1B.1 IMPLEMENTATION OF THE PHASE I OF TIGGE

Baudouin Raoult*, Manuel Fuentes, Jörg Urban

Data and Services Section, ECMWF

Douglas Schuster, David Brown, Steven Worley

Computational & Information Systems Laboratory, NCAR

Tom Yoksas

Unidata, UCAR

1. INTRODUCTION

TIGGE, the THORPEX Interactive Grand Global Ensemble, is a key component of THORPEX: a World Weather Research Programme to accelerate the improvements in the accuracy of 1-day to 2 week high-impact weather forecasts for the benefit of humanity. The key objectives of TIGGE are:

- An enhanced collaboration on development of ensemble prediction, internationally and between operational centres and universities,
- New methods of combining ensembles from different sources and of correcting for systematic errors (biases, spread over-/under-estimation),
- A deeper understanding of the contribution of observation, initial and model uncertainties to forecast error,
- A deeper understanding of the feasibility of interactive ensemble system responding dynamically to changing uncertainty (including use for adaptive observing, variable ensemble size, on-demand regional ensembles) and exploiting new technology for grid computing and high-speed data transfer,
- Test concepts of a TIGGE Prediction Centre to produce ensemble-based predictions of high-impact weather, wherever it occurs, on all predictable time ranges,
- The development of a prototype future Global Interactive Forecasting System.

2. DEFINING TIGGE

The first workshop on TIGGE was held from 1 to 3 March 2005, at ECMWF. The purpose of this workshop was to collect the views of the community on what the TIGGE science aims should be, what the requirements are for use of the TIGGE data and hence what are the infrastructure requirements.

During the first meeting, the goals and requirements of TIGGE were stated as follows:

- Collect in near real-time global ensemble forecast model output, including those from operational centres.
- Organize the common formatted data at several Archive Centres and make it available to researchers in the operational and academic communities.
- Provide easy access to long series of data from multiple centres to enable a broad array of research, e.g. bias correction or the optimal combination of ensembles from different sources
- Use TIGGE as a basis for a development prototype for the future Global Interactive Forecasting System.

The TIGGE database will contain a core dataset consisting of ensemble forecasts generated routinely at different centres (Data Producers) around the world. The core dataset is estimated to be 200 GB/day. These data will be complimented with observational data and existing datasets including re-analyses and re-forecasts (e.g. ERA40), as well as special datasets generated during the TIGGE project for specific research and applications.

The database could also include intermediate verification output data that can be used to calculate individual skill scores.

Access to the database will be through websites that will also support links and pointers to associated regional and user-specific observational data sets.

2.1. Data policy

The leading principal is that all TIGGE data are to be made available to all users for research and educational purposes.

* e-mail: baudouin.raoult@ecmwf.int

User registration is required prior to accessing data and the process for approval will be handled electronically. Requiring simply a valid e-mail address and acknowledgment of conditions of supply.

As the default a 48-hour data access delay from the forecast initialization time will be imposed. Special consideration will be given to cases where real-time access to data is necessary for demonstration projects and field experiments. Real-time access will be handled via the the THORPEX International Program Office.

2.2. Phase I, II, and Partners

TIGGE is planned to have two Phases, Phase I is where archives are used as central collections points and data distribution hubs, and Phase II will be designed following a distributed design without central collection.

CMA (China), NCAR (USA), and ECMWF (Europe) have volunteered to act as Archive Centres for Phase I. During this Phase the Internet will be used to transport the data in near real-time, the provider centres are expected to be the data backup for their own information, and the Centres will largely mirror each other's data holdings. Schematically, this is shown in Figure 1.

The three grey boxes represent the Archive Centres that will receive data from different Data Providers (the coloured arrows). Each Archive Centre has a copy of all the data from all the Data Providers, which form the TIGGE database (shown as multicoloured slabs in the cylinders). A user can extract a selection of products across all Data Providers from any Data Centre. For Phase I NCEP (USA), ECMWF (Europe), UKMO (United Kingdom), JMA (Japan), BMRC (Australia), CPTEC (Brazil), KMA (Korea), and MSC (Canada) will be Data Providers.

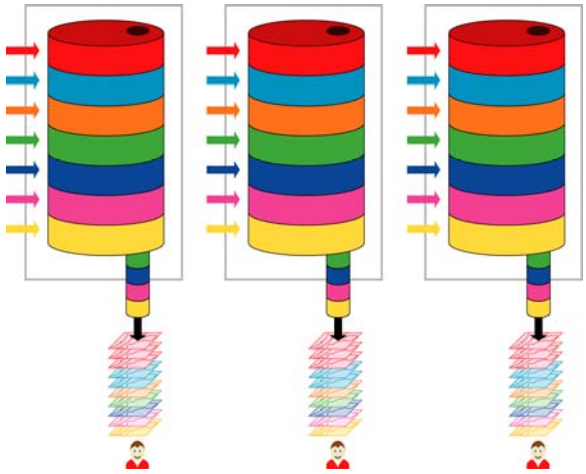

Figure 1: TIGGE Phase I

In Phase II, the TIGGE database will be distributed. A user will access the data through a single interface that connects all the different Data Providers. This is illustrated in Figure 2 where now the grey boxes could be more numerous Archive Centres and individual Data Providers that are willing to provide distributed access.

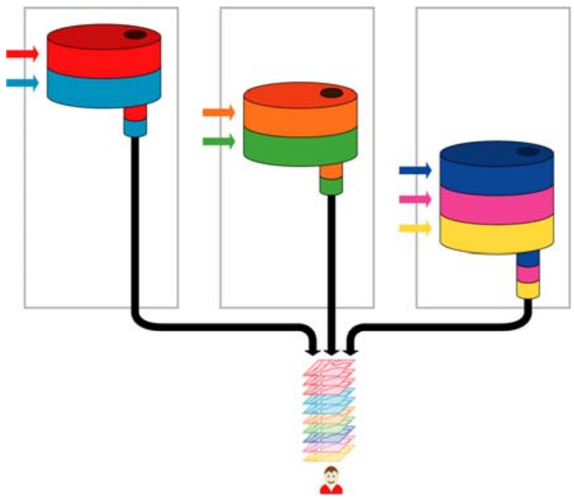

Figure 2: TIGGE Phase II

2.3. Timetable and milestones

Some of the major past and future milestones for the large TIGGE project are summarized below.

2005

- Completed basic planning (including costs and resource) to initiate Phase I,
- Secured commitments from CMA, ECMWF, and NCAR to act as Archive Centres,

- Secured initial commitments from Data Providers to submit data,

2006

- Agreed upon data exchange protocols and data formats,
- Finalised a list of parameters to store in the TIGGE archive,
- Implemented and tested Internet data transport,
- Several Phase I Archive Centres (ECMWF and NCAR) began collecting contributions in near-real time,

2007 and 2008

- Initiate user data access to the available data,
- Begin work on Phase II software developments (once required funding is secured),
- Expand user data access options to include subsetting, regridding and interpolation, collaborate with tool developers to improve user support,
- Insure TIGGE infrastructure is sufficient to provide support for THORPEX research in:
    - o International Polar Year field campaigns,
    - o Beijing 2008 Olympics.

3. TECHNICAL IMPLEMENTATION

Technical planning began at a meeting of a Working Group on Archiving (with representatives from CMA, ECMWF, NCAR and North American Ensemble Forecast System (NAEFS) Project) held at ECMWF 19-21 September 2005. These TIGGE Phase I technical plans were documented in a report[1] and it was presented to the members of the Implementation Meeting held at ECMWF 9-10 November 2005. This latter group represented both Archive Centres and Data Providers. They addressed the technical issues, refined certain aspects, and approved the plan to go forward with the implementation.

The remainder of this paper describes the rational behind the technical choices and outcomes for the implementation of TIGGE Phase I.

3.1. Homogeneity of the TIGGE database

For this project to succeed, it is paramount that the content of the TIGGE database be as homogeneous as possible. This will insure a productive environment that has systematic data management and user access to data from many provider centres. The more consistent the archive the easier it will be to develop applications.

The multi-model seasonal forecast project DEMETER is a successful example where the effort put into creating a homogenous archive led to a variety and some unforeseen useful applications.

3.1.1.Common terminology

All the partners must agree on a common way to reference data. Using common dictionaries all fields should be described with the same attributes (dates, level, step, parameter, etc.). This means common metadata descriptors will be used across all data providing partners. This promotes the most rapid data processing, creates uniform reference catalogues in the access portals, and TIGGE-wide accurate search and discovery capability at all Archive Centres.

3.1.2.Defining the core dataset

A field is uniquely identified within the TIGGE dataset by the following tuple:

(analysis date, analysis time, forecast time step, originating centre, ensemble number, level, parameter)

When using fields to create a "grand ensemble", i.e. when considering all members from several originating centres as a super ensemble, we must make sure that they share the same values for the tuple *(analysis date, analysis time, forecast time step, level, parameter)*.

As a result, a core dataset must be defined in terms of parameters, analysis times, forecast time steps and levels. It is critical that Data Providers adhere to the core dataset definition and use identical physical units on all parameters.

---

1 http://tigge.ecmwf.int/tigge/d/show_archive/table=documents/

All parameters are provided every 6 hours. All accumulations are computed from the beginning of the forecast. There are 26 standard surface and single level fields. Data provider are encourage but not required to provide all fields.

Parameter	Unit
Mean sea level pressure	Pa
Surface pressure	Pa
10 meter u-velocity	$m\ s^{-1}$
10 meter v-velocity	$m\ s^{-1}$
Surface air temperature	K
Surface air dew point temperature	K
Surface air max temperature	K
Surface air min temperature	K
Skin temperature	K
Soil moisture	$kg\ m^{-3}$
Soil temperature	K
Total precipitation (liquid + frozen)	$kg\ m^{-2}$
Snow fall water equivalent	$kg\ m^{-2}$
Snow depth water equivalent	$kg\ m^{-2}$
Total cloud cover	0-100%
Total column water	$kg\ m^{-2}$
Time-integrated surface latent heat flux	$W\ m^{-2}\ s$
Time-integrated surface sensible heat flux	$W\ m^{-2}\ s$
Time-integrated surface net solar radiation	$W\ m^{-2}\ s$
Time-integrated surface net thermal radiation	$W\ m^{-2}\ s$
Time-integrated outgoing long-wave radiation	$W\ m^{-2}\ s$
Sunshine duration	s
Convective available potential energy	$J\ kg^{-1}$
Convective inhibition	$J\ kg^{-1}$
Orography	gpm
Land-sea mask	Proportion

The vertical structure is defined by five parameters on eight pressure levels, i.e. 40 fields. The eight levels are 1000, 925, 850, 700, 500, 300, 250 and 200 hPa and in addition geopotential height is provided at ninth level, 50 hPa.

Parameter	Unit
Temperature	K
Geopotential height	gpm
U-velocity	$m\ s^{-1}$
V-velocity	$m\ s^{-1}$
Specific humidity	$kg\ kg^{-1}$

The following parameter is available on isentropic level Theta=320K.

Parameter	Unit
Potential vorticity	$K\ m^2\ kg^{-1}\ s^{-1}$

The following parameters are available on the potential vorticity level 2 PVU.

Parameter	Unit
Potential temperature	K
U-velocity	$m\ s^{-1}$
V-velocity	$m\ s^{-1}$

Grid and resolution

Data Providers can submit native grids and resolutions to the Archive Centres assuring a high resolution TIGGE database. In order to support multiple model comparisons each Data Providers will supply interpolation routines for conversion to regular latitude-longitude grids and for point extraction. Archive Centres may endeavour to return data in regular grids using these interpolation routines.

3.1.3. Common data format

GRIB edition 2 (GRIB2) was chosen as the standard TIGGE format, because it is the only WMO standard that supports ensemble data without the need of local extensions and also the chosen by the NAEFS community.

Data providers are required to make their data available in the archive format. This assures the accuracy of the TIGGE database at the Archive Centres, provisional data conversions elsewhere could introduce error.

There are several issues related to the use of GRIB2: there is little experience in using this format, the available documentation is ambiguous (e.g. there are three correct ways to describe “2m maximum temperature”) and all TIGGE parameters are not covered.

To address these issues, requests for clarifications and proposal for new parameters were submitted to the WMO Joint Meeting of CT-MTDCF/ET-DRC, Montreal, 8-12 May 2006. As a result, a substantial number of amendments were made to the Guide to the WMO Table Driven Code Form.

The Archive Centres have identified the list of GRIB2 codes, tables and templates to use for each of the fields of the TIGGE database and are providing guidelines (best practices) on how all TIGGE fields should be coded in GRIB2, as well as examples of properly encoded model output[2].

3.2. Data transfers

3.2.1. Network Bandwidth

It is thought that the available bandwidth between Europe and the USA is sufficient to meet the needs of TIGGE, whereas CMA

[2] http://tigge.ecmwf.int/tigge/d/show_archive/table=parameters/

raised concerns about the current bandwidth between China and Europe, as well as between China and the USA (the latter being probably better). The current bandwidth between China and the other partners currently appears to be much lower than projected as TIGGE requirements. Nevertheless, the situation may improve by the end of the year, and the Working Group is interacting with various international networking groups to investigate some of the possible network options (e.g. CTSNET, GLORIAD). This lack in bandwidth is a potential risk for this project.

Tests, using ordinary FTP, between NCAR and ECMWF have established that TCP tuning (matching window/buffer size) improves data transfer rates, but they also shows that a large number of parallel transfers (up to 50) are required in order to be able to meet the TIGGE requirements of exchanging several hundreds of GB daily across the Internet.

3.2.2. Choice of transport mechanism: IDD/LDM

After extensive testing, it was decided that IDD/LDM (Internet Data Distribution system, Local Data Manager), an Internet based distribution system, was will suited for TIGGE, in particular it supports the parallel transfers needed to exchange the large volumes of data. Furthermore, UNIDATA has setup a special support for the TIGGE project.

This system is fundamentally designed for the distribution of real-time data, and is already running operationally in the UNIDATA community at roughly three hundred sites worldwide. It is thought to be scalable, and provides built-in mechanisms for monitoring and providing statistics on data transfer activity.
IDD/LDM is also being employed operationally in the CONDUIT project to provide distribution services for large-volume USA/NCEP forecast products.

LDM is a broadcast system, based on subscription. Every LDM instances has a queue. A “downstream” LDM can subscribe to “products” from an “upstream” LDM by specifying a regular expression for the products names that are desired. When a product is inserted in the upstream LDM, it is automatically transferred the queues of all the downstream LDMs that have subscribed to this product. The downstream LDM then extracts the products from the queue and write them to disk at designated locations. LDM queues have a fixed size and inserting new products will result in removing older products from the queue.

One of the downside of LDM is that products may not be received by the downstream LDM in the following scenarios:

- if the upstream LDM inserts products in its queue faster than the network can transport them to the downstream queues,
- if the downstream LDM cannot extract the product from its queue fast enough,
- if the downstream LDM is down for a period of time, and the upstream LDM continues to insert product in its queue.

To overcome these potential problems, a protocol has been defined and tested for the TIGGE LDM implementation. The protocol uses a data file manifest that identifies all files to be transferred. Actual file transfer is verified against the manifest and retransmission is requested if any files that fail to arrive at the downstream location. A complete description of the protocol and samples are available on the TIGGE web site (http://tigge.ecmwf.int).

The following figure shows the data transfer magnitude versus time typically achieved between ECMWF and NCAR, i.e. about 10 GB/hour.

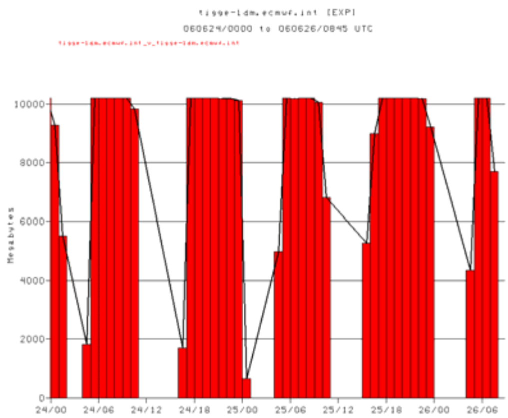

3.3. *Organisation of the collaboration*

TIGGE is a collaborative project with the focal point for data exchange at the Archive Centres during Phase I. The success of TIGGE is directly linked to the degree of commitment of the Data Providers, and the ability of these partners and Centres to work together.

As the Archive Centres will have a global view of the data production, it is proposed that they provide the project technical coordination and take on the responsibility of defining the necessary procedures.

From the Archive Centre's view the following are important aspects for successful operations.

3.3.1.General organisation

Each Data Provider will nominate two contact points:

- a technical contact point, which will be able to address operational and technical issues, such as troubleshooting, networking or timeliness of data delivery,
- a scientific contact point, which will be able to address issues such as forecast performances or numerical errors.

When these contact points are unavailable the partners must nominate alternate contact persons.

The communication will be established through a series of mailing lists, collaborative tools (e.g., Wiki) and a web site, which ECMWF has offered to host and will be mutually cross linked to the other Archive Centres.

3.3.2.Issue of completeness

The objective is to have 100% complete data at the Archive Centres.

Completeness may not be achieved for two reasons:

- the transfer of the data from the Providers to the Archive Centre fails,
- operational activities at the Providers centre are interrupted and back filling past runs is impractical.

Unfortunately, an incomplete dataset is often difficult to use. Most of the current tools used for ensemble data assume a fixed number of members from day to day. These tools will have to be rewritten to be more flexible.

It is recommended that Data Providers endeavour to send missing data, whenever possible, to the Archive Centres, even if this means rerunning a forecast cycle.

3.3.3.Operations

Day to day operations

Tools must be built to monitor the data transfer from all Data Providers within the system. Transfer statistics are required to quantify variations in the performance of the Internet and to enable early detection of anomalies that indicate problems.

Each Archive Centre will set up a web page showing volumes, date of data and date of reception for each Data Provider. This information will be used to cross-validate the content at the three archives.

When problems arise that prevent data delivery to the Archive Centres, the Data Provider will be responsible to notify all the Archive Centres, e.g. by sending an email to the appropriate TIGGE mailing list.

When an Archive Centre does not receive the expected data from a Data Provider, or if the data are incomplete or corrupted, it will first check with other Archive Centres and determine if the failure is an isolated case. If it is an isolated case recovery will be initiated between Archive Centres, if not the Data Provider must re-initiate the data delivery. In any case, the incidents must be investigated and documented. The use of a trouble ticketing system will be investigated to facilitate tracking problems.

Long term operations

The Archive Centres have agreed to define and collect common metrics that can be used to create combined TIGGE-wide reports. This information will be used for future evolutions of the system. Participation in TIGGE must not interfere with the operational activities of Data Providers, i.e. they should be able to upgrade models, introduce higher resolutions, and make all customary changes as needed. Mechanisms should exist that allow new products from the Providers to be easily integrated into the TIGGE Archive Centres. These procedures need to be established and will include ways to test delivery of new products and will likely require version number control, to name just a few features.

On the other hand, Data Providers must take into account their participation into TIGGE when planning changes to their forecasting systems, and must inform Archive Centres accordingly.

3.4. *User access: Data retrieval*

The possibility that each Archive Centre would provide an identical retrieval interface to the TIGGE database was considered, but it was established that such a unified interface was not possible without significant development and therefore is not part of the Phase I implementation.

However, the Archive Centre will guarantee that user interfaces will present the same information (e.g. same variable names), and that similar requests, although expressed differently, should return identical results.

ECMWF will utilise the MARS system initially, and NCAR will build upon its Research Data Archive and Community Data Portal environments in order to serve their respective user communities. CMA is still developing their processes for data delivery. Over time and with additional project support, it is expected that there will be opportunities to further unify the user interface by leveraging developments from the WMO Information System (WIS) effort.

4. CONCLUSIONS

The success of this project will depend greatly on:

- the commitment of each data provider,
- the establishment of a collaboration methods (email, web sites,…),
- the availability of sufficient network bandwidth,
- the homogeneity of the catalogue,
- the adherence to WMO standards (GRIB2) as defined for TIGGE.

Most of the issues and tasks identified by the members of the first technical meeting held in September 2005 have been tackled:

- Definition of the database content (list of variables),
- Choice of a data format (GRIB2),
- Choice of a transport mechanism (LDM),
- Creation of collaboration tools (email, web sites,…).

On the 1st of October 2006, after several months of exchanges of test data, the Archive Centres have started to routinely archive data from three Data Providers. The remaining Data Providers are actively working to make their data available by the end of the year.

The next challenge will be to provide end users with efficient and easy to use tools to access the database.