INTEGRATED FOREST BIOREFINERIES: PRODUCT-BASED ECONOMIC FACTORS

Cooper, Brian L.[1]*, London, Jeffrey R.*[1]*, Mellon, Robert J.*[1]*, Behrens, Michael A.*[1]

ABSTRACT

Chemicals and fuels produced at a biorefinery require complex operations and often specialized equipment, especially in those cases in which a hybrid process includes both biochemical operations (utilizing the carbohydrate content of the feedstock) and thermochemical processes (utilizing the residue after some or all of the carbohydrate has been removed). Given that the products of the two types of operation are not entirely fungible, comparative studies on the optimal balance of products from a biorefinery can be difficult. Because biochemical processes normally require a particular feedstock (monomeric sugars in solution), while thermochemical processes are relatively feedstock- agnostic, it stands to reason that an optimization of an integrated process that includes both types of operation is possible by examining the downstream effects of variations in early process parameters. The question is essentially, can yield sacrifices in feed preparation for biochemical operations be recouped by lower operational costs and/or higher quality thermochemical products? Examining the output of a biorefinery in terms of relative difficulty in processing and subsequent level of operational cost in the design phase can aid in making decisions about process paths and potentially make the case for more operational flexibility.

INTRODUCION

The integrated forest biorefinery has been described as a factory built to exploit the chemistry of natural resources for the manufacture of a wide variety of chemicals, fuels, and other intermediates. Taking advantage of the complex polymeric structures evident in naturally occurring biomass feedstock is a logical goal and one with enticing potential given the reactivity (and digestibility) of carbohydrates as well as the similarity of the variety of phenylpropenyl units extant in lignin residues to valuable and complex chemical commodities. In principle this makes good sense; however, recalcitrant lignocellulosic biomass processing encounters many challenges when moving from plant structures to salable products.

Most cost models put a definitive mark on yields at each step, below which the economics of the process are unfavorable. The greatest return for extra effort (in biorefining and just about any other process) lies in the early steps; as a direct cost driver, the feedstock and associated transport costs can seemingly make or break a venture. Increasing the amount per acre harvested has plant biologists

---

[1] Hazen Research, Inc., Golden, Colorado

across the world attempting to select or engineer biomass that allows for reduction in the cost of feedstock by any means (ignoring for now the complications regarding harvesting and drying procedures, the incremental cost of transporting biomass, and the cost of sizing the biomass). As in most research projects, the last bit of recovery proves to be the most difficult in processing lignocellulosic feed material into monomeric carbohydrate slurries appropriate for fermentation. Given that the resultant sugars will begin to degrade under the same pretreatment conditions in which they are liberated from their polymeric forms, those sugars released early in the process are difficult to collect because the reaction is still underway. The mature pretreatment/saccharification process is therefore a delicate balance of closely controlled aggressive chemistry that will yield the desired product and remove it from the system before the majority of the carbohydrates begin to react into furans and furfurals. All of this is performed in a vessel that can handle high temperatures, pressures, and aggressive chemistry in a reducing environment, normally under constant duty.

As one would expect, this requires very specialized equipment, particularly when dilute acid hydrolysis is the primary pretreatment mechanism. Alternative chemistries have shown promise and may well be superior in terms of carbohydrate liberation; however, each still results in a difficult process to control. Even minor changes in feedstock particle size, not to mention structural chemistry, can render yields and reagent usage unacceptable. Contrast this with the goals of the bulk feedstock supplier. It is in the interest of the forest harvester to transport feedstock to the refinery in bulk form; it is therefore in the interest of the refinery to be able to accept this bulk form without excessive handling. The incremental costs of progressive size reduction eat at the margin of specialized processes and are of particular interest with biomass because regional climate variations can produce very different feeds in a forest, not to mention similar feeds that respond to sizing techniques differently due to their age or relative water content. Can a biorefining process be developed for a particular feedstock that will result in low costs and high recoveries? Absolutely. The issue is when the feedstocks that are evaluated in the design basis of the plant either change or become too expensive; the impact of either development on the process economics is identical in that the cost of goods increases.

Consideration of whole-process design close to the beginning of the experimental stage can yield surprising and potentially advantageous results. While a particular unit operation may initially appear costly and inefficient, the downstream benefit can be shown to offset these perceived drawbacks. This same paradigm can be applied to bulk feedstocks, those that are deemed too variable and unacceptable for one mode of operations can be utilized more readily in a plant designed to accept the greatest variation possible. A small number of experimental test series can be carried out to look at the relationships between processing parameters and give the process engineer hope that managing a number of sensitive unit operations in the plant will yield a profitable, integrated process.

For the business manager of such an enterprise, there is also potential for a partial respite from the hair-raising volatility of energy and fuels markets in supporting the development of a product mix from the biorefinery that can be altered without extensive plant retrofits to suit economic realities. The key lies in a commitment to serve multiple markets with biorefinery products, along with the requisite capability and planning to develop a flexible biorefinery that can modify product output to some extent. This type of design flies in the face of those of us schooled in the science of unit operation optimization, but we should take solace in the knowledge that purposely accepting a low yield in one operation need not impact the overall process efficiency and may be the key to competing in mature energy and chemical markets such as they are.

BIOREFINERY OPERATIONAL PARAMETERS

The biorefinery typically operates according to some variation of the diagram shown in Figure X.1.

Figure X.1 A common biorefinery flowsheet

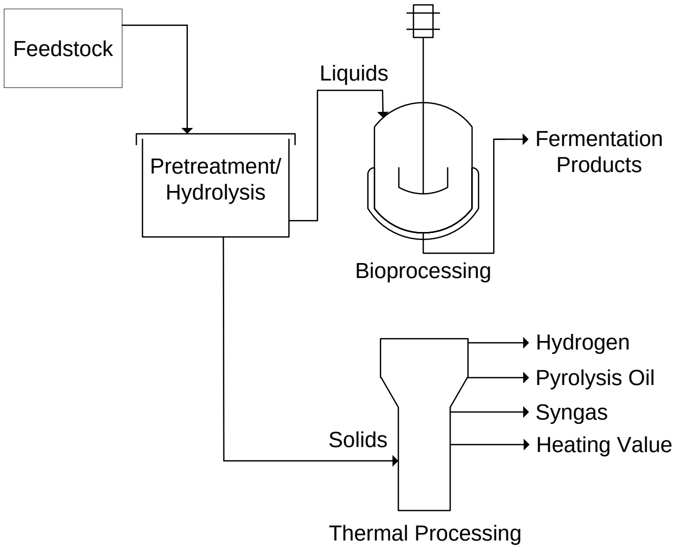

As the variations on this basic flowsheet are many, it behooves one to examine those processes that are shared by all: feedstock supply and pretreatment. In many cases this would be defined as harvesting and sizing the feedstock, while in others it may involve extensive processing in order to supply satisfactory woody materials to the more advanced processes of the biorefinery. In general, if the process being studied is a regional forest-supplied biorefinery, the most important parameters are tonnage fees paid to suppliers and operational expenses at the refinery involved with sizing and chemical treatment of the lignocellulosic materials. Because the chemical makeup of the available feedstocks for biorefineries has been covered in detail elsewhere in this publication, this section will assume the use of softwood for discussion purposes. Though there are slight variations, the basic makeup of the softwoods studied in detail at Hazen Research, Inc. are as shown in figure X.2 (from *Pinus ponderosa* and *P. contorta*):

Figure X.2 The chemical makeup of a generic softwood feedstock

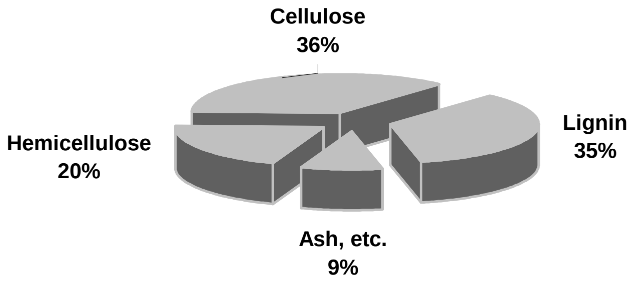

The crucial parameters for the engineer attempting to realize value from these potential sources assume a maximal yield of carbohydrate in the pretreatment/hydrolysis step. Normally a foregone conclusion as a worthy operational goal, this target bears scrutiny.

HYDROLYSIS YIELD IMPACT ON ECONOMIC MODELS

The United States Department of Energy National Renewable Energy Laboratory (NREL) has been developing an integrated biorefinery ASPEN model for general distribution and in hopes of generating interest and providing process knowledge to the biorefining industry in the United States. The model assumes the target product is primarily cellulosic ethanol, but hydrolysis residues contribute to the cost model by way of heating value. This model, in addition to the Laboratory Analytical Procedures (LAP) that are released and updated for the public's use, provide the researcher with some valuable tools in the science of feedstock characterization. Unless otherwise noted, all the analytical results in this chaper as it relates to feedstocks and specific carbohydrates were generated following the NREL LAPs. The most recent distribution of the associated ASPEN model was announced and released in May 2011. It should be noted that the NREL model utilizes both an acid pretreatment and an enzyme hydrolysis; we have focused on the pretreatment here so as to limit the scope of the discussion.

Besides providing a great deal of data surrounding actual yields in biorefining operations, the model provides sufficient information about the cost drivers of the products, so that one can change some of the assumptions made and calculate what the expected impact of the change. Using this type of data manipulation, one can then plot the various outcomes of different yield scenarios. While not entirely linear, the theoretical selling price of ethanol from the model NREL biorefinery increases with a decrease in hemicellulose hydrolysis. This is shown in figure X.2.

Figure X.3 Selling price of ethanol generated utilizing the process modeled by NREL as it is impacted by carbohydrate yields in acid pretreatment

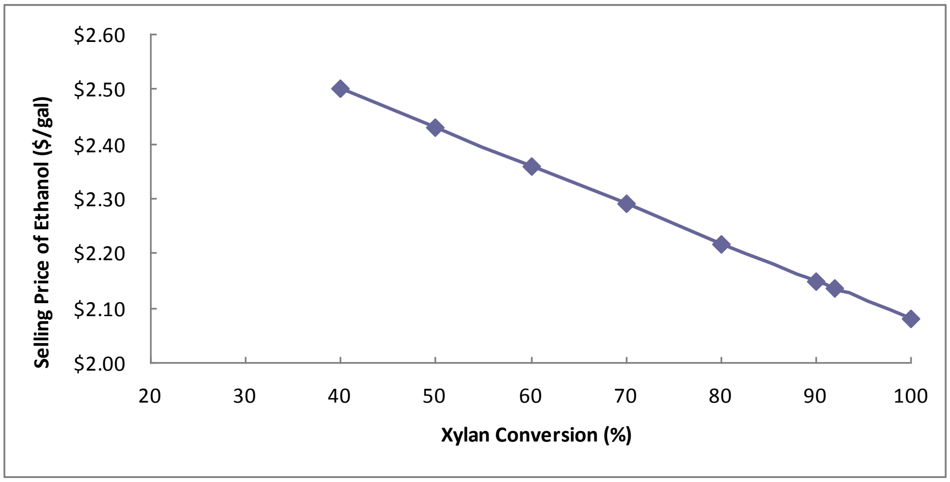

This is a completely predictable result, given the assumptions of the model, which again posit that the residues are contributing to product revenue only in regard to the potential power generation realized by combustion. Given that carbohydrates with their associated oxygen will yield a less energy intensive residue, one would naturally assume that the best use of feedstock is in removing the sugars prior to power generation. It is apparent that increasing acid concentration in pretreatment (to a point) will yield higher carbohydrate recoveries as well as higher heating value residues, as shown in the proximate and ultimate analysis. This is shown below in Table X.1

Table X.1 Biomass residues after acid hydrolysis subjected to proximate and ultimate analysis

First stage temperature (°C)	**175**				**140**			
Acid (wt %)	0.05	0.10	0.20	0.40	0.05	0.10	0.20	0.40
Sugar (% yield)	36.8	54.7	39.7	74.9	37.5	55.3	58.6	67.7
				Proximate				
Volatile	84.2	83.8	81.5	76.2	83.2	84.6	79.8	77.1
Fixed Carbon	15.0	15.5	17.8	23.4	14.4	11.8	18.4	20.9
Btu/lb (HHV)	9023	9149	9295	9542	8913	8972	9157	9484
				Ultimate				
Carbon	52.3	52.1	54.4	56.3	50.9	51.3	53.0	54.2

Hydrogen	5.72	5.62	5.81	5.77	5.61	5.48	5.61	5.42
Oxygen	41.0	41.4	38.0	36.1	40.9	39.4	39.3	37.9
Sulfur	0.05	0.03	0.05	0.12	0.04	0.06	0.09	0.19
F-Factor (DSCF/million Btu)	9091	8861	9369	9505	8931	8960	9123	9004

The combustion exhaust in each case is roughly equivalent when looking at the F-Factors, but in this case the sulfur content of the exhaust (and the associated SO_x) increases with more aggressive application of sulfuric acid. One would assume the same correlation for NO_x and those processes that perform pretreatment with nitric acid. It is clear that looking at the residues as strictly a combustion feedstock, the data steer the process toward complete carbohydrate removal prior to any thermal processing. As such, the sunk costs associated with this removal are assumed to be a given in the cost model of many processes.

These assumptions are sound for the set of conditions specified in most cost models, yet have a fundamental weakness due to the variability of the chemisty of feedstocks. While general assumptions can be made about the makeup of various feeds, the yield assumptions in the models assume a particular chemistry that can be understood and incorporated into process economics, and will remain essentially unchanged. This can be a grievous error, and one that can impact processing costs to a large degree. Just as the agricultural residues available for processes will change year over year depending on many different environmental factors and fertilization regimes, so trees harvested from different regions within the same forest can and do present measurable variation, both within the same harvest year and in year to year comparisons.

The answer to this issue has been to paint feedstock with the widest brush possible, developing processes looking at what are assumed to be the largest deviations from the norm and setting process parameters that can accommodate these. This is in direct conflict with the requirements of the process that need to recover the maximum amount of carbohydrate from these same feedstocks. While some techniques are more robust than others, it is generally understood that pretreatment reagent usage and equipment sizing must be calibrated to a feed that is fairly well understood, unchanged, and available in bulk year after year.

It may well be that this feedstock does not exist. As an example, the Colorado forests of the Rocky Mountains comprise lodgepole pine, ponderosa pine, spruce, limber pine, and Douglas fir (*P. contorta, P. ponderosa, Picea pungens, Pinus flexilis*, *Pseudotsuga menziesii subsp. Glauca,)* as well as many other associated woody trees and shrubs. Bulk harvesting of these forests will include some mix of these, with the primary constituent understood and marketed as such. Therefore a load of “ponderosa pine” mulch will likely contain many different species in some unknown and changing concentration. This disparity of feedstock consistency would of course be greater when harvesting forests still in the Rockies but a thousand miles to the North; one assumes that a bulk harvest from any forested region would present similar diversity. In order to examine what effects this changing mix would have on a biorefinery, Hazen has looked at a number of feedstocks in isolation. While the argument could be made that product effects seen in isolated feeds may not manifest in the same manner in a mixed feed, the discussion must start somewhere.

Our examination included two pines, lodgepole and ponderosa, that are both evergreen American hard pines (Section *Trifoliae*), which are normally harvested for various construction and landscaping purposes on both sides of the North American continental divide. In addition to these two, a subset of lodgepole pine that represents something of a special case was evaluated. In the last few years, large portions of the Rocky Mountain forests in North America have been beset by an infestation of the mountain pine beetle (*Dendroctonus ponderosae)* that has left large swaths of forest literally standing dead. While the causes for this infestation are beyond the scope of the current text, the end result has been that wood harvested from any of these forests will likely contain both desiccated dead wood as well as fresh trees. The appearance of these forests is striking, because the needles of the dead trees turn red and are noticeable even to the casual observer. Given the mood of the population that resides in this region, there has been considerable political and therefore research interest in attempting to make use of this wood. We began our evaluation of biorefinery feedstock using harvests of “beetle kill” trees because the material was readily available. As we engaged in our study, we included the fresh-harvested lodgepole, as well as ponderosa, for comparative purposes. Our initial findings were surprising and became the first few data points in what would evolve into this study of the biorefinery from a product-driven perspective.

Each of the feedstocks was utilized in an as-received state from a lumberyard; higher carbohydrate recovery could certainly have been attained with a finer particle size distribution. It was our intent to examine the possibilities of pretreating/hydrolyzing biomass in bulk harvest to get an idea of how much sugar material could be recovered without optimizing the feed. As such, samples included in our analysis regularly contained rocks, dirt, and other detritus commonly mixed in with large wood piles. This feed was slurried with a dilute acid solution in a titanium pressure vessel, then brought to hydrolysis temperatures utilizing direct steam injection. Figure X.4 below shows the carbohydrate liberation from live harvest lodgepole, dead standing lodgepole (beetle kill), and live harvest ponderosa feed as the acid concentration in pretreatment is increased, all other conditions being equal.

Figure X.4 Total Sugar Recovery from Softwood Feedstocks

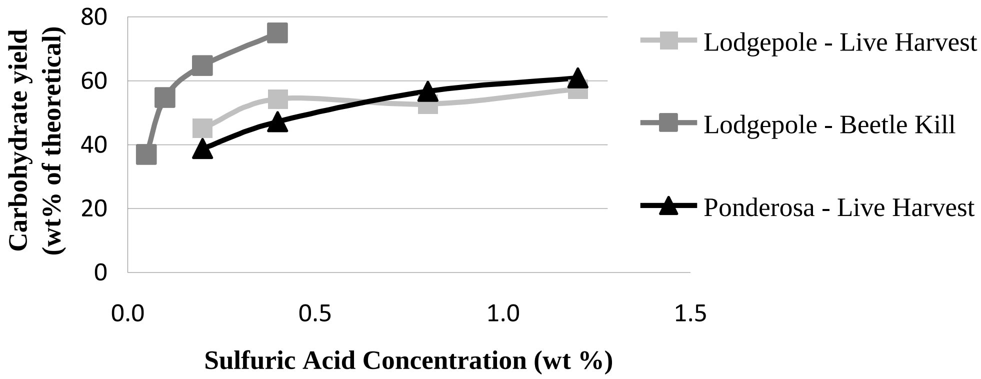

The results were somewhat unexpected, particularly with regard to the beetle kill when compared with the fresh harvest lodgepole. The variables that could explain the observed differences are many, as are the questions that are suggested. Did the beetle kill (standing dead) lodgepole have a greater carbohydrate content because the passive dessication and subsequent loss of any volatile organics? Perhaps this drier feedstock is more friable in a wood chipper, resulting in more fines and a subsequent surface area increase. Is 0.4% acid really more effective on lodgepole than it is on ponderosa? Would it benefit the biorefiner to implement a passive drying of feedstock prior to chipping, and if so, how does this add to the incremental cost of the feedstock? Finally, because the data set was culled from a single year's harvest, will the results be identical as different regions of the forest are harvested in the coming year, or will a year-over-year change in local weather patterns cause the hydrolysis data to show a different trend if the experiments are repeated?

The data are stark in demonstrating the issues, for example – how much time does a bioprocessor have to store biomass before it begins to break down and affect the chemistry of the process? Which conditions should be selected for the commercial plant? How aggressive should the chemistry be, and what equipment (size and materials of construction) will be specified to run the process at scale? These questions have answers, but in general they are qualified and thus not comprehensive. The early goal of the lignocellulosic process to be feedstock-agnostic was discarded along with the processes that would present consistent challenges in order to fit this definition.

BENEFITS OF PRODUCT-DRIVEN OPERATIONAL PARAMETERS

Understanding the point at which many North American biorefinery companies find themselves is aided by a (very) brief summary of recent history. The market factors that have influenced the development of biorefineries in the United States are complex and varied, and they are still sending mixed signals to all interested parties. In the last 12 to18 months, the United States has seen a shift in priorities of biofuel startup companies; the passage of the Energy Independence and Security Act of 2007 encouraged many of these companies to focus on cellulosic ethanol production under the requirements of the Renewable Fuel Standard. The early focus was justifiably set on maximizing the ethanol production out of a biorefinery process, a somewhat challenging task when utilizing a bioagent that respires carbon dioxide as a part of fermentation.

The publication of Policy Research Working Paper 4682 (WPS4682) by the World Bank in 2008 is generally accepted as the first shot in the food-versus-fuel debate, asserting that increasing biofuel production in the United States was a major factor in the worldwide increase in food commodity prices. Despite the subsequent release of WPS5371, which backed away from the initial charge against biofuels, the political winds in the United States had changed, calling into question the continued existence of mandated tax rebates for ethanol producers. While lobbying extensively for their tax-advantaged status, startup companies also surveyed their options, in most cases finding that the potential for alternative products from their processes represented a way to add certainty to their business plans by taking government mandates out of the equation, if only in part. Therefore, a company that had completed its original fundraising as XYZ Biofuels now changed its focus, and changed its name to XYZ Bioproducts, or simply XYZ Technologies, or some similar variation.

The product offerings also changed; where ethanol was the initial primary product with the opportunity to use process side streams for other purposes, chemicals -- particularly replacements for petrochemical-derived intermediates -- became the target. Many of the new process flowsheets included not only a fermentation, but a catalytic reaction that could produce these chemicals. Often the reaction was a hydrogenation, which is not surprising given the aromatic content of most proposed lignin structures.

The production of hydrogen in synthesis gas (syngas) then becomes much more important to the biorefiner than its role as a way of producing process heat. The value of residual carbohydrate in thermal feedstocks then goes up, as the more favorable H/C ratio of carbohydrates (1.4-1.7 as opposed to less than 1 in typical aromatic lignin units) changes the economics of a reduced carbohydrate yield in the pretreatment/hydrolysis stage. As a demonstration of how this can be exploited in the biorefinery, the syngas produced from feedstock subjected to different pretreatment procedures can be evaluated as it is produced by the gasifier.

By way of a brief review, in typical dilute acid hydrolysis experiments the biomass charge is slurried with reagents in a pressure vessel, brought to temperature, and then the contents are flashed into a catch vessel. This allows fairly precise control of the time that the feed is subjected to hydrolysis conditions because the expanding slurry drops in temperature very quickly. A common adjunct to dilute acid hydrolysis is the steam explosion, where the same process is performed without the acid, with the theory that the rapid expansion of water inside the cellular systems in the biomass will rupture recalcitrant structures and facilitate hydrolysis, by reagents or by enzymes.

Figure X.5 shows a plot of the syngas generated from hydrolysis residues that had been subjected to a steam explosion overlain with a plot of those that hadn't.

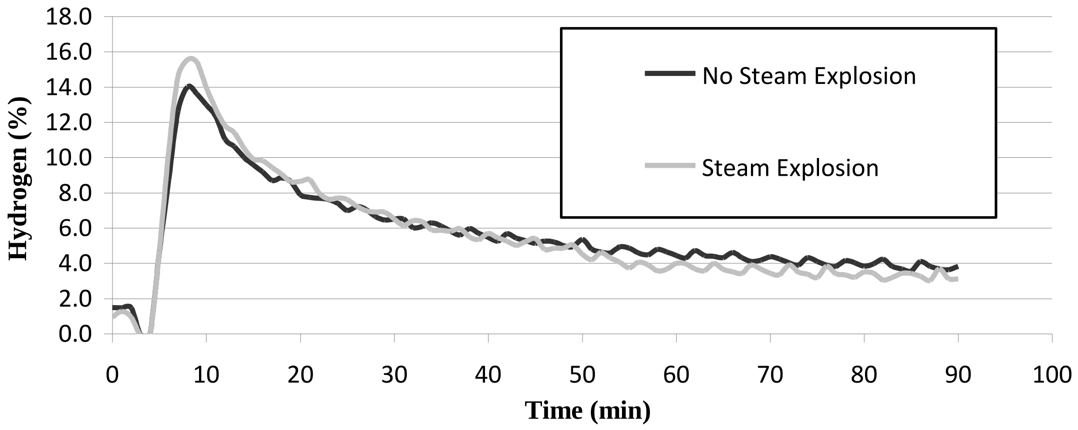

Figure X.5 Hydrogen in Syngas Comparison between Steam Explosion Biomass and Non-steam Explosion Biomass

The difference appears slight; however, the trend was observed in multiple experiments utilizing different biomasses and different hydrolysis conditions While the hydrogen effusion in all these reactions tends to tail off at a similar concentration in the gas (one assumes this is the baseline Gibbs free energy to liberate the hydrogen from the heart of the aromatic lignin), the relative amount of hydrogen can be estimated by the curve maxima as well as the incidence of hydrogen in the early stages of the experiment (essentially the duration). The addition of water in the steam explosion appears to be hydrogenating the residues to some degree, and this effect manifests in a higher hydrogen yield in the gasification. While the additional steam explosion did not appear to affect the carbohydrate yield for the better or for the worse, it may seem a waste of energy and time at first glance. However, the hydrogen yields imparted by the extra pretreatment step likely justify the extra work, provided that gasification is to be performed on the residues generated by the pretreatment.

As with the hydrolysis data, the gasification of residues that have very similar proximate and ultimate analyses yields results that would not be expected from the fuel values. Each of the experiments shown in the Figure X.6 activated the gas analysis equipment when the kiln heating cycle was started.

Figure X.6 Hydrogen Percentage in Syngas from Various Feedstock Hydrolysis Residue

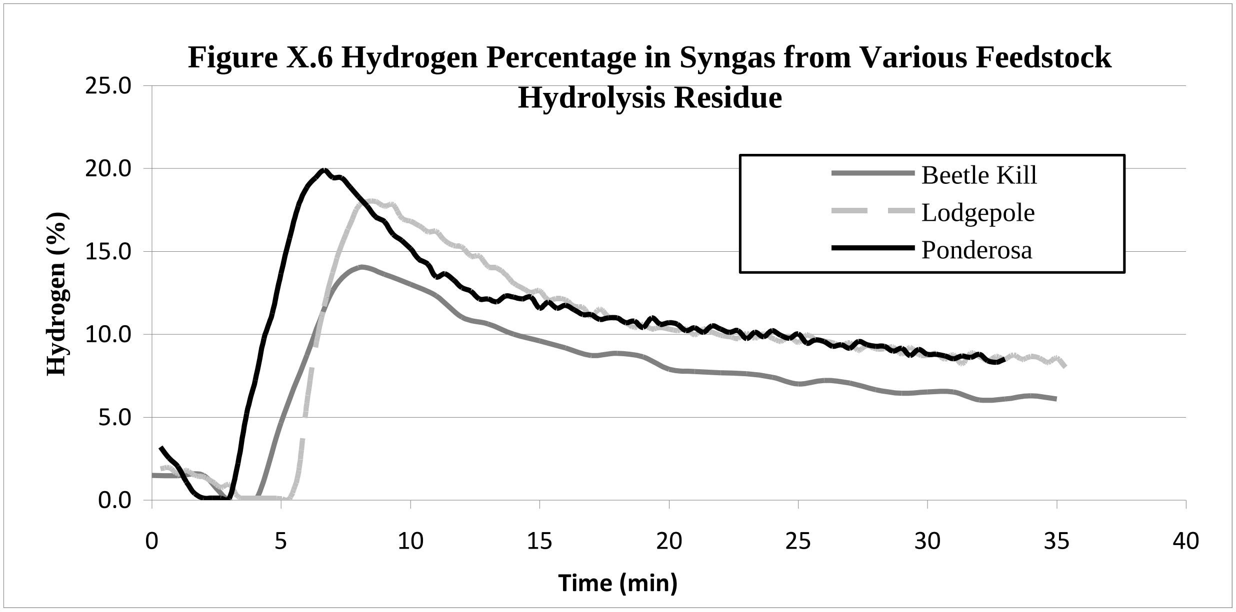

Note the difference in the times that hydrogen concentrations began building. It would appear that ponderosa's structural hydrocarbons are more amenable to gasification and will yield more hydrogen as well. With hydrogen generation in the product mix, one can see the advantage of processing this feedstock when it is available, or at least when it represents the majority of the various woody constituents of an available feed.

VALUE OF RESIDUES

The issue that complicates assigning relative values to the various product streams of a biorefinery is that one is compelled to assume a liquid fuel Btu is equivalent to a boiler fuel Btu. Each energy output from a biorefinery is competing in markets whose variations complicate a comparison to some degree. As a cursory analysis of the liquid fuel market (even in a single geopolitical region) and its contributing factors would likely require a volume as large as this entire textbook, heating value may present a simpler case when evaluating a process stream. In our analysis, it is taken to represent a worst-case valuation of side streams for budgeting purposes. Though an examination of the heating values does not account for all the various factors that may go into a stream's relative worth, it is a sufficient place to begin to build a comparative study.

Table X.2 shows a simple study of the equivalent dollar value of a short ton of feedstock depending on the product ratio from the biorefinery as defined in the NREL models. In this example, it is assumed that the two product options are ethanol and heating value assigned to the resultant pyrolysis oil. Comparing a couple of test cases with a theoretical maximum, and the yield defined by NREL, and assuming that pyrolysis oil could be marketed as a heating oil with a reduced

Btu value, one can see that substantial changes in product output yield revenues at roughly equivalent levels.

Table X.2 Product Revenue Comparison Per Dry Ton of Feedstock

Assumptions	Carbohydrate Yield, %	Ethanol, lb	Pyrolysis Oil, lb	Total Revenues[a], $USD
Theoretical Maximum	100	783	382	269
NREL Model	90	688	549	267
Inefficient Hydrolysis	50	391	1070	262
No Hydrolysis	0	0	1759	256

[a] Assumes $1.79/gallon Ethanol, $1.05/gallon pyrolysis oil, based on proportional Btu content of heating oil.

There is a slight revenue advantage to the conversion of carbohydrates in this model; however, one can imagine that the capital and operational expenditures needed to reach high levels of carbohydrate conversion would likely mitigate this advantage. These data show the diminishing returns in targeting the last bit of yield from a biomass: a $2/st revenue advantage to increasing the carbohydrate yield from 90 to 100%. Knowing what we do about the difficulty of obtaining that last 10%, we are given confidence in the model because it captures the tradeoff very well.

As in the nature of the feedstock itself, it is folly to assume that all pyrolysis oil is equivalent. Take the example of the pyrolysis oil produced from the three feedstocks discussed in the previous section. Examining the resultant oil from these different residues from identical pretreatment experiments, one can see variations that have significant impacts on the process design, Table X.3 shows the fuel analysis of the oil.

Table X.3 Fuel Analysis of Pyrolysis Oil Generated from Hydrolysis Residues

	Single Phase	Heavy Oil		Light Oil	
Residue Type	Beetle Kill	Ponderosa	Lodgepole	Ponderosa	Lodgepole
	Lodgepole	Fresh Harvest	Fresh Harvest	Fresh Harvest	Fresh Harvest
Ultimate (weight %)					
Water	55.6	20.8	14.7	74.7	72.4
Ash	0.31	0.27	0.24	0.02	0.05
Sulfur	0.07	0.32	0.38	0.24	0.13
Carbon	32.1	59.7	60.8	11.1	12.8
Hydrogen	1.75	5.46	6.02	0.71	1.04
Nitrogen	0.02	0.03	<0.01	<0.01	<0.01
Oxygen	10.1	13.4	17.9	13.2	13.6
Proximate (weight %)					
Water	55.6	20.8	14.7	74.7	72.4
Ash	0.31	0.273	0.24	0.02	0.05
Volatile Matter	42.5	72.12	80.2	25.3	27.6
Fixed Carbon	1.58	6.81	4.88	<0.01	<0.01
Calorific Value					
BTU/lb	5238	9904	10147	1100	1934

Note that while the aggregate fuel values are roughly equivalent, there is a phase separation in the product of fresh harvest residue that is not present in the pyrolysed beetle kill. This difference could certainly be exploited in the commercial operation, because the water partitions into the light oil phase in large part, increasing the calorific value of the heavy oil by its subtraction. It is also apparent that whatever volatile organic compounds are lost during the desiccation that occurs after tree death, are likely oxygen-heavy, because the beetle kill lodgepole-derived product produces a more stable oil. An oil that is to be developed for commercial use benefits from the lower oxygen content because the polymerization of the organic acid content would be delayed to some degree. Alternatively, if the oil is to be used as a hog fuel boiler feed, additional oxygen works in favor of the application, imparting a higher heating value. In either application of course, emissions can be an issue and so the F-Factors must be evaluated, and the resultant comprehensive value of the stream could be augmented or discounted appropriately.

One can imagine the impact of the appearance of products with these characteristics in a commercial operation that was not designed to handle a two-phase oil or was designed in expectation of the phase separation that does not occur due to the age of the feedstock. Nearly every measure that is used to put a value on the pyrolysis oil is impacted by the appearance (or lack of) the phase separation. Again, these disparate products are the result of identical experimentation in which the only variable is the condition of the wood at harvest in the field.

In this case, it is clear that a commercial biorefinery utilizing a bulk feedstock that is primarily beetle kill would likely benefit from an increased carbohydrate recovery in pretreatment/hydrolysis because the resultant pyrolysis oil is of lower value due to the water content/lack of phase separation. If process metrics are to require a given carbohydrate recovery with no allowance for process changes based on feedstock, this flexibility would have to come in the form of revenue generated from the pyrolysed residue, which of course in the example would vary wildly with mixtures of feedstocks that produce either of the two products. It would not be unreasonable to look at this revenue picture in some instances and forego pyrolysis entirely, shunting the revenues right to incineration for their heat value alone.

These data represent a couple of series of experiments performed on some relatively pure feedstocks; for the commercial enterprise, a much broader study would obviously be warranted to substantiate the general findings. Given the number of correlated variables discovered in our small study, it could be expected that mixtures of the feeds would introduce even more variability to the product mix. As biorefiners struggle to enter extremely competitive and well-established energy markets with their products, any variability that can be understood and controlled in the commercial setting eliminates a potential barrier to entry.

THERMOCHEMICAL OPTIONS

Utilizing the thermochemical product offerings to drive harvesting and pretreatment process decisions is a legitimate pursuit. Pyrolysis and gasification are the initial process options to consider and thermochemical catalysis, long a staple of fine chemical production from petrochemicals (and upgrading of some of those same chemicals), is currently in development and used in commercial stages in many biorefineries. The product options are vast when one considers the value of carbohydrates, and even their degradation products, in a synthesis operation.

As a brief aside, the majority of the examples presented in this comparison assume that the thermochemical processing will take place after the carbohydrates are removed. Processes do exist whereby feedstocks are gasified as the initial process step in the biorefinery; and the resultant carbon monoxide, carbon dioxide, and hydrogen are used as nutrients in fermentation. One can imagine that feedstock variations would have less of an affect on this type of operation; one also assumes that these processes would have to be carried forward with supplemental hydrogen, which carries its own costs and hazards.

As with pyrolysis oil, syngas has a relative worth based on the chemical makeup, which is more complex than that suggested by the calorific value and which can be influenced by the pretreatment technology. Table X.4 shows the makeup of syngas generated from a 0.4% acid hydrolysis residue from the three feedstocks.

Table X.4 Syngas Constituents from Gasification of Hydrolysis Residue

	Beetle Kill Lodgepole	Fresh Harvest Lodgepole	Fresh Harvest Ponderosa
Hydrogen (%)	14.1	13.1	13
Carbon Monoxide (%)	41.8	43.5	42.7
Carbon Dioxide (%)	16	10.9	14.3
Calorific Value (kJ/kg)	38,869	40,475	38,782

The calorific values of these three gases are essentially the same, as are the hydrogen values and carbon monoxide/carbon dioxide ratios, suggesting that gasification is not affected by the feedstock type to the degree that pyrolysis appears to be. When different pretreatment conditions are introduced to the same biomass, a change is seen in the residue gasification, as can be seen in Table X.5; the hydrolysis experiments that generated these residues were conducted with 1.2% acid, and the samples were collected in Tedlar bags.

Table X.5 Tedlar Bag Sample Analysis of Gasification Products

	Lodgepole		Ponderosa	
Constituent (mole %)	No Steam	Steam Explosion	No Steam	Steam Explosion
Helium	0.0	0.0	0.0	0.0
Hydrogen	10.9	12.0	6.40	15.3
Oxygen	1.32	0.68	9.42	5.67
Carbon monoxide	47.0	42.6	39.6	38.0
Carbon dioxide	9.32	13.31	17.20	13.50
Methane	28.2	27.7	22.7	24.5
Ethane	1.32	1.30	1.33	1.11
Ethylene	0.65	0.74	0.86	0.58
Propane	0.38	0.50	0.67	0.47
N-butane	0.03	0.06	0.08	0.06
C4 Olefins	0.11	0.12	0.16	0.12
N-pentane	0.03	0.09	0.04	0.00
C5 Olefins	0.03	0.03	0.16	0.06
Hexanes	0.73	0.91	1.41	0.64
Calorific Value (kJ/kg)	37,484	38,538	28,034	40,605

It is apparent that the effect of the steam explosion in the pretreatment extends beyond the hydrogen content for these two feedstocks. The trends are intriguing in that the addition of the steam explosion step is beneficial for both types of residue from the standpoint of the calorific value. In the lodgepole experiments, it appears that the higher alkanes are generated at the expense of carbon monoxide, while the ponderosa shows a general reduction in the higher alkane incidence, carbon monoxide is unaffected, and the carbon dioxide is significantly reduced.

This type of data set would suggest a repeat of the experiment, which was performed and yielded similar results. The mole fractions for some of the alkanes are low enough that general conclusions can be made about their presence; the overriding point is that thermochemical processes will produce different products given small changes in the treatment of the feed.

The generation of these products does not always work in favor of the biorefiner. A number of the hydrolysis experiments were conducted using more aggressive chemistry (essentially identical concentrations of acid but higher hydrolysis temperatures in the pressure vessel) in an attempt to yield more carbohydrates. While the carbohydrate yield did increase somewhat, the formation of hydroxymethylfuran and other degradation products also increased. These residues were then gasified, resulting in the data shown in Figure X.7.

Figure X.7 Carbon Monoxide Generation from the Gasification of Residues Subjected to Differing Temperatures in Hydrolysis

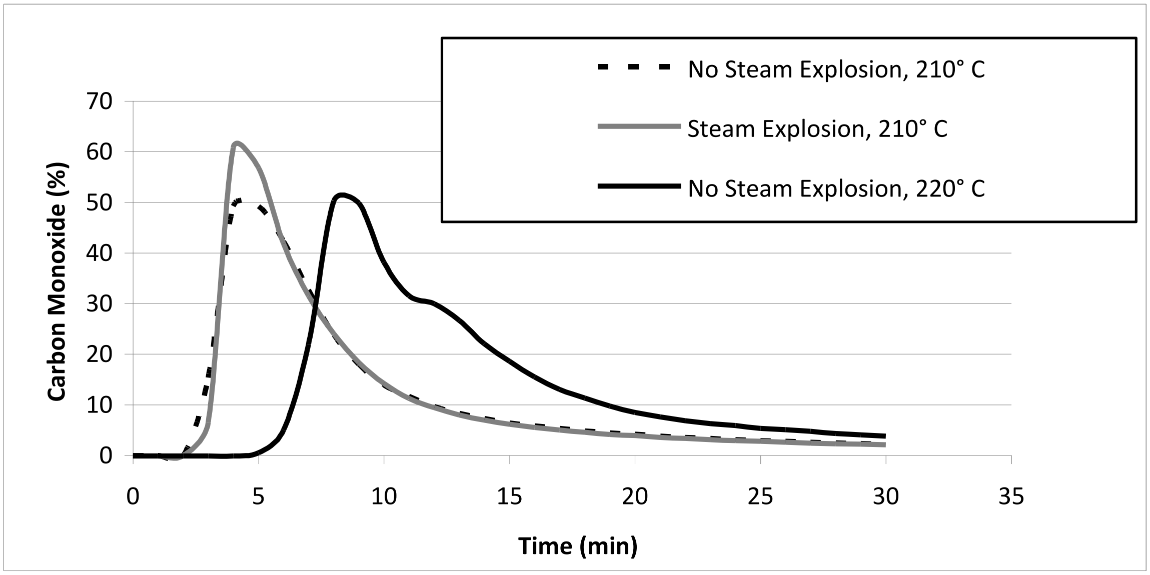

The delayed appearance of carbon monoxide in the syngas was also seen in the plots of data for hydrogen and for carbon dioxide. After verifying that this phenomenon was not an artifact of the experimental technique, it was surmised that the higher temperatures utilized in the hydrolysis began a gasification in the pressure vessel and that the syngas generated was lost in the flash at the end of the experiment. This results in a lower syngas yield, as well as in a shift in the time that syngas generation began. In a commercial operation, this should be avoided because it has the negative effect of degrading more carbohydrates and introducing fermentation inhibitors, as well as reducing yields in gasification. Such a change was seen in all experiments that attempted to utilize the higher temperatures and was consistently observed in plots of all the gasification products that were measured.

INTEGRATED PROCESSING

These data point the process engineer to an integrated biomass refinery where pretreatment can be altered to fit the product demands of the business, which one hopes will be driven by the potential for revenues through the sale of energy producing products. The biorefiner has control over the process design, but he must be cognizant of the need for operational flexibility; this control is largely surrendered once bulk shipments of biomass feedstock begin. The best laid plans will not prevent the situation arriving when material secured from forests differs somewhat in makeup from the feedstock that was used to develop the process. As one of the larger components of the cost of goods, accepting feedstock at preferential price points should always be a goal. In our experience, this results in a heterogeneity that is difficult to predict and therefore must be accommodated to the fullest extent possible.

The biorefiner utilizing the technologies examined in this work has the advantage of building flexibility into the unit operations in order to meet this need. Dilute acid hydrolysis can be run in batch or semibatch fashion, allowing for the addition of a steam explosion step if the overall process calls for it. Likewise, increasing or decreasing acid concentrations can be had without major changes to the installed capacity, assuming the materials of construction are such that the most aggressive conditions are acceptable. Additionally, thermal processing units are often installed with the capability of running in a gasification or a pyrolysis mode depending on the desired product generation. It should also be noted that competing pretreatement technologies that utilize different chemistries would likely see similar benefits with relatively small changes in operations. Just as the conclusions herein regarding dilute acid were drawn from a reasonable set of experiments, the study could be repeated for additional technologies in order to verify the assumption.

The basic flowsheet for a biorefinery capable of the variable processing suggested is less linear and more complex than the basic version presented at the start of the chapter, as seen in figure X.8.

Figure X.8 A basic integrated biorefinery flow sheet

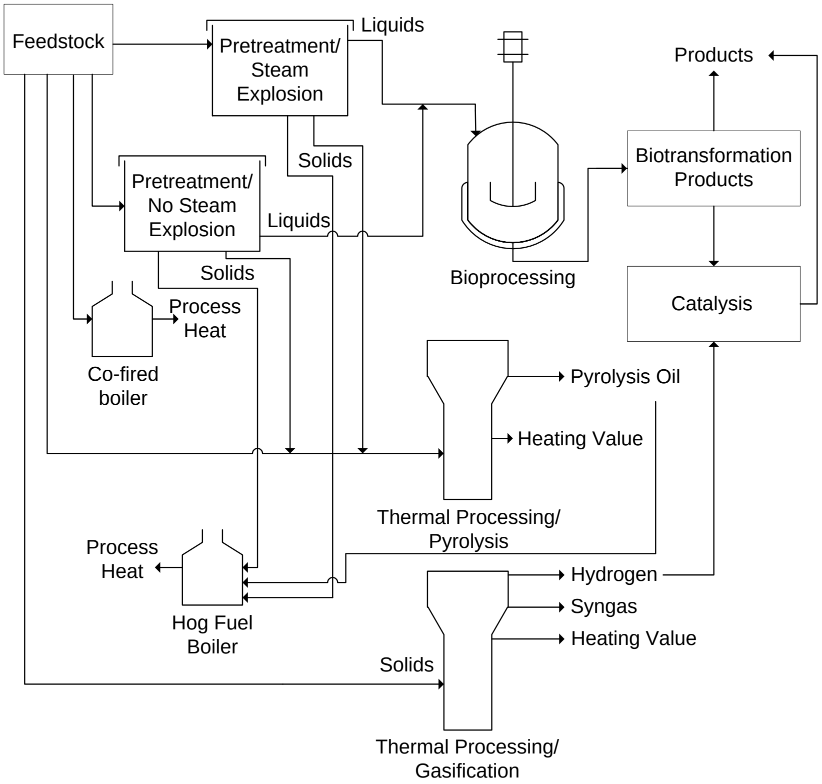

An arrangement of this sort allows a spectrum of options for the feedstocks, from simple co-firing of those that are either of such low quality that further processing is not economically viable to an integrated catalysis that can make use of multiple processing technologies through hydrogenation or other synthetic reactions.

The design of such a process gives the biorefiner an option to divert feed to process heat recovery at just about every step in the process. Alternatively, high grade or easily processed feed can be taken through as many operations as margins allow. In this way, the impact of both feedstock variation as well as market fluctuations can be absorbed by the process to a point; the understanding here is that not running the plant at all is not a viable long-term option. Of course there will certainly be imbalances brought about by the changing mode of operation; our desire is to present a different way of thinking about process development that will avoid too narrow a focus. Often the biotransformation step will utilize proprietary technology that changes the revenue picture and truly forms the backbone of many venture business plans. The goal of the integrated biorefinery model presented is the same, build in flexibility early in order to mitigate risk once you are to market.

CONCLUSION

It is clear from the limited data presented in this chapter that variations in the feed influence the downstream processing in an integrated biorefinery and that the incidence of this variation can take the form of something as simple as the age of the feed. Because others have presented a great deal of information on the heterogeneity of biomass harvested in bulk, we can take it as a given that the integrated biorefinery will have to accept process inefficiencies, unless it is designed with a certain amount of process variability that can mitigate the effect of early processing on the product mix. The studies conducted in support of this writing were able to identify demonstrable correlations early on; one can imagine that a larger research effort on mixed feeds could contribute to this knowledge base, identifying additional cause and effect relationships that could be exploited in the name of stabilizing income from product sales.

The data also present a cautionary tale, in that unit operations that are optimized in development to a maximal recovery may not function this well in commercial practice. The research engineer is then encouraged to weigh the benefit of the effort it takes to squeeze the last bit of product out of early process steps. It may well be that the next year's harvest will not respond to the process in an identical manner, requiring a return to the development cycle. No process designer would want to work on such a process that will continually push the return on research investment back by continually eating up precious development resources.

As the integrated biorefinery is intended (at some level) to replace the traditional petroleum refinery, it stands to reason that a similar market-based manufacturing approach be adopted. Just as the petroleum refiner initially set out to recover kerosene, the biorefiner likely began with the idea that they would be manufacturing fuel ethanol from lignocellulosic sources. The evolution of the modern petroleum refinery has carved the path that the biorefiner can follow, designing truly integrated operations with builtin flexibility to accommodate feedstocks that may or may not be well understood. One should not forget that the petroleum refinery has gone through many iterations in its long history, adapting to not only changing feeds, but changing market demands, changing environmental regulation, and changing geopolitical importance; it only stands to reason that the biorefinery is in for the same treatment. The larger concern is whether the biorefinery will be given sufficient time to adjust to changes in market realities; expecting these new plants to spring up fully integrated and optimized is folly. The researcher and the design engineer know how much effort must go into each unit operation, and are now learning that the race has just started. This same effort must be expended again to ensure that one part of the process is not degrading the performance of another.

A key factor that has not yet been addressed is the upstream push of the market forces of commodity products. The entry of biorefined petroleum replacements into the larger energy and chemical markets will certainly be a lengthy and messy affair. The notion that United States refiners used to dump gasoline in rivers as industrial waste is laughable to the modern engineer. Not quite so amusing is the history of international oil markets with their propensity to encourage collusion amongst multinational corporations (and countries for that matter), politically influenced tariffs and protectionism, and wild boom and bust periods. While some may see the biorefinery as the entity that can end this destructive cycle by returning at least some control over energy production to the local level, it is more likely that biorefiners will find the same market forces in a microcosm as fluctuations will hit regional areas as opposed to the global impact of movements in petroleum manufacturing. It would not be a great stretch to imagine two local energy/chemical producers in conflict over a local biomass supply or inadvertently driving up the cost of agricultural products in their demand for a particular feed. Such effects have only been speculated upon to date, but even hypothetical effects have been enough to influence policy.

A common criticism (often self-directed) of biorefining is that the collective research community is evaluating technologies that were used in earlier phases of industrialization and abandoned for the tempting energy density and availability of petroleum. For all the staggering success in biotechnology and microbiology in the last 50 years, there has been no great breakthrough that suddenly changes the game, allowing the energy/chemical markets the easy choice of product replacement with sustainable alternatives. It may well be that after 150 years of process development and market demands, some of the hardest problems the researcher faces may yield some of the greatest opportunities. As we are forced to deal with the here and now, steadfast design principles must be applied: integrate the process so the energy loss in one portion can be recovered in another, minimize waste, and strive to accept the most diverse feedstock available. Lacking success in those three areas, it is hard to see how a biorefinery can compete in the long term.

Of course advances in technology are not limited to biorefining, and a very real market mover will be the incremental cost of fossil fuels. One would assume that the general consensus among the readers of this book is that the need to find cost-equivalent replacements for fossil sources of energy and chemicals is urgent and real. We would posit that innovation is not limited to any particular branch of science, that there may be a breakthrough in energy production that changes the economic picture permanently, and it may put forward a new favored source: coal, natural gas, petroleum, or nuclear. The point is that one should not assume that expected market phenomena will gradually remove fossil fuels from the picture. Rather, it is best to assume that biorefinery products will be expected to compete on a level field with their intrinsic benefits laid bare. The investment in the research that will allow this has just begun.