Do regions of ALICE matter? Social relationships and data exchanges in the Grid

E D Widmer[1], F Carminati[2], C Grigoras[2], G Viry[1] and G Galli Carminati[3]

[1] Department of Sociology, University of Geneva, Switzerland
[2] CERN, Geneva, Switzerland
[3] Units of Mental Development Psychiatry, Department of Mental Health and Psychiatry, University Hospitals of Geneva, Switzerland

E-mail: Federico.Carminati@cern.ch

Abstract. Following a previous publication [1], this study aims at investigating the impact of regional affiliations of centres on the organisation of collaboration within the Distributed Computing ALICE infrastructure, based on social networks methods. A self-administered questionnaire was sent to all centre managers about support, email interactions and wished collaborations in the infrastructure. Several additional measures, stemming from technical observations were produced, such as bandwidth, data transfers and Internet Round Trip Time (RTT) were also included. Information for 50 centres were considered (60% response rate). Empirical analysis shows that despite the centralisation on CERN, the network is highly organised by regions. The results are discussed in the light of policy and efficiency issues.

1. Introduction

This study aims at measuring the impact of regions on the structure of interactions between the centres of the Distributed Computing ALICE infrastructure, based on social network methods [2]. These centres are part of the Worldwide Large hadron collider Computing Grid (WLCG [3]) They form a large computational network where data and workload are exchanged, but also a large and complex social network with $\sim 3,000$ possible links. The operation experience of this Grid since ten years has shown that the coordination and collaborations between the different centres are as important as the material conditions to ensure the proper functioning of this complex system extended over different time-zones and continents.

In a previous article [1], we found that the centres of the ALICE Distributed Computing Infrastructure derive most of their support from CERN and that various types of interactions were centralised on CERN. However, we also found that there were signs of local organisation of data exchange and collaborations. In this paper, we estimate the impact of regional influences. The research issue considered is the extent to which ALICE is organised with reference to regional anchorage that may be related to cultural, historical, political or network connectivity issues.

2. Data

Information about collaborations were collected using a self-administered questionnaire which was filled by the technical manager of each ALICE centre in the Fall 2009 and early 2010.

Technical managers had to estimate by yes/no questions which centres of the Grid provided their centre with significant help in its work at least once a week (support). They also had to estimate which centres were in regular e-mail contact at least once a week with their own (interactions), and with which centres their centre would like to have more interaction in its work (wished collaborations). These indicators refer to self-reported exchanges. The questionnaire was sent via e-mail to the 73 centres of the ALICE Grid. Answers for 50 centres were received and considered in the empirical analysis (68% response rate).

Additional information were gathered by technical observation such as the theoretical capacity of the network linking the centres (bandwidth), the actual quantity of data exchanged and the Internet Round Trip Time (RTT). These indicators refer to observed exchanges, in contrast with self-reported exchanges provided by responses of technical managers in the self-administered on line questionnaire. Overall, this study focuses on support, interactions, wished collaborations, theoretical capacity of the links between any two centres (bandwidth), quantity of data exchanged, and Internet Round Trip Time (RTT). We investigate to what measure those interconnections are structured by regions.

Centres are dispersed in various regions. Overall, the sample includes 23 centres in Northern Europe, 9 in Southern Europe, 6 in Asia and 8 in Russia. In order to estimate the impact of regions, we could not include Africa (one centre), South-America (one centre) and north America (two centres) due to the very limited number of centres pertaining to each of these regions.

3. Measurements

A variety of measures exist in network analysis ([2, 4, 5]). We focus on density as a measure of cohesion, at the level of the full network and at the level of regions. Density in a directed graph is equal to the number of existing arcs (directed ties) divided by the total number of possible arcs. In addition, two measures of network centralisation were computed, which capture different conceptual dimensions of centrality. In-degree and out-degree centralisation express the degree of variability (or inequality) among centres in the number of ties pointing to and going from a specific centre. For instance, a network characterised by a small number of centres receiving many direct ties, and other centres receiving few ties, has a strong in-degree centralisation. These measures provide information on the local dimension of centralisation as a centre may have many connections within a rather isolated subgroup of centres. Quite distinctly, betweenness centralisation measures the degree of variability (or inequality) among centres in the proportion of interactions in the network captured by a centre. The network is said to be centralised if a small number of centres lie between all other centres´ chains of relationships. Finally, the number of cliques measures the extent to which the network is structured around multiple clusters of centres or, conversely, is composed of a small number of big clusters. Formally, a clique is a sub-set of a network including the maximum number of centres that have all possible ties present among themselves. For the present analyses, a clique includes at least four centres.

Using the software UCINET [6], we compute these parameters on the overall network, with and without CERN in it. This procedure is set up to control for the impact of CERN on the overall structure, as the centrality of CERN in the network was acknowledged [1] as it was to be expected since CERN is the largest laboratory, both source of the experimental data and of most of the software used on the WLCG Grid. Because CERN is by its institutional role lead to be central, it is necessary to estimate regional influences with and without it included in the network. Threshold values for observed exchanges were fixed to determine only the substantial interactions between the centres. Regarding bandwidth and the actual quantity of data exchanged, two centres are considered to be linked when the theoretical capacity and the data transfer from one centre to the other amount at least 100 and 0.1 megabyte per second respectively. Finally, RTT links were considered when the Internet Round Trip Time from one centre to the other was below 20 milliseconds. Table 1 presents various indices measuring the

cohesion and the centralisation of the network. Indices are computed either with CERN included or CERN excluded, in order to estimate the impact of CERN on the overall organisation of the Network.

Table 1. Network indices (with and without CERN).

	Help	E-mail contacts	Wished collaborations	Bandwidth	Data transfers	RTT Capacity
With CERN						
Density	4.3	4.9	3.6	16.2	8.4	18.2
Betweenness centralisation	36.2	31.3	26.3	6.2	8.0	21.0
Indegree centralisation (provided)	72.6	59.5	42.1	34.0	12.9	39.9
Outdegree centralisation (providing)	33.1	17.9	13.0	29.9	31.6	71.1
Number of cliques (min. size 3, sym. max)	24	19	23	38	24	67
Without CERN						
Density	2.2	3.4	2.5	14.9	7.1	16.4
Betweenness centralisation	0.5	3.3	2.4	6.9	6.9	22.8
Indegree centralisation (provided)	12.7	11.5	6.0	33.9	11.8	34.0
Outdegree centralisation (providing)	10.5	9.3	12.4	29.6	26.7	72.3
Number of cliques (min. size 3, sym. max)	10	9	6	38	15	65

Results from the overall network, with CERN included, show a highly centralised network and a low level of density for all interactions. When CERN is excluded from the network, the centralisation becomes much weaker. The wished collaborations have somewhat a lower level of centralisation, thus showing that centre managers value to some extent the development of less centralised interactions within ALICE. For observed exchanges, there is also less centralisation. Getting CERN out of the network does not make such a great difference as in the case of self-reported exchanges.

4. Results

Results from the visual inspection of various graphs confirm that interactions, both self-reported and observed, depend to a large extent on regional influences. Figure 4 shows the interaction between regions and help provided. Figure 4 report responses given to the question about contacts by emails. The different regions are illustrated with the following colours: centres from Northern Europe (yellow, n=23), Southern Europe (green, n=9), Asia (white, n=6), Russia (orange, n=8), Africa (pink, n=1), South-America (blue, n=1) and north America (red, n=2). Both networks show a high level of centralisation on CERN as well as an organisation by regions. Figure 4 refer to the wishes of centre managers for collaborations. In that case, the network is much less centralised on CERN and the structuring by regions is less straightforward compared with current social interactions.

Are observed interactions among centres also significantly associated with their regional affiliations? As figure 4 shows, the theoretical capacity of the network is higher among the centres of the same region than among centres of distinct regions. Also, there are more exchanges

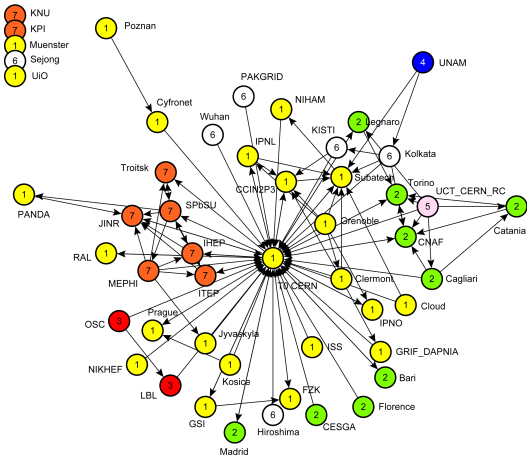

Figure 1. Help provided by regions.

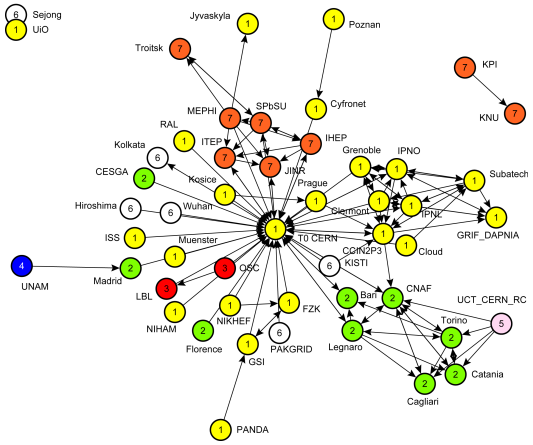

Figure 2. Exchanges of emails by regions.

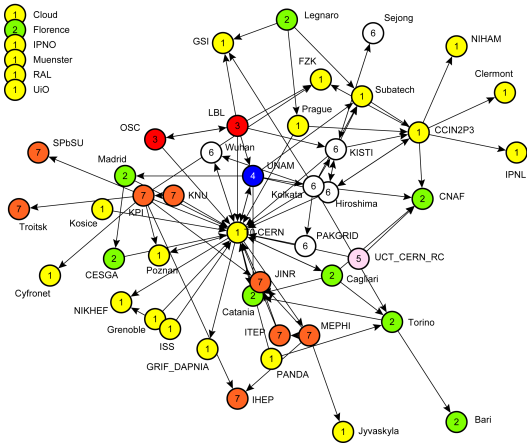

Figure 3. Wished interactions provided by regions.

among centres belonging to the same regions than among centres of distinct regions (figure 4), and the internet round-trip time (figure 4) is in their case lower.

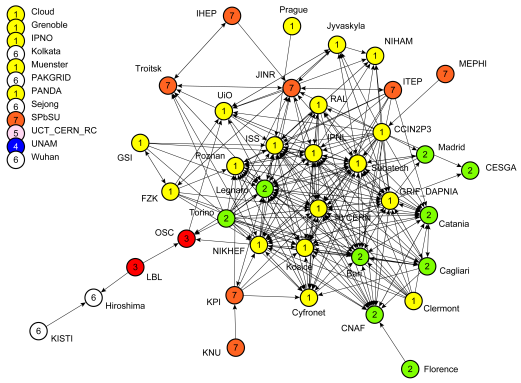

Figure 4. Bandwidth by regions.

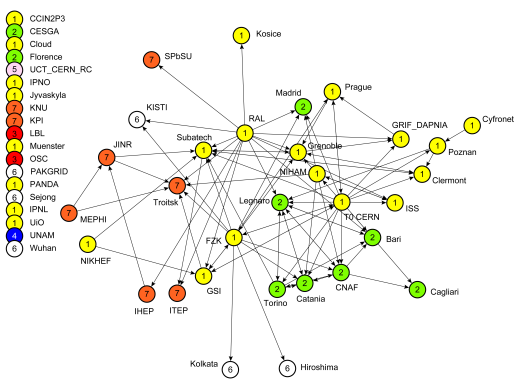

Figure 5. Data transfer by regions.

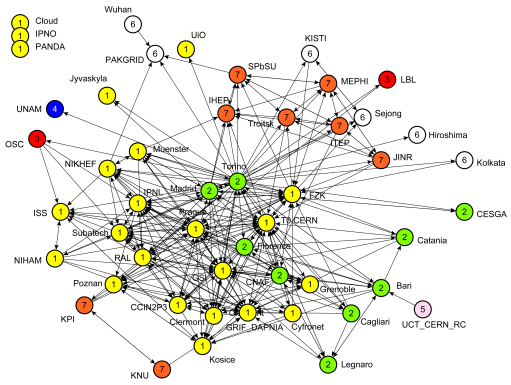

Figure 6. RTT Capacity by regions.

Are those graphical results confirmed by computational analysis? We first present the density within the main regions where most ALICE centres are located. Those regions are Northern Europe, Southern Europe, Russia and Asia. Indices for North America, South America and Africa were not computed as the number of countries associated with each regions is very small or even, in some cases, equals to 1. The density within Northern Europe was computed with and without CERN in it.

Table 2. Density within regions.

	Help	E-mail contacts	Wished collaborations	Bandwidth	Data transfers	RTT Capacity
North Europe (with CERN)	8.1	11.1	3.0	37.7	20.8	45.5
North Europe (without CERN)	3.5	8.2	1.5	35.3	16.8	41.1
South Europe	9.7	23.6	8.3	47.3	34.9	43.9
Asia	3.3	0.0	23.3	4.3	0.0	11.1
Russia	23.2	26.8	16.1	17.8	10.8	52.1
Total (with CERN)	4.3	4.9	3.6	16.2	8.4	18.2
Total (without CERN)	2.2	3.4	2.5	14.9	7.1	16.4

Table 2 shows that overall, the density on almost all indices based on self-reported exchanges is much higher within each region than in the full network. For instance, the density of contacts is twice higher in Northern or Southern Europe than in the full network. Similar patterns of collaboration within regions are found for email contacts. Interestingly, the observed exchanges are also frequently organised at the regional level. The only exception concerns wished collaboration, where there is a clear distinction between centres located in Northern Europe and centres in other regions. Centres in northern Europe are not especially seeking collaboration with other north European centres (average similar to the overall mean). Although the density of internal interaction is a bit higher, Southern Europe presents a quite similar pattern of results. Contrastingly, centres from Russia, and especially Asia, wish to increase their interactions at the regional level. The pattern of density of interactions among Asian centres is somewhat different from that in other regions. The density on self-reported and objective exchanges is low

in comparison with other regions or even the full network, while the wishes of centre managers for collaborations are very high.

We now estimate the extent to which those results are statistically significant. In order to know whether the patterns revealed in table 2 are due to chance or reveal structures, we first build an hypothetical network perfectly structured around regions, i.e., in which a tie exists if and only if the centres are of the same region. We then run a set of permutation models, based on QAP (Quadratic Assignment procedures) [7]. The Quadratic assignment methods compute correlations between the entries of two square matrices and assess the frequency of random and true correlations between them, in order to test whether or not those dimensions are correlated beyond chance. The algorithm proceeds in two steps. In the first step, it computes a Pearson's correlation coefficient between corresponding cells of the two data matrices (corresponding to the observed and hypothetical network data). In the second step, it randomly permutes rows and columns (synchronously) of one matrix (the observed matrix, if the distinction is relevant) and recomputes the correlation. The second step is carried out hundreds of times in order to compute the proportion of times that a random correlation is larger than or equal to the observed correlation calculated in step 1. A low proportion (<.01) results in the rejection of the null hypothesis of independence (a network with permuted centres could have a correlation with the hypothetical network at least as high as the observed network) and suggests a strong relationship between the matrices that is unlikely to have occurred by chance ([6]). Here again, Pearson's correlation coefficients were computed with and without CERN.

Table 3. Qap testing of regional influences (Pearson's correlation coefficient).

	Help	E-mail contacts	Wished collaborations	Bandwidth	Data transfers	RTT Capacity
with CERN	0.15**	0.24**	0.07**	0.30**	0.27**	0.40**
without CERN	0.16**	0.27**	0.10**	0.30**	0.26**	0.39**
**p<.01						

Table 3 shows the average value of Pearson correlations between each empirical matrix and the theoretical matrix constructed on the assumption that ties only exist within regions.It also provides a p-value which reports the proportion of times in which the permuted empirical matrix shows a correlation with the theoretical matrix equal or higher than the original empirical matrix. The table reveals significant effects of regions for all indices considered in this research. Belonging to the same region makes the density of interactions among centres significantly increase. This holds true for both self-reported and observed exchanges. For some indices, the difference is important, as where it is marginal for others.

5. Discussion

The patterns emerging from this study help us understanding the underlying structure of the ALICE Grid within the Worldwide LHC Grid structure. Access to data and resources is actually ubiquitous within the ALICE Grid. Members of the ALICE collaboration can submit work requests (jobs) from any point in the world and this gets executed by any centre available with the only constraint of data locality for reconstruction jobs, while simulation workload gets assigned only on the basis of CPU availability.

The first remark is that the distribution of the links for the different relationships in the network is highly non-random, as it can be seen in table 3. This does not come as a surprise, as

the Grid is a highly structured network with specific functionalities and connections. However this fact implies that the relations found in this paper are significant and can be interpreted as the effect of structuring factors.

If we look at the "physical" layer (RTT and bandwidth) we can "see" some regionalisation in figure 4 and figure 4. This is already a non trivial observation because the presence of large supranational initiatives like Géant in Europe establishing high-speed international networking.

Coming to the role of CERN, from table 1 we note that it does not have a particularly important role in the structure of the physical layer, as the different indicators for RTT and bandwidth show relatively little variation with and without it. Only the substantial reduction of the indegree centralisation in the RTT suggests the role of "hub" the CERN plays in the ALICE physical network. This role is confirmed by the opposite trend of the betweenness centralisation, indicating that the removal of CERN leaves a network where more centres play the role of relay between peers in a more "democratic" network.

If we now move to data transfer, we observe that CERN's role as a "hub" (Indegree centralisation) seems less prominent, while the network of data exchanges seems to be strongly hierarchised around CERN (Outdegree centralisation drops by 30% with CERN's removal). This is consistent with the fact that CERN is the ultimate "source" for the experiment data. This is mitigated by the fact that MonteCarlo data are generated at all the nodes of ALICE's Grid. A large change in the number of cliques also seems consistent with CERN's role of producer of data and receiver of elaborated data for custodial storage which makes of it a preferred vertex of fully connected subsets.

The above suggests that at the "physical" layer the ALICE Grid looks little hierarchised and rather homogeneous. The data transfer is characterised by CERN's special role as source of the data, but otherwise also shows little structure. Regional densities for the "physical layer" shown in table 2 differ largely for the different regions.

If we now turn to the "relational layer" we observe a largely different situation. The three categories Help, E-mail contacts and Wished collaborations are again shown in table 2. The density of the three categories show a large change with and without CERN. The nature of this change can be understood looking at the other parameters. The betweenness centralisation is reduced by an order of magnitude removing CERN and the effect is even more dramatic for the Help category, implying a structure where the largest number of links point to CERN, as can be easily seen in the pictures 4-4. The drop of Indegree centralisation is also very large, clearly indicating the "hub" role of CERN in providing help, email contacts and as a target for wished collaboration. Similar consideration for the number of cliques in wished collaborations. It seems that CERN is seen as the "missing vertex" for multilateral collaborations.

Interestingly enough the drop in outdegree centralisation, which measures the hierarchisation of the system, is less pronounced, and this may suggest that the hierarchical structure remains even if the "top" is removed. Looking at the graphs this hierarchisation appears very clearly to depend on geographical regions. The density of wished collaboration within regions seems to be inversely proportional to the perceived existing level of help and email exchange, which is an expected result.

It is interesting to compare these findings with the hierarchical organisation foreseen by the MONARC model [8] which has largely served as a blueprint for the WLCG Grid. The structure for the MONARC "physical layer" was influenced by the foreseen limitation on the network bandwidth and the projected need to optimise resource utilisation within this constraint. The actual evolution of the network capacity has been beyond all expectations. At the end of last century, when the MONARC proposal was put together, researchers were hoping to have 622 Mb/s links, while now links with a capacity one order of magnitude more are commonplace. As a consequence of this, at the "physical layer" this model has been replaced by a model evolving toward a more "democratic" cloud paradigm. In the case of ALICE, this move has

been accompanied by precise architectural choices in the design of the middleware and in the operation of the Grid, with the aim of optimising the usage of resources. The fact that the data transfer follows this pattern is in this sense revealing that the architectural design is reflected in the actual dataflow.

At the "relational" layer a completely different picture emerges. Geographical regions play a large role in modelling the infrastructure, which looks highly hierarchised in a way similar to the one suggested by the MONARC model. While the machine-to-machine exchanges are designing a "cloud", the human relationships, both actual (email exchange), perceived (help) and wished (wished collaborations) follow a hierarchical schema strongly influenced by geographical and national / cultural elements. This layer looks like an example of dynamic "self organisation" capable to adjust itself in order to "optimise" the Grid usage from the "user's perspective" within the constraints coming from the "physical layer". Of course the "physical layer" itself also evolves as a result of the how the system is really used, aiming at improving efficiency and resource utilisation. Analysing the mutual influence between these two layers could be an interesting subject for future studies. Moreover, it seems that these two layers, although different, can work together in a complementary fashion.

We believe that this work is just the beginning of what could be a much deeper analysis of the complex human and technical relations that link the Grid nodes. We have built a large, highly non-linear system and, in some sense, we are just beginning to understand it. A massive amount of monitoring information is now accumulating, thanks also to the constant improvement of the monitoring tools. The analysis and understanding of these data will be essential to grasp the functioning of the Grid, and this paper can be considered part of this effort. However there is also another, even more exciting perspective that has to be considered, and this is the usage of these data to optimise the working of the Grid, making it more efficient and resilient. For the data that can be collected in real time, the tuning process can happen dynamically, making the Grid a self-adjusting system. Can "relational" data be used as part of this optimisation? This paper does not answer this question, but we hope it shows that these data can at least be collected and analysed meaningfully. In this sense a lot of interesting work lies ahead of us.

How much the results shown relate to the evolution of Internet as a whole outside HEP is also a very relevant question, and it would be important to compare our results to similar studies in different fields. At the same time it would of course be very interesting to repeat this study in the future to see how the situation evolves.

Acknowledgments
The authors wish to thanks Dr. Iosif Legrand for his insightful reading of the draft.

References

[1] E D Widmer F Carminati C G and Carminati G G 2010 How alice computing centres connect? an exploration of cooperative ties among computing centres. *13th International Workshop on Advanced Computing and Analysis Techniques in Physics Research. ACAT 2010* (Proceedings of Science)
[2] Wasserman S and Faust K 1994 *Social Network Analysis: Methods and Applications* (Cambridge University Press)
[3] Worldwide large hadron collider computing grid http://lcg.web.cern.ch/LCG/public last updated on 05-10
[4] Burt R 2001 *The social capital of structural holes* (New York: Russel Sage Foundation)
[5] Scott J 2000 *Social Network Analysis* (Newbury Park, CA: Sage Publications Ltd; 2nd edition)
[6] S Borgatti M G E and Freeman L C 2002 *Ucinet 6 for windows. software for social network analysis.* technical Report
[7] Hubert L J and Schultz J 1976 *British Journal of Mathematical and Statistical Psychology* 190–241
[8] The monarc project http://monarc.web.cern.ch/MONARC last updated on 13-01-2000