Nat. Hazards Earth Syst. Sci. Discuss.,
doi:10.5194/nhess-2016-32-AC1, 2016

Natural Hazards and Earth System Sciences
Discussions

Interactive comment on "VISIR: Technological infrastructure of an operational service for safe and efficient navigation in the Mediterranean Sea" *by* G. Mannarini et al.

G. Mannarini et al.

gianandrea.mannarini@cmcc.it

Received and published: 23 June 2016

(In the following "NHESSD" stands for the paper and under Review. When not specified, all other references to equations, figures, and tables are relative to the present document.)

General comments

The paper describes a decision support system (DSS) for on-demand computation of optimal ship routes in the Mediterranean sea. The focus of the paper is not on the mod-

elling part but on the operational infrastructure and how the set of software components are brought together to realise the DSS. In this sense, the paper clearly describes the objective, role and scope of each software component and how this relates to each other. Such technical description is later accompanied by a couple of examples to demonstrate the operational functioning of the system. Even though some of the client apps devised by the authors are in progress, it is quite clear that the core contribution of the paper is not these client application but the multi-layered architecture that supports asynchronous computation of optimal routes (the model itself). Developed web services adhere to well-known web standards and best practices. [..] In summary I found the paper quite interesting and technically sound. The topic of the paper, and the design and architecture of the system, is acceptable.

–Authors' response:
The Referee correctly focused on the manuscript's scope and we are glad to see that our effort to adhere to the state-of-the-art of technology was acknowledged.

As the paper focuses on the overall functioning of the system, and how the distinct pieces are connected, more details on the assessment methods and performance evaluation are required to truly demonstrate the system..

–Authors' response:
See specific comment **B** below.

Specific comments

A - *The heart of the system seems to be the message-based broker because, to this reviewer, it is intended to make the system more scalable (not mentioned by authors).*

–Authors' response:
The Message Broker (MB) certainly has a central position within the TESSA architecture, in the sense that it brokers between the client-oriented services and the backed model computations. However, each individual VISIR system component owns a given level of scalability; not just the MB, but also the CDAM Gateway and the cluster infrastructure that are managed by the slurm scheduler.

–Authors' changes to manuscript:
A1.
A1) On P5, row 139, to insert:
"Each individual VISIR system component owns a given level of scalability; not just the MB, but also the CDAM Gateway and the cluster infrastructure that are managed by the slurm scheduler."

B - *As there are so many pieces connected acting asynchronously, I am wondering about performance issues of the system. These aspects are not presented or discussed in the paper, so I encourage the authors to specify limitations or performance penalties the system has, either by creating a new section to this matter or by extending sect 4.1 and 4.2 accordingly.*

–Authors' response:
We thank the Referee for this interesting observation. It will lead to a significant improvement of the manuscript.

We actually already provided some hint about the performance of the VISIR operational system in Tab.3 and Fig.4 of NHESSD. However, following the Referee's question, we made new extensive experimental tests for augmenting that information, see "Authors' changes to manuscript" below.

–Authors' changes to manuscript:
B1, B2, B3, B4, B5, B6, B7, B8, B9, B10, B11, B12, B13, B14.

B1) In Abstract, starting from row 6, replace with:
"This paper focuses on the technological infrastructure developed for operating VISIR as a DSS. Its main components are described, the performance of the operational system is assessed through experimental measurements, and a few case studies are presented."

B2) In Introduction, starting from P2 row 53, replace with:
"A detailed presentation of the operational system follows in Sect.3, and its functioning and performance are addressed in Sect.4. A few case studies are discussed in Sect.5, before Conclusions are drawn in Sect.6."

B3) On P9, edit Sect. 3.4 "Execution logic" starting from row 259 as follows:
"Once it is forwarded through the DSS-portlet (3') to the Message Broker (MB), a correlation-ID is provided back to the DSS-portlet (4a) and the client application (4b). At the same time, the MB forwards the request to the CDAM (4c). While the VISIR model is run on the supercomputing facility controlled by the CDAM (6 a-c), a polling activity starts at the client (5a) and is forwarded down to the level of the MB (5a'). Such a querying is iterated until the model results from the CDAM are sent to the MB (7). They are immediately propagated up to the level of the DSS-portlet (8) and, finally, to the client (9)."

B4) Rename NHESSD Section 3 into "Operational infrastructure"

B5) Define a new Section 4 "Functioning of the operational system" and renumber NHESSD Sections 4 and 5 to Sections 5 and 6 respectively.

B6) Renumber Sect. 3.4 to Sect. 4.1 "Execution logic" and Sect. 3.5 to Sect. 4.3 "Use of the system"

B7) On P9, insert a new Sect. 4.2 "System Performance" as follows:
"As outlined in the previous subsections, the TESSA architecture consists of various connected components acting asynchronously.

In order to asses its performance, it is convenient to consider it as a network whose nodes are linked in a directional way, as in Fig.2. The network is characterised by delays along the links and by node internal waiting times. Some logical steps of the flux diagram of Fig.1 can be remapped to the network of Fig.2 via the correspondence provided in Tab.1. The nodes of Fig.2 were chosen because they correspond to the places where it is presently possible to perform time measurements. The names of the nodes correspond to the name of the TESSA tier (cp. Fig.1 of NHESSD) where measurements are taken.

The performance of the VISIR operational system can then be assessed by comparing time measurements at various nodes. The experimental activity was carried out ensuring that the computational cluster was nearly idle. Computations of routes up to 500 NM were requested from the UI. For each given route length, the same departure and arrival were employed in up to 10 routing jobs during test sessions occuring on three different calendar dates. The experimental protocol also included recording

time-stamps at the various nodes of the operational infrastructure of Fig.2.

In Fig.3 we display data collected for both the motorboat (panels a,b) and sailboat (panels c,d) VISIR jobs. In the left panels we display the "waiting times" $\eta(Q) = t(Q_2) - t(Q_1)$, with $Q = \{U, S, C\}$ (see Fig.2), and $\tau(m)$, the duration of the matlab job. The measurement of the $\eta(Q)$ times, being differences of absolute times at a given tier of the TESSA infrastructure (cp. Fig.2 and Tab.1), does not require clock synchronization among the asynchronous system tiers.

For longer routes, $\eta(Q)$ is dominated by $\tau(m)$, that, with respect to route length L, scales as a power law with an exponent of about 3 (cp. Tab.2). The relation between L and the number N of grid points in the selected bounding box (see Sect.3.5 of NHESSD) depends on the aspect-ratio of the latter. For a squared box, $L \sim \sqrt{N}$, for a more elongated one, $L \sim N$. Consequently, the fitted exponent also implies a power-law dependence on N with an exponent comprised between 1.5 and 3. This is compliant with the fact that a quadratic trend with respect to N was fitted in (Mannarini et al., 2016). This in turn agrees with the theoretical performance of Dijkstra's algorithm (Sect.2.2 of NHESSD) in its implementation without any specific data structures. Thus, the performance of the VISIR operational system for long routes mirrors that of the model for ship route optimization.

For shorter routes however, it is apparent that the values of $\eta(Q)$ at the various system layers differ by several seconds. In order to better explore the performance of the VISIR operational system, we subtract the $\tau(m)$ dominant contribution, defining the excess times $\delta(Q) = \eta(Q) - \tau(m)$. Furthermore, we display $\delta(Q)$ vs. a proxy of the size of the payload transferred from the CDAM up to the level of the UI. This corresponds to the sum of the two files containing the geodetic and the optimal route. This way, the pictures displayed in the right panels of Fig.3 are obtained for the motorboat and sailboat cases respectively.

Starting from the CDAM δ data (blue markers and lines), a nearly constant trend is

observed for both motor- and sailboat tests. We find $\delta(C) \approx 2$ seconds. Between the C_2 and C_1 nodes two distinct processes occur on the CDAM: the matlab job is submitted via slurm and, upon completion, the results are uploaded using the callback URL mechanism (cp. Sect.3.2 of NHESSD) for making them available to the SSA platform. In order to assess the contribution from the startup submission time of a slurm/matlab job, we separately tested that the duration of a slurm task that just submits a void matlab job is between 1.4 and 2.0 seconds. Most of the $\delta(C)$ time should be ascribed to this irreducible slurm duration. $\delta(C)$ is independent of payload size, as a consequence of the high-speed internet connection of the computing facility hosting the CDAM.

For the SSA δ data (red markers and lines), again a nearly constant trend is observed for both motor- and sailboat job tests. We find $\delta(S) \approx 3$ seconds. Despite the name (arising from the TESSA tier where the measurements are carried out), $\delta(S)$ contains, besides the SSA-CDAM bidirectional delays, mainly CDAM tasks, and namely: the preliminary validation of the job request parameters (occuring at the CDAM Gateway), the setup of the remote environment on the supercomputing facility (CDAM Gateway), the namelist production (CDAM Launcher), the processes leading to $\delta(C)$ (CDAM Launcher), the creation of an output json file containing both the geodetic and the optimal route. With the exception of the first and the last, these processes are independent of payload size and, thanks to the fair internet bandwidth available at the SSA platform location, they result in a nearly constant $\delta(S) - \delta(C) \approx 1$ sec.

Moving to the UI δ data (black markers and lines), we note first of all that, for a given route length, the sailboat payload is smaller than the motorboat one. This is due to the fact that a smaller number of fields is stored in the sailboat payload (e.g. the flags of the safety constraints are just available for the motorboat case, see Sect.2.3 of NHESSD). Furthermore, the UI measurements are characterized by a large variability (error bar size) among the different tests. Both motorboat and sailboat data show a linear trend with respect to route length. The intercept of the linear fits (e parameter in Tab.2) is

located at about 7 seconds. In order to understand these results, it should be recalled that $\delta(U) - \delta(S)$ includes, besides the UI-SSA bidirectional delays, the waiting time $\tau(p)$ due to the polling interval t_{poll} of the client application (Sect.3.4 of NHESSD). The polling mechanism can be considered as a normally closed gate, opening instaneously any t_{poll} time units. Thus, a signal arriving at the gate at a time σ cannot pass before the additional delay $\tau(p)$, given by:

$$\tau(p) = t_{\mathsf{poll}} - \mathsf{mod}(\sigma, t_{\mathsf{poll}})\,, \tag{1}$$

is elapsed. Eq.1 defines a piece-wise linear function of σ. In the VISIR case σ includes both deterministic and stochastic (due to internet bandwidth and computing resources availability) processes, as described above. Thus, it is convenient to estimate the average of $\tau(p)$, given by

$$\overline{\tau(p)} = \langle \tau(p) \rangle_\sigma = \frac{1}{2}\, t_{\mathsf{poll}} \tag{2}$$

In VISIR case, $t_{\mathsf{poll}} = 5$ sec. Thus, we find that, on the average, it must be $\delta(U) - \delta(S) \geq t_{\mathsf{poll}}/2 = 2.5$ sec. For this reason, we display in Fig.3b,d the line $\delta(S) + t_{\mathsf{poll}}/2$ (dashed grey line) and, as expected, we find that $\delta(U)$ is never smaller than this reference.

Finally, we discuss the slope found for the linear fits of $\delta(U)$ (Fig.3b,d). We note that it is different for the motorboat and sailboat cases (Tab.2). We recall that the δ signals displayed are not affected by the VISIR model internal time $\tau(m)$, that in fact differ between motorboat and sailboat computations, as shown by the values in Tab.2. Thus, the reason for the different slopes must be searched either in the internet connection or in some UI processing. The error bars in Fig.3b,d and the uncertainty in the fitted parameters (Tab.2) confirm a large variability of the internet speed during the tests. However, there are also specific data checks and rendering peculiarities in the linechart (Fig.6b,c of NHESSD) of the VISIR sailboat interface. Thus, at the moment we cannot rule out any of the possible explanations. Both are likely to be at play, though other tests and measurements would be needed for assessing their relative weight.

Ultimate performance limit of VISIR
The experimental findings above and their analysis enable us to assess the ultimate performance limit of the VISIR technological infrastructure.

There are limiting factors of two types: on the one hand, there are parameters (such as t_{poll} and the slurm duration for sumbitting a job) that could be tuned for obtaining some minor improvement (less than 5 seconds in the best case). On the other hand, there are larger penalties paid to the present architecture representing the actual bottlenecks of performance, namely:

a) the computational cost for generating the ship routes;

b) the internet-based comunication among asyncronous system components.

In particular, a) dominates the UI waiting time for long routes (cp. panels a,c in Fig.3) while b) is the limiting factor for the duration of the short routes.

Fixing item (a) corresponds to a situation where there are dedicated computational resources and negligible competition within the supercomputing facility and the performance of the computing algorithm has significantly improved. It can be simulated by subtracting the duration $\tau(m)$ of the matlab job, as we did in Fig.3b,d.

If further item (b) were fixed by a fast link (optical fiber or so), the ultimate limit for the VISIR user waiting time would be given by $e + t_{\text{poll}}/2$, where e is the intercept of $\delta(S)$ (cp. Tab.2). Using current values for both SSA e and t_{poll}, this would imply a total waiting time of about 6 seconds after a route request is submitted from the user interface."

B8) On P12, row 370 insert:
"We analyzed in detail the performance of the infrastructure enabling VISIR operations (Sect 4.2). In general, the CDAM and the SSA platform appear to communicate quite

efficiently, while the client internet connectivity may act as a bottleneck. Though this can easily be improved for terrestrial accesses, it could represent an even more severe issue in case of offshore use of VISIR. This is why a caching system for the imagery of the environmental fields could be developed for using the system in case of poor internet connectivity. The other bottleneck outlined in Sect.4.2 is the performance of the ship routing model and it is already being addressed by an ongoing research activity at CMCC. The ultimate performance limit of the present architecture is a user waiting time of about 6 seconds."

B9) On P12, row 379 remove:
"Also, a caching system for the imagery of the environmental fields could be a viable solution for using the system in case of poor internet connectivity."

B10) replace Fig.3 of NHESSD with Fig.1 of this comment.

B11) replace Fig.4 of NHESSD with Fig.2 ("Graph Structure..") of this comment.

B12) replace Fig.5 of NHESSD with Fig.3 ("Performance of VISIR operational system..") of this comment.

B13) replace Tab.2 of NHESSD with Tab.1 of this comment.

B14) replace Tab.3 of NHESSD with Tab.2 of this comment.

C - *Please also revise the citation style of Mannarini et al. (2015) in the body of the*

paper because I was not able to distinguish which one (of two in the reference list) the authors actually referred to.

–Authors' response:
Surely we'll do this. In the meanwhile, the GMD-Discussions manuscript evolved into the peer-reviewed paper (Mannarini et al., 2016) and the accepted paper on sailboats has been published (Mannarini et al., 2015).

–Authors' changes to manuscript:
C1, C2.
C1) Using reference to (Mannarini et al., 2016)
C2) Using reference to (Mannarini et al., 2015)

Table 1. Correspondence between nodes of the VISIR system (cp. Fig.1) and steps within the Data flux diagram (Fig.3 of NHESSD). The step part (s: start, e: end) is also given when relevant ("-" otherwise).

	U_0	U_1	S_1	C_1	m	C_2	S_2	p	U_2
Data flux diagram (Fig.3 of NHESSD) step #	1	5a	4c	6a	6b	6c	7	5a-8	9
step part	s	s	s	s	-	e	e	s-e	e

References

Flannery, B. P., Press, W., Teukolsky, S., and Vetterling, W. T.: Numerical Recipes in FORTRAN 77: The Art of Scientific Computing, 1992.

Mannarini, G., Lecci, R., and Coppini, G.: Introducing sailboats into ship routing system VISIR, in: Information, Intelligence, Systems and Applications (IISA), 2015 6th International Conference on, pp. 1–6, IEEEXplore, doi:10.1109/IISA.2015.7387962, 2015.

Mannarini, G., Pinardi, N., Coppini, G., Oddo, P., and Iafrati, A.: VISIR-I: small vessels – least-time nautical routes using wave forecasts, Geoscientific Model Development, 9, 1597–1625, doi:10.5194/gmd-9-1597-2016, http://www.geosci-model-dev.net/9/1597/2016/, 2016.

Interactive comment on Nat. Hazards Earth Syst. Sci. Discuss., doi:10.5194/nhess-2016-32, 2016.

Table 2. Parameters and scores for the least-square fits shown in Fig.3. For the η and τ signals (a,c panels of Fig.3) all fits are of type $ax^b + c$, while for the δ signals (b,d panels of Fig.3) the fits are of type $dx + e$. The b exponent is given with the 95% uncertainty bounds, while the e intercept is given with the estimated variance (among brackets). Least-χ^2 linear fits (Fig.3b,d) employ formulas provided in (Flannery et al., 1992, Chap.15) for the coefficients of the linear regression and their uncertainties.

			[units]	motorboat	sailboat
UI	η	a	[s·NM^{-b}]	$3.9 \cdot 10^{-7}$	$2.4 \cdot 10^{-5}$
		b	[-]	3.3 (2.6 - 3.9)	3.0 (2.5 - 3.5)
		c	[s]	19.8	14.1
		R^2	[%]	98.5	99.5
	δ	d	[kB/s]	90(100)	30(18)
		e	[s]	7.1(1.2)	6.3(0.7)
		R^2	[%]	80.6	59.3
SSA	η	a	[s·NM^{-b}]	$9.3 \cdot 10^{-7}$	$1.8 \cdot 10^{-6}$
		b	[-]	3.1 (2.6 - 3.7)	3.0 (2.5 - 3.5)
		c	[s]	14.5	10.5
		R^2	[%]	98.3	99.6
	δ	e	[s]	3.2(0.2)	3.1(0.2)
		R^2	[%]	74.1	94.2
CDAM	η	a	[s·NM^{-b}]	$1.1 \cdot 10^{-6}$	$2.0 \cdot 10^{-6}$
		b	[-]	3.1 (2.6 - 3.7)	3.0 (2.5 - 3.5)
		c	[s]	12.6	8.6
		R^2	[%]	99.6	99.6
	δ	e	[s]	2.1(0.1)	2.1(0.1)
		R^2	[%]	98.7	94.8
matlab	τ	a	[s·NM^{-b}]	$1.7 \cdot 10^{-6}$	$3.7 \cdot 10^{-6}$
		b	[-]	3.0 (2.5 - 3.6)	2.9 (2.4 - 3.4)
		c	[s]	9.3	5.5
		R^2	[%]	97.9	99.5

C13

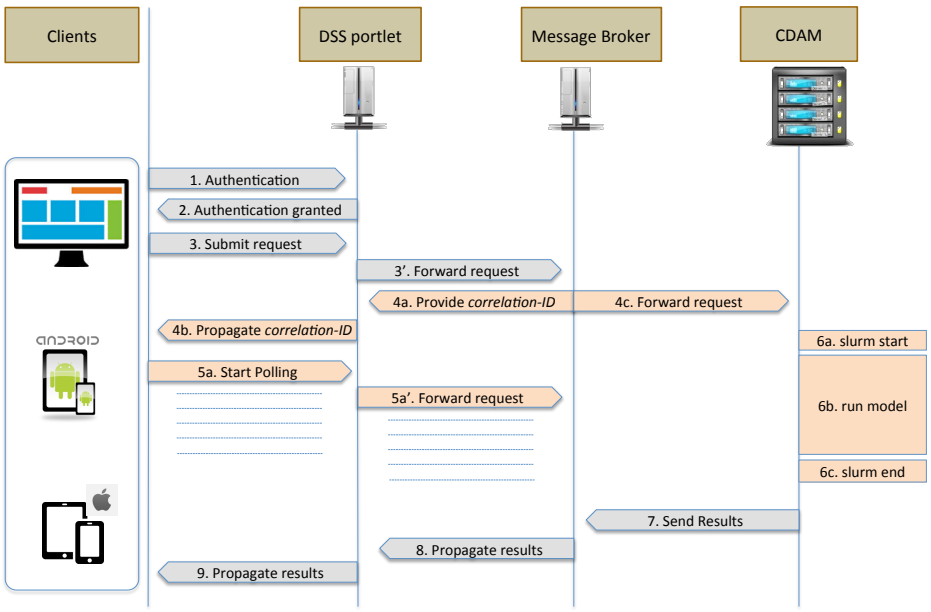

Fig. 1. Data flux diagram for the operational system of VISIR.

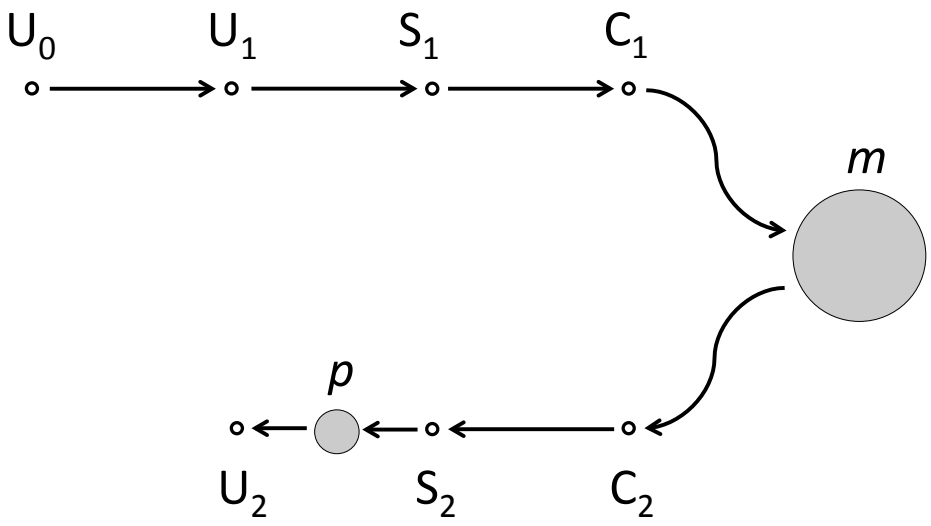

Fig. 2. Graph structure of the TESSA system for running DSSs like VISIR. The nodes refer to points of time measurements.

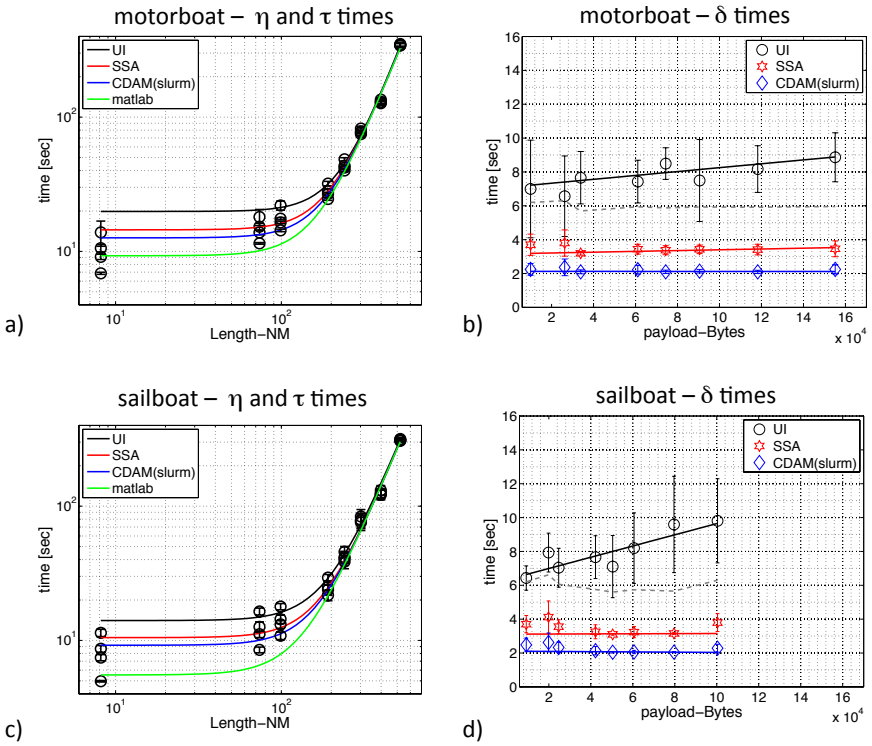

Fig. 3. Performance of VISIR operational system, distinguishing between motorboat (a,b) and sailboat mode (c,d).