FLASH CENTER RAMP-DOWN PROPOSAL

Submitted to the National Nuclear Security Administration
Advanced Simulation and Computing
Academic Strategic Alliances Program

22 February 2008

1 Introduction

Major recent achievements of the Flash Center include (1) development and release of FLASH 3.0, a highly capable, fully modular, extensible, community code; (2) breakthrough 3-D simulations of the gravitationally confined detonation (GCD) explosion mechanism for Type Ia supernovae that show that this mechanism (which was discovered by the Flash Center through 2-D cylindrical simulations carried out several years ago) robustly detonates for a range of initial conditions; and (3) demonstration that the GCD model of Type Ia supernovae can produce a broad range of nickel masses and peak luminosities, as observed.

The Flash Center is unanimous in thinking that the strategic plan put forward in its proposal to the ASC Predictive Science Academic Alliance Program (PSAAP) was and is the right plan for the Center. The Center will therefore execute a de-scoped version of this strategic plan, despite the rapid decrease in funding from the ASC that it will experience over the next two years. Execution of the de-scoped strategic plan will (1) give the young scientists in the Center accomplishments that will maximize their chances of moving on to excellent positions elsewhere, (2) maximize the ability of the Center to obtain funding from other federal agencies and programs, and (3) maximize the scientific impact of the Center on the Type Ia supernova and Dark Energy field.

The de-scoped strategic plan entails (1) completion of a major verification study of buoyancy-driven turbulent nuclear burning; (2) rigorous, systematic, but limited global validation of a subset of the four current models of Type Ia supernovae; and (3) development and implementation of a bare-bones predictive science code to enable use to compare the output of the simulations with observed Type Ia supernova light curves. We describe these three elements of the de-scoped strategic plan below.

2 Verification of SN Ia Simulations

The Flash Center has underway a major verification study to answer the question: Does buoyancy-driven turbulent nuclear burning occur mostly as a result of small-scale or large scale structure in the flame surface?

Current simulations of Type Ia supernovae by various groups use different subgrid models of turbulent burning:

1. The Munich SN Ia simulation group uses a complicated subgrid model based on assumed existence of fully developed turbulence with parameters of the model set by the properties of the flow at resolved scales – the model increases $s_{\rm flame}$ and produces a large increase in the burning rate [1,2].

2. The NRL group used a simple subgrid model based on properties of R-T–driven turbulent burning in a box – the model greatly increases the effective flame speed $s_{\rm flame}$ and therefore produces a large increase in the burning rate that depends on the value of a parameter m, where $L = m\Delta$ and δ is the grid size, with m taken to be 1 or 2 [3,4].

3. Earlier Flash Center simulations used a local implementation of the NRL subgrid model with m taken to be 1 [5,6,7].

4. Current Flash Center simulations use no subgrid model. This approach is motivated by three pieces of evidence that come from verification simulations the Center has done that suggest that the rate of buoyancy-driven turbulent nuclear burning is dominated by the gemoetry of the flame surface at large scales: (1) The time-averaged burning rate in our "astrophysical flame laboratory" (AFL) simulations of turbulent nuclear burning is *independent of resolution* [8]; (2) Based on this, our most recent simulations do not use a subgrid model for turbulent nuclear burning, yet the burned mass $M_{\rm burn}$ and released nuclear energy $E_{\rm nuc}$ as a function of time in our 3-D simulations of the GCD mechanism show convergence with resolution *without introducing a subgrid model for turbulent nuclear burning* [9];[1] (3) The semi-analytic model of bubble growth and rise to breakout agrees best with the breakout time $t_{\rm breakout}$ in our 3-D simulations *assuming that the burning occurs mostly on large scales* [10].

---

[1] We do introduce a "floor" under the flame speed s to keep the advection-diffusion-reaction (ADR) flame model from being torn apart by R-T–driven turbulence at low densities ($\rho < 1 \times 10^7$ g cm^{-3}). This has an important but modest effect, since burning ceases at densities $\rho < 5 \times 10^7$ g cm^{-3}.

Figure 1: SN Ia validation flow chart, showing the SN Ia simulation pipeline in detail. Yellow circles represent scientific analysis software and/or data archive. Blue boxes are computational codes and algorithms. Yellow boxes are initial conditions and observations.

If most burning does occur on large scales and if we can demonstrate this, it would produce *a major paradigm shift* in the Type Ia supernova field. It would remove a major source of uncertainty in simulations of Type Ia supernovae, and would mean that simulations of Type Ia supernovae can be done with much greater confidence. This would stimulate more comparisons between the predictions of simulations and observations. For these reasons, answering this scientific question will be a major objective of the Flash Center during the coming year.

We will address this question in two ways: (1) We are carrying out a series of AFL simulations of fully-developed R-T turbulent nuclear burning, using the Endeavour (100 Tflop) and Intrepid (500 Tflop) IBM BG/P machines at Argonne National Laboratory (ANL). The details of these series of simulations are the following:

1. **3-D AFL simulations for planar initial conditions.** A sequence of $64^2 \times 256$, $128^2 \times 512$, $256^2 \times 1024$, and $512^2 \times 2048$ simulations are underway on Endeavour at ANL. These will lead up to several $1024^2 \times 4096$ simulations on Intrepid at ANL this spring and summer, each of which will take $\sim$ 10M cpu-hrs.
2. **3-D simulations of the transition from the flamelet nuclear burning regime to the distributed nuclear burning regime due to fully developed R-T turbulence.** A sequence of $64^2 \times 256$, $128^2 \times 512$, and $256^2 \times 1024$ simulations will be done on Endeavour at ANL. These will lead up to several $512^2 \times 2048$ simulations (and possibly a $1024^2 \times 4096$ simulation) on Intrepid at ANL in late 2008. Each of the former simulations will take $\sim$ 0.7M cpu-hrs, while the latter would take $\sim$ 10M cpu-hrs.

3 Validation of SN Ia Simulations

3.1 SN Ia Validation Flow Chart

Figure 1 shows a schematic diagram of the integrated SN Ia pipeline. The pipeline consists of data defining the initial conditions and data from observations of SNe Ia (yellow rectangles), computer codes (blue rectangles), and permanent storage of simulation data (yellow circles). *The data stored permanently will be made available to the scientific community for research purposes.* The major elements of the pipeline are the following:

1. **Initial Conditions and Physics Input.** The initial conditions used in the simulation will be permanently stored in this archive. Simulations of the SN Ia smoldering phase are underway, using low-Mach number implicit codes, to address the questions of what the number and the location of ignition points are [**?**, **?**], but definitive answers are unlikely to be available before the end of the PSAAP. Therefore, we will not simulate the smoldering phase of SNe Ia. Rather, we will take as physical input parameters the number and location of ignition points, as well as the C/O composition of the white dwarf, the convective flow pattern in the core of the white dwarf, whether a detonation is posited (DDT) or followed via simulations (GCD), and rotation. The initial conditions also include simulation parameters like the choice of parameters for the subgrid flame model, the adaptive mesh refinement (AMR) criteria, and the resolution. The interface between the initial parameters is well-defined; in fact, almost all of the input parameters (variables) are FLASH run-time parameters.
2. **SN Ia Explosion Reactive Flow Fluid Dynamics.** We will use the FLASH hydro/reactive flow code to predict the amount of nuclear energy that is released in the explosion, and the density and velocity distributions of the resulting freely expanding ejecta. The release of nuclear energy is over in $\sim$ 3 s, after which the expansion is very nearly homologous (and so can be computed analytically). Therefore it suffices to calculate the density and velocity distributions only to this time. These will

Figure 2: SN Ia verification and validation roadmap.

be permanently stored. The simulations will contain Lagrangian tracer particles, and the temperatures and densities of these particles will be saved as the simulation progresses. These will also be permanently stored.

3. **Nucleosynthesis Post-Processing.** We will use detailed nuclear reaction networks, and the temperature and density histories of the Lagrangian tracer particles computed by the FLASH code, to predict the chemical composition throughout the ejecta [?].
4. **Free Expansion Model Assembly.** We will use the density and velocity distribution at ~ 3 s in Eulerian form provided by the FLASH code and the final chemical composition provided in Lagrangian form by the nucleosynthesis codes at ~ 3 s to assemble the free-expansion model of the ejecta. These will be permanently stored.
5. **Light Curves, Spectra, and Polarization.** Since the free expansion of the ejecta after the release of nuclear energy ends is very nearly homologous, and so can be calculated analytically, the radiation transport codes SEDONA and *Phoenix* will use the free-expansion model at ~ 3 s as input. These codes will output multi-waveband light curves, spectra, and polarization as a function of time and viewing angle.
6. **SN Ia Model Assembly.** The initial conditions, free expansion models, and observational properties of a SN Ia model (velocity, density, chemical composition, light curves, spectra, and polarization) will be integrated into a single dataset and permanently stored.
7. **Predictive Science Code SIMPI.** The predictive science code SIMPI will be used to compare SN Ia simulations and observations, as already described.
8. **Large-Scale Structure (LSS) Simulations.** The large-scale structure (LSS) of the universe distorts the relationship between the distance and the redshift of individual SNe Ia, introducing errors into SN Ia-based cosmological measurements. However, these effects are modest (1 %) and therefore couple weakly to the SN Ia simulation pipeline. We will therefore neglect them in the work undertaken in the next two years.

The Center already has in hand a basic, low fidelity version of the full, integrated SN Ia simulation pipeline. This version is capable of 3-D, full-star hydro/reactive flow simulations for all four SN Ia models from ignition through detonation (if this happens) and into the free expansion phase [?, ?], using the FLASH code [?], and estimates the abundances of chemical elements using the density and temperature of the expanding ejecta [?]. The SEDONA radiation transfer code and output data from the FLASH code have already been used to calculate the predicted light curves, spectra, and polarizations for the GDC model in 2-D (see Figure 4) [?]. We will incorporate a second radiation transfer code, *Phoenix* into the pipeline during the first year.

3.2 SN Ia Verification & Validation Roadmap

Figure 2 shows the SN Ia global validation roadmap we will follow to conduct a rigorous, systematic, but limited global validation of a subset of the current SN Ia models. The four current models are the pure deflagration (D), the deflagration-to-detonation transition (DDT), the gravitationally confined detonation (GCD), and the pulsationally driven DDT (P/DDT) models. Initially, we will treat all models as equally likely.

In globally validating these models, we will use an iterative approach in both the space of output parameters and the space of input parameters. Thus, we will start by using only two output variables, $E_{\rm nuc}$ and $M_{\rm Ni}$, which we have seen are fundamental to SNe Ia: $E_{\rm nuc}$ powers the explosion, while $M_{\rm Ni}$ powers the light curve. We will use existing methods to estimate $E_{\rm nuc}$ and $M_{\rm Ni}$ from observed SN Ia light curves, and in this approximate way we will dramatically reduce the size of the output space that we will initially consider. In the second year, we will add data from SN Ia light curves.

From the beginning, we will validate the existing SN Ia models using not only the properties of individual SNe Ia, but also the properties of *ensembles* of SNe Ia. We will use cross validation of subsets of the data for individual SNe Ia and of the data for subsets of the total sample of SNe Ia to assess model quality. The accuracies of the simulation output that we will require are 10% in $E_{\rm nuc}$ and $M_{\rm Ni}$, and 30% RMS in light curves.

We will start by considering only two input variables: the number of ignition points and their locations, but will quickly go on to consider also the effect of the subgrid model for turbulent burning, as characterized by the effective flame velocity $v_{\rm flame}$. The major verification study of buoyancy-driven turbulent nuclear burning that is underway, as described above, ties into the latter. In the second year, we will consider the effects of convection in the core of the white dwarf star and DDT's, specified by time and location.

Thus the SN Ia validation activities and deliverables during the next two years of the Center will be the following:

1. The major objective of the first year will be to dramatically reduce the space of input parameters, and possibly the space of models, that we will need to consider in Years 2-5 by using a grid of ~ 500 high-resolution (8 km) simulations to determine which SN Ia models are able to match the observed values of $E_{\rm nuc}$ and $M_{\rm Ni}$, and for what values of their input variables. The deliverables at the end of Year 1 will be a set of models plus input parameter values that remain viable, using their ability to match observed values of $E_{\rm nuc}$ and $M_{\rm Ni}$ as the metric, and a much larger set of models plus input parameter values that do not. We estimate that this will require ~ 3000 2-D simulations and ~ 1000 3-D simulations. The latter take 50K cpu-hrs, and so will require $\sim$ 50M cpu-hours, well within the computational resources we expect to have in the first year (see below).

2. We will use a combination of high-resolution (8 km) and very high resolution (2 km) simulations; the predictive science code; and $E_{\rm nuc}$ and $M_{\rm Ni}$, and light curves of individual SNe Ia to rigorously and systematically validate the set of viable models and parameter values from Year 1. We will evaluate the effects of the convection in the core of the white dwarf star and a posited DDT on the output parameters of the simulations, using high-resolution (8 km) simulations. The deliverables at the end of the second year will be a set of comparisons of very high resolution simulations and for a dozen or so SNe I for which high-quality light curve data has been obtained, and a possibly reduced set of models and parameter values that remain viable. We estimate that this will require ~ 2000 2-D simulations and ~ 500 3-D simulations. We estimate that each of the latter will take 100K cpu-hrs, including the simulation of light curves and spectra, and so will require $\sim$ 50M cpu-hours in total, which is well within the computational resources we expect to have in the second year.

3.2.1 SN Ia Global Validation Data

In order to validate the SN Ia models, we need high quality (i.e., high signal-to-noise) light curvesa. The light curves provide information on the amount and distribution of ^{56}Ni and iron peak nuclei in the ejecta; the spectra provide information about the kinetic energy of the explosion (and thus the amount of nuclear energy E_{nuc}) and the amount and distribution of intermediate mass elements.

We have formed collaborations with three world-class supernova observing teams to obtain such data. The three teams are the Harvard-Smithsonian Center for Astrophysics (CfA) supernova group, led by Robert Kirshner; the Sloan Digital Sky Survey (SDSS) Supernova Survey/Dark Energy Survey (DES) team, led by Joshua Frieman; and the Carnegie Supernova Project (CSP), led by Mark Phillips.

CfA Supernova Group. The CfA supernova group is responsible for the detection and spectroscopic identification of more than 2/3 of the very low redshift ($z < 0.1$) SNe Ia that are known [**?**]. In favorable cases, the CfA supernova group is able to obtain high quality light curves from the UV to the near infrared (NIR) and a comprehensive series of spectra.

SDSS Supernova Survey/DES Group. The SDSS Supernova Survey/DES group is detecting and observing SNe Ia at low redshifts ($z < 0.3$). The group expects to have ~ 250 SNe Ia with good quality light curves and spectra by the end of 2007. The Center and the SDSS Supernova Team have been working together for nearly a year, and have recently begin a sustained effort to apply the Gaussian Process method to the characterization of SN Ia light curves and spectra.

Analysis of SDSS Supernova Survey data will continue into 2009. The Dark Energy Survey will begin in late 2010 and will detect 1600-2000 during the following five years, of which ~ 250 will be low redshift ($z < 0.3$) SNe Ia. These SNe Ia will have good quality light curves.

CSP. By the end of 2009, the CSP team will have obtained very high quality spectra as well as very high quality light curves from the UV to the NIR of 100 very low redshift ($z < 0.1$) SNe Ia.

4 Validation Science Software

The Center will use the output of its Type Ia supernova simulations and time-resolved SN Ia light curves to (1) validate a subset of the current SN I models; (2) discriminate among different models, both physical and computational; (3) calibrate and quantify the uncertainties in the model parameters; and (4) predict – including fully calibrated errors – correlations among SN Ia observables that may make it possible to make SNe Ia more accurate standard candles and thus better constrain the properties of dark energy. The Center will use a bare-bones predictive science code to do this.

For each set of values of the input parameters, SN Ia simulations provide a *time-dependent spectrum* – a flux that impinges upon detectors at the Earth and that we want to compare with observed multi-color light curve and spectral data, after convolving it with detector responses. While this comparison is, in principle, well-understood for analytic models, it is *not* a straightforward matter to do this when the models are computationally demanding computer simulations for which the input parameter space can therefore be only be sparsely sampled.

Three main challenges must be overcome in order to make a successful attack on this problem:

1. **High-Dimensional Input Space.** This spans the number and locations of the ignition points, the pattern of convection in the core of the white dwarf star, where, when, and if a DDT occurs, the differential rotation of the white dwarf, and so forth (see Figure 2); and code-related parameters such as the subgrid model of turbulent nuclear burning and adaptive mesh refinement (AMR) criteria.

2. **High-Dimensional Output Space.** This spans SN Ia light curves, spectra, and spectropolarimetry sampled at high resolution which the three partnering observing teams will provide for > 200-300 SNe Ia.

3. **Sparsely Sampled Response Surface.** The response surface – which is the predicted time-dependent SN Ia spectrum as a function of the input parameters – has a complex structure that can only be sparsely sampled because we cannot afford high-fidelity simulations on as many different sets of input parameters as desired. This exacerbates the problems caused by the large dimensionality of the input and output spaces, and gives rise to a very challenging problem of statistical inference.

4.1 Validation using Predictive Science

The methodological approach we adopt is Bayesian statistics. Bayesian statistics is a broad class of statistical methods providing a powerful and flexible framework for rigorous accounting of various sources of uncertainty in making inferences and decisions. A key advantage of the Bayesian approach is the ability to combine model outputs and multiple data sources in a natural way. In particular, Bayesian models are helpful in formulating and analyzing the complex dependencies that may exist between various sources of information. The ability to use all sources of information reduces the uncertainties in the models and makes possible a more realistic quantification of these uncertainties. Both random and systematic errors can be naturally accommodated in this approach [**?**, **?**, **?**].

We now describe in more detail the steps needed in order to compare the output of the Center's SN Ia simulations and observed SN Ia light curves.

1. **Real and Simulated Data:** Data, and simulation output at various input parameter values and fidelity levels (e.g. dimensionality, resolution, physics) are read into the Gaussian Process Emulator (GPE). Numerous low-fi simulations are used to probe broad swathes of parameter space, and fewer higher-fidelity simulations are used to calibrate the low-fi/high-fi response surface discrepancy.

The simulation output and experimental/observational data are read by the Data Interface module, which provides a layer of abstraction separating the predictive science codes from the detailed structure of the data and simulation output. From this point onwards, the predictive science code has no knowledge whatever of the physics embodied by the simulation codes or of the experimental or observational application. This information has been encapsulated in an abstract response surface which maps the input parameter space to the output vector space.

The problem is now a generic one: given a data vector in output space (with measurement uncertainty), and given a sparse sampling of response surface points, (a) supply a GP emulation of the response surface including model uncertainties, (b) determine what region of the response surface best resembles the data vector, (c) exhibit the nearby characteristics of the emulated response surface, (d) compare the data and the model in this neighborhood, and (e) make decisions about future simulations with a view to optimally reducing the uncertainties in the above inferences and predictions. These high-level concerns transcend the specific nature of the data under examination. They are the generic predictive science problems of prediction, inference, and decision.

2. **Data Reduction Services:** DRS lower the dimensionality and simplify the structure of the output space, by coarsening or otherwise remapping the data and the simulation output. Some reduction strategies are generic, such as Principal Component Analysis (PCA) [**?**]. Others are specialized to the type of data and simulation output at hand.

3. **Response Surface Emulation:** Based on the simulations read in at the previous stage, the GPE supplies a GP emulation of the response surface — basically, a "fuzzy" interpolation of the surface that incorporates its own source of uncertainty. This is composed of two separate GP models. The first represents the "cheap" low-fidelity simulations, which probe the parameter space in greatest detail. The other is calibrated at high-fi simulation points and represents the low-fi/high-fi discrepancy. It corrects the inadequacies of the low-fidelity simulations. The two GP models have probabilistic error terms that are naturally compounded.

These GP emulation models are not unique, but may rather be chosen from a wide range of models that are available from the Gaussian Process Model Library. The Library will be designed to explore covariance models of varying complexity. Simple models are expected to be valid only for the coarsest surrogate models and data reductions. More complex models can free the emulator from strong assumptions such as variability isotropy, separability, or stationarity. The tradeoff is greater computational complexity and cost [**?**]. Any available automatic derivatives of the simulation output are incorporated into the GP model, allowing the derivative information to further constrain the emulation.

4. **Prediction:** The GPE computes values of the response surface and their uncertainties at input parameters supplied by the UQ. These probabilistic predictions can be output as science output products. They are useful for plotting model/data comparisons to assess best-fit model quality. They are also useful for estimating prediction uncertainty, through cross-validation. For example, SN Ia model quality may be assessed by "training" the model on light curves in all but one photometric band, and comparing the predictive distribution in the remaining band to the corresponding data.

Sensitivities are derivatives of the response surface with respect to input parameters; they characterize the local structure of the response surface. They are useful for propagating model uncertainties and for assessing valuable directions to explore in parameter space. Sensitivities are also a science output of the GPE.

5. **Model/Data Comparison:** The internal comparison of the emulated model and the data is carried out by means of the likelihood function, which the GPE computes by incorporating both model uncertainties due to the emulation and measurement uncertainties inherent in the data. The likelihood function is computed at parameter space points requested by the UQ.

In addition, the GPE can supply a measure of information borne by the data and the existing simulations: the local predictive entropy. This is essentially an estimate of the information that a future simulation at a proposed parameter setting might contribute to the likelihood calculated at current emulator parameter settings. This quantity can be used to help assess where future simulations should be conducted for optimal uncertainty reduction [**?**, **?**].

6. **Uncertainty Quantification:** The UQ samples the likelihood function using its MCMC unit. At each parameter space point, the UQ obtains the current likelihood function value from the GPE. The likelihood is multiplied by the configured prior density to produce the posterior density. The MCMC unit probes and integrates the posterior density over parameter space, locating the highest probability value (the best-fit, or mode), establishing Bayesian confidence regions of probability on configured parameters, and supplying low-

dimensional marginal distributions resulting from integrating the posterior density over all but one or two parameters. The mode, confidence regions, and marginal confidence regions are all output science products of the UQ.

The agnosticism in this approach concerning the nature of model and data implies that beyond the analysis of individual SNe Ia, it can just as well jointly analyze larger *ensembles* of SNe Ia, embedding the explosion model in a metamodel specifying the population distribution of model parameters over SNe Ia. Forward propagation of this population model, with due accounting of uncertainty, should result in predictive distributions that manifest correlations among observational parameters (within the quantified uncertainties) reflecting the inferred structure of the metamodel.

5 Required Elements of Statement of Work

5.1 Personnel Management

Current students in the Center will be used as much as possible to perform the task required in the de-scoped strategic plan. All graduate students will either have completed their degrees or have transitioned to support from Teaching Assistantships by the end of the two-year ramp-down period.

The Center's personnnel management plan will retain to the maximum extent possible and for as long as possible key personnel needed to execute the de-scoped strategic plan.

5.2 FLASH Legacy Plan

Final release of FLASH 3.0 is scheduled for 29 February 2008 – well before the end of the two-year ramp-down period. Per an agreement made long ago, the University of Chicago will support 2-3 staff members to maintain the FLASH code should no other sources of support materialize during the two-year ramp-down period. However, this is not sufficient to assure the continuation of FLASH as a community code for several reasons: (1) the task of merely maintaining FLASH will not be interesting enough to retain the best people in the code group; (2) major developments of the FLASH code will be necessary in order for it to work well on new architectures (e.g., threading of the FLASH code for multi-core architectures); and (3) all codes decay if they are not actively used and growing in response to the needs of local users.

The Center is therefore vigorously pursuing funding for the Flash Center from other federal agencies and programs, including; (1) funding by the DOE Office of Science in support of the Joint Dark Energy Mission; (2) funding by the DOE NNSA to support development of FLASH for use as the code for the open NIF user community; and (3) funding by NSF to support development of FLASH for use as a community code for part of the computational fluid dynamics community.

5.3 Capstone Simulations

The largest simulations of buoyancy-driven turbulent nuclear burning and the GCD model of Type Ia supernovae that were described above will constitute the capstone simulations carried out by the Center. Global validation of the latter will be carried out, also as described above.

5.4 Project Reviews

The Center will respond to any required project reviews.

5.5 Final Report

The results of the study of buoyancy-driven turbulent nuclear burning and global validation of the GCD model of Type Ia supernovae will be reported in papers published in refereed journals. These results, as well as lessons learned, will also be described in a final report.

Table 1: **COMPUTATIONAL RESOURCES**

Category	Current Year	Year 1	Year 2
Total CPU Used (Proj.) or Needed(Est.)	**17**	**50**	**70**
CPU Used/Identified:			
ASC	2	2	2
INCITE (NERSC)	5	8	8
ANL BG/P Early User	10	40	
INCITE (ANL BG/P)		20	60
Total CPU Identified (Est.)	**17**	**70**	**70**
Mass Storage Used (Proj.) or Needed (Est.)	1	$\sim$ **3**	$\sim$ **3**
Mass Storage Identified (Est.)	**1**	$\sim$ **3**	$\sim$ **3**

CPU: in millions of hours
Mass Storage: in petabytes

5.6 Computational Resources

Local Computational Resources. The Flash Center has a large number of laptop computers and 21 powerful workstations, as well as other equipment. In addition, the Flash Center has Ellipse, a 16-node (32-processor) Linux cluster, and a RAID array with more than 50 TB of storage.

The Flash Center is a participant in the CI SAN mass storage, which currently has 60 TB of high-speed storage connected by multiple GigE to Ellipse, as well as to Teraport, a 240-processor Linux cluster in the CI to which the Flash Center has access. These computational resources are used for code development, and visualization and scientific analysis requiring large computational resources. Examples include the production of movies of the Flash Center's simulations of the GCD model of SN Ia, and analysis of the 25 TB of data generated by the simulation of homogeneous, isotropic, driven, weakly compressible turbulence carried out by the ASAP Flash Center on the LLNL BG/L. The turbulence data set resides on the CI SAN. The data set and the parallel analysis tools needed in order to analyze it, and Teraport are all being available to the turbulence community, so that they can either transfer limited data sets to their home institutions for analysis locally, or interrogate the whole data set or large parts of it remotely, using Teraport. The Flash Center for Predictive Science will also be a participant in all of these computational resources.

The Flash Center is connected by a 10 GigE link to ANL, with GigE connections to Fermilab, and all of the Big Ten Universities. The Flash Center therefore has excellent Internet connections to the DOE DP labs and the DOE Office of Science labs.

Remote Computational Resources. Table 1 lists the projected remote computational resources that the Flash Center needs and expects to have available this year. It also lists estimates for the computational resources that the Flash Center will need and has identified for the next two years.

We have identified a combination of programs and sources from which the Center can obtain the computational resources it will need. These include NNSA ASC resources in the open, time that the Center will have access to by virtue of FLASH's early user status at the DOE Leadership Facility at ANL on the BG/P, and time on the NERSC machines at LBNL and on the BG/P at ANL that we will compete for through the DOE INCITE.

The Flash Center has already begun using 8 M cpu-hrs available to it on Endeavor, the 100 Tflop IBM BG//P machine at ANL, that is available to it as an early user. In addition, the Flash Center was recently awarded the full 4M cpu-hrs it requested for the 2008 calendar year on Franklin under the ERCAP at NERSC at LBNL, and the full 21M cpu-hrs it requested for the 2008 calendar year on Endeavour, the 110 Tflop IBM BG//P machine at ANL under the INCITE program. In addition, it will have available 50-60 M cpu-hrs during th 2008 calendar year as an early user of Intrepid, the 500T IBM BG/P machine at ANL.

The Center will continue to be constrained in the foreseeable future in the V&V it can do by the computational resources available to it. We have therefore thought carefully about the most effective and efficient way to use the resources we expect to have in order to carry out rigorous, systematic, but limited global validation of a subset of the current SN Ia models. Consequently, the needed and identified resources in Table 1 match. Should more resources become available, it will be relatively easy to design an expanded SN Ia V&V plan that will make valuable use of them. Should fewer resources be available, the Center will reconfigure the SN Ia V&V plan to take this into account as intelligently and thoughtfully as we can, based on the circumstances at that time and the long experience it has had in successfully dealing with similar challenges.

6 References