Multiprocessor Scheduling Implementation of the Simultaneous Multiple Volume (SMV) Navigator Method

Vladimir Kolmogorov,[1] Thanh D. Nguyen,[2] Anthony Nuval,[2] Pascal Spincemaille,[2] Martin R. Prince,[1] Ramin Zabih,[1] and Yi Wang[1,2]*

The simultaneous multiple volume (SMV) approach in navigator-gated MRI allows the use of the whole motion range or the entire scan time for the reconstruction of final images by simultaneously acquiring different image volumes at different motion states. The motion tolerance range for each volume is kept small, thus SMV substantially increases the scan efficiency of navigator methods while maintaining the effectiveness of motion suppression. This article reports a general implementation of the SMV approach using a multiprocessor scheduling algorithm. Each motion state is regarded as a processor and each volume is regarded as a job. An efficient scheduling that completes all jobs in minimal time is maintained even when the motion pattern changes. Initial experiments demonstrated that SMV significantly increased the scan efficiency of navigator-gated MRI. Magn Reson Med 52:362–367, 2004. © 2004 Wiley-Liss, Inc.

Key words: navigator; simultaneous multiple volume; multiprocessor scheduling; motion suppression; MRI

Physiologic motion is a major source of artifacts in MRI. Difficulties in modeling or predicting physiologic motion have frustrated efforts to suppress motion artifacts. This challenge is addressed by the navigator approach, which measures motion parameters directly from MR signal (navigator echo) and modifies image sampling accordingly (1–3). A widely used navigator algorithm is gating, which constructs the final image from data acquired only when the motion parameter is in a given tolerance range (2,3). This navigator-gating approach has been shown to be effective for suppressing respiratory motion artifacts in MRI.

Two of the main practical issues in applying the navigator-gating method are the selection of the allowed motion range and the increase in scan time. The allowed motion range is preferably near the most frequent position, such as end of expiration, but the drifting behavior of respiration makes it difficult to pin down the most frequent position. Two elegant solutions have been devised to address this problem. One solution is the diminishing variance algorithm (DVA) that automatically selects the smallest motion range for a given acquisition time (4). Another is the phase-encoding ordering with automatic window selection (PAWS) algorithm that introduces view ordering to reduce residual motion artifacts within the

tolerance range and attempts to minimize data acquisition time (5). However, these algorithms only use data acquired in a small motion range for the final image reconstruction. Data acquired outside the used motion range are discarded, resulting in a substantial increase in scan time. The effect on scan time by a navigator method can be characterized by its navigator efficiency E_{nav}:

$$E_{nav} = \frac{\text{time spent on acquiring data used in the final reconstructed image}}{\text{total acquisition time}}.$$

The long scan time of navigator-gating methods is undesirable in clinical practice and it is highly valuable to develop means to improve navigator efficiency (6).

Recently, the simultaneous multiple volume (SMV) acquisition algorithm was introduced to improve navigator efficiency (7). In SMV, several different volumes are acquired simultaneously at various positions in the motion range, such that most of the scan time is utilized for the final image reconstruction. In the ideal case of motion with an unchanging histogram, a set of volumes with sizes proportional to the position probabilities can be completed simultaneously, achieving 100% efficiency. In reality, the histogram of respiration changes during the scan, diminishing the navigator efficiency. The major determinant for navigator efficiency in the SMV approach is the assignment of positions in the motion range to k-spaces of image volumes.

In the preliminary experiment for validating the feasibility of the SMV navigator approach, three large volumes of equal sizes were acquired simultaneously (7). An average of 22.5% (0–50%) increase in navigator efficiency was observed in six subjects. This encourages us to investigate more effective methods for assigning motion positions to image volumes. In this article we report a general implementation of the SMV approach using a multiprocessor scheduling algorithm. Our in vivo experiments demonstrated that the SMV method allows the navigator efficiency to reach 60% when acquiring many (five or more) volumes with a 2 mm tolerance range of diaphragm position.

MATERIALS AND METHODS

Multiprocessor Algorithm

For SMV acquisition, each position bin (or bin group) is considered a processor, each volume a job. The problem is to assign volumes to bins such that the scan time is minimal under the constraint that, upon completion, each volume is acquired at a single bin (or bin group). This is a
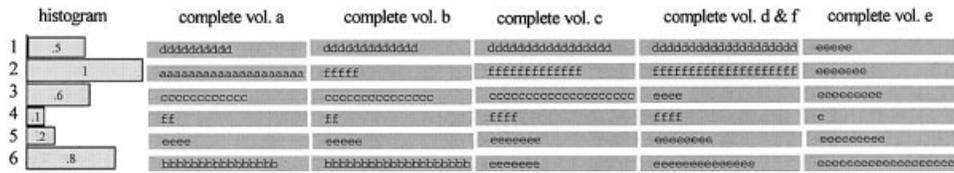
FIG. 1. A simple example of a schedule to complete six volumes a, b, c, d, e, and f using SMV acquisition. The vertical position indicates the position bins and the letters indicate the data acquired in the correspondingly labeled volumes. The histogram is scaled such that the frequency of the fastest bin (bin 2) is 1. The time sequence of the decided schedule may consist of the following steps. 1) Complete volume a using the fastest motion bin 2 while acquiring volume b at bin 6, volume c at bin 3, volume d at bin 1, volume e at bin 5, and volume f at bin 4 (second column). 2) Immediately after completion of volume a, reacquire volume f at the fastest bin 2 (because bin 4 is the slowest), while continuing acquisitions at other bins with volume b completing at bin 6 (third column). 3) Immediately after completion of volume b at bin 6, reacquire volume e at bin 6 (because bin 5 is the second slowest) while completing volume c at bin 3 (fourth column). 4) Complete volume d at bin 1 and volume f at bin 2 (fifth column). 5) Complete volume e at bin 6. In this example, bins 4 and 5 are not really used, but it is useful to assign some volumes to them in case the histogram changes in the future.

multiprocessor scheduling problem, which has been well studied in computer science (8). For reference, we restate the problem as follows. There are a number of parallel processors and a number of independent jobs. Each job is assigned to one of the processors and has to be completed by the same processor. Completing job j on the processor p requires time T(p,j). The problem is to find a schedule that minimizes the completion time of all jobs.

Figure 1 illustrates a simple schedule for completing six volumes under a stable histogram. We formulate here an implementation that accounts for changing histograms. Let N(v) be the total number of *k*-space data points (such as views in standard Cartesian *k*-space sampling) in a volume v, let h(b,t) be the frequency of the instantaneous histogram at bin b and time t, and let n(b,v,t) be the number of points of volume v already acquired at bin b and time t. Then the remaining time needed to complete volume v at bin b and time t is:

$$T(b,v,t) = (N(v) - n(b,v,t))/h(b,t). \qquad [1]$$

An optimal schedule is computed using the procedure of depth-first search with backtracking and pruning (Appendix). Because the motion histogram varies with time, future histograms (and hence the completion times in Eq. 1) are only known approximately. We decided to minimize the mean completion time over all volumes.

As indicated in Fig. 1, the scheduling algorithm may leave certain bins unassigned. These bins are assigned heuristically to certain volumes, which is useful when the histogram changes. Since the respiratory histogram h(b) varies with time in most patients, the histogram is updated periodically (every 25 navigator echoes), and the scheduling algorithm is rerun with updated T(b,v) in Eq. 1 to adapt the volume acquisition to these changes.

For implementation of the SMV method, the whole range of motion is divided into several bins of equal size. Each volume is acquired within a single bin. Data sampling in the bin is in accordance with the 2-bin PAWS algorithm as follows (5). Each bin is further divided into two subbins (Fig. 2a). When the position measured from the navigator is in subbin 1, data is acquired such that *k*-space is sampled from the two edges towards the center symmetrically. When the position is in subbin 2, *k*-space is sampled from the center towards the edges symmetrically.

Experiment

The SMV algorithm was implemented with a 2D fast gradient echo sequence (TE/TR/flip = 1.9/8.5/30, 256 × 160 matrix, 40 cm FOV, 7 mm slice thickness) to acquire multiple sections. Following each navigator echo, two disdaqs (execution of RF and gradient pulses with data sampling disabled) were used to allow the real-time system to update the scanner pulse waveforms and to establish spin equilibrium at the same time. Next, a *k*-space segment of 10 image views was acquired for a volume selected by the real-time SMV algorithm. The real-time software was executed on a workstation (Ultra 1, Sun Microsystems, Palo Alto, CA) connected to a 1.5 T whole-body MR scanner (Signa CV/I with 9.1 software, GE Medical Systems, Waukesha, WI) via a BIT3 cable and a socket protocol. The navigator echo was acquired from a cylinder of tissue through the dome of the right hemidiaphragm using a 2D spatially selective excitation RF pulse (9–11). The diaphragm position was extracted from the navigator profile using an image space least squares algorithm (12). Bin size of 2 mm (subbin size = 1 mm) was used for all acquisitions. For comparison, imaging with the PAWS navigator gating (sequentially acquiring all volumes), imaging under free-breathing without gating, and imaging during breathhold were performed using the same pulse sequence. The SMV navigator acquisition, PAWS navigator acquisition, free-breathing acquisition, and breathhold acquisition of 10 coronal and sagittal sections were obtained on five healthy subjects (age 23–35, all male).

Image Analysis

Images from SMV navigator acquisition, PAWS navigator acquisition, free-breathing acquisition, and breathhold acquisition were randomized and presented to an experienced radiologist (MRP) who was blinded to the acquisition techniques. The radiologist quantified the quality of each section image from the study using the following scoring system: 0 = no motion artifacts, 1 = marginal motion artifacts, 2 = considerable motion artifacts but organs in image can be evaluated, 3 = major motion artifacts and organs in image can not be evaluated, 4 = no definable structure. Then average scores for the coronal and sagittal orientations for all acquisitions were derived from the corresponding 10 individual sections. A two-
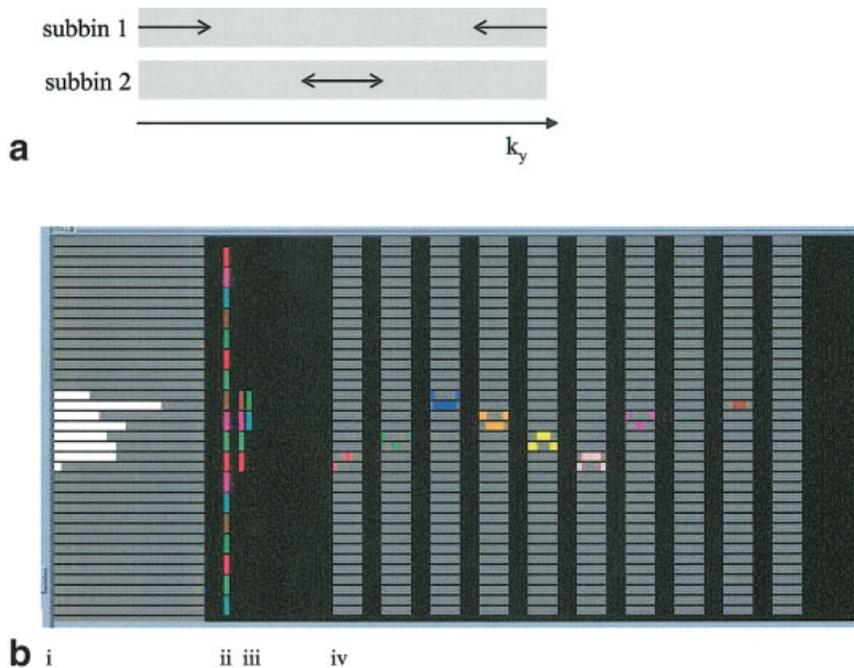
FIG. 2. **a:** *k*-space sampling within each bin using a 2-subbin PAWS algorithm. **b:** Main window of the real-time scheduling program. Different colors correspond to different volumes. Vertical axis shows different bins (diaphragm positions). Horizontal axis: i) instantaneous histogram; ii) assignment of bins to volumes; bins which are not needed according to (iii) are assigned heuristically); iii) assignment of volumes to bins (computed using the multiprocessor scheduling algorithm); iv) volumes 1–10; volumes 3–6 (blue, orange, yellow, and pink) are completed.

tailed paired sample *t*-test was performed on the average scores from the five subjects to determine the statistical significance of the difference between the SMV navigator acquisition and the PAWS navigator acquisition, the difference between the SMV navigator acquisition and the free-breathing no-gating acquisition, and the difference between the SMV acquisition and the breathhold acquisition.

RESULTS

The results are summarized in Table 1. The SMV algorithm doubled the navigator efficiency of the PAWS on average. The SMV, PAWS, and breathhold acquisitions consistently generated image quality acceptable for evaluation (score = 2), while the free-breathing acquisition generated image quality mostly unacceptable for evaluation (score = 3). The scores over 10 sections from the SMV navigator acquisition were consistently better than that from the free-breathing acquisition for all five subjects in both coronal and sagittal sections ($P < 0.01$). The image quality scores for the SMV acquisition, PAWS acquisition, and breathhold acquisition were similar, with no statisti-

cal difference. For all methods, the coronal sections demonstrated slightly higher quality than the sagittal sections (the difference was significant only for the SMV and PAWS acquisitions ($P < 0.05$)), possibly because of susceptibility to the anterior–posterior amplitude of respiration.

An example of coronal imaging is illustrated in Fig. 3 and an example of sagittal imaging is illustrated in Fig. 4 for the free-breathing with no gating, SMV navigator, PAWS navigator, and breathhold acquisitions. The SMV, PAWS, and breathhold images were similar to each other and were consistently better than the free-breathing ungated acquisition.

The navigator efficiency for the SMV algorithm increased with the number of volumes (2D sections), according to simulation using the in vivo navigator record as the motion input (Fig. 5). The rate of increase was the greatest for the first five volumes, where the scan efficiency reached 60%, doubling that of the PAWS algorithm.

DISCUSSION

Our preliminary results demonstrate that navigator efficiency can be increased significantly using the SMV algo-

Table 1
Summary of Image Quality (Mean ± Standard Deviation Over 5 Subjects of Image Scores Averaged Over 10 sections)

Section	Technique	Image quality (navigator efficiency)	*P*-value
Coronal	SMV navigator	1.16 ± 0.34 (66% ± 4%)	
	PAWS navigator	1.14 ± 0.39 (33% ± 6%)	0.87
	Free-breathing	3.02 ± 0.30	0.002
	Breathhold	1.36 ± 0.42	0.45
Sagittal	SMV navigator	1.52 ± 0.29 (72% ± 8%)	
	PAWS navigator	1.52 ± 0.24 (36% ± 5%)	1
	Free-breathing	3.22 ± 0.48	.006
	Breathhold	1.56 ± 0.26	0.48

Score scale: 0 = no motion artifacts, 1 = marginal motion artifacts, 2 = considerable motion artifacts but organs in image can be evaluated, 3 = major motion artifacts and organs in image cannot be evaluated, 4 = no definable structure.

FIG. 3. One of the 10 coronal sections depicting the abdomen, including liver, stomach, and kidneys. **a:** free-breathing with no gating (image quality score IQ = 2). **b:** SMV gating (IQ = 1). **c:** PAWS gating (IQ = 1). **d:** Breathhold (IQ = 1). The SMV, PAWS, and breathhold acquisitions provided similar image quality, which was much superior to the free-breathing acquisition without gating.
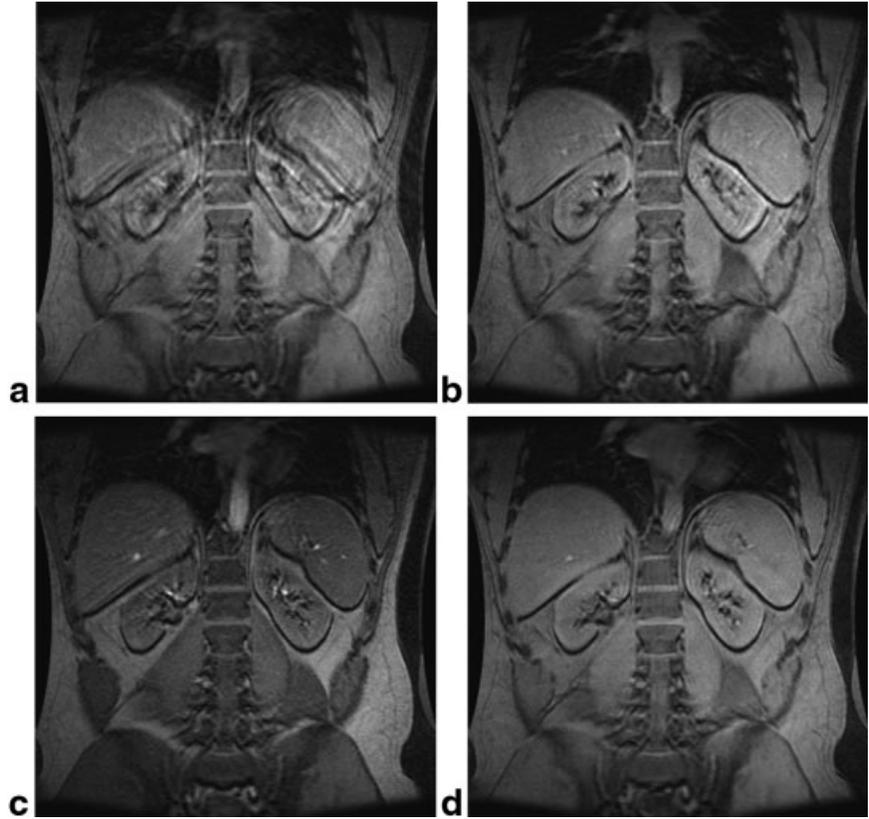
rithm by simultaneously acquiring multiple volumes at various positions. The average efficiency improves as more volumes are acquired (~60% at five volumes). Each volume is acquired with the PAWS algorithm that is efficient and effective for single volume acquisition. This SMV algorithm increases the navigator efficiency while maintaining effective motion suppression. The SMV navigator acquisition provides consistent diagnostic image quality
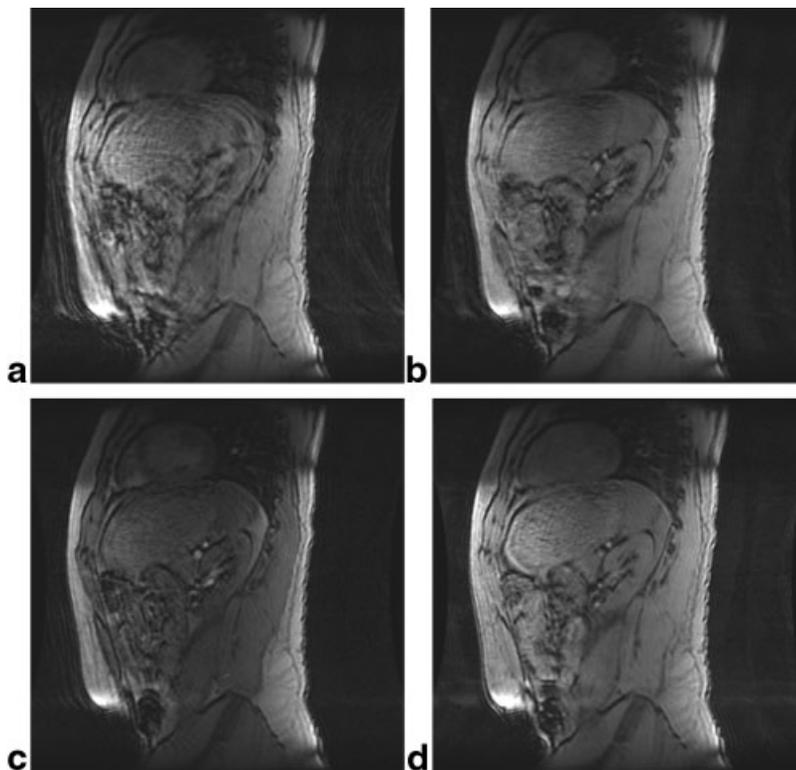

FIG. 4. One of the 10 sagittal sections depicting the abdomen, including liver, stomach, and kidneys. **a:** Free-breathing with no gating (image quality score IQ = 3). **b:** SMV gating (IQ = 1). **c:** PAWS gating (IQ = 2). **d:** Breathhold (IQ = 2). The SMV, PAWS, and breathhold acquisitions provided similar image quality, which was much superior to the free-breathing acquisition without gating.
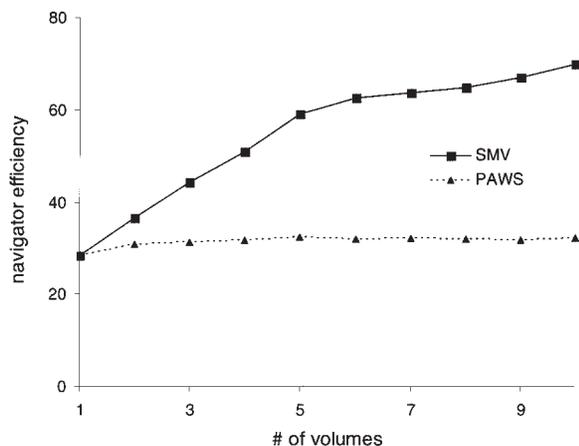
FIG. 5. Navigator efficiency (averaged over five subjects) vs. the number of volumes (volume sizes are kept the same) for both the SMV and PAWS algorithms. The rate of increase in navigator efficiency by SMV is the greatest up to five volumes (~60% efficiency at five volumes and 2 mm diaphragm range) and tapers off. Overall the efficiency of SMV increases with the number of volumes, but the efficiency of PAWS stays approximately the same.

and can be a valuable alternative to breathhold acquisition when patients cannot suspend their respiration for the required scan duration.

Previously, the PAWS algorithm was regarded as the optimal navigator technique because it combines the benefit of automatically selecting the gating window to minimize scan time for a given motion tolerance range (the essence of the DVA algorithm) with the benefit of view ordering to further reduce artifacts from residual motion (4,5). However, the PAWS algorithm does not address the issue of long scan time, the longstanding problem caused by the small motion tolerance in navigator gating. The SMV algorithm solves this problem by acquiring multiple volumes simultaneously at various motion levels, with each volume acquired with the PAWS algorithm. The scheduling implementation of the SMV algorithm maintains the benefit of the PAWS algorithm while drastically improving its efficiency.

The fundamental concept of the SMV algorithm is that the navigator efficiency is increased by acquiring additional volumes at different respiratory positions. Images acquired at different respiratory levels may be used directly without further processing in clinical applications, similar to the routinely performed multiple-section imaging with separate breathholds. These multiple sections are invariably acquired at different respiratory levels because of the inconsistency in breathholding. In many clinical evaluations, contiguous volume coverage is not necessary and gaps between sections are prescribed routinely. Images acquired at different respiratory levels may also be processed to provide contiguous anatomic coverage. Respiratory motion of organs in the body trunk is predominantly in the superior–inferior direction. For both coronal and sagittal acquisition, organs move in plane and will not escape between sections. The anatomy of an organ can be encompassed contiguously and may be patched together using morphing algorithms (13). In typical clinical prac-

tices, radiologists can read through sections acquired at different respiratory levels and form an assessment of the entire organ.

In this SMV navigator implementation with the 2D fast gradient echo sequence, disdaqs (playing RF pulses without collecting data) were used while waiting for the scanner sequence board to update the commands after the execution of the navigator pulse. In future MR scanners with fast computer boards, this waiting may be substantially shortened. The spin equilibrium may be disturbed by the navigator pulse if it induces a nonuniform excitation in the section under acquisition and can be restored by disdaqs. The SMV algorithm itself does not necessarily require the use of disdaqs. For example, disdaqs can be eliminated by alternatively assigning two sections to the same respiratory bin to ensure that a fresh section is always acquired after the navigator echo and by using centric or reverse-centric view orders to smooth k-space signal modulation.

This preliminary study demonstrates that the SMV navigator algorithm provides consistent image quality with acceptably minimal motion artifacts, that the SMV navigator acquisition is consistently and significantly superior to the free-breathing acquisition, and that the image quality of the SMV navigator acquisition is similar to that of the PAWS acquisition and also similar to that of the breathhold acquisition. It is expected that both SMV and PAWS provide the same image quality, because each volume in SMV is acquired with the same motion tolerance window and the same view ordering as in the PAWS algorithm. Since the image quality of the SMV navigator acquisition is not affected by patient effort, the SMV acquisition may be a better choice for imaging patients with difficulties in breathholding. This potential encourages further clinical evaluation of the SMV navigator acquisition.

In conclusion, the simultaneous multiple volume (SMV) algorithm can significantly increase the scan efficiency while maintaining the image quality of navigator-gated acquisition.

APPENDIX

MultiProcessor Scheduling Algorithm

Multiprocessor scheduling is a well-known problem in computer science, and is widely believed to have no fast algorithm for its solution if the number of jobs is large. We were able to achieve a real-time execution of the multiprocessor scheduling (less than a second) if the number of jobs (volumes) was less than 15. In this section we describe our implementation of the multiprocessor scheduling algorithm.

The input to the algorithm is the matrix T describing completion times: T(p,j) is the time to complete the job j on the processor p. The output is an assignment of jobs to processors that minimizes the mean completion time (or, equivalently, the sum of completion times for all jobs).

Let us consider a search tree whose nodes are partial assignments of jobs to processors. The root of this tree is the node with no assignment. All nodes except leaves have child nodes with one more job assignment, and leaves are nodes with complete assignments of all jobs to processors.
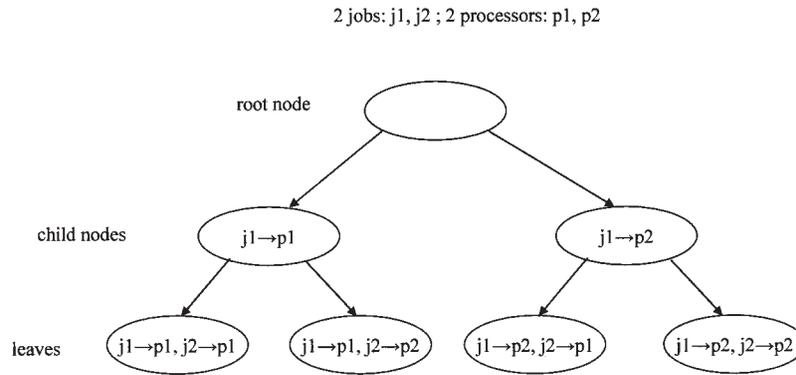
2 jobs: j1, j2 ; 2 processors: p1, p2

FIG. 6. An example search tree used in the scheduling algorithm to search for the optimal assignment of jobs to processors. Nodes contain partial assignments of jobs to processors (the root node represents no jobs being assigned). Leaves contain assignments for all jobs. The schedule is initialized by traversing the tree until arriving at a leaf; a particular branch is taken by choosing the child with the shortest partial completion time. The schedule is updated by traversing the tree towards the leaves. A branch is eliminated from consideration if either of the following two criteria is met: 1) the partial completion time of the associated child node exceeds either the total completion time as determined by the current schedule; 2) the partial completion time associated with the child node constructed by exchanging the processors associated with two jobs.

Links between nodes and their children are edges. Such a tree for the case of two jobs assigned to two processors is illustrated in Fig. 6.

For all nodes we can compute the sum of completion times of assigned jobs. Let us call this sum the partial completion time. For the root node, for example, this time is zero. Our goal is to find the leaf node (complete assignment of all jobs to processors) with the smallest partial completion time.

The number of leaves is too big to search through all of them (for example, if the number of jobs is 15 and the number of processors is 10, then there are 10^{15} leaves). Fortunately, we can safely skip most of them by using pruning (14), as described below.

At all times we maintain a current best schedule (best leaf found so far). If during the search we encounter a leaf which is better than the current best schedule, than we update it. We initialize this schedule using a simple algorithm: we start from the root and go toward the leaves, picking each time the child with the smallest partial completion time. The schedule found by this algorithm is not guaranteed to be the best, but it may be a good estimate.

Our search method is a depth-first search with some pruning. We start from the root of the search tree and go toward the leaves until a leaf is reached or certain pruning criteria are satisfied. In this case we backtrack—we go back along the path and try exploring other children. We use two pruning criteria. The first criterion is based on the fact that we minimize the sum of all completion times: if the partial completion time for a node is greater or equal than the sum for the current best schedule, then we backtrack.

The second criterion can be described as follows. Let us pick two assigned jobs j_a, j_b ($j_a > j_b$) assigned to different processors p_m, p_n, respectively. Let us compare our node with the node obtained by swapping jobs j_a and j_b (i.e., j_a is assigned to p_n and j_b is assigned to p_m). If $T(p_m,j_b) \leq T(p_m,j_a)$, $T(p_n,j_a) \leq T(p_n,j_b)$ and at least one of the inequalities is strict, then the second node is better; thus, we can backtrack. One special case is when both inequalities are actually equalities. In this case the nodes are equivalent, so we can explore either of them; we backtrack if $p_m < p_n$.

REFERENCES

1. Ehman RL, Felmlee JP. Adaptive technique for high-definition MR imaging of moving structures. Radiology 1989;173:255–263.
2. Sachs TS, Meyer CH, Hu BS, Kohli J, Nishimura DG, Macovski A. Real-time motion detection in spiral MRI using navigators. Magn Reson Med 1994;32:639–645.
3. Wang Y, Rossman PJ, Grimm RC, Riederer SJ, Ehman RL. Navigator-echo-based real-time respiratory gating and triggering for reduction of respiration effects in three-dimensional coronary MR angiography. Radiology 1996;198:55–60.
4. Sachs TS, Meyer CH, Irarrazabal P, Hu BS, Nishimura DG, Macovski A. The diminishing variance algorithm for real-time reduction of motion artifacts in MRI. Magn Reson Med 1995;34:412–422.
5. Jhooti P, Gatehouse PD, Keegan J, Bunce NH, Talor AM, Firmin DN. Phase ordering with automatic window selection (PAWS): a novel motion-resistant technique for 3D coronary imaging. Magn Reson Med 2000;43:470–480.
6. Kim WY, Danias PG, Stuber M, Flamm SD, Plein S, Nagel E, Langerak SE, Weber OM, Pedersen EM, Schmidt M, Botnar RM, Manning WJ. Coronary magnetic resonance angiography for the detection of coronary stenoses. N Engl J Med 2001;345:1863–1869.
7. Kolmogorov NK, Watts R, Prince MR, Zabih R, Wang Y. Simultaneous multiple volume (SMV) acquisition algorithm for real time navigator gating. Magn Reson Imag 2003;21:969–975.
8. Brucker P. Scheduling algorithms. New York: Springer; 1998.
9. Pauly J, Nishimura D, Macovski A. A k-space analysis of small-tip-angle excitation. J Magn Reson 1989;81:43–56.
10. Hardy CJ, Cline HE. Broadband nuclear magnetic resonance pulses with two-dimensional spatial selectivity. J Appl Phys 1989;66:1513–1516.
11. Liu Y, Riederer SJ, Rossman PJ, Grimm RC, Debbins JP, Ehman RL. A monitoring, feedback, and triggering system for reproducible breath-hold MR imaging. Magn Reson Med 1993;30:507–511.
12. Wang Y, Grimm RC, Felmlee JP, Riederer SJ, Ehman RL. Algorithms for extracting motion information from navigator echoes. Magn Reson Med 1996;36:117–123.
13. Brown L. A survey of image registration techniques. ACM Computing Surveys 1992;24:325–376.
14. Russell SJ, Norvig D. Artificial intelligence. Englewood Cliffs, NJ: Prentice Hall; 1994. p 84.