloud los tiempos totales de ejecución no se ven afectados por la variabilidad de los tiempos de espera en cola, implícita en computación Grid, lo que permite a sus usuarios poder realizar una mejor estimación del tiempo total necesario para completar un trabajo desde que se solicita su ejecución.

Cloud ofrece diferentes modelos de servicio que definen el nivel de abstracción en el que el usuario interactúa con el entorno, entre ellos los más importantes son: Software como Servicio (SaaS), Plataforma como Servicio (PaaS) e Infraestructura como Servicio (IaaS).

En concreto, en el modelo Cloud que ofrece infraestructuras como servicio (IaaS), los recursos elásticos y virtualizados se ofrecen bajo demanda para construir clústeres virtuales según los requerimientos de los usuarios que los demandan. El modelo IaaS puede permitir un importante ahorro de costes y ofrece gran flexibilidad para administrar y variar las configuraciones de los recursos según la necesidad en cada momento. Los servicios IaaS permiten a sus usuarios, por ejemplo, configurar diferentes entornos y experimentar con servicios en desarrollo sin necesidad de disponer de diferentes tipos de recursos físicos específicos, para luego de manera sencilla pasar dichos desarrollos a un entorno de producción específicamente configurado para el servicio final.

En el segundo caso presentado en esta memoria ha sido desarrollado dentro del proyecto europeo BonFIRE-VCOC (*Bonfire-Virtual Clusters on Federated Cloud Sites*), en él se investigaron las características de los entornos de computación Cloud para el modelo IaaS. En este caso se trabajó principalmente en torno a la propiedad de elasticidad horizontal del Cloud, que junto a la virtualización de recursos, permite variar la configuración inicial de un servicio para cumplir con las restricciones temporales requeridas por una aplicación. Además, se analizaron las capacidades de entornos Cloud distribuidos para proteger el servicio frente fallos en uno de los proveedores de un clúster, con objetivo de reubicar automáticamente parte del clúster y recuperar las condiciones necesarias para mantener calidad del servicio.

Como base del estudio se empleó una aplicación para la verificación de dosis en tratamientos de radioterapia. Esta aplicación había sido desarrollada previamente en el marco del proyecto de investigación eIMRT y está basada en métodos estadísticos de Monte Carlo. Para obtener un buen resultado con este tipo de métodos es imprescindible modelar un elevado número de historias (cada partícula primaria y sus productos secundarios) para así obtener suficiente precisión estadística. El caso de simulaciones Monte Carlo es buen ejemplo de simulaciones adecuadas para ser realizadas en Cloud, ya que requieren de comunicaciones poco frecuentes (o ninguna) entre los procesos que modelan las diferentes historias y que pueden ejecutarse en diferentes nodos de cálculo (desde segundos a horas).

Como resultado de la investigación se demostró la posibilidad de implementar mecanismos, basados en las características de virtualización y la elasticidad horizontal del Cloud, para mejorar la calidad de los servicios y cumplir con las restricciones temporales indicadas para una aplicación, implementado incluso una mecanismo para la tolerancia a fallos en una solución Cloud operando con recursos provistos por diferentes proveedores.

Las técnicas de gestión y provisión de recursos, que permitan establecer una mejor calidad de servicio, continúan siendo un reto en Cloud en la actualidad y es frecuente encontrar trabajos relacionados con este tipo de investigaciones en la literatura reciente, tanto desde el punto de vista del usuario, que demanda calidad de servicio, como de los proveedores, que tratan de encontrar la mejor solución para minimizar los tiempos de ejecución y el coste del servicio. Por su parte, la novedosa solución presentada en el último caso de uso permite cumplir con ambos objetivos: minimizar el tiempo necesario para la obtención de resultados y maximizar el aprovechamiento de los recursos, al tiempo que proporciona soluciones de valor añadido que pueden ser implementadas sobre funcionalidades básicas de los sistemas de computación de altas prestaciones de propósito general y de uso compartido.

Cuando el factor tiempo pasa de ser importante a crítico, como es el caso de servicios operacionales, la calidad del servicio se convierte en un elemento indispensable y los sistemas descentralizados, distribuidos geográficamente o virtualizados dejan de ser la opción más adecuada. Factores como la necesidad de una alta velocidad en transferencia de datos, alta capacidad almacenamiento, uso de sistemas de ficheros distribuidos o la obtención de la máxima velocidad de ejecución del código, requieren de una infraestructura dedicada específicamente a la computación de altas prestaciones. Sin embargo, por sus características, este tipo de infraestructuras son costosas, más aún cuando solo se necesita hacer un uso periódico y discontinuado de ellas, lo que hace que su adquisición y mantenimiento resulten poco viables económicamente. El uso de un servicio de computación sobre recursos dedicados supone un elevado coste que pocas organizaciones pueden asumir, además del inconveniente de que se dispone de un número limitado de recursos ante situaciones de fallo, al tiempo que se produce un desaprovechamiento de recursos y de energía durante los tiempos de inactividad. Por ello el uso de servicios de supercomputación resulta una opción más adecuada. Sin embargo, los supercomputadores de propósito general y de uso compartido presentan un reto a la hora de garantizar la ejecución de servicios de tipo operacional empleando los mismos recursos que se usan diariamente para otras aplicaciones de investigación.

El tercer caso de uso fue desarrollado en colaboración con Puertos del Estado en el marco de la serie de proyectos europeos MyOcean. El principal objetivo fue realizar la integración de un flujo de trabajo formado por servicios de computación y otros procesos necesarios para asegurar el correcto funcionamiento del servicio operacional, que ejecuta costosas simulaciones oceanográficas basadas en el modelo oceanográfico NEMO para la zona IBI (Iberia-Vizcaya-Irlanda).

Para hacer posible la correcta ejecución del servicio operacional, y la diseminación los productos resultantes en tiempo, se establecieron una serie de mecanismos y herramientas de control del flujo de trabajo que se desarrolla sobre un entorno de supercomputación de propósito general y no dedicado, el supercomputador Finis Terrae del Centro de Supercomputación de Galicia. Estas herramientas incluyen aspectos como la monitorización del rendimiento de la aplicación, en tiempo de ejecución, o la gestión de alarmas vinculadas a los diferentes pasos del flujo de trabajo, que permiten la anticipación y resolución de problemas que puedan suponer un retraso del servicio.

El trabajado desarrollado en este último caso demuestra que es posible usar con éxito una infraestructura de supercomputación de uso compartido para dar solución al mismo tiempo a trabajos de investigación de diversa índole y garantizar la calidad de servicio en casos operacionales, donde la limitación en el tiempo de obtención de resultados es un factor crítico.

Las tecnologías de computación continúan avanzando a un ritmo que hará posible llevar nuevamente la capacidad de cálculo a otros horizontes. Las familias de procesadores, las redes de interconexión y los sistemas de almacenamiento siguen ofreciendo mejoras constantes o nuevas versiones que están disponibles para su uso en mercado antes de que las últimas tecnologías adquiridas hayan podido ser rentabilizadas por sus usuarios. Desde un punto de vista económico, esta carrera continua supone inversiones constantes en nueva tecnologías y equipamientos compatibles, así como en su mantenimiento y en personal especializado. Esta es una de las razones por las que el uso de servicios externos de computación distribuida, ya sea en plataformas Grid, Cloud o supercomputadores, es clave especialmente en entornos de investigación y de servicios basados en aplicaciones científicas, que intentan resolver problemas cada vez más complejos. Sin embargo, el principal objetivo de estos avances deberá ser siempre el de su aplicabilidad, conseguir resolver nuevos retos o problemas actuales en tiempos menores. Así, y para que esto resulte posible, no se puede perder de vista el hecho de que la capacidad de análisis y resolución de los grandes problemas científicos a abordar reside justamente en los usuarios finales de las aplicaciones científicas: investigadores e ingenieros que deben poder acceder a las nuevas capacidades de cálculo de forma sencilla, sin necesidad de especializarse en la tecnología que las soporta ni de conocer su complejidad.

TABLE OF CONTENTS

TABLE OF FIGURES

TABLE OF TABLES

1 INTRODUCTION

This initial chapter aims to introduce the context, and describe the motivation and the objectives for this research work.

1.1 MOTIVATION

The advances in technology and the reduction of computer equipment costs have made possible, in recent decades, the development and continuous improvement of computer systems, especially the capacity of distributed computing systems, which are key to solve both commercial problems and scientific applications with great computational demands.

The first local area networks (LAN) gave rise to the first distributed computing systems, and these, together with the TCP / IP protocols, had an enormous impact on the computing world allowing the interconnection of local heterogeneous systems. The need for connecting computers located a long distance from one another was an immediate need and then the scope of LAN was extended to wide area networks (WAN).

Distributed systems, interconnection networks and communication protocols are essential in business, academic and research environments, and currently also in most of the everyday services and technologies. However, distributed computing has gone through different phases and evolved in various ways, giving rise to several distributed paradigms that can be more or less appropriate according to the type of problem and the specific conditions of each use case.

Many scientific problems, that are interesting to both researchers and engineers, cannot be solved using a regular computer due to the limitation of resources or the amount of time required to obtain results. Therefore, very often research groups are hindered in their ability to perform their projects because they lack the required computational resources. In addition to hardware costs, the investment of time and money needed to train staff to deploy, maintain, and effectively use such systems presents a significant barrier. In some cases researchers can have access to their own in-house dedicated systems, however in most cases these resources may end up being insufficient, inadequate or difficult to maintain updated. In addition, those local resources that can be optimal for running a type of application may not be scalable or suitable for other types of problems, which are additional disadvantages for achieving good multipurpose in-house computing systems. Therefore, very often researchers have to choose among different types of remote distributed computing systems to be able to execute their applications. Selecting the best system to solve a research problem, and to obtain the results in time, can be a complex task because specific scientific applications have different requirements and restrictions. In the selection of the most appropriate system many factors can influence the final user's decision: ease of use, final cost, parallelization level, data transfer frequency, queue times, fault tolerance, time to solution, etc.

This dissertation goes over the evolution of computing technology and some of the paradigms associated with the distributed computing architectures over the last decade, through the analysis of three real use cases developed and implemented at the Galicia Supercomputing Centre: Grid, Cloud and supercomputing infrastructures based on cluster architectures. Each case explores a different distributed computing model applied to a specific scientific problem and seeks to develop solutions valid to be implemented also for other problems. To do this, different aspects of each technology were studied, requirements and improvements were analysed, and specific solutions were designed and implemented.

All the developments presented in this work were carried out in the framework of different research projects funded through public open calls. As a result of the work presented in this dissertation it has been possible to improve several services and processes based on distributed computing technologies.

During the development of this dissertation the author has participated in several research and innovation projects (see section 8), also related to distributed and High Performance Computing, and she has published other scientific publications that complement this work and extend the results presented here. These contributions will be cited throughout this document.

1.2 RESEARCH FRAMEWORK: DISTRIBUTED COMPUTING

Since the advent of sequential computers in the market, the computing capacity has increased progressively following Moore's Law [1] which predicts that processors' performance will double every two years. This law has driven the development of semiconductors for decades. However, the possibility of increasing the computing power without increasing the chips' size has a limit that is close to be reached. In spite of this, processors' manufacturers continue developing processors that comply with Moore's Law, and it is expected that this trend will continue over the next 10 years [2][3], while providers face the challenge of commercializing 7 nanometres transistors.

The evolution of High Performance Computing has been marked by the challenge of obtaining the best performance. The computing capacity has improved significantly in the last decade, not only due to the increase in the computing power provided by new microprocessors, but also due to the possibility of combining several processors from a single or several computers connected by high speed networks. The development of local area interconnection and communications networks allows for the computers connection with the possibility of high speed data transfer. This ability to use a collection of independent computers, interconnected through a network, working as a single coherent system to solve a single large problem is called distributed computing.

Distributed computing brought many advantages, opening the possibility of addressing previously unapproachable problems, but also posed great challenges such as the coordination and management of resources, synchronization of tasks, new protocols, failures' handling, etc.

Local resource management systems (LRMS), also known as queue managers, were one of the tools that solved some of these new issues presented by distributed computing. These systems allow for the allocation of resources among different users, planning and managing the jobs that will be executed in the different computing nodes. Through the LRMS, users send their computing tasks to a distributed system, specifically to a queue where they wait to be executed until the required resources are available; often these systems follow policies based on priorities. Some of the most popular queue managers are Grid Engine (SGE) [4], Portable Batch System (PBS) [5] or Simple Linux Utility for Resource Management (SLURM) [6].

Figure 1.1 shows a comparison [7], in terms of scale and applicability, of the different types of distributed computing systems that are discussed in this work, which will be described in the following sections. In addition to the evolution of each distributed system, there have been also efforts and various research works aimed at obtaining mechanisms that allow these systems to be combined efficiently. The combination of computing clusters, Grid and Cloud infrastructures, has been called Jungle Computing [8] in the literature.

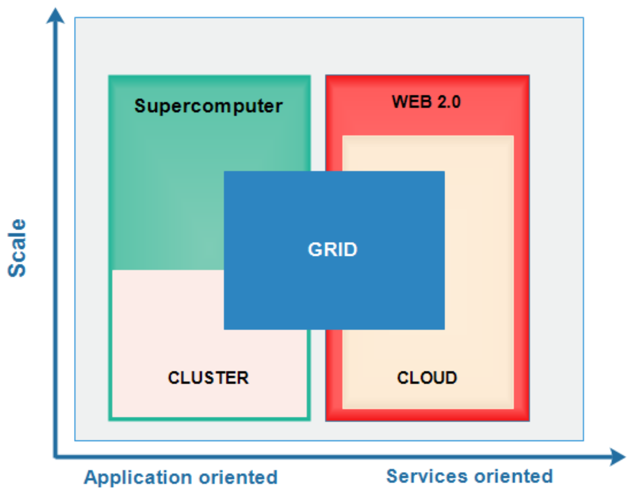

Figure 1.1 Distributed systems overview. (Figure adapted from [7])

1.2.1 Virtualization

One of the key technologies commonly used in most types of computing infrastructures, and that has special relevance in this dissertation, is virtualization. This concept appears in the 1960's and early 70's due to the need of implementing a time-sharing system that would let multiple users access the same computer simultaneously, allowing them to share the mainframes resources. In that decade IBM developed supercomputers that used this technology. With the creation of model IBM S/360-67 the virtualization of all hardware interfaces, through a virtual machine monitor, was possible for the first time. A decade later, this monitor became known as *hypervisor* due to its ability to run different systems over the one running on the underlying hardware.

However, in 1978 the first computer of the x86 architectures' family was introduced by Intel and then a new era of personal computers, multi-user operating systems, and client-server applications began. This put large scale virtualization on the back burner for several years. However, as x86-based data centres were growing the old idea of virtualization started to become important again. The existing hardware was at that time very powerful and highly efficient but there were many underexploited resources and additionally an excess of time was spent managing servers. The main problem was that, in many cases, it was not possible to take advantage of these powerful resources to host several services simultaneously or to run different types of applications, mainly due to compatibility and security issues, network configurations, or simply because each case required a different operating system. All these issues caused virtualization to recover its importance and make a strong come back with the objective of allowing for the splitting of the available hardware so that physical resources could operate as multiple independent systems. Thus, in 1998 VMware [9] was founded, it was the first company entirely dedicated to this technology that a year later brought to market its first product, the VMware workstation. Not long after this, other large suppliers such as Intel or AMD introduced their own technologies to the market, followed by many others in the coming years. Currently there is a wide range of products related to virtualization that have been incorporated into the best known computing services and infrastructures.

There are different types of virtualization, such as server virtualization, storage, applications or network virtualization. The type of virtualization that has been used in the use cases presented in this dissertation, and probably the most commonly known, is the server virtualization where different virtual machines can operate on the same physical resource. Server virtualization can be implemented in different ways [10], among the different possibilities are full virtualization, para-virtualization or operating system virtualization. All these techniques allow for a better use of available resources by dividing or combining them to provide multiple services based on the capacity of the existing equipment, according to the needs of each case.

Benefits of server virtualization technology include substantial savings in hardware, energy consumption, space requirements, and hardware administration tasks because it allows reducing the number of physical servers needed to provide a set of services. In addition, virtualization provides other important advantages [11] such as:

- **Isolation**: a failure in a virtual resource will only affect the services running on that virtual resource and not other services running on the same physical equipment.
- **Administration and security**: each virtual resource can define independent users with specific access configurations, different from other virtual machines running on the same physical resource.
- **Flexibility**: virtual resources with different characteristics of memory, number of processors, disk size, network configuration and operating system can be deployed for each service.
- **Agility**: the deployment of a new virtual resource is a fast process that can be typically done in the order of seconds.
- **Portability**: the virtual system configuration resides in one or several files, which makes copying or moving virtual resources between physical computers an easy task.

- **Rapid recovery in case of failure**: deploying a new resource, as a copy of the original one, provides a fast recovery from failures and contributes to achieve a better quality of service (QoS).
- **Elasticity**: allowing for resource reconfiguration (increase or decrease) while the system is running, according to the service load conditions

The aforementioned features have made virtualization a key technology in computing and services infrastructures due to its flexibility as well as the important savings in equipment, energy and maintenance tasks. All the case studies presented in this report have made use of server virtualisation, both for the optimization of computing resources and also for improving the management of the servers used for users interface, monitoring, and dissemination services.

Server virtualization provides a variety of benefits, one of the biggest being the ability to consolidate applications onto a single system, but this approach had its drawbacks. A virtual machine includes a separate operating system image, which adds an overhead in memory and storage footprint, and sometimes also limits its portability between different computing systems and traditional data centres. Containers are an emergent technology that attempt to simplify and accelerate the process of building and isolating applications. Applications that are run in containers provide the entire runtime image, including libraries and any other dependencies. Containers, unlike virtual machines, do not require a hypervisor; they can offer many of the benefits of virtualisation with faster start up and lower overhead. Therefore containers could reduce bottlenecks in computing, and input and output operations to achieve near bare metal performance, making them a suitable technology for running HPC applications.

1.2.2 Grid Computing

A new paradigm of distributed computing called Grid [12] arose in the early 90's as a natural evolution of the distributed systems construct emerged in previous decades. Grid computing emerged with the main aim of providing collections of distributed computing resources over a network and making them available to the end users as one large virtual computing system.

The Grid is typically a geographically distributed infrastructure composed of heterogeneous resources provided by different organizations that are not under a centralized control. Its operation is based on the interaction of a large number of services in charge of different required tasks: discovery of available resources, finding the suitable resources to execute a given job, providing access authorizations, etc.

The aforementioned services can be orchestrated in different ways, according to the specific objective of each infrastructure. This gives rise to different types of Grid: computational Grid, data Grid, collaborative Grid services, etc.

The main Grid proposal is based on the idea of creating specific and well-defined services to offer distributed computational resources through the execution of batch jobs, ensuring heterogeneous clusters interoperability.

Grid infrastructures strongly depend on hardware and communications networks, however a software layer that allows users to access computers distributed across the network is needed. This software is called "middleware" because it is located between the computers operating systems and the scientific applications used to solve the user problems. Grid middleware provides a common base between heterogeneous systems and facilitates interoperability through the use of several software packages that have different functionalities (access control, job submission, etc.) There are numerous examples of Grid middleware: Globus toolkit [13], gLite [14], Unicore [15], etc.

At the core of workload management for Grid computing is a software component, called meta-scheduler or Grid resource broker that provides a virtual layer on top of heterogeneous Grid middleware, schedulers, and resources. In this way, it is possible to achieve interoperability among different Grid middleware and provide a single client that can support multiple submission systems while users and applications only have to deal with a single protocol.

In order to achieve the interoperability among different Grid infrastructures and services the development and adoption of open standards were needed. The Open Grid Forum (OGF) [16] was the main organization leading the global standardisation effort for Grid computing. OGF accelerated the evolution and adoption of Grid by providing an open forum for Grid innovation and developing open standards for its software interoperability. Nowadays OGF operates also in the area of Cloud and other advanced distributed computing systems.

Services provided from a computational Grid were initially accessed through a command line interface (CLI). However working with Grid based applications through the CLI commands required specific knowledge that regular users (researchers from different areas) did not need to acquire to carry out their research. This is why Grid portals emerged, with the main purpose of serving as a single-entry point to hide the underlying complexity, and to provide researchers and engineers with customized tools. These portals, based on web technologies, act as bridges between the infrastructure and the user, hiding the underlying complexity of the intermediate layers and allowing less computer-savvy users to focus solely on their research tasks. Some examples of solutions that allowed for the development of these Grid portals are among others GridSphere [17], PortalLab [18] and GridPort [19].

Middleware, open standards, and Grid portals have been key tools in the development of projects related to Grid technology. Europe has led numerous projects in this field, some of them dedicated to the creation and development of large infrastructures while others were focused on the accessibility or on its use for specific fields of Science. Among the major Grid projects, it is worth mentioning TeraGrid [20] or DEISA [21] . The latter of these integrated computing and storage resources from eleven different organizations. The DataGrid project [22], initiated in 2001, demonstrated the potential of the Grid in the fields of High Energy Physics, Biology and Earth Sciences. Later, other large-scale initiatives would take the lead, using the results of DataGrid as a basis, as the series of EGEE projects [23] that encompassed scientists and engineers from more than 250 institutions and from up to 48 countries. EGEE initiative continued over time with the EGI-InSPIRE [24] and the EGI-Engage [25] projects in charge of continuing and improving the initial infrastructure, and expanding its services. Currently this infrastructure is still active and it is called EGI [26]. EGI endorses the principles of the European Open Science Cloud (EOSC), an initiative of the European Commission which is projected to be in place by 2020. EOSC aims to be the Europe's virtual environment for researchers to store, manage, analyse and re-use data for research, innovation and educational purposes. As a first step, the EOSC-hub project [27] mobilises providers from the major European research infrastructures to create a central contact point for European researchers and innovators to discover, access, use and reuse a broad spectrum of resources for advanced data-driven research.

The author has participated in the European project EGI-InSPIRE [P.7] as well as in the national project RETELAB [P.1]. The work done under the RETELAB project is described in the first case of use of this dissertation (chapter 3), which describes how a Grid infrastructure, virtualization techniques, open standards and a Grid portal were integrated in order to create a virtual laboratory able to hide the complexity of the Grid system for non-expert users.

For years, important advances in Grid technologies were made that demonstrated its power and usefulness, becoming a technology widely adopted by researchers and research projects such as those previously mentioned that have produced important success stories in different areas [28].

Despite these successes in the academic and research environments, Grid has not yet achieved the same success in the commercial field, partly due to the lack of simplicity in its use and the difficulty to guarantee the quality of service, because of its decentralized nature and the volatility of its resources.

Grid definition shows that this type of systems has a great advantage in terms of computing power by allowing many distributed heterogeneous resources, however this advantage has intrinsic drawbacks such as latency in communications or the impossibility of predicting the time needed to complete an execution. In addition to the latency issue, it is important to consider the response times required by each of the services involved in the process of executing a job in a Grid (resources management, authorization, monitoring, or data transfer). Heterogeneity in computer systems, local management systems, storage and accumulated latency make the execution time prediction of an application in any given Grid a complex task.

Additionally, because the organizations that contribute resources to a Grid maintain autonomy over their resources, these infrastructures are highly variable environments. This variability makes guarantying a given level of quality of service (QoS) in Grid very difficult to achieve. A successful commercial service offering must both provide the needed functionality and provide the needed level of QoS. The commercial providers did not find interest in Grid from their customers; therefore they focused their attention on other options such as the Cloud.

In general, a large batch of jobs, that can contain even millions of independent jobs, loosely coupled parallel applications, parametric studies and some types of workflows that assemble and orchestrate different components, are suitable applications for Grid computing. Among others, these type of applications include those from the High Energy Physics community, which frequently use parametric applications, those from Bioinformatics, which usually use workflows, or the treatment of extensive data files like the digital smoothing of satellite images in Earth observation.

In the current era of Cloud computing, Grid computing is considered as the option to create virtual supercomputers for research. One of its best known applications is currently deployed in the largest experiment in the world, the Large Hadron Collider (LHC) [29].

Some recent works on Grid that can be found in the literature explore the idea of applying Grid techniques in next generation mobile phones [30][31], using them as computational nodes of a Grid system. The evolution that mobile phones have experienced in terms of their processing capacity, connectivity, storage or the number of sensors that they incorporate have aroused interest in their potential application in distributed computing.

1.2.3 Cloud computing

The concepts behind Cloud Computing can be attributed to John McCarthy, recognized as one of the "founding fathers" of Artificial Intelligence, who in 1961 foresaw computing in the form of a global network that would act as a public service:

"... computing may someday be organized as a public utility just as the telephone system is a public utility....The computer utility could become the basis of a new and important industry".

McCarthy was correct in his predictions and his vision was truly realized. A public computing utility, allowing users of any scope and scale to access the infrastructure they need, the Cloud.

The different definitions of Cloud vary fundamentally depending on the technology use or on the main actor role (providers, integrators or final users). There is no universally accepted definition; however, there are international organizations whose objectives are the standardization of Information Technology (IT) and, in particular, Cloud computing. One of the most recognized organisms is the National Institute of Standards and Technology (NIST) [32] that, after several years and 15 previous drafts [33], finally in 2011 defined Cloud computing as:

"Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."

The advances and improvements in virtualization over the last decade boosted the development of Cloud computing, which usually uses virtualized sets of physical resources. Virtualization is a fundamental technology in the Cloud since, as it was mentioned above, it allows providers to save energy and costs, and enables the execution of different applications and operating systems in the same physical resource. However, many of these advantages are of interest not only for Cloud but also for the resources in a Grid infrastructure and even in HPC clusters or supercomputers, allowing the deployment of single resources to perform specific actions: system front-ends, data access interfaces, monitoring services, results dissemination, etc.

The NIST definition identifies five essential characteristics of Cloud computing:

- on-demand self-service
- broad network access
- resource pooling
- rapid elasticity
- accounting

It also includes three service models (software, platform and infrastructure) and four implementation models (private, community, public and hybrid) that together categorize the ways of offering Cloud services.

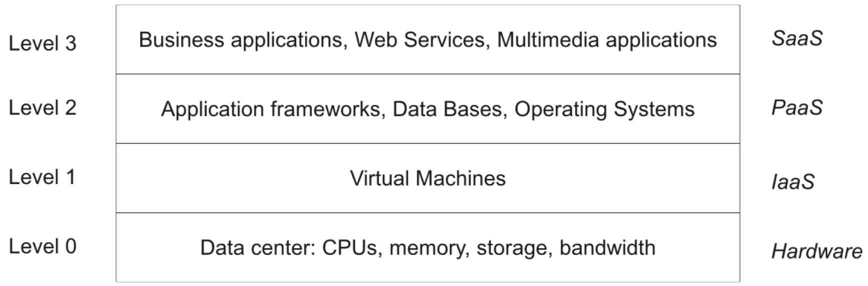

Figure 1.2 Cloud computing architecture levels.

Cloud service models define different levels of abstraction for the user interaction with a Cloud computing environment (Figure 1.2). The three models are: Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS).

In the Cloud computing environment, virtualization solutions are widely used, these solutions can include multiple layers of virtualization technologies (hardware, network, applications, etc.), which are combined in a flexible way to produce different solution models according to the specific environment required by an application.

Specifically, in the IaaS model elastic and virtualized resources are offered on demand to build virtual solutions as exact patterns of customers' requirements. The IaaS model for HPC applications can allow for cost savings and offers great flexibility to advanced users to manage and modify the configurations of computing resources according to their needs at any time. The IaaS services allows, for example, simulating different environments and testing services at different stages without the need of maintaining different physical resources, so that the system can be transferred, in a more simplified and effective way, from the testing stage to the production environment. In the same way IaaS services allow for the on demand deployment of production environments specifically configured and optimized for a service.

Grid and Cloud share similar characteristics and purposes in many aspects. Both distributed systems are composed of typical elements and processes with similar objectives: data, metadata and client nodes, as well as procedures for replication, monitoring and load balancing. However, Grid and Cloud follow different modes of operation and interaction with the user. The provision of services and resources in Cloud is done on demand while in the case of the Grid the allocation is managed through queue systems and batch jobs.

Unlike Grid Computing, Cloud Computing provides a high level of abstraction, offers real-time services, non-shared virtual resources and greater transparency for the user [7].

A series of open standards were also needed in Cloud computing in order to achieve the interoperability of their components. Organizations such as the OpenStack Foundation, the Open Group and the Open Grid Forum among others have worked in this direction. Open Grid Forum deals with the Open Cloud Computing Interface

(OCCI) [34], which includes a set of cloud-based interaction specifications. OCCI provides a protocol and API design components applicable to different Cloud management tasks with a strong focus on integration, portability, interoperability and innovation, while offering a high degree of extensibility. This standard was originally intended to create a remote administration API for services based on the IaaS model, which allows for the development of interoperable tools for common tasks, including implementation, scaling and monitoring. This was the case of the work described in chapter 4: a Cloud IaaS model using an API based on OCCI. Currently, the Open Cloud Computing Interface is also suitable for other models other than IaaS, including PaaS and SaaS.

The advances in virtualization and distributed computing technologies, and the adoption of standards dedicated to the interoperability between different systems have led to new lines of work on computational infrastructures called hybrids [35]. Supercomputers, combined with dynamic and elastic virtual resources that Cloud platforms can provide on demand. The objective of these systems is to achieve a combined use of the different computing paradigms and the attenuation of the concrete disadvantages of each one.

The exponential growth in the amount of data that can be collected from different processes or devices is currently a key factor that affects the design and operation task of the infrastructures. Technologies that allow to perform efficient calculations and analyses when large volumes of data are used become necessary. In this case, there is a change from the traditional approach of moving the data where the computing capacity is, to move the calculation to where the data are, avoiding the performance loss due to the movement of data. Technologies such as Google MapReduce, Google File System (GFS), Hadoop and Hadoop Distributed File System (HDFS), Microsoft Dryad and CGL-MapReduce adopt this approach. Both the aforementioned technologies and hybrid infrastructures have not been the objective of this work.

During the last two decades, numerous research projects, dedicated to the development and improvement of Cloud computing systems were carried out, in both academic and industrial computer systems. In the latter case, the projects dedicated to the adoption of Cloud computing into the industrial processes have acquired special importance. The author has participated in several European projects related to Cloud computing technologies, such as the BonFIRE project [P.5], for the construction of a European Cloud infrastructure for the scientific community or the Fortissimo project series [P.8][P.9] for the adoption of simulation in HPC-Cloud environments by small and medium enterprises (SME) [36].

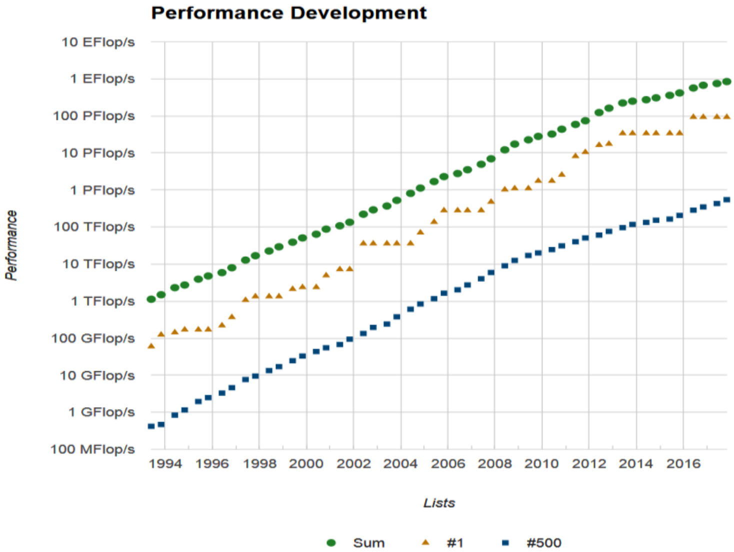

Figure 1.3 Top 500 (Nov. 2017) - Performance development.

1.2.4 Supercomputing based on cluster architectures

Supercomputers have evolved with the main aim of providing the maximum performance: the maximum computing power needed to solve a single large problem in the shortest amount of time. This trend is reflected in the periodic statistics provided by the TOP500 project, which is in charge of compiling the list of the 500 most powerful supercomputers in the world [37]. In this list the supercomputers performance is measured using the benchmark known as High

Performance Linpack (HPL). In the TOP500 List table, the computers are ordered first by their *Rmax* (the maximum performance achieved in the HPL). In the case of equal performances for different computers, they will be ordered by their *Rpeak* value (the theoretical peak performance). For sites that have the same performance, the order is by memory size and then alphabetically. Figure 1.3 shows the evolution of the supercomputers performance in this list from 1994 to 2016 based on three values: the global power (Sum), considering the accumulated of all the components of the list, the corresponding performance to infrastructures in first (# 1) and last position (# 500).

With the proliferation of general-purpose computing and affordable prices of processors some supercomputer manufacturers moved from the initial architectures based on a set of processors in shared memory architecture to large cluster architectures. Thus, most of the supercomputers that make up the Top 500 list are now cluster architectures (Figure 1.4), including accelerators such as Xeon Phi or GPUs.

Architecture	Count	System Share (%)	Rmax (GFlops)	Rpeak (GFlops)	Cores
Cluster	437	87,4	550,486,396	913,489,746	47,911,936
MPP	63	12,6	294,634,109	425,847,877	21,092,704

Figure 1.4 Top 500 (Nov. 2017) - Architecture System Share

The main differences between a cluster-based supercomputer and a conventional cluster are the larger number of available compute nodes, the high performance and low latency interconnect [38][39], and the high performance distributed file systems that are available in a supercomputer. These main features allow supercomputers to provide better performance than a conventional cluster and execute large parallel applications including those with high memory demand or frequent and fast read and write operations. On the other hand, supercomputers are usually shared systems, although it is usual that part of a supercomputer can be used exclusively during a task execution, they are generally shared general-purpose systems, not dedicated to a single type of application, as it could be the case of some clusters.

Something similar happens when Cloud computing and supercomputers are compared. Cloud infrastructures are generally based on Ethernet networks and virtual machines, while supercomputers are based on bare metal resources connected locally with low latency networks that are able to provide fast communications between thousands of distributed processes. These networks together with distributed file systems provide high performance in jobs that require or produce large volumes of data. For these reasons, it is usually considered that Cloud computing is indicated for parallel applications in which the different processes perform part of the calculation independently, but not for large-scale parallel processes with critical time constraints. The case of Monte Carlo simulations is a good example of suitable simulations to be performed in the Cloud, since such applications require infrequent communications (or none) between the processes that run on different nodes (from seconds to hours). On the other hand, supercomputers can handle parallel problems on hundreds of nodes with frequent and fast communications (in the order of milliseconds), and provide distributed and parallel file systems [40][41] that optimize operations on large volumes of data or reading operations and simultaneous writing by several processes.

Currently, hybrid commercial models are being introduced in which the advantages of a supercomputer can be accessed from a Cloud environment. Cray and Microsoft have formed an alliance to offer a service that gives access to dedicated Cray supercomputers that are accessible through the Azure platform [42]. Both providers announced previously unimaginable performance and scalability for an important set of applications executed in the public Cloud scenario.

Common use cases in supercomputing systems are those related to the meteorological and oceanographic forecast areas. These types of cases, in addition to research and analysis projects [43][44], demand services in operational mode that require periodic executions, generate large amounts of data and must produce and disseminate results with strict time constraints. For these types of services, obtaining maximum performance is a critical factor for achieving the lowest possible execution time, ensuring results within the mandatory time frame. For this, the applications to be executed on the HPC systems must be properly configured and optimized to find the optimal parameters for obtaining the maximum performance in the system.

Supercomputers, like many clusters, use local resource management systems to coordinate and distribute tasks execution. These systems can be configured to set advance reservations of the required resources for the execution of an application or a workflow, adapting, for example, to the computation requirements and deadlines of an operational service. Chapter 5 will describe how the execution of an operational oceanographic service is carried out on a general purpose supercomputer coexisting with the execution of other types of scientific workloads and regular computing services.

1.3 OBJECTIVES

This dissertation presents different distributed computing use cases that focus on three specific scientific applications. The works developed and the solutions achieved are not applicable only to these cases but also to other scientific applications with similar characteristics, both for research environments and for operational services.

The overall objective has been to simplify the execution of scientific applications on distributed environments with a special emphasis in obtaining the in-time results, according to the requirements of each case. Different cases of distributed computing applied to the resolution of different problems have been studied. In each case, the necessary developments were designed and implemented to solve weaknesses, improve the systems functionalities, and the application and workflows performance on the selected technology.

Table 1.1 shows a comparison of the main characteristics of the three use cases that are described in the following sections.

Table 1.1 Use cases comparison

	CASE 1	CASE 2	CASE 3
Infrastructure	Grid	Cloud	Supercomputer
Resources	Heterogeneous	Heterogeneous	Homogeneous
Virtualized resources	Computing and user interface	Computing, user interface and monitoring	Monitoring and results dissemination
Local resource manager	Torque	Non-applicable	SGE
Middleware	Globus Toolkit and GridWay	OpenNebula and OpenStack	OpenNebula (dissemination server)
User interface	Web portal	CLI and web portal	CLI and web portal
Main applications	ROMS and primary production model	EIMRT (Monte Carlo)	NEMO
Parallelization	OpenMP and MPI	SPMD	MPI
Monitoring	Jobs	Jobs and application	Jobs, application and resources.
Availability	Batch mode	On demand	In advance reservation
Resources type	Physical and virtual resources (shared)	Virtual resources (exclusive)	Physical (non-dedicated) and virtual resources (exclusive).
Time to solution	Not urgent	Important	Critical

2 USE CASES OVERVIEW

This section briefly introduces the cases of study that will be described in detail in the publications presented in chapters 3 to 5 and summarizes their main results. Additionally, references to other publications of the author are cited, these contributions are not part of the group of papers presented in this dissertation, however they do indeed complement the work carried out in each one.

2.1 A VIRTUAL LABORATORY BASED ON GRID TECHNOLOGY.

The first case of study was developed within the framework of the RETELAB project [P.1]. The objective of this first case was the creation of a collaborative and distributed work environment configured as a virtual laboratory for the development of interdisciplinary research projects related to oceanographic remote sensing, within the framework of the National Oceanographic Remote Sensing Network (RETEMAR). The aim of the virtual laboratory was to provide the computing capacity and storage, required by this scientific community, through the use of Grid technologies and an user friendly interface.

Laboratory development focused on providing a single secure web interface for researchers to computational resources (computing, storage, data, and applications) and other services, while hiding the complexity of the underlying hardware and software of the distributed computing environment. The laboratory infrastructure was composed by several computational resources provided by USC and CESGA and combined as a Grid platform. The final solution was based on the use of web technologies, through Gridsphere solution, different components of the Globus Toolkit middleware [45] and the adoption of open standards such as DRMAA [46]. The use of virtualization techniques allowed the optimization of the available computing and storage resources as well as a suitable management of the virtual laboratory web server

Additionally, an improvement in the management of local user accounts was implemented which avoided the overload of creating several local *Unix* accounts in each local resource involved in the Grid. The solution uses only one single *Unix* account, associated with the portal, which delegates the local identification of researchers exclusively in their credentials of X509 certificate. This feature was reported as a potential functionality to be incorporated into the GridWay meta-scheduler, in order to improve its use in combination with Grid portals. This proposal was positively considered by the GridWay development team and the new functionality was incorporated both in the production version of the RETELAB portal and also in the following public GridWay versions (5.6).

The virtual laboratory was successfully validated by the partners of the RETELAB project through the use of different applications such as the oceanographic model ROMS (Regional Ocean Modeling System) [47][48] the application "*sentinazos*" [49] and other custom codes for ocean studies that were developed by researchers from the Canary Institute of Marine Sciences (ICCM).

2.1.1 Complementary publications

Find hereafter additional complementary papers related to the work done in this case study published with the contribution of the author:

- D. Mera, J. M. Cotos, **C. Cotelo**, Y. Sagarminaga, J. Pérez, "Desarrollo e Implementación de un Laboratorio Virtual para la Teledetección Oceanográfica basado en Grid", XIII Congreso de la Asociación Española de Teledetección (AET09), Calatayud 23-26 September, 2009.
- A. Gómez, **C. Cotelo**, J. I. López, D. Mera. "Metaschedulers in the environment of eScience portals: a case study with GridWay". EGEE'09 conference, 21-25 September 2009. Barcelona, Spain.
- David Mera, José M. Cotos, José Varela, **Carmen Cotelo** and J. Ignacio López, "An Integration of Several Technologies in the Architecture Definition and Deployment of a Geospatial Grid Web Portal", The 2009 International Conference on Grid

Computing and Applications (GCA’09), pp. 86-91, Las Vegas 13-16 July, 2009.

- David Mera, José Manuel Cotos, Joaquín A. Trinanes, **Carmen Cotelo**, "An Integrated Solution to Store, Manage and Work with Datasets Focused on Metadata in the RETELAB Grid Project". IWANN (2) 2009: 491-494.
- **Carmen Cotelo Queijo**, Andrés Gómez Tato, J. Ignacio López Cabido and David Mera Pérez. ”Metaschedulers in the environment of eScience portals: a case study with GridWay”, 3rd Iberian Grid Infrastructure Conference Proceedings, Eds, Vicente Hernández García, Gaspar Barreira, Ignacio Blanquer Espert and Jorge Gomes, pp 410-419, ISBN 978-84-9745-406-3, 2009.
- David Mera, José M. Cotos, José R. R. Viqueira and **Carmen Cotelo.** ”Software Integration in the Development of a Spatial Data Grid Prototype based on Metadata”, 3rd Iberian Grid Infrastructure Conference Proceedings, Eds, Vicente Hernández García, Gaspar Barreira, Ignacio Blanquer Espert and Jorge Gomes, pp 305-314, ISBN 978-84-9745-406-3, 2009.

2.2 CLOUD-BASED SIMULATIONS on FAULT-TOLERANT VIRTUAL CLUSTERS

The experiment "VCOC - Virtual Clusters in Federated Sites" was one of four experiments that were executed in the first part of the European project BonFIRE. The objective of this experiment was to investigate the feasibility of using multiple Cloud resource providers for the provision of services that require the assignment of clusters of virtual computing nodes intended for the execution of applications that can divide a problem into separate groups of tasks, a type of calculation that fits well to the characteristics of a Cloud infrastructure.

In applications where a significant amount of computing resources is required, and the time to solution becomes important, the possibility of failures at some point in the system is high and the effects counterproductive. The horizontal elasticity of the Cloud was used as the basis for implementing a fault tolerant system in the provision of a software as a service (SaaS), allowing providers to improve the Quality of Service (QoS). As in the first use case, this one can also be applied to research environments, where researchers (from public and private institutions) need to configure sets of resources for the development of scientific or technical projects with specific characteristics that can vary along the course of their projects.

As for the scientific software, an application for the verification of doses in radiotherapy treatments was used. This application had previously been developed under the eIMRT project [50][51] and based on Monte Carlo methods [52]. This statistical method has proven to be very accurate for the calculation of doses in radiotherapy treatments. Through mathematical models, Monte Carlo allows for the reproduction of particles' interactions and trajectories from the knowledge of the probability distributions of the physical processes involved. To obtain a good result with this type of methods it is essential to model a high number of events (each primary particle and its secondary products) in order to obtain sufficient statistical precision. The application selected for this case of study executes a workflow consisting of five stages. Two of these are responsible for executing hundreds of computing jobs.

During the development of this experiment several questions related to the use of virtual clusters in a distributed and multi-supplier environment, were studied. A first set was related to those questions dealing with the time necessary for cluster deployment, the variation of the number of nodes that compose it, as well as other factors that could affect these times (interference with other users, network, etc.). The main objective was to obtain a specific solution and greater knowledge applied to the management of clusters to guarantee timely results. In order to achieve this, Cloud elasticity was analysed in detail and used in conjunction with information from the application performance monitoring. Based on the dose application performance monitoring information the system knows when the cluster needs to be expanded to guarantee in-time results.

With the same objective, the study was extended to a new scenario; the same cluster was deployed in two Cloud infrastructures from different providers, in order to obtain a final configuration that would allow for the provision of a fault-tolerant service (provider, network...). For this, several scenarios were simulated in which part of the cluster was forced to fail or the application performance was decreased until reaching the pre-established execution time became endangered.

The results of the project, in addition to being incorporated as an open source toolkit [53] for the BonFIRE platform, were subsequently taken as the basis for the design and development of service quality policies for other services offered by CESGA.

2.2.1 Complementary publications

Find hereafter additional complementary papers related to the work done in this case study published with the contribution of the author:

- Valín R., Carril L.M., Mouriño J.C., **Cotelo C**. and Fernández C. "Enabling the Deployment of Virtual Clusters on the VCOC Experiment of the BonFIRE Federated Cloud" CLOUD COMPUTING 2012, The Third International Conference on Cloud Computing, GRIDs, and Virtualization pp. 237-242. ***Best Paper Award.***
- Gómez A., Carril L.M., Valín R., Mouriño, J.C. and **Cotelo C.** Experimenting Virtual Clusters on distributed Cloud environments using BonFIRE. Fire Engineering Workshop 2012.

2.3 OPERATIONAL SERVICES RUNNING ON A GENERAL-PURPOSE SHARED SUPERCOMPUTER

The third case study was focused on the development, commissioning and improvement of an operational oceanographic service within the framework of the MyOcean projects [54].

EuroGOOS [55] defines Operational Oceanography as the activity of systematic and long-term routine measurements of the seas and oceans and atmosphere, and their rapid interpretation and dissemination. An Operational Oceanography service usually starts with the transmission of observational data to data assimilation centres, where powerful computers process the data, run numerical forecasting models and produce outputs. These outputs are used to generate final products, such us warnings (coastal floods, ice and storm damage, harmful algal blooms and contaminants, etc.), prediction of primary productivity, ocean currents, etc. Finally, the forecasts products should be distributed rapidly to industrial users, government agencies, and regulatory authorities.

According to the TOP500 list, as of November 2016 there were 23 supercomputers on this list dedicated to either meteorological or oceanographic forecasting. That doesn’t include a rather large number of systems housed at national labs that devote a portion of their resources to supporting climate and weather research. In fact, many countries without a specific meteorological infrastructure rely on these general-purpose research supercomputers to perform this work. The use of a computer service on dedicated resources implies a high cost that few organizations can assume, a limited number of resources in situations of failure, while a waste of resources and energy during downtime between successive executions.

The use of supercomputing for atmospheric and oceanographic operational services allows a quick turn-around time for these simulations so that severe events can be predicted well ahead of time. General-purpose research supercomputers can support operational services and they are a good option for those centres without their own dedicate infrastructures. However, in a shared environment, the required mechanisms for achieving a successful coexistence of these operational services with other regular workloads must be implemented. Chapter 5 describes the design and implementation of a centralized workflow, and the necessary mechanisms, that guarantee obtaining results in time for a daily oceanographic operational service running in a non-dedicated HPC environment (Finis Terrae supercomputer) [56].

The execution of different applications and processes on a non-dedicated system, by different users, means that the system state can be altered between two operational executions due to many different factors and thereby introduce errors or delays in the results delivery. The use of non-dedicated nodes for this kind of services brings advantages for both the service supplier and the supercomputing centre. The operational service is not limited to use the same set of nodes to execute the daily simulations, so it is possible to replace one or several nodes in case of failure without a service readjustment or spending time in the configuration of a new node. In a similar way, the HPC Centre can maintain the management and maintenance of all nodes in a homogeneous way, removing and adding nodes without special consideration when a maintenance task is required.

The objective of this work was to perform the integration of a complete workflow composed by computing and storage services, and different virtual services to ensure the correct operation of the service which executes large simulations based on the oceanographic model NEMO [57]. Although the part of the process that consumes most of the time and resources is undoubtedly the NEMO simulation, the complete workflow covers different required steps before and after the execution of the simulation. All of these steps are important to achieve the dissemination of results in time. Dissemination is the last step of the workflow and the one that establishes the deadline. To make possible the correct execution of the service, a series of control mechanisms and some specific tools were developed and integrated in the workflow over a non-dedicated HPC environment.

Based on the information provided by the NEMO outputs during its execution, and the deadline set for the results dissemination, an intermediate layer of applications was developed that also integrated monitoring and warning systems, allowing researchers and service operators to react when a decrease in the performance exceeds the accepted threshold.

The workflow of this service has been already migrated from the Finis Terrae system to the new Finis Terrae II supercomputer, keeping most of the tools and methods originally used in the previous supercomputer and which are detailed in chapter 5 of this dissertation. The results of this use case are therefore still active in the daily execution of the operations of *Puertos del Estado* (the Spanish Port Authority) at CESGA, using the same computing resources that are used by other researchers for different purposes on an everyday basis.

2.3.1 Complementary publications

Find hereafter additional complementary papers related to the work done in this case study published with the contribution of the author:

- Sotillo M.G., Cailleau S., **Cotelo C**., Lorente P., Levier B., Aznar R., Amo-Baladrón A. , Reffray G. , Alvarez-Fanjul E. Dynamics of the Iberian waters derived from MyOcean IBI products (forecasted and reanalyzed) for the period 2002-2012. EOF 2014, Las Palmas de Gran Canaria 11-13 July, 2014.
- Sotillo M.G., Cailleau S., Lorente P., Reffray G., Drevillon M., Queralt S., Chanut J., **Cotelo C**., Levier B., Álvarez Fanjul E., El Sistema Operacional IBI-MFC: Balance de un año de servicio de predicción oceanográfica regional y perspectivas futuras (MyOcean 2). II Encuentro de la Oceanografía Física Española (EOF-2012) November 2012.
- M.G. Sotillo, S. Queralt, J. Chanut, G. Reffray, P. Lorente, **C. Cotelo**, E. Álvarez-Fanjul. El sistema MyOcean de predicción oceanográfica operacional para la zona IBI (Iberia-Vizcaya-Irlanda): una nueva herramienta de oceanografía operacional a disposición del sistema portuario. XI Jornadas Españolas de Ingeniería de Costas y Puertos La Palmas de Gran Canaria 5-6 May, 2011.

3 USE CASE 1: GRID

Future Generation Computer Systems 26 (2010) 1157_1164

https://doi.org/10.1016/j.future.2010.05.018

RETELAB: A Geospatial Grid Web Laboratory for the Oceanographic Research Community

Carmen Cotelo[a,*], Andrés Gómez[a], J. Ignacio López[a], David Mera[b], José M. Cotos[b], J. Pérez Marrero[c], Constantino Vázquez[d]

a Galicia Supercomputing Centre (CESGA),Avda. de Vigo s/n Campus Sur, 15705 Santiago de Compostela, Spain.

b Systems Laboratory, Electronics and Computer Science Department, Technological Research Institute, University of Santiago de Compostela, Campus Sur, 15782 Santiago de Compostela, Spain.

c Canary Institute of Marine Sciences (ICCM), 35200 Telde Gran Canaria, Spain.

d Departamento de Arquitectura de Computadores y Automática, Facultad de Informática, Universidad Complutense de Madrid, Spain.

4 USE CASE 2: CLOUD

Future Generation Computer Systems 34 (2014) 17–25

https://doi.org/10.1016/j.future.2013.12.027

Fault-tolerant virtual cluster experiments on federated sites using BonFIRE

A. Gómez [a], L.M. Carril [a], R. Valin [a], J.C. Mouriño [a], C. Cotelo [a,b]

a Supercomputing Center of Galicia, Campus Vida, Avda. de Vigo S/N, 15705 Spain

b Departamento de Electrónica e Computación, Universidade de Santiago de Compostela, Campus Vida, 15782 Spain

5 USE CASE 3: SUPERCOMPUTERS

The International Journal of High Performance Computing Applications Vol 32, Issue 2, pp. 302 - 313 (2018).

https://doi.org/10.1177/1094342017692045

On the successful coexistence of oceanographic operational services with other computational workloads

Carmen Cotelo[1,2], María Aránzazu Amo Baladrón[3], Roland Aznar[3], Pablo Lorente[3], Pablo Rey[1] and Aurelio Rodríguez[1]

1 Galicia Supercomputing Centre (CESGA), Spain.

2 Departamento de Electrónica e Computación, Universidade de Santiago de Compostela, Spain.

3 Puertos del Estado (PdE), Spain.

6 CONCLUSIONS AND FINAL DISCUSSION

6.1 CONCLUSIONS

This dissertation work has being driven by three key ideas: computing requirements and time restrictions of scientific applications, optimization of computing resources' usage, and ease of use for non-expert users.

The present work has analysed, designed and implemented the use cases of three scientific applications deployed in different types of distributed infrastructures within a period of time in which there were great advances in research and development of computing technologies, communications networks and virtualization.

Firstly a solution based on Grid computing was designed and implemented while there was a boom of Grid technology and related tools. **The work developed in chapter 3 showed that it is possible to provide a friendly interface for researchers that allows them to access a remote computing infrastructure based on complex integrations of different distributed resources**. The implemented solution provided a simple web portal to access a complex infrastructure for the researchers of the RETEMAR Network without missing key aspects as security. This portal allowed RETEMAR users to carry out research works and operational services that they could not perform in their local resources. Computing capacity was improved and the time for obtaining results was reduced.

Although the solution was customized for the RETEMAR community and for applications in the oceanographic area, as its design and implementation were based on open standards, this solution provides a suitable base to be applied to other areas and applications. In fact, solutions developed some years later adopted the scheme implemented by the RETELAB virtual laboratory [58].

However, resources in Grid are allocated through batch systems; this means that researcher's jobs will be waiting in a queue until enough resources are available to run their simulations. Additionally, due to the decentralized and dynamic nature of the Grid there is no guarantee to allocate the same resources already used in previous jobs. The dynamism and the heterogeneity of Grid resources mean difficulties for those users with time restrictions for the achievement of their results. Applications with important temporal restrictions did not find in the Grid a reliable solution. Cloud technologies provided an alternative solution for those types of applications that, being suitable for execution in the Cloud, have important time restrictions for their execution. Cloud computing provides scientists with a revolutionary model for using computing infrastructure and applications. Based on virtualization, computing resources and middleware can be accommodated, and delivered as an ubiquitous-on-demand service according to any given application's requirements.

Grid and Cloud computing are two paradigms that share the idea of offering computing capacity through remote resources that can be used together to solve a problem. However, in Cloud computing the resources are provided on demand which offers faster provisioning, interactivity and homogeneity. In this way, the execution times in Cloud are not affected by the variability of waiting times in the batch systems queue implicit in Grid computing, and it allows a better estimate of the total time needed to complete a job.

In the second case presented in this document, the characteristics of Cloud environments were investigated with a special focus on horizontal elasticity. Cloud elasticity together with the virtualization of resources, allows providers to modify a service configuration to comply with the time restrictions for an application. **As a result of this research, the possibility of implementing specific mechanisms was demonstrated (chapter 4) based on virtualization and the horizontal elasticity of the Cloud. These mechanisms allow for the improvement of QoS and comply with the time restrictions indicated for an application, taking into account even the fault tolerance capabilities.**

On demand resource management and provisioning are key concerns in Cloud computing environments, very important in order to achieve better QoS. Both issues are still a challenge in Cloud research and it is common to find research works related to these issues in recent literature [59] both from the point of view of the user (that demands quality of service) as well as from the point of view of the providers (that try to find the best solution to minimize the execution times and the cost of the service) [60].

Finally, the novel solution presented in chapter 5 allows meeting several objectives: to minimize service time and maximize the use of resources; to meet the imposed QoS; and to provide user-friendly tools for supporting the researchers work while providing added value services that can be implemented in most HPC systems.

When time becomes a critical factor, as in the case of operational services, QoS becomes essential. Decentralized, geographically distributed or virtualized systems are no longer the best option. Factors such as the need for high speed data transfer, storage capacity, distributed file systems, or obtaining the maximum performance running applications require an infrastructure dedicated specifically to high performance computing. However, due to its characteristics, an HPC infrastructure is very expensive, even more when only a discontinued use of the resources is needed, which makes inviable its acquisition and maintenance**. The work presented in chapter 5 demonstrates that it is possible to use a shared supercomputing infrastructure for providing, at the same time, a platform for running regular research jobs and for operational services with critical time limitations meeting the required QoS.**

6.2 DISCUSSION

Computing technologies are continuously evolving. Processor families, interconnection networks and storage systems, among others, continue to offer constant improvements or new versions that sometimes are available in the market before previous ones could have been profitable. From an economic point of view, this means continuous investments in new technologies and compatible equipment and infrastructures, as well as in their maintenance and specialized personnel. This is one of the reasons why the use of external distributed computing services, whether on Grid, Cloud or supercomputer platforms, is essential for research and services based on scientific applications. From a computational point of view, the reasons are multiple: ability to solve problems faster, higher accuracy results, more complex and new types of scientific problems.

The evolution of these computing capabilities has allowed techniques such as artificial intelligence, and in particular machine learning techniques, to experience remarkable progress. There are multiple tools that attempt to simplify the developments of the models used in these fields, reducing the barriers in their implementation and adoption. Even so, experiments in deep learning require considerable time and numerous computational resources [61] in order to find the best models and deliver accurate results in affordable time, simple but broad tasks that are often only feasible through the use of multiple computing resources (CPUs or GPUs).

Projects and products involving machine learning often have varying priorities, ranging from speed to accuracy to reliability. These requirements introduce a significant amount of non-trivial tasks into the researchers' workflow, such as GPU assignment, distribution of code and data on the computing resources, etc. As a result, researchers often have to spend too much time dealing with issues unrelated to their research. Thus, the simplification of user environments is still needed to ease the interaction of researchers with new computing models allowing them to obtain maximum performance without spending too much time learning non-trivial auxiliary tasks that are not central to their research. Currently, as already happened with Grid and Cloud, the area of machine learning is working on the development of these user environments that attempt to simplify the use of these emerging technologies [62][63].

A new generation of computing driven by large companies such as IBM, Google, Intel or Rigetti is emerging strongly, Quantum Computing. It promises incredibly powerful machines that take a new approach using the principles of quantum mechanics. IBM, one of the current leaders in quantum computing, estimates that in five years the effects of this technology will reach beyond the research laboratories and will be widely used by new categories of professionals and developers as the alternative to solve problems that were considered unsolvable in classical computing [64]. Extremely complex problems, such as those that are addressed at the level of machine learning; financial risk analysis, genetics or cryptography are some of the first target areas. Quantum computing takes off but the technology and its associated concepts are still strange and the access to these computers strongly relies on a highly specialized scientific community. Thus, initiatives as the "IBM Q Experience" emerge, a platform that through the Cloud facilitates the open use of the first series of quantum computers in the world from the web. This portal allows users to easily experiment with a quantum processor, learn to program it, and create and test algorithms.

Computing technology has been growing exponentially over decades and it has enabled impressive improvements in cost, speed, and accuracy, in addition to providing predictive capabilities. These advances will continue and computing technology will open up new horizons. However, the main objective should be always to ensure that applications benefit from these advances and from the computing power emerging technologies can provide to solve new challenges or current problems in shorter times. To reach this objective, we cannot lose sight of a key issue; the capacity for analysing and solving the big problems resides in the final users of scientific applications, researchers and engineers who must be able to access the computing infrastructures without having to specialize in the technology that supports them or to know their intricacies or underlying complexities.

7 BIBLIOGRAPHY

[1] G. E. Moore, ‘Cramming more components onto integrated circuits (Reprinted from Electronics, pg 114-117, April 19, 1965)’, *Proc. Ieee*, 1965.

[2] IBM, ‘IBM press release, 09 Jul 2015’. [Online]. Available: https://www-03.ibm.com/press/us/en/pressrelease/47301.wss. [Accessed: 21-May-2018].

[3] M. B. (Intel), ‘Technology Leadership - Technology and manufacturing day’. [Online]. Available: https://newsroom.intel.com/newsroom/wp-content/uploads/sites/11/2017/09/mark-bohr-on-intels-technology-leadership.pdf. [Accessed: 21-May-2018].

[4] W. Gentzsch, ‘Sun Grid Engine: Towards creating a compute power grid’, in *Proceedings - 1st IEEE/ACM International Symposium on Cluster Computing and the Grid, CCGrid 2001*, 2001.

[5] Henderson R.L, ‘Job Scheduling Under the Portable Batch System’, *Lect. Notes Comput. Sci.*, vol. 949, 1995.

[6] D. G. Feitelson and L. Rudolph, *Job scheduling strategies for parallel processing : IPPS ’95 Workshop, Santa Barbara, CA, USA, April 25, 1995 : proceedings*. 1995.

[7] I. Foster, Y. Zhao, I. Raicu, and S. Lu, ‘Cloud computing and grid computing 360-degree compared’, *Grid Comput. Environ. Work. 2008 (GCE ’08)*, 2008.

[8] F. J. Seinstra *et al.*, ‘Jungle Computing: Distributed Supercomputing beyond Clusters, Grids, and Clouds’, *Grids, Clouds Virtualization, Comput. Commun. Networks*, 2011.

[9] ‘VMware’. [Online]. Available: https://www.vmware.com. [Accessed: 21-May-2018].

[10] VMware, ‘Understanding Full Virtualization, Paravirtualization, and Hardware Assist’, *Memory*, 2007.

[11] ‘Virtualization Essentials: VMware, Inc.’ [Online]. Available: https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/VMware_paravirtualization.pdf. [Accessed: 21-May-2018].

[12] I. Foster and C. Kesselman, *The Grid: Blueprint for a New Computing Infrastructure*. 1999.

[13] I. Foster, 'Globus toolkit version 4: Software for service-oriented systems', *J. Comput. Sci. Technol.*, 2006.

[14] E. Laure *et al.*, 'Programming the Grid with gLite', *Comput. Methods Sci. Technol.*, 2006.

[15] D. W. Erwin, 'UNICORE - a Grid computing environment', *Concurr. Comput. Exp.*, 2002.

[16] 'Open Grid Forum (OGF)'. [Online]. Available: http://www.ogf.org. [Accessed: 21-May-2018].

[17] J. Novotny, M. Russell, and O. Wehrens, 'GridSphere: an advanced portal framework', in *30th Euromicro Conference*, 2004.

[18] M. Lil, P. Van Santen, D. W. Walker, O. F. Rana, and M. A. Baker, 'PortalLab: A web services toolkit for building semantic grid portals', in *Proceedings - CCGrid 2003: 3rd IEEE/ACM International Symposium on Cluster Computing and the Grid*, 2003.

[19] M. Thomas, S. Mock, M. Dahan, K. Mueller, D. Sutton, and J. R. Boisseau, 'The GridPort toolkit: A system for building Grid Portals', *IEEE Int. Symp. High Perform. Distrib. Comput. Proc.*, 2001.

[20] N. Wilkins-Diehr, D. Gannon, G. Klimeck, S. Oster, and S. Pamidighantam, 'Teragrid science gateways and their impact on science', *Computer (Long. Beach. Calif).*, 2008.

[21] W. Gentzsch *et al.*, 'DEISA: E-Science in a collaborative, secure, interoperable and user-friendly environment', in *eChallenges e-2010 Conference*, 2010.

[22] L. Momtahan and A. Martin, 'E-science experiences: Software engineering practice and the EU DataGrid', in *Proceedings - Asia-Pacific Software Engineering Conference, APSEC*, 2002.

[23] Erwin Laure and Bob Jones, 'Enabling Grids for e-Science: The EGEE Project', Geneva, Switzerland, 2009.

[24] 'EGI-InSPIRE'. [Online]. Available: https://wiki.egi.eu/wiki/EGI-InSPIRE:Main_Page. [Accessed: 21-May-2018].

[25] 'EGI-engage'. [Online]. Available: https://wiki.egi.eu/wiki/EGI-Engage. [Accessed: 21-May-2018].

[26] 'EGI'. [Online]. Available: http://www.egi.eu. [Accessed: 21-May-2018].

[27] 'EOSC-hub project'. [Online]. Available: https://www.eosc-hub.eu. [Accessed: 04-Jul-2018].

[28] EGI.eu, 'EGI case studies', 2017. [Online]. Available: http://ec.europa.eu/newsroom/dae/document.cfm?doc_id=42596. [Accessed: 21-May-2018].

[29] 'LHC'. [Online]. Available: http://wlcg.web.cern.ch. [Accessed: 21-May-2018].

[30] S. C. Shah, 'Recent Advances in Mobile Grid and Cloud Computing', *Intell. Autom. Soft Comput.*, pp. 1–13, Feb. 2017.

[31] R. García, L. Flórez-Valencia, and M. Curiel, 'On existing mobile grids for Android devices', in *2016 8th Euro American Conference on Telematics and Information Systems, EATIS 2016*, 2016.

[32] 'NIST'. [Online]. Available: https://csrc.nist.gov/projects/cloud-computing.

[33] 'Cloud definition by NIST -'. [Online]. Available: http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf. [Accessed: 21-May-2018].

[34] 'OCCI - Open Cloud Computing Interface'. [Online]. Available: http://oci-wg.org. [Accessed: 21-May-2018].

[35] G. Mateescu, W. Gentzsch, and C. J. Ribbens, 'Hybrid Computing-Where HPC meets grid and Cloud Computing', *Futur. Gener. Comput. Syst.*, 2011.

[36] Carmen Cotelo, Gonzalo Ferro, Andrés Gómez, Jose Otero, 'Enabling SMEs to benefit from HPC based OpenFOAM simulation services', in *OpenFOAM 2nd Iberian Meeting*, 2018, p. 26.

[37] 'Lista TOP 500'. [Online]. Available: http://www.top500.org.

[38] N. J. Boden *et al.*, 'Myrinet: A Gigabit-per-Second Local Area Network', *IEEE Micro*, 1995.

[39] A. Ghiasi, 'InfiniBand-The Cluster Interconnect', in *Handbook of Fiber Optic Data Communication*, 2008.

[40] C. F. Sy and I. Stems, 'Lustre: A Scalable, High-Performance File

System', 2002.

[41] F. Schmuck and R. Haskin, 'GPFS: A Shared-Disk File System for Large Computing Clusters', *Proc. First USENIX Conf. File Storage Technol.*, 2002.

[42] 'Cray and Azure alliance', *azure.microsoft.com*, 2017. [Online]. Available: https://azure.microsoft.com/es-es/blog/cray-supercomputers-are-coming-to-azure. [Accessed: 21-May-2018].

[43] Antonella Galizia and Carmen Cotelo Queijo, 'Exploiting the FinisTerrae supercomputer to improve the characterization of spatio-temporal predictability of severe weather events', in *Geophysical Research Abstract. EGU General Assembly*, 2011, p. Vol. 13 EGU2011-13791.

[44] M. R. Marcos Cobas, Elena Vázquez Cendón, Andrés Gómez, Breogán Hombre, Joaquín Triñanes, Carmen Cotelo, Pablo Carracedo, Pedro Costa, 'First Approach to the application of operational 4dvar data assimilation to the ROMS model', in *Numerical methods for hyperbolic equations: Theory and Applications*, 2013, pp. 393–396.

[45] 'GridWay'. [Online]. Available: http://www.gridway.org. [Accessed: 21-May-2018].

[46] P. Troger, H. Rajic, A. Haas, and P. Domagalski, 'Standardization of an API for Distributed Resource Management Systems', in *Seventh IEEE International Symposium on Cluster Computing and the Grid (CCGrid '07)*, 2007.

[47] 'MyROMS'. [Online]. Available: http://www.myroms.org. [Accessed: 21-May-2018].

[48] C. C. Queijo, A. G. Tato, I. L. Cabido, and J. M. C. Yañez, 'Adapting ROMS to execute on GRID using a hybrid parallelization model', in *Proceedings - The 2nd International Conference on Advanced Engineering Computing and Applications in Sciences, ADVCOMP 2008*, 2008.

[49] D. Mera, 'Desarrollo, Despliegue y Validación de un Laboratorio Virtual Oceanográfico basado en Computación GRID', Santiago de Compostela, 2013.

[50] M. P. C. José Carlos Mouriño Gallego, Andrés Gómez, Carlos

Fernández Sánchez, Francisco J. González Castaño, Daniel Amós Rodríguez Silva, Javier Pena García, Faustino Gómez Rodríguez, Diego González Castaño, ‘The eIMRT Project: Planning and Verification of IMRT treatments on Grid’, in *Proceedings of Ibergrid*, 2007.

[51] ‘e-IMRT project’. [Online]. Available: http://eimrt.cesga.es/. [Accessed: 21-May-2018].

[52] A. Gómez *et al.*, ‘Remote radiotherapy planning: the eIMRT project.’, *Stud. Health Technol. Inform.*, 2006.

[53] ‘BONFIRE Demo Kit’. [Online]. Available: http://forge.cesga.es/projects/bonfiredemokit. [Accessed: 21-May-2018].

[54] ‘MyOcean’. [Online]. Available: http://www.copernicus.eu/projects/myocean. [Accessed: 21-May-2018].

[55] ‘EuroGOOS - European Global Ocean Observing Systems’. [Online]. Available: http://eurogoos.eu/. [Accessed: 21-May-2018].

[56] ‘Finis Terrae supercomputer’. [Online]. Available: https://www.top500.org/system/9156. [Accessed: 21-May-2018].

[57] ‘NEMO’. [Online]. Available: https://www.nemo-ocean.eu/. [Accessed: 21-May-2018].

[58] B. Arunachalam, B. Kalasagar, V. S. Arackal, and B. B. Prahlada Rao, ‘Open Source Job Submission Portal for Grid’, in *2013 National Conference on Parallel Computing Technologies (PARCOMPTECH)*, 2013, pp. 1–5.

[59] G. Galante, L. C. Erpen De Bona, A. R. Mury, B. Schulze, and R. da Rosa Righi, ‘An Analysis of Public Clouds Elasticity in the Execution of Scientific Applications: a Survey’, *J. Grid Comput.*, 2016.

[60] G. Mariani, A. Anghel, R. Jongerius, and G. Dittmann, ‘Predicting Cloud Performance for HPC Applications: A User-Oriented Approach’, in *Proceedings - 2017 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, CCGRID 2017*, 2017.

[61] G. F. and C. Cotelo, ‘Hyper-parametric search using HPC

infrastructures for Tensorflow', in *Machine Learning Workshop Galicia*, 2017.

[62] N. Sung *et al.*, 'NSML: A Machine Learning Platform That Enables You to Focus on Your Models', Dec. 2017.

[63] F. Nasoz and C. Shrestha, 'A web-based user interface for machine learning analysis', in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2017.

[64] '5 in 5 - IBM Research'. [Online]. Available: http://research.ibm.com/5-in-5/quantum-computing/. [Accessed: 13-Jun-2018].

8 RESEARCH AND INNOVATION PROJECTS

[P.1] **RETELAB** - *Un Laboratorio Virtual para la Red Nacional de Teledetección Oceanográfica*. Ministerio de Educación y Ciencia. ESP2006-13778-C04.
http://citius.usc.es/investigacion/publicacions/listado/retelab-geospatial-grid-web-laboratory-oceanographic.
October 2006 - September 2009.

[P.2] **i-MATH** - *Ingenio MATHEMATICA*. Proyecto Consolider Ingenio Mathematica. Ministerio de Educación y Ciencia. CSD2006-32.
http://mathematica.nodo.cesga.es.
October 2006 - April 2010.

[P.3] **MeteoSIX** - *Difusión de datos meteorológicos y oceanográficos.* Xunta de Galicia. 09MDS034522PR.
http://meteosix.cesga.es.
October 2009 - October 2012.

[P.4] **GIS-OCEANO** - *Xeración, procesado e distribución de campos oceanográficos operacionais multipropósito usando estándares abertos e servizos web*. Xunta de Galicia. 09MDS009CT.
http://oceano.cesga.es.
October 2009 - October 2012.

[P.5] BonFIRE-VCOC - *Bonfire-Virtual Clusters on Federated Cloud Sites*. European Commission - European Union's Seventh Framework Programme for research, technological development (FP7). Grant agreement No 257386 and No 287938.
http://www.bonfire-project.eu/innovation/virtual-clusters-on-federated-cloud-sites.
January 2011 - November 2012

[P.6] **Acceso y mejora de Instalación Científica y Técnica Singular (ICTS) Finis Terrae.** Ministerio de Ciencia e Innovación (MICINN). No ICTS-2009-40.
http://icts.cesga.es.
November 2009 - December 2011.

[P.7] **EGI-Inspire** - *Integrated Sustainable Pan-European Infrastructure for Researchers in Europe.* European Union's Seventh Framework Programme for research, technological development (FP7). Grant agreement No RI-26132.
https://www.egi.eu.
October 2012 - April 2014.

[P.8] **FORTISSIMO** - *Factories of the Future Resources, Technology, Infrastructure and Services for Simulation and Modelling.* European Commission - European Union's Seventh Framework Programme for research, technological development and demonstration (FP7). Grant agreement No 609029.
https://www.fortissimo-project.eu.
July 2013 - December 2016.

Within the Fortissimo project more than 50 experiments were conducted. Experiments were dedicated subprojects with 18 months duration executed within the Fortissimo scope. They were end-user relevant case studies demonstrating the use of Cloud-based HPC and the benefits it brings to the value chain from the end-user to the HPC infrastructure provider.

The researcher has participated in ten Fortissimo experiments:

[P8.1] Improved mechanical design of metal flanges.
https://www.fortissimo-project.eu/experiments/404

[P8.2] Cloud4Maps - Cloud-based HPC platform for maps conversion.
https://www.fortissimo-project.eu/experiments/501

[P8.3] CastINcloud - Sustainable cloud Services for bringing High Performance Casting Simulations to the SMEs.
https://www.fortissimo-project.eu/experiments/503

[P8.4] Improved and optimized design of high temperature exhaust gases concentric chimneys.
https://www.fortissimo-project.eu/experiments/505

[P8.5] Experimentation of VIrtual Metrology HPC Simulation Services for SME Production Process Control in Sustainable and Competitive CAmshaft Manufacturing.
https://www.fortissimo-project.eu/experiments/509

[P8.6] StamHPC - High Performance Computing for the metal stamping industry.
https://www.fortissimo-project.eu/experiments/516
[P8.7] Cloud-based Optimization platform for reinforcement steel cut industries.
https://www.fortissimo-project.eu/experiments/517
[P8.8] Multi-physics simulation of high-temperature superconducting devices.
https://www.fortissimo-project.eu/experiments/521
[P8.9] HPC-SHEAKS - HPC enabled system for enhanced seakeeping and station-keeping design.
https://www.fortissimo-project.eu/experiments/605
[P8.10] FLASH - Enabling fatigue life assessment in HPC-Cloud to SMEs.
https://www.fortissimo-project.eu/experiments/605

[P.9] **FORTISSIMO 2** - *Factories of the Future Resources, Technology, Infrastructure and Services for Simulation and Modelling.* European Commission - European Union's Horizon 2020 research and innovation programme - Grant agreement No 680481.
https://www.fortissimo-project.eu.
October 2015 - October 2018.

Within the Fortissimo 2 project more than 35 experiments were conducted. They are dedicated subprojects with 18 months duration executed within the project scope. Fortissimo 2 experiments are end-user relevant case studies demonstrating the use of Cloud-based HPC, Big Data and HPDA, and the benefits they bring to the value chain from the end-user to the HPC infrastructure provider.

The researcher has participated in five Fortissimo 2 experiments:

[P9.1] CyPLAM - Cyber-Physical Laser Metal Deposition
https://www.fortissimo-project.eu/experiments/707
[P9.2] Advance dimensional data analytics for knowledge generation in camshaft manufacturing.
https://www.fortissimo-project.eu/experiments/708
[P9.3] FREIGHPC - Enhancing Manufacturing cost efficiency through the utilization of HPC for the estimation of freight rates on maritime shipping markets.

https://www.fortissimo-project.eu/experiments/803
[P9.4] OTEAres - Improvement of the remote expert system based on software OTEA.
https://www.fortissimo-project.eu/experiments/904
[P9.5] Digester 2.0: Optimisation of the anaerobic digestion process for biogas generation.
https://www.fortissimo-project.eu/experiments/905