

A parallel framework for reverse search using `mts`*

David Avis¹ and Charles Jordan²

¹School of Informatics, Kyoto University, Kyoto, Japan and School of Computer Science,
McGill University, Montréal, Québec, Canada, avis@cs.mcgill.ca

²Graduate School of Information Science and Technology, Hokkaido University, Japan,
skip@ist.hokudai.ac.jp

October 14, 2016

Abstract

We describe `mts`, which is a generic framework for parallelizing certain types of tree search programs, that (a) provides a single common wrapper containing all of the parallelization, and (b) minimizes the changes needed to the existing single processor legacy code. The `mts` code was derived from ideas used to develop `mplrs`, a parallelization of the reverse search vertex enumeration code `lrs`. The tree search properties required for the use of `mts` are satisfied by any reverse search algorithm as well as other tree search methods such as backtracking and branch and bound. `mts` is programmed in C, uses the MPI parallel environment, and can be run on a network of computers.

As examples we parallelize two simple existing reverse search codes: generating topological orderings and generating spanning trees of a graph. We give computational results comparing the parallel codes with state of the art sequential codes for the same problems.

Keywords: reverse search, parallel processing, topological orderings, spanning trees
Mathematics Subject Classification (2000) 90C05

1 Introduction

Although multicore hardware is now pervasive there are relatively few legacy codes that can profit from it. In some cases the underlying algorithm seems inherently sequential, rendering parallelization difficult if not impossible. In other cases the underlying algorithm may in principle be easy to parallelize. Nevertheless, existing widely-used and debugged sequential codes may be extremely complex and adding parallelization directly into such codes can be difficult and error prone. A simpler approach would be to design a wrapper that provides the parallelization and hopefully requires very little modification of the existing code. This approach was successfully applied by the authors to the vertex enumeration code `lrs`, resulting in the parallel `mplrs` code [4]. The ideas used in `mplrs` can be used for a wide variety of tree search algorithms including reverse search, backtracking, branch and bound, etc. This motivated us to develop a generic wrapper, `mts`, for parallelizing them. In this note we explain how to modify a reverse search code to enable parallelization by `mts`. We then describe in detail how to parallelize two simple existing reverse search codes. Finally, we give computational results showing nearly linear speedup on a cluster with 192 cores.

2 Background

Reverse search is a technique for generating large, relatively unstructured, sets of discrete objects [3]. Some simple C implementations were given in the tutorial [2] which we will extend in this note to allow

*This work was partially supported by JSPS Kakenhi Grants 16H02785, 23700019 and 15H00847, Grant-in-Aid for Scientific Research on Innovative Areas, ‘Exploring the Limits of Computation (ELC)’.

parallel processing via the `mts` package. We first give an outline of reverse search and the necessary modifications required for parallelization. This description is essentially the one given in [4].

In its most basic form, reverse search can be viewed as the traversal of a spanning tree, called the reverse search tree T , of a graph $G = (V, E)$ whose nodes are the objects to be generated. Edges in the graph are specified by an adjacency oracle, and the subset of edges of the reverse search tree are determined by an auxiliary function, which can be thought of as a local search function f for an optimization problem defined on the set of objects to be generated. One vertex, v^* , is designated as the *target* vertex. For every other vertex $v \in V$ repeated application of f must generate a path in G from v to v^* . The set of these paths defines the reverse search tree T , which has root v^* .

A reverse search is initiated at v^* , and only edges of the reverse search tree are traversed. When a node is visited, the corresponding object is output. Since there is no possibility of visiting a node by different paths, the visited nodes do not need to be stored. Backtracking can be performed in the standard way using a stack, but this is not required as the local search function can be used for this purpose. This implies two critical features that are useful for effective parallelization. Firstly, it is not necessary to store more than one node of the tree at any given time and no database is required for visited nodes. Secondly, it is possible to *restart* the enumeration process from any given node in the tree using only a description of this one node.

In the basic setting described here a few properties are required. Firstly, the underlying graph G must be connected and an upper bound on the maximum vertex degree, Δ , must be known. The performance of the method depends on G having Δ as low as possible. The adjacency oracle must be capable of generating the adjacent vertices of some given vertex v sequentially and without repetition. This is done by specifying a function $\text{Adj}(v, j)$, where v is a vertex of G and $j = 1, 2, \dots, \Delta$. Each value of $\text{Adj}(v, j)$ is either a vertex adjacent to v or null. Each vertex adjacent to v appears precisely once as j ranges over its possible values. For each vertex $v \neq v^*$ the local search function $f(v)$ returns the tuple (u, j) where $v = \text{Adj}(u, j)$ such that u is v 's parent in T . Pseudocode is given in Algorithm 1 and C implementations for several simple enumeration problems are given at [2]. For convenience later, we do not output the root vertex v^* in the pseudocode shown. Note that the vertices are output as a continuous stream.

Algorithm 1 Generic Reverse Search

```

1: procedure RS( $v^*$ ,  $\Delta$ , Adj,  $f$ )
2:    $v \leftarrow v^*$    $j \leftarrow 0$    $depth \leftarrow 0$ 
3:   repeat
4:     while  $j < \Delta$  do
5:        $j \leftarrow j + 1$ 
6:       if  $f(\text{Adj}(v, j)) = v$  then ▷ forward step
7:          $v \leftarrow \text{Adj}(v, j)$ 
8:          $j \leftarrow 0$ 
9:          $depth \leftarrow depth + 1$ 
10:        output  $v$ 
11:       end if
12:     end while
13:     if  $depth > 0$  then ▷ backtrack step
14:        $(v, j) \leftarrow f(v)$ 
15:        $depth \leftarrow depth - 1$ 
16:     end if
17:   until  $depth = 0$  and  $j = \Delta$ 
18: end procedure

```

Observe that Algorithm 1 does not require the parameter v^* to be the root of the search tree. If an arbitrary node in the tree is given, the algorithm reports the subtree rooted at this node and terminates. This is the key property that allows parallelization, as we discuss next.

3 Maximum depth and budgeting

In order to parallelize Algorithm 1 we need to break up the search of the entire tree T into searching a set of subtrees. We do this in two ways: by limiting the search depth and by limiting the number of nodes visited. In both cases, unexplored subtrees are generated and these must be flagged for later processing. To implement this we will supply three additional parameters:

- *start_vertex* is the vertex from which the reverse search should be initiated and replaces v^* .
- *max_depth* is the depth at which forward steps are terminated.
- *max_nodes* is the number of nodes to generate before terminating and reporting unexplored subtrees.

Both *max_depth* and *max_nodes* are assumed to be positive, for otherwise there is no work to do. The modified algorithm is shown in Algorithm 2.

Algorithm 2 Budgeted Reverse Search

```

1: procedure BTS(start_vertex,  $\Delta$ , Adj, f, max_depth, max_nodes)
2:    $j \leftarrow 0$    $v \leftarrow \text{start\_vertex}$    $\text{count} \leftarrow 0$    $\text{depth} \leftarrow 0$ 
3:   repeat
4:     unexplored  $\leftarrow$  false
5:     while  $j < \Delta$  and unexplored = false do
6:        $j \leftarrow j + 1$ 
7:       if  $f(\text{Adj}(v, j)) = v$  then ▷ forward step
8:          $v \leftarrow \text{Adj}(v, j)$ 
9:          $j \leftarrow 0$ 
10:         $\text{count} \leftarrow \text{count} + 1$ 
11:         $\text{depth} \leftarrow \text{depth} + 1$ 
12:        if  $\text{count} \geq \text{max\_nodes}$  or  $\text{depth} = \text{max\_depth}$  then ▷ budget is exhausted
13:          unexplored  $\leftarrow$  true
14:        end if
15:        put_output ( $v$ , unexplored)
16:      end if
17:    end while
18:    if  $\text{depth} > 0$  then ▷ backtrack step
19:       $(v, j) \leftarrow f(v)$ 
20:       $\text{depth} \leftarrow \text{depth} - 1$ 
21:    end if
22:  until  $\text{depth} = 0$  and  $j = \Delta$ 
23: end procedure

```

Comparing Algorithm 1 and Algorithm 2, we note several changes. Firstly an integer variable *count* is introduced to keep track of how many tree nodes have been visited, a process we call *budgeting*. Secondly, a flag *unexplored* is introduced to distinguish the tree nodes whose subtrees have not been explored. It is initialized as *false* on line 4. The flag is set to **true** in line 13 if either the budget of *max_nodes* has been exhausted or a depth of *max_depth* has been reached. Each node encountered on a forward step is output via the routine put_output on line 15. In single processor mode the output is simply sent to the output file with a flag added to unexplored nodes. In multi-processor mode, the output is synchronized and unexplored nodes are returned to the controlling master process.

Backtracking is as in Algorithm 1. After each backtrack step the *unexplored* flag is set to **false** in line 4. If the budget constraint has been exhausted then *unexplored* will again be set to **true** in line 13 after the first forward step. In this way all unexplored siblings of nodes on the backtrack path to the root are flagged in put_output. If the budget is not exhausted then forward steps continue until either it is, *max_depth* is reached, or we reach a leaf.

To output all nodes in the subtree of T rooted at v we set $start_vertex = v$, $max_nodes = +\infty$ and $max_depth = +\infty$. This reduces to Algorithm 1 if $v = v^*$. To break up T into subtrees we have two options that can be combined. Firstly we can set the max_depth parameter resulting in all nodes at that depth to be flagged as unexplored. Secondly we can set the budget parameter max_nodes . In this case, once this many nodes have been explored the current node and all unexplored siblings on the backtrack path to the root are output and flagged as unexplored.

Consider the tree in Figure 1 which has 25 nodes, $\Delta = 6$ and is rooted at vertex 0. For convenience the nodes are numbered 0,1,...,24 in reverse search order but this is in no way essential. If we set $max_depth = 1$ and $max_nodes = +\infty$ then only nodes 1,7,18,22 are visited and output with $unexplored = \mathbf{true}$. Now suppose we set the parameter $max_nodes = 13$ and $max_depth = +\infty$. Firstly nodes 1,...,12 are output with $unexplored = \mathbf{false}$. Then, nodes 13,15,16,18,22 are output with $unexplored = \mathbf{true}$. Alternatively, if we set $max_nodes = 8$ then nodes 1,2,...,7 are output with $unexplored = \mathbf{false}$ and nodes 8,9,10,11,15,16,18,22 are output with $unexplored = \mathbf{true}$.




Figure 1: Tree with 25 nodes and $\Delta = 6$

4 An example mts interface: topsorts

In the tutorial [2] a C implementation (*per.c*) is given for the reverse search algorithm for generating permutations. A small modification of this code generates all linear extensions of a partially ordered set that is given by a directed acyclic graph (DAG). Such linear extensions are also called topological sorts or topological orderings. The code modification is given as Exercise 5.1 and a solution to the exercise (*topsorts.c*) appears at the URL [2]. In this section we describe in detail how to modify this code to allow parallelization via the *mts* interface.

It is convenient to describe the procedure as two phases. Phase 1 implements max-depth and budgeting and organizes the internal data in a suitable way. This involves modifying an implementation of Algorithm 1 to an implementation of Algorithm 2 that can be independently tested. We also prepare a global data structure *bts.data* which contains problem data obtained from the input. In Phase 2 we build a node structure for use by the *mts* wrapper and add necessary routines to allow initialization and I/O in a parallel setting. In practice this involves sharing a common header file with *mts*. The resulting program can be compiled as a stand-alone code or as a parallel code with no change in the source files. The code shown in the two subsections below is essentially that as given at [2] with a few nonessential parts deleted.

4.1 Phase 1 : atopsorts.c

- Internal data grouped in `bts_data` structure

```
struct bts_data {
    int n;           /* number of nodes and edges in graph */
    int m;
    int A[100][100]; /* the graph */
    int countonly;  /* TRUE means nodes are counted but not output */
};
```

- Initialization of `bts_data` from the input file via `bts_init`:

```
bts_data *bts_init(int argc, char **argv)
{
    int i, j, k, m, n;
    bts_data *b = malloc(sizeof(bts_data));
    b->countonly=FALSE;

    /* process command line arguments */
    for (i=1; i<argc; i++)
        if (strcmp(argv[i], "-countonly") == 0)
            b->countonly=TRUE;

    /* read input file and build bts_data */
    scanf("%d %d", &b->n, &b->m);
    n=b->n;
    m=b->m;
    for (i=1; i<=n; i++)
        for (j=1; j<=n; j++)
            b->A[i][j]=0;

    for (k=1; k<=m; k++)
        {
            scanf("%d %d", &i, &j);
            b->A[i][j]=1;
        }
    return b;
}
```

- `max_nodes` and `max_depth` parameters are introduced to the reverse search code as described in Algorithm 2:

```
long bts(bts_data *b, perm v, long max_depth, long max_nodes)
{
    int j=0, depth=0, unexplored;
    long count=0;
    int n = b->n;
    int maxdeg=n-1;
    do
    {
        unexplored = FALSE;
        while (j < maxdeg && !unexplored)
        {
            j++;
            if (reverse(b, v, j))
                /* forward step */
                {
                    Adj(v, j);
                    depth++;
                    count++;
                    if (count >= max_nodes || depth == max_depth)
                        unexplored=TRUE;
                    put_output(b, v, depth, unexplored);
                    j = 0;
                }
        }
    } /* end while */
}
```

```

        if(depth > 0)
            {
                j = backtrack(b, v);
                depth--;
            }
    } while(depth > 0 || j < maxdeg);
    return count;
}

```

- *main()* is modified to call *bts_init*, process command line arguments to test budgeting and call *bts()* to perform the tree search:

```

int main(int argc, char **argv)
/* generate topsorts with budgeting and bts_data structure */
{
    bts_data *b;
    perm v;
    long count;
    int i, n;
    long max_nodes=LONG_MAX, max_depth=LONG_MAX;

    b=bts_init(argc, argv);
    n=b->n;

    for (i=1; i<=n; i++)
        v[i] = i;

    for (i=1; i<argc; i++)
    {
        if (strcmp(argv[i], "-maxd") == 0)
            max_depth = atoi(argv[i+1]);
        if (strcmp(argv[i], "-maxnodes") == 0)
            max_nodes = atoi(argv[i+1]);
    }

    /* the tree search is now initiated from root */
    put_output(b, v, 0, 0); /* first output */
    count = bts(b, v, max_depth, max_nodes);

    /* finish up */
    printf("\nnumber of permutations=%ld\n", count+1); /* count the root */

    free(b);
    return 0;
}

```

4.2 Phase 2 : btopsorts.c

In the second phase we add the ‘hooks’ that allow communication with *mts*. This involves defining a Node structure which holds all necessary information about a node in the search tree. The roots of unexplored subtrees are maintained by *mts* for parallel processing. Therefore whenever a search terminates due to the *max_nodes* or *max_depth* restrictions, the Node structure of each unexplored tree node is returned to *mts*. As we do not wish to customize *mts* for each application, we use a very generic node structure. The user should pack and unpack the necessary data into this structure as required. The Node structure is defined in the *mts.h* header. We now give some more specific details. Most of the code is common to single and multithread mode. The few exceptional cases are handled by the compile switch *#ifdef MTS*.

- The generic Node data structure in *mts.h* is:

```

typedef struct node_v {
    long *vlong; /* tree node longs (for user) */
    unsigned int size_vlong; /* number of longs in vlong */
    int *vint; /* tree node ints (for user) */
}

```

```

    unsigned int size_vint; /* number of ints in vint */
    char *vchar;           /* tree node chars (for user) */
    unsigned int size_vchar; /* number of chars in vchar */
    float *vfloat;        /* tree node floats (for user) */
    unsigned int size_vfloat; /* number of floats in vfloat */

    int unexplored; /* this subtree is unexplored */
    long depth;     /* depth of v in tree */
    struct node_v *next; /* for list of unexplored nodes */
} Node;

```

This is not user modifiable without changing `mts.c` itself.

- The tree search begins from a single root node, which is allocated and initialized in:

```

Node *get_root(bts_data *b)
/* return a node which is the root of the search tree */
{
    int i, n;
    Node *root = malloc(sizeof(Node));

    n=b->n;
    root->vlong = malloc(sizeof(long)*(n+1));
    root->size_vlong = n+1;
    root->vint = NULL; root->size_vint = 0;
    root->vchar = NULL; root->size_vchar = 0;
    root->vfloat = NULL; root->size_vfloat = 0;
    root->depth = 0;
    root->unexplored = FALSE;

    for (i=1; i<=n; i++) /* root permutation for topsort is 1,2,...,n */
        root->vlong[i] = i;

    if (b->countonly == FALSE)
        put_output(b, root); /* first output ! */
    return root;
}

```

In `btopsorts.c` a call to `get_root()` is made immediately after the call to `bts_init()`. When using `mts`, a call is made to `get_root()` as part of the initialization process. This is called only once, `mts` deals with transferring nodes between processors.

- Options for `btopsorts.c` are collected in the structure:

```

const mtsoption bts_options [] = {
    {"-countonly", 0}, /* -countonly option has no parameters */
    {"-prune", 1},    /* -prune option has one parameter */
};
const int bts_options_size = 2;

```

which has currently two user defined options. The first option, `-countonly`, takes no parameters and the second, `-prune` (described in Section 4.3), takes one. Additional options can be added here as long as `bts_options_size` is updated and the options do not conflict with the `mts` option list (described in Section 5).

- The input file is handled by `mts` which converts it to a string pointed to by file pointer `f`. The user simply needs to convert all `scanf(...)` statements to `fscanf(f, ...)`. `mtslib.c` contains `get_input()` to convert input to Data for internal use. No user modification is required or recommended.

Output is handled by `mts` in multiprocessor mode. Output to `stdout` and `stderr` is controlled by the values set for `b->output` and `b->err` in `bts_init()`. This is transparent to the user, who needs only replace all instances of `printf(...)` by `tsprintf(b->output,...)` and also instances of `fprintf(stderr,...)` by `tsprintf(b->err,...)`.

These I/O handling considerations require a slight modification of *bts_data* and associated declarations: Note in particular the extra two lines required in the definition of *bts_data* and the compile switch.

```
#ifndef MIS
#define tsprintf stream_printf
#define tstream mts_stream
#else
#define tsprintf fprintf
#define tstream FILE
#endif

typedef long Perm[100];

/* bts_data is built from the input and can be user modified */
/* except where mentioned */

struct bts_data {
    int n; /* number of nodes and edges in graph */
    int m;
    int A[100][100]; /* the graph */
    int countonly; /* TRUE means nodes are counted but not output */
    int prune; /* mark unexplored only if nchild > prune */
    tstream *output; /* output stream for stdout */
    tstream *err; /* output stream for stderr */
};
```

- A few small changes to *put_output()*, making use of the compile switch, are required to be sure the output goes in the right place:

```
void put_output(bts_data *b, Node *treenode)
{
    int i;
    if (!b->countonly)
    {
        for (i=1; i<=b->n; i++)
            tsprintf(b->output, " %ld", treenode->vlong[i]);
            tsprintf(b->output, " d=%ld", treenode->depth);
#ifdef MIS
        if (treenode->unexplored)
            tsprintf(b->output, " *unexplored");
#endif
        tsprintf(b->output, "\n");
    }
#ifdef MIS
    if (treenode->unexplored)
        return_unexp(treenode);
#endif
}
```

The *return_unexp()* routine is supplied in *mts.c*.

- A few small changes are required to *bts_init()*. Omitting code already given above, this becomes:

```
bts_data *bts_init(int argc, char **argv, Data *in, int proc_no)
{
    int i, j, k, n,m;
    FILE *f;

    bts_data *b = malloc(sizeof(bts_data));
    b->countonly=FALSE;

    /* process command line arguments to btopsorts - as above, omitted */

    f = open_string(in->input);
#ifdef MIS
    b->output = open_stream(MISOUT);
```



```

        b->err = open_stream(MTSERR);
#else
        b->output = stdout;
        b->err = stderr;
#endif
        i = fscanf(f, "%d %d\n", &b->n, &b->m);
        if (i == EOF)
        {
            tsprintf(b->err, "no input found\n");
            fclose(f);
#ifdef MTS
            close_string();
#endif
            return NULL;
        }

        n=b->n;
        m=b->m;

        /* read input data containing the input graph - as above, omitted */

        return b;
}

```

- Finally there is *bts_finish()* which is called by the master program once after all worker processors have terminated. This routine is not required for reverse search applications and is included only for future compatibility. It is intended to be used by applications that use *shared_data*, such as branch and bound, to provide a final solution to a problem.

```

void bts_finish(bts_data *bdata, shared_data *sdata, int size, int checkpointing)
{
    return;
}

```

By itself *btopsorts.c* compiles to the standalone program *btopsorts*, but it becomes the parallel code *mtopsorts* if compiled with the *mts.c* wrapper using the MTS flag.

4.3 Extras

There are some additional features implemented in *mts* that are not essential but may prove useful. They are documented in the code available at [2] and include:

- *cleanstop()* and *emergencystop()*. These allow a user process to cleanly close down the entire parallel execution.
- Shared data. This allows user processes to pass data back to *mts* which may be useful for other processes. This is not required for reverse search applications but will be useful for applications such as branch and bound, satisfiability, game tree search, etc.
- Output blocks. Although *mts* guarantees that output from each call to *stream_printf* is synchronized with other processes, it may be desirable to synchronize a block of *stream_printf* output. If the application opens an *output block*, all output produced until the application closes this *output block* (or *bts* returns) will be printed as a block without interruption from other workers.
- Pruning. The efficiency of *mts* depends on keeping the job list non-empty until the end of the computation, without letting it get too large. Depending on the application, there may be a substantial restart cost for each unexplored subtree. Surely there is no need to return a leaf as an unexplored node, and the *prune=0* option checks for this. Further, if an unexplored node has only one child it may be advantageous to explore further, terminating either at a leaf or at a node with two or more children, which is returned as *unexplored*. The *prune=1* option handles this condition, meaning that no isolated nodes or paths are returned as unexplored. Note that

pruning is not a built-in `mts` option; it is an example option that applications may wish to include. An example of pruning is implemented in `mtopsorts`, as mentioned in Section 4.2.

4.4 A second example : spanning trees

In the tutorial [2] a C implementation (`tree.c`) is given for the reverse search algorithm for all spanning trees of the complete graph. An extension of this to generate all spanning trees of a given graph is stated as Exercise 6.3. Applying Phase 1 and 2 as described above results in the codes `atree.c` and `btree.c`. The `mts` wrapper may be directly compiled with `btree.c` to provide the parallel implementation `mtree`. All of these codes are given at the URL [2].

5 Using the `mts` interface

5.1 Building `mts`

Applications may have additional requirements, but the requirements for `mts` are fairly simple. The most important is an MPI implementation. We develop and test on Open MPI and Intel MPI, but any reasonably modern implementation should work. We use the common `mpicc` compiler wrapper to build `mts`. This is installed as part of the MPI implementation; note that when installing MPI via binary packages, the development package (if separate from the runtime support) will likely be needed to compile. `mts` should work on any SMP workstation, properly-configured cluster or general-purpose supercomputer. Configuring MPI for clusters is beyond the scope of this document and depends on the implementation chosen.

5.2 `mts` options

Similar to the `bts_options` array shown in Section 4.2, the built-in `mts` options are collected in an `mts_options` array in `mts.c`:

```

const mtsoption mts_options [] = {           /* external linkage for */
    {"-maxd", 1},                             /* curious applications */
    {"-maxnodes", 1},
    {"-scale", 1},
    {"-lmin", 1},
    {"-lmax", 1},
    {"-maxbuf", 1},
    {"-freq", 1},
    {"-hist", 1},
    {"-checkp", 1},
    {"-stop", 1},
    {"-restart", 1}
};
const int mts_options_size = 11;

```

which currently has eleven built-in options, which are similar to options supported by `mplrs` [4]. The first option `-maxd` takes a parameter that specifies `max_depth`, and likewise `-maxnodes` takes a parameter specifying `max_nodes`. The `max_nodes` parameter is scaled by the scaling factor given by the `-scale` option. The `-lmin` and `-lmax` parameters specify when the budget is modified based on the number of unexplored subtrees available. The idea is to use smaller budgets to quickly grow the job list when it is small, and larger budgets to avoid overhead when the job list is sufficiently large.

More precisely, the `max_depth` is used only if the number of unexplored subtrees available ($|L|$ in [4]) is less than the number of processes in use times the `lmin` value. The `max_nodes` value is scaled by the scaling factor if the number of unexplored subtrees available is larger than the number of processes times the `lmax` value.

While these parameters have default values used in this paper, it is very likely that different applications will require different values to attain good performance. The next two parameters, `-hist` and `-freq` allow one to specify *histogram* and *frequency* files. These collect statistics on the degree of

parallelization obtained, and can easily be plotted using `gnuplot`. This is intended to help the user tune `mts` to their application. We explain the format and usage of these files in Section 7.

The last three options relate to checkpoint files and restarts. The `-checkp` file is used to specify a *checkpoint* file. If the run is interrupted (e.g., by the application using `cleanstop()` to request a checkpoint), `mts` will save the current list of unexplored nodes, shared data, and other information needed to restart. The `-restart` option is used to specify a checkpoint file to restart from. The `-stop` option is used to specify a *stop* file. When this option is used, `mts` will periodically check for the existence of the specified file. If it exists, `mts` will checkpoint and exit when possible.

Finally, the option `-maxbuf` controls the “streaminess” of the output, where the parameter specifies a number of bytes. If the buffer is larger than the specified number of bytes and ends in a newline, it is flushed if possible. Smaller values increase streaminess but may increase overhead. Currently the default parameters are:

$$\text{max_depth} = 2 \quad \text{max_nodes} = 5000 \quad \text{scale} = 40 \quad \text{lmin} = 1 \quad \text{lmax} = 3 \quad \text{maxbuf} = 1048576.$$

6 Computational results

The tests were performed at Kyoto University on `mai32`, a cluster of 5 nodes with a total of 192 identical processor cores, consisting of:

- `mai32abcd`: 4 nodes, each containing: 2x Opteron 6376 (16-core 2.3GHz), 32GB memory, 500GB hard drive (128 cores in total)
- `mai32ef`: 4x Opteron 6376 (16-core 2.3GHz), 64 cores, 256GB memory, 4TB hard drive.

6.1 Linear extensions: `mtopsorts`

The tests were performed using the following codes:

- `VR`: obtained from the Combinatorial Object Server [1], generates linear extensions in lexicographic order via the Varol-Rotem algorithm [8] (Algorithm V in Section 7.2.1.2 of [6])
- `Genle`: also obtained from [1], generates linear extensions in Gray code order using the algorithm of Pruesse and Rotem [7]
- `btopsorts`: derived from the reverse search code `topsorts.c` [2] as described in Section 4
- `mtopsorts`: `mts` parallelization of `btopsorts`.

For the tests all codes were used in count-only mode due to the enormous output that would otherwise be generated. All codes were used with default parameters given at the end of Section 5.2.

The problems chosen were the following graphs which are listed in order of increasing edge density:

- `pm22`: the partial order $a_1 < a_2, \dots, a_{21} < a_{22}, a_1 < a_3 < \dots < a_{21}$ that generates all perfect matchings on 22 elements
- `cat42`: the partial order $a_1 < a_2, \dots, a_{41} < a_{42}; a_1 < a_3, \dots, a_{39} < a_{41}; a_2 < a_4, \dots, a_{40} < a_{42}$ that generates all the 2 x 21 Young Tableaux, whose cardinality is the 21st Catalan number
- `K8,9`: the partial order $a_1 < a_9, \dots, a_1 < a_{17} \quad \text{ldots} \quad a_8 < a_9, \dots, a_8 < a_{17}$ corresponding to the complete bipartite graph $K_{8,9}$ with all edges directed from the smaller part to the larger.

The constructions for the first two partial orders are well known (see, e.g., Section 7.2.1.2 of [6]).

The computational results are given in Table 1. First observe that the reverse search code `btopsorts` is very slow, over 900 times slower than `Genle` and over 70 times slower than `VR` on `pm22` for example. However the parallel `mts` code obtains excellent speedups and is faster than `VR` on all problems when 192 cores are used.

Graph	m	n	No. of perms	VR	Genle	btopsorts	mtopsorts				
							12	24	48	96	192
<i>pm22</i>	22	21	13,749,310,575	179	14	12723 (1)	1172 (11)	595 (21)	360 (35)	206 (62)	125 (102)
<i>cat42</i>	42	61	24,466,267,020	654	171	45674 (1)	4731 (9.7)	2699 (17)	1293 (35)	724 (63)	408 (112)
$K_{8,9}$	17	72	14,631,321,600	159	5	8957 (1)	859 (10)	445 (20)	249 (36)	137 (65)	85 (105)

Table 1: Linear extensions: mai32, times in secs (speedup on btopsorts)

6.2 Spanning trees: mtree

The tests were performed using the following codes:

- **grayspan**: Knuth’s implementation [5] of an algorithm that generates all spanning trees of a given graph, changing only one edge at a time, as described in Malcolm Smith’s M.S. thesis, *Generating spanning trees* (University of Victoria, 1997)
- **grayspspan**: Knuth’s improved implementation of **grayspan**: “This program combines the ideas of **grayspan** and *spspan*, resulting in a glorious routine that generates all spanning trees of a given graph, changing only one edge at a time, with ‘guaranteed efficiency’—in the sense that the total running time is $O(m + n + t)$ when there are m edges, n vertices, and t spanning trees.” [5]
- **btree**: derived from the reverse search code *tree.c* [2] as described in Section 4.4
- **mtree**: mts parallelization of **btree**.

Both **grayspan** and **grayspspan** are described in detail in Knuth [6]. Again all codes were used in count-only mode and with the default parameters given at the end of Section 5.2.

The problems chosen were the following graphs which are listed in order of increasing edge density:

- *8-cage* : the Tutte-Coxeter graph, which is the unique minimal graph of girth 8
- P_5C_5 : 5X5 cylinder $P_5 \square C_5$
- C_5C_5 : 5X5 cylinder $C_5 \square C_5$
- $K_{7,7}$: the complete bipartite graph with 7 vertices in each part
- K_{12} : the complete graph on 12 vertices.

The latter 4 graphs were motivated by Table 5 in [6]: P_5C_5 appears therein and the other graphs are larger versions of examples in the table.

Graph	m	n	No. of trees	grayspan	grayspspan	btree	mtree				
							12	24	48	96	192
<i>8-cage</i>	30	45	23,066,015,625	3166	730	10008 (1)	1061 (9.4)	459 (21)	238 (42)	137 (73)	92 (109)
P_5C_5	25	45	38,720,000,000	3962	1212	8918 (1)	851 (10)	455 (20)	221 (40)	137 (65)	122 (73)
C_5C_5	25	50	1,562,500,000,000	131092	41568	230077 (1)	26790 (8.6)	13280 (17)	7459 (31)	4960 (46)	4244 (54)
$K_{7,7}$	14	49	13,841,287,201	699	460	2708 (1)	259 (10)	142 (19)	68 (40)	51 (53)	61 (44)
K_{12}	12	66	61,917,364,224	2394	1978	3179 (1)	310 (10)	172 (18)	84 (38)	97 (33)	148 (21)

Table 2: Spanning tree generation: mai32, times in secs (speedup on btree)

The computational results are given in Table 2. This time the reverse search code is a bit more competitive: about 3 times slower than `grayspan` and about 14 times slower than `grayspspan` on *8-cage* for example. The parallel `mts` code runs about as fast as `grayspspan` on all problems when 12 cores are used and is significantly faster after that. Near linear speedups are obtained up to 48-cores but then tail off. For the two dense graphs $K_{7,7}$ and K_{12} the performance of `mts` is actually worse with 192 cores than with 96.

7 Evaluating and improving performance

The amount of work contained in a node can vary dramatically between applications, and so `mts` includes features intended to help tune its budgeting to a particular application. We briefly mentioned histogram and frequency files in Section 5; here we take a closer look at using these to improve performance. We focus on `mtopsorts` and the $K_{8,9}$ instance introduced in Section 6.1.

When `mts` is used with the `-hist` option, it periodically writes a line to the specified histogram file. The line contains (in order and separated by whitespace):

1. Time in seconds since execution began (floating point)
2. The number of workers currently busy
3. The number of unexplored nodes currently in the joblist
4. The number of workers owing the master a message about unexplored nodes
5. Unused (currently 0)
6. Unused (currently 0)
7. The total number of unexplored nodes (jobs) the master has seen since execution began.

For example, the following is part of a histogram file generated by `mtopsorts` using 128 cores and default parameters:

```
1.003161 114 1780 114 0 0 6297
2.005773 120 1675 120 0 0 9946
3.006864 117 1679 117 0 0 13665
```

This indicates that at (approximately) 3 seconds, there were 117 busy workers, 1679 jobs in the list, 117 workers owing messages, and 13665 total jobs had existed (i.e. completed jobs, currently executing jobs, and the jobs currently in the list).

The number of workers owing messages to the master is always at least the number of busy workers (since any busy worker must report the number of unexplored nodes it will return). It could become larger, for example if these messages are being delayed due to a saturated interconnect. However, the other entries are likely to be of more interest.

The `mts` distribution includes a script (`plotL.gp`) to produce graphical plots from these histogram files using `gnuplot`. The script produces three plots, with time on the horizontal axis and the second to fourth columns on the vertical axis. Figure 2 contains the plots produced for the `mtopsorts` run with default parameters in the example above.

In Figure 2 using the default parameters, we see that the master is struggling to keep all workers busy despite having jobs available. This suggests that we can improve performance by using better parameters. In this case, a larger `-scale` or `-maxnodes` value may help, since it will allow workers to do more work (assuming a sufficiently large subproblem) before contacting the master.

For example, we might try increasing `-maxnodes`. Figure 3 shows¹ the result of using the (much larger) value 1000000. There we see the opposite problem; many workers become idle because the job list becomes empty. This suggests that `-maxnodes` 1000000 is not suitable in this case; even when the job list is nearly empty the master schedules jobs with large budgets.

¹We omit the third plot; the number of messages owed normally tracks the number of busy workers.

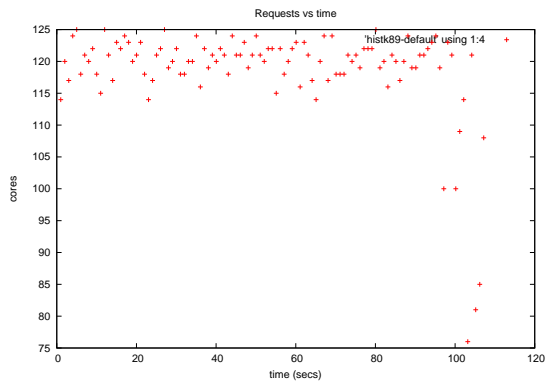




Figure 2: Histograms for mtopsorts on $K_{8,9}$: busy workers, job list size, messages owed

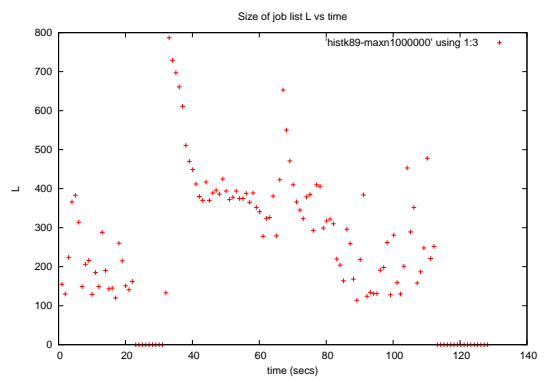
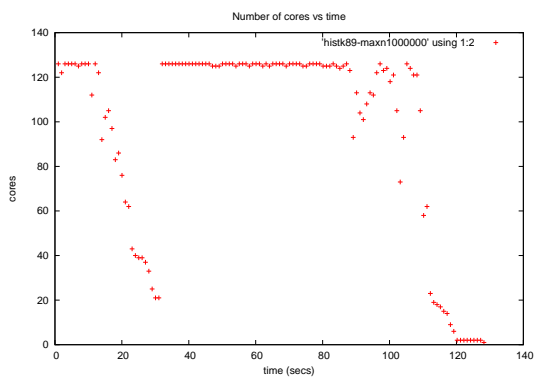


Figure 3: Histograms with $-maxnodes$ 1000000: busy workers (l), joblist size (r)

Instead of dramatically increasing *-maxnodes*, we might try a larger *-scale* together with a modest increase in *-maxnodes*. This will leave budgets relatively low when the job list is small, and increase the budget by a larger amount when the job list is large. Figure 4 shows the result of using a value of 200 for scaling and 10000 for *-maxnodes*. These parameters produce less than half the number of total number of jobs compared to the default parameters, and increase overall performance by about five percent on this particular input.




Figure 4: Histograms with *-scale* 200 *-maxnodes* 10000 on $K_{8,9}$: busy workers (l), joblist size (r)

In addition to the performance histograms, mts can generate *frequency* files. These files contain one number per line, giving the values that the application's *bts()* function returns on each job. In our sample applications, this value is the number of nodes that *bts()* visited – and so the frequency file contains the actual size of each job. This can provide statistical information about the tree that is helpful when tuning the parameters for better performance. For example, it may be particularly helpful to implement and use pruning if many jobs correspond to leaves. Likewise, increasing the budget will have limited effect if only few jobs use the full budget.

The mts distribution contains a script (*plotD.gp*) that uses *gnuplot* to plot the distribution of subproblem sizes. Figure 5 shows the distribution of subproblem sizes that was produced in a run of *mtpsorts* on $K_{8,9}$ with default parameters.




Figure 5: Subproblem sizes in $K_{8,9}$, default parameters: all (l) small subproblems only (r)

8 Conclusions

We have presented a generic framework for parallelizing reverse search codes requiring only minimal changes to the legacy code. Two features of our approach are that the parallelizing wrapper does not need to be user modified and the modified legacy code can be tested in standalone single processor mode. Applying this framework to two very basic reverse search codes we obtained comparable results to that previously obtained by the customized `mplrs` wrapper applied to the the complex `lrs` code [4]. We expect that many other reverse search applications will obtain similar speedups when parallelized with `mts`.

Ongoing work is to extend the use of `mts` to more complex search tree algorithms such as branch and bound, SAT solvers, game tree search, etc. Since these applications require some sharing of common data it will be interesting to see if the same sorts of speedups can be obtained.

References

- [1] Combinatorial Object Server : Linear extensions. <http://theory.cs.uvic.ca/inf/pose/LinearExt.html>
- [2] Avis, D.: Tutorial on reverse search with C implementations (2000). <http://cgm.cs.mcgill.ca/~avis/doc/tutorial>
- [3] Avis, D., Fukuda, K.: Reverse search for enumeration. *Discrete Applied Mathematics* **65**, 21–46 (1993)
- [4] Avis, D., Jordan, C.: `mplrs`: A scalable parallel vertex/facet enumeration code. CoRR [abs/1511.06487](https://arxiv.org/abs/1511.06487) (2015). URL <http://arxiv.org/abs/1511.06487>
- [5] Knuth, D.E.: Programs to read. <http://www-cs-faculty.stanford.edu/~uno/programs.html>
- [6] Knuth, D.E.: *The Art of Computer Programming, Volume 4A*. Addison-Wesley Professional (2011)
- [7] Pruesse, G., Ruskey, F.: Generating the linear extensions of certain posets by transpositions. *SIAM Journal on Discrete Mathematics* **4**(3), 413–422 (1991)
- [8] Varol, Y.L., Rotem, D.: An algorithm to generate all topological sorting arrangements. *The Computer Journal* **24**(1), 83–84 (1981)