Condor and the Football Pool Problem

Or, How Condor Can Make you Rich Beyond Your Wildest Dreams

Jeff Linderoth

COR@L Lab
Lehigh University

The Annual Condor Clambake
University of Wisconsin-Madison
April 25, 2006

Thanks: NSF (OCI-0330607) and DOE (DE-FG02-05ER25694) for continuing support of this work

Application — Code Design

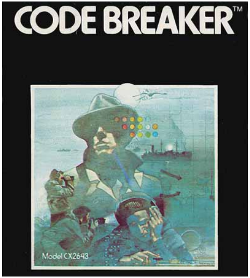

- $W(v, \alpha)$: Set of all "words" of length v from alphabet $\{0, 1, \ldots \alpha - 1\}$.
- $|W(\alpha, v)| = \alpha^v$
- We will abbreviate $W(v, \alpha) = W$
- A code is a subset $C \subseteq W$
- Hamming distance: $a \in W, b \in W$, $\text{dist}(a, b) = |\{i \mid a_i \neq b_i\}|$

Code Applications

Error Correcting Code

- Find $C \subset W$ such that $a \in C, b \in C \Rightarrow \text{dist}(a, b) \geq 2d + 1$
- Maximize $|C|$
- Application: Words in C submit over a "noisy" channel on which at most d letters are changed can be "self-corrected."

Covering Code

- Find a code $C \subset W$ such that every word $w \in W$ is at most a distance d away from at least one word in C
- $(\text{dist}(w, C) \leq d \ \forall w \in W)$
- Minimize $|C|$
- Application: Something far more practical

Are You Ready for Some Football!

- Predict the outcome of v soccer matches
- $\alpha = 3$
  - 0: Team A wins
  - 1: Team B wins
  - 2: Draw
- You win if you miss at most $d = 1$ games

The Football Pool Problem

What is the minimum number of tickets you must buy to assure yourself a win?

How Many Must I Buy?

Known Optimal Values

v	1	2	3	4	5
$\lvert C_v^* \rvert$	1	3	5	9	27

The Football Pool Problem

What is $\lvert C_6^* \rvert$?

- Despite significant effort on this problem for > 40 years, it is only known that

$$65 \leq C_6^* \leq 73$$

But It's Trivial!

- For each $j \in W$, let $x_j = 1$ iff we word j is in code C
- Let $A \in \{0,1\}^{|W| \times |W|}$
- $a_{ij} = 1$ iff word $i \in W$ is distance $\leq d = 1$ from word $j \in W$

IP Formulation

$$\min e^T x$$

$$\begin{array}{lrcl} \text{s.t.} & Ax & \geq & e \\ & x & \in & \{0,1\}^{|W|} \end{array}$$

Solving IPs in a Nutshell

- Problem is in general $\mathcal{NP}$-Hard
- Loads of theory and techniques going back > 40 years
- Workhorse algorithm is a tree-search procedure known as branch-and-bound.
- But really, branch-and-bound or its souped-up cousin branch-and-cut have been replaced in the most part by the new technique: give-it-to-CPLEX
- CPLEX: A commerical IP package that is putting integer programmers out of business.

CPLEX Can Solve Every IP

```
        Nodes                                        Cuts/
   Node  Left     Objective  IInf  Best Integer     Best Node    ItCnt     Gap

      0     0       56.0769   729                     56.0769     2200
*     0+    0                   0      243.0000       56.0769     2200   76.92%
*     0+    0                   0      110.0000       56.0769     2200   49.02%
                    56.5164   729      110.0000     Fract:  56    2542   48.62%
*     0+    0                   0      107.0000       56.5164     2542   47.18%
                    56.5279   729      107.0000      Fract:  6    2673   47.17%
*     0+    0                   0       94.0000       56.5279     2673   39.86%
*     0+    0                   0       93.0000       56.5279     2673   39.22%
Elapsed time =  90.03 sec. (tree size =  0.00 MB)
*    50+   50                   0       91.0000       56.5285    12242   37.88%
Elapsed time = 6841.16 sec. (tree size = 14.12 MB)
  31100 30002       60.1690   544       87.0000       57.1864  5467339   34.27%
  31200 30102       77.7888   216       87.0000       57.1864  5499451   34.27%
* 31200+28950                   0       86.0000       57.1864  5499451   33.50%
  31300 29044       58.9809   611       86.0000       57.1870  5511005   33.50%
Elapsed time = 9500.15 sec. (tree size = 18.70 MB)
  42700 39098       78.3242   197       85.0000       57.2845  7623200   32.61%
* 42740+36552                   0       83.0000       57.2845  7626440   30.98%
Elapsed time = 117349.90 sec. (tree size = 202.88 MB)
Nodefile size = 74.98 MB (61.52 MB after compression)
 465100 434311       66.8425   410       80.0000       58.0439 92473005   27.45%
```

Bounds

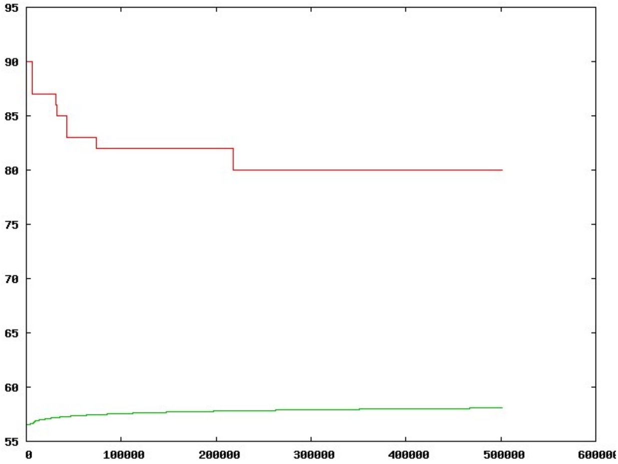

- Roughly 10^{30} seconds $\approx 10^8$ universe lifetimes in order to establish that $|C_6^*| \geq 73$

Jeff's Research Advice for the Day

How to solve problems that other people can't...

1. Be smarter than other people
   - Sadly, this strategy doesn't work for me
2. Get a more powerful computer than other people
   - Thank you Condor!!

The Optimal Solution

1. Collaborate with (mooch off of) smart people

AND

2. Get a more powerful computer

The Brains of the Operation

Carnegie Mellon

{ François Margot
Carnegie Mellon

THE UNIVERSITY of WISCONSIN MADISON

{ Greg Thain
Master Worker
University of Wisconsin-Madison

A Conversation With Miron

"What do you think I should talk about at the upcoming Condor Week, Miron?"

"Don't talk about your research, Jeff. Nobody cares about your research."

Miron Was Wrong!

"I love you, son. I think you should talk about your research if that makes you happy."

A Rock and a Hard Place

Should I listen to Miron or my Mom?

My Mom, Naturally!

Solutions for $v = 3$

Answer #1

1	2	3
0	0	0
1	1	0
1	0	1
0	1	1
2	2	2

Answer #2

1	2	3
2	0	0
1	1	0
1	0	1
2	1	1
0	2	2

- These solutions are isomorphic.
  - For first component: $2 \leftrightarrow 0$
- There are LOTS of isomorphic solutions:
  1. "Rename" W,L,D for any subset of the matches: $(\alpha!)^v$
  2. Reorder the matches: $v!$
- There are $(\alpha!)^v v! = 1296$ equivalent solutions for $v = 3$

What's the Problem!?: Symmetry

- π : Permutation of $\{1, 2, \ldots n\}$
- $\pi(x) = \pi(x_1, x_2, \ldots x_n) = (x_{\pi(1)}, x_{\pi(2)}, \ldots x_{\pi(n)})$
- π is a symmetry of an IP if
  1. x feasible $\Leftrightarrow \pi(x)$ feasible
  2. $c^T x = c^T \pi(x)$
- $G(IP)$: Set of all symmetries of IP
- For covering design, $|G(IP)| = v!(\alpha)^v$
- $6! \times 3^6 = 524880$

Isomorphism Pruning

- For some permutation $g \in G(IP)$ and set of indices $S \subset \{1, 2, \ldots, n\}$, let

$$g(S) = \{g(i) \mid i \in S\}$$

- At a node a of the branch-and-bound tree
  - $F_1^a = \{i \mid x_i \text{ fixed to 1 at } a\}$
  - $F_0^a = \{i \mid x_i \text{ fixed to 0 at } a\}$
- Nodes a and b are isomorphic if

$$\exists g \in G(IP) \text{ with } g(F_1^a) = F_1^b, g(F_0^a) = F_0^b$$

- You may prune one of a or b. [Bazaraa, Kirca 83]

Minimum Index Branching

[Butler, Ivanov, Lam, Margot, McKay, Read, Stinson, ...]

- The set $\{g(S) \mid g \in G(IP)\}$ is an equivalence class of all equivalent "relabelings" of $S \subseteq \{1, 2, \ldots, n\}$
- Choose one representative for each potential set of variable fixings.
- For example, in Minimum Index Branching, A set S is a representative of its equivalance class if

$$S = \text{lexmin}\{g(S) \mid g \in G(IP)\}$$

- Isomorphism Pruning:
  - If F_1^a is not a representative, then prune node a.

Results (All Thanks to François)

The Good

- For $d = 1, v = 5, \alpha = 3$, Isomorphism Pruning can establish $|C_5^*| = 27$ in 1409 nodes, 82 seconds.
- CPLEX (v9.1) does not solve the problem in more than 4 hours.

The Bad and Ugly

- For $d = 1, v = 6, \alpha = 3$, Isomorphism Pruning gets nowhere..
- $|C_6^*| \geq 61$ after *long* running time.

Subcodes

- Partition W into words that start with each letter
- $W(6, 3) = W_0 \cup W_1 \cup W_2$
- $w \in W_0$ covers 11 words in W_0
- $w \in W_1$ covers 1 words in W_0
- $w \in W_2$ covers 1 words in W_0
- An optimal code has
  - $C_0^* \subset W_0, |C_0^*| \stackrel{\text{def}}{=} y_0$
  - $C_1^* \subset W_1, |C_1^*| \stackrel{\text{def}}{=} y_1$
  - $C_2^* \subset W_2, |C_2^*| \stackrel{\text{def}}{=} y_2$
- So if a code of size $|C^*| = M$ exists, then it must satisfy

Covering System

$$\begin{aligned} 11y_0 + y_1 + y_2 &\geq 243 \\ y_0 + 11y_1 + y_2 &\geq 243 \\ y_0 + y_1 + 11y_2 &\geq 243 \\ y_0 + y_1 + y_2 &= M \end{aligned}$$

Repeat!

- This same idea applies to fixing $m > 1$ component

Nice Idea! (Östergård and Weakley)

- Pick an $M = |C^*|$
- Enumerate all non-isomorphic solutions to covering system for $m = 1, 2, \ldots 6$
- If for some m, there are no solutions, then $M + 1$ is a valid lower bound on $|C^*|$
- Östergård and Weakley (2000): Able to show $M = 62$ is optimal code length for $d = 1, v = 9, \alpha = 2$
- Östergård and Wassermann (2002): Able to show $M \geq 65$ for $d = 1, v = 6, \alpha = 3$. Required over 1 CPU year!

Extending the Idea

A New Idea!

Combine the subcode fixing with IP.

- For some (small) m, and optimal code size M, enumerate all non-isomorphic solutions to the covering system
- This gives a list of possible y values, e.g. for $m = 1$ you get a list of triples (y_0, y_1, y_2)
- For each member of the list, solve the "Sequence IP"

Sequence IP (M, y_0, y_1, y_2)

$$\begin{aligned} \min\ & e^T x \\ \text{s.t.}\ & Ax \geq e \\ & \sum_{i \in W_0} x_i = y_0 \\ & \sum_{i \in W_1} x_i = y_1 \\ & \sum_{i \in W_2} x_i = y_2 \\ & e^T x \leq M \\ & x \in \{0, 1\}^{|W|} \end{aligned}$$

Improving the Lower Bound on C_6^*

If you find no solution, then $M + 1$ is a valid lower bound

Great Research

The End Result

We need to perform many smaller branch-and-bound calculations

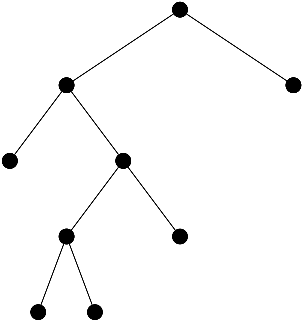

Grid Programmers Do It In Parallel

- Nodes in disjoint subtrees could be evaluated independently
- But this is not a ~~embarassingly~~ pleasantly parallel operation

Master-Worker!

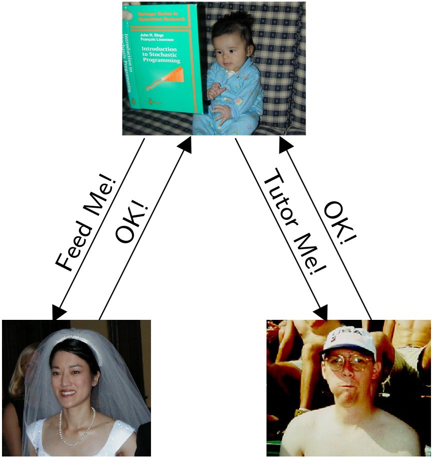

Important Points

- Master-Worker is a flexible, powerful framework for Grid Computing
- The MW software library can help you quickly write master-worker applications that run over Condor-provided computational grids
- Please come talk to Greg, Steve Wright, or myself about setting you up to run with MW

Mechanism #1—Flocking

Building My Grid

- Currently flock to UW-Madison CS Pool
- Large Condor Pool at NCSA
- Small pool in our lab at Lehigh: COR@L

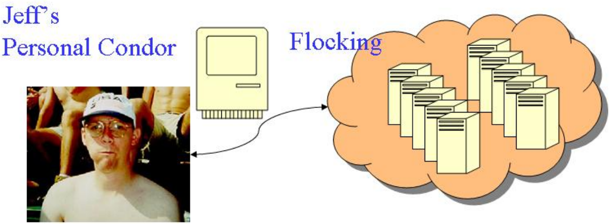

Mechanism #2: Hobble-In

1. Log in to Teragrid site
2. Install condor_startd and condor_starter
3. Configure condor to report to your condor master
4. Write PBS or PSF scripts that run "preconfigured" condor startd's for limited about of time.
5. Submit away. Machines will appear in your personal pool when local scheduler deigns to run them.
6. Currently working at the following Teragrid sites:
   - Argonne/UC, NCSA, SDSC, TACC

More Questions?

I'm happy to send anyone my hobble in scripts...

Building Your Grid: Flocking + Hobble-In

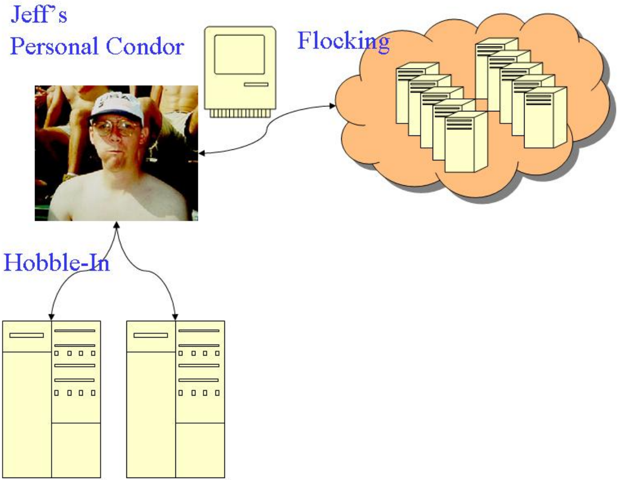

Mechanism #3: `sshIdle-In`

- 300 Machines at Lehigh are on running Condor on a private network: And I wanted to use them!

`sshIdle-In`

- A GCB-free mechanism for running on private networks
- For use currently with `MWSocket` RMComm only

1. Look up port number where master is listening: default 8997.
2. Log into "landing point" – machine with connections to both private network and outside world – preferably condor_master
3. Transfer (or build) worker executables on landing point
4. Forward traffic:

```
ssh -l fmargot -g -N -f -L 8997:0.0.0.0:8997 meisterbrau.cs.wisc.edu
```

5. Submit worker executables in private network

Sample sshIdle-In Script

```
Universe = Vanilla
Executable = mw_exec0.$$(Opsys).$$(Arch).exe
arguments = 1512 8997 8997 192.168.1.1
should_transfer_files = Yes
when_to_transfer_output = ON_EXIT
rank = Mips
on_exit_remove = false
nice_user = true
queue

arguments = 1513 8997 8997 192.168.1.1
queue
arguments = 1514 8997 8997 192.168.1.1
queue
arguments = 1515 8997 8997 192.168.1.1
queue
```

Flocking + Hobble-In + sshIdle-In

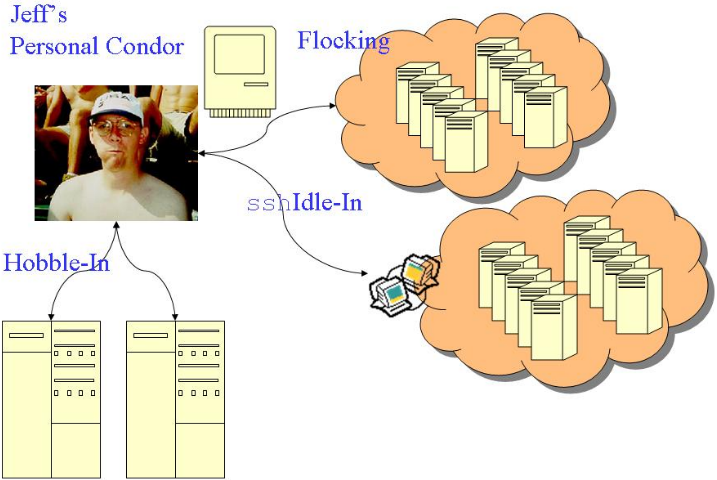

Other Mechanisms You Should Probably Use

- Condor-Glidein: That's what it was designed for!
- Gridshell: GridShell/Condor creates virtual Condor pools across TeraGrid clusters connected through Globus.
    - `http://gridshell.net/`
    - Talk to Ed Walker

It Begs The Question

- Before you ask, I don't have a good excuse for using hobble-in other than it works for me.
    - I don't need to rely on local admins to configure gateway properly
    - Scripts for various Teragrid sites (to get condor startd's talking to condor master) were already configured/set-up before I knew about Gridshell.

Mechanism #3: IceCream-In

Stardate: Date: Thu, 20 Apr 2006 16:22:28 -0500

A Conversation Over Ice Cream...

"I really like all these mechanisms, Miron, but where and how can I get some more workers?"

"There's this thing called the Open Science Grid. I'll see what I can do..."

Greg is My Hero!

Stardate: Date: Thu, 20 Apr 2006 17:15:47 -0500

```
Jeff:
The Open Science Grid has asked us for some more jobs
to demonstrate some of our Grid technologies, so I'm
working on getting some of the football pool jobs routed
over there.  Hang on for some Super Grid Power!  I need
to reboot the personal condor first, so don't worry if
the ff3 output looks a little wonky.
-greg
```

Workers in Solving One Batch of Subproblems

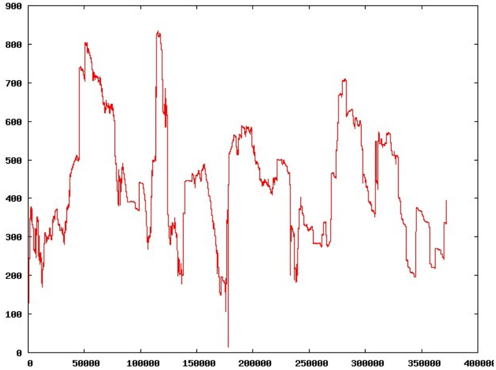

- This is the closest thing I've seen to a true "Grid" reality: I'm eating an ice cream cone, and I get 300 new workers.

Jeff's Super-Duper Computational Grid

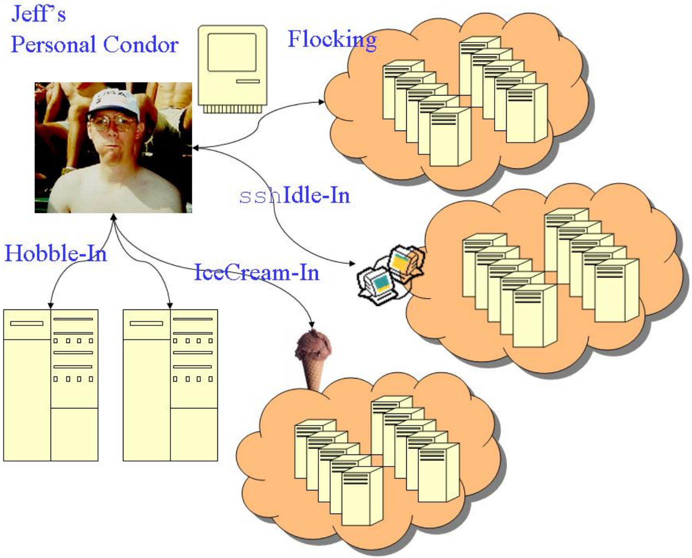

$|C_6^*| \geq X???$

- To establish that $|C_6^*| \geq 70$ Final total: 346 sequence IPs
- François and Jeff implemented these ideas using the old MW-FATCOP framework for MILP.
- Many of the engineering/tuning ideas from the QAP experience were used in this implementation
- Greg has made numerous improvements in MW's robustness and has been great at scavenging cycles in an effort to solve this problem

Woo Hoo!!!!!

Breaking News

Deal colleagues, we are pleased to annouce that we have improved the lower bound on the cardinality of a covering code of radius one for the Hamming Space $\mathbb{Q} = \mathrm{F}_3^6$ to 70.

Or Have We!?!?!?!

Stardate: Tue, 11 Apr 2006 12:41:29 -0400 (EDT)

From: Francois Margot <fmargot@andrew.cmu.edu>
To: Jeff Linderoth <jtl3@Lehigh.EDU>
cc: gthain@cs.wisc.edu
Subject: Re: BooYah!

Now, remove all sharp objects from your desk. I noticed a mistake in the codbt06_fix2_reg01236_s.mylp.gz file (a permutation of the inequalities). I fixed the file, but I am afraid that we will have to rerun the *69s_*.seq problems. However, I can reduced the number of sequences in the files, taking into account the fact that every 5-subcode must have at least 22 words. I will let you steam a little bit before doing so ...

Doh!


```
jeff@pbr:  diff correct.lp incorrect.lp
1c1
< 738 729 1
---
> 739 729 1
```

Condor Allows You To Make Mistakes

- We (read François) wasted "only" 19.76 CPU Years.
- If we had to rely only on machines for which I had to write grants (i.e. Teragrid), we never would have been able to make progress

We're Still Working

- I really,really, really wanted to announce that $C_6^* \geq 70$ today.
- But it hasn't been for a lack of trying

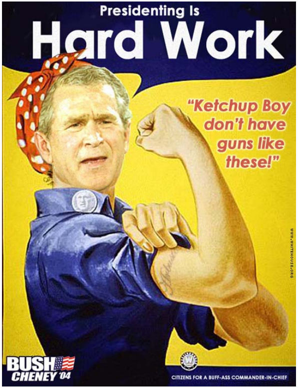

Statistics so far...

Wall Time	47.15 days
CPU Time	57.31 years
Avg Workers	467.3
Max Workers	1253
Total Nodes	8.37×10^8
Total LP Pivots	6.53×10^{11}
Parallel Performance	94.9%

Conclusions

- There's lots of CPU power out there for the taking
- There are lots of mechanisms for you to build your Grid via Condor—Flocking, Glide-in, GridShell, Hobble-In, `ssh`Idle-In, IceCream-In (or schedd on the side)
- Many computations can be put in a master-worker framework: Maybe yours can too!
- I'd be happy to help you in using MW.

The MWAgenda

The MWPresent

- MW works on Windoze (cygwin socket)
- MW has "Black Box" capabilities, so you need not write any code other than code for the MWDriver.

The MWFuture

- Making `ssh`-idle in more automatic
- Easier Interfaces?:
- MWBOINC?

We wants you to help us steer the MWShip

Let's Tawlk

- I'm lonely—Please come talk to me about MW or my research.
- MW: `http://www.cs.wisc.edu/condor/mw`
- COR@L: `http://coral.ie.lehigh.edu`
- `mailto:jtl3@lehigh.edu`

- Thanks (again) to NSF (OCI-0330607) and DOE (DE-FG02-05ER25694) for continuing support of this work