

How ALICE computing centres connect? An exploration of cooperative ties among computing centres

Eric D. Widmer

Sociology department, University of Geneva, Switzerland

E-mail: Eric.Widmer@unige.ch

Federico Carminati*

CERN, Geneva, Switzerland

E-mail: federico.carminati@cern.ch

Costin Grigoras

CERN, Geneva, Switzerland

E-mail: costin.grigoras@cern.ch

Giuliana Galli Carminati

Psychiatry Department, Adult Psychiatry, UPDM University of Geneva, Switzerland

E-mail: Giuliana.GalliCarminati@hcuge.ch

This study aims at characterizing the collaboration ties between the centres of the Distributed Computing ALICE infrastructure based on social networks methods. A selfadministered questionnaire was sent to all center managers about support, email interactions, conflict and future collaborations in the infrastructure. Answers for 50 centers were received and considered in the empirical analysis (68% response rate). Empirical analysis shows that the network is centralized on CERN. There are however a series of centers which form cohesive subgroups, and a willingness to promote more interactions at the regional level. Interestingly, social relationships are only loosely correlated with computer exchanges.

*13th International Workshop on Advanced Computing and Analysis Techniques in Physics Research,
ACAT2010*

February 22-27, 2010

Jaipur, India

*Speaker.

1. Introduction

This study aims at characterising the collaboration ties between the centres of the Distributed Computing ALICE infrastructure based on social networks methods [Wasserman and Faust, 1994; Scott, 2000]. The computing needs of the ALICE experiment range in the 200 kHEP-SPEC06/year [HEPiX, 2006] and 10 Peta-Bytes (10^{15} Bytes) of disk storage. This is provided by 73 computer centres worldwide arranged in Tiers (1 T0, 6 T1s¹, 65 T2s). These centres are part of the Worldwide Large hadron collider Computing Grid (WLCG [WLCG, 2010]) These centres 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 social relations between the different centres is as important as the material conditions to ensure the proper functioning of this complex system extended over different timezones and continents.

Social network analysis is a discipline of the social sciences which from its start in the nineteen sixties has been dedicated to the study of complex patterns of interactions in large groups of social actors. Based on the assumptions that human action is better understood as embedded in series of interdependencies with others, it has developed a range of computing and visualising methods [Granovetter, 1973; Scott, 2000; Borgatti et al., 2002; Wasserman and Faust, 1994] that have been applied on a large series of issues, from businesses exchanges, kinship ties, friendship interactions, to the Internet. Based on such approach, we estimate the extent to which various dimensions of social relationships (support, interactions, future collaborations) follow similar structural patterns. We wonder how central CERN is within the network and we estimate the extent to which centres are interrelated beyond their direct connections with CERN. We also measure the extent to which social relationships are interrelated with computing activities within the network. In other words, we assess the extent to which social relationships are autonomous from computing tasks and data transfer networks.

2. Data and measures

The sociometric data were collected using a self-administered questionnaire which was sent to 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 make it difficult for their centre to carry out their job responsibilities (conflict), 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 (collaborations). Aside managers had to estimate on a five-item scale how they rated the functioning of their centre. 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 such as the theoretical capacity of the network linking the centres, the actual quantity of data exchanged, the physical distance (in kilometres) between centres, and the Internet Round Trip Time (RTT). Topology information (both geographical and Internet network) was extracted from

¹One of them is a two-site T1, each one of the site having provided an answer.

the MonALISA MonALISA [2003] monitoring tool that has one instance deployed in each of the ALICE sites. Part of the monitoring suite is a Round Trip Time test between all sites and a bandwidth estimator based on the Fast Data Transfer (FDT MonALISA [2008]). The connectivity between the sites is continuously assessed and the average values were used in determining the correlations. We focus this exploration on support, interactions and collaborations and see how they compare with objective data such as physical distance and actual quantity of data exchanged.

In addition to visualising the network of ties [Batagelj and Mrvar, 1998], various measures are proposed by network analysis ([Burt, 2001; Scott, 2000; Wasserman and Faust, 1994]). Density in a directed graph is equal to the number of existing arcs divided by the total number of possible arcs. Transitivity is calculated as the proportion of triads in which $i \rightarrow k$ exists when $i \rightarrow j$ and $j \rightarrow k$ exist. In addition, two measures of network centralisation were computed, which capture different conceptual dimensions of centrality: in-degree and out-degree centralisation refer to the number of ties pointing to and going from a specific centre. It provides 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 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. We compute these parameters in the overall network, with and without CERN in it. This procedure is set up to estimate the impact of CERN on the overall structure. If the number of components greatly increases when CERN is removed from the network, it shows that its impact is high, i.e. its removal makes the whole network significantly less centralised

3. Results

Figure 1 considers which centres are support resources for other centres. Arrows originate from centres which support others (support providers) and point to centres which are supported by others (recipients of support).

Figure 1 shows that Support is perceived as coming essentially from CERN with some signs of presence of local support structures in Italy, France and Russia. When CERN is removed from the network, it becomes split in two rather large components and a series of isolates. The centralisation of the network is very large and mostly associated with the role of CERN.

Interactions through regular e-mail contacts shows similar patterns than support (see Section 3)

Overall CERN remains the major actor within the network. When CERN is removed, the network splits in that case into three clearly bounded groups, with a high density of internal interactions and a rather low connectivity to other centres, except for CERN. Interestingly, those subgroups have a clear territorial and geographical basis. This seems to confirm the fact that the original MONARC [MONARC, 2000] model of hierarchical relations between centres has evolved into a so-called "Cloud" model, and in particular one where geographical relations dominate the support and exchange relations between centres.

According to the MONARC model, relations between the computing centres should have been hierarchically arranged, with the T2s receiving support from, and exchanging communications mainly with their reference T1. This model was based on a precise definition of the roles of the different centres, which should have been accompanied by a correspondingly deterministic data

Figure 1: Support relations among ALICE Centres

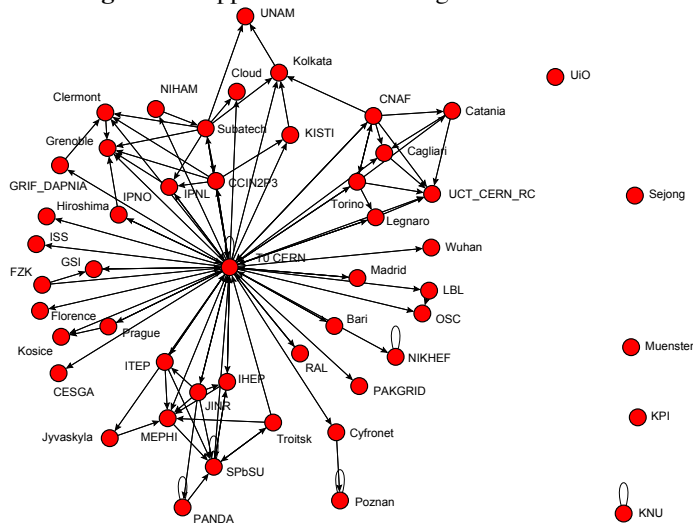
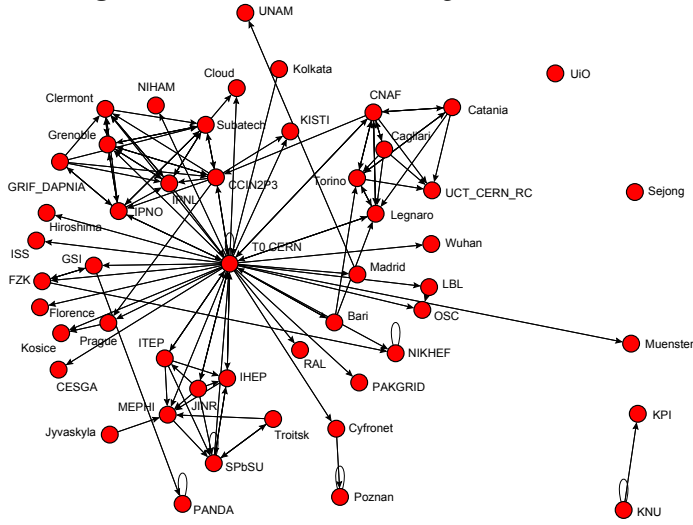


Figure 2: Email interactions among ALICE Centres



placement and job execution. One of the reasons for such a model was the outlook that was available at the time (1990's) on the evolution of the network capacity, which suggested a hierarchical arrangement to optimise a resource (the network) that was anticipated to be rather scarce compared to the needs. The evolution of the network capacity, both for commodity networks and for research, lead to much larger capacities than expected, albeit still at a considerable cost. This has lead to a relaxation of the hierarchical structure, which has allowed other resources which are now more critical than network capacity, such as CPU and data storage, to be better used. This has also allowed the adoption of a more flexible resource procurement policy, as T1 and T2 resources are, to a large extent, interchangeable on the Grid, and therefore is more and more true that what counts

is the total capacity (CPU and storage) available and not so much where it is provided.

All this has led to a more “democratic” computing model, evolving from a hierarchical structure to a more peer-to-peer (P2P) paradigm. This trend is not specific to High Energy Physics (HEP), but it is a general one, and the new P2P paradigm is generally indicated as “Cloud” computing. In some sense this is a return to the original Grid manifesto [Kesselman and Foster, 1998, 2003], which was proposing a complete “delocalisation” of the computing resources. Although the present study is concerned with social interactions and not with data or job placement, its results are symptomatic of this evolution, as the MONARC model was foreseeing that support and communication relations would follow the same hierarchical arrangement as data and processing.

The trend to privilege geographical proximity is clearly visible for the Italian and French cluster, where it seems that e-mail and support relationships follow very similar patterns. The close relationship between the Russian (and nearby) centres well reflects their plan to establish and operate a “distributed T2”. This is also reflected in the Table 3, where geographical distance is clearly related to support and to interactions. The exceptions to this pattern reflect special support relationships established via international collaboration agreements. An example of such collaboration is the KISTI (Republic of Korea) to CCIN2P3 (France) relationship.

Unfortunately only two of the NorduGrid centres have responded to our survey, in spite of the active participation of the NorduGrid complex to the ALICE Grid. The presence of the other Nordic centres would have been interesting to see how their distributed T1, a unique structure within the WLCG, interacts.

The case of Germany is also interesting to note, because, contrary to other regional communities, there seem to be relatively little activity of support or relationship between these centres, which relates mostly to CERN.


Another fact which is worth noticing is that there seems to be no hint of a special collaboration between centres in Asia (India, Japan, China and Korea), in spite of them being somewhat closer to each other than to CERN. This reflects the fact that, in spite of several attempts to create a forum for Asian-Pacific LHC computing resources, this has never really materialized.

Interestingly, the graph about expected collaborations provides a quite distinct pattern of interactions (see Figure 3).

When asked about whom they wished to develop more interactions with in the near future, respondents provided responses less structured by subgroups organised on geographical proximity. Their responses also granted less centrality to CERN compared with support and interactions. Two rather large subgroups coexist, with CERN as the connection node between them. Some major tendencies can be detected here.

On one side there seems to be a clear desire for the reinforcement of the connections between the sites of the Russian distributed T2. According to our operational experience, and also to the results shown in the previous two pictures, these relationships are already rather good. This call for further integration shows that the operation of a multi-centric site requires an even higher level of communication of what is there now.

The second tendency that can be detected is the aspirations of the T1s to a greater communication. While we have seen that communication and support are dominated by the geographic or proximity factor, T1s find themselves to solve very similar problems, in particular in the area of custodial storage of data. It is also our empirical remark that an enhanced collaboration among

Figure 3: Expected collaboration among Alice Centres

them, for instance in the area of mass storage, could have led to an increased standardisation, with a substantial scale economy. This has been several times evoked, but with no real action. Still the TIs appear to be willing to increase their communication and mutual support relations.

A third strong trend is that the “peripheral” centres (with respect to CERN) seem to express a strong request for support. However these requests are directed to European centres with whom they happen to have had relations in the past. This can be due to a tutorial or a visit by an expert of another centre or to institutional collaboration agreements that however did not fully develop into a technical collaboration yet. This seems to be true for centres both in the Asian-Pacific zone and in the US. This would suggest not only that they indeed need more support, but also that this needs to be organized in a more systematic way. Again it seems that the establishment of regional support and exchange infrastructures for Computing Centres should be pursued more vigorously.

The parameters of the graphs presented in Table 3 provide additional information. They first present density and centralisation measures for the whole network. Then, it computes the same parameters when the CERN is removed from the network, in order to fully estimate the structural role that CERN plays within the ALICE infrastructure. Measures on the full network shows the tremendous centralisation of the ALICE infrastructure for support provision. Interestingly, the reception of support is much more equally shared. The network of interactions is significantly less centralised than the network for support. As for future collaboration, they are meant to be developed in a less centralised way than the current situation stands. Indeed, the degree and betweenness centralities are lower for wished collaborations than for support and email interactions.

When CERN is removed from the network, the parameters related to centralisation highly decrease, making the network almost non-centralised. This confirms that CERN plays a key role of centralising support and exchanges in the ALICE infrastructure. This effect of CERN is however weaker in the network of wished collaborations.

Are there significant correlations between support, interactive and wished collaborative ties in the ALICE network? Also, how do those dimensions relate to computing activities such as the ac-

Table 1: Relational parameters of support, interactions and future collaborations

	Support	Interactions	Future collaborations
With CERN			
Density	4	5	4
Betweenness centralisation	36	31	26
Indegree centralisation (provided)	33	18	13
Outdegree centralisation (providing)	74	60	42
Without CERN			
Density	2	3	3
Betweenness centralisation	1	3	2
Indegree centralisation (provided)	11	9	6
Outdegree centralisation (providing)	13	12	13

tual quantity of data exchanged, the physical distance between centres, and the Internet Round Trip Time (RTT). In using Quadratic Assignment ([Hubert, L. J. and Schultz, J., 1976]), we test whether or not those dimensions are correlated beyond chance. The Quadratic Assignment methods compute correlation between entries of two square matrices, and assess the frequency of random and true correlations between them. The algorithm proceeds in two steps. In the first step, it computes Pearson's correlation coefficient (as well as simple matching coefficient) between corresponding cells of the two data matrices. 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 ($<.05$) suggests a strong relationship between the matrices that is unlikely to have occurred by chance ([Borgatti et al., 2002]).

Table 2: QAP Correlations among relational parameters

	Support	Interactions	Future collaborations	Capacity	Bandwidth	Distance
Support	0	0.68**	0.3**	-0.07	0.03	-0.1*
Interactions		0	0.25**	-0.1	0.13*	-0.16**
Future collaborations			0	0.08	0	-0.02
Capacity				0	-0.13*	0.42**
Bandwidth					0	-0.26
Distance						0

** $p < .01$; * $p < .05$

Table 3 shows that there are strong correlations between support, interactions and wish to de-

velop further collaborations in the future. Interactions and support are especially associated. Future collaborations are more loosely correlated with interactions and supports. All three questions are weakly correlated with objective measurements. Only interactions show significant associations with bandwidth and distance. On average, the more distance between two centres there is, the fewer interactions there are, and the more bandwidth there is, the more interactions there are. The coefficients show however only average correlations.

4. Discussion

Overall, the centres of the ALICE Distributed Computing Infrastructure derive most of their support from CERN. Interactions through emails are also centralised on CERN, although to a lesser extent. Communication is clustered mostly according to geographic locality as foreseen by the WLCG computing model and TIs seem to fulfil to some extent their role of support of smaller centres.

There is however a clear request for increased collaboration at the local level. Centres of areas with a small concentration of sites (Asia, South America), tend to privilege existing relations due to personal contacts or collaboration. agreements rather than regional communication.

The data also reveal that social relationships are only loosely correlated with the computing activities, which suggest that the impact of other variables should be assessed. Cultural, political and historical proximities among countries may play a crucial role for the understanding of social interactions within the ALICE infrastructure. Future research will estimate the impact of such variables.

References

- Batagelj, V. and Mrvar, A. (1998). *Pajek. Program for analyzing large networks*, pages //vlado.fmf.uni-lj.si/pub/networks/pajek.
- Borgatti, S., Everett, M. G., and Freeman, L. C. (2002). *Ucinet 6 for windows. software for social network analysis*. Technical report.
- Burt, R. (2001). *The social capital of structural holes*. Russel Sage Foundation, New York.
- Granovetter, M. S. (1973). The strength of weak ties. *American Journal of Sociology*, 78(6):1360–1380.
- HEPiX (2006). HEP-SPEC06 Benchmark. <http://w3.hepix.org/benchmarks/doku.php>. Last updated on 19-11-09.
- Hubert, L. J. and Schultz, J. (1976). Quadratic Assignment as a General Data Analysis Strategy. *British Journal of Mathematical and Statistical Psychology*, 29:190–241.
- Kesselman, C. and Foster, I. (1998). *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann Publishers.

- Kesselman, C. and Foster, I. (2003). *The Grid 2: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann Publishers.
- MonALISA (2003). In *Conference for Computing in High-Energy and Nuclear Physics*, La Jolla, California, USA.
- MonALISA (2008). Fast Data Transfer. <http://monalisa.cern.ch/FDT>.
- MONARC (2000). The MONARC Project. <http://monarc.web.cern.ch/MONARC>. Last updated on 13-01-2000.
- Scott, J. (2000). *Social Network Analysis*. CA, Newbury Park.
- Wasserman, S. and Faust, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge University Press.
- WLCG (2010). Worldwide Large Hadron Collider Computing Grid. <http://lcg.web.cern.ch/LCG/public>. Last updated on 05-10.

POS (ACAT2010) 045