Making a Case for an Irish500 List

Brett A. Becker [1], John Regan [2], Michael Salter-Townshend [3], Kevin Casey [4]

[1] College of Computer Training, Dublin 2, Ireland
brett.becker@cct.ie

[2] Helsinki Institute of Information Technology, Aalto University, Helsinki, Finland
john.regan@hiit.fi

[3] Complex and Adaptive Systems Laboratory, University College Dublin, Dublin 4, Ireland
michael.salter-townshend@ucd.ie

[4] School of Computing, CloudCORE, Dublin City University, Dublin 9, Ireland
kcasey@computing.dcu.ie

Abstract

Ireland is uniquely positioned in the global High Performance Computing arena. Among the reasons for this are a small population coupled with diverse groups of world-class academic institutions, well-funded government bodies, and world-leading research groups. In addition, Ireland is a location of choice for the European/EMEA/Global headquarters of a large number of multinational computing and technology companies. Nonetheless, the Irish High Performance Computing (HPC) landscape is difficult to define domestically and even more difficult to place amongst the international community. This is largely due to the lack of a central repository of Irish HPC resources, cataloguing their capabilities, and application areas. A natural way to organise such resources is to rank them in terms of performance. This format also serves to command wider interest from outside the HPC community. As HPC resources have become a strategic asset, the ability to quickly identify them qualitatively and quantitatively is a powerful resource, particularly when this information can be searched, partitioned and tailored for specific uses. Projects such as the Top500 list have served this purpose for years, but only for the countries with the budgets and resources to place on such lists with a statistically significant number of systems.

This paper makes a case for an "Irish500" list. The distinguishing feature of this list is that all systems must be geographically located in Ireland. The two primary questions this project seeks to answer are: What is the landscape of HPC in Ireland today? Where does this landscape fit globally? The mission of the Irish500 list is the promotion, advocacy and advancement of HPC in Ireland. We persist with the convention of using 500 in the name for consistency, as have most other lists since the Top500 started; we are not aiming to create a list that actually comprises as many as 500 machines. The list ranks Irish HPC installations using the Linpack performance benchmark which is also used by the Top500 and other lists. In this paper we also make the case for the introduction of a flops/watt metric, used by the Green500 list, to rank systems in terms of computational energy efficiency. We discuss the motivation behind and benefits of the Irish500 list and explain its relation to existing lists. We then explore work based on, and interest in, other HPC lists. Finally, we describe the structure of the proposed list, including benchmarks and metrics used, along with a schedule of releases and anticipated participant sectors. We then conclude with a discussion on the implementation of the list and the irish500.org website.

Keywords: Top500, Green500, Irish500, Supercomputing, High Performance Computing

1 Introduction and Motivation

The Top500 Supercomputers (www.top500.org) has been maintaining a list of the most powerful computers in use for over 20 years. The Top500 has proven to be extremely important for several reasons. At a very high level, HPC resources have become a strategic asset. Being able to quickly identify them qualitatively and quantitatively is a powerful resource, particularly when this

information can be searched, partitioned and customised for particular uses. At a very basic level, it has been a driver for HPC procurements and provides funding attractiveness and enhanced recruitment power. These benefits apply not only to those at the top of the list but arguably for most sites that rank on the Top500. It has been shown that funding attached to Top500 ranking systems has a direct link to their performance, and conversely that high ranking attracts more funding and is also associated with an increase in the number of publications [1]. The list also has value in a public relations role, with journalists in many major newspapers reporting the latest list releases, drawing public attention and opinion to HPC in general. Again there is a link to funding in this regard, as politicians and policy-makers have taken an interest in the Top500 [2]. In practice, the Top500 has influenced the HPC industry to the extent that new supercomputer installation announcements are often clustered around the releases of list updates in June and November of each year [3]. The Top500 has even trickled down into general public awareness to the point that a recent *New York Times* article featuring a Top500 founder discussed building a Top500 machine using networked iPads [4]. HPC capabilities have even become a source of national pride [5].

In 2006 the second HPC ranking list, the Green500 (www.green500.org) was announced, ranking installations using the flops/watt metric [6]. The Green500 was a direct response to both the public interest in environmental friendliness, and more directly, to the rapidly growing power consumption problem in HPC. A third global list, the Graph500 was announced in 2009 (www.graph500.org). This was a response to criticism that the Linpack [7] benchmark which is utilised by the Top500 is not the most representative benchmark for predicting the performance of the actual applications that HPC installations execute. We discuss this further in Section 4.2. The Top, Green and Graph lists share one thing in common – their scope is limited to the top 500 performing systems *globally*. The only list known that ranks a different domain is Top Supercomputers India which ranks only those systems located in India [8].

1.1 Motivation for an Irish Supercomputer List

Ireland is uniquely positioned in the global HPC arena. Despite a small population (less than seven million, including Northern Ireland) there are large, diverse groups of world-class academic institutions, well-funded government bodies, and world-leading research groups. Moreover, Ireland is a favoured European/EMEA/Global headquarters location for several multinational computing and technology companies. Ireland has hosted eleven Top500 ranked installations, with ten in the past decade. Figure 1 charts these systems and their rank history. Table 1 shows installation owner/site, brief architecture details, and Top500 information for these systems.

This history is evidence of Ireland's ability to build and attract world-class HPC facilities. However, this hardly approaches a complete picture of HPC in Ireland, particularly at the present time. Many world-class systems in Ireland are not represented here simply because they fall short of the Top500 global ranking. In addition, five Irish Top500 systems have been decommissioned. It is also known that there are Irish companies and institutions that currently possess hardware that will rank highly domestically, but their performance data is not publicly available at present, or benchmarking of any kind has not yet been done. It is a motivation of the Irish500 list to compile a more representative and current picture of the Irish HPC landscape.

As the Irish HPC infrastructure is difficult to define domestically, it is even more difficult to place it amongst the international community. This is largely due to the lack of a central repository of Irish HPC resources and their capabilities. A natural way to organise such resources is to rank them in terms of performance or capability. Such repositories (such as the Top500) have served many purposes for years, but primarily for the countries with the budgets and resources required to place systems on them. The Top500 can yield useful, detailed and important information for countries such as Canada for example, which places seventh globally in terms of HPC resources. Such information is used for many purposes, including national strategy and budgetary decisions, funding applications, and even matters such as a sense of community amongst computer scientists. Currently Ireland does not have enough representation on the Top500 to serve these purposes, nor has it a domestic platform which

can do so. It is a chief motivation of the Irish500 to provide such a platform for Ireland. First and foremost, the Irish500 seeks to answer:

- *What is the landscape of HPC in Ireland today?*
- *Where does this landscape fit globally?*

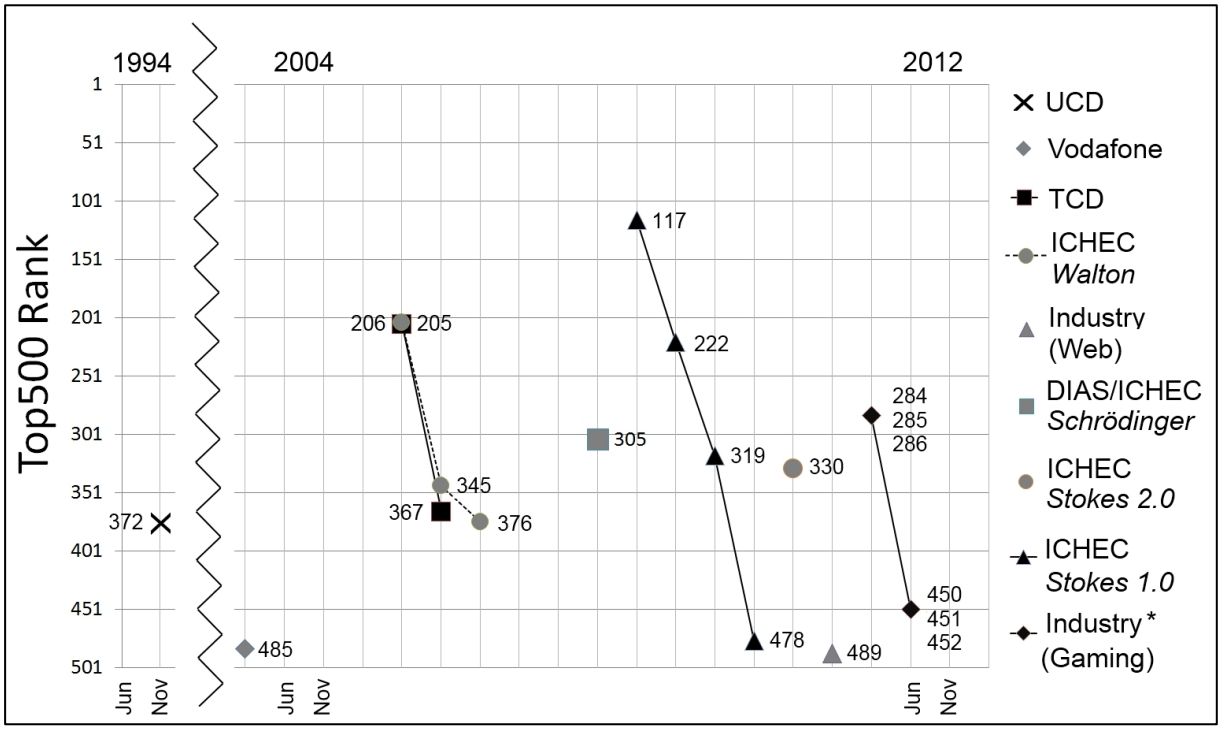

Figure 1: Ireland's history of Top500-class HPC installations. *(three identical clusters)

Installation	Max Top500 Rank	Top500 Dates	Hardware Specification
ICHEC	117	11/08 – 6/10	SGI Altix ICE 8200EX, Xeon 4C 2.8G
ICHEC	205	6/05 – 6/07	IBM eServer Opteron 2.4G, 1G Eth
TCD	206	6/05 – 11/05	IBM eServer Opteron 2.4G, Infiniband
Industry (Gaming)*†	284/285/286	11/11 – 6/12	HP Cluster Platform 3000 BL460c G7, Xeon X5660 6C 2.80G, 10G Eth
DIAS/ICHEC	305	6/08	IBM BlueGene/P
ICHEC	330	11/10	SGI Altix ICE 8200EX, Xeon 6C 2.66G
UCD	372	11/94	MasPar MP-2216
Vodafone	485	6/03	HP SuperDome 875 MHz, HyperPlex
Industry (Web)†	489	6/11	HP DL160 G6, Xeon X5650 2.66G, 1G Eth

Table 1: Site and architecture information for Ireland's Top500-class HPC installations in increasing order of highest Top500 rank achieved.
***(three identical clusters)** **†Anonymous submissions to Top500 list**

The only country known to have a dedicated domestic list is India [8]. Despite having only eight systems on the latest Top500 list, the December 2012 Top Supercomputers India list has 26 operational systems, giving a robust picture of the Indian HPC landscape. Given the relatively good finances of the Irish research sector, and abundance of multinational IT companies, the Irish list could easily gain as many systems. Additionally, a relatively small population is seen as an advantage – in particular it brings easily identifiable and contactable HPC providers/users in addition to making comprehensiveness a realistic possibility. If each university, institute of technology, and a handful of multinationals were to enter just one system each, the number of installations on the Irish500 could exceed 30. If submissions are made by individual faculties, schools, research groups and smaller

companies (not to mention multiple submissions by each) the number of systems becomes substantial enough to be valuable for substantial analysis, including market trends and hardware/interconnect usage, as well as application information. We explore this further in Section 4.3 and present a feasibility study demonstrating that the Irish500 list could realistically equal or surpass the Indian list, both in system numbers and performance in the short-term. This would be a positive result for the Irish HPC community, and a major step towards addressing the questions posed above.

Another motivating factor in maintaining an Irish list is to promote accuracy and representation. In November 2011, Amazon Web Services benchmarked a system in the U.S. made of 1062 EC2 cc2.8xl (Cluster Compute Eight Extra Large) instances. The result was 240 Tflops, placing it at #42 on the June 2011 Top500 list, and currently at #102 [9]. This cluster contained 17,024 cores with 66 TB of RAM. In 2012 Amazon announced that these same instances were available in their Irish (EU-West) region [10]. The combined capability of these instances in the Irish datacentre is not currently known, but unless the number of machines is significantly less than that in the U.S., the Irish Amazon region could also place in the Top500 list and very possibly top the Irish500.

Additionally, the Irish Amazon region has been used in two global projects conducted by the software firm Cycle Computing. The first, *Nekomata*, is a 30,000 core homogeneous effort consisting of machines in three Amazon regions. The exact Irish contribution is known, therefore allowing a very reliable estimate of the performance of the Irish machines involved [11]. The second, *Naga* is a 50,000 core, seven region effort. Unfortunately the system is made of highly heterogeneous instances and although the number of Irish machines is known, the breakdown of instance types is unknown and therefore a reasonable performance share estimate is not feasible [12].

It is widely believed, based on experiments that probe EC2 and calculate numbers of racks by various means, that the Amazon EU-West (Ireland) region is the second largest of the seven EC2 regions, second only to US-East (Virginia) [13]. The fact that these resources exist in Ireland, and have paying customers conducting hard science [11] [12] [14] is encouraging, yet the true capabilities of these resources remain unknown. It is a motivation of the Irish500 list to identify such capabilities, while respecting any installation owner's right to anonymity and/or privacy of certain specifications when required.

Finally, the Irish500 list will help to fill the void left with the closure of Grid Ireland at the end of 2012 [15]. Amongst many hardware based services, Grid Ireland also functioned as a central networking point for the Irish HPC community, something that the authors believe needs to be fostered to maintain that sense of community and forward national momentum. It is believed that the Irish500 can help serve as one such central networking body.

2. Existing Lists and Projects

There are four related lists now in operation: The Top500, Green500, Graph500 and Top Supercomputers India. In addition, the HPC Challenge Benchmark serves a similar purpose to the traditional list format, as we discuss below. It is included here, as together these are the only known implementations of HPC rankings.

2.1 Top500

The Top500 list ranks the top 500 supercomputers currently in operation. It was established in 1993 and is updated twice per year – in June coinciding with the International Supercomputing Conference (always held in Europe) and in November coinciding with the ACM/IEEE Supercomputing Conference (always held in the US). The Top500 continues to grow in prominence, attracting widespread media and public attention, not to mention serious attention from researchers, universities, politicians and funding bodies.

The Top500 uses a single benchmark, Linpack, most often in the form of HPL, a portable implementation of the high-performance Linpack benchmark for distributed-memory computers [7]. This benchmark calculates the floating point operations per second (flops) achieved during the solution of a dense system of linear equations. The actual metric is Rmax in Tflops, where Rmax is the maximum sustained performance during the benchmark run. This is related to Rpeak, the theoretical maximum performance, though the relationship varies from installation to installation. Summing all 500 Rmax and Rpeak values in the latest Top500 list reveals an Rpeak/Rmax ratio $\approx \sqrt{2}/2$. This is useful to estimate the Rmax value of a machine from the Rpeak, in cases where a machine has not yet done any benchmarking but has a known theoretical performance. Admittedly Linpack has drawn some criticism of late, for not being conceptually representative of real-world applications executed by HPC systems [16]. Nonetheless, it is simple, relatively easy and cheap to run, is very scalable, and has no direct replacement as of yet [17]. Most importantly, the Top500 has grown in influence year-on-year since inception, and continues to use Linpack, with no sign of this changing in the near future [2].

2.2 Graph500

The Graph500 list was started in 2010 as a response to the criticism of Linpack [16]. It utilises a suite of data-intensive benchmarks, using a metric of TEPS (Traversed Edges Per Second) which is designed to emphasise the importance of the communication network as well as the number-crunching capabilities of an installation. As of November 2012, there are 124 systems on the Graph500.

2.3 Green500

As the power consumption of supercomputers has risen well into the megawatts, energy efficiency has come to the forefront of HPC. The Green500 list ranks the current Top500 installations not in terms of flops but flops/watt [6]. The list is generated with Top500 data and always has 500 systems listed – a list equivalent to the Top500 except utilizing a different ranking metric.

2.4 Top Supercomputers India

Top Supercomputers India (TSI) was started in November 2008, but now lags the release schedule of the Top, Graph and Green lists by one month. This is most likely due to the possible overlap of Top500 and TSI installations (which is of course allowed, but needs to be consistent). It is the closest analogue to the proposed Irish500 list in that it is limited to installations in a specific geographical area, and performance need not be amongst the top 500 globally. The benchmark used, as for the Top500, is Linpack. The November 2012 TSI list has 27 ranked installations. One of the likely differences between the proposed Irish500 list and the TSI list is that all TSI installations, barring one from commercial giant Tata are owned or administered by national/governmental bodies or research/education/university institutions. Such a lack of industry presence would not be expected to characterise the Irish list, as almost half of Irish Top500 systems have been in the industry sector, including the fastest Irish installation ever. Section 4.3 shows that the proposed Irish500 list and the TSI are similar in number of systems and performance.

2.5 The HPC Challenge Benchmark

The High Performance Computing Challenge (HPCC) is a departure from the traditional "500-style" list [18]. It is a series of awards aligned to a suite of seven benchmarks with a similar aim to those of the Graph500 – to test a system's performance with benchmarks that are representative of real-world HPC applications. It is worthwhile to note that Linpack (HPL) is one of the seven benchmarks in the HPCC suite, and that some of the others are closely related to it. Due to the nature of the benchmarks and the small number of awards, the HPCC looks poised to always be in the domain of the supercomputing super-elite, perhaps the global top 50 or so.

3 Related and Resulting Work

The existing lists have been the topic of, and source of data for, research and debate in the literature. Being updated twice per year, a substantial amount of data is compiled including where the installation is located, the installation performance history, architecture, manufacturer, processor family, age, benchmark performance, core/node/processor count, interconnect, power consumption, OS, accelerator information and possible application areas. This makes for a very good dataset upon which to identify current trends and predict future ones. Some of the latest data analysis from the Top500 was featured in the June 2012 cover story of Scientific Computing [19]. One of these insights revealed that the research and enterprise system technology gap is growing. Without such hard data, this would be impossible to demonstrate quantitatively.

In a 2009 study, the question "How high can a cloud computing service get in the Top500 list?" was addressed [17]. The conclusion was that cloud computing platforms (specifically Amazon EC2) were not yet mature enough for HPC applications. Nonetheless, this paper did give an indication that, at some point in the near future, cloud platforms would be viable for HPC applications. The authors' assertions were validated just two years later when Amazon had a cloud platform place at #42 on the Top500 [9]. In addition the paper demonstrated the viability of the economics involved in running the Linpack benchmark on pay-on-demand cloud platforms, with single runs of Linpack costing only tens of dollars. In this case the Top500 list was the perfect venue to see where a platform new to HPC/Supercomputing fitted, and to make predictions about the future of the platform in HPC. Similar stories are unfolding in the GPU/Accelerator arena, as the reality of systems running more than one million cores has arrived [20].

Other studies have sought to extract more qualitative data from the current lists including identifying and predicting trends in HPC. In [5] the authors identify invariant trends in Top500 data, predicting growth rates, and identifying limiting factors. On the Green500 front, recent forecasts indicate that the "performance at any cost" paradigm is no longer sustainable as HPC moves towards the exascale [21]. Such studies are valuable to the HPC community and in some cases have already served as indicators of things to come.

Top Supercomputers India has also spawned some work of its own, in the form of a project that collects additional information about the machines on the Indian list including numbers of jobs per month, how many cores jobs use, application categories, fault-tolerance, and other performance information. Currently, the preliminary outputs of this project are pending [8].

4 The Irish500 List

An advisory committee has been formed for the Irish500 list, including personnel at Irish and foreign bodies (for international and unbiased perspective). As efforts gain pace, the size and diversity of this group will be encouraged to grow. Having laid out the motivation and justification for an Irish500 list, we will now present initial work on the Irish500 project.

4.1 Mission Statements

The founding mission statements of the Irish500 follow, and will be regularly reviewed:

1. To form a central point for HPC installation operators and users in Ireland across all sectors, public and private, including (but not limited to) academic, research, industry and government.
2. To improve the awareness of HPC amongst the Irish public.
3. To represent and enhance the profile of the HPC landscape domestically and internationally.
4. To serve as a globally-facing resource representing Irish HPC, in keeping with both global and national conventions.

5. To identify current and emerging trends in Irish HPC, which can then be compared to those identified in global lists.
6. To provide a list whose statistics are representative of the Irish HPC landscape, for use by anyone for any purpose.
7. To maintain a list whose statistics are not dominated and skewed by a small number of systems at the top end – systems that also often feature expensive, exotic or custom hardware not available to most users of HPC systems.
8. To allow any university department, vendor, research group or company to rank on a list of peers, to quantitatively determine their rank, and qualitatively identify and analyse the properties and trends of similar systems, by similar organisations, with similar goals and means.
9. To provide a lifetime-long, not time-on-list-long history of system performance. This will be achieved by tracking system performance until decommissioning, not only until failure to attain a specified rank or performance mark.
10. To identify hardware and cooling trends that lead to cost effective performance and maximum energy-efficiency.

4.2 Choice of Metrics

The Irish500 will ultimately feature two separate rankings. The first and higher priority is based on computational performance. As the Irish HPC community becomes more involved, the second, ranking energy efficiency will be introduced. The energy efficiency metric will be flops/watt, measuring the average power consumption of a system while running at full capacity, identical to the Green500. The performance metric used by the Irish500 list will be that of the Top500 list which calculates system floating point operations per second (flops), using Linpack. The principal reasons for this are that it is cheap and simple to run, produces a single figure, is very scalable and has longstanding historical precedence. On a more technical level, we also choose Linpack because it gives a good indication of performance, despite some recent criticism. Most of this criticism has been directed at Linpack not stressing enough system components. Indeed, it was in response to this criticism that Graph500 was created. Although it is not the purpose of this paper to support or criticise either Linpack or proposed alternatives, a brief case to justify the use of Linpack as the Irish500 computational performance metric beyond the recent criticism is presented. This is done by demonstrating that good Linpack performance can be shown to correlate with good Graph500 performance (and vice-versa), within a very acceptable margin.

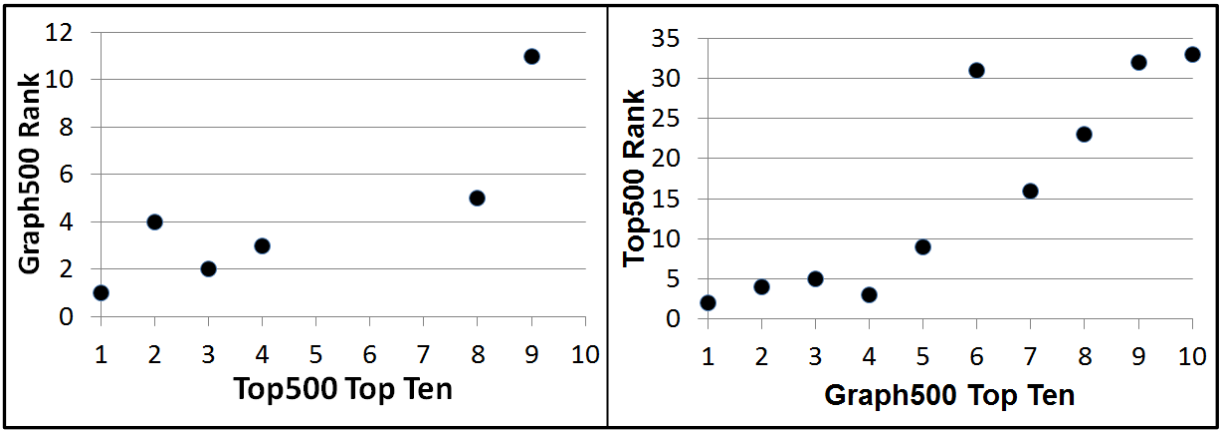

Figure 2: The top ten machines on the Top500 and Graph500 lists and their corresponding ranks on the other list (November 2012).

Figure 2 shows the top ten installations on the November 2012 Top500 and Graph500 lists, and their corresponding rank on the other list. All six machines in the Top500 top ten which entered the Graph500 benchmarks (not all did) are in the top eleven on the Graph500. Similarly, the top five

Graph500 machines are in the top ten on the Top500, and all top ten Graph500 are in the top 33 of the Top500. It is important to point out that this is not necessarily a like-for-like comparison, as machines may have run each benchmark under different configurations. Either way, it is interesting and important to see that good Top500 performance can be an indicator of good Graph500 performance, and vice-versa. This supports the decision making the cheaper, faster, simpler, and historically proven Linpack the benchmark for the Irish500. Finally, the Top500 website states "It is very unlikely that another benchmark will replace Linpack as basis for the Top500 lists in the near future" [2].

4.3 Moving Forward - Feasibility

A pre-release feasibility list as of February 2012 can begin to take shape now, based on information in the current and past Top500 lists and other publicly available information [22] [23] [11] [24]. It should be noted that a goal of the Irish500 is to develop direct contact with all installations that feature on future lists, therefore not relying on public information, but verifiable, direct-from-source information. The feasibility list is shown in Table 2, and was compiled with the sole purpose of determining the feasibility of the Irish500 project, having stated the motivation and justifications behind it.

Rank	Top500 Rank	Rmax (Tflops)	Installation	Hardware Specification
1 (**1**)	285*	64.9	Industry (Gaming)	HP Cluster Platform 3000 BL460c G7, Xeon X5660 6C 2.80 GHz, 10G Eth
1 (**2**)	286*	64.9	Industry (Gaming)	HP Cluster Platform 3000 BL460c G7, Xeon X5660 6C 2.80 GHz, 10G Eth
1 (**3**)	287*	64.9	Industry (Gaming)	HP Cluster Platform 3000 BL460c G7, Xeon X5660 6C 2.80 GHz, 10G Eth
2 (**4**)	489*	40.5	Industry (Web)	HP DL160 G6, Xeon X5650 2.66 GHz, 1G Eth
3 (**5**)	330* 117**	36.6	ICHEC – Stokes	SGI Altix ICE 8200EX, Xeon X5650 2.67GHz, ConnectX Infiniband DDR
4 (**6**)		31.9†	Cycle Computing – Nekomata	500 Amazon EC2 c1.8xl instances, 10G Eth
5 (**7**)		8.9	TCD – Lonsdale	Opteron 2.3GHz, 1232 Core, Infiniband DDR
6 (**8**)		8.9†	TCD – Kelvin	Intel 2.66Ghz, 1200 core, Qlogic Infiniband QDR
7 (**9**)		8.3†	TCD – Parsons	Intel 2.5Ghz, 1104 core, Voltaire Infiniband (QDR switching and Connectx DDR hosts)
8 (**10**)		5.1††	ICHEC – Stoney	Bull Novascale R422-E2, Intel Xeon X5560 2.8GHz, ConnectX Infiniband (DDR)
9 (**11**)		1.0†	TCD – Crusher	Intel2.5Ghz, 144 cores, Infiniband
Total		**280.7**		

Table 2: The Irish500 pre-release (feasibility) list as it stands in February 2013, using only publicly available information, all systems currently in commission. *Highest historical rank **Highest historical rank (Pre-upgrade) †Estimated Linpack Rmax ††CPU only (no GPUs)

With 11 systems overall (three identical), the performance ranges from 1 to 64.9 Tflops, and includes six past Top500 machines. The sectors span academia, industry, research, and government. With such a spread of performance and sectors already identified, a combination of a call for participation, public awareness, and the advisory committee contacting identified bodies, this list will surely grow. Only one university (TCD) features on this list, because they are the only one with publicly available figures. There is little doubt that UCD, UCC, UCG, NUIM, DCU, QUB and other academic institutions could significantly contribute to this list. In fact, it is known that most of these institutions have hardware in place that *will* rank highly in this list. Industry has a strong presence in the feasibility list with five installations. Certainly Dell, Intel, IBM, Google, Microsoft and other multinationals present in Ireland can help the list grow further.

There are four decommissioned machines that would also rank on the feasibility list in terms of performance. It is envisaged that the Irish500 will also serve as a repository for systems as they are taken out of commission. This is in contrast to the other lists that stop tracking decommissioned machines as soon as they fail to rank. The Irish decommissioned machines (performance in Tflops) are: ICHEC [Schrodinger (11.1), Lanczos (4.7), Walton (3.14)] and TCD [IITAC (2.7)]. As the number of decommissioned machines grows, their data, combined with machines still in use will together help create a broader, more historical picture of the Irish HPC landscape as it evolves.

Finally, it is encouraging looking at the Top500 feasibility list coupled with the fact that the threshold of Top500-class is ~75Tflops, and that many machines on the Irish list above are very close to this. It is also interesting to see that even though this list which will certainly be added to, the top eleven Irish installations sum to a healthy 280 Tflops. Additionally, comparing this list to the current TSI list is encouraging. The Irish feasibility list has 11 systems compared to 27 for TSI and both the highest and lowest ranking systems on each list are comparable in performance (well within small multiples).

4.4 Irish500.org

The Irish500 website (www.irish500.org) will primarily function as a forum to present results to the public and participants but will also accept submissions. All submissions will be followed up by personal contact from the committee. The information collected from participants will include, but is not limited to: location, segment, system model, manufacturer, Rmax, power consumption, total cores, accelerator/GPU (cores/type/family), processor (generation/speed), RAM (per core/total), total system storage, OS, OS family, interconnect (speed/family), application area(s), total purchase cost, running costs and average utilisation.

4.5 Call for Participation

The Irish500 Advisory Committee has issued a public call for participation, located on the website, targeted at Irish HPC installation owners/organisers from all sectors. The Advisory Committee is also actively working on engaging these sectors to discover and confirm existing HPC installations. Later this year, the first Irish500 list will be released. Participation will occur both through the Irish500.org website and personal contact by the committee. The committee is also keen to expand by taking on new members, particularly from other institutions and private industry. We envisage a twice-annual release schedule similar to Top500, in June and November of each year. Within one year of the initial (performance-based) list, we envisage the first release of energy efficiency rankings.

4.6 Privacy and Anonymity

The Irish500 is committed to improving the public awareness of HPC domestically and internationally, and therefore all submissions of Irish-based installations, regardless of sector, are welcome. Due to the small size of the Irish market, the committee does recognise the possibility that, particularly in private industry, some companies may be reluctant to release certain information to the public for competition reasons. The Irish500 therefore provides for anonymous submissions, where only performance data and limited (agreed) hardware specifications are released to the public. The committee will need to verify other details of the installation. However these will not be released at the request of the installation owner.

5 Conclusion

This paper presents a case for an Irish500 list, ranking the top HPC installations geographically situated in Ireland by performance, and in the near future energy efficiency. The Irish500 project is motivated by two considerations. Firstly, it seeks to be a vehicle for the goals of advancing, advocating and promoting Irish HPC domestically and internationally. Secondly, Ireland has a unique

position in the global HPC arena and there is a lack of a central and independent body to help achieve these goals. In particular the Irish500 seeks to answer the questions:

- *What is the landscape of HPC in Ireland today?*
- *Where does this landscape fit globally?*

The proliferation of the Top500 list format and the importance and success of other lists, justify adopting a similar format for the Irish500 list. Evidence that such lists have significantly impacted stakeholders in HPC, from practitioners to the general public further strengthens the case for the Irish500. A feasibility list is presented, *made up entirely of publicly-available information* describing eleven systems with a performance profile similar to the Top Supercomputers India list, and containing systems of global importance. Some of these systems are formerly ranked on the Top500 list, and every major HPC sector in Ireland is represented. This demonstrates that the Irish500 list could quickly be as successful as the Top Supercomputers India list (the only other known geographically-linked list), both in system numbers and performance. This would be a positive development for the Irish HPC community, a major step towards addressing the two questions posed above, and a significant factor in moving the Irish HPC community forward. We have formed an advisory committee and developed a prototype website. Following further expansion of the committee and direct contact between the committee and identified Irish HPC owners and maintainers, we will issue a general call for participation. After this, the first Irish500 list will be announced. It is intended that the Irish500 will be continuously open to submission, constantly updated and refined, and released on a bi-annual basis.

Acknowledgements

The authors would like to thank Catherine Mooney, Kevin Hely, Liam Power and Justin Smyth for useful assistance and advice.

References

[1] A. Apon, S. Ahalt, V. Dantuluri, C. Gurdgiev, M. Limayem, L. Ngo and M. Stealey, "High Performance Computing Instrumentation and Research Productivity in U.S.," *Journal of Information Technology Impact,* vol. 10, no. 2, pp. 87-98, 2010.

[2] "Top500 Supercomputer Sites - Project," [Online]. Available: http://www.top500.org/project/. [Accessed 20 February 2013].

[3] R. Smith, "Cineca's Tesla K20-Based "Eurora" Supercomputer Unveiled; Water Cooling Unlocks Extra Efficiency," 31 January 2013. [Online]. Available: http://www.anandtech.com/show/6717/cinecas-tesla-k20based-eurora-supercomputer-unveiled-water-cooling-unlocks-extra-efficiency. [Accessed 17 February 2013].

[4] J. Markoff, "The iPad in Your Hand: As Fast as a Supercomputer of Yore," *New York Times,* 9 May 2011.

[5] D. G. Feitelson, "The supercomputer industry in light of the Top500 data," *Computing in science & engineering ,* vol. 7, no. 1, pp. 42-47, 2005.

[6] S. Sharma, C.-H. Hsu and W.-c. Feng, "Making a case for a green500 list," in *Proceedings of the 20th International IEEE Parallel and Distributed Processing Symposium*, 2006.

[7] A. Petitet, R. C. Whaley, J. Dongarra and A. Cleary, "HPL - A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers, Version 2.0," Innovative Computing Laboratory, University of Tennessee, 10 September 2008. [Online]. Available: http://www.netlib.org/benchmark/hpl/. [Accessed 17 February 2013].

[8] S. Vadhiyar, "Top Supercomputers India," Supercomputer Education and Research Centre, Indian Institute of Science, 27 December 2012. [Online]. Available: http://topsupercomputers-india.iisc.ernet.in/. [Accessed 17 February 2013].

[9] J. Barr, "Next Generation Cluster Computing on Amazon EC2 - The CC2 Instance Type,"

Amazon Web Services, 14 November 2011. [Online]. Available: http://aws.typepad.com/aws/2011/11/next-generation-cluster-computing-on-amazon-ec2-the-cc2-instance-type.html. [Accessed 25 February 2013].

[10] J. Barr, "High Performance Computing Heads East - EC2 CC2.8XL Instances in EU West (Ireland)," Amazon Web Services, 18 June 2012. [Online]. Available: http://aws.typepad.com/aws/2012/06/high-performance-computing-heads-east-.html. [Accessed 17 February 2013].

[11] "Cycle Computing Blog," Cycle Computing LLC, 29 September 2011. [Online]. Available: http://blog.cyclecomputing.com/2011/09/new-cyclecloud-cluster-is-a-triple-threat-30000-cores-massive-spot-instances-grill-chef-monitoring-g.html. [Accessed 17 February 2013].

[12] "Cycle Computing Blog," Cycle Computing LLC, 19 April 2012. [Online]. Available: http://blog.cyclecomputing.com/2012/04/cyclecloud-50000-core-utility-supercomputing.html. [Accessed 17 February 2013].

[13] H. Liu, "Amazon Data Center Size," 13 March 2012. [Online]. Available: http://huanliu.wordpress.com/2012/03/13/amazon-data-center-size/. [Accessed 17 February 2013].

[14] T. Trader, "Utility Supercomputing Heats Up," Tabor Communications, 28 February 2013. [Online]. Available: http://www.hpcinthecloud.com/hpccloud/2013-02-28/utility_supercomputing_heats_up.html. [Accessed 2 March 2013].

[15] "Grid-Ireland Closure on 31 December 2012," Grid-Ireland, 24 October 2012. [Online]. Available: http://www.grid.ie/closure.html. [Accessed 17 February 2013].

[16] R. C. Murphy, K. B. Wheeler, B. W. Barrett and J. Ang, "Introducting the Graph 500," Cray User's Group (CUG), 2010.

[17] J. Napper and P. Bientinesi, "Can cloud computing reach the top500?," in *Proceedings of the ACM combined workshops on UnConventional high performance computing workshop plus memory access workshop* , 2009.

[18] "HPC Challenge," University of Tennessee Innovative Computing Laboratory, 17 February 2013. [Online]. Available: http://icl.cs.utk.edu/hpcc/. [Accessed 17 February 2013].

[19] E. Strohmaier, "Insights from the Top500," *Scientific Computing*, pp. 14-16, August 2012.

[20] T. Trader, "Stanford Lights Up One Million Sequoia Cores," Tabor Communications, 28 January 2013. [Online]. Available: http://www.hpcwire.com/hpcwire/2013-01-28/stanford_lights_up_one_million_sequoia_cores.html. [Accessed 17 February 2013].

[21] B. Subramaniam and W. C. Feng, "Understanding power measurement implications in the green500 list," in *Green Computing and Communications (GreenCom), IEEE/ACM Int'l Conference on & Int'l Conference on Cyber, Physical and Social Computing (CPSCom)* , 2010.

[22] E. Strohmaier, S. Horst and J. Dongarra, "Top 500 Supercomputer Sites, November 2012," Top500.org, 12 November 2012. [Online]. Available: http://www.top500.org/lists/2012/11/. [Accessed 17 February 2013].

[23] "Trinity Centre for High Performance Computing," Trinity College Dublin, 15 February 2013. [Online]. Available: http://www.tchpc.tcd.ie/resources/clusters. [Accessed 17 February 2013].

[24] "High-Performance Computing (HPC) Infrastructure," Irish Centre for High-End Computing, January 2012. [Online]. Available: http://www.ichec.ie/infrastructure/. [Accessed 17 February 2013].

[25] G. Wang and T. E. Ng, "The Impact of Virtualization on Network Performance of Amazon EC2 Data Center," in *INFOCOM, 2010 Proceedings IEEE*, 2010.