

Rapid and Accurate Peptide Identification from Tandem Mass Spectra

Christopher Y. Park,^{†,#} Aaron A. Klammer,^{†,#} Lukas Käll,[†] Michael J. MacCoss,[†] and William S. Noble^{*,†,‡}

Departments of Genome Sciences and Computer Science and Engineering, University of Washington, Seattle, Washington, 98195

Received February 19, 2008

Mass spectrometry, the core technology in the field of proteomics, promises to enable scientists to identify and quantify the entire complement of proteins in a complex biological sample. Currently, the primary bottleneck in this type of experiment is computational. Existing algorithms for interpreting mass spectra are slow and fail to identify a large proportion of the given spectra. We describe a database search program called Crux that reimplements and extends the widely used database search program SEQUEST. For speed, Crux uses a peptide indexing scheme to rapidly retrieve candidate peptides for a given spectrum. For each peptide in the target database, Crux generates shuffled decoy peptides on the fly, providing a good null model and, hence, accurate false discovery rate estimates. Crux also implements two recently described postprocessing methods: a *p* value calculation based upon fitting a Weibull distribution to the observed scores, and a semisupervised method that learns to discriminate between target and decoy matches. Both methods significantly improve the overall rate of peptide identification. Crux is implemented in C and is distributed with source code freely to noncommercial users.

Keywords: Mass spectrometry • peptide identification • proteomics • bioinformatics

1. Introduction

Tandem mass spectrometry is the method of choice for many protein identification studies. However, this technology suffers from an analysis bottleneck, with a need for more efficient and more accurate methods of mapping from the observed fragmentation spectra to the corresponding peptides.

The most widely used methods for peptide identification, such as SEQUEST,¹ MASCOT,² X! Tandem,³ Inspect⁴ and Lookup Peaks,⁵ exploit a database of known protein sequences. For each observed spectrum, these methods search the database for the peptide whose theoretical spectrum best matches the observed spectrum. The resulting peptide–spectrum matches (PSMs) can be ranked using a predefined score function, or by using machine learning methods such as linear discriminant analysis,⁶ support vector machines^{7,8} or decision trees.⁹

In this work, we describe a computational tool called Crux that solves the peptide identification problem efficiently and accurately. Figure 1 is a schematic diagram of Crux's work flow. Crux's key features include the following:

- **Efficient retrieval of candidate peptides.** Given a spectrum with a specified precursor mass-to-charge ratio (*m/z*) and an assumed charge, Crux uses a precomputed peptide

database to retrieve efficiently all peptides whose mass lies within a user-specified window around the target mass (called *candidate peptides*). The database is sorted by peptide mass and is stored on disk, with mass indices stored in memory. We show that, relative to SEQUEST's strategy of reading the protein sequence database from disk for each new spectrum, the indexed database decreases candidate peptide retrieval time from the human protein database by 79.5%–97.2%, depending upon the width of the selected mass window. When run under Windows, Crux's running time is comparable to that of TurboSEQUEST.

- **On-the-fly generation of decoy peptides.** In evaluating the statistical significance of a PSM, mass spectrometrists frequently employ a *decoy database* comprised of protein sequences that have been reversed,¹⁰ shuffled¹¹ or generated from a Markov chain derived from the given *target database*.¹² The number of matches to the decoy database yields an estimate of the false discovery rate associated with a collection of target PSMs. Crux uses this target-decoy strategy, generating decoys by shuffling the target peptides. This approach ensures that each decoy peptide exhibits precisely the same amino acid composition and total mass as the corresponding target decoy. To avoid the memory overhead associated with storing shuffled, non-overlapping decoy peptides in memory, Crux generates decoy peptides on the fly.

* To whom correspondence should be addressed. E-mail: noble@gs.washington.edu.

[†] Department of Genome Sciences, University of Washington.

[#] These authors contributed equally to this work.

[‡] Department of Computer Science and Engineering, University of Washington.

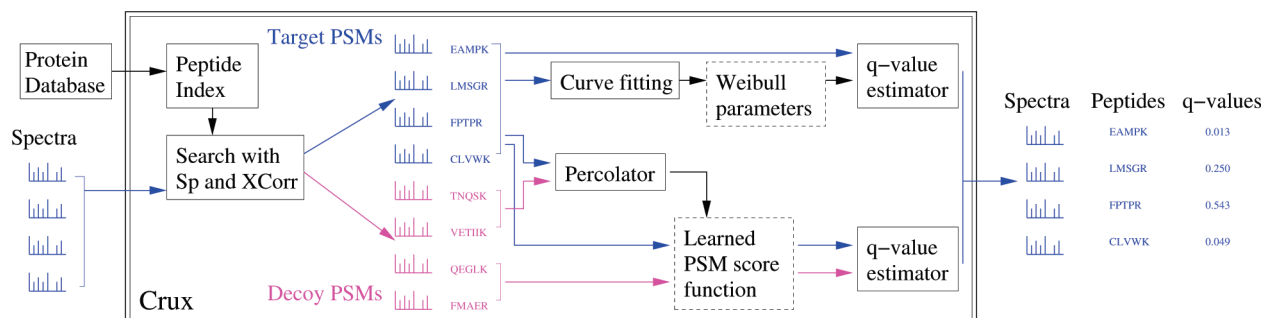


Figure 1. The Crux algorithm. Crux takes as input a collection of fragmentation spectra and a target protein sequence database, and produces a list of peptide–spectrum matches, each with an associated q value, a measure of false discovery rate.

- Two state-of-the-art PSM reranking algorithms.** After scoring all peptides with respect to a given spectrum, the top-ranked PSM must be ranked with respect to other PSMs from the same data set. Some score functions, such as SEQUEST's cross-correlation score (X_{corr}), have been used to carry out both relative ranking and absolute ranking; however, numerous studies have demonstrated that better absolute rankings can be achieved by using machine learning methods.^{6,7,9,8} Crux incorporates two recently described methods for reranking PSMs. The first method, called Percolator,⁸ trains a machine learning method to discriminate between target and decoy PSMs. This dynamically trained model incorporates specific characteristics of the given data set. The second method fits the observed X_{corr} scores to a Weibull distribution and uses the resulting distribution to compute accurate p values.¹³
- Accurate false discovery rate estimates.** Crux uses well-established statistical methods to estimate false discovery rates based upon decoy PSMs.¹⁴ Crux reports, along with each PSM, a q value, which is defined as the minimal false discovery rate threshold at which a given PSM is deemed correct.¹⁵ Depending upon the postprocessing method, Crux estimates these q values either using decoy PSMs or, for the Weibull curve fitting, directly from the estimated p values. Perhaps most significantly, from the perspective of the mass spectrometry research community, Crux provides all of the above in a stand-alone C program which is distributed, with source code, free for noncommercial use (<http://noble.gs.washington.edu/proj/crux>). Below, we demonstrate that Crux can accurately emulate SEQUEST, and that Crux is both efficient and accurate, yielding a significant improvement in running time and an increase in the number of correctly identified spectra.

2. Materials and Methods

All Linux experiments were performed on a using a single CPU of a 2 x Intel Xeon 2.33 GHz 4-core processor machine with 4MB cache and 16GB of RAM running RedHat Linux. All Windows experiments were performed using a Dell OptiPlex 745 Windows XP machine with an Intel Core 2 Duo E6400 processor (2.13 GHz, 2MB L2) and 2GB RAM with OpenSSH/Cygwin installed.

Using Crux, we precomputed a peptide database for two protein sequence databases: the predicted open reading frames from yeast *Saccharomyces cerevisiae* (released 2004-04-02, 3.6MB, 6298 proteins) and human (released 2004-04-21, 16.8MB, 27 209 proteins). The indexes contained tryptic peptides of length 6–50 amino acids and mass 200.0–7200.0 Da while

allowing missed cleavage sites. The indexing procedures required, respectively, approximately 2 and 4 min of real time and produced indices of 69 and 240 MB in size.

For validation, we use a publicly available tandem mass spectrometry data set,¹⁶ available as the 60cm data set at <http://noble.gs.washington.edu/proj/retention/data/data.html>. The data set consists of a collection of 18 149 spectra derived from analysis of soluble yeast whole-cell lysate on an LTQ ion trap mass spectrometer as described previously.¹⁶

3. Results

3.1. Efficient Retrieval of Candidate Peptides. When presented with a new query spectrum and its associated precursor mass, Crux must first retrieve from the sequence database all of the candidate peptides, that is, peptides whose masses lie within a specified range of the precursor mass. A straightforward retrieval method is to read the entire database from disk for each query, evaluating candidate peptides as they are encountered. This approach, which is used by SEQUEST, is space-efficient but can be slow for large databases or large numbers of query spectra.

In addition to implementing the above approach, Crux offers an alternative strategy that allows more efficient candidate peptide retrieval. In this approach, the user must preprocess a given protein database to produce a binary protein database and a binary peptide index. The preprocessing step sorts all of the peptides in the database by mass and stores pointers to their locations in the protein database (sequence index, start position, and peptide length). Crux can then quickly retrieve a set of candidate peptides via a simple range query on the peptide index. Upon execution, Crux accesses the binary protein database and the peptide index as memory-mapped files. This approach offloads to the operating system the decision about how much of the database to store in memory, ensuring that small databases are read fully into memory but large databases are paged into and out of memory only as needed.

To demonstrate the efficiency of Crux's candidate peptide retrieval strategy, we searched with a collection of 100 spectra against predicted human protein sequences from the human genome, generating candidate peptides from a series of increasingly wider mass windows. Figure 2A plots the average search time per spectrum running under Linux as a function of mass window width. Analogous plots with similar trends for a search against the yeast *S. cerevisiae* are shown in Supporting Information. In general, Crux runs more quickly than SEQUEST, and the improvement is larger when Crux uses a peptide index. The difference in running time is most pronounced when the

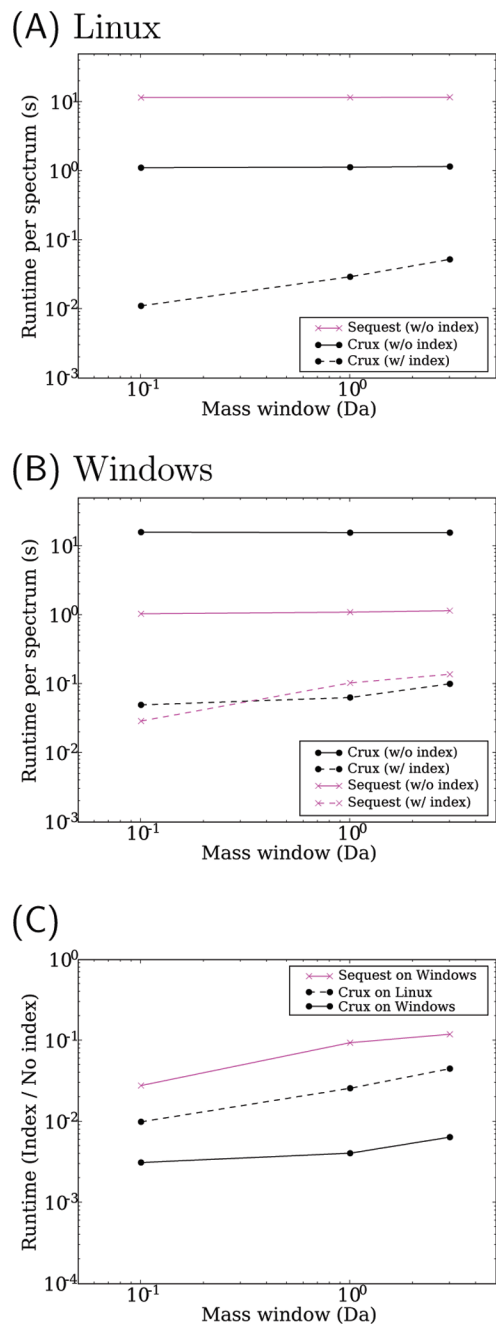


Figure 2. Rapid retrieval of candidate peptides. (A and B) The average running time required to search 100 tandem mass spectra against the human protein databases on computers running the Linux (A) or Windows (B) operating systems, using SEQUEST and Crux with and without indices. Running time is plotted as a function of the mass tolerance used to define candidate peptides. The Linux plot includes only three series because we do not have a Linux implementation of TurboSEQUEST. (C) The figure plots ratio of running times for indexed versus nonindexed searches for Crux on Windows and Linux and for SEQUEST on Windows.

mass tolerance window is small, because Crux does not need to scan the entire database to retrieve a small collection of candidate peptides. For example, with a mass tolerance window of 0.1, on Linux SEQUEST requires an average of 11.6 s clock time per spectrum, whereas Crux requires an average of 10 ms. This represents a 99.9% reduction in running time. However, even

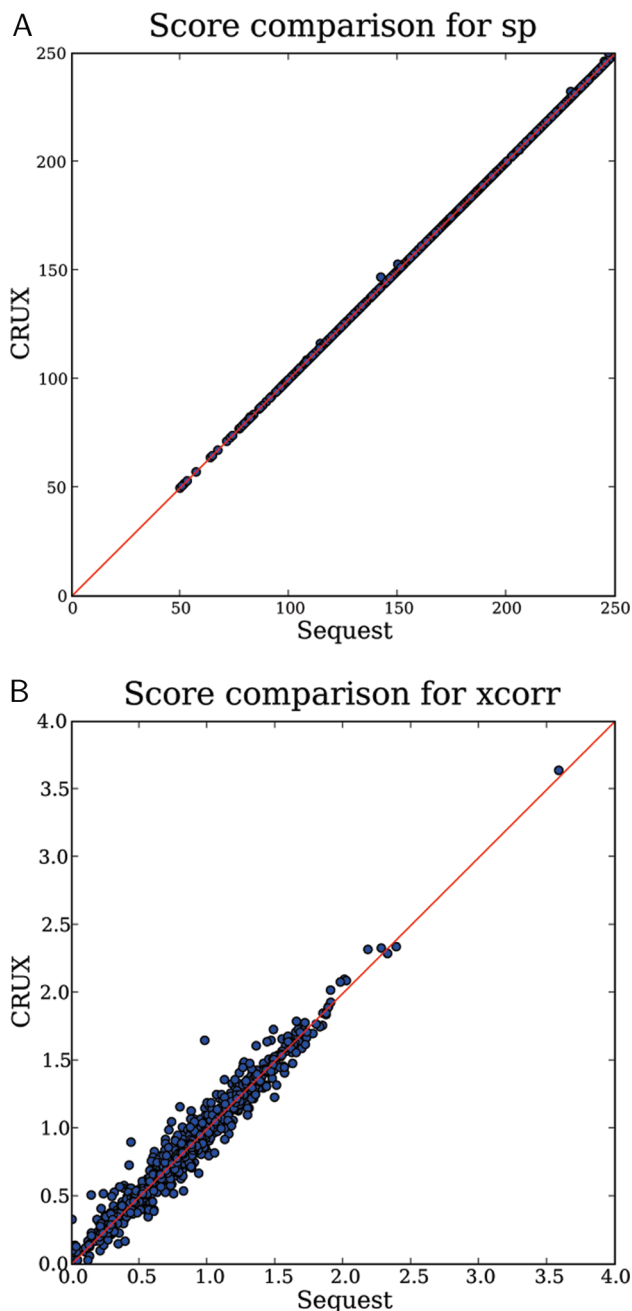


Figure 3. Reimplementation of S_p and X_{corr} scoring functions. The figure plots, for a collection of peptide–spectrum matches, the S_p (A) and X_{corr} (B) scores as computed by Crux as a function of the same scores as computed by SEQUEST.

when we use a mass tolerance of 3, Crux yields a 99.5% speed improvement (52 ms versus 11.7 s).

Crux is not the first attempt to improve the speed of candidate peptide retrieval. Thermo Scientific (Waltham, MA) markets a version of SEQUEST called TurboSEQUEST as part of the Bioworks 3.1 software, which implements a peptide index using a proprietary scheme. In Figure 2B, we compare Crux to TurboSEQUEST under Windows XP. Execution of Crux on Windows was performed from within the Unix emulator Cygwin, which induces a slight overhead. This may explain, in part, the reversal between the nonindexed versions of Crux and SEQUEST between panels A and B: under Windows in this mode, SEQUEST runs significantly more quickly than Crux. Nonetheless,

Rapid, Accurate Peptide Identification

with indexing, Crux's speed is comparable to that of TurboSEQUEST. In Figure 2C, we directly assess the speedup due to indexing, plotting the ratio of the running time with and without indexing for Crux under Linux and Windows and for SEQUEST under Windows. For all three mass window sizes, Crux achieves a much larger speedup than SEQUEST.

A surprising property of Figure 2A,B is the flatness of the two series corresponding to SEQUEST and Crux without an index. This flatness indicates that the running time for retrieving candidate peptides is dominated by the fixed cost associated with scanning the file from disk. Relative to this fixed cost, the time required to calculate the preliminary score S_p is negligible.

3.2. Accurate Reimplementations of S_p and X_{corr} . SEQUEST is the first and one of the most widely used database search methods for peptide identification from tandem mass spectra. Crux, therefore, begins with a reimplementation of the core of SEQUEST. After retrieving all of the candidate peptides for a given spectrum, the candidates are ranked according to the preliminary SEQUEST score S_p . Subsequently, the 500 top-scoring peptides are reranked according to the cross-correlation score X_{corr} .

For S_p scoring, Crux applies a series of preprocessing steps to the observed spectrum. We summarize the main operations here. First, Crux takes the square root of each peak's intensity and rounds each peak's m/z to the nearest integer. Second, Crux normalizes the peak intensities to sum to 100. Third, Crux extracts the top 200 highest intensity peaks. Crux compares the resulting observed spectrum v with the b and y fragment ions u from the candidate peptide to compute S_p as follows:

$$S_p(u, v) = \frac{\left(\sum_{i=1}^N v_i \delta(u_i) \right) \left(\sum_{i=1}^N \delta(u_i, v_i) \right) (1 + 0.075R(u, v))}{\sum_{i=1}^N \delta(u_i)} \quad (1)$$

where N is the maximal m/z value, $R(u, v)$ is the maximum number of consecutive b- or y-ions from the theoretical spectrum that appear in the observed spectrum, and $\delta(x) = 0$ if $x = 0$ and $\delta(x) = 1$ otherwise.

X_{corr} scoring also requires some preprocessing of the observed spectrum. As before, Crux first takes the square root of each peak's intensity and rounds each peak's m/z to the nearest integer. Second, Crux divides the spectrum into 10 regions and normalizes the spectrum intensity in each region to maximum of 50.0. To create the theoretical spectrum from the candidate peptide, Crux constructs a spectrum with peak intensity 50.0 for b- and y-ions, 25.0 for the ± 1 m/z flanking peaks of each b- and y-ion, and 10.0 for ions resulting from a neutral loss of ammonia from b- or y-ions or neutral loss of water from b-ions. Crux then computes X_{corr} as follows:

$$X_c(u, v) = \frac{0.015 \sum_{i=1}^N u_i v_i}{\sum_{\tau=-75}^{75} \sum_{i=1}^N u_i v_{i-\tau}} \quad (2)$$

Figure 3 shows that the Crux-calculated versions of S_p and X_{corr} closely match the SEQUEST-calculated versions of these scores. We also analyzed MALDI-TOF/TOF and QTOF spectra¹⁷ against the same database, in order to more fully sample the space of possible PSMs, and again our Crux-calculated scores are similar to the SEQUEST versions (see Supporting Informa-

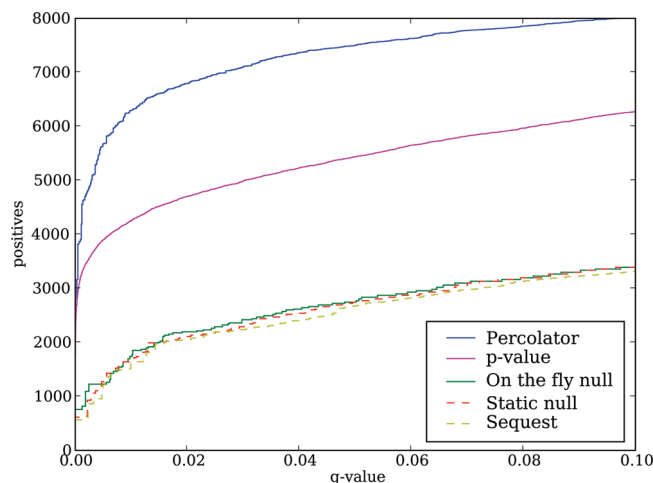


Figure 4. Improved peptide identification. The figure plots, for a variety of database search algorithms, the number of PSMs as a function of the estimated q value. The five series correspond to SEQUEST, Crux's implementation of SEQUEST using a static shuffled decoy database, Crux using decoys generated on-the-fly, Crux with the p value calculation enabled, and Crux with the Percolator postprocessor enabled.

tion). A vast majority of the PSMs receive identical scores using either method. Furthermore, as shown in Figure 4 (described in section 3.5), the Crux and SEQUEST algorithms perform very similarly in practice.

3.3. Null Peptide Generation. When identifying peptides from tandem mass spectra, a commonly used null model is to search the given set of spectra against a *decoy database*. A decoy database is a database of amino acid sequences that is derived from the original protein database (called the *target database*) by reversing the target sequences,¹⁰ shuffling the target sequences,¹¹ or generating the decoy sequences at random using a Markov model with parameters derived from the target sequences.¹² Ideally, the decoy database should contain peptide-like amino acid sequences that are not in the target database.

Crux uses a decoy strategy; however, rather than generating a fixed decoy database, Crux generates decoys on the fly. For each comparison between a spectrum and a target peptide, Crux generates a corresponding decoy peptide by shuffling the nonterminal amino acids of the target peptide. Consequently, Crux reports for every target PSM one or more corresponding decoy PSMs. On-the-fly decoy generation introduces a small computational overhead but saves space, because the decoy database does not need to be stored.

X! Tandem³ will automatically generate a reversed decoy database from the given database, but by definition, a reversed decoy database can only be generated once for a given database. Crux, in contrast, generates a new shuffled decoy peptide for every peptide-spectrum comparison. Thus, if a particular peptide is selected as a candidate for 100 different spectra, then it will be shuffled 100 different times. This approach eliminates the possibility that a single decoy peptide will match well against many spectra.

An additional advantage of generating decoys on the fly is that the resulting decoys share exactly the same total mass and amino acid composition as the target. When using a static decoy database, ensuring that the target and decoy peptide mass and amino acid composition stays the same is very difficult. For example, shuffled and Markov chain-generated

decoy databases will not retain these properties. A reversed database will retain these properties, but only when trypticity is not enforced. On-the-fly decoy generation has the further advantage that, unlike reversed decoys, multiple decoys can be generated for a single target. These additional decoys can be useful if the user wants more accurate confidence estimates.

3.4. Methods for Postprocessing Peptide–Spectrum Matches. Identifying peptides from tandem mass spectra via database search is, fundamentally, a ranking procedure, in which good PSMs are ranked above poor PSMs. This ranking task can be usefully separated into two phases: a *relative ranking* task, in which all of the candidate peptides are ranked relative to a single spectrum to identify a single best PSM for that spectrum, and an *absolute ranking* task, in which the top-ranked PSMs for multiple spectra are ranked relative to one another. Absolute ranking is, by definition, more difficult than relative ranking. Furthermore, empirical evidence suggests that, although X_{corr} does a good job at relative ranking, it performs more poorly when attempting to rank PSMs absolutely.^{6,7} Crux, therefore, implements two recently described methods for postprocessing PSMs to improve their absolute ranking.

The first postprocessor is a semisupervised machine learning method called Percolator that dynamically learns an absolute PSM ranking function.⁸ Percolator uses a subset of the high-scoring target PSMs as positive examples and all of the decoy PSMs as negative examples to train a discriminative classification algorithm called a support vector machine (SVM).^{18,19} Each PSM is characterized using a collection of features that capture characteristics of the spectrum, the peptide, and the quality of the match between the spectrum and the peptide. Details of the Percolator algorithm are given in ref 8.

The second postprocessing method fits the observed distribution of candidate PSM X_{corr} scores to a Weibull distribution. To do this fitting in an unbiased way, Crux computes 500 additional X_{corr} scores at random from the candidate set, in addition to scoring the top 500 candidates selected by S_p . The resulting Weibull distribution is then used to compute a p value for the observed maximum score, and the score is further corrected to account for the total number of candidates that were considered. Details of this p value estimation procedure are given in ref 13.

Crux produces as output a ranked list of target PSMs, one per spectrum in the given data set. Along with each PSM, Crux reports a q value, which is defined as the minimal false discovery rate at which this PSM is called positive.^{15,20} False discovery rates are estimated using a previously described method.¹⁴ When running without postprocessing, the null distribution is estimated using the decoy PSMs. When Percolator is used, half of the decoy PSMs are used to train the classifier, and half are reused to estimate the final q values. For the empirical p value calculation, decoy PSMs are not necessary, because the FDR can be estimated directly from the computed p values. Eliminating the decoys leads to an approximately 50% decrease in running time for this search mode versus the standard target-decoy search approach.

3.5. Experimental Comparison of Methods. To compare the performance of SEQUEST and the various algorithms implemented in Crux, we used a collection of 18 149 spectra derived from a yeast whole-cell lysate (see Materials and Methods). Each peptide identification method searches against the yeast proteome and produces a list of 35 676 PSMs ranked by q value, computed as described in section 3.4 (the number of spectra differs from the number of PSMs because some spectra are

searched multiple times, once for each possible charge state). For SEQUEST and for Crux with a static decoy database, we use a decoy database of shuffled proteins. Figure 4 plots, for each method, the number of accepted PSMs as a function of q value threshold.

Three of the series in Figure 4 lie very close to one another. The similarity between the results given by SEQUEST and the Crux reimplementation of SEQUEST provides evidence that Crux accurately reimplements the SEQUEST algorithm. The series corresponding to on-the-fly decoy generation is also essentially identical to the two series generated using static decoys. Although the on-the-fly decoy generation does not lead to better results, we believe that it is a better null model, for the reasons given in section 3.3. We therefore use this null model by default in Crux.

Finally, consistent with previously reported results,^{8,13} the p value calculation yields a large improvement over simple ranking by X_{corr} , and using Percolator yields an even more dramatic improvement. At a q value threshold of 0.01, X_{corr} ranking yields 1609 PSMs, ranking by X_{corr} p values yields 4904 PSMs, and ranking by Percolator score yields 6294 PSMs. It is not surprising that Percolator performs better than the p value calculation, because Percolator exploits a wide variety of PSM features, whereas the p value is based only on the observed X_{corr} distribution and the number of candidate peptides. On the other hand, the p value calculation is at least twice as fast as all of the other methods reported here, because the q values are calculated analytically, rather than by searching a decoy database.

4. Discussion

One practical challenge impeding research on mass spectrometry peptide identification algorithms is that source code for two of the most widely used algorithms—SEQUEST and MASCOT—are not publicly available. For noncommercial users, Crux remedies this situation, while also making available a fast candidate peptide retrieval algorithm and several state-of-the-art postprocessing algorithms. Although Crux is implemented in C, the code design is object oriented. The core of the Crux search engine is currently S_p and X_{corr} , but the code modules for candidate retrieval or postprocessing could be easily modified to work with alternative searching algorithms.

Recently, several groups have described an alternative approach to candidate peptide retrieval using metric space embeddings.^{21,22} These approaches promise even larger speed-ups at the expense of failing to retrieve a small fraction of the candidate peptides.

The comparison of postprocessing methods in Figure 4 shows that Percolator significantly improves upon the p value-based ranking. However, computing p values has the advantage of cutting the overall search time in half by eliminating the need for a decoy database. One could imagine coupling these two postprocessing algorithms by providing p values as inputs to Percolator; however, doing so would require decoy PSMs to train Percolator. Hence, we believe that decoy database searching will continue to be a valuable component of database searching, even in the presence of analytically calculated significance measures.

Crux is under development, and several significant enhancements are planned for the near future. Some are straightforward, including allowing alternative file formats for input and output (currently, Crux uses the MS2 and SQT file formats²³), improved support for searches that allow multiple masses per

amino acid, and parallelization of the code. Our indexing strategy was specifically designed with the latter in mind: the parallelization will divide the database among processors, sending each observed spectrum to the processor that stores the appropriate mass range. We also plan to extend the peptide index to allow storage of additional pieces of precomputed information, such as predicted peak intensities, predicted peptide retention time, or predicted proteotypic peptides.²⁴ Finally, we have plans for algorithmic improvements, including increasing the variety of features used by Percolator, improving the accuracy of our q value estimates using existing methods,²⁵ and implementing algorithms for protein-level identification.

Acknowledgment. We thank Jimmy Eng for useful discussions. Funding: NIH awards P41 RR11823 and R01 EB007057.

Supporting Information Available: Figures of reimplementation of S_p and X_{corr} scoring functions and rapid retrieval of candidate peptides. This material is available free of charge via the Internet at <http://pubs.acs.org>.

References

- Eng, J. K.; McCormack, A. L.; Yates, J. R., III. An approach to correlate tandem mass spectral data of peptides with amino acid sequences in a protein database. *J. Am. Soc. Mass Spectrom.* **1994**, *5*, 976–989.
- Perkins, D. N.; Pappin, D. J. C.; Creasy, D. M.; Cottrell, J. S. Probability-based protein identification by searching sequence databases using mass spectrometry data. *Electrophoresis* **1999**, *20*, 3551–3567.
- Craig, R.; Beavis, R. C. Tandem: matching proteins with tandem mass spectra. *Bioinformatics* **2004**, *20*, 1466–1467.
- Tanner, S.; Shu, H.; Frank, A.; Wang, L. C.; Zandi, E. InsPecT: Identification of posttranslationally modified peptides from tandem mass spectra. *Anal. Chem.* **2005**, *77*, 4626–4639.
- Bern, M.; Goldberg, D.; Cai, Y. Lookup peaks: A hybrid de novo sequencing and database search for protein identification by tandem mass spectrometry. *Anal. Chem.* **2007**, *79*, 1393–1400.
- Keller, A.; Nesvizhskii, A. I.; Kolker, E.; Aebersold, R. Empirical statistical model to estimate the accuracy of peptide identification made by MS/MS and database search. *Anal. Chem.* **2002**, *74*, 5383–5392.
- Anderson, D. C.; Li, W.; Payan, D. G.; Noble, W. S. A new algorithm for the evaluation of shotgun peptide sequencing in proteomics: support vector machine classification of peptide MS/MS spectra and SEQUEST scores. *J. Proteome Res.* **2003**, *2*, 137–146.
- Käll, L.; Canterbury, J.; Weston, J.; Noble, W. S.; MacCoss, M. J. A semi-supervised machine learning technique for peptide identification from shotgun proteomics datasets. *Nat. Methods* **2007**, *4*, 923–925.

- Elias, J. E.; Gibbons, F. D.; King, O. D.; Roth, F. P.; Gygi, S. P. Intensity-based protein identification by machine learning from a library of tandem mass spectra. *Nat. Biotechnol.* **2004**, *22*, 214–219.
- Moore, R. E.; Young, M. K.; Lee, T. D. Qscore: An algorithm for evaluating sequest database search results. *J. Am. Soc. Mass Spectrom.* **2002**, *13*, 378–386.
- Klammer, A. A.; MacCoss, M. J. Effects of modified digestion schemes on the identification of proteins from complex mixtures. *J. Proteome Res.* **2006**, *5*, 695–700.
- Colinge, J.; Masselot, A.; Giron, M.; Dessingy, T.; Magnin, J. OLAV: Towards high-throughput tandem mass spectrometry data identification. *Proteomics* **2003**, *3*, 1454–1463.
- Klammer, A. A.; Park, C. Y.; Noble, W. S. Statistical calibration of peptide-spectrum score functions. *J. Proteome Res.*, submitted for publication, **2007**.
- Benjamini, Y.; Hochberg, Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc. B* **1995**, *57*, 289–300.
- Storey, J. D.; Tibshirani, R. Statistical significance for genome-wide studies. *Proc. Natl. Acad. Sci. U.S.A.* **2003**, *100*, 9440–9445.
- Klammer, A. A.; Yi, X.; MacCoss, M. J.; Noble, W. S. Improving tandem mass spectrum identification using peptide retention time prediction across diverse chromatography conditions. *Anal. Chem.* **2007**, *79*, 6111–6118.
- Klimmek, J.; Eddes, J. S.; Hohmann, L.; Jackson, J.; Peterson, A.; et al. The standard protein mix database: a diverse data set to assist in the production of improved peptide and protein identification software tools. *J. Proteome Res.* **2008**, *7*, 96–1003.
- Boser, B. E.; Guyon, I. M.; Vapnik, V. N. A Training Algorithm for Optimal Margin Classifiers. In *5th Annual ACM Workshop on COLT*; Haussler, D., Ed.; ACM Press: Pittsburgh, PA, 1992; pp 144–152.
- Noble, W. S. What is a support vector machine? *Nat. Biotechnol.* **2006**, *24*, 1565–1567.
- Käll, L.; Storey, J. D.; MacCoss, M. J.; Noble, W. S. Assigning significance to peptides identified by tandem mass spectrometry using decoy databases. *J. Proteome Res.* **2008**, *7*, 29–34.
- Ramakrishnan, S. R.; Mao, R.; Nakorchevskiy, A. A.; Prince, J. T.; Willard, W. S. A fast, coarse filtering method for peptide identification by mass spectrometry. *Bioinformatics* **2006**, *22*, 1524–1531.
- Dutta, D.; Chen, T. Speeding up tandem mass spectrometry database search: metric embeddings and fast near neighbor search. *Bioinformatics* **2007**, *23*, 612–618.
- McDonald, W. H.; Tabb, D. L.; Sadygov, R. G.; MacCoss, M. J.; Venable, J.; et al. MS1, MS2, and SQT-three unified, compact, and easily parsed file formats for the storage of shotgun proteomic spectra and identifications. *Rapid Commun. Mass Spectrom.* **2004**, *18*, 2162–2168.
- Mallick, P.; Schirle, M.; Chen, S. S.; Flory, M. R.; Lee, H.; et al. Computational prediction of proteotypic peptides for quantitative proteomics. *Nat. Biotechnol.* **2006**, *25*, 125–131.
- Storey, J. D. A direct approach to false discovery rates. *J. R. Stat. Soc.* **2002**, *64*, 479–498.

PR800127Y