2013 IEEE INTERNATIONAL WORKSHOP ON M&N Measurements and Networking
October 7-8, 2013 | Naples, Italy

Practical Early Detection of Performance Degradation in Aggregated Traffic Links

D. Fernández-Hermida†, M. Rodelgo-Lacruz‡, C. López-Bravo†, F.J. González-Castaño†, V. Estévez-David‡
† Grupo de Tecnologías de la Información – AtlantTIC, Universidad de Vigo, Spain {davidfh,clbravo,javier}@gti.uvigo.es
‡ Gradiant, Vigo, Spain {mrodelgo,vestevez}@gradiant.org

Abstract

The growth of Internet traffic and the many different traffic classes that exist make network performance control extremely difficult for operators. The methods available rely on complex or costly hardware. However, recent research on bandwidth sharing has introduced methods that require only basic statistics of aggregated link utilization, such as mean and variance. This data can be easily obtained through SNMP calls, lowering the cost of monitoring systems. Unfortunately, to the best of our knowledge, no tools have yet been developed to implement these methods. This paper presents the implementation of a plugin for the monitoring environment Nagios and the validation of a degradation detection tool from link utilization traces. The plugin does not require complex or costly hardware for acquiring data. Instead, it employs basic SNMP data about link utilization.

Method

$$I(\rho) = \frac{1 - \frac{V_U(\tau)\rho}{V_0}}{1-\rho}$$

$$V_0 = \lim_{\rho \to 0} \frac{V_U(\tau)}{\rho}$$

In [1] the relationship between the variance of link utilization and the degradation of the TCP flows is formulated. It is possible to derive the degradation index $I(\rho)$, where $V_U(\tau)$ is the variance of τ samples, ρ is the link utilization, and V_0 is the variance of the traffic when the link is not saturated.

The result of $I(\rho)$ is proportional to the delay factor f_r, which represents the ratio between the transfer times of a flow with and without link saturation. However, the proportionality between $I(\rho)$ and f_r cannot be calculated, so $I(\rho)$ can be used to detect variations in link degradation but not degradation in a particular time window. We can monitor the deviation of $I(\rho)$ when the link utilization is low (e.g. $\rho < 0.5$) and the performance is not degraded and setting a thershold that the value $I(\rho)$ rarely exceeds, such as three times the standard deviation from its mean.

[1] Ishibashi, K.; Kawahara, R.; Asaka, T.; Aida, M.; Ono, S.; Asano, S., "Detection of TCP performance degradation using link utilization statistics," IEICE Transactions on Communications, 89(1), pp. 47-56.

Structure

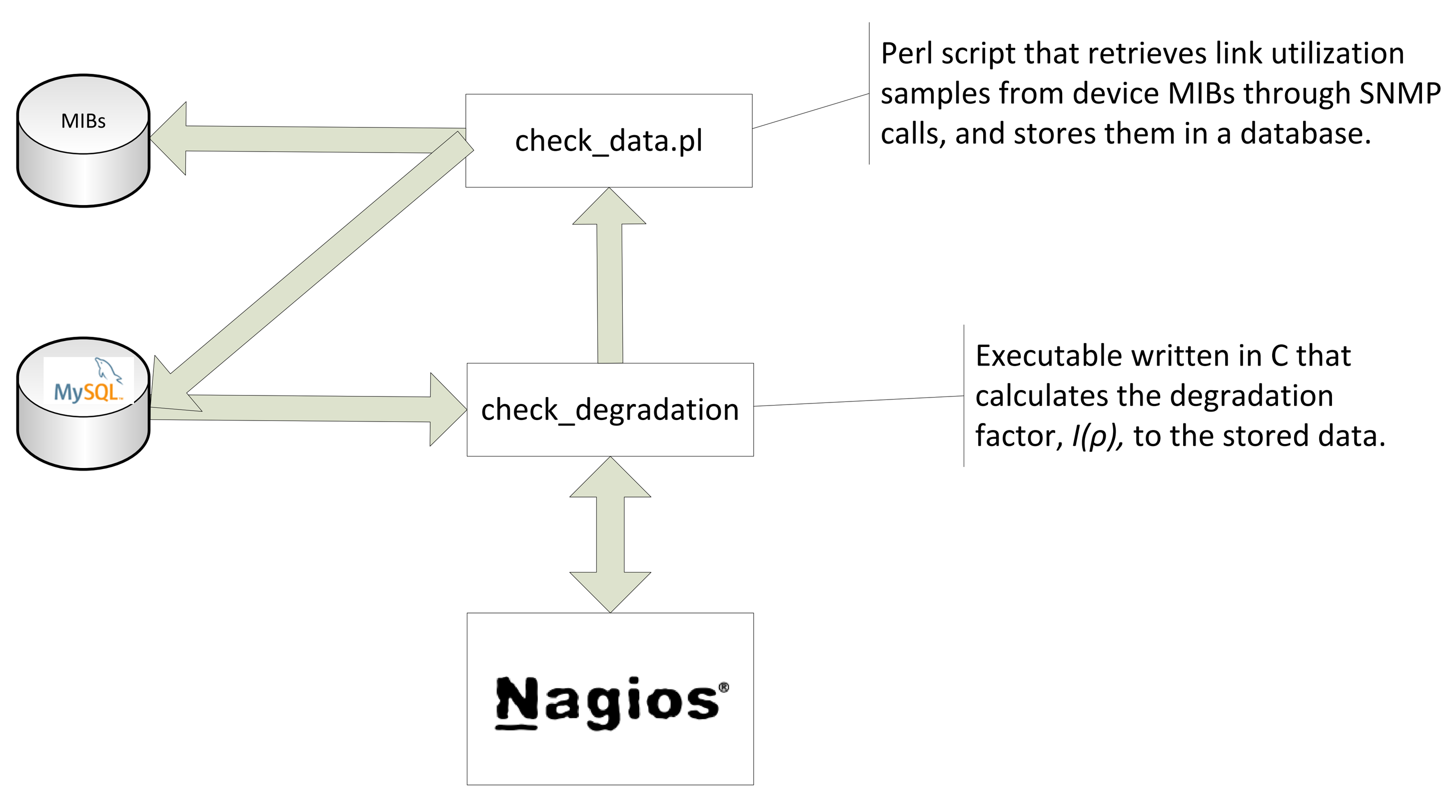

Simulator

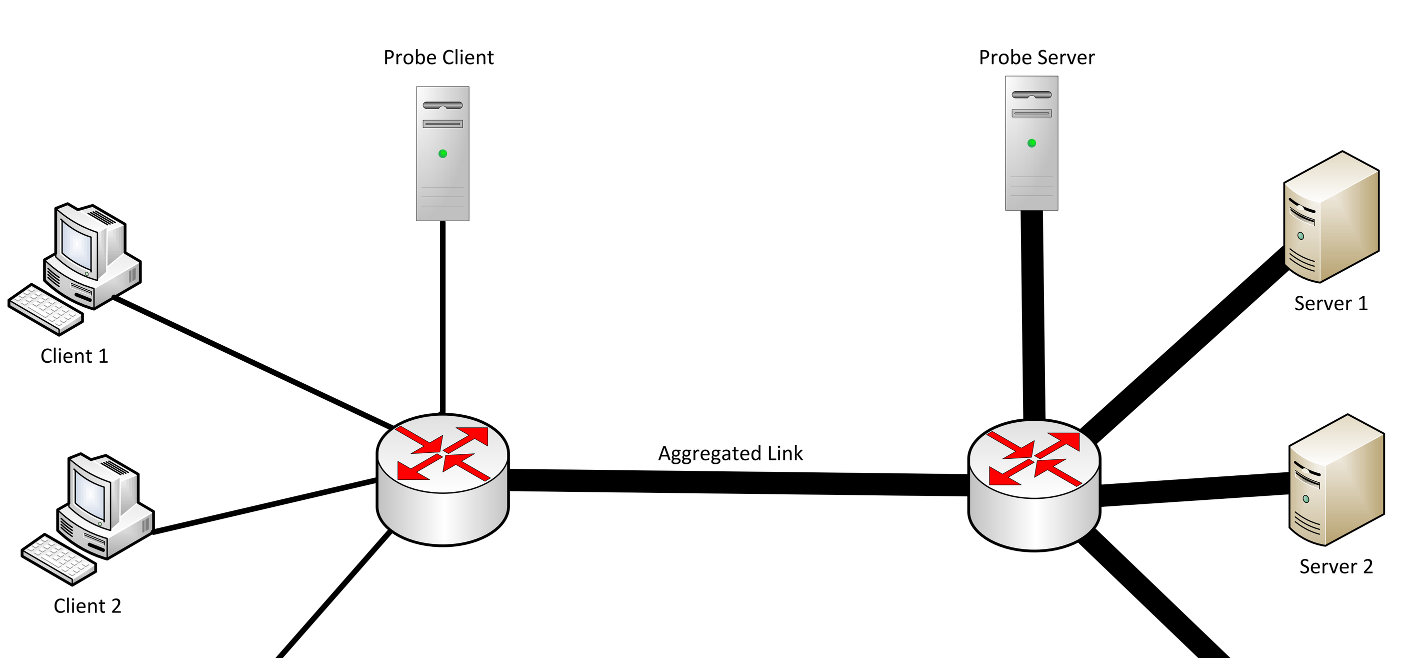

The plugin estimates traffic degradation in scenarios that meet the M/G/R processor sharing model assumption, that is, in links aggregating the traffic of many users with lower access capacity (the usual in telco operator links). Thus, utilization samples from aggregated links would be necessary to validate the plugin functionality, in addition to measures of real traffic degradation for each trace, to be compared with the plugin results. We developed our own event-driven simulator using the OMNET++ libraries and the INET Framework.

Simulation Results

We considered three different scenarios consisting of a constant traffic load with low, medium, and high degradation. These figures show the changes in aggregated link utilization and $I(\rho)$ over 6 and 50 hours. The mean delay factors for these scenarios were 1, 1.5 and 6, respectively. The figures for 50 hours indicate the $I(\rho)$ thresholds for saturation detection at 1.52, 1.92 and 2.06, respectively. Figures for 6 hours are detailed plots of the 50 hours showing the effect in $I(\rho)$ of changing the number of τ samples used in $V_U(\tau)$.

In the low saturation scenario the degradation index is close to 1, always below the threshold (set at 1.52), indicating no link saturation. In the medium saturation scenario the degradation index increases, with a mean value of 2 and peak values of up to 8. These values indicate that a file transfer took up to eight times longer in the case of saturation. The high saturation scenario reveals strong degradation with a mean degradation index value of 10 and peak values of 30.

In the medium and high saturation scenarios the degradation index was highly variable. Note that it is an instantaneous value, whereas the delay factor value is an average. Due to the high variability of the graphs, the best way of interpreting them was to take into account the maxima of the plot, or to smooth them with a low pass or exponential filter.

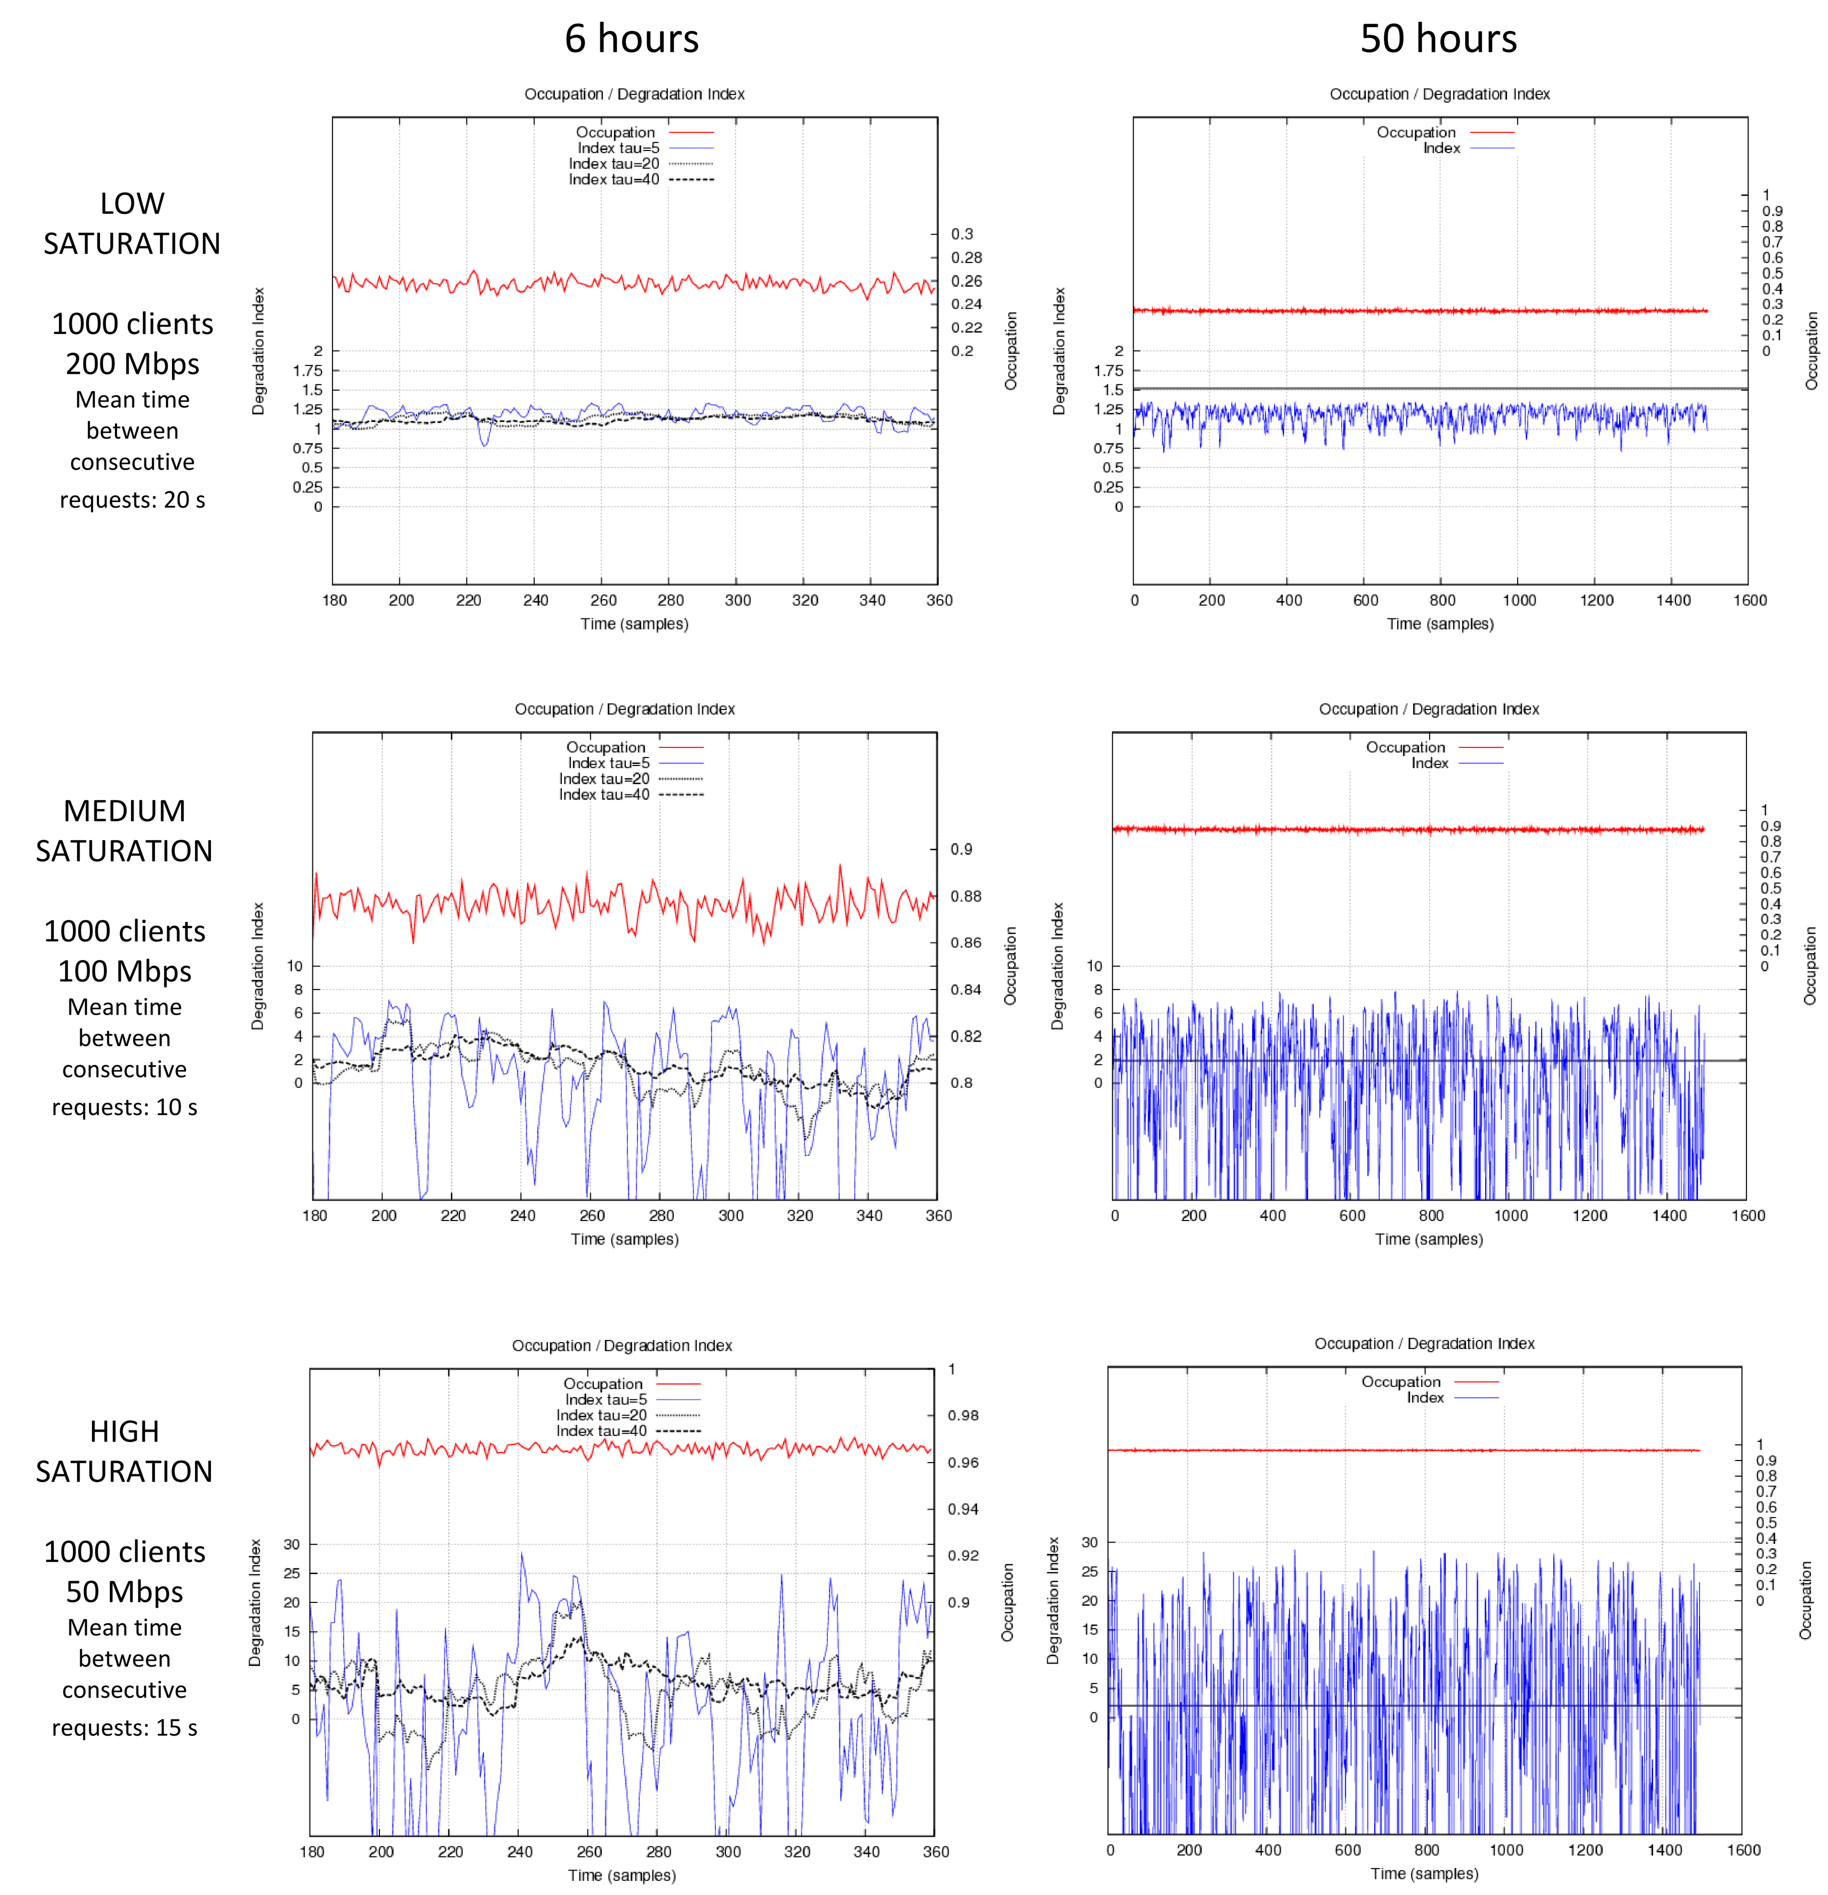

Real Scenario Results

As an example of a real scenario, we obtained SNMP utilization traces from an operator, corresponding to a window of 50 hours at a sampling rate of 1 sample every 120 seconds. The problem with these utilization traces is that we ignored the instantaneous level of saturation. However, since we knew that most clients had an access link capacity of 20 Mbps and the aggregated link had a capacity of 160 Mbps, we were able to estimate a degradation threshold for the utilization with the model in [2]. Assuming a capacity ratio of 8, and that TCP flow degradation would start at a delay factor of approximately 1.1, the model estimated a threshold of 68.6% aggregated link occupation.

Figure shows the result of applying the plugin to those traces with the threshold for $I(\rho)$ set to 1.99. Note the daily traffic pattern, with minimum utilization at night and maximum utilization in the evening. There are two intervals in the graph where occupation exceeds the 68.6% occupation threshold and traffic is presumably degraded: approximately between samples 30 and 300 and between 750 and 980 (highlighted in yellow in the figure). We can see that the degradation index increases significantly in these intervals and reaches values of up to 4.5; however it stays close to 1 in the rest of the graph. In fact, the intervals where the plugin alerts of saturation are between 10 and 350 and between 720 and 1000. Thus, the results of both methods are very similar, although the plugin does not need to know user access capacity.

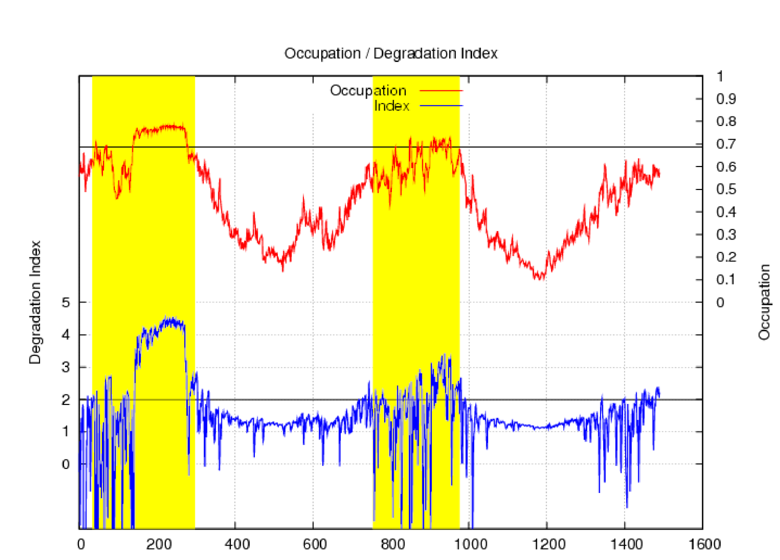

[2] Kawahara, R.; Ishibashi, K.; Asaka, T.; Ori, K., "A method of IP traffic management using TCP flow statistics," IEEE Global Telecommunications Conference, 2003, GLOBECOM '03, vol. 7, pp. 4059-4063.

This research has been supported by projects:
CALM (TEC2010-21405-C02-01), Mineco, Spain
AtlantTIC (CN 2012/260), European Regional Development Fund (ERDF) and Galician Regional Government, Spain
MEFISTO (10TIC006CT), Xunta de Galicia, Spain

Accessible at: Nagios Exchange and sourceforge

AtlantTIC | Universidade de Vigo | Gradiant Centro Tecnolóxico de Telecomunicacións de Galicia | Xunta de Galicia | European Union European Regional Development Fund