

Modeling Languages and Condor: Metacomputing for Optimization *

Michael C. Ferris Todd S. Munson[†]

October 21, 1998

Abstract

A generic framework for utilizing the computational resources provided by a metacomputer to concurrently solve several optimization problems generated by a modeling language is postulated. A mechanism using the Condor resource manager and the AMPL and GAMS languages is developed and applied to a technique for solving a mixed integer programming formulation of the feature selection problem. Due to the method's computational requirements, the ability to perform optimizations in parallel is necessary in order to obtain results within a reasonable amount of time. We provide details about our simple and easy to use tool and implementation so that other modelers with applications generating many independent mathematical programs can take advantage of it to significantly reduce duration.

1 Introduction

One branch of the machine learning community, those researching supervised learning [3, 18, 19], attempts to construct a process based upon historical data for the purpose of forecasting. The feature selection problem chooses a small number of the data characteristics with the best predictive capability. This problem is applicable in numerous situations and is becoming increasingly important, especially in data mining. Many approaches to solving the problem have been postulated and used [5, 6, 7, 20, 21, 22, 23, 29]. The method presented in this paper generates a large number of independent mixed integer programs (MIP). To make this technique practical, we need to perform the individual optimizations in parallel. Rather than require a large parallel computer, we utilize a metacomputer, a confederation of heterogeneous computing resources including, but not limited to, supercomputers, workstations, and specialized machines connected

*This material is based on research supported by the Air Force Office of Scientific Research under grant F49620-98-1-0147 and National Science Foundation Grants CCR-9619765 and CDA-9726385.

[†]Computer Sciences Department, University of Wisconsin – Madison, 1210 West Dayton Street, Madison, Wisconsin 53706 (ferris,tmunson@cs.wisc.edu)

through a network. The tool developed for concurrently solving optimization problems exploits one vast, and largely untapped, part of this metacomputer, a pool of pre-existing, “off-the-shelf” workstations via Condor.

Condor [12, 24] is a management tool based upon the premise that most workstations are severely under-utilized. Typically, the workstation owner does some computation and then the machine idles for long stretches of time. Condor notices these idle machines and schedules tasks on them. When the owner wants to perform additional work, Condor migrates the task to another available computer. In essence, Condor uses resources that would have otherwise been lost.

Our mechanism allows Condor to be used from within a modeling language [4, 16] which provides a natural, convenient way to represent mathematical programs. These languages typically have efficient procedures to handle vast amounts of data and can quickly generate a large number of models. For this reason, modeling languages are heavily used in practical applications.

Condor and modeling languages form a synergistic combination. Linking them together gives us expressive power and allows us to easily generate simple parallel programs. Successful applications of our mechanism should possess two key properties; they should generate a large number of independent tasks and each individual task should take a long time to complete. Applications with the above two properties cannot be reasonably performed serially. Furthermore, the model generation time and scheduling overhead are ameliorated by the resources spent solving each individual task.

The programming paradigm envisioned is the master/worker model in which a master program generates a large number of independent subproblems, the workers, that can be solved in parallel. Once the subproblems finish, the master program performs additional computations and creates a new set of subproblems. This cycle repeats.

The framework we postulate uses two functions. The first is a way to spawn tasks for solution on a metacomputer. The corresponding retrieve function is used to obtain the results. We implemented these commands using the Condor resource manager and the AMPL and GAMS modeling language. We then use the tool for solving the feature selection problem. The key contribution of this paper is the general framework provided for solving such problems in parallel. The feature selection problem developed is only one example that uses this mechanism. All of the code and examples developed in this paper are available for download at [14].

Our formulation of the feature selection problem, described in Section 2, is an ideal candidate for parallelization. The mechanism that we develop is simple and easy to use and can also be adapted for use in other applications. Section 3 provides simple examples in both the AMPL [17] and GAMS [8] modeling languages of how the tool is used. The implementation of the feature selection problem in this framework is also discussed. We present technical details of the tool in Section 4. Results of the feature selection using several representative datasets are found in Section 5. They demonstrate that significant improvements in overall solution time can be reliably achieved by running

state-of-the-art MIP solvers such as XPRESS [10] in parallel. Finally, we draw some conclusions and point out relevant extensions of the work in Section 6.

2 Feature Selection

Researchers in the machine learning community studying supervised learning [3, 18, 19] attempt to construct a scheme based upon known historical data to classify unknown observations. While no guarantee can be made that the forecast is correct, a best possible guess is desired. We consider the two category case, in which the given data consists of measurements taken from elements in each of two categories. Many important applications can be posed within this framework [2, 6, 26].

Let $\mathcal{A} \subseteq \mathfrak{R}^F$ and $\mathcal{B} \subseteq \mathfrak{R}^F$, be finite sets of observations taken from elements in category 1 and 2 respectively. Each point measures the same set of F characteristics, with elements in the same category sharing similar properties. By $A \in \mathfrak{R}^{|\mathcal{A}| \times F}$ and $B \in \mathfrak{R}^{|\mathcal{B}| \times F}$ we denote the matrices formed from all elements in \mathcal{A} and \mathcal{B} , where $|\cdot|$ is used to denote the cardinality of a set.

The approach considered in this paper attempts to quantify differences between the two categories by constructing a separating hyperplane [1, 25], $P := \{x \in \mathfrak{R}^F \mid x^T w = \gamma\}$ with $w \in \mathfrak{R}^F$ and $\gamma \in \mathfrak{R}$ such that for all $a \in \mathcal{A}$, $a^T w > \gamma$ and for all $b \in \mathcal{B}$, $b^T w < \gamma$. If $y \in \mathfrak{R}^F$ is an unknown observation we want to classify, we use the following process to categorize it:

1. If $y^T w > \gamma$ then y likely belongs to category 1.
2. If $y^T w < \gamma$ then y likely belongs to category 2.
3. Otherwise, we cannot make any determination.

Typically, we cannot construct such a separating hyperplane, so we choose one minimizing a measure of the misclassification on the known data. Furthermore, recent work [5, 6, 23] has shown that it is typically beneficial to use a subset of the measured features for the classification. If we employ too few features, we will not have enough information to correctly classify future data. However, if we have too many features, we will over-classify the system and perform poorly on the unseen data. We want to choose the optimal number of features, i.e. that with a minimal expected misclassification percentage.

In the next subsection, we discuss a mechanism for constructing a hyperplane minimizing the one-norm of the mean violation for a given number of features. The following subsection describes how 10-fold cross validation [30] is used to determine the effectiveness of the hyperplane by calculating the expected misclassification percentage. The final subsection discusses the issues involved in finding the best number of features to use and the computational aspects that make this application ideal for our parallelization tool.

2.1 Classification

We use linear programming to construct a hyperplane such that a measure of the violation of the misclassified points is minimized. In the best case, we want to choose w and γ such that:

$$Aw > e\gamma \text{ and } Bw < e\gamma$$

where e is a vector of ones of the appropriate dimension. We can remove the strict inequalities by normalizing the problem [1] to obtain an equivalent system:

$$Aw \geq e(\gamma + 1) \text{ and } Bw \leq e(\gamma - 1).$$

We want to minimize the mean of the one-norm violations to this condition [1]. Stated as a linear program, we have:

$$\begin{aligned} \min \quad & \frac{1}{|\mathcal{A}|}e^T s + \frac{1}{|\mathcal{B}|}e^T t \\ \text{subject to} \quad & -Aw + e(\gamma + 1) \leq s \\ & Bw + e(1 - \gamma) \leq t \\ & s, t \geq 0 \end{aligned}$$

We note that if $w_i = 0$, then feature i is not used in the construction of the hyperplane, and is hence irrelevant to the classification process.

Since we want to vary the number of features selected, we use a mixed integer programming reformulation that forces some of the components of w to zero. If we want to use at most f of the F features when constructing the hyperplane, we can solve the following problem:

$$\begin{aligned} \min \quad & \frac{1}{|\mathcal{A}|}e^T s + \frac{1}{|\mathcal{B}|}e^T t \\ \text{subject to} \quad & -Aw + e(\gamma + 1) \leq s \\ & Bw + e(1 - \gamma) \leq t \\ & -\alpha y \leq w \leq \alpha y \\ & e^T y = f \\ & s, t \geq 0, y \in \{0, 1\}^F \end{aligned}$$

where α is a large positive constant. For $y_i = 0$, we have that $w_i = 0$. Since $e^T y = f$ also holds, at most f features can be used. We note that this model does not preclude the selection of fewer than f features. If $y_i = 1$, then we can still set $w_i = 0$, i.e. feature i is not used.

The w and γ resulting from solving this mixed integer program provides one hyperplane minimizing the one norm of the mean misclassification using at most f features. We now want to determine the expected misclassification percentage, which is done using 10-fold cross validation.

2.2 Validation

10-fold cross validation [30] is a statistical technique used to measure generalization ability. In our case, we use the technique to determine the expected misclassification percentage when f features are used to construct a hyperplane minimizing violation. The procedure starts by splitting the given data, \mathcal{A} and \mathcal{B} , into 10 random subsets of approximately equal size, $\mathcal{A}_1, \dots, \mathcal{A}_{10} \subseteq \mathcal{A}$ and $\mathcal{B}_1, \dots, \mathcal{B}_{10} \subseteq \mathcal{B}$. Each element of \mathcal{A} is placed into exactly one of the subsets \mathcal{A}_j ; similarly for each element of \mathcal{B} . Let $\mathcal{A}_j^c := \mathcal{A} \setminus \mathcal{A}_j$ and let $A_j^c \in \mathbb{R}^{|\mathcal{A}_j^c| \times F}$ denote the matrix formed from the elements of \mathcal{A}_j^c . Furthermore, let $\mathcal{B}_j^c := \mathcal{B} \setminus \mathcal{B}_j$ and let $B_j^c \in \mathbb{R}^{|\mathcal{B}_j^c| \times F}$ denote the matrix formed from the elements of \mathcal{B}_j^c . For a fixed $j \in \{1, \dots, 10\}$, the set $\mathcal{A}_j^c \cup \mathcal{B}_j^c$ is known as the training set, while $\mathcal{A}_j \cup \mathcal{B}_j$ is known as the testing set. We note that approximately 90% of the data will be used for training and 10% for testing.

For each $j \in \{1, \dots, 10\}$, we solve the following classification problem:

$$\begin{aligned} \min \quad & \frac{1}{|\mathcal{A}_j^c|} e^T s + \frac{1}{|\mathcal{B}_j^c|} e^T t \\ \text{subject to} \quad & -A_j^c w + e(\gamma + 1) \leq s \\ & B_j^c w + e(1 - \gamma) \leq t \\ & -\alpha y \leq w \leq \alpha y \\ & e^T y = f \\ & s, t \geq 0, y \in \{0, 1\}^F \end{aligned}$$

for a fixed f giving (w^j, γ^j, y^j) . Using the hyperplanes generated for the training sets, we calculate the percentage of misclassified elements on the corresponding testing data. We define the following error functions:

$$\begin{aligned} \text{error}_{\mathcal{A}}(x, w, \gamma) &:= \begin{cases} 1 & \text{if } x^T w \leq \gamma, \\ 0 & \text{otherwise.} \end{cases} \\ \text{error}_{\mathcal{B}}(x, w, \gamma) &:= \begin{cases} 1 & \text{if } x^T w \geq \gamma, \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

These functions indicate if a point has been misclassified with regard to the hyperplane defined by w and γ and the sets \mathcal{A} and \mathcal{B} respectively. We then determine the percent misclassified as follows:

$$p^j := \frac{1}{|\mathcal{A}_j| + |\mathcal{B}_j|} \left(\sum_{a \in \mathcal{A}_j} \text{error}_{\mathcal{A}}(a, w^j, \gamma^j) + \sum_{b \in \mathcal{B}_j} \text{error}_{\mathcal{B}}(b, w^j, \gamma^j) \right).$$

The final number generated is the average percentage of misclassified elements over all 10 folds, $p := \frac{1}{10} \sum_{j=1}^{10} p^j$. This number indicates how well we can be expected to perform on unseen data, assuming that \mathcal{A} and \mathcal{B} accurately represent the real world. Since we use random subsets, we need to run the validation several times to truly obtain a good estimate of expected performance.

2.3 Selection

In order to select the best number of features to use, we vary f from $1, \dots, F$ and perform the 10-fold cross validation procedure L times for each selected f , tracking the expected misclassification percentage for each validation run. We calculate the mean and standard deviation of the measure over all L validation runs. We choose the number of features producing the smallest mean value as our optimal solution.

From a computational perspective, each of the independent tasks spawned is a mixed integer program. These problems can be very difficult to solve and can require a large amount of computational resources. For a fixed f , we can run $10L$ independent tasks concurrently; each run of 10-fold cross validation generates 10 mixed integer programs and we perform L such validations at a time. The feature selection problem therefore satisfies our candidate criteria and is ideal for the parallelization envisioned. We now turn our attention towards how our tool is used and the implementation for the feature selection problem.

3 Examples

Our tool for solving multiple models in parallel is flexible and powerful, yet can be easily incorporated into pre-existing applications written using the AMPL or GAMS modeling languages. New versions of these packages are not required in order to use our software. Two procedures were developed for our implementation; one to spawn jobs and the other to retrieve the results. Modelers do not have to entirely rewrite their code to exploit inherent parallelism, they simply replace the solve with a spawn command and later issue a retrieve. The commands are implemented within the modeling language framework. The spawn is non-blocking, which means we immediately resume processing the application code after the command is issued, thus enabling us to put multiple models into a job queue before waiting for results. Users are also able to interleave solves performed under Condor and those run on the host machine. For example, problems that contain confidential data can be solved locally, while less sensitive subproblems can be solved elsewhere. Furthermore, Condor allows us to utilize the full resources of a heterogeneous computing environment and allows the state of long running jobs to be saved and restored at a later time.

In the rest of this section, we provide some simple examples illustrating the basic features of our mechanism. We then discuss the issues encountered when coding the feature selection problem.

3.1 Simple Examples

The simplest way to use our tool is to replace a single solve command with a spawn/retrieve combination. This technique illustrates the basic method, but fails to exploit the full capabilities of the mechanism. Two simple examples, one within the AMPL modeling language and the other from GAMS, highlight how we can use the spawn and retrieve to solve multiple models in parallel.

3.1.1 AMPL

The AMPL implementation consists of two routines called via the `commands` function, `condor_spawn` and `condor_retrieve`. Instead of issuing a `solve`, the modeler simply replaces it with a `condor_spawn`, and later issues the corresponding `condor_retrieve`.

The example selected for exposition in AMPL is an implementation of Dantzig-Wolfe decomposition for a multi-commodity transportation problem. The original code is available at:

<http://www.ampl.com/cm/cs/what/ampl/NEW/LOOP2/multi2.run>

We only show the portion of the source code modified to run under Condor.

The original code uses a looping construct to solve several subproblems serially:

```
for {p in PROD} { printf "\nPRODUCT %s\n\n", p;
    solve SubI[p];
    display Artif_Reduced_Cost[p];
    ...
};
```

We can solve each individual subproblem via Condor by replacing the `solve` with the following `spawn/retrieve` combination:

```
for {p in PROD} { printf "\nPRODUCT %s\n\n", p;
    problem SubI[p];
    commands condor_spawn; commands condor_retrieve;
    display Artif_Reduced_Cost[p];
    ...
};
```

In this case, one problem is submitted to Condor and then we immediately wait for the result. This approach is likely to give worse performance than the serial code because of the additional scheduling overhead. A better technique is to repeat the loop, once to spawn all the jobs and the second time to retrieve all the results:

```
for {p in PROD} {
    problem SubI[p]; commands condor_spawn;
};

for {p in PROD} { printf "\nPRODUCT %s\n\n", p;
    problem SubI[p]; commands condor_retrieve;
    display Artif_Reduced_Cost[p];
    ...
};
```

We note that the first loop just spawns the jobs, while the second does all of the output relevant to the application. In this case, all the subproblems are spawned out to separate processors (i.e. machines) before any results are collected into AMPL.

3.1.2 GAMS

Performing the spawn and retrieve in GAMS involves “solving” the model twice. The first time, we replace the solver with one that submits the job to Condor. The second uses a “solver” that retrieves the result from Condor. A simple example of the process is carried out within the GAMS code given below. In this application, we have two trivial problems, `probA` and `probB`. The models are spawned to Condor to run independently and their results are retrieved after performing a solve task locally. Although seemingly trivial, this example shows how a master/worker program can be formulated and solved in parallel using Condor. Note in particular, that some (sub)problems can be spawned and solved elsewhere, during which time other optimizations (for example, master or synchronization problems) can be carried out on the host machine.

```
variables obj; equation f, g;

f.. (obj - 1) =g= 0.;
g.. sqr( obj ) =l= 2;

model probA / f /; model probB / g /;

* Spawn jobs
option nlp = mns5con;
solve probA using nlp minimizing obj;
solve probB using nlp minimizing obj;

* Perform other tasks
option nlp = minos5;
solve probA using nlp minimizing obj;

* Retrieve results
option nlp = rescon;
solve probA using nlp minimizing obj;
solve probB using nlp minimizing obj;
```

In this case, the modeler just needs to know which solvers are available on the substituting machine (e.g. `minos5`) and what the Condor solvers are called (e.g. `mns5con`). The special `rescon` solver is provided by us to retrieve the results.

3.2 The Feature Selection Problem

The implementation of the feature selection problem presented in Section 2 is quite sophisticated. We chose the GAMS modeling language in which to implement it. Before running the application, we need to generate data readable by GAMS. To accomplish this task, we wrote a simple program that converts the data set given into one ready for importation by GAMS. At the same time we also calculate random permutations which are used to split the data for validation. We note that the permutations can be created within GAMS, but we found it easier to calculate them outside of the modeling language.

Initially, the code declares appropriate sets and imports the observation data. We give a small section of this portion of the code below:

```
set a / 1 * 358 /;
set o / 1 * 30 /;

parameter a_data(a, o) /
#include "a_data.inc"
/;
```

We read data for the second category in a similar fashion.

We then declare dynamic subsets for the testing and training data and define some equations:

```
set a_test(a);
set a_trai(a);

positive variable s(a);
variables w(o), gamma;

equations a_def(a);

a_def(a_trai)..
    -sum(o, a_data(a_trai, o) * w(o)) + gamma + 1 =l= s(a_trai);
```

In looping statements, we dynamically determine the elements in `a_test` and `a_trai` and submit a mixed integer program to Condor. Since the equations are just defined over the training set, only the relevant ones are generated.

To make the times for each task more reasonable, we only perform the validations for a fixed f at a time. This technique allows us to use information obtained for f features when we solve the mixed integer program for $f + 1$ features. In particular, the objective value for the f features problems is an upper bound on the objective value for $f + 1$ features. Therefore, we have two groups of looping statements. The first set is simple and generates all of the problems needed for the one feature case and submits them to Condor using the aforementioned syntax. The second set is slightly more complex because it needs to retrieve the results, check for errors and then generate the problem for $f + 1$ features. The inner part of the loop consists of the following:

```

option mip = rescon;
solve train using mip minimizing c;

if (execerror or train.solvestat eq 7 or train.modelstat eq 11,
    execerror = 0;
    option mip = xpress;
    solve train using mip minimizing c;
);

* Generate report data here.

if (ord(f) lt 8,
    features = ord(f) + 1;
    option mip = xprcon;
    solve train using mip minimizing c;
);

```

The first solve retrieves the solution to the problem from Condor. The first if-statement checks for execution errors and licensing problems and solves the same problem locally if such a condition exists. An execution error is generated for example if the solver fails to write a solution file. Because of some limitations in the way Condor migrates tasks and the nature of the code used, this error condition is sometimes encountered. The problem arose primarily due to simultaneously reading and writing to the same file. Another potential difficulty is that we can run out of licenses for a particular solver. We check this condition in GAMS by looking at the solver and model status reported. The last if-statement spawns the next problem in the sequence if required. Between the two if-statements, we calculate and store the percent misclassification on the testing data. At the end of all the loops, we calculate the average misclassification percentage and report our findings to the user.

The complete listing of the GAMS program and all relevant data can be downloaded by anonymous ftp [14]. This example demonstrates many of the features incorporated into our tool. Clearly, by adding statements to the AMPL or GAMS languages, we could implement the process more efficiently. However, we believe the strength of our approach derives from its simplicity and the fact that operations researchers can easily and in many practical instances effectively exploit parallelism using it.

4 Technical Details

The spawn procedure is a non-blocking call used to submit a task for solution. Before delving into how the mechanism is implemented, we need to first examine how AMPL and GAMS behave when a normal model is solved. The solve statement first writes an internal representation of the model to disk. We note that both modeling languages considered can write a machine independent representation, thus allowing heterogeneous solves. The particular solver executable is

then run with appropriate command line arguments. The implementation of the spawn command replaces the solver in this two-step process with a more complex series of statements that implement the following steps:

1. Create a new temporary directory.
2. Copy the internal model representation to the temporary directory.
3. Submit the task to Condor and obtain the associated Condor job number.
4. Place the user job number and Condor job number on a job queue.
5. Return control to the modeling language.

The job queue plays a critical role when we want to retrieve the results.

To submit a job to Condor, we create a job description file. One such file for a GAMS model is provided below:

```
Executable = gamsxprcon.out
Initialdir  = condor_home/temporary_directory
Arguments   = gamscntr.scr
Log         = condor_home/log
Notification = never
Queue
```

We then submit the job using the command `condor_submit -v`. This gives “verbose” output which we process to obtain the Condor job number. If for some reason the submission fails, we wait a small amount of time and then reattempt the submission.

The retrieve is a blocking call used to obtain the results. When the solver executable finishes, a solution file is written. Essentially, the retrieve places the solution file in the correct location for the modeling language to read. The problem whose solution is desired needs to be in scope before issuing the retrieve command, otherwise problems can occur when the modeling language tries to import the solution. Therefore, we require that results for the models be retrieved in precisely the order they were spawned. While this may be inefficient in terms of load balancing, it is simple to implement and ensures the modeler knows precisely the ordering of jobs. To retrieve a job, we perform the following steps:

1. Take a user job number and Condor job number off the job queue.
2. Wait for the Condor job to finish.
3. Copy the solution file into the correct location.
4. Remove the temporary directory.
5. Return control to the modeling language.

We check to see whether the Condor job has finished by periodically checking the log specified in the job description file. Whenever the Condor job changes state, an appropriate event is noted in this log. We look for the “job completed” event which tell us that the Condor job has finished.

The basic mechanism is the same for the procedures in both AMPL and GAMS. However, differences in the modeling language lead to different implementations. In the next two subsections we discuss the relevant features of the implementation for each of the modeling languages considered. We finish the section by mentioning some Condor related issues.

4.1 AMPL Details

AMPL allows us to execute a script via the `commands` feature. The `condor_spawn` is implemented as the following script:

```
write bcondor;  
shell 'condor_sub minos x86';
```

The first command writes out a (binary) file called `condor.nl` which is a representation for the problem currently in scope. The second command executes a perl script called `condor_sub` that submits the model to Condor for execution. In this particular case, the solver used is MINOS and we run it on the x86 machines in the pool.

Similarly, the `condor_retrieve` is implemented as:

```
shell 'condor_results .';  
solution condor.sol;
```

Here, the `condor_results` script waits for the model to finish processing before returning back to the modeling language. The `solution` command reads in the `condor.sol` file generated by the solver.

4.2 GAMS Details

The solver GAMS uses can be any executable. Therefore, for the submission call, we created a perl script that is invoked instead of an actual solver. However, GAMS presents some unique issues that needed to be resolved.

The first is that when we copy the model, we also need to copy the options file if one exists. Otherwise, we can run into some file permission problems with Condor. At the same time, we update the `gamsctr.scr` file so that it looks in the appropriate temporary directory for the model and options.

Another detail is that if we fail to provide a solution to GAMS from any solve statement, we would get an execution error. This is undesirable since the user would have to perform additional checks to remove the error. Therefore, we have implemented a very simple solver that writes a solution file containing the initial, user-specified point and the function evaluation at that point. This simple solver is called at the end of the submission script. In this way, we avoid the problem.

The retrieve is implemented as a special solver that waits for the results and then copies the solution to the correct location, essentially as in Section 4.1.

4.3 Condor Issues

The Condor tool is designed for high-throughput computing in a dynamic environment. Machines which are currently idle are available for our use. When the owner of the resource returns and wants to work, the machine is removed from the pool and any job currently running on it is stopped and migrated to another available computer. The same mechanisms are used to provide fault tolerance when machines unexpectedly become unavailable. Condor provides two modes of operation for such cases. The ‘vanilla’ universe is primarily used for jobs that do not require a long time to complete. When ‘vanilla’ jobs are migrated, the computational effort already expended is lost and the job is restarted from the beginning. All that is required in this case is an executable. For example, existing executables for solver links to AMPL and GAMS could be used in this ‘vanilla’ universe without any modification whatsoever.

The ‘standard’ environment is preferred for long running programs and enables a job to be checkpointed and restarted without losing all of the time spent in the code. The checkpoint appears as a normal executable but contains a “snapshot” of the job at its current state. When the checkpoint is run, it resumes the computation from the exact state at which the checkpoint was made. In order to submit a ‘standard’ job, the program called need only be relinked using the ‘condor_compile’ command. This command incorporates the appropriate Condor libraries into the code to perform checkpointing.

Some limitations are placed upon the jobs that we want to checkpoint. The main one we encountered is that we are not allowed to simultaneously read and write to the same file. The results are unpredictable if a code does this type of i/o. Another problem has to do with the use of sockets, which are not allowed under Condor. This could cause potential problems with some license managers. Other limitations have to do with signal handlers and forking processes; these are detailed in [9].

There are several advantages that accrue to a Condor user. First of all, we note that we are not required to have an account on all of the machines in the Condor pool. Condor can start jobs on machines for which we do not have permission to directly use. Condor makes the processor and memory on these machines available to the job submitter. Data is read from and results are written to the submitter’s disk either through a “shared file system” or by using remote callbacks. These mechanisms make Condor an ideal “intranet” tool to facilitate large scale computations over a local network of workstations. Internet usage is also possible as documented in [13]; however, in this environment other security issues also arise.

A second key feature is that Condor enables us to run our solvers in a heterogeneous environment. This power comes with one caveat. Since the checkpoints use system dependent information, once we start a job on a particular architecture, we cannot migrate the job to a different architecture without restarting

the job from the beginning. For ‘standard’ universe jobs, we need to specify the exact platform on which to run the jobs because of this limitation.

Finally, we note that all machines in the pool are not equal. Some of them may have fast processors with a large amount of memory, while others may be slower. The classified ad mechanism provided by Condor allows us to prioritize the available machines. We consequently rank the machines based upon processor speed and available memory. Documentation on the Condor supplied commands, job universes, checkpointing, and classified ads can be obtained from the Condor website [9].

5 Results

Two standard datasets were used for our tests with different characteristics. The galaxy dataset [27] contains a large number of observations, but has a small number of features. There are 1505 and 957 elements in the two categories, dim and bright, respectively. Each has 14 measured features. The WDBC database [26] contains observations of malignant and benign breast tumors. The set consists of observations for 358 benign tumors and 212 malignant tumors. Each observation measures 30 attributes.

Instead of performing the validation for all of the features in each dataset, we varied the number of features selected from 1 to 8. The difficulty of the mixed integer programs solved for the WDBC data set increases dramatically as the number of features selected increases. A model is solved for f features before moving on to selecting $f + 1$ features. The objective function value for f features is an upper bound on the objective function value for $f + 1$ features. The branch and bound code used, in this case XPRESS [10], can use this information to significantly limit the size of the tree generated, making the problem more tractable. For each of the runs, we prove the optimality of our solution. We plot the average number of nodes explored and the average time per solve in Figures 1 and 2. While in practice it may be sufficient to solve the problems to within a prespecified error tolerance, this was not carried out here since both the quality of the solution and utility of our tool were deemed more important.

To maintain efficiency, we attempt to keep as many tasks in the job queue as possible. To do this we immediately spawn $10L$ tasks. Whenever we retrieve results for a model, we immediately add another job to the queue. We decided to perform 10-fold cross validation 20 times, leading to a maximum of 200 jobs in the job queue at a time. We use up to 8 features, so a total of 1600 jobs were created and solved. Each of the jobs can take from several seconds up to several hours to complete.

Because of the way the branch and bound code is written, it is possible that a license might not be available. Therefore, we need to detect that an error occurred in the GAMS code and recover from it. Fortunately, in GAMS, we can check the model and solver status to determine when such a case was encountered. We also check for execution errors and reset the value of the `execerror` flag. Problematic subproblems are then solved on the local machine.

Figure 1: Plots of Average Number of Nodes Explored

Figure 2: Plots of Average Time Per Problem

Figure 3: Plots of Average Misclassification Percentage

In this way, we can detect and recover from the problems encountered.

The Condor pool at the University of Wisconsin-Madison, containing around 300 workstations, was used as our testbed. Our critical resource with this model is disk space on the submitting machine. Due to this limitation, we only allowed 200 jobs to be concurrently running.

A serial version of the code ran in 3.3 days for the galaxy set and 6.9 days for the WDBC dataset. 8% and 3% of this total time was spent generating the models in GAMS. Using Condor we were able to get the results back in 0.6, 1.4 days respectively. The total speed up is then 5.5 and 4.9 time. A plot of the misclassification error versus the number of features is found in Figure 3. We can see from the plot that the optimal number of features to choose for the galaxy dataset is 5 and the best number of features to select for the WDBC database is 3.

6 Conclusion

In this paper, we have described an easy to use tool that provides enormous computational resources to a mathematical modeler. The computational power is harnessed from an existing network of workstations using the Condor resource manager to detect and deliver idle cycles within this network.

To facilitate high throughput within this metacomputing environment, we have built an interface to the two most popular modeling languages in optimization, namely AMPL and GAMS. Our system requires no changes to these languages, but is provided as an add-on feature, built within the confines of the existing language structures. Once our system is installed, only very minor and simple changes are required to an optimizer's model file.

The tool developed is suitable for solving multiple optimization models in parallel from within AMPL and GAMS. We were able to obtain significant speed-up for a particular application, namely the feature selection problem. Other applications, such as multimedia architecture design [11] and radiation

therapy [28], can also use the mechanisms provided to improve performance. The only limitation to the number of concurrent jobs is the availability of computational resources.

Our facility can be used in other situations where we can exploit a larger metacomputer concept. For example, we can have one machine dedicated to visualization of results and another pool of workstations for solving problems. Furthermore, due to the tool's generality, we can also exploit other resources, such as supercomputers. The only facilities required are a mechanism to spawn an optimization task on a remote machine and a way to retrieve the results. A key contribution of our work is that we provide a template for these mechanisms in the Condor implementation. Other utilities, such as GLOBUS [15], could be used for a similar implementation in a broader metacomputing environment.

Finally, our work would benefit by Condor providing additional facilities. The ability to add tasks to a pre-existing Condor cluster would help to reduce the resources spent scheduling. Another feature would be the ability to obtain a resource and use it to solve several subproblems without relinquishing it. Both of these features would help to reduce Condor's overhead. An additional facility that would simplify the spawn would be for Condor to allow us to submit jobs to multiple platforms with Condor deciding which one is used based upon our current ranking of the machines.

References

- [1] K. P. Bennett and O. L. Mangasarian. Robust linear programming discrimination of two linearly inseparable sets. *Optimization Methods and Software*, 1:23–34, 1992.
- [2] K. P. Bennett, D. Wu, and L. Auslender. On support vector decision trees for database marketing. Technical report, Department of Mathematical Sciences, Rensselaer Polytechnic Institute, 1998.
- [3] C. M. Bishop. *Neural Networks for Pattern Recognition*. Clarendon Press, Oxford, England, 1995.
- [4] J. Bisschop and A. Meeraus. On the development of a general algebraic modeling system in a strategic planning environment. *Mathematical Programming Study*, 20:1–29, 1982.
- [5] P. S. Bradley and O. L. Mangasarian. Feature selection via concave minimization and support vector machines. In J. Shavlik, editor, *Machine Learning Proceedings of the Fifteenth International Conference (ICML '98)*, pages 82–90, San Francisco, California, 1998. Morgan Kaufmann. <ftp://ftp.cs.wisc.edu/math-prog/tech-reports/98-03.ps.Z>.
- [6] P. S. Bradley, O. L. Mangasarian, and W. N. Street. Feature selection via mathematical programming. *INFORMS Journal on Computing*, 10(2):209–217, 1998. <ftp://ftp.cs.wisc.edu/math-prog/tech-reports/95-21.ps.Z>.

- [7] E. Bredensteiner and K. Bennett. Feature minimization within decision trees. *Computational Optimization and Applications*, forthcoming, 1998.
- [8] A. Brooke, D. Kendrick, and A. Meeraus. *GAMS: A User's Guide*. The Scientific Press, South San Francisco, CA, 1988.
- [9] Computer Sciences Department, University of Wisconsin, Madison, Wisconsin. *Condor homepage*. <http://www.cs.wisc.edu/condor>.
- [10] Dash Associates, Blisworth House, Blisworth, Northants, UK. *XPRESS-MP User Guide*, 1997.
- [11] D. L. Eager, M. C. Ferris, and M. K. Vernon. Optimized regional caching for on-demand data delivery. Technical Report 1385, Computer Sciences Department, University of Wisconsin, Madison, Wisconsin, 1998. Forthcoming in the proceedings of Multimedia Computing and Networking 1999 (MMCN99).
- [12] D. H. J. Epema, M. Livny, R. van Dantzig, X. Evers, and J. Pruyne. A worldwide flock of condors: Load sharing among workstation clusters. *Journal on Future Generations of Computer Systems*, 12, 1996.
- [13] M. C. Ferris, M. P. Mesnier, and J. Moré. NEOS and condor: Solving nonlinear optimization problems over the Internet. Mathematical Programming Technical Report 96-08, Computer Sciences Department, University of Wisconsin, Madison, Wisconsin, 1998. Also available as ANL/MCS-P708-0398, Mathematics and Computer Science Division, Argonne National Laboratory.
- [14] M. C. Ferris and T. S. Munson. Condor modeling language interface. <ftp://ftp.cs.wisc.edu/math-prog/condor>.
- [15] I. Foster and C. Kesselman. Globus: A metacomputing infrastructure toolkit. *International Journal of Supercomputer Applications*, 1998. To appear.
- [16] R. Fourer, D.M. Gay, and B.W. Kernighan. A modeling language for mathematical programming. *Management Science*, 36:519–554, 1990.
- [17] R. Fourer, D.M. Gay, and B.W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming*. Duxbury Press, 1993.
- [18] S. I. Gallant. *Neural Network Learning and Expert Systems*. MIT Press, Cambridge, Massachusetts, 1993.
- [19] J. Hertz, A. Krogh, and R. G. Palmer. *Introduction to the Theory of Neural Computation*. Addison-Wesley, Redwood City, California, 1991.

- [20] G. H. John, R. Kohavi, and K. Pflieger. Irrelevant features and the subset selection problem. In *Proceedings of the 11th International Conference on Machine Learning*, pages 121–129, San Mateo, CA, 1994. Morgan Kaufmann.
- [21] K. Kira and L. Rendell. The feature selection problem: Traditional methods and a new algorithm. In *Proceedings of the Tenth National Conference on Artificial Intelligence*, pages 129–134, San Mateo, CA, 1992. Morgan Kaufmann.
- [22] J. Kittler. Feature selection and extraction. In T. Y. Young and K.-S. Fu, editors, *Handbook of Pattern Recognition and Image Processing*. Academic Press, New York, 1986.
- [23] Y. le Cun, J. S. Denker, and S. A. Solla. Optimal brain damage. In D. S. Touretzky, editor, *Advances in Neural Information Processing Systems II (Denver 1989)*, pages 598–605, San Mateo, California, 1990. Morgan Kaufmann.
- [24] M. J. Litzkow, M. Livny, and M. W. Mutka. Condor: A hunter of idle workstations. In *Proceedings of the 8th International Conference on Distributed Computing Systems*, pages 104–111, June 1988.
- [25] O. L. Mangasarian. Linear and nonlinear separation of patterns by linear programming. *Operations Research*, 13:444–452, 1965.
- [26] P. M. Murphy and D. W. Aha. UCI repository of machine learning databases. Department of Information and Computer Science, University of California, Irvine, California, <http://www.ics.uci.edu/AI/ML/MLDBRepository.html>, 1992.
- [27] S. Odewahn, E. Stockwell, R. Pennington, R. Hummphreys, and W. Zুমach. Automated star/galaxy discrimination with neural networks. *Astronomical Journal*, 103(1):318–331, 1992.
- [28] D. M. Shepard, M. C. Ferris, G. Olivera, and T. R. Mackie. Optimizing the delivery of radiation to cancer patients. Mathematical Programming Technical Report 98-07, Computer Sciences Department, University of Wisconsin, Madison, Wisconsin, 1998.
- [29] W. Siedlecki and J. Sklansky. On automatic feature selection. *International Journal of Pattern Recognition and Artificial Intelligence*, 2(2):197–220, 1988.
- [30] M. Stone. Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society*, 36:111–147, 1974.