Report of the TCAN Review Panel

K. Heitmann, G. Laughlin, M. Mac Low, A. Mezzacappa, R. Rosner (chair)

February 2, 2016

Following the empanelment of our review committee in December 2015, and distribution of the various materials from the TCAN networks, including our charge, our committee took the following steps:

a. Each reviewer wrote a short description of their impressions of the TCAN program, based on their readings of the materials sent out to us jointly by NASA and the NSF.
b. The chair assembled a first draft, based on these contributions, which was distributed to both all panelists and to the TCAN program managers at NASA and the NSF. This draft document followed the format suggested by the TCAN program managers, responding to three specific questions regarding the program's management. These three questions define the framework of our Report, and are discussed below.
c. We then jointly held a videoconference on January 8, 2016, in which all panelists participated, together with the agency program managers. This videoconference was used to exchange impressions based on the circulated draft documents, and to clarify our charge with the TCAN management.
d. A new review document was then drafted, which took into account our discussions during the videoconference, as well as subsequent comments sent to the chair by the panel members. This revised draft was then re-circulated, and a final document was then prepared. This present document represents the final state of our Panel's deliberations. Note that we did not comment on the scientific quality of the results obtained by the present complement of TCAN networks – such comments were not within our charge. Nevertheless, the Panel was impressed by the overall quality of this program, especially given the severe funding constraints that TCAN has operated under from its inception.

1. *Does the TCAN program, as currently configured and funded, address the objectives laid out in the 2010 Decadal Survey of Astronomy and Astrophysics?*

   It seems evident that all 6 networks were able to establish functional collaborations, collaborations that did end up yielding results that were not likely to have come about in the absence of such networks. On the other hand, there is a clear gap between the current modest program and the ambitions expressed by the 2010 Decadal Survey, who stated, "Appropriately focused and led research collaborations and networks are 'efforts of scale' that can make long-term investments in personnel, computing, and scientific networking uniquely effective in tackling some of the most difficult problems in modern astrophysics." TCAN as a program marks an important step in this direction but has started off severely underfunded, given its ambitions. It seems that an important component underpinning the successes that have been achieved by these networks is the existence of substantial other funding (outside of TCAN) that in fact supported much of the research. That is, there is good evidence that TCAN funds served

as an enabler of the research, and drove the formation of new collaborations among groups working on different aspects of major problems, but they were insufficient to fully support the work that was actually carried out. The modest scope of TCAN funding has also meant that the kind of computer science research and development that DOE-supported programs have enabled for astrophysics (viz., via SciDAC and ASC) cannot yet be carried out under the aegis of this program.

2. *Taking into account the increased workload for both the proposers and the reviewers, is there a continued need for a dedicated TCAN funding line? Or would this time, effort, and funding be better spent on theory proposals funded through individual-investigator programs: i.e. the Astronomy and Astrophysics Research Grants at NSF, and the Astrophysics Theory Program at NASA?*

   The TCAN program has done something that other research funding streams at NSF and NASA have not had great success at accomplishing, namely enabling effective collaborations between groups that share ultimate scientific interests but attack astrophysics problems in complementary ways. TCAN has successfully performed a function that is referred to as "nudging" by behavioral economists – by providing 'enabling' funds on the condition that the funded networks demonstrate collaborative success, TCAN has in fact managed to drive highly collaborative networks. Thus, even if overall funding constraints at NASA and NSF prevent an expansion of TCAN to levels that could fully enable attaining the goals of the 2010 Decadal Survey, the existin modest TCAN funding has already proved to be of substantial value to the astrophysics community: TCAN should continue.

3. *If the TCAN program were to be continued, what changes should be made to better serve the astrophysics community, given the constraints imposed by future flat budgets for both NSF/AST and NASA/APD?*

   We offer four kinds of suggestions for improving the TCAN program, under the assumption that future budgets for both NSF/AST and NASA/APD remain flat, relating to funding, administrative procedures, computing resources, and the mix of research topics.

   First, we address the issue of funding. We note that the TCAN program resembles in many ways the Dept. of Energy's Office of Science (DOE/SC) SciDAC program: Both are designed to advance high-performance computing within scientific disciplines. In the case of SciDAC, DOE/SC has encouraged funding contributions from both the computing program office of DOE/SC (ASCR) and the various science program offices (e.g., high energy physics (HEP), nuclear physics (NP), basic energy sciences (BES), which includes materials science, and biological and environmental research (BER), which includes climate science). This has led to successful collaborations between disciplinary scientists (funded by the DOE/SC science program offices) and computer scientists and applied mathematicians (funded by ASCR); and we therefore suggest that NSF/CISE might similarly contribute to the TCAN program in order to ensure the necessary

involvement of computer scientists in development of the required algorithms as leading-edge computers move more and more in the direction of many-core-based hybrid architectures rather than homogeneous massive parallelism. Moreover, we also suggest that NASA/NSF consider funding slightly fewer networks (e.g., 4 versus 6) in order to better enable those selected. This, coupled with investment from NSF/CISE, would go a long way toward better instantiating the original TCAN vision and the intent of the Decadal Survey, and enabling the TCAN program to be successful in the way the DOE ASC and SciDAC programs have been.

Second, we address possible administrative changes in TCAN. That is, there are a few modest administrative changes that TCAN could implement that have the potential for considerably improving the value of TCAN for the astrophysics community. These changes address three issues related to the funded networks:

- TCAN should require that all code products supported via TCAN be open source, with defined access rules; this change has the potential for substantially increasing TCAN's impact on high-performance computing within the astrophysics community.
- TCAN should require that every network maintain a public web presence for the benefit of the larger astrophysics community that provides (a) regular updates on network accomplishments, including links to papers submitted/published and (b) updated descriptions of computational tool capabilities and links to code repositories and detailed code user documentation. (Writing of such documentation then needs to be included in budget narratives explicitly.) This modification of current procedures would enhance the visibility of TCAN within the astrophysics community.
- TCAN should consider network proposals whose duration is justified by the nature of the collaborative work being proposed. The current approach of "one size fits all" fails to take into account that some challenge problems are simply more difficult than others, so that not all networks can fully reach their goals in the same time frame – insisting that all networks complete their projects in 3 years fails to take this into account.

Finally, there are two additional overarching issues that are primarily the province of NASA and NSF to deal with:

- Computing Access: The challenge of access to the major computer resources required by these networks remains an issue, since successful TCAN projects then must apply separately for computer time in an extremely competitive environment. This is an issue faced by both NSF and NASA. This problem will get worse because of the required re-engineering of codes that were developed and optimized for massively parallel architectures rather than many-core hybrid architectures. Again, following the DOE SciDAC model, we suggest that NSF and NASA consider providing a base level of computing cycles to its awarded TCAN networks. To facilitate this, NSF and NASA could request information from the proposing teams regarding NSF and/or NASA platforms on which they have already run and their expected needs for the TCAN-proposed work.

- Mix of Research Topics and Participants: TCAN should seek to broaden the mix of science areas and institutions supported by its networks. The present mix of science in the TCAN program is remarkable for its lack of breadth – three of the six networks focus on basically the same topic, namely the physics of black hole accretion. Some of the most exciting areas of modern astrophysics in which computing plays a major role are missing entirely: cosmology and structure formation; star formation and the evolution of (proto)planets; and the end points of stellar evolution. Furthermore, few institutional types (e.g., elite private, flagship state, and other private and state institutions, particularly smaller and potentially minority-serving institutions) are participating, possibly limiting access to a diverse cohort of theorists and simulators.