

# Evaluation of the parallel performance of the TRIGRS v2.1 model for rainfall-induced landslides

M. Alvioli<sup>\*1</sup>, D. Spiga<sup>2</sup>, R. L. Baum<sup>3</sup>

<sup>1</sup>CNR IRPI, via Madonna Alta 126, 06128, Perugia, Italy

<sup>2</sup>INFN, Sezione di Perugia, via A. Pascoli 1, 06123, Perugia, Italy

<sup>3</sup>US Geological Survey, P.O. Box 25046, Mail Stop 966, Denver, CO 80225-0046, USA

\* Massimiliano.Alvioli@irpi.cnr.it

The widespread availability of high resolution digital elevation models (DEM) opens the possibility of applying physically based models of landslide initiation to large areas. With increasing size of the study area and resolution of the DEM, the required computing time for each run of the model increases proportionally to the number of grid cells in the study area.

The aim of this work is to present a new parallel implementation of TRIGRS (Alvioli and Baum, 2016), an open-source FORTRAN program (software available for download at <http://geomorphology.irpi.cnr.it/tools/trigrs> and <https://github.com/usgs/landslides-trigrs>) designed for modeling the timing and distribution of shallow, rainfall-induced landslides, and to discuss its parallel performance. The spatial distribution of landslides is obtained in TRIGRS by computing transient pore-pressure changes, and attendant changes in the factor of safety due to rainfall infiltration, and using a simple infinite-slope description on a cell-by-cell basis. Time dependence is implemented in the model by time-dependent rainfall infiltration, resulting from storms that have durations ranging from hours to a few days.

The motivation for having a parallel, and substantially faster, code is generally found in the need to perform multiple simulations, for example to calibrate parameters, or to increase the spatial extent of simulations, or both (Alvioli *et al.*, 2014; Raia *et al.*, 2014; Mergili *et al.*, 2014). In the specific case of TRIGRS, we have shown that the simulation of a large test area is possible with great detail in many respects: the high resolution DEM, the dense rainfall pattern in time, the number of output maps. Using the new version of the code, we were able to reduce run time from about one day to about an hour for a 1340-km<sup>2</sup> test area. This allowed us to obtain a detailed time series of model outputs and match them with the field observations, providing a robust understanding of landslide phenomena in the area (Alvioli and Baum, 2016).

The previous version of the code (Baum *et al.*, 2008) was improved with new features (Iverson, 2000; Baum *et al.*, 2010) and parallelized within the Message Passing Interface (MPI) framework (MPI forum, 2012). We show the performance gain with respect to the serial code on commercial hardware (CNR IRPI, Perugia), on a high-performance multi-node machine (Galileo cluster, CINECA, Bologna) and on the OpenStack cloud environment (CERN, Geneva). The performance of the parallel code was assessed both measuring overall running time  $T_{Tot}$  and using standard quantities such as speedup

$$S(N_p) = \frac{T_{Tot}(N_p=1)}{T_{Tot}(N_p)}, \quad (1)$$

46

47 the ratio of serial ( $T_{Tot}(N=1)$ ) to parallel ( $T_{Tot}(N)$ ) running times as a function of the number of  
 48 processes  $N_p$ , and efficiency

49

50

$$E(N_p) = \frac{S(N_p)}{N_p} = \frac{T_{Tot}(N_p=1)}{N_p T_{Tot}(N_p)}. \quad (2)$$

51

52


53

54

55

56

For a better assessment of the parallel performance, we have split the overall running time into three terms: the time required for data reading and scattering to all the computing processes,  $T_R$ , the time required for core computing and gathering of the partial results from each computing process,  $T_C$ , and time required from writing the results on disk,  $T_W$ ; this is summarized in Fig. 1. The other quantities we have considered, speedup  $S$  and efficiency  $E$ ,



**Fig. 1.** Decomposition of the overall running time  $T_{Tot}$  into three intervals, as described in the text.

57

58

59

60

61

62

63

64

65

66

67

68

69

70

71

have been calculated according Eq. (1) and (2), respectively, considering separately the overall running time  $T_{Tot}$  and replacing  $T_{Tot}$  with the partial running time  $T = T_{Tot} - T_W$ . The corresponding results are shown in Fig. 2. When run on a single-mode, high-end multi-core machine, the code offers performance gain up to about  $N_p = 50$ , for our case study. The major limitation for a further performance gain are the operations summarized with  $T_R$  and  $T_W$ , the data read and scatter operations presenting the most severe drawback when  $N_p$  increases over 50. When run on a multi-node, high-performance machine, the limitations due to read-write and data scatter-gather operations are largely removed and we do not see performance degradation within the  $0 < N_p \leq 256$  range. We conclude that the parallel performance is very satisfactory, since the code provides sufficient performance improvement on a single node machine with a number of computing cores which is well beyond the typical number of cores available on common desktop machines. A further performance increase can be obtained running the code on a large machine equipped with high-performance input-output devices and high-speed interconnection network such as the Galileo cluster.

72

73

74

75

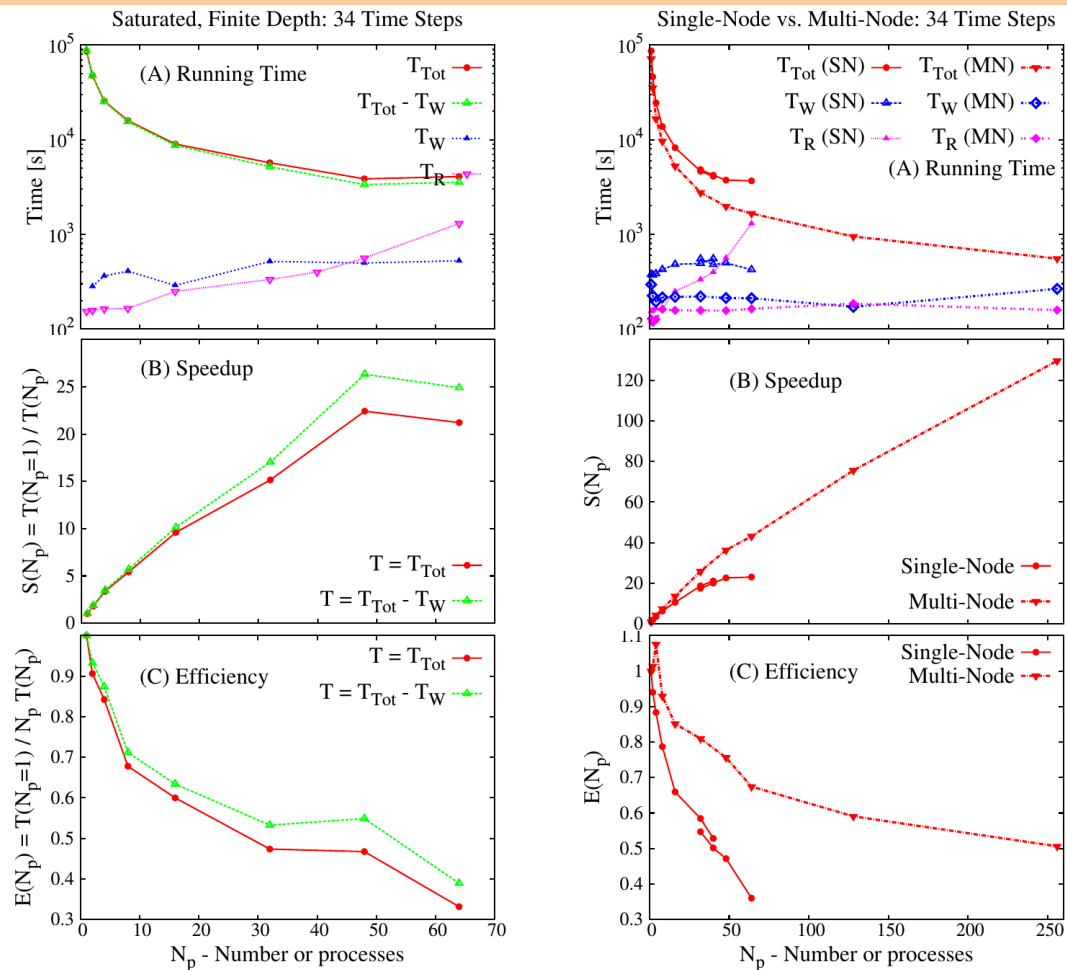
76

77

78


79

We have investigated the dependence of the parallel performance upon the problem size, which, in turn, is practically linearly proportional to the number of grid cells. We show in Fig. 3 the total computing time required for running the code with the study area used in Fig. 2, and also with study areas of twice, one-half, one-fourth, one-eighth and one-sixteenth of the original size. We find, as expected, that performance degradation when running on our single-node machine occurs for smaller  $N_p$  in the larger problem with respect to the original one, and we do not reach the point at which performance starts getting poorer for the smaller-size cases.



**Fig. 2.** Results for the parallel performance of TRIGRS v2.1, as a function of the number of processes in the MPI pool,  $N_p$ . A: total ( $T_{Tot}$ ), read + data scattering + compute + data gathering ( $T_{Tot} - T_W$ ), write ( $T_W$ ) and read + data scattering ( $T_R$ ) running times. B: parallel speedup with respect to  $T_{Tot}$  and  $T_{Tot} - T_W$  running times. C: parallel efficiency with respect to the same partial running times as in B. The left column refers to execution on a single-node machine equipped with 64 computing cores, while the right column shows results of execution on the Galileo cluster at CINECA; we considered  $N_p \leq 256$ ; the results of the left column are also shown for comparison. The 34 time-steps referred to in the Figures are the number of time intervals of the storm considered in the simulation.


80



**Fig. 3.** Dependence of the total running time on the size of the study area. “1” refers to the problem size considered in Fig. 2, “2” refers to a problem with double size, “1/2” to one-half the size, and so on.

81 Subsequently, we have investigated the parallel performance of the code on a typical cloud  
 82 environment setup. In particular, we have used virtual machines with sixteen computing cores  
 83 each, using the OpenStack cloud environment installed at CERN. Results for the total and  
 84 read-scatter running times are shown in Fig. 4, compared to the ones obtained on the single-  
 85 node and the multi-node machines considered earlier. We find that the performance gain is  
 86 comparable to the single-node case when the code is run within one node. We find severe  
 87 limitations mostly due to the read-scatter fraction when the code is run across multiple nodes  
 88 in OpenStack: the performance appears to be degraded and depending on the network status  
 89 in each particular run, as expected.

90



**Fig. 4.** Total (A,  $T_{Tot}$ ) and read-scatter (B,  $T_R$ ) running times obtained on OpenStack cloud environment compared to the ones obtained on the single-node and multi-node machines considered in Fig. 2.

91 In conclusion, users of the code can run the new parallel version using any kind of multi-core  
 92 machine equipped with a modern FORTRAN compiler and MPI libraries. We have optimized  
 93 both the serial and parallel part of the code in a way that aims at improving performance  
 94 without requiring extra expertise from the users while keeping a simple, easy-to-maintain  
 95 code. In particular, we did not make use of native MPI I/O (Prabhat and Koziol, 2014), which  
 96 will require major code design revision, extra system requirements and user expertise; we  
 97 leave the MPI I/O implementation for a future release of the code. Performance results show a  
 98 nice balance between the point at which file input and output dominates parallel speedup and  
 99 the great performance improvement when using a high-performance machine. We show that  
 100 the performance on a general-purpose cloud environment might be limited by network  
 101 communications between the computing nodes, which should be tuned for this specific  
 102 purpose. Most users would find the possibility of running the code on demand on MPI-  
 103 optimized cloud resources more appealing than on massively multi-core machines.  
 104 Availability of a parallel version of TRIGRS enables simulation of landslide initiation from  
 105 storm events affecting large heterogeneous regions. This greatly aids analysis of the effects of  
 106 rainfall patterns and terrain complexity on landslide initiation.

107

## 108 Acknowledgments

109

110 M. Alvioli acknowledges the CINECA award under the ISCRA initiative, for the availability  
 111 of high performance computing resources and support. M. Alvioli thanks CNR for a “Short  
 112 Mobility Grant”, 2014, during which part of this work was completed.

113

114

115

116 **Essential bibliography**

117

118 Alvioli, M., Guzzetti, F., Rossi, M., 2014. Scaling properties of rainfall-induced landslides  
119 predicted by a physically based model. *Geomorphology* 213, 38-47;  
120 doi:10.1016/j.geomorph.2013.12.039

121 Alvioli, M. and Baum, R.L., 2016. Parallelization of the TRIGRS model for rainfall-induced  
122 landslides using the message passing interface. *Environ. Modell. Softw.* 81, 122-135;  
123 doi:10.1016/j.envsoft.2016.04.002

124 Baum, R.L., Savage, W.Z., and Godt, J.W., 2008, TRIGRS—A Fortran program for transient  
125 rainfall infiltration and grid-based regional slope-stability analysis, version 2.0: U.S.  
126 Geological Survey Open-File Report, 2008-1159, 75 p;  
127 URL:<http://pubs.usgs.gov/of/2008/1159/>

128 Baum, R.L., Godt, J.W., Savage, W.Z., 2010. Estimating the timing and location of shallow  
129 rainfall-induced landslides using a model for transient, unsaturated infiltration. *J.*  
130 *Geophys. Res. Earth Surf.* 115 (F3), F03013; doi:10.1029/2009JF001321

131 Iverson, R.M., 2000. Landslide triggering by rain infiltration. *Water Resour. Res.* 36 (7),  
132 1897-1910; doi:10.1029/2000WR900090

133 Mergili, M., Marchesini, I., Alvioli, M., Metz, M., Schneider-Muntau, B., Rossi, M., Guzzetti,  
134 F., 2014. A strategy for GIS-based 3-D slope stability modelling over large areas.  
135 *Geosci. Model Dev.* 7, 2969-2982; doi:10.5194/gmd-7-2969-2014

136 MPI Forum, 2012. Message Passing Interface (MPI) Forum Home Page.  
137 URL:<http://www.mpi-forum.org/> (Sep. 2015).

138 Prabhat, Koziol, Q., 2014. High Performance Parallel I/O. Chapman and Hall/CRC, Boca  
139 Raton, USA; ISBN-13:978-1466582347

140 Raia, S., Alvioli, M., Rossi, M., Baum, R.L., Godt, J.W., Guzzetti, F., 2014. Improving  
141 predictive power of physically based rainfall-induced shallow landslide models: a  
142 probabilistic approach. *Geosci. Model Dev.* 7, 495-514; doi:10.5194/gmd-7-495-2014

143