

Fault Tolerance in APPSPACK

Asynchronous Parallel Pattern Search
for Derivative-Free Optimization

TAMARA G. KOLDA
Sandia National Labs

USING PATTERN SEARCH

Our goal is to solve the unconstrained optimization problem

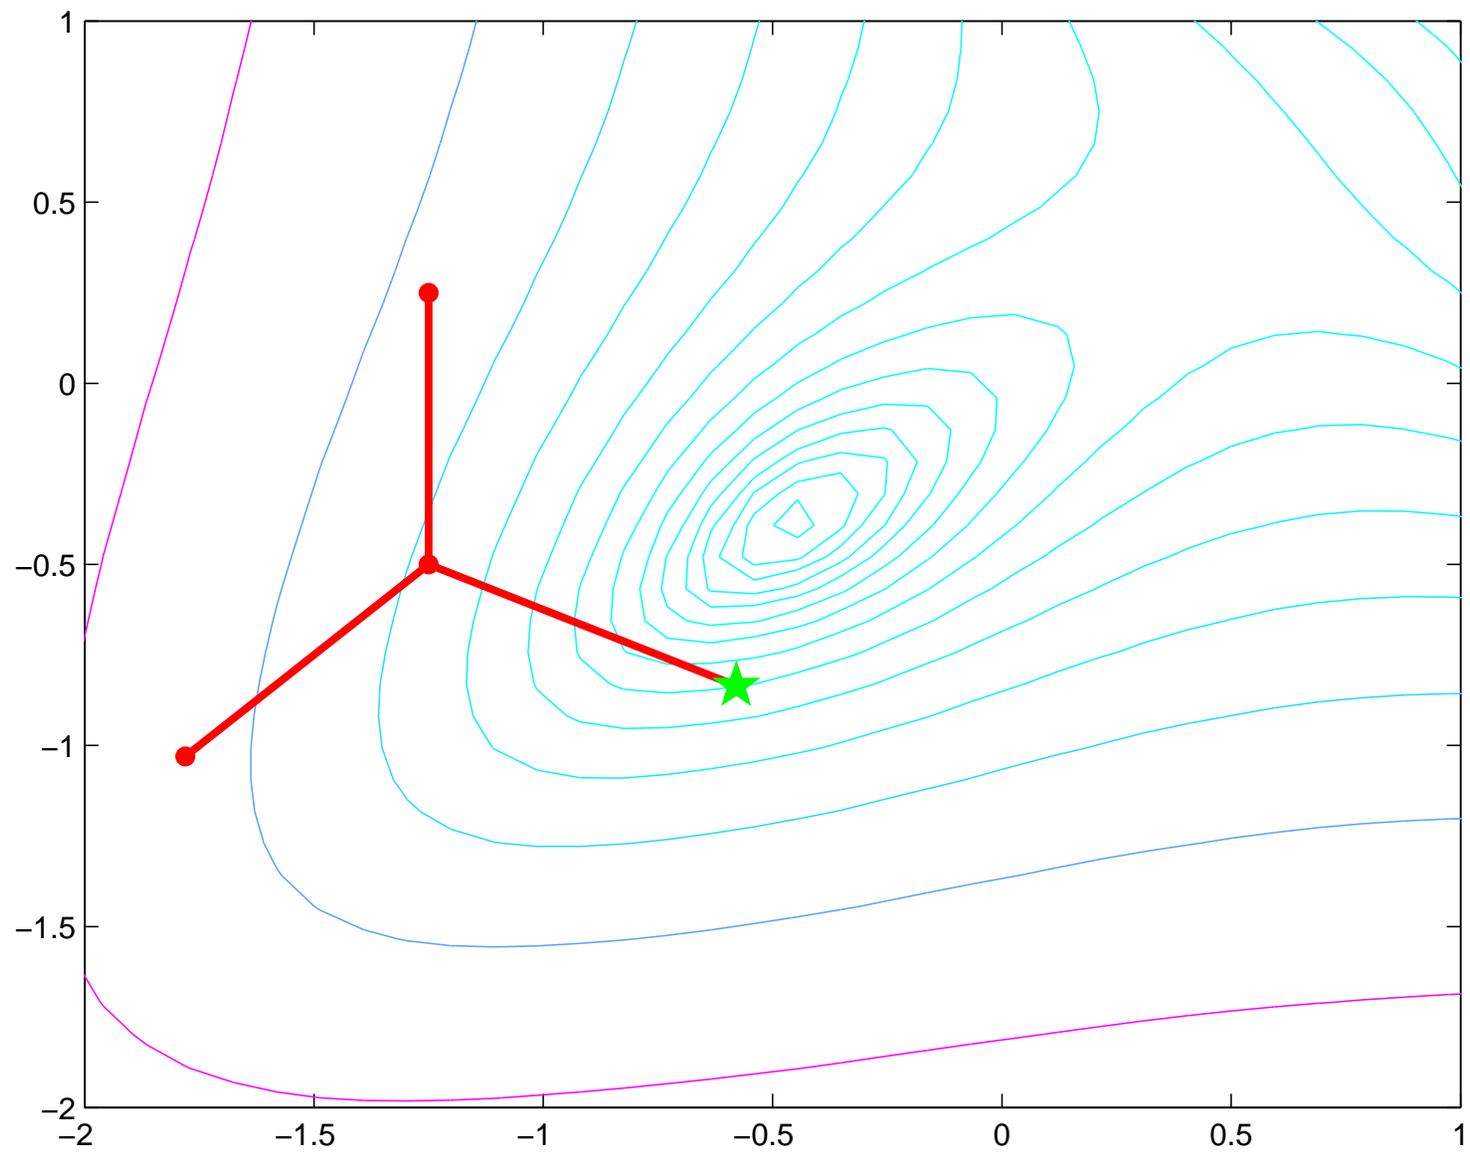
$$\min f(x) \quad x \in \mathbb{R}^n$$

without derivative information. Pattern search is a popular *derivative-free* method which is most useful when...

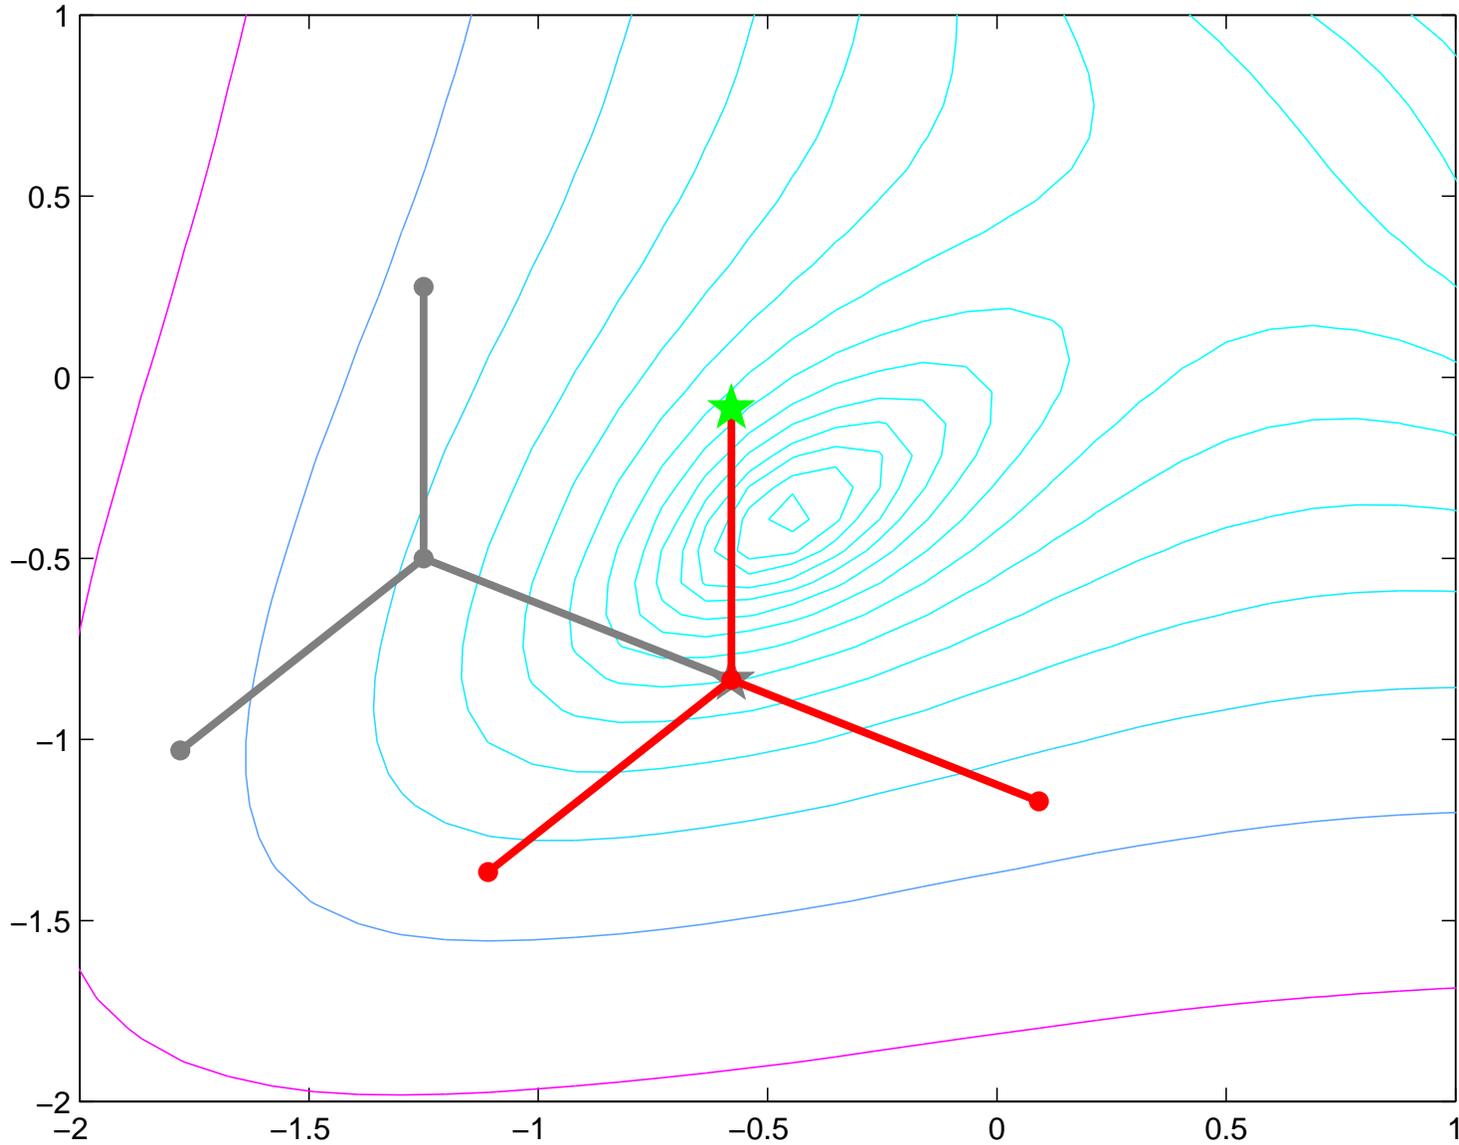
- The function is expensive to calculate.
- The gradient cannot be calculated.
- Numerical approximation of the gradient is too slow, or the function values are too noisy to yield reliable gradient approximations.

We are motivated by engineering optimization applications.

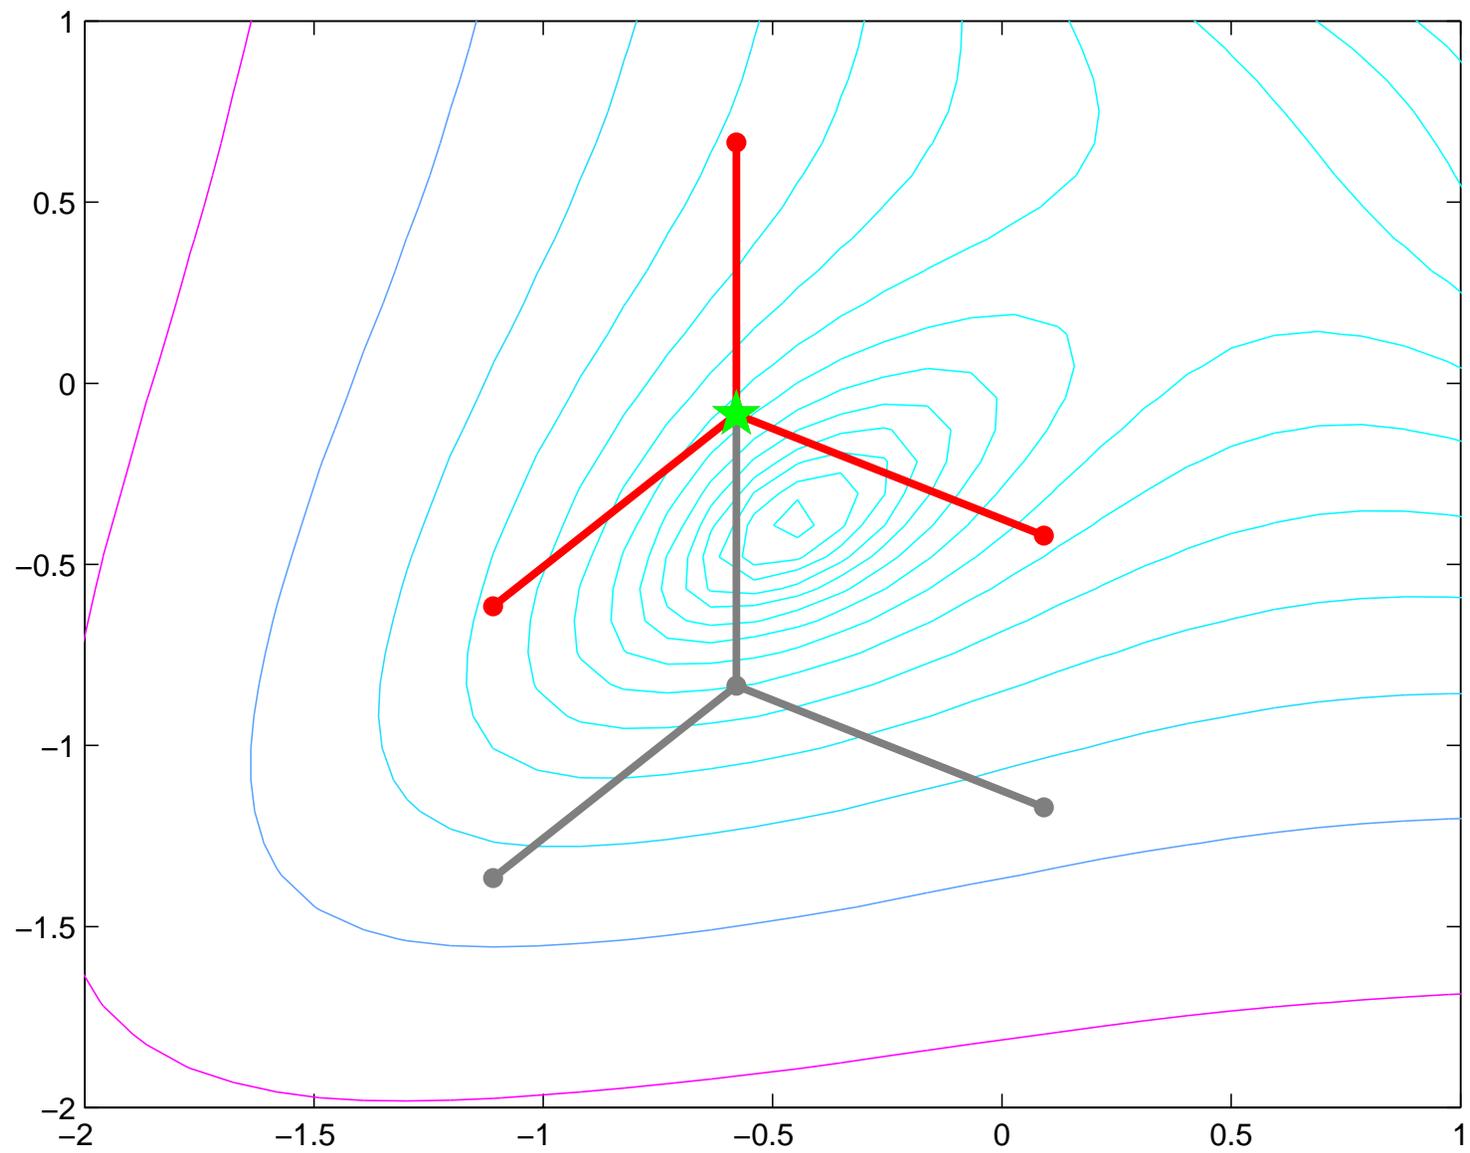
PARALLEL PATTERN SEARCH



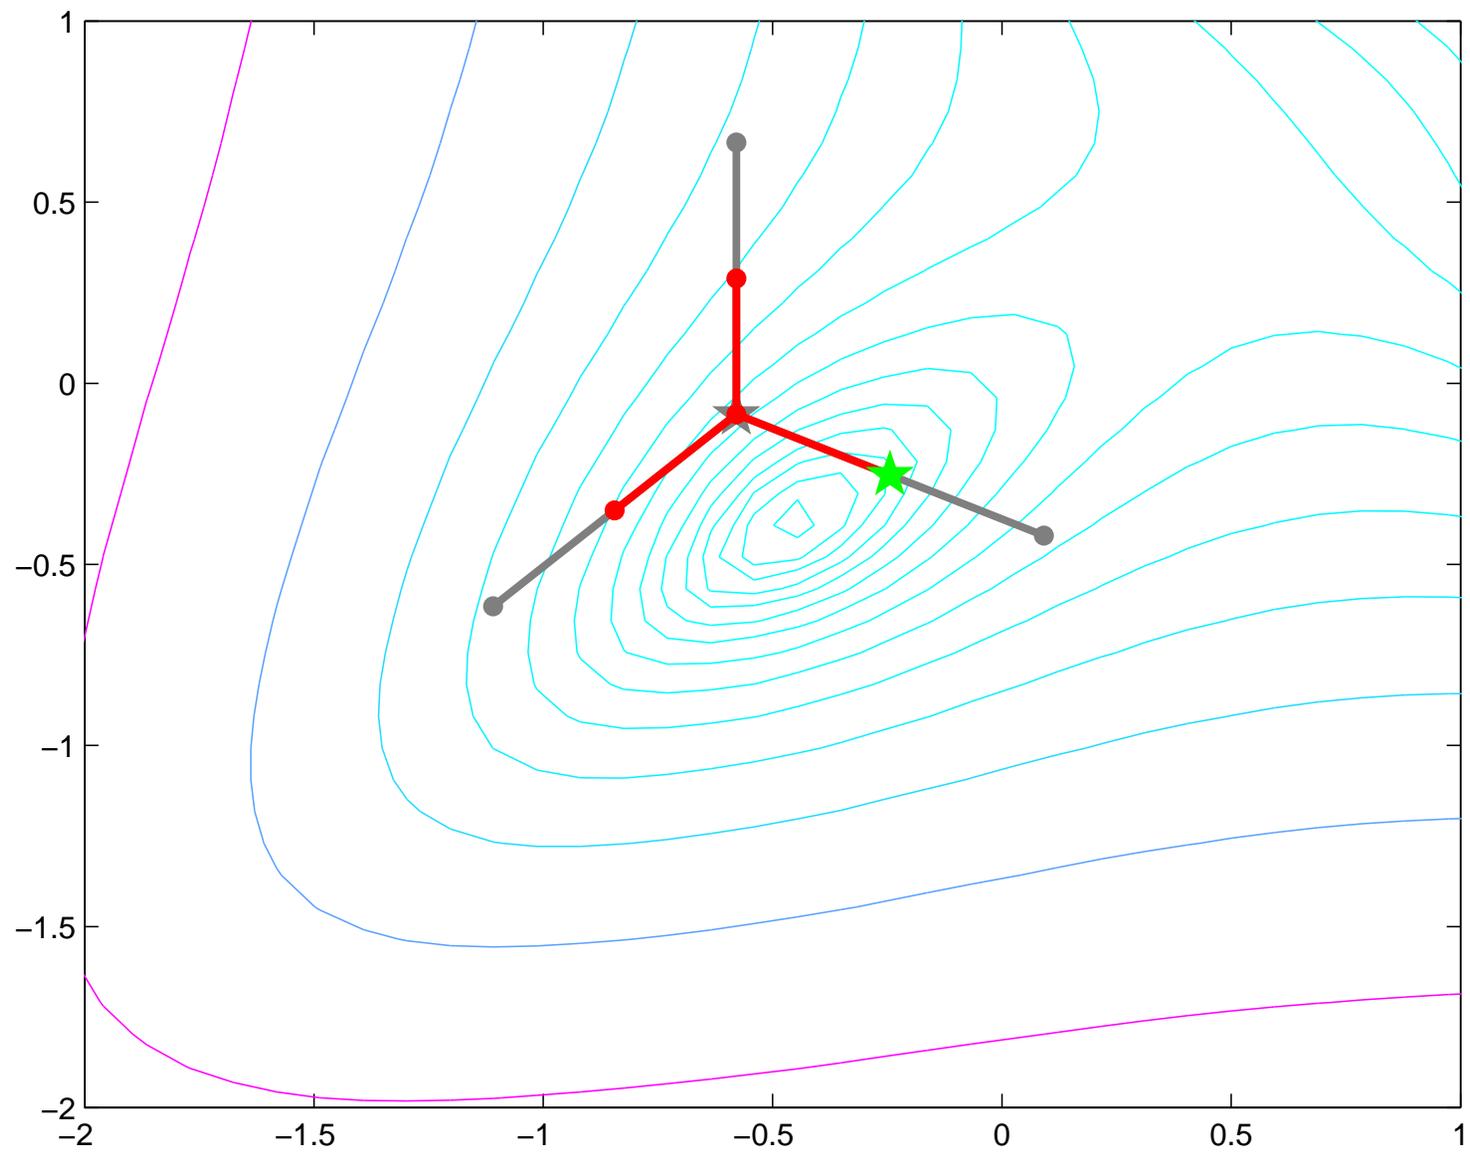
PARALLEL PATTERN SEARCH



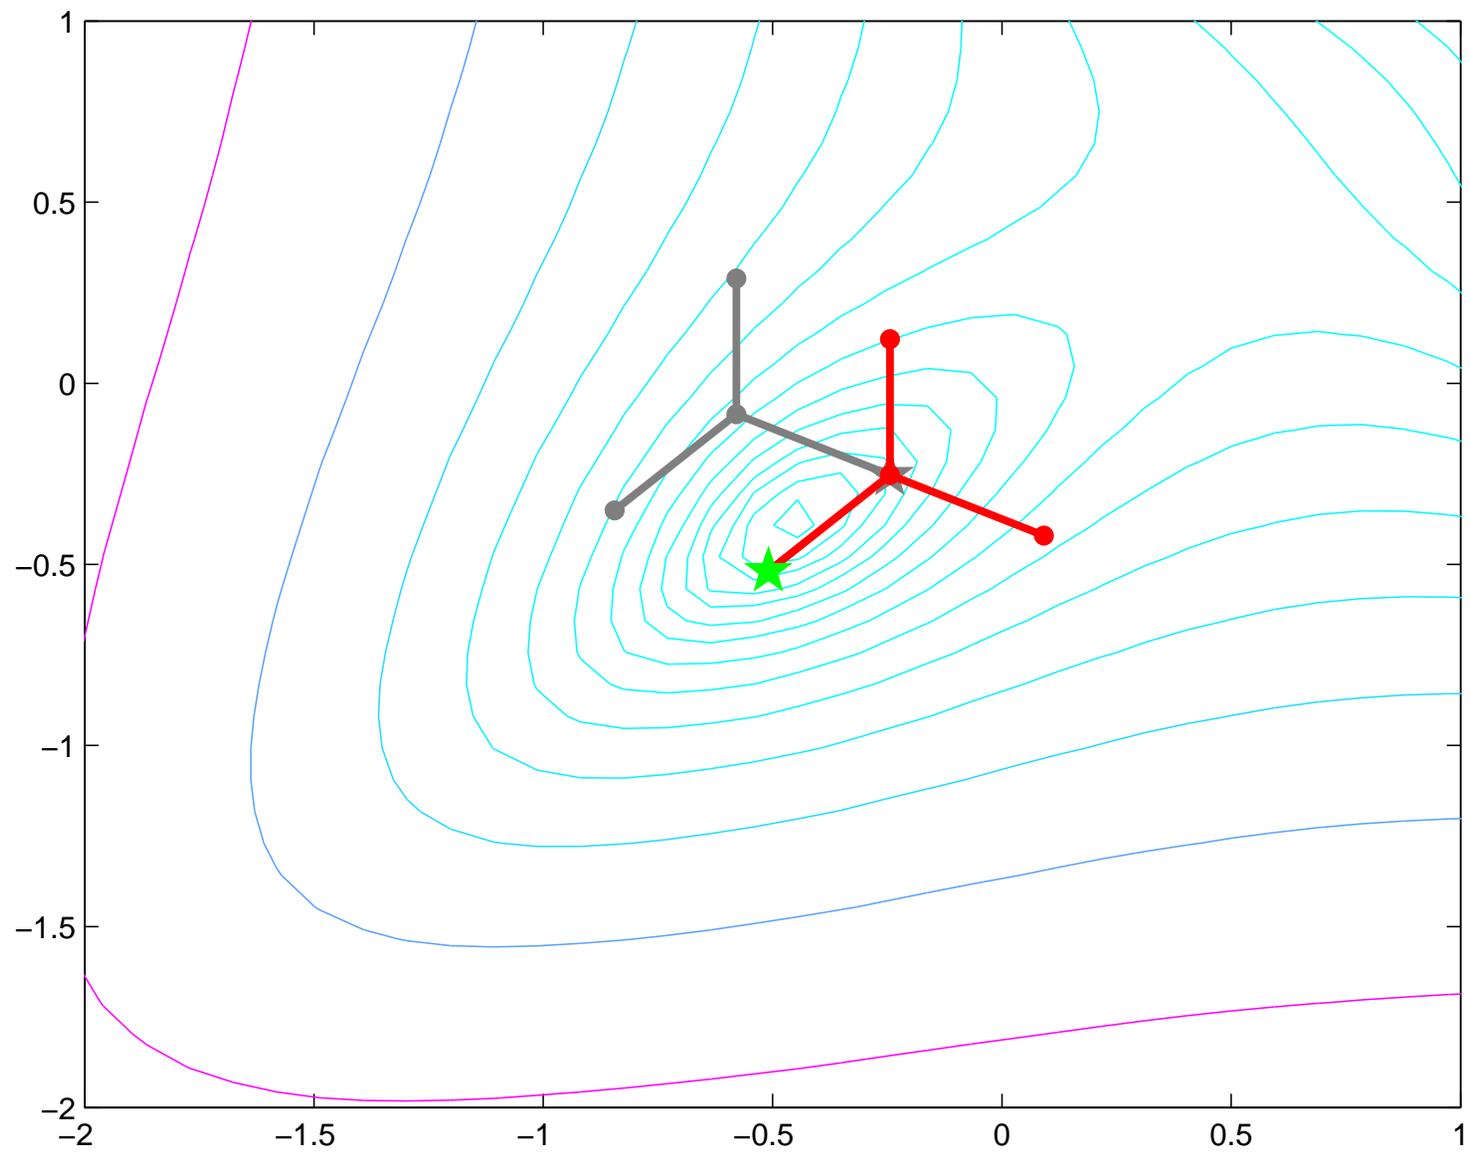
PARALLEL PATTERN SEARCH



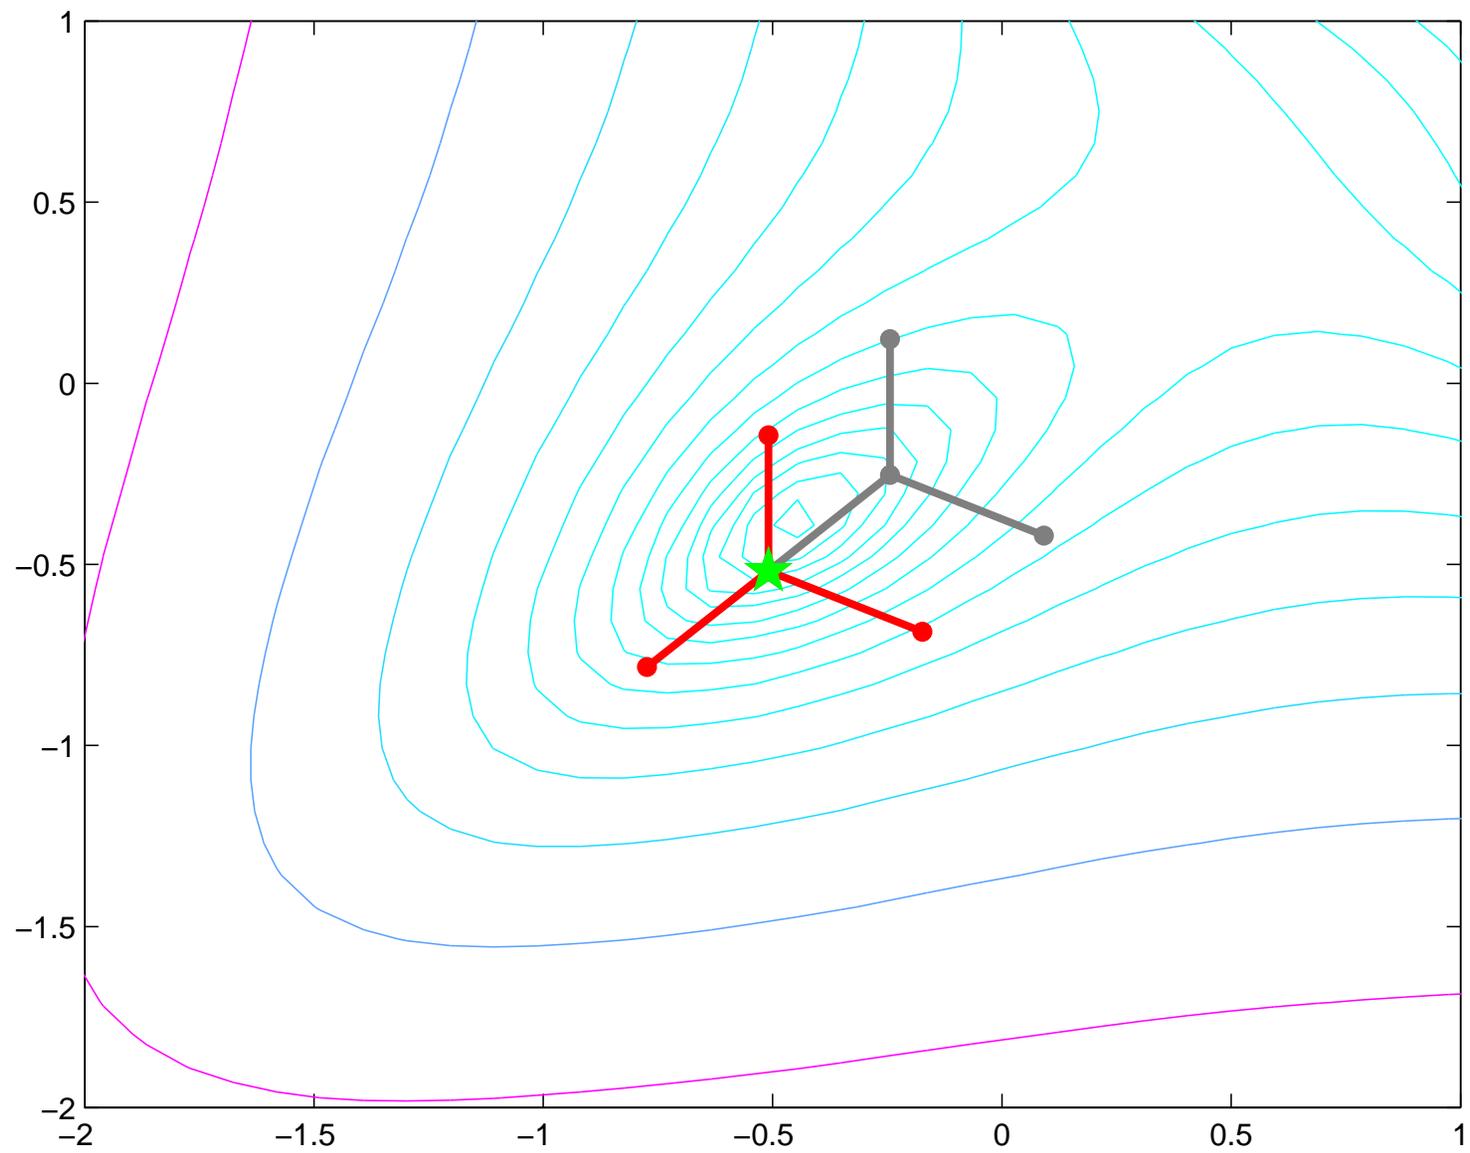
PARALLEL PATTERN SEARCH



PARALLEL PATTERN SEARCH



PARALLEL PATTERN SEARCH



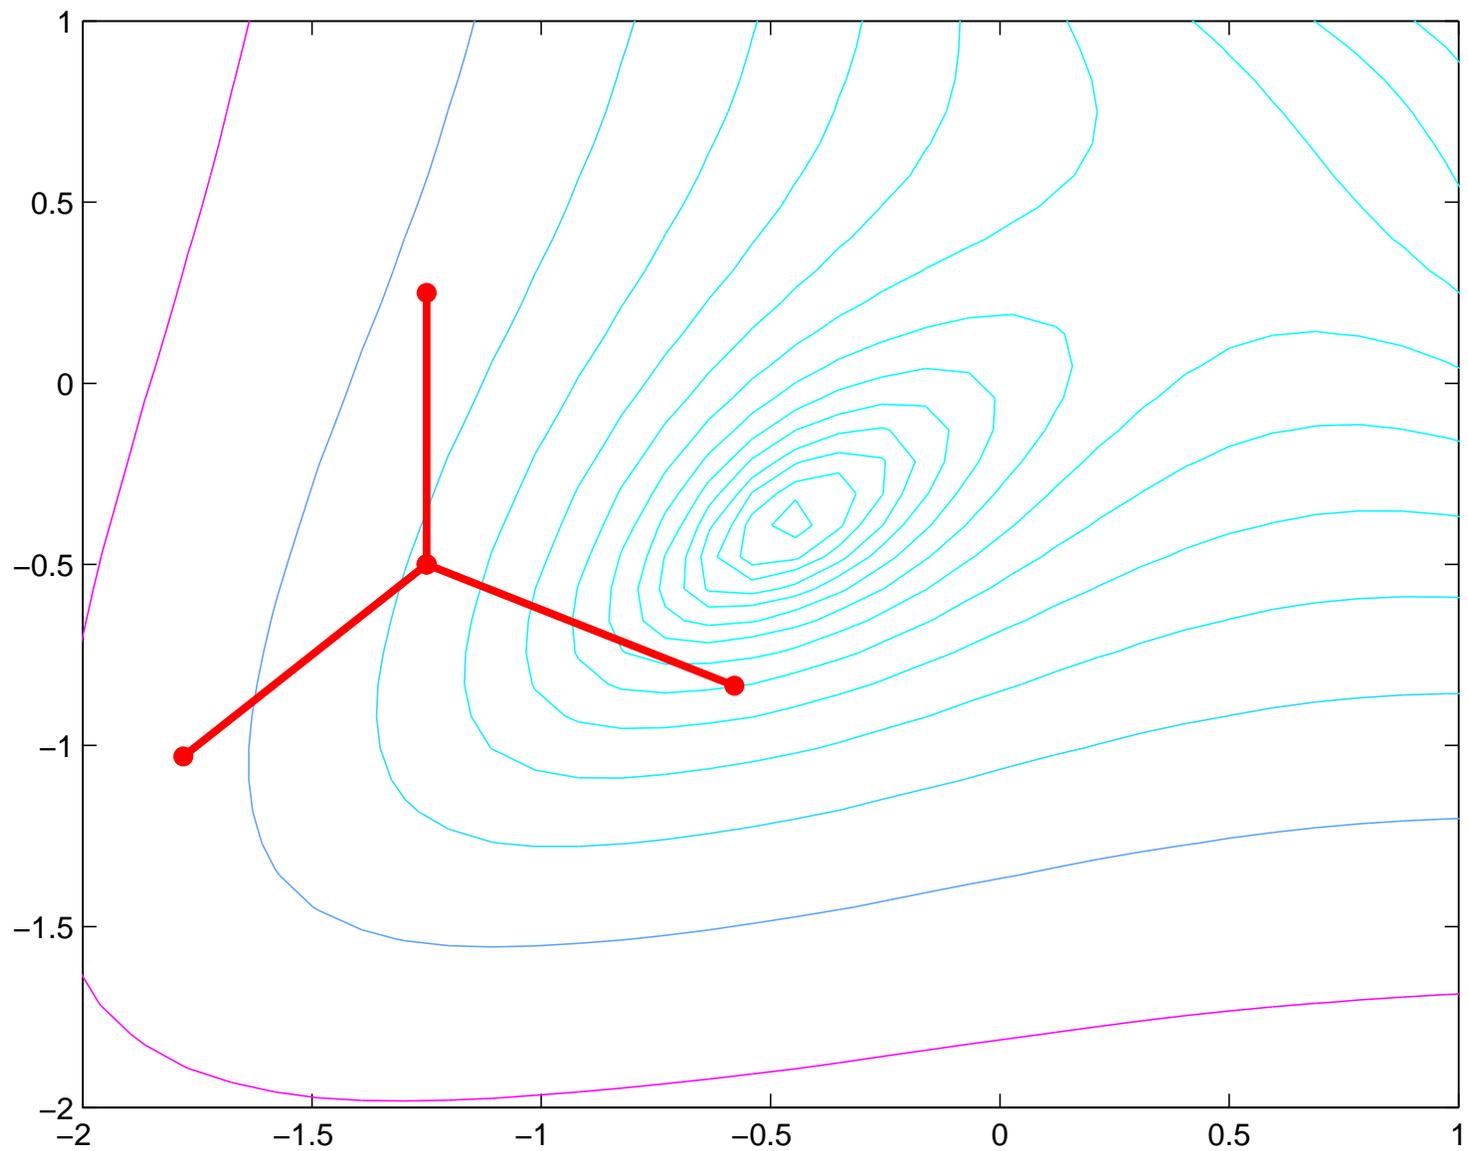
ASYNCHRONOUS PPS REDUCES IDLE TIME

- Remove synchronization bottleneck
 - Different running times for simulations
 - Varying processor loads
 - Heterogeneous architectures
- Easier to incorporate fault tolerance
 - COTS systems

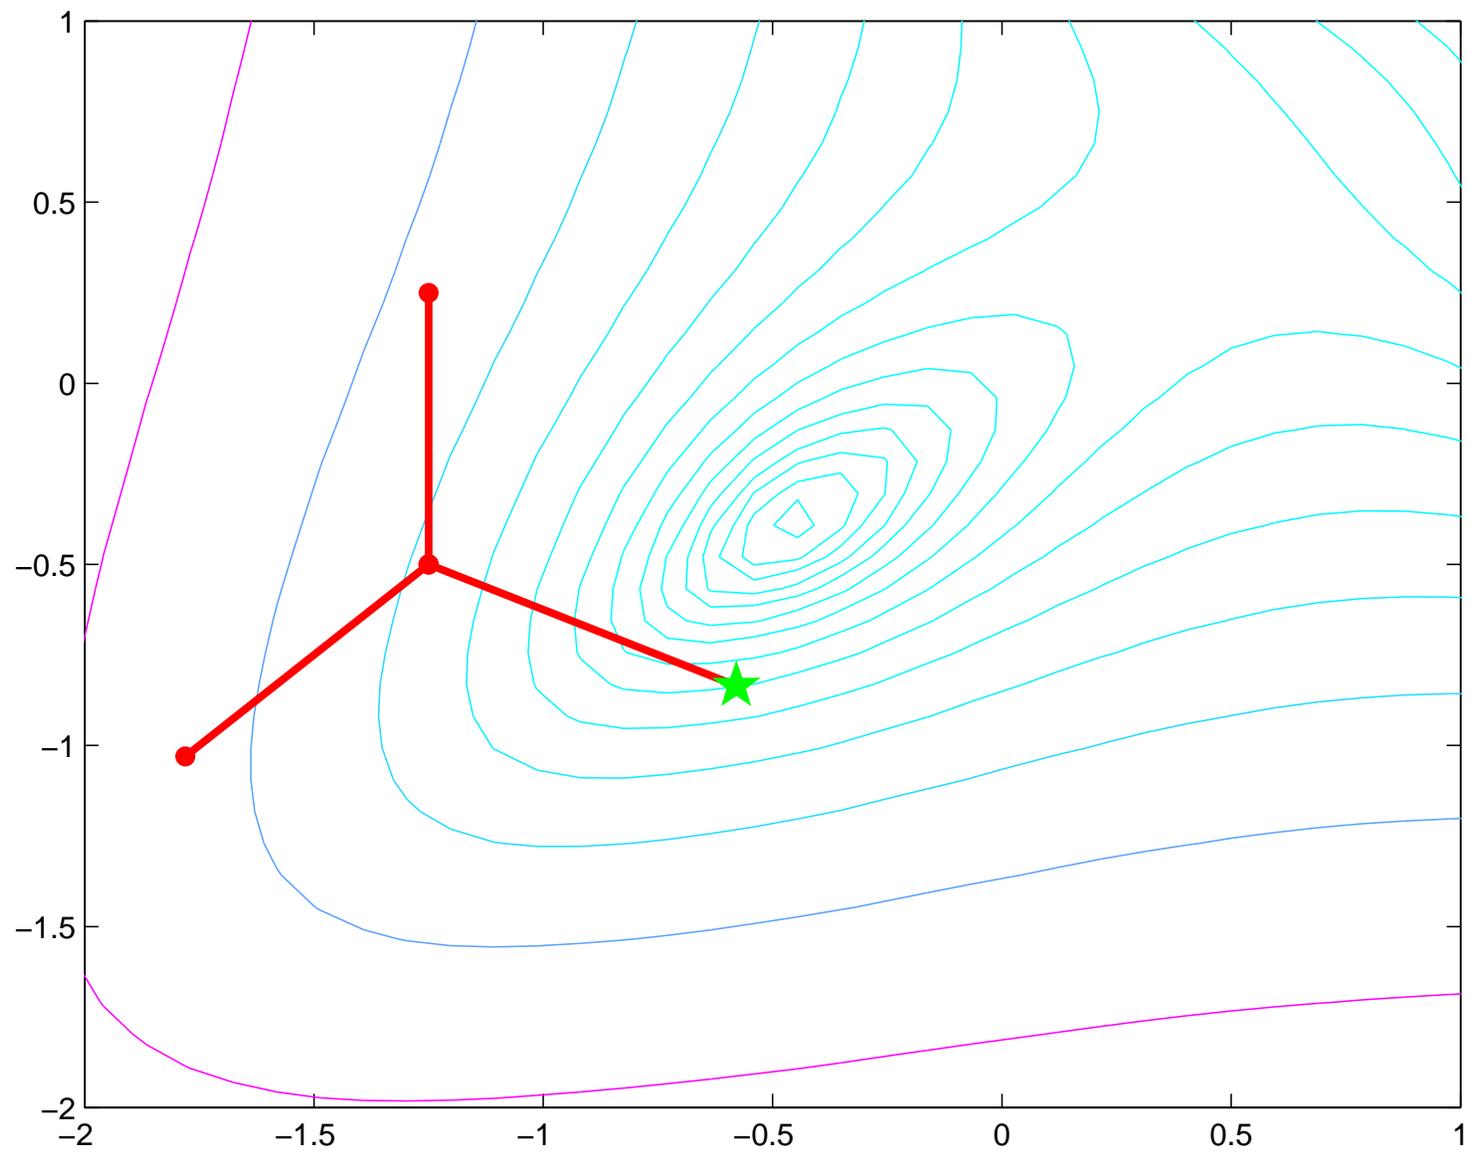
APPS/APPSPACK COLLABORATORS

- Patty Hough
- Virginia Torczon (William & Mary)
- Alton Patrick (Summer Intern)
- Sarah Brown (Summer Intern)

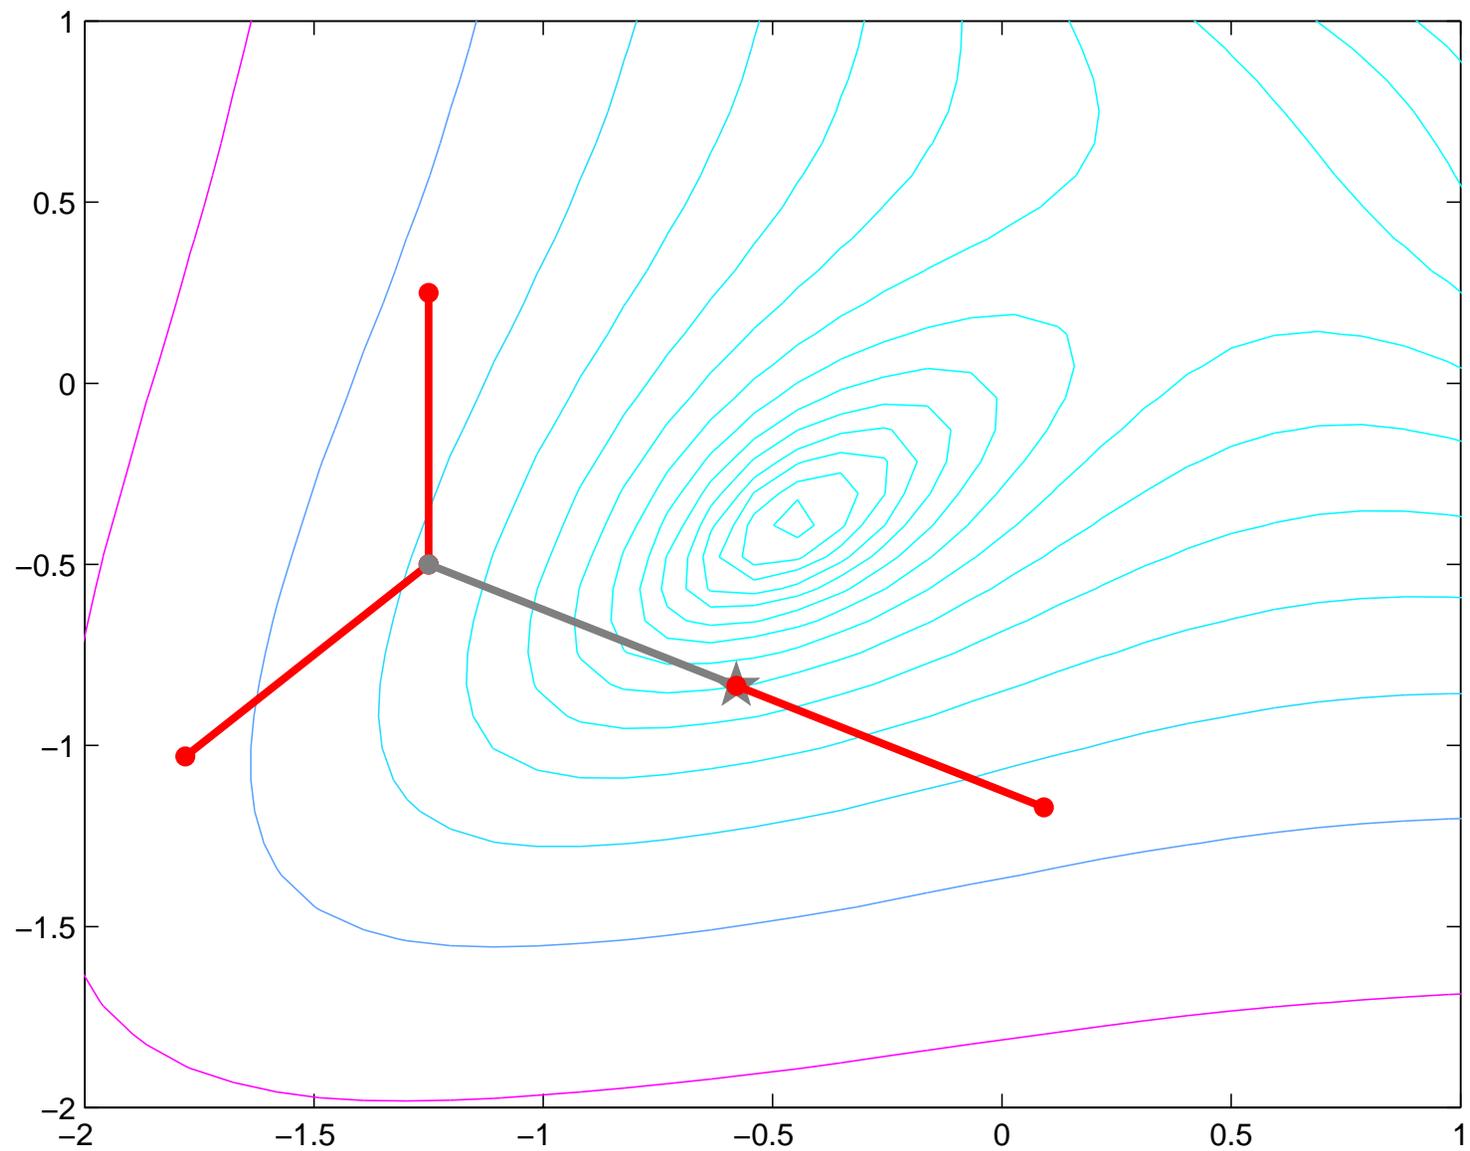
ASYNCHRONOUS PARALLEL PATTERN SEARCH



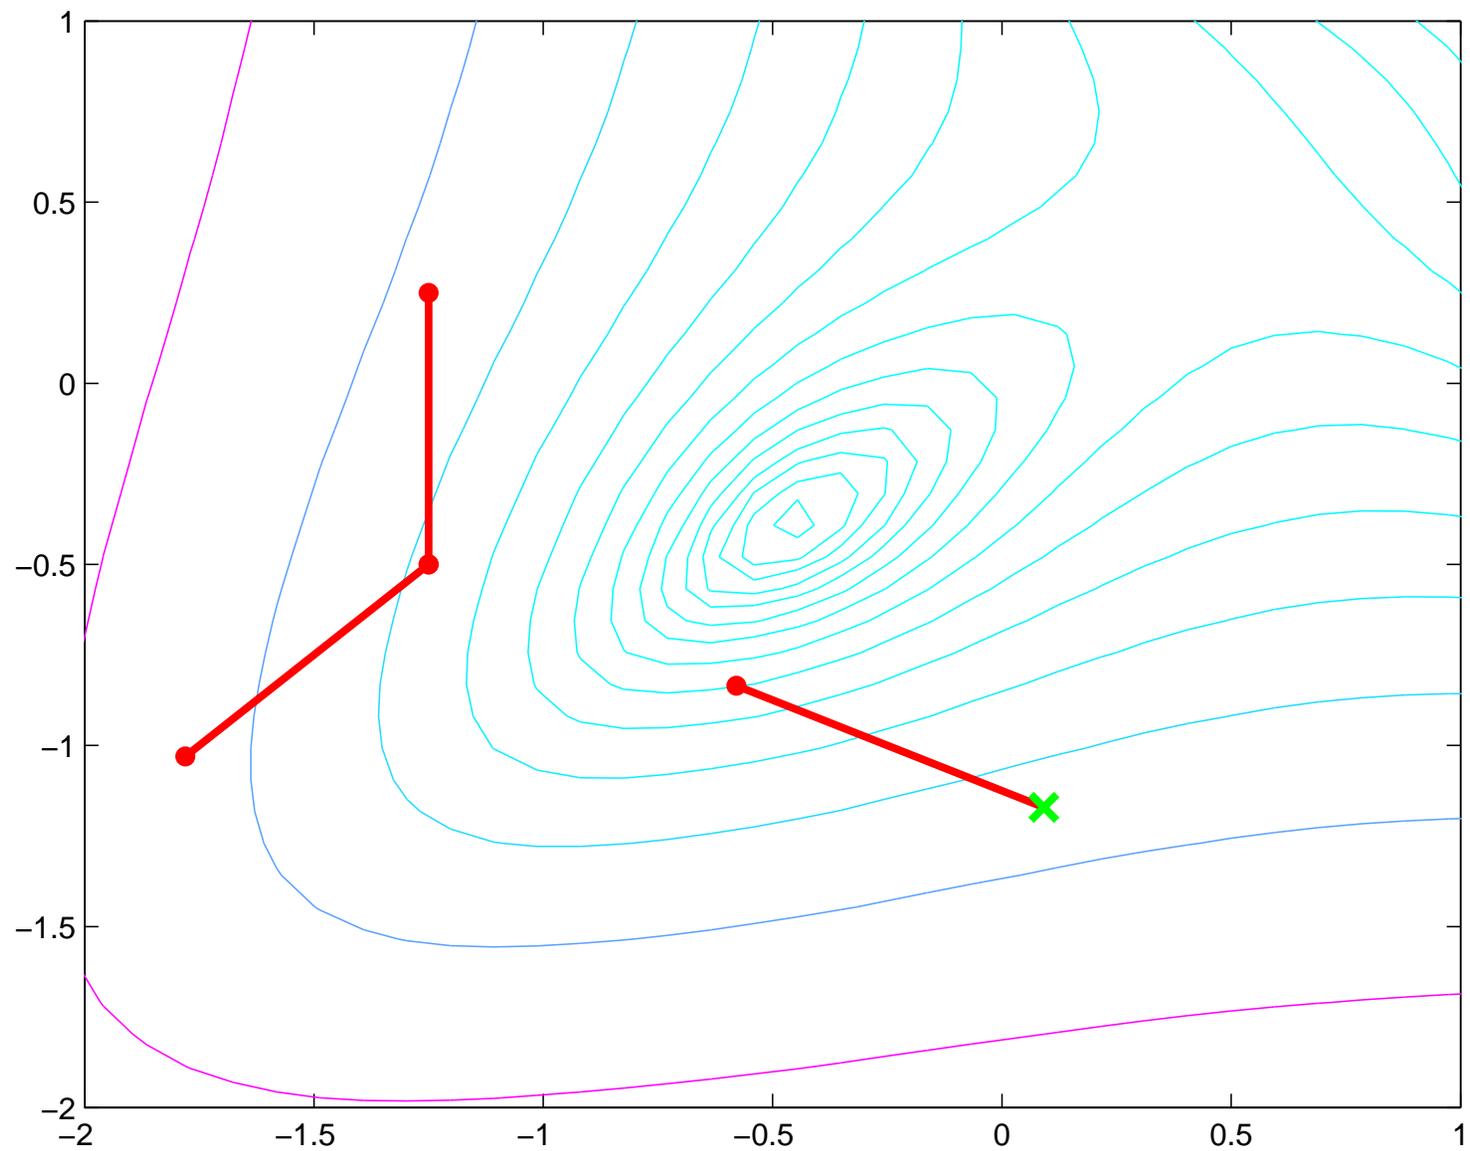
ASYNCHRONOUS PARALLEL PATTERN SEARCH



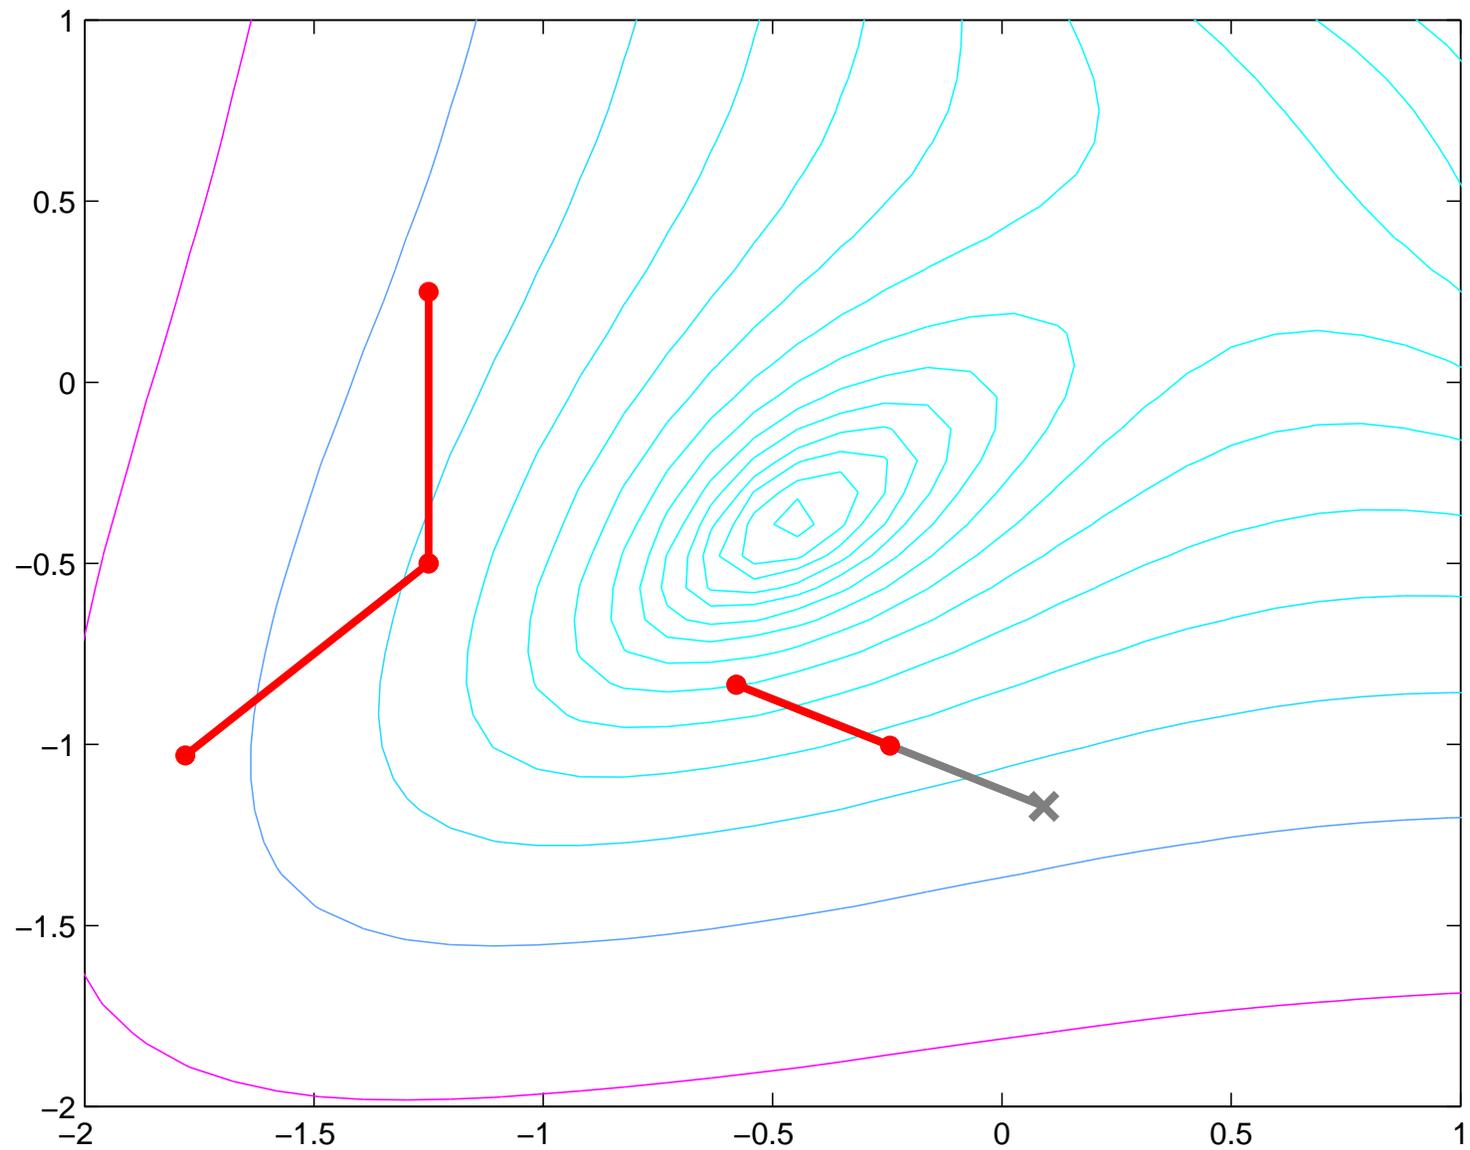
ASYNCHRONOUS PARALLEL PATTERN SEARCH



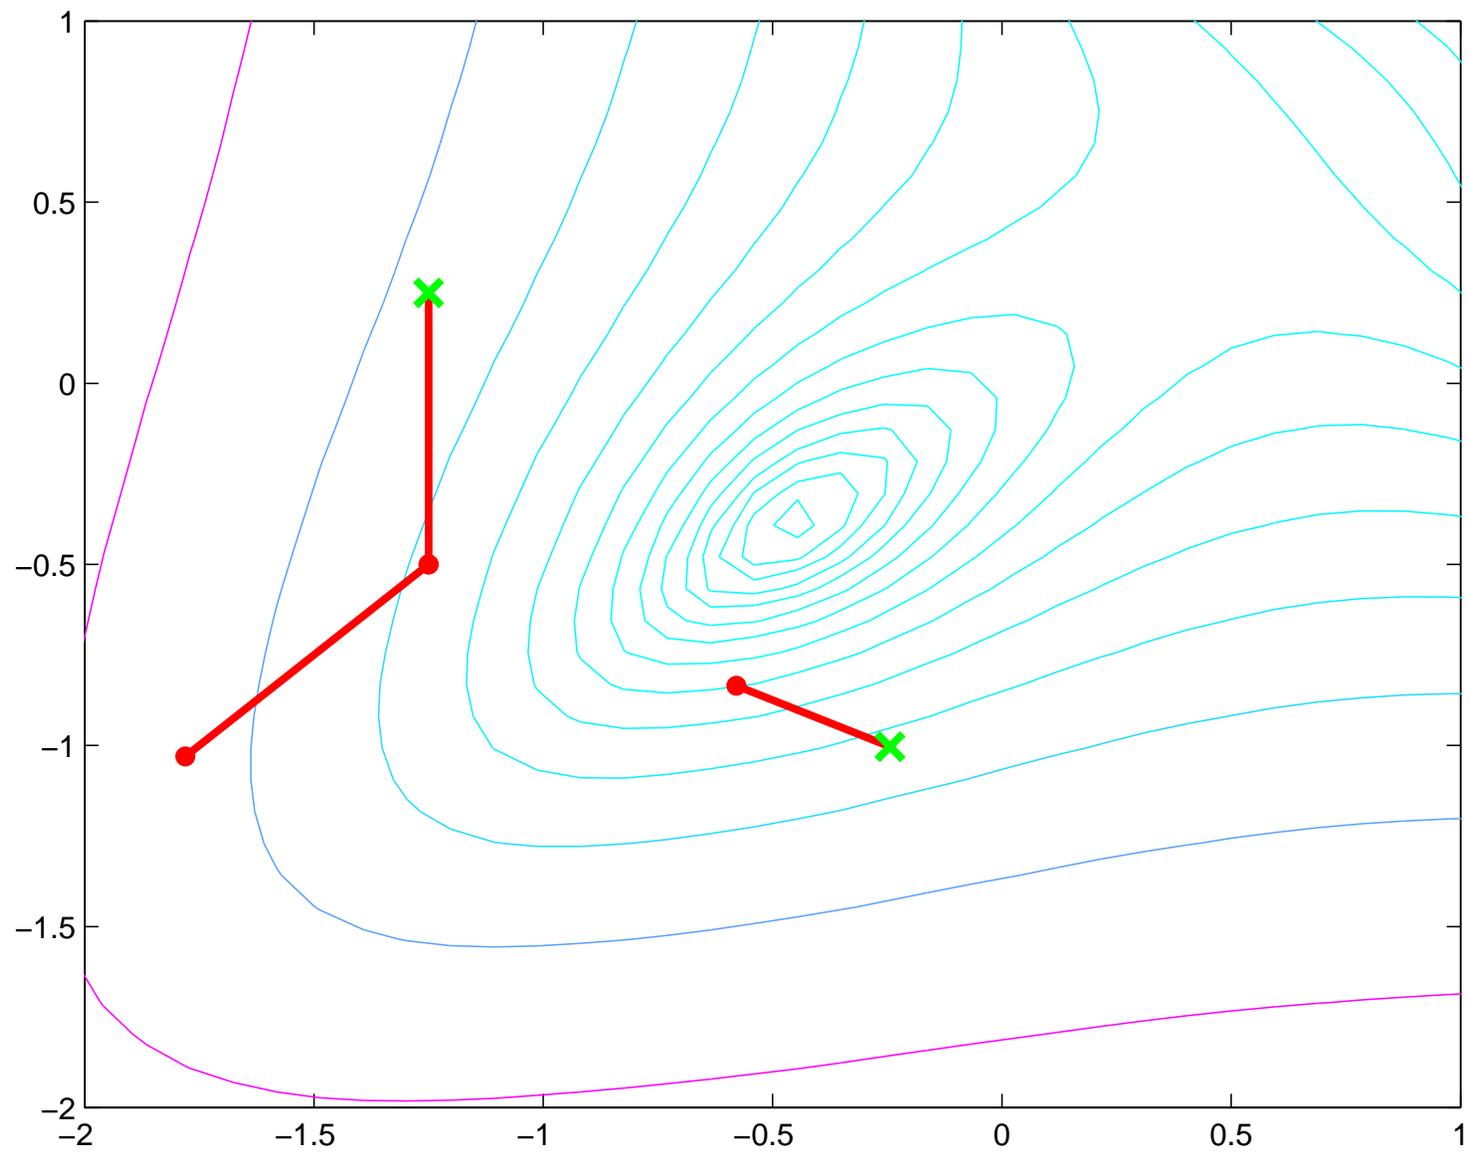
ASYNCHRONOUS PARALLEL PATTERN SEARCH



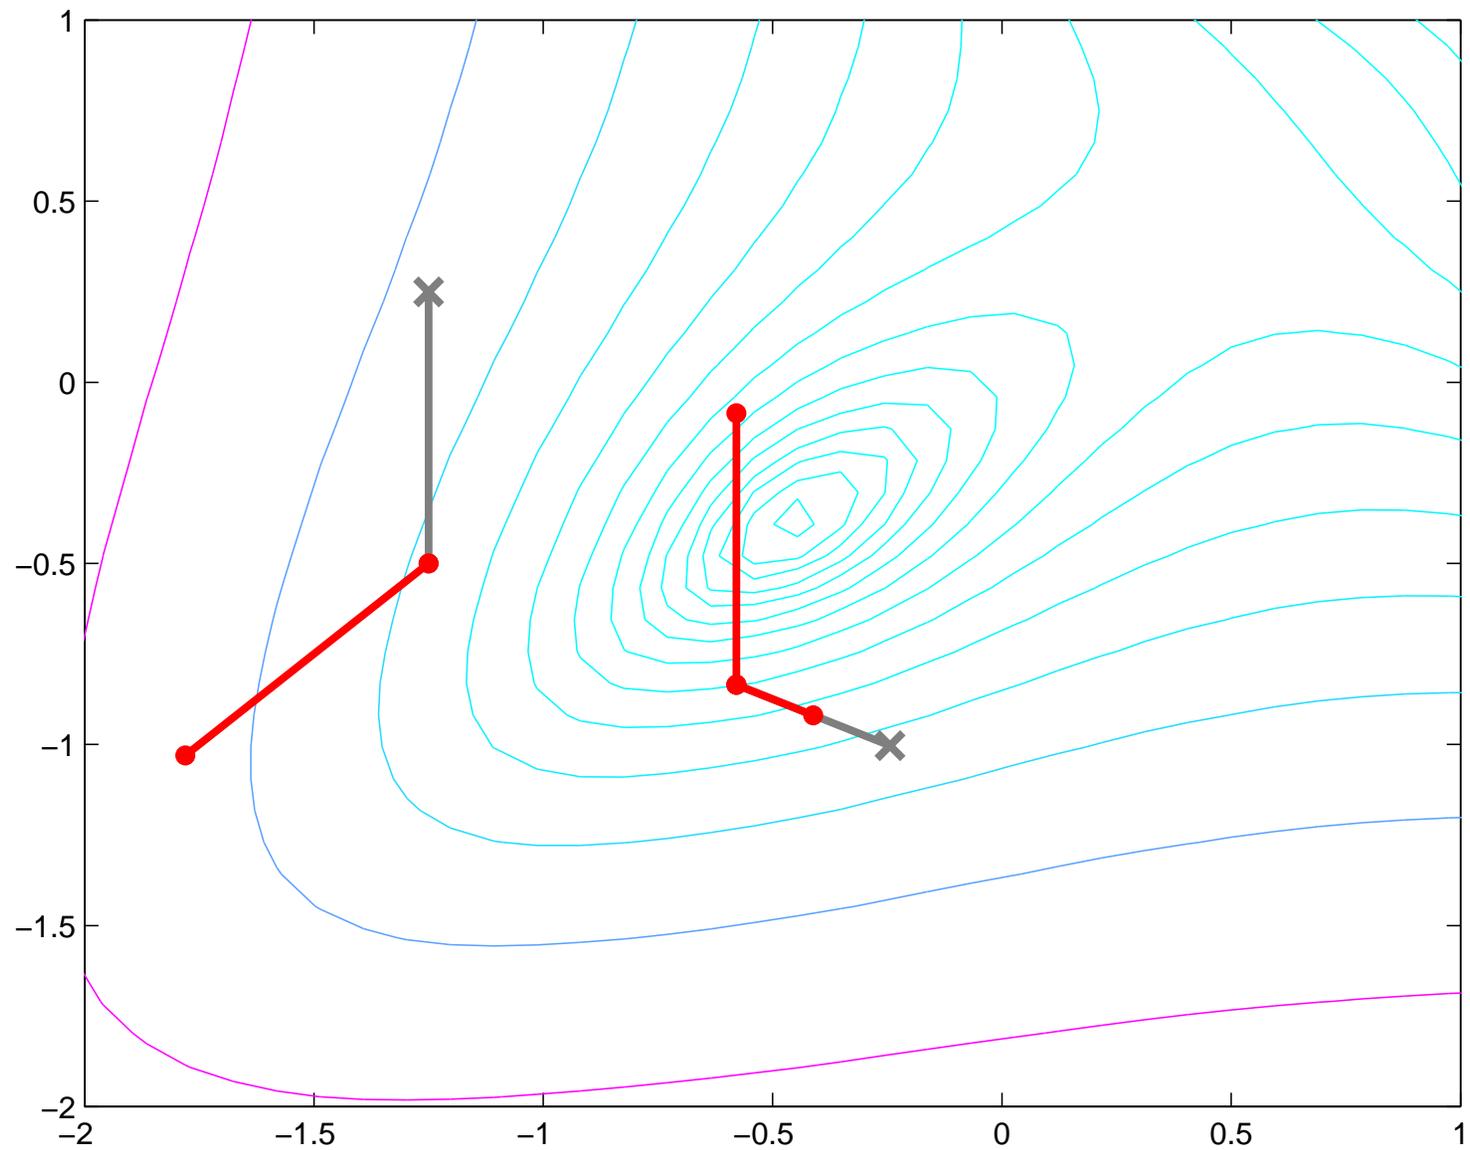
ASYNCHRONOUS PARALLEL PATTERN SEARCH



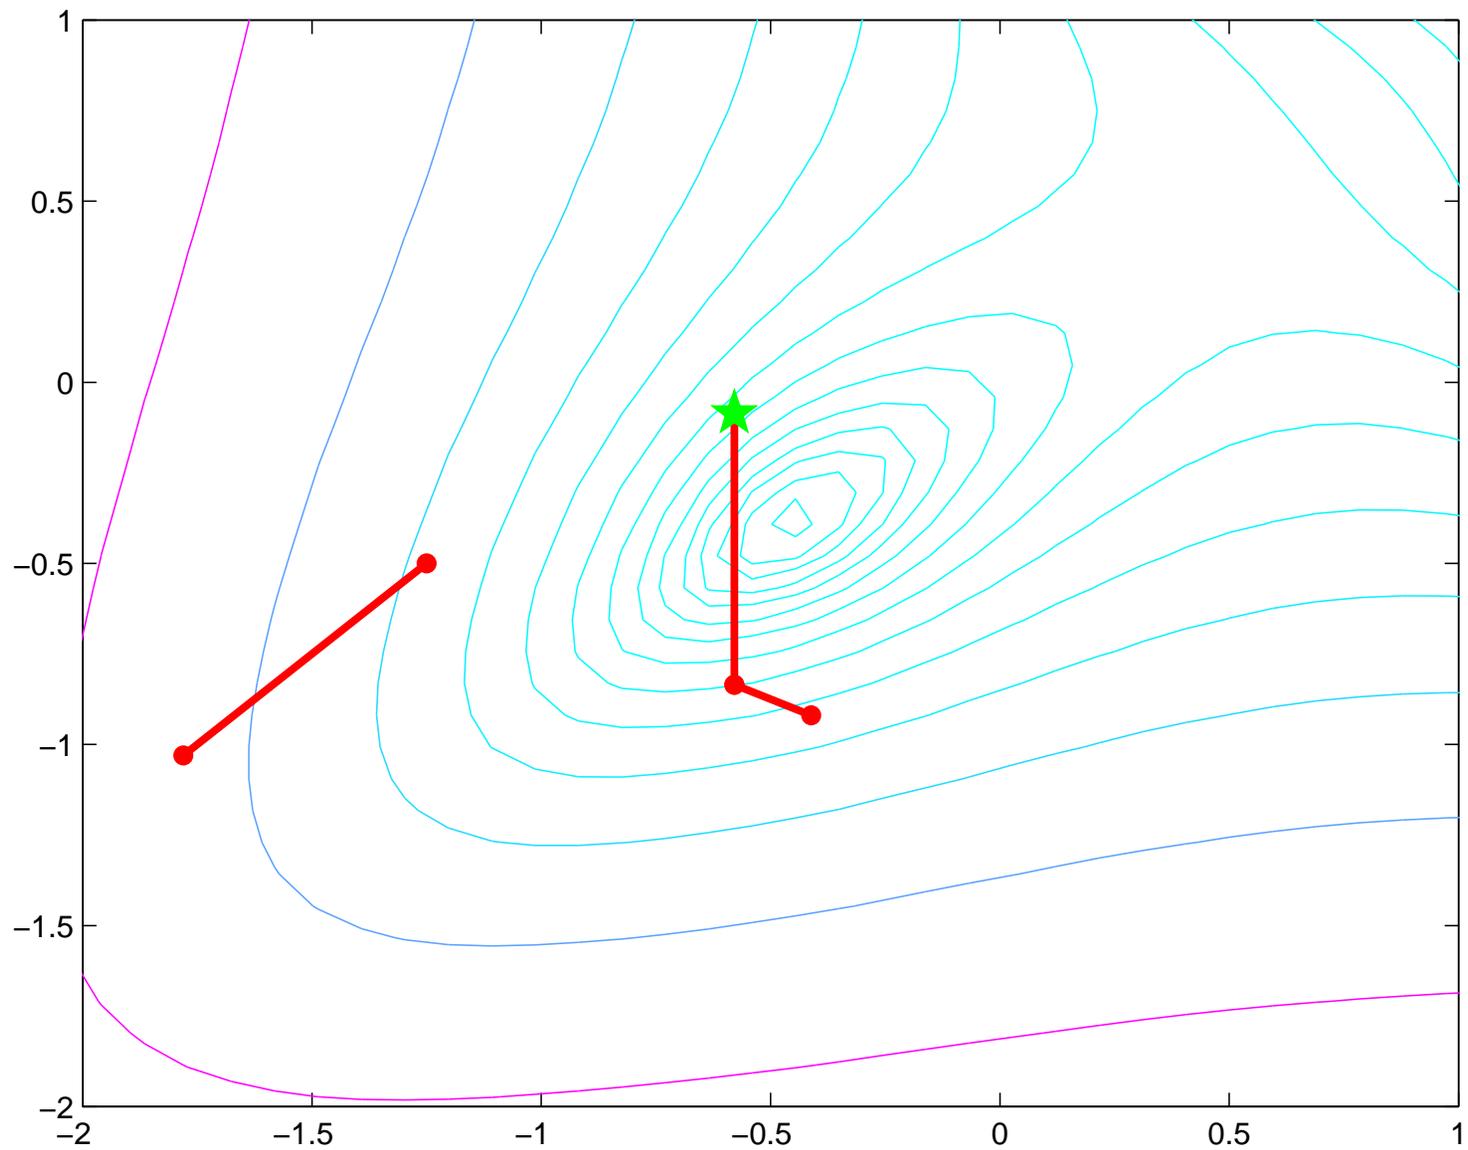
ASYNCHRONOUS PARALLEL PATTERN SEARCH



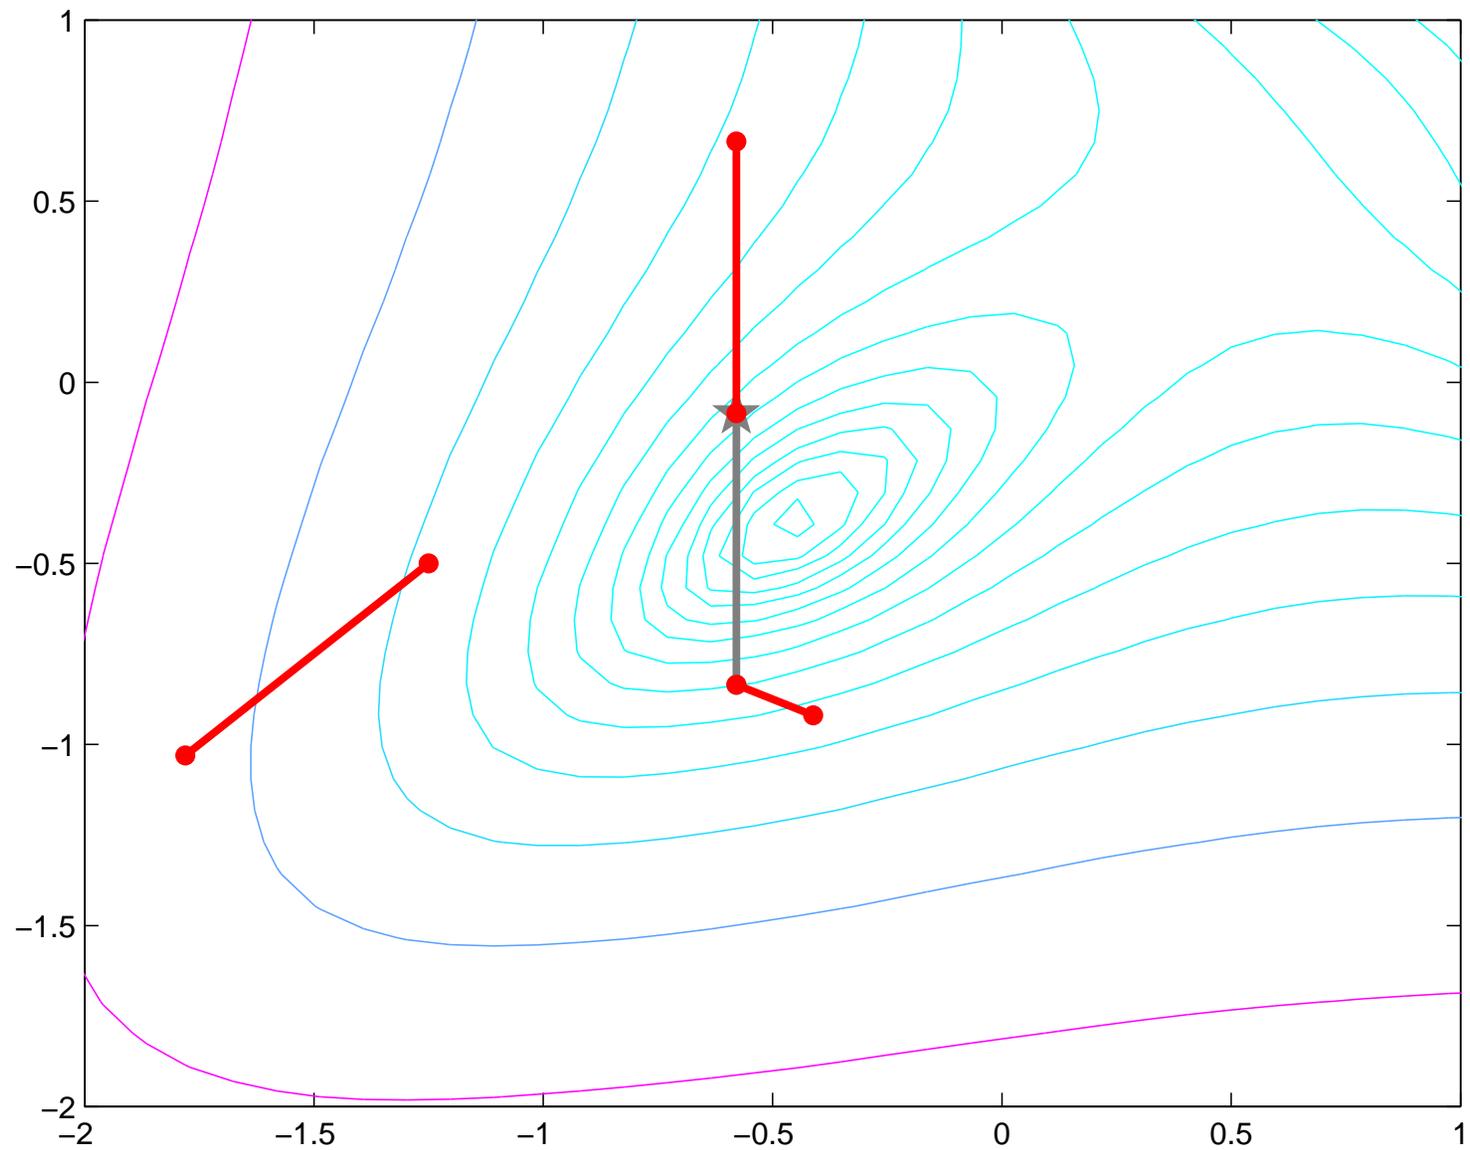
ASYNCHRONOUS PARALLEL PATTERN SEARCH



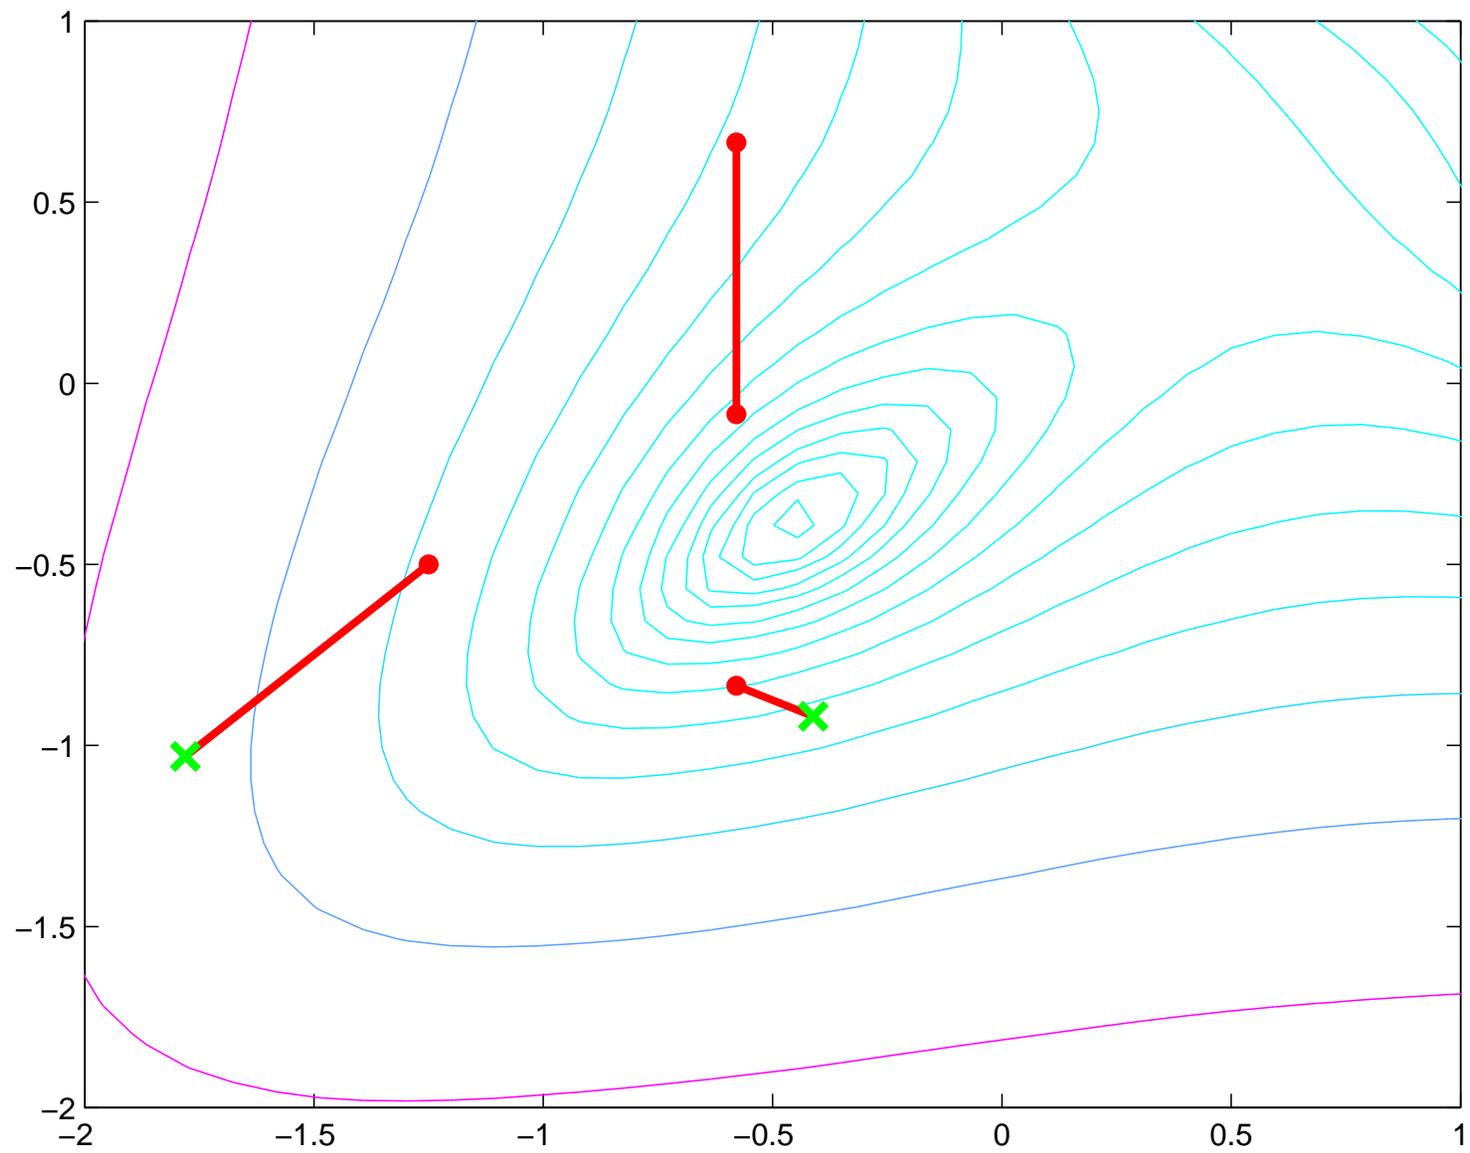
ASYNCHRONOUS PARALLEL PATTERN SEARCH



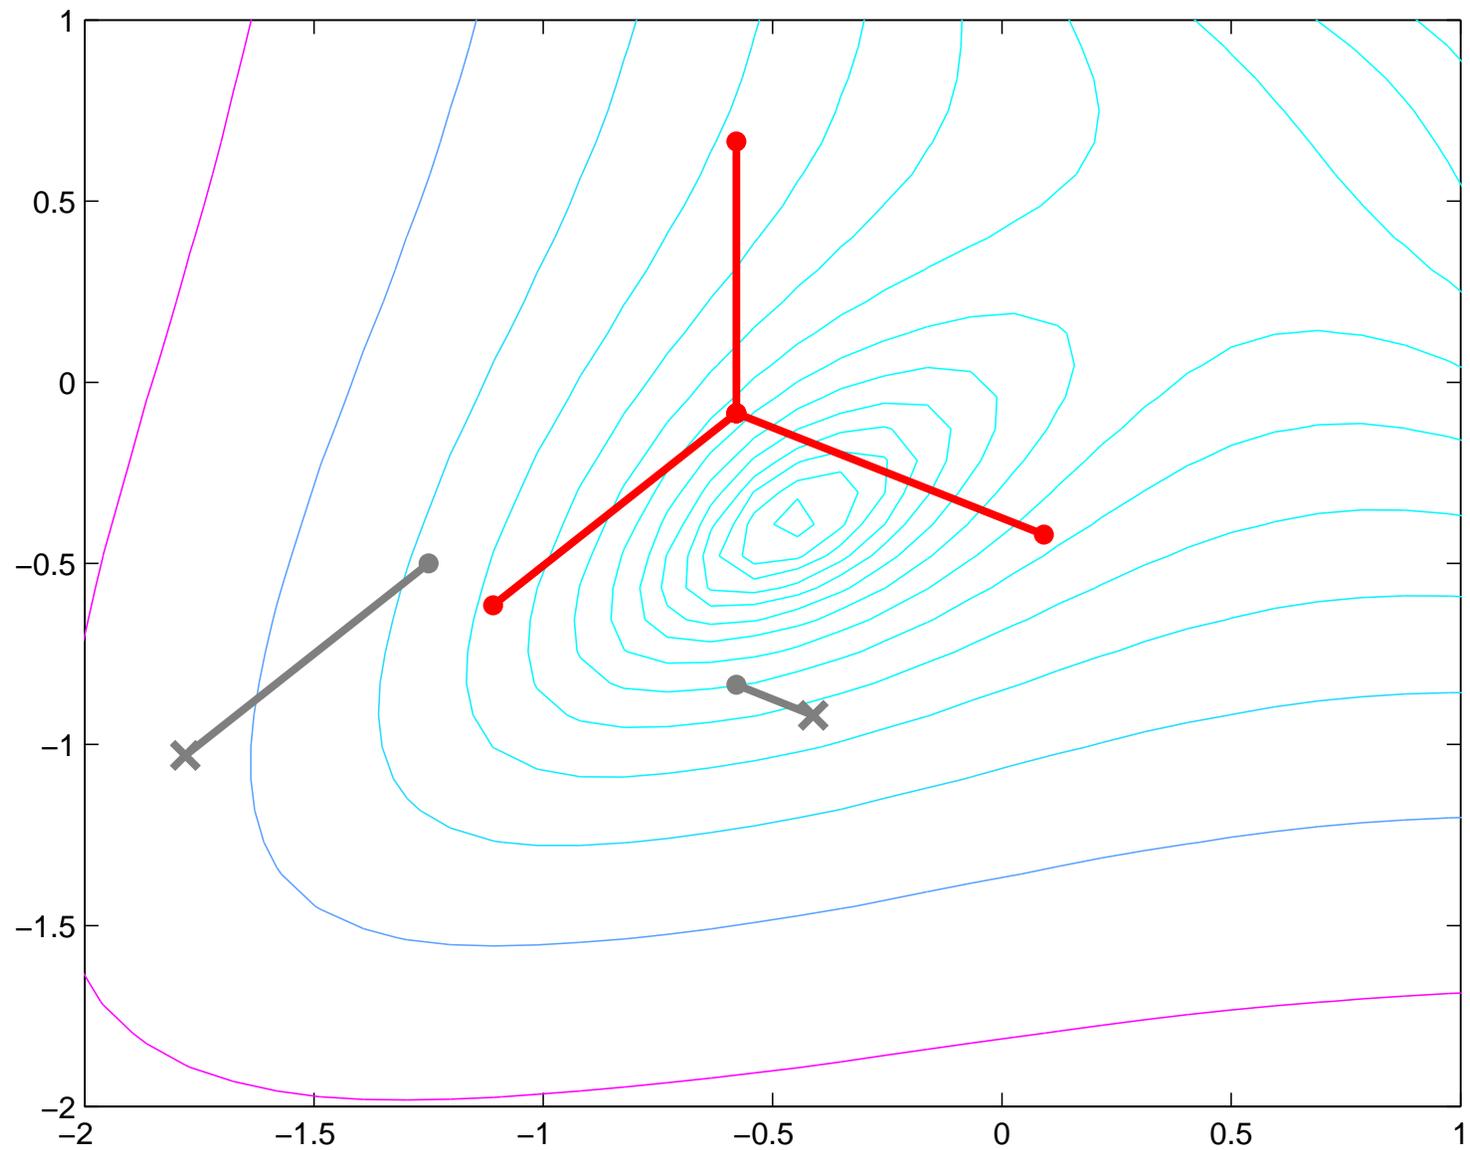
ASYNCHRONOUS PARALLEL PATTERN SEARCH



ASYNCHRONOUS PARALLEL PATTERN SEARCH



ASYNCHRONOUS PARALLEL PATTERN SEARCH



ENGINEERING EXAMPLE: DETERMINING THE CHARACTERISTICS OF A CIRCUIT

- **Collaborator:** Ken Marx (Sandia)
- **Variables:** inductances, capacitances, diode saturation currents, transistor gains, leakage inductances, and transformer core parameters
- **Simulation Code:** SPICE3

$$f(x) = \sum_{t=1}^N (V_t^{\text{SIM}}(x) - V_t^{\text{EXP}})^2,$$

x = 17 unknown characteristics

$V_t^{\text{SIM}}(x)$ = Simulation voltage at time t

V_t^{EXP} = Experimental voltage at time t

N = Number of time steps (2700)

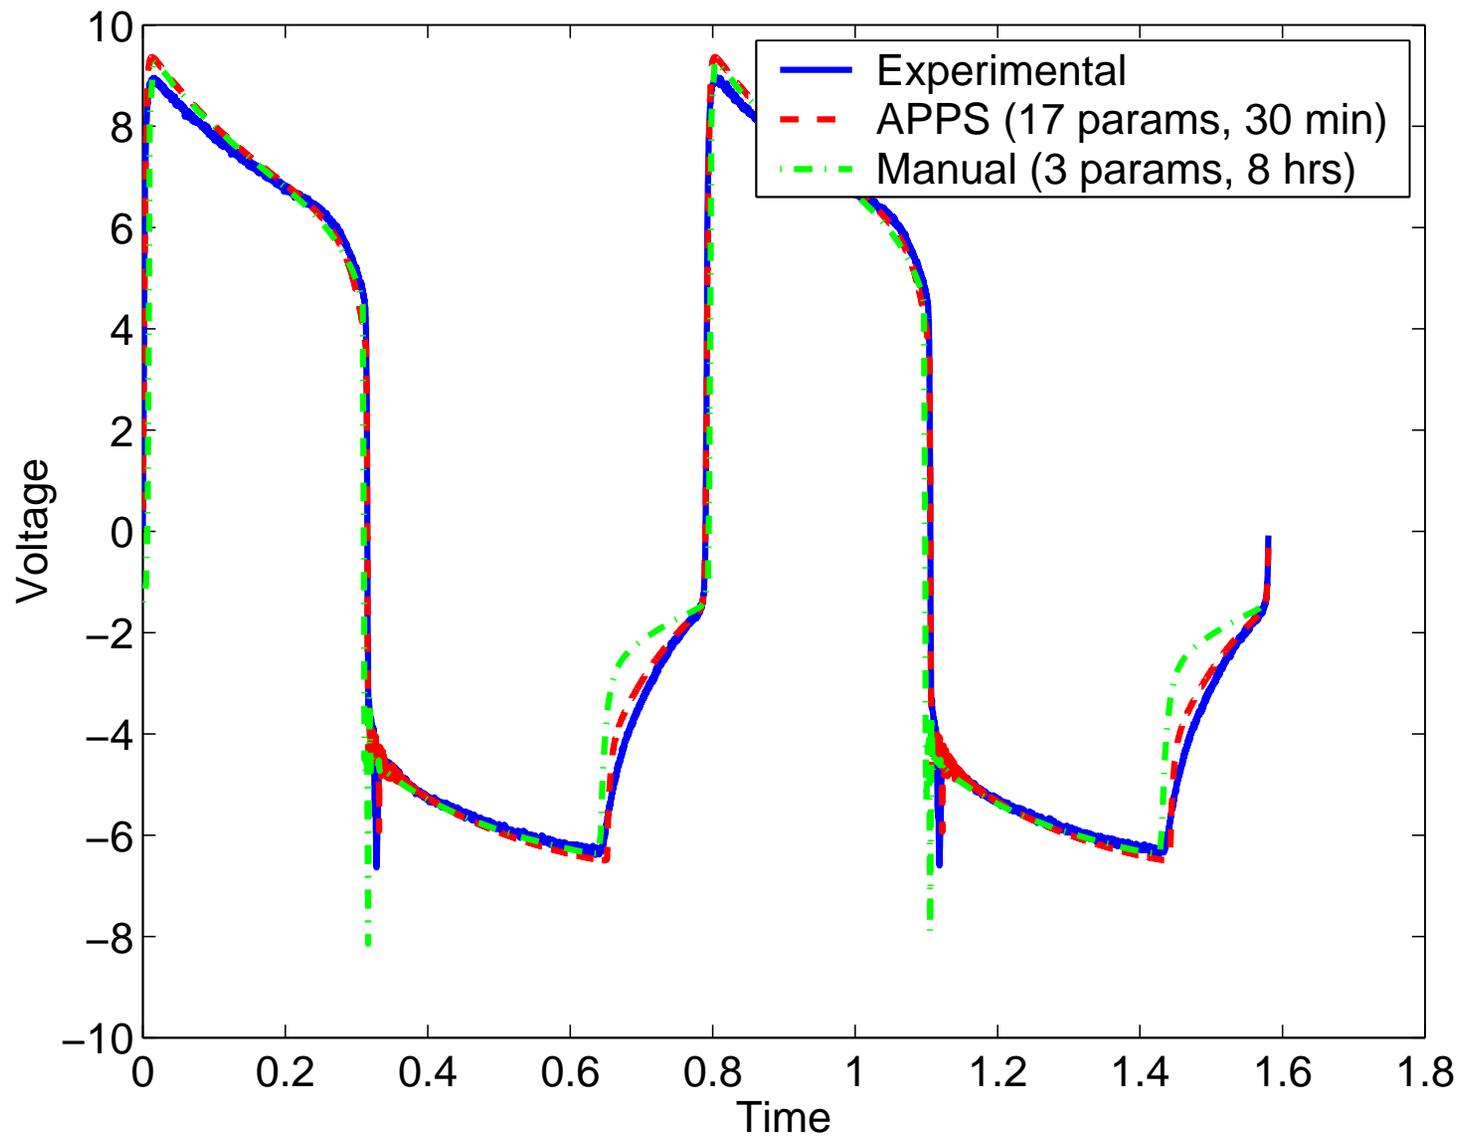
APPS DRAMATICALLY REDUCES IDLE TIME FOR THE CIRCUIT SIMULATION PROBLEM

$$f(x) = \sum_{t=1}^N (V_t^{\text{SIM}}(x) - V_t^{\text{EXP}})^2 \quad (17 \text{ variables})$$

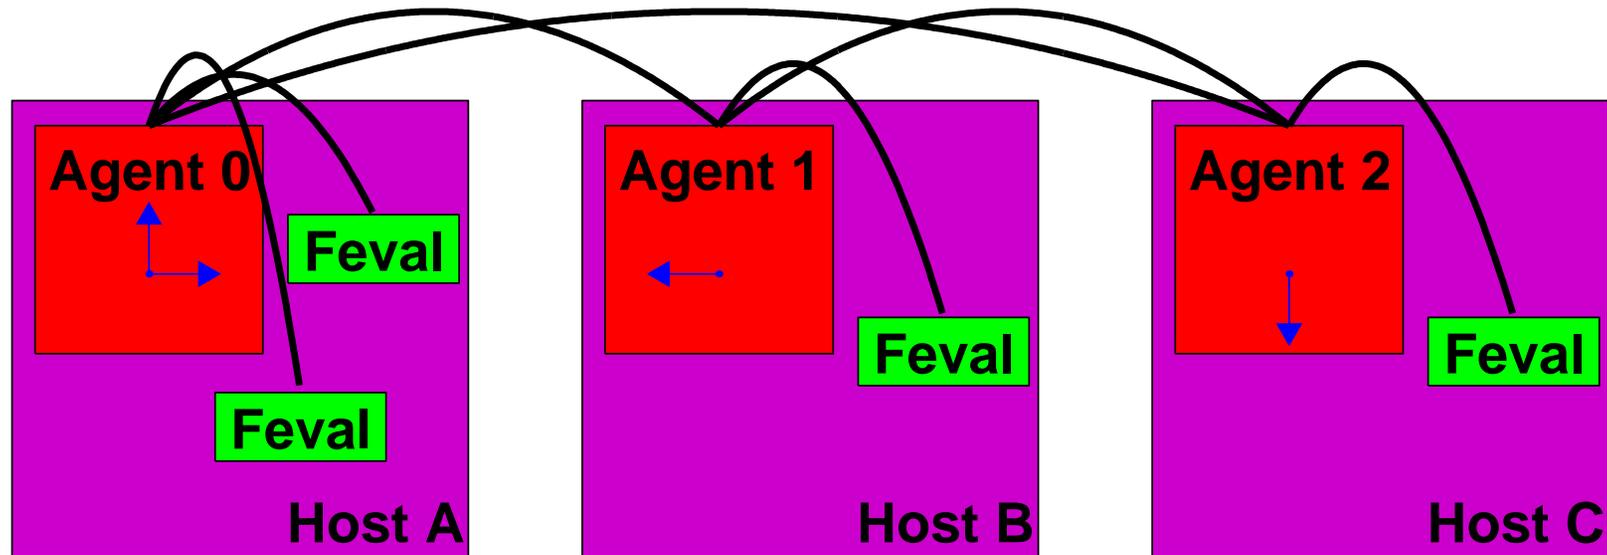
- Search directions are \pm Unit Vectors (34) plus Random
- The variables are bound constrained; i.e., $l_i \leq x_i \leq u_i$

Method	Procs	$f(x^*)$	Function Evals	Idle Time	Total Time
APPS	34	26.2	57.5	111.92	1330.55
APPS	50	26.9	50.6	63.22	807.29
PPS	34	28.8	53.0	521.48	1712.24
PPS	50	34.9	47.0	905.48	1646.53

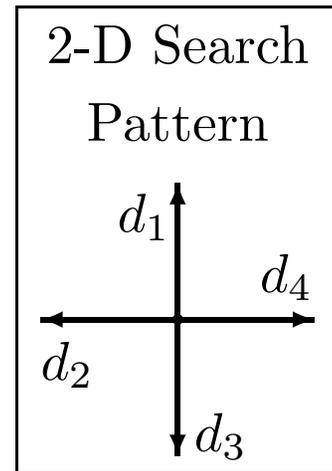
APPS SOLUTION IS FASTER AND MORE ACCURATE THAN HAND-TUNING



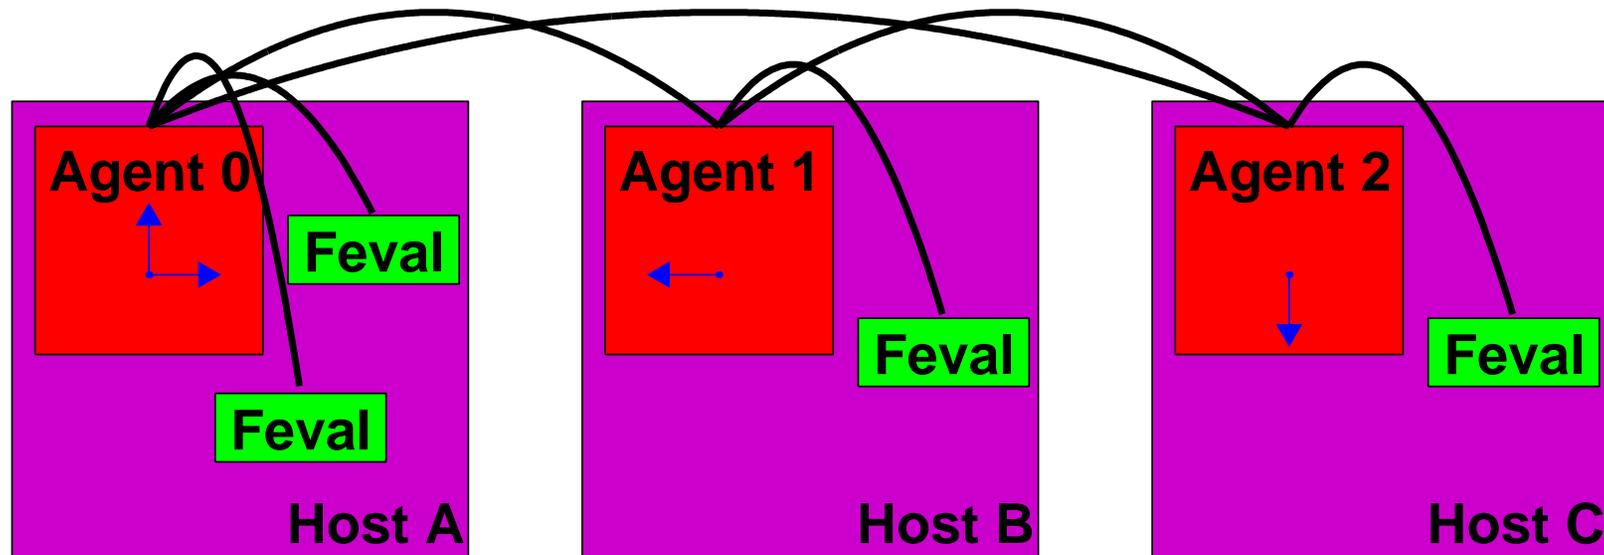
APPS HAS AGENT AND FEVAL PROCESSES



- Each host has an **Agent**, and each agent owns a subset of the **Search Pattern**.
- Each agent dynamically spawns **Function Evaluations**, and the maximum number of function evaluations per agent is equal to the number of search directions it owns.



APPS DISTRIBUTES CONTROL VIA AGENTS



- Each Agent has a **unique integer id**, starting with 0 and numbered consecutively.
- Agents communicate with each other and with any function evaluations that they own.
- The agents contain the logic of the search, but the bulk of the computation work is done by the function evaluations.

AGENT INFORMATION

Agent	Alive?	HostId	TaskId	Speed	Hostname
0	true	0x40000	0x4006f	1	su1cn1
1	true	0x80000	0x8006e	1	su1cn2
2	true	0xc0000	0xc0068	1	su1cn3

Job Assignment = [0 1 2 0]

- Agent numbers are never re-used.
- The HostId and TaskId are generated by PVM and are never reused.
- The speeds are used to balance the job assignments.
- The job assignments map search directions (4 in this example) to agents.

TWO KEYS TO FAULT TOLERANCE

Notify

```
GCI::notify(MsgTag::F_EXIT, ftid);
```

```
GCI::notify(MsgTag::AGENT_EXIT, agentid);
```

PVM automatically sends a message with the appropriate flag when the specified process exits, either successfully or unsuccessfully. The message body contains the process id.

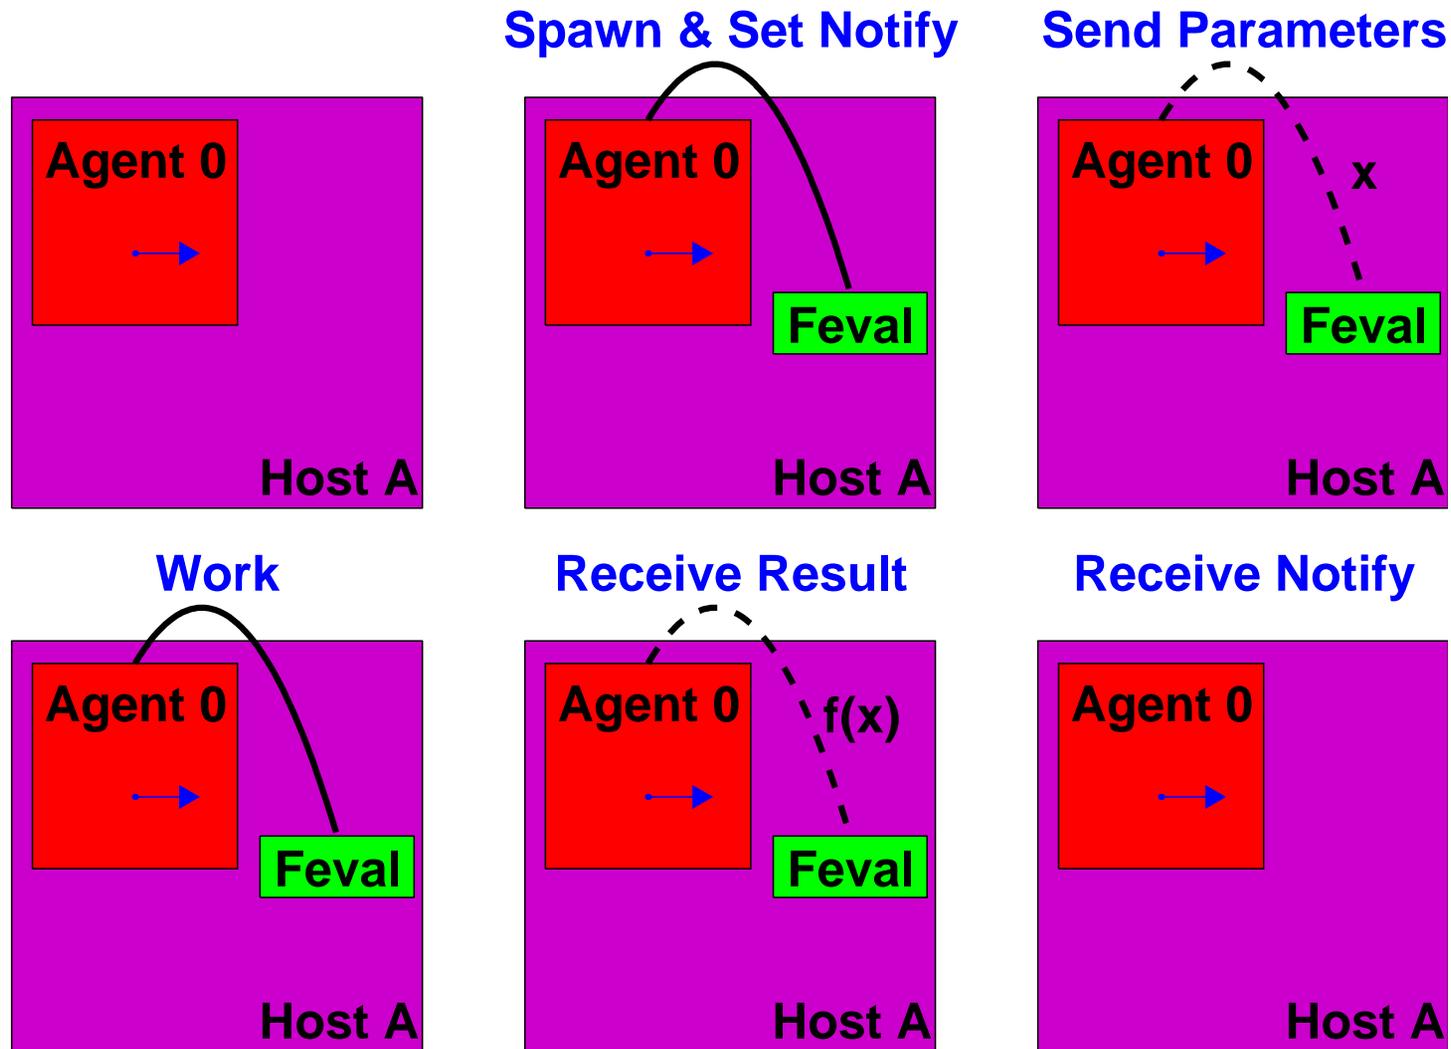
```
GCI::notify(MsgTag::AGENT_NEW_HOST);
```

PVM automatically sends a message with the appropriate flag if a new host is added.

GUID (Globally Unique Id)

A two-part identification that is used to differentiate and prioritize information. The first part of the GUID is the agent id, and the second part is produced by a counter which increments each time a new GUID is produced by that agent.

FEVAL FAULT TOLERANCE IS EASY



A fault can happen in any step, in which case we set $f(x) = +\infty$.

AGENT FAILURE

1. All the other agents get a notify message containing the TaskId of the failed agent, and they each update their “Agent Information” to show that the agent is dead.
2. Each agent determines whether or not it is the *temporary master*, which is defined to be the agent with the lowest number.

If Agent 0 fails, then Agent 1 is the temporary master.

Agent	Alive?	HostId	TaskId	Speed	Hostname
0	false	0x40000	0x4006f	1	su1cn1
1	true	0x80000	0x8006e	1	su1cn2
2	true	0xc0000	0xc0068	1	su1cn3

Job Assignment = [0 1 2 0]

We see that the job assignment must be updated...

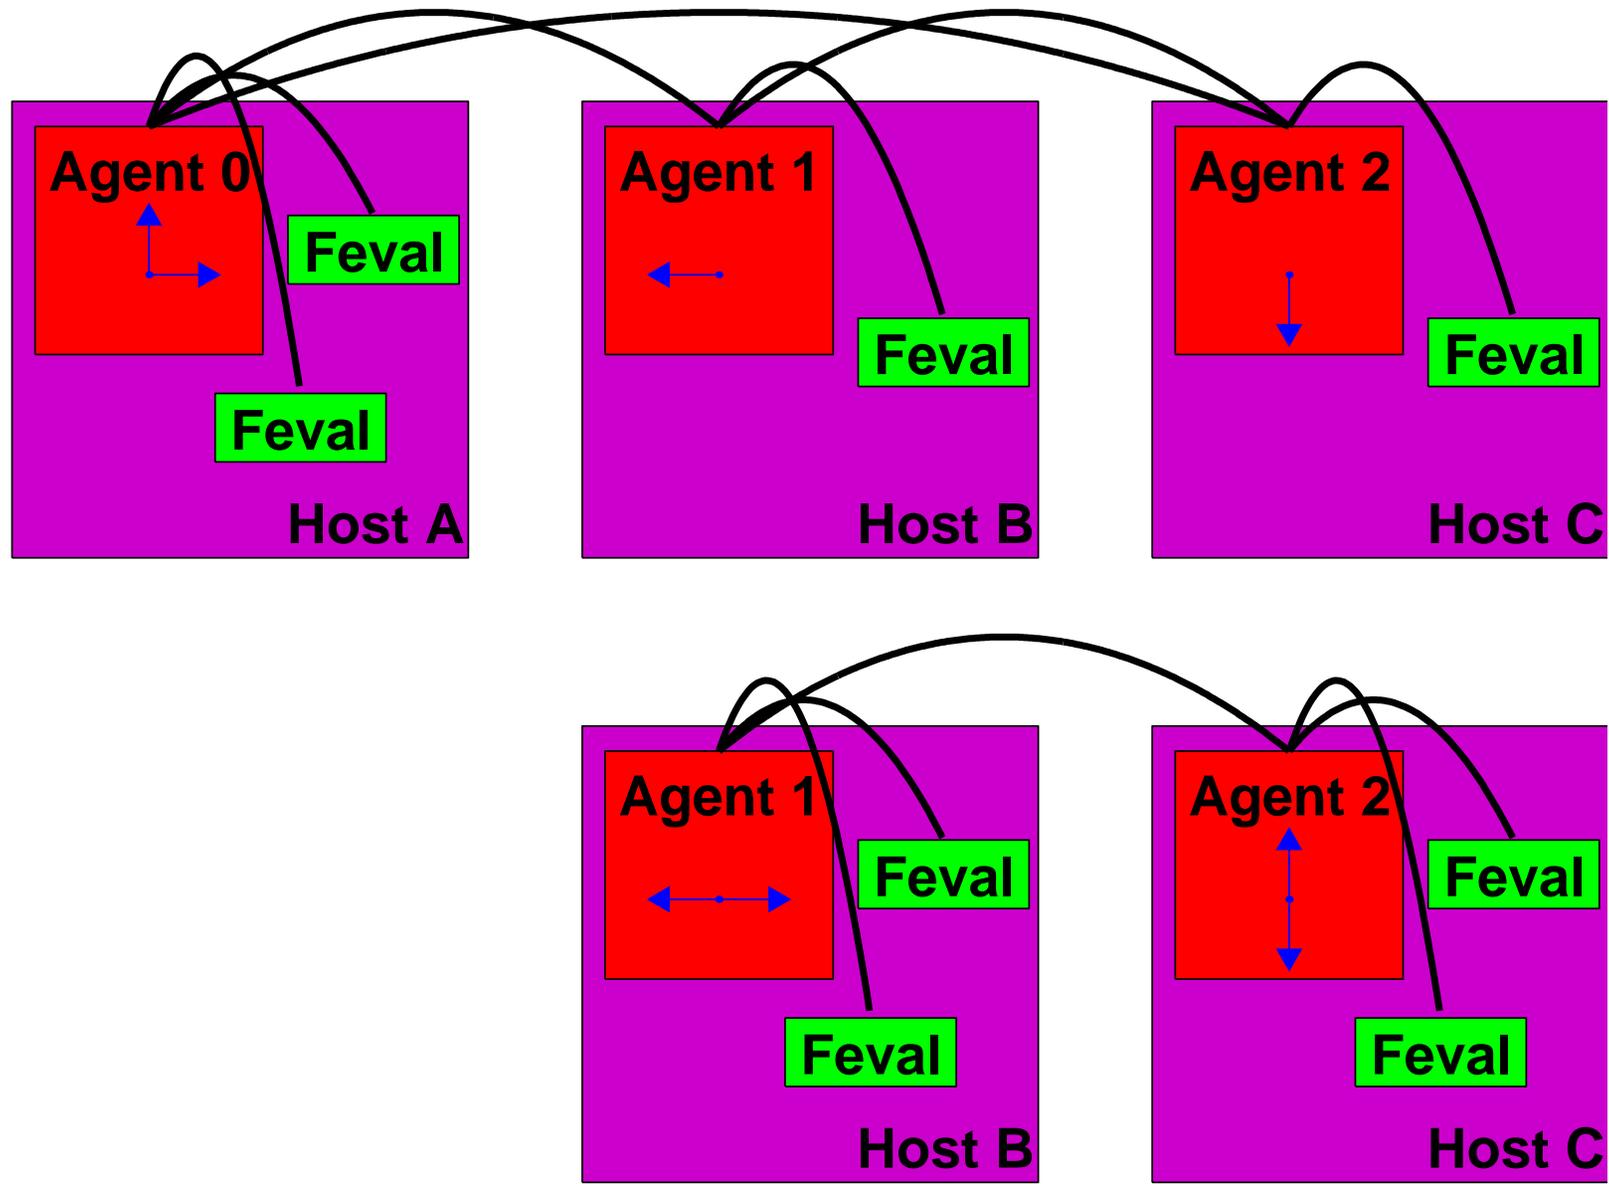
TEMPORARY AGENT RESPONSIBILITIES

1. Check for convergence.
2. Re-verify which of the agents/hosts are alive.
3. Reset search direction assignments, if necessary.
4. Generate a new GUID for the agent information.
5. Send an update message to all other agents containing the updated agent information.

Agent	Alive?	HostId	TaskId	Speed	Hostname
0	false	0x40000	0	1	su1cn1
1	true	0x80000	0x8006e	1	su1cn2
2	true	0xc0000	0xc0068	1	su1cn3

Job Assignment = [1 1 2 2]

RESETTING SEARCHES AFTER UPDATE



NEW HOST

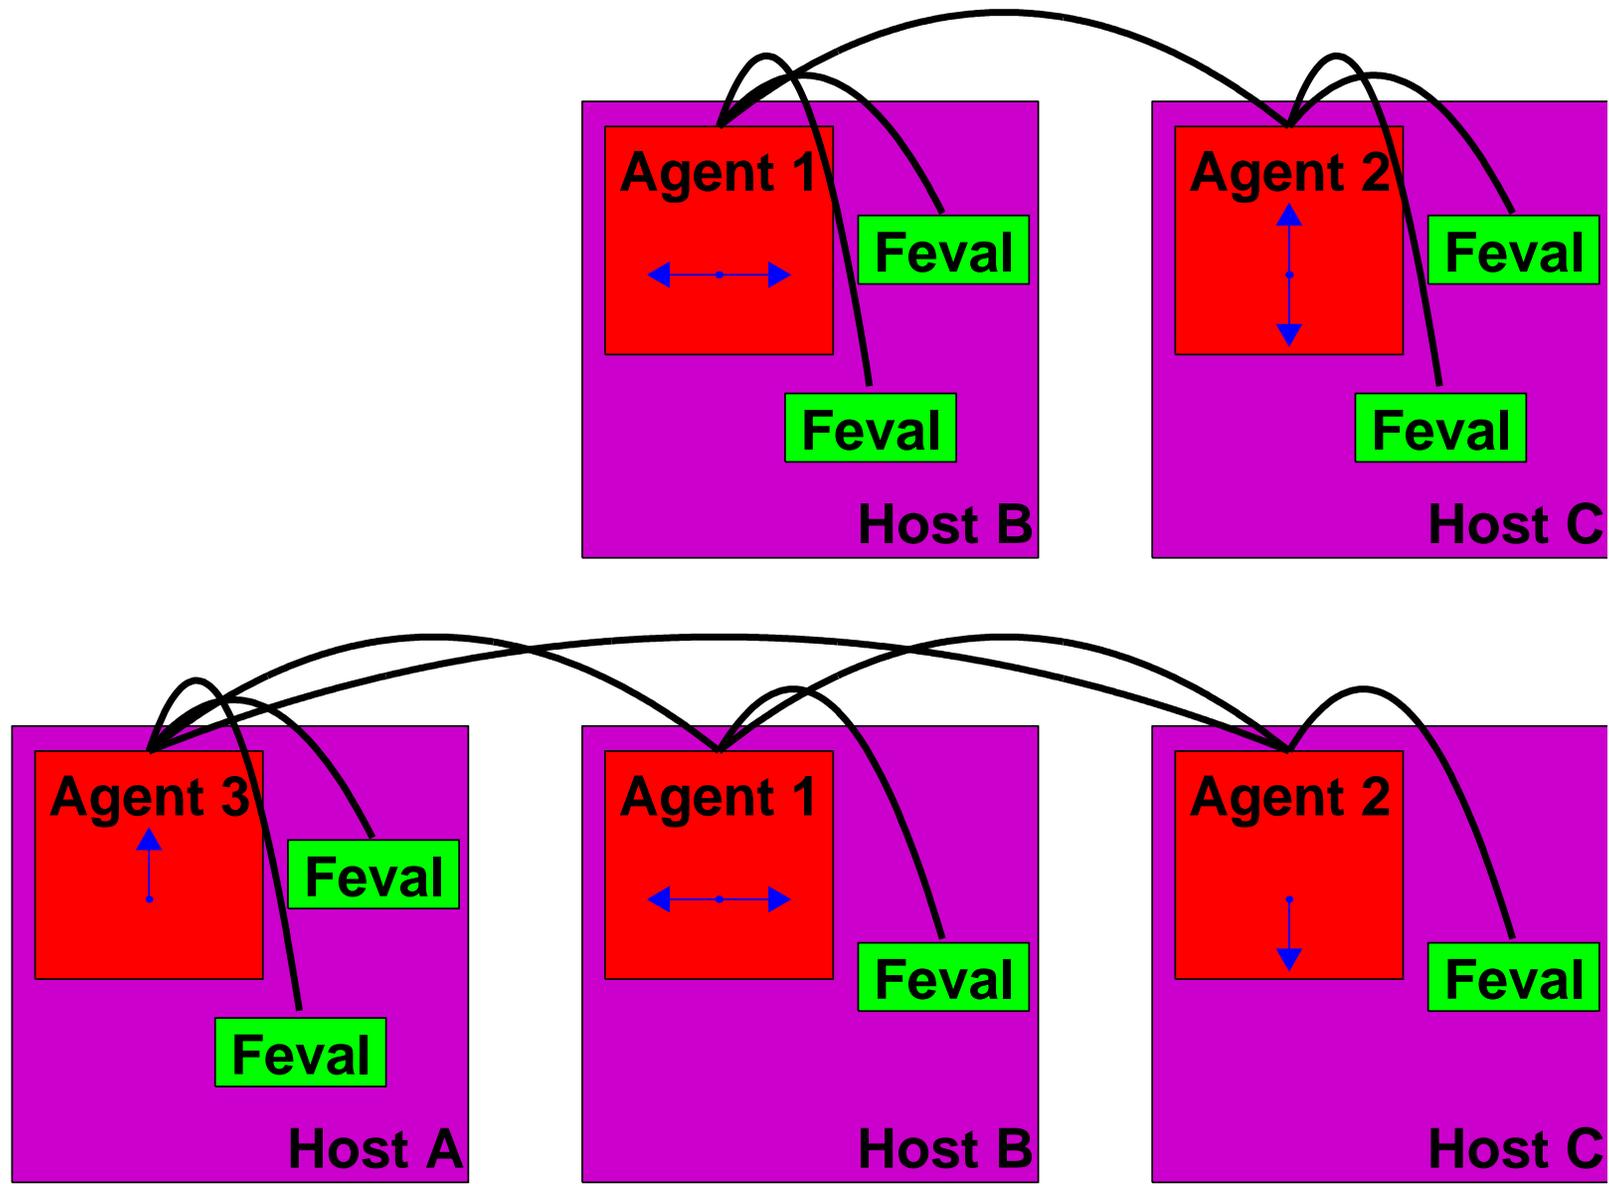
Every agent gets a notify message. The temporary master must...

1. Re-verify which of the old agents/hosts are alive.
2. Spawn agents on new hosts.
3. Reset search direction assignments, if necessary.
4. Generate a GUID for the agent information.
5. Send initialization message to new agents (including “best” point and search parameters)
6. Send an update message to old agents.

Agent	Alive?	HostId	TaskId	Speed	Hostname
0	false	0x40000	0	1	su1cn1
1	true	0x80000	0x8006e	1	su1cn2
2	true	0xc0000	0xc0068	1	su1cn3
3	true	0x100000	0x100004	1	su1cn1

Job Assignment = [1 1 2 3]

RESETTING SEARCHES AFTER UPDATE



SORTING AGENT INFORMATION

Concern: What if multiple agents fail in succession and we receive multiple agent information updates. How do we know which one to use?

Solution: The GUIDs and Temporary Master concepts allow us to sort out the information.

- The first part of the GUID is the agent number. Since the lowest numbered agent that is alive is the temporary master, we know that a higher first part of the GUID indicates a more recent message.
- The second part of the GUID is produced by a counter, so we know that the higher the counter, the more recent the information.

APPSPACK IS FAULT TOLERANT!

Circuit Simulation Example: The “fault” versions have a failure in the simulation or agent every 30 seconds.

Method	Initial Procs	$f(\mathbf{x}^*)$	Total Time
APPS	34	26.2	1330.55
APPS	34-faults	27.8	1618.46
PPS	34	28.8	1712.25
APPS	50	26.9	807.29
APPS	50-faults	54.2	1041.14
PPS	50	34.9	1646.53

There is no version of PPS that supports faults.

WHERE WE ARE

- Fault tolerance for function evaluations and agents.
- No dependence on a “master” agent.
- Automatic redistribution of jobs.

WHAT I DIDN'T MENTION

- Function Value Cache Process
- Surrogate Model Process

WHAT'S MISSING?

- Independence from PVM master daemon.
- Fault tolerance in MPI version.
- Disk check-pointing and restart.
- Method to detect that a process is “slow”.

TO OBTAIN APPSPACK...

<http://csmr.ca.sandia.gov/projects/apps.html>

FOR MORE INFORMATION...

Tammy Kolda

tgkolda@sandia.gov

<http://csmr.ca.sandia.gov/~tgkolda/>

925-294-4769

*APPSPACK is freely available
under the terms of the GNU L-GPL.*