科学研究費補助金 特定領域研究

新世代の計算限界 — その解明と打破 —

平成17年度成果報告書

平成18年4月

領域代表者 岩間 一雄京都大学 大学院 情報学研究科 教授

目 次

1. 総括班活動報告		-
総括班活動報告	 	. :
教科書シリーズ	 	
電子ジャーナル	 	. !
ニュースレター	 	. (
ミニプロジェクト....................................	 	• '
電子情報通信学会和文論文誌 Α 「 新世代の計算限界 招待解説論文特集号 」.....	 	. 8
電子情報通信学会英文論文誌 D 「 新世代の計算限界 招待論文特集号 」	 	. 9
国際ワークショップ iETA の企画	 	. 10
Randomness and Computation	 	. 1
NHC 国際会議	 	. 13
第18回 回路とシステム軽井沢ワークショップ 特別講演セッション	 	. 10
研究集会(モデル班)....................................	 	. 1'
ミニ研究集会 (組合せゲーム・パズル)	 	. 18
ミニ研究集会 (複雑ネットワーク, ウェブグラフ)	 	. 20
ミニ研究集会 (列挙合宿)	 	. 2
ミニ研究集会 (Complexity)	 	. 22
ミニ研究集会 (暗号理論)	 	. 24
ミニ研究集会 (ジオメトリ)	 	. 20
電子情報通信学会 COMP-NHC 学生シンポジウム................	 	. 2'
第19回 回路とシステム軽井沢ワークショップ 講演企画	 	. 28
2. 研究課題別活動報告		29
A01: オンライン予測の手法を用いた意思決定モデルに関する研究	 	. 3
A02: 計算理論的設計による知識抽出モデルに関する研究	 	. 4
A03: 論理関数表現のモデルとシンボリックアルゴリズム	 	. 54
A04: 組み合わせ最適化における指数サイズ・多項式時間近傍の設計	 	. 6
A05: ハードウェアアルゴリズムの性能評価に関する研究	 	. 7
A06: 幾何的配置問題の計算量限界と近似可能性の研究	 	. 78
A07: 現実的な状況下での量子計算の能力に関する研究	 	. 8
A08: 動的な構造をもつネットワーク上の資源割当て問題の研究	 	. 103
A09: 非明示的表現に対するアルゴリズムの開発	 	. 114
A10: 変移する要素間の関係を条件とする組合せ最適化モデル	 	. 12
B01: 代数的および確率的手法による離散構造の限界の究明	 	. 13
B02: 暗号解析手法の計算量理論による改良とそれに基づく暗号方式	 	. 14
B03: 量子論理回路の最適化に関する研究	 	. 15'
B04: 回路計算量の下限の研究とその応用	 	. 16

	B05:	ブール理論に基づく離散システムの構造解析と計算限界の研究
	B06:	暗号システムに対する実装攻撃の適用と限界に関する計算的研究
	C01:	グラフ描画アルゴリズムとその Web 情報検索への応用
	C02:	連続と離散の統合によるロバストアルゴリズム構築
	C03:	ネットワーク上での社会的効用と個人的効用の対立問題に対するアルゴリズム的研究 222
	C04:	実践的な列挙アルゴリズムの理論構築
	C05:	幾つかの画像関連問題の計算複雑度の解析と効率的な解決法の提案
	C06:	グラフ構造を有する問題に対する近似アルゴリズムの設計
	C07:	Web コンテンツ活用に関連した離散最適化問題の研究 273
	C08:	高性能近似アルゴリズムの設計法に関する研究
	C09:	ネットワーク問題のモデル化とアルゴリズムの研究
	C10:	情報基盤アルゴリズムとしてのメタヒューリスティクスの研究
	C11:	大量データ処理のための領域効率の良いアルゴリズム
3. ⁻	各種資	341
	<u></u>	ースレター 第3号
	<u></u>	ースレター 第4号
	平成	17 年度 第1回 全体会議 講演資料

平成 17 年度 第 2 回 全体会議 講演資料	381
NHC Spring School and Workshop on Discrete Algorithms 講演資料	413

3. 各種資料

特定領域研究

「新世代の計算限界」ニュースレター 第3号 2005/7/11

はじめに

このニュースレターは,特定領域・新世代の計算限界のメンバーの情報交換と交流を目的とした 情報発信誌です.毎回,いくつかの研究関連の記事と,特定領域のスケジュール・活動報告と,各 研究者の活動予定などをお送りいたします.今回は,東北大学徳山豪先生の経験談と,6月に行 われた全体会議,幹事会の報告を行います.

- 1.「群れ派」の研究の仕方 徳山 豪 (東北大学)
- 第2回全体会議を振り返って 浅野 哲夫 (JAIST)
- 3. 2005年度第1回幹事会議事録 事務局
- 4. イベントカレンダー + 事務連絡
- 5. このニュースレターについて
 - 「群れ派」の研究の仕方 徳山 豪 (東北大学)

最近は「引きこもり」や「ネットが友達」など一人でいるのが居るのがいいという方も結構いるらしい が、私は典型的な「群れ派」である。特定研究のおかげでいろいろな人の考えが聞けて、おしゃべ りが出来るのは楽しくて仕方が無い。そういう場では少々躁状態になるので、声がでかくて「うざっ たい」(特に岩間さんや西関先生と一緒だと)が、暖かくお許しいただけるとありがたい。大体が教 授室に1時間とじっとしていられないで学生研究室をうろうろしているので、学生には(だれも口に 出して言わないが)うっとうしい存在である。そんな性格なので、一人でコツコツ本や論文を読んで 勉強するのは苦手で、常に人との関りが幸運をもたらしている。宇野さんのリクエストもあり、そん な話(記憶が薄れているので、意図せずフィクションが入ったらごめんなさい)をしてみる。

食事とドライブ

楽しいおしゃべりの相手というと、浅野哲夫先生(二人の取り決めにより、「浅野さん」と書く)がなんといっても第一である。1988年にはじめての海外出張でイリノイ大学でのACMのSCG(計算幾何 学シンポジウム)に行ったのがお付き合いのはじまり、日本でもお会いしたことはあったのだが、 当時私は純粋数学(代数的群論)から計算機科学に転向して1年経たないほどの新米で、大先生 相手に気後れして直接お話をしたことがなかった。

イリノイ大学のあるアーバナ・シャンペーンというところは平らな大地以外に何にも無いところで,し かも会議参加者は見知らぬ人ばかり.大体が食事をする場所も近くにない.うろうろしていると,う しろから「こんにちは」の声がして,黒いサングラスを掛けたヒゲの先生がニコニコしている.「タご 飯, どうします? みんなと一緒に行きません? 僕, レンタカーしてるから. 」うれしかったなあ. なんといっても人間は食べることが一番. おまけに,後でIBMのWatson研究所でホストをしてくれ たAlok Aggarwalや,研究での転換のきっかけをくれたMicha Sharirら有名人たくさんに紹介しても らった. 結局会議期間中ずっと浅野さんの車に同乗したのだが,気に入ってもらえたらしく,「せっ かくだから,一緒に研究しましょう. 2次元のハッシュ問題だけど,興味ある?」この研究は後で 「トポロジカルウォーク」というアイデアを生み,私の出世作(1991SCG)になったのだが,浅野さん との出会いが無ければ今の半分も研究人生が楽しくなかっただろうと思う.

話のついでだが、この会議で、「ちょっといいか」といって、話しかけてきた二十歳くらいのパデュー 大学の中国人学生が居た.テーブルで無理やり向かい合わせたかと思うと、いきなり「計算幾何 学の並列アルゴリズムに将来性はあると思うか、返答は如何に」と質問してくる.なんだなんだと 思いながら適当に答えていると、「Thank you、参考になった」といって握手してバイバイである.こ の生意気な学生はDanny Chen.当時はまだ論文も無い学生だったが、今ではトップクラスの研究 者で、貴重な共同研究者である.

講演依頼は喜んで受けよう

STOC-FOCS-SODAとよく言うが、こういった会議に最初に論文を通すのが難しい.私の場合は、 山下雅史先生と渡辺治先生のおかげである.といっても共同研究ではなく、駆け出しの研究者だった私に、「ランダムアルゴリズムについての一時間講演をお願い.絶対いい講演ができるから.」 とおだてながら迫ったのがこのお二人(詳細は忘れたが、IEICE関係で、聴衆はずいぶんたくさん 居た). 当時ランダムアルゴリズムが世界的にブレイクし、反面日本ではあまり研究者が居なかったことは事実だが、ともかく自分でランダムアルゴリズムを設計したことが無い人間が「良い講 演」をするのは難しいし、オリジナルな結果がまったく無いチュートリアルをするのは恥ずかしいと 思った.講演期日までの間、柄にもなく一生懸命にさせられた結果が、割り当てや輸送問題を計 算幾何学のランダムアルゴリズム手法で解く(問題自体は研究中だったが、ランダムというのは考 えても居なかった)というもの.幸運なことにいい問題に当たった(1991SCG,1992SODA)が、講演 を断っていたら運はめぐってこなかったろう.

人の話に首を突っ込むと

誰かが面白そうな(研究の)話をしているときに,話に加わりたいが邪魔をすると悪いなと思うかも しれない. IBMの東京基礎研究所で入社4年目くらいの頃,当時の上司の岩野さんと加藤直樹先 生(これも加藤さんと呼ぶ)は頻繁に共同研究をしていた.ちょうど私の研究も軌道に乗って,翌年 からアメリカ赴任という状況だったが,加藤さんが岩野さんを訪れて来た.ご挨拶でもしようと二人 の話している会議室に(呼ばれもしないのに)行くと,なにやら面白げな話をしている.ぐずぐずし ている間に,「徳山さんは専門家だから,聞いてよ」というのでありがたく参加すると,幾何学的な 動的最小木とパラメトリックなマトロイドについての加藤さんの知識と発見の話題である. O(n⁴) がマトロイドの理論で O(n³)になるらしい. その場で考えてみると,計算幾何学のテクニックでも O(n³)になるし,両方を組み合わせると O(n²)にも出来るようである. わずか1時間ほどの Discussionで出来た論文が,「動く点集合の最小木について」(1992FOCS). さらに,関連するテ ーマで,たくさん(1995SODA,1999SODA,1999FOCS,2001SCG)論文を書いた. 結論として,人の話 には遠慮なく首を突っ込むべきである.

朝目覚めると

「ワーム食べませんか?」 玉木さんと最初に会ったときに彼が言った言葉である. 1992のFOCS で, ワームホールラウティングの講演をした直後, 講演に使った(したがってOHP のガラスの上を のたくりまわった)グミキャンディーのミミズのおもちゃを半分にちぎって, 片方を銜えながら. こん な面白い人がIBMの研究所にやってきて、2年間同室だったのだから、いい研究ができないわけが ない.まず手始めにということで、Micha Sharirと学会のlunchで一緒したときに「解くべきなのだ が、まったく手掛りがない」といわれて挑戦し続けていた「放物線のレベル問題」について玉木さん に話すと、Lovaszのマッチングとカバーリングの双対定理という道具を使って、あららといううちに 解けてしまったのはおどろいた(1995SCG).私一人では3年間うまくいかなかったのに、また或る 時、国際会議で同室して、時差のせいで朝(?)3時に目覚めると、玉木さんなにやら思索にふけっ ている.「目が覚めましたか、ちょっと一緒に考えませんか?」、こちらは目覚めただけで起きたわ けではないのだが、一緒につきあう、これで出来たのが1997年のISAACで最優秀論文賞をもらっ た研究(なんだか、私はサンチョ・パンザみたい).

さて、玉木さんと私ではタイプはだいぶ違う、共通の趣味である囲碁を打つと良くわかるのだが、 私は直感派で諦めが良く、玉木さんは論理(読み)を積み重ねないと納得しない、囲碁は個人ゲー ムだが、研究は協力して出来るので、違うタイプが重なると相乗効果で大変よろしい、考えてみる と、「浅野ー加藤-玉木-徳山」は4人ともタイプはそれぞれ違うように思う(どう違うかは、皆さん で判断してください)、浅野さんがリーダーシップを取って、この4人で合宿をしたりして、随分楽しく 共同研究をしてきたものである。

研究のエピソードは書き出すと尽きないのだが,残りはまたの機会にしようと思う.

第2回全体会議を振り返って 浅野 哲夫 (JAIST)

去る6月16日から2日間に亘って開催された第2回の全体会議について報告します. 昨年, 特定 研究がスタートして以来, 平成16年10月に東北大学で開かれた第1回の全体会議, 今年の2月 に京都で開かれた国際会議に引き続いての会合ですが, 最初の2回についてはまだ成果を出せ る状況にないということもあって,招待講演を主とする構成になっていました.代表幹事の伊藤先 生から、東京電機大学の築地先生と一緒に次回の全体会議を計画するように言われたとき、どの ようなコンセプトにするか色々と考えました.まず考えたことは、特定研究にとって最も重要なこと は何かでした、とにかく偉大な成果を出すことが要求されている訳ですが、特定研究としての成果 というのは、単独では難しかった問題を特定研究のグループの力を借りて初めて解くことに成功し たというものであることが望ましい訳です.つまり,うまくコラボレーションができたかどうかが問わ れる訳です.もちろん.理論のことですから.何度も集まったから必ず成果が出るという保証は何 もありませんが、社会に対しては、少なくともコラボレーションがうまく行くように最大限の努力はし たということが見えなければならないと考えた訳です. では. そのためにはどうすればいいかです が、コラボレーションを成功させるためには、まず我々のグループの中でどんな成果があるのか、 誰がどんな研究を行っているのかを知ることが肝腎です.ということで.まずは全ての研究班に自 分たちの研究を紹介してもらおうという考えに至りました. もちろん. 全員に話をしてもらうのが最 適ですが、そうすると時間が2日では全く足りないので、ポスターの形式にしました、コラボレーショ ンの切っ掛けは、日常会話に近い、ちょっとした会話の中から生まれるものだと個人的に信じてい るので(とにかく、興味ある共通の話題がないことにはやる気が起きないので)、とにかくインフォ ーマルな会話の時間を大切にしたかったのです.

もう一つの視点は,私がこのコミュニティーが好きなのは,日本の社会にしては珍しく年齢の壁が高くないことです.理論をやっていると,どうしても若い人の集中力には負けるから,経験が豊富だというだけで威張ることができません.これが年齢を余り感じることなく誰とでも話せる雰囲気を醸し出しているのだと思います.いつの間にか自分も年寄りの部類に入ってしまっているのでしょう

が,気分はまだ20代です(言い過ぎです,ごめんなさい).この雰囲気を大事にしたくて,偉い人に だけ講演をお願いするというスタイルを取りませんでした.

最後に、4人の方にだけ普段より長めの時間を用意して講演をしてもらいました. いずれも活躍が 伝わってくる講演で、改めて本グループの質の高さを実感しました. 今回の4人の講演者は私の ー存で決めましたが、4人のご講演を聞いただけでも本グループの全体像が浮かび上がる、うま い人選になっていたのではないかと自画自賛しています. もちろん、講演を快諾して頂いた4人の 講師の講演の質の高さが成功に導いたことは言うまでもありませんが. 次回はどんな趣向で全体 会議が企画されるか今から楽しみです.

2005年度第1回幹事会議事録 事務局

日時:2005.6.16 12:00--13:50

場所:国立情報学研究所12階会議室

[報告事項]

○ 教科書シリーズの進行状況(杉原,山下,室田,渡辺)

(仮題)アルゴリズムサイエンスシリーズ 全16巻 A5版 約200ページ

- 電子ジャーナル (徳山)
 - ・ECCCの成功(ドイツの大学)を見習ってECCCのアルゴリズム版を作る?
 - ・維持が大変(お金と人手)
 - 本屋を巻き込めば継続可能だが
 - -> 全体討論で議論
- 成果報告書 (堀山)
 - ・印刷体のものを各研究代表者に送る
- 確率シンポジウム(徳山)
 - ・資金は基本的に折半
- ニューズレター (宇野・牧野)
 - •7月出版を目安に準備

[検討事項]

○NHC国際会議に続いて、今年度も国際会議を開催するか?

・若手の育成を重視し、サマースクール(またはスプリングスクール)を実施

・前半スクール後半チュートリアル?

一 全体討論で議論

- ○次回の全体会議の場所, 担当者
 - •名古屋大,平田先生
 - •11/21,22予定
- ○林先生のコメント
 - ・審査が厳しくなっている.
 - きちんと評価ー>反映 という傾向.
 - 中間審査の結果で減額されたり、最終報告の結果で次期プロジェクトに

反映されたりするようになってくる.

- この特定は情報基盤を支える重要なものである。
- ・現実的な解法,具体的な成果が求められる.
- ・領域の連携が重要:単なる集まりでは無い(連携することでどんな結果が得られたか?)

目標に対してどこまで到達したか?)

- 「情報学」が終わるのでこの特定が情報系の代表格になる
- ・グランドチャレンジを期待する.

「新世代の計算限界」イベントカレンダー + 事務連絡

特定領域研究も2年目に入り,段々と盛り上がってきて事務局も嬉しい忙しさです.春学校,ミニプ ロジェクトなど新しい試みも始まります.皆様,これからもご協力の程よろしくお願いします.

★ 報告書の作成について皆様のご協力有難うございました. 作成方法やフォーマットなどについてご意見をお持ちでしたら, 伊藤, 堀山までお寄せ下さい.

★ ミニ研究集会はどなたがどんな思いつきで企画されても結構ですから, どしどし行って下さい. 総括班から費用のサポートもいたします. **全体会議**(予定) 11/21(月),22(火)名古屋大学

列举合宿(宇野毅明·中野眞一·上原隆平)

3日間, 泊り込みで列挙・数え上げを中心に組み合わせアルゴリズムに関する雑談をします.

9/27(火)13:00 - 29(木)11:30 群馬県伊香保町 群馬大学セミナーハウス

7/11(月)-15(金) ICALP, ポルトガル・リスボン

7/11(月)-15(金) IFORS, ハワイ(参加者:伊藤大雄, 巳波)

7/18(月)-21(木) <u>ランダムネスと計算に関わるワークショップ</u> 仙台国際センター (参加者: 元木光 雄)

7/19(火)-22(金) Combinatorial Pattern Matching 韓国済州島

- 7/25(月)-27(水) LAシンポジウム 福岡県宗像市
- 8/15(月)-17(水) WADS, カナダ・ウォータールー
- 8/16(火)-29(金) COCOON, 中国·昆明
- 8/29(月)-9/2(金) MFCS, ポーランド・グダニスス
- 10/1(土)-5(水) CP2005 スペイン・バルセロナ (参加者:元木光雄)
- 10/3(月)-6(木) ESA, スペイン・イビザ
- 10/21(金)-10/22(土) 第17回RAMPシンポジウム 青森県弘前市 シティ弘前ホテル
- 10/23(日)-25(火) FOCS 2005 アメリカ・ピッツバーグ
- 12/7(水)-12/9(金) <u>I-SPAN2005</u> アメリカ・ラスベガス
- 12/19(月)-21(水) ISAAC, 中国·海南島
- 1/22(日)-1/24(火) <u>SODA2006</u> アメリカ・マイアミ

このニュースレターについて

ニュースレター各号は電子メールで配布する予定です.短い記事や連絡事項は全て掲載します が,長い記事,イベントの詳細などはwebページに掲載する予定です.webページには詳細まで全 てを載せた完全版を掲載して,目次,あるいは各記事の末尾のURLを参照すると,web版の同じ 記事を参照できるようにいたします.

記事は、各回、1つの研究課題に担当をお願いする予定です。各研究課題で2000-4000字程度、研究に関わる記事を書いていただければと思います。通常、このようなニュースレターでは、 研究成果を報告するのが一般的だと思われますが、この特定領域では「研究者の交流」に焦点を 当てたいため、「研究の成果以外」の記事を面白く解説していただければと思います。例えば、最 近参加した国際会議の情報を、どのようなものが流行っていたか、何が面白かったか、などの主 観的な解説を交えて報告したり、最近考えている問題、あるいはオープン問題を、この辺までは解 けるがここがうまくいかない、といった解説を交えて紹介する、という形です。 また,研究者間の交流を促進するため,各研究者の,国内外の会議への出席予定を集約して 掲載していこうと考えています.研究者の交流には,顔をあわせる回数を増やすことが肝要です. 他の研究者の参加予定がわかれば,会議への出席のモチベーションを高めることにもつながり, それがディスカッションや研究成果を生むきっかけにもなるでしょう.特定領域メンバーの皆さんに は,自分のわかる範囲で,国内外の会議・研究会の情報と,自分の参加予定を教えていただけれ ばと思います.

この他, 個人からの寄稿を募集いたします. 100-1000字程度で, 情報宣伝されたいことを自由な 形式で書いて送っていただければ, 掲載いたします. メールで配布する関係上, テキスト形式のも のしか扱えませんが, そこはご了解お願いいたします.

次回は10月ごろを予定しています.

★ ニュースレター編集委員では、皆様からのご意見をお待ちしております. 編集方針や内容の追加など編集全体にかかわることから細かいことまで、幅広いご意見をお願いいたします.

■■ 新世代の計算限界 ニュースレター ■■

編集委員長 宇野 毅明 uno@nii.jp (問合せ先)

副編集委員長 牧野 和久 makino@sflab.sys.es.osaka-u.ac.jp

特定領域研究

「新世代の計算限界」ニュースレター 第4号 2006/1/27

はじめに

このニュースレターは,特定領域・新世代の計算限界のメンバーの情報交換と交流を目的とした情報発信誌 です.毎回,いくつかの研究関連の記事と,特定領域のスケジュール・活動報告と,各研究者の活動予定など をお送りいたします.今回は,名古屋大学高木直史先生の研究科紹介と,11月に行われた全体会議,幹事会 の報告を行います.

- 中心?隙間?それとも辺境? 高木 直史(名古屋大学)
- <u>第3回全体会議の報告</u> 平田富夫(名古屋大学)
- 3. 2005年度第2回幹事会議事録 事務局
- 4. イベントカレンダー + 事務連絡
- 5. このニュースレターについて
 - 中心?隙間?それとも辺境? 高木 直史(名古屋大学)

■ 集積システム論

主要大学の中では少し遅れをとりましたが、名古屋大学でも2003年4月に情報系の大学院研究科として、 情報科学研究科が設立されました。計算機数理科学専攻、情報システム学専攻、メディア科学専攻、複雑系 科学専攻、社会システム情報学専攻の5つの専攻からなっています。本特定領域研究には、計算機数理科 学専攻の平田先生、柳浦先生(昨年10月に転入)も参加されています。(藤戸先生は豊橋技術科学大学に転 出されました。)私と高木一義先生は情報システム学専攻の集積システム論講座(大講座:3研究室)に所属 しています。

集積システム論というと、半導体回路工学等をイメージされるかもしれませんが、集積回路で構成されるシス テムが対象であり、いわゆるシステムLSIや組込みシステムのハードウェア(アーキテクチャ)やソフトウェアに 関する研究を行う学問と考えています。現在、年間100億個以上のマイクロプロセッサ(4ビット以上)が出荷さ れていますが、そのうち、パソコンやワークステーション等の汎用のコンピュータに用いられるのは2%程度で あり、ほとんどが組込み用に使われています。実際、あらゆる家電製品にマイクロプロセッサが組み込まれて おり、一般家庭で40~50個はあるといわれています。また、自動車にも多く使われており、高級車では70~80 個使われています。

以前は、集積回路を部品として汎用コンピュータを作るために、汎用機のアーキテクチャ、汎用機のOS、汎用 機のソフトウェア、そして部品としてのIC設計の研究が行われていました。しかし、今は、コンピュータ(プロセ ッサ)は部品となり、メモリや種々の周辺回路とともにボード上、さらには、一つのLSI上に集積されています。 集積(組込み)システムのアーキテクチャ、組込み用プロセッサのアーキテクチャ、組込みOS、組込みソフトウェア、そしてシステムとしてのLSI設計の研究が必要です。

私は、集積システム論は、電気・情報分野の「中心」に位置すると思い(たく)、実際、産業界はシステムLSIや 組込みソフトウェアの高度技術者を必要としているのですが、残念ながら、我国の大学では電気電子の分野 と情報の分野の「隙間」になっています。企業では、半導体技術者等を社内再教育でこの分野に振り向けて いるのが現状です。

■ 集積システムの設計支援とハードウェアアルゴリズム

これまでは、本特定領域研究の皆様にはピンと来ない話だったかと思います。しかし、実は、集積システムの 設計過程は、組合せ最適化問題の宝庫です。処理のどの部分をハードウェアで行いどの部分をソフトウェア で行うか等を決めるシステムレベルの設計から、LSIのレイアウトまで、最適化問題を解く作業の連続です。最 適化の対象は、計算時間や面積に加え、消費電力が重要になっています。ちょっと意外かもしれませんが、よ く現れる処理は、ソフトウェアで行うよりハードウェアで行った方が消費電力が少なくて済みます。製造手順の 最適化もあります。少し前の話になりますが、IBMの岩野さんらが、プリントボードに多数のホールを開ける作 業を、巡回セールスマン問題の解法を応用して効率化されたことを覚えておられる方もおられると思います。

本特定領域研究には、「ハードウェアアルゴリズムの評価に関する研究」という課題で参加させて頂いていま す。ハードウェアアルゴリズムの評価は、それの基づく回路の性能によって行います。そのためには、回路の 安定したモデルが必要です。従来、組合せ回路モデル上での段数および素子数、さらには、配線領域も考慮 したVLSIモデル上での面積が評価基準として用いられています。段数は計算時間の指標ですが、現実には、 論理ゲートでの遅延よりも配線での遅延の方が大きくなっており、配線遅延を含めた計算時間を評価できる 回路モデルと評価基準の確立が必要だと考えています。また、消費電力を評価するための回路モデルと評価 基準も必要だと考えています。

近年、LSIの配線層数が増えてきており、さらに、先日、東北大学でLSIチップを貼り合わせた3次元LSIが開発されました。3次元のVLSIモデルも興味深い研究課題ではないかと思います。

我々の研究は、本特定領域研究では「辺境」に位置しますが、皆さんにとっても、いろいろ面白い課題がある と思います。

第3回全体会議の報告 平田富夫(名古屋大学)

昨年の11月21日、22日に開かれました第3回全体会議についてご報告いたします。

これまで科研の集会というと研究発表のみというスタイルが多かったと思いますが、前回6月の第2回全体会 議では、メンバーのコラボレーションを促進することを主目的として各研究班を紹介するポスターセッションが もたれました。(そのコンセプトは前回の全体会議企画担当の浅野哲夫先生の記事(ニュースレター第3号)を ご覧ください。)

第3回の全体会議もこのコラボレーション路線を引き継ぎました。今回はオープン問題を持ち寄り共同でチャレンジすることをメインにしましたが、この企画は、第2回全体会議の直後に浅野哲夫先生と岩間先生にご相談して決めたものです。皆さんがお持ちのオープン問題で全体会議に話題提供できるものを募集したところ5人の先生方が応じてくださいました。この場を借りてこの5人の先生方にお礼申し上げます。私の印象としては、だれもがすぐ理解でき、しかもアイデアが出せそうな問題で非常によい問題だったと思います。中には他のメンバーの結果と結びついて発展した話題もあったようです。このような試みは短期的な成果を求めるのが目的ではありませんが、これを機会にひとつでも成果に結びつくものが出てくれば今回の全体会議は大成功

といえます。

2日目の講演はSODA、FOCSで発表を予定しているメンバーを中心に、私と柳浦先生が興味をもっている話 (は皆さんも興味があるだろうと勝手に考えて)で人選させてもらいました。内容はレベルが高く楽しめたので はないかと思います。ご講演いただいた5人の先生方にこの場を借りてお礼申し上げます。

場所は初日が野依記念学術交流館で、2日目がIB電子情報館大講義室でした。野依記念学術交流館は野 依良治先生のノーベル賞受賞を記念して建てられた瀟洒な建物です。カンファレンスホール(200席)と会議 室8室(1室10名~20名、間仕切り撤去可)があり、オープン問題セッションの後のディスカッションにはうって つけでした。この会場は人気があり早くからの予約が必要ですが、今回はたまたま1日だけあいていたので使 わせてもらいました。

6月に次回会議の日程を決めたにもかかわらず、私の不手際で会議の直前まで案内をだすことができません でした。伊藤先生はじめ事務局の先生方にはご心配をおかけしました。幸い柳浦先生が10月に名古屋大学 に着任されてからは進捗著しく無事会議を終えることができました。ご協力いただいた皆様に感謝します。

2005年度第2回幹事会議事録 事務局

日時:2005.11.21, 16:00--17:30

場所:名古屋大学 野依学術交流館 カンファレンスホール

[活動報告]

○ミニ研究集会

- ・組合せゲーム・パスル(9/12,京大)
- ・複雑ネットワーク・ウェブグラフ(9/13, 関学)
- ・列挙アルゴリズム(9/27--29,群馬大)
- •Complexity(10/11, 東工大)

〇ポスドク

- ·西村治道(10月~,京大)
- ○招聘研究者
- ○国際会議渡航補助

·FOCS(河原林[東北大])

○事務員

·宮本希(8月~)

○国際会議 ランダムネス [渡辺]

7/18--21, 東北大学

[進捗状況報告/討論]

○教科書シリーズ [杉原]

順調。脱稿2件

- ○ニュースレター [牧野・宇野]
 - ・これまでに3回発行。4回目は来年の1月。
 - ・年に2回程度の予定で進めたい。->了承(必要ならば、号外を適宜。)

○電子ジャーナル [定兼]

- ・引き続き、検討を続ける
- ・協力が必要な場合は、後日に相談
- ○春学校、ワークショップ (2/27(月) ~ 3/3(金))
 - ・オーガナイザ・プログラム:浅野孝夫
 - ・ローカルアレンジ:西野・垂井
 - ·2/27(月)~3/1(水)春学校
 - 電通大図書館 AV教室(301号室, 306号室)
 - 講師は5名
 - ·3/2(木) ~ 3/3(金) ワークショップ
 - 調布クレストンホテル
 - ・春学校 … 初学者の道案内を目指す(定員 90名)

(※ 参考: 春学校に参加させてみたい学生数 … 会場だけで 30名以上)

近隣の諸国にもアナウンスを流す

・春学校/ワークショップではregistration fee を徴収する(高額ではない)

○ミニプロジェクト

・ゲーム [伊藤]

- ミニ研究集会を開催

- 今後、スクールみたいなのを開催したい
- •Complexity [垂井]
 - ミニ研究集会を開催、今後も開催予定
 - 学生にも声をかけたい
- ・ジオメトリ[加藤]
 - 3月後半、学生対象の mini school
 - 講師 加藤・浅野・徳山・(David Avis)
- ·列挙[中野]
 - ミニ研究集会を開催
- ○電子情報通信学会 特集号 [増山]
 - ・来年6月に、和文論文誌、特定の解説論文の特集
- ○ICALP ワークショップ
 - ・SATの厳密解法関係のミニワークショップの提案(渡辺先生)
 - ヨーロッパの人々とも相談する必要あり
 - (後日、渡辺先生によって申請、採録された。)
- ○電子情報通信学会 英文論文誌 [和田]
 - ・来年の8~10月頃、〆切は2月頃
 - ・基本は研究論文(解説論文も ok かも)
 - ・もう少し検討
- ○予算の使途
 - ・国際会議発表のための海外渡航の補助
- ○次回の全体会議
 - ·九州地区(山下、定兼)
 - ・2006年、5月か6月頃
- ○電子情報通信学会 全国大会 企画 [山下]

·3月24日(金)~27日(月)国士舘大学

- ・若手(主に博士課程の学生)のためのシンポジウム
- ・参加費と旅費は、総括班から支援

○次回以降の春学校/ワークショップ

- ・来年度からは、時期を夏休みぐらいにしたい
- ・早急に内容の相談を始める必要がある

OEATCS記事募集 [牧野]

・特定内外で、記事があれば、牧野まで

○特定領域としての成果

- ・どんな集まりで、どんなアイデアが出て、どんな進展・研究があったか
- どんなに細かいことでも良いので、事務局まで教えてほしい

○次の特定領域

- ・来年は、調査研究の申請をする
- ・そのために、今の特定領域で、領域として集まった成果を挙げる

「新世代の計算限界」イベントカレンダー + 事務連絡

成果報告書の執筆依頼

今年度の成果報告書を領域として作成する時期となりました.詳細は別途メール・ウェブで御連絡いたしま すが,昨年度と同じ要領で,A01からC11の研究課題ごとに「研究活動報告」と「発表文献リスト」を御執筆い ただけますようお願い申し上げます.なお,提出〆切は3月31日(金)です.

http://keisan-genkai.lab2.kuis.kyoto-u.ac.jp/members/reports.html

旅費の補助

総括班予算より国際会議等の旅費の補助を行います。補助希望者は伊藤までお申し出ください。特に締め切りなどは設定しておりません。

·ジオメトリ合宿 (加藤直樹 浅野哲夫、徳山豪、David Avis)

2006年3月15日昼から17日まで

コミュニティ嵯峨野(京都府勤労者研修センター)(<u>http://www.com-sagano.com/</u>)

〒616-8372 京都市右京区嵯峨天龍寺広道町3-4

COMP-NHC 学生シンポジウム(田中圭介)

2006年3月24日(金)~26日(日)のいずれか一日、国士舘大学世田谷キャンパス(東京都世田谷区)

電子情報通信学会 2006 年総合大会 シンポジウム講演 (B) DS-1「COMP-NHC 学生シンポジウム」

(総合大会は2006年3月24(金)~27(月))

3/27(月)-29(水) EWCG'06 (22nd European Symposium on Computational Geometry) Delphi, Greece 3/31(金)-4/2(日) Algebra and Combinatorics 2006 Manila, Philippines 4/24(月)-25(火) 第19回 回路とシステム軽井沢ワークショップ 軽井沢プリンスホテル西館 5/21(日)-23(火) STOC'06 Seattle, Washington, USA 6/5(月)-7(水) SoCG'06 Sedona, Arizona, USA 6/20(火)-22(木) The Second International Conference on Algorithmic Aspects in Information and Management 香港 6/22(木)-24(土) WG2006 (グラフアルゴリズムの会議です) ノルウェー, ベルゲン 6/25(日)-28(水) INFORMS International Hong Kong 2006 香港 7/3(月)-5(水) SIROCCO2006 Chester, United Kingdom 7/6(木)-8(土) SWAT 2006 リガ (ラトビア) 7/9(日)-16(日) ICALP 2006 ベネチア (イタリア) 7/30(日)-8/4(金) ISMP2006 Rio de Janeiro, Brazil 8/15(火)-18(金) COCOON 2006 台北(台湾) 8/22(火)-24(木) IFIP/TCS 2006 Santiago, Chile 9/11(月)-15(金) ALGO 2006 チューリッヒ (スイス)

このニュースレターについて

ニュースレター各号は電子メールで配布する予定です.短い記事や連絡事項は全て掲載しますが,長い記 事,イベントの詳細などはwebページに掲載する予定です.webページには詳細まで全てを載せた完全版を掲 載して,目次,あるいは各記事の末尾のURLを参照すると,web版の同じ記事を参照できるようにいたします.

記事は、各回、1つの研究課題に担当をお願いする予定です。各研究課題で2000-4000字程度、研究に関わる記事を書いていただければと思います。通常、このようなニュースレターでは、研究成果を報告するのが 一般的だと思われますが、この特定領域では「研究者の交流」に焦点を当てたいため、「研究の成果以外」の 記事を面白く解説していただければと思います。例えば、最近参加した国際会議の情報を、どのようなものが 流行っていたか、何が面白かったか、などの主観的な解説を交えて報告したり、最近考えている問題、あるい はオープン問題を、この辺までは解けるがここがうまくいかない、といった解説を交えて紹介する、という形で す。

また,研究者間の交流を促進するため,各研究者の,国内外の会議への出席予定を集約して掲載していこうと考えています.研究者の交流には,顔をあわせる回数を増やすことが肝要です.他の研究者の参加予定

がわかれば, 会議への出席のモチベーションを高めることにもつながり, それがディスカッションや研究成果を 生むきっかけにもなるでしょう. 特定領域メンバーの皆さんには, 自分のわかる範囲で, 国内外の会議・研究 会の情報と, 自分の参加予定を教えていただければと思います.

この他,個人からの寄稿を募集いたします.100-1000字程度で,情報宣伝されたいことを自由な形式で書い て送っていただければ,掲載いたします.メールで配布する関係上,テキスト形式のものしか扱えませんが, そこはご了解お願いいたします.

次号は6月ごろを予定しています.

★ ニュースレター編集委員では,皆様からのご意見をお待ちしております.編集方針や内容の追加など編集 全体にかかわることから細かいことまで,幅広いご意見をお願いいたします.

■■ 新世代の計算限界 ニュースレター ■■

編集委員長 宇野 毅明 uno@nii.jp (問合せ先)

副編集委員長 牧野 和久 makino@sflab.sys.es.osaka-u.ac.jp

平成17年度第1回全体会議

日時 平成 17 年 6 月 16 日 (木), 17 日 (金)

会場 国立情報学研究所 12 階会議室

プログラム

6月10日(不))
----------	---

10:00 - 10:50	A01-A10 班の研究紹介

- 11:00 12:00 ポスターセッション (A01–A10 班)
- 12:00 14:00 幹事会
- 14:00 14:30 B01–B06 班の研究紹介
- 14:40 15:40 ポスターセッション (B01-B06 班)
- 16:00 16:55 C01-C11 班の研究紹介
- 17:05 18:00 ポスターセッション (C01–C11 班)

6月17日(金)

()		
9:30 - 12:00	全体会議	
13:30 - 14:15	時間枠つき配送計画問題に対するメタ戦略アルゴリズム	
	柳浦睦憲 (京都大学)	59
14:15 - 15:00	孤立したクリークの線形時間列挙	
	伊藤大雄 (京都大学)30	64
15:30 - 16:15	Compact Encoding of Plane Triangulations with Efficient Query Support	
	中野眞一 (群馬大学)	69
16:15 - 17:00	オンライン予測の理論と応用	
	瀧本英二 (東北大学)3'	73





































number of customers		best known	MB (2003)	GH (2002)	LC (2003)	our ILS
	CNV	692	694	696	694	69
200	CTD	169281	168537	179328	173061	17033
	time (min)		5.88	3.83	21.66	33.
	CNV	1386	1389	1392	1390	138
400	CTD	392444	390386	428489	408281	40128
	time (min)		12.49	12.95	43.32	66.
	CNV	2076	2082	2079	2088	207
600	CTD	799355	796172	890121	836261	82719
	time (min)		29.39	23.53	64.97	10
	CNV	2754	2760	2765	2766	275
800	CTD	1429914	1535849	1361586	1475281	142613
	time (min)		42.32	106.53	86.63	133.
	CNV	3438	3446	3446	3451	343
1000	CTD	2106125	2078110	2290367	2225366	216945
	time (min)		440.82	361.2	108.29	166.
NV: cum	ulative nui	nber of v	ehicles 1	MB: Mes GH: Gebi	ter and B	sräysy Hombe







計算の複雑さに関する結果 ・時刻依存の移動時間および移動コスト ・移動時間関数の総複雑度に対して線形時間のDP ・移動時間関数の複雑度が各客ペアでO(1) →移動時間が定数の場合と同程度の計算時間 ・移動時間が可変,移動時間に応じたコスト ・移動時間コスト関数が一般 → NP困難;擬多項式時間のDP ・移動時間コスト関数が凸(時間ペナルティ関数は 一般) → 多項式時間のDP

まとめ

- ・時間枠に対する要求を一般的なペナルティ関数として 表現できる汎用的な問題の定式化
 - → 時間ペナルティ最小化問題を解く必要性
 - →動的計画法の提案
- 高速化のさまざまな工夫
- Gehring & Homberger のベンチマーク問題例に対する 良好な結果
- ・移動時間に対する汎用的な定式化

今後の課題

- ・異なるタイプの汎用化
- ・配送ルートの探索能力の向上
- ・大規模問題例への適用
- ・時間枠の取り扱い →
 より実装しやすい単純なアルゴリズムで
 高速かつ実用的なものの開発
 cf. 提案手法はいずれも非常に複雑で実装困難



• L. Ed **







Observation for Proving Strategy of enumeration Theorem 0.1 · Check each vertex whether or not it can be a • Lemma 1. An isolated clique has a vertex pivot. (called a pivot) that has no outgoing edge from the clique. · Sort and renumber all vertices as $d(v_1) \leq d(v_2) \leq \ldots \leq d(v_n).$ pivots • If $\exists j > i, v_i \in N(v_i)$ (adjacent vertices of v_i), v_i can be ignored (we can consider that the vertex having the minimum index in a isolated clique is the pivot): Test (a). • If v is the pivot of S, then $S=N(v) \cup \{v\}$.







Observation 3

- Lemma 4. If S={v=w₁, ..., w_k} is an isolated clique and v is the pivot of S (d(w₁)≦...≦ d(w_k)), then S_i={w₁, ..., w_i} has at most i-1 outgoing edges from S.
- Proof. Assume that $|E(S_i, V-S)| \ge i$. Then $d(w_i) \ge d(v)+1$, and hence $d(v_j) \ge d(v)+1$ for all j=i+1, ..., k. Therefore $\Sigma_{w \in S}$ $sd(w) = \Sigma_{w \in Si}d(w) + \Sigma_{v \in S-Si}d(w)$ $\ge i+(k-i)=k$, contradiction.

Strategy 3

- Assume that v passed Tests (a) and (b) (primal tests).
- Let $S=N(v) \cup \{v\}=\{v=w_1, w_2, ..., w_k\}$ $(d(w_1) \le d(w_2) \le ... \le d(w_k)).$
- Clique test: From i=1 to k,
 - check whether (1) w_i is adjacent to $w_1,\,...,\,w_{i\text{-}1}$ (i.e., $S_i\text{=}\{w_1,\,...,\,w_i\}$ is a clique) and
 - (2) S_i has at most i-1 outgoing edges from S.
 - If not, v is not a pivot and then skipped (finish checking v).

Running time

- Sorting vertices by their degrees: O(m).
- Primal tests: O(d(v)) for each v∈V, i.e., O(m).
- Clique test: Assume the test breaks at w_p . $d(w_1)+d(w_2)+...+d(w_p)=O(w_p^2) \rightarrow O(m^2)$?
 - → More precise estimation!









Results (cont.)

- Theorem 2. Let *c*, *x*, and *y* are functions of *n* s.t. c=xy. There is a graph with *m* edges for which the # of maxmal *c*-isolated cliques is Ω ($(2x^{y/c^2})m$).
- Cor. 3. If *c*= ω(1), there is a graph with *n* vertices consisting of super-linear # of maximal *c*-isolated cliques.
- Cor. 4. If *c*= ω(log*n*), there is a graph with *n* vertices consisting of super-polynomial # of maximal *c*-isolated cliques.

Results (cont.)

- Cor. 1. If c=O(1), all maximal *c*-isolated cliques can be enumerated in linear time.
- Cor. 3. If *c*= ω(1), there is a graph consisting of super-linear # of maximal *c*-isolated cliques.
- c=Θ(1) is the tight bound for enumerating all maximal c-isolated cliques in linear time.

Results (cont.)

- Cor. 2. If *c*=O(log*n*), all maximal *c*-isolated cliques can be enumerated in polynomial time.
- Cor. 4. If *c*=ω(log*n*), there is a graph consisting of super-polynomial # of maximal *c*-isolated cliques.
- c=Θ(logn) is the tight band for enumerating all maximal c-isolated cliques in polynomial time.



Proof of Corollaries 3 and 4

- If $c=\omega(1)$, then by letting x=2, y=c/2 $(2x^{y}/c^{2})m$ becomes super-linear.
- If $c = \omega (\log n)$, then by letting $x = c/\log n$, $y = \log n (2x^{y}/c^2)m$ is super-polynomial.

Other Results: Pseudo-Clieques

- Let α (k) and β (k) are functions.
 Pseudo-Clique PC(α, β) is a vertexproper-subset S⊂V (|S|=k) s.t.
- $\operatorname{av}_{v \in S} d_{G(S)}(v) \ge \alpha(k)$ and
- $\min_{v \in S} d_{G(S)}(v) \ge \beta(k)$.

Results for PC

- Theorem 3. Suppose f(k)= Ω (1) and 0< ε <1 is a constant.
 - There is a graph including super-poly. # of maximal isolated PC(k-f(k),k^ε).
 - There is a graph including super-poly. # of maximal isolated PC(k-k^ε,k/f(k)).
- Proposition 1. All maximal isolated PC(α, c₁k) and PC(k-c₂,k^ε) are enumerated in poly. time for constant c₁<1 and c₂≧ 1.

Results for PC (Cont.)

- Theorem 4. There is a graph including super-poly. # of maximal isolated PC(k-(logk)^{1+ε}, k/(logk)^{1+ε}) for any 0<ε.
- Theorem5. All maximal isolated PC(k-logk,k/logk) can be enumerated in poly. time.

Summary

- Introduce f-isolated cliques with parameter function c.
- All c-isolated cliques can be enumerated in linear time for any constant c.
- $c=\Theta(1)$ is the tight bound of linear time enumeration.
- All c-isolated cliques can be enumerated in poly. time if c=O(logn).
- c=Θ (logn) is the tight bound of poly. time enumeration.









既知の結果その2					
クエリサポートなし ・一般の平面グラフ ・一般の平面グラフ ・極大平面グラフ ・極大平面グラフ ・極大平面グラフ	4m bit 3.58m bit 1.53m bit 1.33m bit 1.33m bit 1.08m bit ⁻	[Turan 84] [Keeler他 95] [同上] [He他 99] [Poulalhon 03] ICALP 2003 下限			
グラフを符号化して送り、複合化して使用					

既存の結果その3						
クエリサポートあり	J					
・木	2m +o(n)bit	[Jacobson 89] FOCS 89				
・一般平面グラフ	2m+8n+o(n)bi	t [Munro他97] FOCS 97				
・一般平面グラフ	2m+2n+o(n)bi	t [Chiang他01] SODA 01				
・極大平面グラフ	2m+n+o(n)bit	[Chuang他]				
一般平面グラフ	2m+5n+o(n)bi	it ICALP 98				
・極大平面グラフ	2m+o(n)bit	[中野 2005]				
グラフを符号化して記憶し、複合化 <mark>せずに</mark> 使用						





















隣接点のリスト

・隣接点のリストをO(d(v))時間で作成できる

略

まとめ

[中野 05]

既知のベスト(クエリサポートつき) 極大平面グラフ 2m+n+o(n)bit [Chuang他 ICALP98]

<mark>今回</mark>の結果 極大平面グラフ 2m+o(n)bit

主なアイディア リ<mark>アライザ</mark>の3本の木を 1本 + 2本 とみなすこと

おしまいです

- 御静聴ありがとうございます。
- ・質問・コメント・アドバイス等歓迎します。
- •応用?

中野眞一(群馬大学)






二分法(Halving Algorithm)

■ クイズ王(エキスパート)N人のうち少なくとも1 人は全問正解すると仮定

二分法: 「これまで一度も間違えていないエキスパート たちの予測の多数決に従う」

誤り回数は高々[log₂N]





重みつき平均アルゴリズム --- 全問正解者がいない場合 ---

■ 重み更新

v_i ← v_i b_i / 規格化

ただし, **b**_i = 1 ・・・ エキスパート i が正しかったとき **0**¹¹・・・ エキスパート i が間違えたとき

誤り回数 $\leq m^* + O\left(\sqrt{m^* \ln N} + \ln N\right)$ m* は最も成績の良かったエキスパートの誤り回数

パーセプトロン ∈ 加算型重み更新

- 学習目標 v* = (0, ..., 1, ..., 0)の線形分離関数 (シングルトン)の学習
- パーセプトロンの出力:重みつき多数決 y' = sgn(∑_{i=1}^N v_i x_i ≥ 0)
- 重み更新 v_i ← v_i + η (y−y') x_i ただし, y は真の値, η は学習定数
- 誤り回数 = Ω(N) cf. 二分法は [log N]

パーセプトロン vs Winnow 加算型重み更新 乗算型重み更新 学習目標が v* = (1, 0, 1, 0, 0, ..., 1) のとき ただし, k = ∑; v*;

- パーセプトロン
 - $v_i \leftarrow v_i + \eta (y-y') x_i$ 誤り回数 = Ω(kN) カーネル手法と整合する

Winnow

v_i ← v_i exp(η (y−y') x_i) / <mark>規格化</mark> 誤り回数 = O(k log N) ・・・ attribute efficient 一般に, カーネル手法と整合しない

この問題と解析の特徴

■ オンライン性

- 予測と結果の提示が交互に繰り返される
- ユニバーサル性 - エキスパートの予測や結果の系列に対して何の 仮定もおかない
- 最悪値評価 - アルゴリズムの性能を最悪の場合で評価する
- ■相対評価
 - アルゴリズムの性能を, 最適なエキスパートの成 績を用いて相対的に評価する

オンライン予測

■ モデル化

- エキスパート統合モデル
- 統合アルゴリズム
 - Winnow (1988)
 - Aggregating Algorithm (1990)
 - Weighted Majority, Weighted Averaging (1993,1994,1999)
 - GD, EG family (1994)
 - Potential-Based (2001)
 - Following Perturbed Leader (2003)

■ 応用

- 学習, 符号化, 統計的推定, 投資, ゲーム理論, 最適化, …





クイズ王に匹敵する問題
• Y = {0,1}
• Z = {0,1}
• L(y, z) =
$$|y - z|$$

· に相当
LOSS^T_A = アルゴリズム A の誤り回数
LOSS^T_i = エキスパート i の誤り回数
 $m^* = \min_i \text{LOSS}^T_i$
ヨA, $R_A^T \leq 4\sqrt{m^* \ln N} + 2.5 \log_2 N$













さまざま Loss $_{AA}^{T} \leq$ Loss $_{WAA}^{T} \leq$	にな凸損失関数と c _A ⊆Loss <mark>⊼ + c_{AA} ln N</mark> ⊆Loss ⊼ + c_{WAA} ln N	$_A$, $c_{ m W}$	VAA	
損失関数	$L(oldsymbol{y},oldsymbol{z})$	c_{AA}	c_{WAA}	
2次損失	$(y - z)^2$	1/2	2	
相対エントロ ピー損失	$(1-y)\ln\frac{1-y}{1-z}+y\ln\frac{y}{z}$	1	1	
Hellinger 損失	$\frac{\frac{1}{2}((\sqrt{1-y}-\sqrt{1-z})^2)}{+(\sqrt{y}-\sqrt{z})^2)}$	1/√2	1	
絶対値損失	y - z	凸間 では	周数 ない	

















無限のエキスパートの統合

■ Y = {0,1} ■ エキスパート $\theta \in [0,1]$ $x_{t,\theta}(1) = \Pr(1 \mid y_{1,t-1}, \theta) = \theta$ $x_{t,\theta}(0) = \Pr(0 \mid y_{1,t-1}, \theta) = 1 - \theta$ ■ AAの出力 = KT-Estimator $z_t(y) = \int v_{t,\theta} x_{t,\theta}(y) d\theta$ $= \frac{\sum_{s=1}^{t-1} y_s + 1/2}{t}$ ■ 相対損失 = O(log T)









オンライン投資

- ポートフォリオ問題において、毎時刻全額投資する かわりに、定額(例えば1万円)投資
- 株価の上昇率に上限 B > 0 を仮定 (相対損失を有限に抑えるため)

以下の問題 <mark>(Y, Z, L)</mark> と等価

- Y = [0, B]^K ■ Z = { $z \in [0,1]^{K} | z(1) + \dots + z(K) = 1$ }
- L(y, z) = y · z •••内積損失

WAAの適用

■ 内積損失 L(y, z) = y · z は凸ではない ■ 適当な η に対して 相対損失 = $O\left(\sqrt{2BL^* \ln N} + B \ln N\right)$ L* = 最適なエキスパートの損失





















まとめ

- ■さまざまなオンライン最適化問題が, 共通の 統合スキーム(AA, WAA)を用いて解ける
- ■一般の最適化問題への適用
 - 各時刻 t における損失が, それまでの履歴に依存する
 - オフラインで最適な解を求めるのが困難
- ■構造的エキスパートの統合の可能性
- 決定リストの効率の良いオンライン学習
- ■モデルのさらなる拡張
 - リスク情報を用いた統合

平成17年度第2回全体会議

日時 平成 17 年 11 月 21 日 (月), 22 日 (火)

会場 名古屋大学 野依記念学術交流館 / IB 電子情報館 中棟 IB 大講義室

プログラム

- 11月21日(月)
 - 10:30-12:00 未解決問題セッション

Domain-Specific Image Segmentation and Shape Matching

浅野哲夫 (北陸先端科学技術大学院大学)

柔体を扱う計算幾何学

浅野哲夫 (北陸先端科学技術大学院大学)

- 線的施設配置問題
 - 加藤直樹 (京都大学)
- 外面が四角形 (以上) である格子凸描画を求める線形時間アルゴリズム
- の開発
 - 三浦一之(福島大学)
- 幾何的巡回セールスマン問題の厳密アルゴリズムについて 岡本吉央(豊橋技術科学大学)
- 文字列検索における時間と領域のトレードオフ 定兼邦彦(九州大学)
- 12:00 13:30 幹事会
- 13:30 15:50 未解決問題に対するグループ討論
- 16:00 17:30 全体会議
- 11月22日(火)

10:15 - 11:00	透過的データ圧縮
	定兼邦彦 (九州大学)
11:10 - 11:55	確率伝搬法の可能性について
	渡辺治 (東京工業大学)
13:30 - 14:15	RNA 配列の比較アルゴリズム
	浅井潔 (東京大学)
14:30 - 15:15	3つの資源節点集合を持つ 4 点連結グラフを均等分割する問題について
	石井利昌 (豊橋技術科学大学)
15:30 - 16:15	平面グラフ、曲面上のグラフ、マイナーに関して閉じているグラフに関す
	る彩色問題
	河原林健一 (東北大学)408
16:30 - 17:30	未解決問題グループ討論の結果報告

봡 롶 透過的データ圧縮 Transparent Data Compression データ圧縮の目的(過去) - ディスク容量の節約 保存用 K. Sadakane and R. Grossi: Squeezing Succinct Data Structures into Entropy Bounds, Proc. ACM-SIAM SODA 2006, to appear. - 通信コスト(料金,時間)の削減 データ圧縮の目的(現在) - アクセスの高速化 (CPU速度 > ディスク速度) 九州大学システム情報科学研究院 - 連続的なアクセスに限られる 定兼 邦彦 ランダムアクセスができたらどうなるか 2005年11月22日 2

透過的データ圧縮 もし圧縮データの任意部分を高速に復元できれば ・データを圧縮したまま保存できる - ディスク容量の節約 - 高速なアクセスが可能 • 圧縮されていることを意識しなくていい

















木のSuccinctデータ構造の圧縮

- FIDでは不可能
 - $-2n + O(n \log \log n / \log n)$ bits [GRR04]
- EIDでは可能
 - 木を表現する括弧列 S の H_k まで圧縮可 (k=O(log log n))

14

- $-2nH_k$ + O(*n* log log *n*/log *n*) bits
- 問合せの計算量は圧縮前と同じ
- 構造が同じ部分木があると H_k は小さくなる (接尾辞木で特に有効)



従来の圧縮法 16

従来	の圧縮法	
辞書式圧縮法 LZ77 [Ziv, Lempel 77] LZ78 [Ziv, Lempel 78] LZW[Welch 84] compress	統計的圧縮法 PPM[Cleary, Witten 84] PPMD [Howard 93] PPM*[Cleary, Teahan, Witten 95] block sorting	 文字列を話 辞書=す a compressea array, a Pat trea
LZSS [Storer, Szymanski 82] gzip	[Burrows, Wheeler 94] 高圧縮率、PPMより高速 (bzip2) context tree weighting	a compressea array, a Pat [l≕
	[Willems, Shtarkov, Tjalkens 95] PPMより高圧縮率 ¹⁷	







従き	来の圧縮法	との比較
	漸近的圧縮率	log <i>n</i> ビットの復号時間
LZ77 [ZL77]	nH_k	O(<i>n</i>)
LZ78 [ZL78]	nH_k	O(<i>n</i>)
PPM [CW84]	nH_k	O(<i>n</i>)
CTW [WST95]	nH_k	O(<i>n</i>)
Block Sorting [BW94]	nH_k	$O\left(\frac{\log^2 n}{\log\log n} + \log n \cdot \log \sigma\right) [GGV03]$
本研究	nH_k	O(1)
		21























得られた結果

定理. Planted Solution Model のもとでアルゴリズムを実行させたとき, $\Pr[\text{ the algorithm yields the planted solution }] \ \ge \ 1 - 2n \cdot e^{-c_1 n \cdot \frac{(p-r)^4}{p^2}}.$

系.

For any $\delta > 0$ and any p and r, cp < r < p < 1, $p - r \ge c_2 n^{-1/2} \log(1/(\delta(p - r))))$ \Rightarrow the algorithm arswers correctly with probability $> 1 - \delta$.

これまでのアルゴリズムとの比較:

 $\begin{array}{ll} (\text{Boppana, 1987}) & p-r = \Omega(n^{-1/2}) \; (\text{almost}!) \\ (\text{Dyer and Frizz, 1989}) & p-r = \Omega(n^{-1/2}) \; \text{when p is large, e.g., $\Omega(1)$} \\ (\text{Jerrum and Sorkin, 1998}) & p-r = \Omega(n^{-1/6}) \\ (\text{Cendon and Karp, 2001}) & p-r = \Omega(n^{-1/2+\varepsilon}) \end{array}$

その他にも...

• Most Likely Partition 問題(with Onsjoe)

- MAX-2SAT 問題に対する同様のアプローチ (with 山本)
- 固有値を求める手法にできないか?

未解決問題:本当のBPの妥当性?

参考文献

R. McElice, D. MacKay, and J.F. Cheng, Turbo decoding as an instance of Pearl's "Belief Propagation" algorithm, *IEEE J. on Selectend Areas in Communications*, 16(2):140–152, 1998.

10

































SCFGのCYK アル	ゴリズム
$\mathcal{F}(i,j)$: maxmum probability of $(xi, \dots xj)$ is parsed by S	S = xS $S = Sx$ $S = xSy$ $S = SS$
$\begin{aligned} \gamma(i,j) &= \max \begin{cases} \gamma(i+1,j) + \log p(x,S) \\ \gamma(i,j-1) + \log p(Sx_j) \\ \gamma(i+1,j-1) + \log p(y_i) \\ \max_{i \in \mathcal{A}} \varphi_i(i,k) + \gamma(i,k) \end{cases} \end{aligned}$; ; ; $x_i S x_j);$ b = 1 - () = b = e(CC)





RNA 配列比較の課題

РНММТS ЦССО

PSTAG

- 2本配列の比較
- 局所アラインメント(検索)
 ・ 配列群のマルチプルアラインメント Marlet
- ・ 配列群のマルナノルアラインメント
 ・ 配列群の共通二次構造予測
- ・ に勿辞の共通二次構造ア別
 ・ 二次構造付配列への構造アラインメント
- 局所アラインメント(検索)
- 二次構造への構造アラインメント
 局所アラインメント(二次構造モチーフ検索)
- 局所アラインメント
 ・ 配列クラスタリング
- 配列クラスタリンク • 特定RNAのモデル
- 特定RNAのモナル
 局所アラインメント(特定RNA発見)
- 比較ゲノム保存領域からの機能性RNA発見

共通2次構造予測

- Sankoff (1985)
 - 2(N)本のRNAの折畳み。 $O(L^{3N})$ time, $O(L^{2N})$ space
- Gorodkin, Heyer & Stormo (FOLDALIGN, 1997)
 2本局所整列、分岐構造なし、塩基対の最大化、O(L⁴) time
 多重整列: トーナメント法+ Greedy アルゴリズム
- Mathews & Turner (Dynalign, 2002)
- Perriquet, Touzet & Dauchet (CARNAC, 2003)
 局所配列類似性 + エネルギー + 塩基対共変





























Acknowledgments

- "Functional RNA Project" of METI
- Grant-in-Aid for Scientific Research on Priority Area "Comparative Genomics"
- AIST, CBRC, BIRC
- JBIC
- The University of Tokyo













3-bipartition

- (1) 5-vertex-connectivity does not suffice!
- (2) 4-vertex-connectivity suffices if K4 is contained.
- (3) For the edge version of *k*-bipartition (*k*=1,2,3), (*k*+1)-edge-connectivity suffices.























Our Results

3-bipartition

- 1) 5-vertex-connectivity does not suffice!
- (2) 4-vertex-connectivity suffices if K4 is contained.
- 3) For the edge version of k-bipartition (k=1,2,3),

(*k*+1)-edge-connectivity suffices.

Theorem

If *G* is a 4-vertex-connected graph and contains K4, then there exists a 3-bipartition, and moreover, a 3-bipartiton can be found in $O(n^3 \log n)$ time.



























































Our Results

3-bipartition

- (1) 5-vertex-connectivity does not suffice!
- (2) 4-vertex-connectivity suffices if K4 is contained.
- (3) For the edge version of k-bipartition (k=1,2,3),
 - (*k*+1)-edge-connectivity suffices.











What we have done is ...

Every 4-vertex-connected graph G admits a 3-bipartition if G has a K_4

5-vertex-connecitivity does not suffice for 3-bipartition 5-vertex-connecitivity does not suffice for 4-bipartition 5-vertex-connecitivity does not suffice for 5-bipartition

The vertex version implies the edge version.

Every (k+1)-edge-connected graph *G* admits a *k*-bipartition of E(k=1,2,3).

Open Problems

·Sufficient condition for which a k-bipartition exists

Conjecture

Every (k+1)-vertex-connected graph with Kk+1 admits a k-bipartition.

the edge version

Conjecture Every (*k*+1)-edge-connected graph admits a *k*-bipartition.

Open ProblemDefine f(k) be the smallest p such that every
p-vertex-connected graph admits a k-bipartition.
f(1)=2, f(2)=3For k>5, prove $f(k) \ge k+1$. $f(3) \ge 6, f(4) \ge 6, f(5) \ge 6$ For k>3, bound f(k) from above by k+constant. $f(k) = O(\Sigma | T_i|)$ The same questions for the edge version.

Approximating graph coloring of minor-closed graphs

Joint Work with Erik Demaine, Mohammad Hajiaghayi, Bojan Mohar, Robin Thomas Partially joint Work with Neil Robertson and Paul Seymour

Ken-ichi Kawarabayashi

Tohoku University

E-mail: <u>k_keniti@dais.is.tohoku.ac.jp</u>

http://www.dais.is.tohoku.ac.jp/~k_keniti

Contents (Mostly, FOCS paper)

- Motivation (FOCS paper)
- 2-approx. of the chromatic number of minorclosed graphs (FOCS paper)
- Tree-width, Grid-minor, RS-structure.
- Overview of Algorithm (Robertson-Seymour)
- Approx. the list-chromatic number of minorclosed graphs.
- Toward Structural Theorem

Why is it accepted in FOCS?

- It is building on math deep theory. (although NOT AT ALL practical.)
- Minor-closed graphs are natural. (a generalization of planar graphs.)
- It tells how to use RS' main structural theorem.
- It is a bit easier to access (than RS' papers)
- Nice approx. for graph coloring of minor-closed graphs.
- Lucky.

Motivation

Mathematical Motivation

- 1. Hadwiger's Conjecture. (A far generalization of 4CT)
- 2. Graph Minor Theory (Robertson-Seymour) Algorithmic Motivation
- 1. Chromatic number is hard to compute. NP-complete even for deciding 3-colorability of Planar graphs.
- 2. Even hard to approx.
- NP-hard to approx. within constant factor.
- NP-complete to decide the chromatic number of minor closed graphs. (Even for planar graphs) Can you approx. ?

Algorithmic Results

- Theorem (Demaine, Hajiaghayi, KK, FOC\$2005) There exists a 2-approx. algorithm for the chromatic number in minor-closed graphs. (graphs with no Kk-minor)
- The best known result was O(k $\sqrt{\log k}$) approx.
- Proof uses the whole graph minor papers....
- Robertson-Seymour theory consists of 23 papers. Most of them are published in JCTB.

Why is it 2-approx ?

- The main theorem says that if G is Kk-minorfree graphs, then it can be decomposed into two graphs G₁,G₂ such that both G₁ and G₂ have tree-width at most f(k).
- If tree-width is bounded, one can compute the chromatic number in the linear time.
- It remains to give an algorithm for the main theorem...
Proof depends on

- Robertson-Seymour theorem
- It gives a structural theorem for minor-closed graphs.
- Once we have this structure, the rest of proof is not so hard (but not trivial.)
- The main challenge is how to obtain RSstructure.
- It depends on the whole graph minor papers.



Graph Minor Theory [Robertson & Seymour 1984 20041

Seminal series of ≥ 20 papers

- Powerful results on excluded minors: • Every minor-closed graph property
 - (preserved when taking minors) has a finite set of **excluded minors** [Wagner's Conjecture] Every minor-closed graph property
 - can be decided in polynomial time
 - For fixed graph H, graphs minor-excluding H have a special structure: drawings on bounded-genus surfaces + "extra features"

Highlights of Graph Minor Theory

Theorem(The disjoint paths problem) For fixed k, there is a polynomial time algorithm for deciding the disjoint paths problem.

- Minor testing can be done.
- Tree-width and grid-minors are discovered.
- Many mathematical and algorithmic applications.









- Tree-width at most 1 < = > G is a forest.
- Tree-width at most 2 < = > G is series parallel.
- Tree-width at most 3 < = > G has no minor isomorphic to K₅, Octahedron, 5-prism, V₈.
- Tree-width of the complete graph of order n is n-1.
- Tree-width is minor-monotone.
- The (k × k)-grid minor has the tree-width k.







The optimal bound. w(r) = O(r) [Demaine, Hajiaghayi KK 2005]

Grids certify large treewidth in H-minor-free graph

Huge Grid is important

- Routing problem
- The disjoint paths problem and its generalization.
- Actually, Robertson-Seymour use this idea.

Structure of H-minor-free Graphs [GM16—Robertson & Seymour 2003] Main result of RS

- Every H-minor-free graph can be written as O(1)-clique sums of graphs
- Each summand is a graph that can be O(1)-almost-embedded
- into a bounded-genus surface
- O(1) constants depend only on |V(H)|





But

There cannot be so many crossings that are far apart.
The genus addition process stops quite soon.

Otherwise, we would get a desired minor, a contradiction.



- If there is no crosscap in the small area, then it is either vortex or planar graph.
- There cannot be many non-planar small areas that are far apart. This tells us that there are bounded number of vortices.

In summary

- 1. Stating with huge grid H.
- 2. As long as there is a long jump, we shall detect handles.
- 3. Otherwise the graph is embedded into a surface such that all the non-planar graphs are in small areas.
- 4. We shall look at each small area, and detect either vortex or crosscap.
- 5. There are only finitely many vortices and crosscaps. So the process stops.



Approx. list coloring

 Theorem[Mohar and KK]
 There is an O(k)-approx. for graphs without Kk-minor, I.e., minor-closed graphs.

Actually, it is "almost" O(√logk)-approx.
It is approximating within O(√logk)c + O(k), where c is optimal.
The best know appox. was O(k √ logk) approx.

Open: O(1)? (Maybe NP-hard.)

Algorithm for List-coloring Theorem(KK & M) There is an O(n³) algorithm for the following: Input : A graph G, vertex set Z with |Z| <= 4k, precoloring of Z and each vertex in G has 16k-colors available in each list. Output : Determine either G has a Kk-minor, or Precoloring of Z can be extended to the whole graph G, or G has a subgraph H such that H has no Kk-minor and has a vertex set Z' with |Z'| <=4k such that some precoloring of Z' cannot be extended to H.

Algorithm for List Coloring

Corollary

- There is an $O(n^3)$ algorithm for deciding the following:
- (1) G has a Kk-minor
- (2) G has a 16k-list-coloring
- (3) G has a subgraph H such that H has no Kk-minor and no 12k-list-coloring.

It is easy to list-color by O(k $\sqrt{\log k}$) colors

NHC Spring School and Workshop on Discrete Algorithms

Feb. 27th – Mar. 3rd, 2006

University of Electro-Communications / Chofu Creston Hotel

Spring School

Feb. 27th (Mon.)	
$09{:}00 - 12{:}30$	Data-Driven Computing
	Bernard Chazelle (Princeton University)
12:30 - 14:00	Lunch
14:00 - 17:30	Sensor Networks: A Digital Bridge to the Physical World
	Leonidas J. Guibas (Stanford University)
Feb. 28th (Tue.)	
$09{:}00-12{:}30$	Games in Networks: Routing, Network Design and Potential Games
	Eva Tardos (Cornell University)
12:30 - 14:00	Lunch
14:00 - 17:30	Polynomial Time Algorithms for Market Equilibria
	Vijay V. Vazirani (Georgia Institute of Technology)473
Workshop	
Mar. 1st (Wed.)	
09:00 - 12:30	Random Sampling Techniques and Approximation of MAX-CSP
	Marek Karpinski (University of Bonn)
12:30 - 14:00	Lunch
14:00 - 17:30	Discussion
Mar. 2nd (Thur.)	
09:30 - 10:30	Games in Networks, Equilibria, and Inefficiency
	Eva Tardos (Cornell University)
10:30 - 11:00	Break
11:00 - 12:00	Approximation Schemes for Metric Clustering and Partitioning
	Marek Karpinski (University of Bonn)
12:00 - 13:30	Lunch
13:30 - 14:30	Discrete Optimization and VLSI-Design
	Bernhard Korte (University of Bonn)
14:30 - 15:00	Break
15:00 - 16:00	Approximation Algorithms for Facility Location
	Jens Vygen (University of Bonn)
16:00 - 16:30	Break
16:30 - 17:30	Algorithms for a Networked World
	Magnus M. Halldorsson (University of Iceland)

Algorithms for String Manipulation and Related Problems
D. T. Lee (Academia Sinica)530
Break
Dynamic Data Structures in Computational Geometry
Timothy M. Chan (University of Waterloo)
Lunch
Geometric Networks: Integer Linear Programming and Combinatorial
Algorithms
Alexander Wolff (University of Karlsruhe)
Break
Geometric Embeddings and Graph Expansion
James R. Lee (UC Berkeley)
Break
Distance Trisector and Voronoi Diagram with Neurtal Zone
Takeshi Tokuyama (Tohoku University)575

Copyright © <u>Bernard Chazelle</u> | Princeton, January 2006 | <u>home</u> | To appear in "Math Horizons," 2006 | AAAS '06 slides

Could Your iPod Be Holding the Greatest Mystery in Modern Science?

by Bernard Chazelle

Tuturologists are an amiable bunch, so it is a puzzle why the future has been so cruel to them. From flying cars and self-cleaning houses to that bugaboo of workaholics—the leisure society—the soothsayer's trail is littered with the carcasses of pet predictions turned roadkill.

Gordon Moore need not worry. The co-founder of Intel tried his hand at crystal gazing once—and struck gold. His celebrated "law" makes the outlandish prediction that every 18 months, like clockwork, white-clad technicians will huddle in a silicon wafer clean room and cram twice as many transistors onto a microchip.

Moore's Law has ruled the roost for the last 40 years. All the oohs and aahs you hear about the digital revolution are nothing but the squeals humans emit when tickled pink by Moore's Law. From the nice (medical imaging, e-commerce, whole-genome sequencing) to the vital (Xbox, IM, iPod), its rule has been a veritable ticklefest. Moore's Law has been the sizzling cauldron in which savvy cooks have whipped up a dazzling variety of tasty dishes. Without it, the Information Superhighway would be a back alley to Snoozeville; the coolest thing about a computer would still be the blinking lights.

Moore's law has had a good run but, alas, its days are numbered. By mid-

century, a repeal is all but certain. With the heady days of the Incredible Shrinking Chip receding in the past, expect the revolution to grind to a halt; expect pioneers to give way to tinkerers. Bye-bye ticklefest, hello slumber party.

No tears please. Perched atop their towering achievements, computer scientists (the cooks, remember?) will bask in the soothing certainty that their glorious science died at its peak. With a tinge of sadness but not a little pride, they'll chime in unison "There is nothing new to be discovered in computer science now."

If you think you've seen this movie before, you have. A few short years before Einstein turned our world upside down, the great Lord Kelvin bloviated this gem for the ages: "There is nothing new to be discovered in physics now." Not his lordship's finest hour.

Moore's Law has fueled computer science's sizzle and sparkle, but it may have obscured its uncanny resemblance to pre-Einstein physics: healthy and plump—and ripe for a revolution. Computing promises to be the most disruptive scientific paradigm since quantum mechanics. Unfortunately, it is the proverbial riddle wrapped in a mystery inside an enigma. The stakes are high, for our inability to "get" what computing is all about may well play iceberg to the Titanic of modern science.

Brilliant foresight or latest tripe from the Kelvin school of prophecy?

Computing is the meeting point of three Big Ideas: universality; duality; self-reference. To this triad, the modern view adds the concept of tractability and the revolutionary algorithmic paradigm. Here's how it works:

Universality Few would mistake your iPod for an IBM Blue Gene/L—the world's fastest computer. Yet, fundamentally, the two are the same. Why is that? At the heart of your iPod is a written document made of two parts: "program, data". The data section stores the songs as long sequences of 0s and 1s. The program section explains in words (again, 0s and 1s) how to read the data and turn it into sound. Add to this mix a smattering of hardware, the control, to read the program and follow its instructions, and voilà: you've got yourself an iPod. The beauty of the scheme is that the control need not know a thing about music. In fact, simply by downloading

the appropriate program/data document, you can turn your iPod into an earthquake simulator, a word processor, a web browser, or a paperweight. Your dainty little MP3 player is a *universal* computer.

Separating *control* (the hardware) from *program* (the software) was the major insight of Alan Turing—well, besides this little codebreaking thing he did in Bletchley Park that helped win World War II. The separation was the key to universality. No one had seen anything quite like it before. At least not since the Chinese philosopher opined: "Give a man a fish and you feed him for a day. Teach a man to fish and you feed him for a lifetime." In Confucius's hands, the specialized view of fishing = river + fisherman finds itself replaced by a universal one: fishing = river + fishing manual + you. There you have it, computing = data + program + control. The control part of your iPod is a marvel of electronics, but the shocker is that it need not be so: universal computers can be built with control boxes vastly simpler than a cuckoo clock. For all purposes, computing = data + program.

Duality Consider the iPod document "Print this, Let 'em eat cake". Push the start button and watch the words "Let 'em eat cake" flash across the screen. Note how the program part of the document, Print this, is interpreted as a command—printing is what it wants and printing is what it gets. Contrast this with the data part, Let 'em eat cake, which is treated as plain text: no one's eating anything (to Marie-Antoinette's later chagrin). Strings of 0s and 1s are interpreted in one of two ways: as form (data) or as content (program). Tapping into the comic, artistic, and academic potential of this duality, great minds went to work: Abbott and Costello ("Who's on First?"), Magritte ("Ceci n'est pas une pipe"), and Saussure ("signified vs. signifier"). Staring at the sublime will, of course, send the deeper thinkers among us rushing for the classics—such as Homer Simpson's immortal quip: "Oh Marge, cartoons don't have any deep meaning; they're just stupid drawings that give you a cheap laugh."

Self-Reference Write the iPod document "Print this twice, Print this twice" and press the start button. The screen lights up with the words: "Print this twice, Print this twice". Lo and behold, the thing prints itself! Just like a computer virus (remember, I did not teach you this). The magic word is "twice." For example, the iPod document "Print this, Print this" prints this: "Print this"—more Dr. Seuss than self-replication.

The "Big Ideas" were the air that the Gang of Four, Princeton branch,

breathed all day—that would be Alonzo Church, Alan Turing, Kurt Gödel, and John von Neumann. Mother Nature, of course, figured it all out a few billion years earlier. Reformat your genome by lining up the two strands of DNA one after the other, so it looks like a regular program-data iPod document (billions of letters long though):

"ACAAGAT...GCCATTG, TGTTCTA...CGGTAAC".

The base pairings (A,T) and (C,G) ensure that the two strands spell the same word with different letters. So, we lose no genomic information if we translate the "data" part and rewrite the whole document as the duplicated text

"ACAAGAT...GCCATTG, ACAAGAT...GCCATTG".

This is the biological analog of "*Print this twice, Print this twice*". Life's but a walking shadow, Macbeth warned us. Not quite. Life's but a self-printing iPod! Offended souls will bang on preachily about there being more to human life than the blind pursuit of self-replication—though Hollywood's typical fare would seem to refute that. Existential angst aside, *duality* is the option we have to interpret the word *ACAAGAT...GCCATTG* either as genes (the "form" encoding our genome) or as proteins (the "content" mediating the DNA replication). *Self-reference* is the duplication embodied in the base pairings. Viewed through the computing lens, life = duality + selfreference.

In the 1953 *Nature* article that unveiled to the world the structure of DNA, Watson and Crick signed off with this lovely understatement: *"It has not escaped our notice that the specific pairing we have postulated immediately suggests a possible copying mechanism for the genetic material.* "Duality and self-reference embedded in molecules: what sweet music to Turing's ears this must have been! Alas, our war hero was a little distracted at the time, busy as he was enjoying the rewards that the British authorities had lavished upon him for saving millions of lives during World War II—generous rewards like a court conviction for homosexuality with a sentence of forced estrogen injections. Almost one year to the day of Watson and Crick's triumph, Alan Turing went home, injected cyanide into an apple, ate it, and died. His mother preferred to believe it was an accident. **Tractability** The genesis of this fourth Big Idea was the ho-hum observation that checking the validity of a math proof tends to be much easier than finding the proof in the first place. But is it really? Amazingly, no one knows.

Welcome to the most important open question in all of computer science!

Ever wondered if your iPod's 5000-tune library is rich enough to let you compile a playlist of a thousand songs, no two which have ever been played back-to-back on MTV? Let's hope not, because not even an IBM Blue Gene/L could do the job in less time than has passed since dinosaurs were last seen roaming the earth. To find such a playlist (proof-finding) seems hopelessly hard, even on a computer, but to test whether a tentative playlist fits the bill (proof-checking) is a cinch: simply match all possible pairs against MTV's complete playlist, which is readily accessible on the web.

The twin reality of hard proof-finding and easy proof-checking is hardly an MTV-induced aberration. Computer scientists have catalogued over 1000 problems just like it. Of course, courtesy of Murphy's Law, these "Jurassic-1K" include all of the questions humanity is desperate to answer—in artificial intelligence, computational biology, resource allocation, rational drug design, etc.

OK, so life is tough. But since when has that observation qualified as a Big Idea?

Since 1970, roughly. Just as Einstein rebuilt Newtonian mechanics around the constancy of the speed of light, Cook, Edmonds, Karp, and Levin set out to rebuild computing around the notion of *tractability*. A problem is tractable if it can be solved in time growing polynomially in the input size, which is a fancy way of saying 'reasonably fast.' None of the Jurassic-1K appear to be tractable. At least those in the know believe they are not—of course, not so long ago, those in the know believed the earth was flat. Sadly, the great promise of computing seems to lie with problems afflicted with exponentialitis: the dreaded ailment that places even small-size problems beyond the reach of any computer.

This much we know: it's genetic. If a single one of the Jurassic-1K is tractable then, wonder of wonders, all of them are. These tough puzzles are nothing but different translations of the same Shakespeare play. Heady stuff! The day your playlist question can be answered in a few hours will be the day public-key cryptography dies, bringing down with it all of ecommerce. That day will see biology conquer its highest peak, protein folding, and mathematicians contemplate early retirement. Indeed, the day the Jurassic-1K are shown to be tractable (P=NP in computer parlance), proof-finding will be revealed to be no more difficult than proof-checking. Andrew Wiles, the conqueror of Fermat's Last Theorem, will be found to deserve no more credit than his referees. (Note that this says nothing about *understanding* the proof.) To be P or not to be P, that is NP's question. It is likely that P=NP would do for science what the discovery of the wheel did for land transportation. Little wonder no one believes it.

To discover the wheel is always nice, but to roll logs in the mud has its charms, too. Likewise, the intractability of proof-finding would have its benefits. When you purchase a book from Amazon, the assurance that your transaction is secure is predicated on more than your endearing naiveté. For one thing, it relies critically on the intractability of factoring a number into primes.

Just as modern physics shattered the platonic view of a reality amenable to noninvasive observation, tractability clobbers classical notions of knowledge, trust, persuasion, and belief. No less. For a taste of it, consider the great *zero-knowledge* (ZK) paradox: two mutually distrusting parties can convince each other that each one holds a particular piece of information without revealing a thing about it. Picture two filthy-rich businessmen stuck in an elevator. Their immediate goal is (what else?) finding out who's the wealthier. ZK dialogues provide them with the means to do so while revealing zero information about their own worth (material worth, that is—the other kind is already in full view).

Here is a ZK question for the State Department: can a signatory to the Nuclear Non-Proliferation Treaty demonstrate compliance without revealing any information whatsoever about its nuclear facilities? Just as game theory influenced the thinking of cold war strategists, don't be surprised to see ZK theory become the rage in international relations circles.

Tractability reaches far beyond the racetrack where computing competes for speed. It literally forces us to think differently. The agent of change is the ubiquitous *algorithm*.

The Algorithmic Revolution An algorithm is an iPod program with a human face. If a computer could wash your hair, its program would look like "0110001100100110..." but the algorithm behind it might read: "*Rinse, lather, repeat.*" (Don't try this at home if you're a computer scientist.) An algorithm is a list of instructions that tells the computer what to do. It may loop around and entertain alternatives, as in "*Rinse, lather, repeat if unhappy, dry, go to office, answer question: why didn't you rinse the shampoo off your hair?*" An algorithm is, in essence, a work of literature.

The library's bottom shelves might stack the one-line zingers—algorithmic miniatures that loop through a trivial algebraic calculation to produce *fractals* (pictures of dazzling beauty and infinite intricacy) or print the transcendentally mysterious digits of π . Algorithmic zingers can do everything. For the rest, we have the sonnets on the middle shelves. With names like FFT, RSA, LLL, AKS, they are short and crisp, and tend to pack more ingenuity per square inch than anything else in the computing world. The top shelves hold the lush, richly textured, multilayered novels.

Give it to them, algorithmic zingers know how to make a scientist swoon. No one who's ever tried to compute the digits of π by hand can remain unmoved by the sight of its decimal expansion flooding a computer screen like lava flowing down a volcano. And that's not even the awesome part. For that, one must turn to the infamous Brazilian butterfly whose evil wing flaps cause typhoons in China. Zingers embody the potential of a local action to unleash colossal forces on a global scale: complexity emerging out of triviality. Cellular automata, chaos theory, dynamical systems, and all that.

For a glimpse of the fiction genre on the top shelves, check out *PCP*. Suppose that, after popping the genius pill, you wake up one bright morning with a complete proof of the Riemann hypothesis in your head (that's the "Notorious B.I.G." of math rap: the biggest open problem in the field). Few number theorists are likely to listen to your story. That is, until you offer them the *PCP deal*. You'll write down your proof in an agreedupon format, and then let a "verifier" pick 10 lines at random. On the basis of these 10 lines alone, the verifier will decide whether your proof is correct. The shocker: beyond any reasonable doubt, she will be right! (Randomness plays a key role, but the chance of erring is less than that of the proverbial monkey typing all of Hamlet flawlessly.) The mind reels. If your proof is fine, then it will pass any test the verifier can throw at it. But, based on only 10 lines, how can she know that you've proven the Riemann hypothesis and not a baby cousin like 2+2=4? If your proof is bogus, the intuition does not help much either. Presumably, the agreed-upon format is designed to smear any bug all across the proof. But how will the verifier be sure that you didn't play fast and loose with the formatting rules? So many ways to cheat; so little evidence to check. The PCP algorithm upends basic notions of evidence and persuasion, and accomplishes what is usually philosophy's prerogative: to turn the comprehended into the incomprehensible. Somewhere, Wittgenstein must be smiling.

Moore's Law has put computing on the map. Algorithms will now unleash its true potential. Physics, astronomy, and chemistry are all sciences of *formulae*. Chaos theory moved the algorithmic zinger to centerstage. The quantitative sciences of the 21st century (eg, genomics, neurobiology) will complete the dethronement of the *formula* by placing the *algorithm* at the core of their modus operandi. Algorithmic thinking is likely to cause the most disruptive paradigm shift in the sciences since quantum mechanics. And yes, you may trust the future to be kind to this prediction.

Sensor Networks: A Digital Bridge to the Physical World





Smart Sensors and Lecture Outline Sensor Networks Part 1: Introduction to Wireless Sensor **Networks** Distributed monitoring applications; Sensor network hardware; Research issues in sensor networks; Naming and routing; Sensor tasking and control. Environmental sensing Traffic, habitats, pollution. Part 2: Structure Discovery and Information hazards, security Brokerage Industrial sensing Machine monitoring and Morphological analysis (boundaries, holes, bridges); diagnostics (IC fab) Power/telecom grid monitoring Landmarks and local coordinates; Information diffusion; Hierarchical hashing. Human-centered Part 3: Lightweight Spatio-Temporal Reasoning computing Configuration spaces; Collaboration groups; Identity Smart, human-aware spaces and environments Horst management; Occupancy tracking; Conclusion Stormer in Business Week, 8/23-30, 1999. 3 4





















Sensor Network Challenges

- Power management
 - communication 1000s of times more expensive than computation
 - load balancing across nodes
 coordinated sleeping/awake
 - schedules correlated sensor data
- In-network processing
 data aggregation
- overcounting of evidence
 Difficult calibration
 - localization
- time-synchronization
 Constant variability
 networking
 - sensing



[Picture from CACM June 2004]











In-Network Processing

- Information aggregation can happen on the way to a destination and provide substantial energy savings
- Need to balance quality of paths with quality of information collected
- But aggregation makes data lineage harder to ascertain
- Can we have "applicationindependent" paradigms of information aggregation?



Power-Aware Sensing, Computing, and Communication

- Variable power systems
 Let most sensors sleep most of the time; use paging channels
- Exploit correlation in readings between nearby sensors
- Load-balance, to avoid depleting critical nodes



26

28

Sensor Tasking and Control

- Decide which sensors should sense and communicate, according to the high-level task – a non-trivial algorithmic problem
- Direct sensing of relations relevant to the task – do not estimate full world state



27

Enable Data-Base Like Operations Data only available right after sensing operation Dense data streams must be sampled, or otherwise summarized

- Must deal with distributed information storage – "where is the data?"
- Large flash memory availability can make innetwork storage possible







- Resource constraints require close coupling between the networking application layers
- Can we define *application- independent* programming abstractions for sensor networks?

A sensor net stack?

4	User queries, external databases
ľ	In-network: application processing, data aggregation, query processing
к <u>с</u>	Data dissemination, storage, caching
ج ج	Adaptive topology, geo-routing
84) 	MAC, time and location services
1	Phy: comm, sensing, actuation, SP









Geographic Routing

- In sensor networks, naming and routing is frequently based on a node's attributes and sensed data, rather that on some pre-assigned network address.
- Geographic routing uses a node's location to name the node and discover paths to that node
- We assume that
 - nodes know their geographic location
 - nodes know their 1-hop neighbors
 - routing destinations are specified geographically (a point, a region)
 - each packet can hold a small amount (O(1)) of additional routing info to record where it has been on the network
 meet of the time we will medal the connectivity graph of the
 - most of the time we will model the connectivity graph of the nodes as a unit distance graph

Routing Desiderata

- Guaranteed delivery
- Path quality
- Energy awareness
- Robustness to low-level link volatility

























 Alternatively, we want to discover optimal paths without searching the whole connectivity graph G

٠

- If the optimal path between s and t has length L, then every node in that path is within an ellipse with foci s and t defined by L. This ellipse limits the part of G to be searched.
- If L is not known, it can be guessed, approximately



In general, finding a path of length L requires $O(L^2)$ work.















Collaborative Processing in Sensor Networks



- What information is critical for the high-level tasks?
- What is the cost of accessing the information?
- Which nodes should participate in sensing, processing, or communication?
- How should the information be migrated?
- What is routing or querying in this context?

62

66

The information processing needs largely determine the roles of nodes, as well as the required support by other layers of a sensor network





Structure Discovery and Information Brokerage in Sensor Networks



Structure Discovery

- A sensor network is a novel type of computing device -- a sensor computer
- One of its first tasks is to discover its own structure and establish
 - information highways
 sensor collaboration groups
- as well as adapt to its
- signal landscape



[From D. Estrin]

Talk Outline Information Brokerage Information providers Naming and Routing (sources, producers) and information seekers Landmarks and local coordinates (sinks, consumers) need ways to find out about + Hierarchical landmarks 0 Sensor Layout Analysis and rendez-vous with Boundary/hole detection each other Information Dissemination and Aggregation Challenges: Neither knows where the other is Highly dynamic environment Sweep Information Brokerage Hierarchical geographic hashing Limited computation and communication resources Double rulings Information gradients



Such Greedy Protocols Get Stuck, May Overload Critical Nodes



GOAFR+ [Kuhn, Wattenhofer, Zhang, and Zollinger, '03] These require building a planar subgraph of the connectivity graph - not a robust process

Global Embedding Challenges



- Routing on geographic coordinates
 - Only works in 2-D space
 - Planarization is tricky (CLDP, etc.)
 - Sensitive to location inaccuracy
 - Routing on virtual coordinates
 - Requires a global embedding of the link connectivity graph in the plane
 - Forcing a 2-D layout on a 3-D deployment may ignore much of the actual connectivity

Naming an Routing Based on **Connectivity Information Only**

- A two-tier approach utilizing combinatorial Delaunay complexes and local coordinates (GLIDER)
- A hierarchical approach using the Discrete Center Hierarchy (DCH)

I. Using Landmarks and Local Coordinates: GLIDER

- Given a communication graph on sensor nodes with distances defined by hop counts
- Perform structure discovery:
- Select a set of landmarks
- Construct the Landmark Voronoi Complex (LVC)
 Extract the Combinatorial Delaunay Triangulation (CDT) graph on landmarks



G is connected \Leftrightarrow CDT D(L) connected D(L) is compact –

- topology capture has complexity dependent on the number of largescale features in the environment
- D(L) is stable low- level link volatility unlikely to affect the combinatorial complex



be known to all landmarks, or even all nodes

Local Routing with Global Guidance

Global Guidance

the D(L) encodes global connectivity information that is accessible to every node for proactive route planning on tiles.

Local Routing

high-level routes on tiles are realized as actual paths in the network by using local reactive protocols.

Information Stored at Each Node

- The parents on the shortest paths to its home landmark, and its neighbor landmarks
- A bit to record if the node is on the boundary of a tile
- Its coordinates and those of its neighbors for intra-tile greedy routing
- Landmark nodes store the atlas D(L)



<section-header> GLIDER -- Routing Nouting Global route plan Local route intra-tile















- 437 -











HD Names and Routes Summary

- HD effectively discovers the intrinsic geometry of the network
- Provides a hierarchy-based scheme with provable approximation quality on the routing paths
- Node/link failures affect mostly the low levels of the hierarchy

Sensor Layout Analysis

- Boundary/hole detection
- Robust planarization of the communication graph



Information Dissemination and Aggregation

- Most sensor network applications need a robust and efficient implementation of certain basic data operations
- In such a data operations library one needs to include:
 - data dissemination (code images, parameter settings, etc.)
 - •in-network data aggregation

































- 442 -

Information Retrieval **Brokerage Costs** Query time avg std 1.08 0.66 Size # Info. servers std avg 9.95 • Consumer v looking for a 200 0.21 particular information 0.67 400 15.9 0.26 0.89 examines hash locations 600 25.4 1.02 0.77 4.00 of that information in 1000 31.5 0.83 0.57 larger and larger clusters 2000 47.0 0.73 0.47 11.3 0.73 0.42 4000 61.2 containing v 0.79 0.77 6000 69 5 13.6 0.46 8000 79.0 16.5 0.39 • Thm: The length of 10000 84.3 15.6 0.76 0.39 retrieval path from v is at most $4 \cdot d_{uv}$ where u is Storage cost grows slowly Cost of query is distance sensitive the producer (unknown to v Query time = Path length from consumer to hash location Shortest path length to producer



- HD effectively discovers the intrinsic geometry of the network
- Provides a hierarchy-based scheme with provable approximation quality on the routing paths
- Node/link failures affect mostly the low levels of the hierarchy
- Enables distance sensitive information brokerage

Resource Discovery Using Local Double Rulings

Associate with each node two connected `1-d' structures, call them roads – the red and blue

- All red roads together cover the network in a load-balanced fashion
- All blue roads together cover the network in a load-balanced fashion
- For any pair of nodes A and B, the red road from A has to intersect the blue road from B

A double ruling derived via a Morse function ≡ distance to boundary



Information providers and seekers can meet by following blue and red roads respectively





Double-Ruling Scheme in Hashed Tile



- Producers and consumers are guaranteed to meet by following the two sets of curves.
- The consumers may not need to reach the hashed tile to fetch the data as the data are available at some transit tiles

Reducing Producer Cost - En Route Data Aggregation



- Producers of the same content type share the shortest path tree (on CDT) rooted at the hashed tile.
- Data of the same type can be aggregated Inside the tile if two
- producers share one
- Inside the tile of their common ancestors



GLIDER Brokerage Summary

- Distance-sensitive information brokerage is possible with very modest data replication
- Information discovery is closely coupled with the network node naming and routing
- In some ways, geometric methods and tools can be effectively used even when the connectivity graph is all we got




<section-header> Diffusion can be slow to converge What if there are multiple sources with the same type of information? What if there are many different types of sources? What about discretization effects?

Dealing with Many Potential Sources: Bloom Filters for Membership Testing

- Given a set S = {x₁, x₂, x₃,...,x_n} on a universe U, want to answer queries of the form: Is y ε S?
- Example: a set of detection attributes
- Bloom filter provides an
- answer in • "Constant" time (time to hash).
- Small amount of space.
 But with some probability of space.
- But with some probability of being wrong.

- To check if y is in S, check B at B(y). All k values must be 1. B 0 1 0 0 1 0 1 0 1 0 0 1 1 1 0 1 1 0

Dealing with Many Potential Sources: Network Coding to Save Storage

Each node can compute a random linear combination of all the potentials it hears



- If there are *L* active potential sources, then a neighborhood of size *k* around a node provides Ω(*k*²) equations relating the local potential values
- But in a neighborhood of size k, the O(k) boundary values determine all the interior values for each potential (harmonic function property)
- So we have enough equations to recover the unknown potentials if k > L (k^2 constraints vs. kL unknowns)



Information Diffusion Summary

- Diffused information potentials can guide both virtual and physical information seekers to the appropriate sources
- Multiple sources can be handled by having their potentials co-mingled and then decoded as necessary
- Sources may move and the potentials adapt in a smooth manner

Conclusions

- Structure discovery and information brokerage are fundamental problems for WSNs
- With light preprocessing we can extract certain global quantities that can significantly help with local decisions
- These quantities reflect an understanding of the geometry or topology of the sensor filed and do not require localization
- The same quantities are also robust to local volatility in the network connectivity
- Such approaches integrate very well with current `narrow waist' sensor net protocols, such as SP (Berkeley)

Lightweight Spatio-Temporal Reasoning in Sensor Networks



<section-header> More Demanding Sensor Detwork Applications Pervent Applications







Some Lessons and Issues

- An appropriately chosen configuration space can transform a wide-area phenomenon into a localized one.
- Only a small fraction of the nodes in the field need be active at any one time.
- Most message traffic is along the shadow boundary: the physical phenomenon dictates the communication paths.





Goal: count and track the significant signal peaks in the field

- Signal attenuation rate, and the spacing and communication range of sensors have big impact on "signal resolution"
- Number of detected peaks may not equal number of targets due to sampling artifacts and/or noise.









2-D View of a Sensor Field with Cluster Trees Formed Using DFP

















Renormalization, Given Local Evidence



What Do We Want?

- The belief matrix represents a probability distribution. The matrix *A* represents our *a priori* belief, but violates sum constraints.
- We would like to find the sum-constrained (feasible) matrix *B* that is the *closest distribution* to *A* (which is infeasible).
- Use Kullback-Leibler distance (measure of distance between distributions):

$$KL(B:A) = \sum_{j=1}^{n} \sum_{i=1}^{m} b_{ij} \log \frac{b_{ij}}{a_{ij}}$$



 $KL(B:A) = \sum_{j=1}^{n} \sum_{i=1}^{m} b_{ij} \log \frac{b_{ij}}{a_{ij}}$

Theorem: Given a prior matrix $A \in \mathbb{R}^{m \times n}$, the matrix *B* that satisfies the row and column sum constraints, and minimizes the KL-distance from the prior matrix *A* is **always** the solution of the Sinkhorn scaling process. [Balakrishnan, Hwang,Tomlin '04]













Approach B: Information Matrix

- $L = \begin{bmatrix} l_{11} & \cdots & l_{1j} & \cdots & l_{1N} \\ l_{21} & \cdots & l_{2j} & \cdots & l_{2N} \\ & & & & \\ l_{N1} & \cdots & l_{Nj} & \cdots & l_{NN} \end{bmatrix}$ Keep unnormalized beliefs by simply adding log-likelihoods after each local $l_{ij} = \sum \log(p(Z^t = (k,j) | x_j = i), k \in \{1,\cdots,N\}$ evidence event
- No communication necessary, except when mixing events occur or queries are made





- Collaboration groups need not always be physically clustered.
- Different attributes of a phenomenon can be tracked at different rates (target location, identity).
- A change of information representation can have a deep impact on cost trade-offs.
- How do information providers and information seekers locate each other?

V4. Image Sensor Networks

- CMOS technology enables the production of small, lowcost and low-power integrated image sensors
- Cameras (still or video) and other image sensors are becoming cheaper, smaller, and nearly ubiquitous
- However, truly distributed networked systems of image sensors are still not here





Current Multi-Imager Networks

- Data is transported over a wired network to a central location
- Human operators look at the data





automatic ways to filter the data are needed

Distributed Imager Challenges

- Imagers are high data rate sensors; therefore data must be compressed and summarized
 compression must take into
 - account shared data
- goal of compression need not be reconstruction
 Vision algorithms can be
- vision algorithms can be expensive to run on weak capability, low power devices
- Visibility is non-local and discontinuous (occlusions, etc)
- discontinuous (occlusions, etc
- Issues of privacy, etc.



Collaborative, Task-Driven Image Sensing

- Large numbers of simple, inexpensive cameras collaborate over a wireless network to accomplish a task
- Data is compressed locally and aggregated within the network
- Cameras are only tasked as the situation demands
- The system can be expanded incrementally to large numbers of nodes



The goal is to estimate certain high-level, global attributes of the environment.

The Initial Effort

- Use a camera network to obtain information about space occupancy by people.
- Useful for aggregate tracking, counting, etc.
- Crowd density implies multiple occlusions – no one camera by itself can do this.
- No image reconstruction --just high-level distributed spatial reasoning.



Packard 013





- Web cameras:
 - 16 firewire webcams with 49 degree FOV
 - Placed around a 22 x 19 foot room
- Linux computers
 - A PC is connected to 2 webcams
 - A separate process is running for each webcam to simulate an individual camera node
 - All processes can communicate with each other over the network

































The End

Another kind of sensor network exhibiting lightweight spatiotemporal reasoning





Part I

what is a game?
Pure and randomized equilibria
Load balancing and routing as games







Games: setup

Deterministic (pure) or randomized (mixed) strategies?

Pure: each player selects a strategy. simple, natural, but stable solution may not exists

Mixed: each player chooses a probability distribution of strategies.

- equilibrium exists (Nash),
- but pure strategies often make more sense

















Model of Routing Game

- A directed graph G = (V,E)
- source-sink pairs s_i,t_i for i=1,...,k
- rate r_i ≥ 0 of traffic between s_i and t_i for each i=1,...,k



- Load-balancing jobs wanted min load
- Here want minimum delay: delay adds along path edge-delay is a function $\ell_e(\cdot)$ of the load on the edge e



























Sharper results for non-atomic games

Theorem 2 (Roughgarden'02):

- In any network with any class of convex continuous latency functions
- the worst price of anarchy is always on two edge network



Sharper results for non-atomic games

Theorem 2 (Roughgarden'02):

- In any network with any class of convex continuous latency functions
- the worst price of anarchy is always on two edge network



Corollary: price of anarchy for degree d polynomials is O(d/log d).

36













Morale for IP versus ATM?

Corollary: with M/M/1 delay fns: l(x)=1/(u-x), where u=capacity

Nash w/cap. $2u \le opt w/cap. u$

Doubling capacity is more effective than optimized routing (IP versus ATM)



44

- Discrete potential games:
- network design
- price of anarchy stability











Discrete potential game: there is a function $\Phi(f)$ so that change in potential is same as function change perceived by one user

Theorem [Monderer Shapley'96] Discrete potential games if and only if congestion game (cost of using an element depends on the number of users). Proof of "if" direction $\Phi(\mathbf{f}) = \Sigma_e (\mathcal{L}_e(1) + ... + \mathcal{L}_e(\mathbf{f}_e))$

Corollary: Nash equilibria are local min. of $\Phi(f)$

49

Why care about

Best Nash/Opt ratio?

50

Papadimitriou-Koutsoupias '99 Nash = outcome of selfish behavior

⇒worst Nash/Opt ratio: Price of Anarchy

Non-atomic game: Nash is unique... Atomic Nash not unique!









Network Design as Potential Game

Given: G = (V,E), costs $c_e(x)$ for all $e \in E$, k terminal sets (colors)

Have a player for each color.

Each player wants to build a network in which his nodes are connected.

Player strategy: select a tree connecting his set.











































Why stable solutions?

Plan: analyze the quality of Nash equilibrium. But will players find an equilibrium?

- Can a stable solution be found in poly. time?
- Does natural game play lead to an equilibrium?
- We are assuming non-cooperative players, what if there is cooperation?

Why stable solutions? Plan: analyze the quality of Nash equilibrium. But will players find an equilibrium? • Can a stable solution be found in poly. time? • Does natural game play lead to an equilibrium? • We are assuming non-cooperative players, what if there is cooperation?

Answer 1: A clean solution concept and exists ([Nash 1952] if game finite) Does life lead to clan solutions?

78

Why stable solutions?

• Finding an equilibrium?

Nonatomic games: we'll see that equilibrium can be found via convex optimization [Beckmann'56]

Atomic game: finding an equilibrium is polynomial local search (PLS) complete [Fabrikant, Papadimitriou, Talwar STOC'04]

79

Why stable solutions? Does natural game play lead to equilibrium? we'll see that natural "best response play" leads to equilibrium if players change one at-a-time

See also:

Fischer¥Räcke¥Vöcking'06, Blum¥Even-Dar¥Ligett'06 also if players simultaneously play natural learning strategies

































Algorithmic Game Theory

- The main ingredients:
 - Lack of central control like distributed computing
 - Selfish participants game theory
- Common in many settings e.g., Internet

Most results so far:

- Price of anarchy/stability in many games, including many I did not mention
- e.g. Facility location (another potential game) [Vetta FOCS'02] and [Devanur-Garg-Khandekar-Pandit-Saberi'04]:

99

Some Open Directions:

- Other natural network games with low lost of anarchy
- Design games with low cost of anarchy
- Better understand dynamics of natural game play
- Dynamics of forming coalitions

Polynomial Time Algorithms For Market Equilibria

Vijay V. Vazirani



























 The study of market equilibria has occupied center stage within Mathematical Economics for over a century.



This talk: Historical perspective& key notions from this theory.

2). Algorithmic Game Theory

 Combinatorial algorithms for traditional market models

3). New Market Models

Resource Allocation Model of Kelly, 1997

3). New Market Models

- Resource Allocation Model of Kelly, 1997
- For mathematically modeling TCP congestion control
- Highly successful theory

















Arrow-Debreu Model, 1954 Exchange Economy

- Second fundamental market model
- Celebrated theorem in Mathematical Economics



























Arrow-Debreu Theorem

 For continuous, monotonic, strictly concave utility functions, market clearing prices exist.
















• Let
$$f(p) = p'$$

• $S(j) < d(j) \Rightarrow p'(j) = \frac{p(j) + [d(j) - s(j)]}{N}$
• $S(j) > d(j) \Rightarrow p'(j) = \frac{p(j)}{N}$
• N is s.t. $\sum_{j} p'(j) = 1$

2		
	$\forall i : u_i$	is a cts. fn.
=>	$\forall i : B(i)$	is a cts. fn. of <u>p</u>
=>	$\forall j : d(j)$	is a cts. fn. of <u>p</u>
=>	f is a cts. fn.	of <u>p</u>









Combinatorial Algorithms for Market Equilibria

Vijay V. Vazirani

Arrow-Debreu Theorem: Equilibria exist.

Arrow-Debreu Theorem: Equilibria exist.
Do markets operate at equilibria?

- Arrow-Debreu Theorem: Equilibria exist.
- Do markets operate at equilibria?
- Can equilibria be computed efficiently?

Arrow-Debreu is highly non-constructive

Arrow-Debreu is highly non-constructive

- "Invisible hand" of the market: Adam Smith
- Scarf, 1973: approximate fixed point algs.
- Convex programs:
 Fisher: Eisenberg & Gale, 1957
 Arrow-Debreu: Newman and Primak, 1992



Algorithmic Game Theory

- Use powerful techniques from modern algorithmic theory and notions from game theory to address issues raised by Internet.
- Combinatorial algorithms for finding market equilibria.

Two Fundamental Models

Fisher's model

 Arrow-Debreu model, also known as exchange model

Combinatorial Algorithms

- Primal-dual schema based algorithms
 Devanur, Papadimitriou, Saberi & V., 2002 Combinatorial algorithm for Fisher's model
- Auction-based algorithms
 Garg & Kapoor, 2004
 Approximation algorithms.

Approximation

- Find prices s.t. all goods clear
- Each buyer get goods providing at least (1-ε)×optimal utility.

Primal-Dual Schema

 Highly successful algorithm design technique from exact and approximation algorithms

Exact Algorithms for Cornerstone Problems in P:

- Matching (general graph)
- Network flow
- Shortest paths
- Minimum spanning tree
- Minimum branching

Approximation Algorithms

set cover Steiner tree Steiner network *k*-MST scheduling ... facility location *k*-median multicut feedback vertex set

Main new idea

 Previous: problems captured via linear programs

DPSV algorithm: problem captured via a nonlinear convex program

Fisher's Model

- *n* buyers, with specified money, m(i) for buyer *i*
- *k* goods (unit amount of each good)
- Linear utilities: *u_{ij}* is utility derived by *i* on obtaining one unit of *j*
- Total utility of *i*,

$$u_i = \sum_j u_{ij} x_{ij}$$
$$\chi_{ij} \in [0,1]$$

Fisher's Model

- *n* buyers, with specified money, *m(i)*
- *k* goods (each unit amount, w.l.o.g.)
- Linear utilities: *u_{ij}* is utility derived by *i* on obtaining one unit of *j*
- Total utility of *i*,

$$u_i = \sum_j u_{ij} x_i$$

Find prices s.t. market clears

Eisenberg-Gale Program, 1959

$$\max_{i} \sum_{i} m(i) \log u(i)$$

s.t.
$$\forall i : u(i) = \sum_{j} u_{ij} x_{ij}$$

$$\forall j : \sum_{i} x_{ij} \le 1$$

$$\forall ij : x_{ij} \ge 0$$













































buyers <u>m</u>	equality subgraph	goods <u>p</u> ensure Invariant	

<u>m</u>		<u>p</u> x			
$x = 1, x \uparrow$					





$\Gamma(S)$		S	frozen
		<u>p</u> x	active
	$x = 1, x \uparrow$		

$\Gamma(S)$		S	frozen
		<u>p</u> x	active
ſ	$x = 1, x \uparrow$	1	





































total surplus decreases
 flow becomes more balanced





Weak gross substitutability

 Increasing price of one good cannot decrease demand for another good.

Weak gross substitutability

- Increasing price of one good cannot decrease demand for another good.
- => never need to decrease prices (dual variables).

Weak gross substitutability

- Increasing price of one good cannot decrease demand for another good.
- => never need to decrease prices (dual variables).
- Almost all primal-dual algs work this way.

Arrow-Debreu Model

- Approximate equilibrium algorithms:
 - □ Jain, Mahdian & Saberi, 2003: Use DPSV as black box.
 - Devanur & V., 2003: More efficient, by opening DPSV.







Arrow-Debreu Model

- Start with all prices 1
- Allocate money to agents (initial endowment)
- Perform outbid and update agents' money
- Any good with price >1 is fully sold

Arrow-Debreu Model

- Start with all prices 1
- Allocate money to agents (initial endowment)
- Perform outbid and update agents' money
- Any good with price >1 is fully sold

 $\frac{\max price}{\min price} \le \frac{\max u_{ij}}{\min u_{ij}}$

• Eventually every good will have price >1

■ Garg, Kapoor & V., 2004:

.

Auction-based algorithms for additively separable concave utilities satisfying weak gross substitutability ■ Kapoor, Mehta & V., 2005:

Auction-based algorithm for a (restricted) production model

Q: Distributed algorithm for equilibria?

- Appropriate model?
- Primal-dual schema operates via local improvements

New Market Models Resource Allocation Markets

Vijay V. Vazirani

Fisher's Model

- *n* buyers, with specified money, m(i) for buyer *i*
- *k* goods (unit amount of each good)
- Linear utilities: *u_{ij}* is utility derived by *i* on obtaining one unit of *j*
- Total utility of *i*,

$$u_i = \sum_j u_{ij} x_{ij}$$
$$\chi_{ij} \in [0,1]$$

Fisher's Model

- *n* buyers, with specified money, *m(i)*
- *k* goods (each unit amount, w.l.o.g.)
- Linear utilities: *u_{ij}* is utility derived by *i* on obtaining one unit of *j*
- Total utility of *i*,

$$u_i = \sum_j u_{ij} x_{ij}$$

Find prices s.t. market clears

Eisenberg-Gale Program, 1959

$$\max_{i} \sum_{i} m(i) \log u(i)$$

s.t.
$$\forall i : u(i) = \sum_{j} u_{ij} x_{ij}$$

$$\forall j : \sum_{i} x_{ij} \le 1$$

$$\forall ij : x_{ij} \ge 0$$

Via KKT Conditions can establish:

- Optimal solution gives equilibrium allocations
- Lagrange variables give prices of goods

Eisenberg-Gale program helps establish:

- Equilibrium exists (under mild conditions)
- Equilibrium utilities and prices are unique

Eisenberg-Gale program helps establish:

- Equilibrium exists (under mild conditions)
- Equilibrium utilities and prices are unique
- Rational!!

Kelly's resource allocation model, 1997

Mathematical framework for understanding TCP congestion control











TCP Congestion Control

■ *f(i):* source rate

p(e): prob. of packet loss (in TCP Reno) queueing delay (in TCP Vegas)

TCP Congestion Control

- *f(i):* source rate
- *p(e)*: prob. of packet loss (in TCP Reno) queueing delay (in TCP Vegas)
- Kelly: Equilibrium flows are proportionally fair: only way of increasing an agent's flow by 5% is to decrease other agents' flow by at least 5%

TCP Congestion Control

- *f(i)*: source rate
- *p(e):* prob. of packet loss (in TCP Reno) queueing delay (in TCP Vegas)
- Low, Doyle, Paganini: continuous time algs. for computing equilibria (not poly time).

TCP Congestion Control

- f(i): source rate
- prob. of packet loss (in TCP Reno) queueing delay (in TCP Vegas)
- Low, Doyle, Paganini: continuous time algs. for computing equilibria (not poly time).
- AIMD + RED converges to equilibrium primal-dual (source-link) alg.





- Devanur, Papadimitriou, Saberi & V., 2002: for Fisher's linear utilities case
- Kelly & V., 2002: Kelly's model is a generalization of Fisher's model.

Find comb. poly time algs!



























































Eisenberg-Gale-Type Convex Program $\max \sum_{i} m(i) \log u(i)$ s.t. packing constraints

Eisenberg-Gale Market

 A market whose equilibrium is captured as an optimal solution to an Eisenberg-Gale-type program

• Megiddo, 1974: Let T = set of sinks (agents)

- For $S \subseteq T$ define v(S) to be the max-flow possible from *s* to sinks in *S*.
- Then v is a submodular function, i.e., for

$$A \subseteq B \subseteq T,$$

$$t \notin A,$$

$$v(B+t) - v(B) \le v(A+t) - v(A)$$

Simpler convex program for single-source market

$$\max_{i} \sum_{i} m(i) \log f(i)$$

s.t.
$$\forall S \subseteq T : \sum_{i \in S} f(i) \le v(S)$$

$$\forall i : f(i) \ge 0$$



 Any market which has simpler program and v is submodular

Submodular Utility Allocation Market

- Any market which has simpler program and v is submodular
- Theorem: Strongly polynomial algorithm for SUA markets.

Submodular Utility Allocation Market

- Any market which has simpler program and v is submodular
- Theorem: Strongly polynomial algorithm for SUA markets.
- Corollary: Rational!!

Theorem: Following markets are SUA: 2 source-sink pairs, undirected (Hu, 1963) spanning tree (Nash-William & Tutte, 1961) 2 sources branching (Edmonds, 1967 + JV, 2005)

■ 3 sources branching: irrational

- Theorem: Following markets are SUA:
 2 source-sink pairs, undirected (Hu, 1963)
 spanning tree (Nash-William & Tutte, 1961)
 2 sources branching (Edmonds, 1967 + JV, 2005)
- 3 sources branching: irrational
- Open (no max-min thoerems):
 2 source-sink pairs, directed
 - □2 sources, network coding

Theorem: Following markets are SUA: 2 source-sink pairs, undirected (Hu, 1963) spanning tree (Nash-William & Tutte, 1961) 2 sources branching (Edmonds, 1967 + JV)

- 3 sources branching: irrational
- Open (no max-min thoerems):
 2 source-sink pairs, directed
 2 sources, network coding
 - Chakrabarty, Devanur & V., 2006





• Theorem: Strongly poly alg for Comb EG[2].



Other properties:

- Efficiency
- Fairness (max-min + min-max fair)
- Competition monotonicity

Open issues

- Strongly poly algs for approximating

 nonlinear convex programs
 equilibria
- Insights into congestion control protocols?

Random Sampling Techniques and Approximation of CSP Problems

MAREK KARPINSKI

UNIVERSITY OF BONN

(NHC Spring School Lectures, Tokyo, March 1, 2006)

Abstract. We present some recent results and new sampling techniques for absolute and relative approximation of general Constraint Satisfaction Problems (CSP). The methods used are threefold and based on: Smooth or Linearized Integer Programs, combinatorial arguments, and special linear algebraic techniques. In particular we apply those techniques to construct polynomial time approximation schemes (PTASs) for certain instances of both MAX- and MIN-CSP including dense and subdense instances and general metric and quasimetric instances of those problems. In that context we study the generic sample complexity for approximating arbitrary CSP instances and try to establish tight upper bounds for their underlying core-sample sizes. We go also beyond CSP optimization problems and design first PTASs for general metric and quasimetric size-constraint Partitioning Problems.

Selected References

- N. Alon, W. F. de la Vega, R. Kannan and M. Karpinski, *Random Sampling and Approximation of MAX-CSPs*, J. Computer and System Sciences 67 (2003), 212-243.
- S. Arora, D. Karger and M. Karpinski, Polynomial Time Approximation Schemes for Dense Instances of NP-Hard Problems, J. Computer and System Sciences 58 (1999), 193-210.
- W. F. de la Vega, R. Kannan, M. Karpinski and S. Vempala, Tensor Decomposition and Approximation Schemes for Constraint Satisfaction Problems, Proc. 27th ACM STOC (2005), 747-754.
- [4] W. F. de la Vega and M. Karpinski, *Approximation Complexity of Nondense Instances of MAX-CUT*, in preparation, 2006.
- [5] W. F. de la Vega, M. Karpinski and C. Kenyon, Approximation Schemes for Metric Bisection and Partitioning, Proc. 15th ACM-SIAM SODA (2004), 506-515.
- [6] W. F. de la Vega, M. Karpinski, C. Kenyon and Y. Rabani, *Approximation Schemes for Clustering Problems*, Proc. 35th ACM STOC (2003), 50-58.
- [7] M. Karpinski, *Polynomial Time Approximation Schemes* for Some Dense Instances of NP-Hard Problems, Algorithmica 30 (2001), 386-397.

APPROXIMATION SCHEMES FOR METRIC BISECTION AND PARTITIONING

MAREK KARPINSKI UNIVERSITY OF BONN

APPROXIMATING METRIC BISECTION AND RELATED PARTITIONING PROBLEMS

MAREK KARPINSKI UNIVERSITY OF BONN

(Joint work with W.F. DE LA VEGA and Claire Kenyon)

1



0













5







$$\left(\begin{array}{l} W_L = \sum_{L \rtimes L} d(x,y) \ , \\ W_R = \sum_{R \rtimes R} d(x,y) \ . \end{array} \right.$$















$$\begin{cases} A PTAS FOR GENERAL \\ METRIC BISECTION: \\ (VIA LINEARIZED QUADRATIC PROGRAMS) \\ \end{cases}$$
$$\begin{cases} P_{(v,d)} = \frac{1}{2} \sum_{i \neq j} w_{ij} (x_i (1-x_i) + x_j (1-x_i)), \\ \sum_i x_i = \frac{n}{2} \end{cases}$$






























$$\left(\begin{array}{c} \underline{RUNNING TIME}: \underline{LP(n) \cdot 2^{o(\frac{1}{c_{2}})}}\\ (\underline{COMBINATORIAL VERSION}: \underline{n^{2} \cdot 2^{o(\frac{1}{c_{2}})}}\\ (\underline{EXTENSION TO}\\ \underline{METRIC (n_{1}, n_{2}, ..., n_{k}) MIN-PARTITIONING}\\ (\underline{EXISTENCE OF PTASs}) \end{array}\right)$$



IF d(x,y) IS METRIC AND $\alpha > 0$, THEN $(d(x,y))^{\alpha}$ IS A QUASIMETRIC.

[FURTHER RESEARCH: 1. DESIGN THE PTASS (?) FOR "BALANCED" METRIC (n, n, n,) CLUSTERING (EXTENSION OF [FKKR03]-TECHNIQUE ?) 2. IMPROVE EFFICIENCY





28

Approximation Algorithms for Facility Location

Jens Vygen

University of Bonn

Outline

Introduction

Uncapacitated Facility Location

Capacitated and Universal Facility Location

Facility Location and Network Design with Service Capacities

Facility Location: Applications

- manufacturing plants
- storage facilities, depots
- warehouses, retail stores
- libraries, fire stations, hospitals
- servers in the internet
- base stations for wireless services
- buffers distributing signals on a chip
- ► ...

Goal: Optimum service for clients at minimum cost

Common features of facility location problems

- ► Two sets: clients and potential facilities
 - Each client must be served.
 - A potential facility can be opened or not.
 - Clients can only be served by open facilities.
- ▶ Two cost components: facility cost and service cost.
 - Opening a facility involves a certain cost.
 - Serving a client from a facility involves a certain cost.
- ▶ The total cost is to be minimized.

But there are many variants

- Can a client's demand be satisfied by more than one facility?
- Are there constraints on the total demand, or total service cost, that a facility can handle?
- Do the service costs satisfy the triangle inequality?
- Are there finitely or infinitely many potential facilities?
- > Do the facility costs depend on the total demand served?
- Is it allowed to serve only a subset of clients, and pay for those that are not served?
- Is there a bound on the number of facilities that we can open?
- Does the total service cost of a facility depend on the sum of the distances to its clients, or the length of a shortest tour, or the length of an optimal Steiner tree?
- Are we interested in the sum of all service costs, or rather in the maximum service cost?
- ▶ Do we need to serve facilities by second-stage facilities (etc.)?

Example 1: Fermat-Weber Problem

The most prominent example for continuous facility location

Locating a single facility in \mathbb{R}^n : Given $a_1, \ldots, a_m \in \mathbb{R}^n$ and weights $w_1, \ldots, w_m \in \mathbb{R}_+$, find $p \in \mathbb{R}^n$ minimizing

$$\sum_{i=1}^m w_i ||p-a_i||.$$

- For ℓ_1 -norm solvable in linear time (Blum et al. 1973)
- ▶ ℓ₂-norm, n = 2, m = 3: Simple geometric solution (Fermat, Torricelli, Cavalieri, Simpson, Heinen)
- \blacktriangleright For $\ell_2\text{-norm:}$ construction by ruler and compasses impossible (Bajaj 1988)
- Approximate solution for l₂-norm: Weiszfeld's algorithm (Weiszfeld 1937, Kuhn 1973, Vardi and Zhang 2001, Rautenbach et al. 2004)



Instance:

- a finite set D of clients;
- ▶ a finite set *F* of potential facilities;
- ▶ a fixed cost $f_i \in \mathbb{R}_+$ for opening each facility $i \in \mathcal{F}$;
- ▶ a service cost $c_{ij} \in \mathbb{R}_+$ for each $i \in \mathcal{F}$ and $j \in \mathcal{D}$.

We look for:

- ▶ a subset S of facilities (called open) and
- \blacktriangleright an assignment $\sigma: \mathcal{D} \rightarrow \mathcal{S}$ of clients to open facilities,
- such that the sum of facility costs and service costs

$$\sum_{i\in S} f_i + \sum_{j\in \mathcal{D}} c_{\sigma(j)j}$$

is minimum.

More examples discussed later

- Capacitated Facility Location
- Universal Facility Location
- ► Facility Location and Network Design with Service Capacities

These are more general and more realistic in many applications.

Approximation Algorithms: Definition

Let f be a function assigning a real number to each instance. An f-approximation algorithm is an algorithm for which a polynomial p exists such that for each instance I:

- the algorithm terminates after at most p(size(1)) steps,
- the algorithm computes a feasible solution, and
- ► the cost of this solution is at most f(I) times the optimum cost of instance I.

f is called the **approximation ratio** or **performance guarantee**. If f is a constant, we have a **(constant-factor) approximation algorithm**.

Uncapacitated Facility Location is as hard as Set Covering

SET COVERING: Given a finite set U, a family S of subsets of Uwith $\bigcup_{S \in S} S = U$, and weights $w : S \to \mathbb{R}_+$, find a set $\mathcal{R} \subseteq S$ with $\bigcup_{R \in \mathcal{R}} R = U$ with minimum total weight $\sum_{R \in \mathcal{R}} w(R)$.

- No o(log |U|)-approximation algorithm exists unless P = NP. (Raz, Safra 1997)
- ► Greedy algorithm has performance ratio 1 + ln |U|. (Chvátal 1979)
- ▶ SET COVERING is a special case of UNCAPACITATED FACILITY LOCATION: define $\mathcal{D} := U$, $\mathcal{F} := S$, $f_S = w(S)$ for $S \in S$, $c_{Sj} := 0$ for $j \in S \in S$ and $c_{Sj} := \infty$ for $j \in U \setminus S$.
- Conversely, the greedy algorithm for SET COVERING can be applied to UNCAPACITATED FACILITY LOCATION:
 Set U := D, S = F × 2^D, and w(i, D) := f_i + ∑_{j∈D} c_{ij}. (Hochbaum 1982)

Therefore we assume henceforth **metric service costs**: $c_{ij} \ge 0$ and $c_{ij} + c_{i'j} + c_{i'j'} \ge c_{ij'}$ for all $i, i' \in \mathcal{F}$ and $j, j' \in \mathcal{D}$. Equivalently, we assume c to be a (semi)metric on $\mathcal{D} \cup \mathcal{F}$.

A natural assumption: metric service costs

Motivation:

- ▶ The general problem is as hard as SET COVERING.
- In many practical problems service costs are proportional to geometric distances, or to travel times, and hence are metric.

```
But: Greedy algorithm has performance guarantee \Omega(\log n / \log \log n) even for metric instances. (Jain et al. 2003)
```

Integer Linear Prog	rammiı	ng Fo	rmulatio	n
minimize $\sum f_i$	$y_i + \sum$	$\sum c_{ij}$	(ij	
$i \in \mathcal{F}$	$\overline{i\in\mathcal{F}}$	$\overline{j\in D}$		
subject to	¥	<	<i>V</i> :	(ic FicD)
	×ij	\geq	yî	$(I \in J, J \in D)$
	$\sum x_{ii}$	=	1	$(i \in \mathcal{D})$
	$i \in F$			0)
	Xii	\in	$\{0, 1\}$	$(i \in \mathcal{F}, j \in \mathcal{D})$
	5		((,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,
	Уi	\in	$\{0,1\}$	$(i \in \mathcal{F})$
(Balinski 1065)				





First Approximation Algorithm: LP Rounding Analysis of the LP Rounding Approximation Algorithm • Compute an optimum solutions (x^*, y^*) and (v^*, w^*) to the primal and dual LP. • By complementary slackness, $x_{ii}^* > 0$ implies $v_i^* - w_{ii}^* = c_{ij}$, and thus $c_{ij} \leq v_i^*$. ▶ Let *G* be the bipartite graph with vertex set $\mathcal{F} \cup \mathcal{D}$ containing an edge $\{i, j\}$ iff $x_{ii}^* > 0$. Assign clients to clusters iteratively as follows. • In iteration k, let j_k be a client $j \in \mathcal{D}$ not assigned yet and with v_i^* smallest. • Create a new cluster containing j_k and those vertices of G that have distance 2 from j_k and are not assigned yet. Continue until all clients are assigned to clusters. For each cluster k we choose a neighbour i_k of j_k with f_{i_k} minimum, open i_k , and assign all clients in this cluster to i_k . Theorem This is a 4-approximation algorithm for metric UFL. (Shmoys, Tardos and Aardal 1997)

I he service cost for client j in cluster k is at most
$c_{i_k j} ~\leq~ c_{i j} + c_{i j k} + c_{i_k j_k} ~\leq~ {\sf v}_j^* + 2 {\sf v}_{j_k}^* ~\leq~ {\sf 3} {\sf v}_j^*,$
where <i>i</i> is a common neighbour of <i>j</i> and <i>j_k</i> . ► The facility cost <i>f_{ik}</i> can be bounded by
$f_{i_k} \leq \sum_{i \in \mathcal{F}} x^*_{ij_k} f_i = \sum_{i \in \mathcal{F}: \{i, j_k\} \in E(\mathcal{G})} x^*_{ij_k} f_i \leq \sum_{i \in \mathcal{F}: \{i, j_k\} \in E(\mathcal{G})} y^*_i f_i.$
As j_k and $j_{k'}$ cannot have a common neighbour for $k \neq k'$, the total facility cost is at most $\sum_{i \in \mathcal{F}} y_i^* f_i$.
The total cost is at most
$3\sum_{j\in\mathcal{D}}v_j^*+\sum_{i\in\mathcal{F}}y_i^*f_i,$
which is at most four times the LP value. Hence we get:
고

technique	ratio	RT	authors	year
LP-Rounding	3.16	-	Shmoys, Tardos, Aardal	1997
LP-Rounding+Greedy	2.41	-	Guha, Khuller	1998
LP-Rounding	1.74	-	Chudak	1998
Local Search	5.01	0	Korupolu, Plaxton, Ra-	1998
			jaraman	
Primal-Dual	3.00	+	Jain, Vazirani	1999
Primal-Dual+Greedy	1.86	+	Charikar, Guha	1999
LP-Rounding+Primal-	1.73	-	Charikar, Guha	1999
Dual+Greedy				
Local Search	2.42	0	Arya et al.	200
Primal-Dual	1.61	+	Jain, Mahdian, Saberi	2002
LP-Rounding	1.59	-	Sviridenko	2002
Primal-Dual+Greedy	1.52	+	Mahdian, Ye, Zhang	2002

Primal-Dual Algorithm by Jain, Mahdian and Saberi (2002)

Start with U := D and time t = 0. Increase t, maintaining $v_i = t$ for all $j \in U$. Consider the following events:

- ▶ $v_j = c_{ij}$, where $j \in U$ and i is not open. Then start to increase w_{ij} at the same rate, in order to maintain $v_j w_{ij} = c_{ij}$.
- $\sum_{i \in D} w_{ij} = f_i$. Then open *i*. For all $j \in D$ with $w_{ij} > 0$: freeze v_j and set $w_{i'j} := \max\{0, c_{ij} - c_{i'j}\}$ for all $i' \in \mathcal{F}$, and remove j from U.
- ▶ $v_j = c_{ij}$, where $j \in U$ and i is open. Then freeze v_j and set $w_{i'j} := \max\{0, c_{ij} - c_{i'j}\}$ for all $i' \in \mathcal{F}$, and remove j from U.

Improvement by Mahdian, Ye and Zhang (2002)

- Multiply all facility costs by 1.504.
- Apply the Jain-Mahdian-Saberi algorithm.
- Now consider the original facility costs.
- Apply greedy augmentation (Charikar, Guha 1999): Let g_i be the service cost saving induced by adding facility i. Iteratively pick an element i ∈ F maximizing g_i as long as this ratio is greater than 1.

Theorem

This is a 1.52-approximation algorithm for metric UFL.

Lower bound on approximation ratios

Theorem

There is no 1.463-factor approximation algorithm for metric UFL unless P = NP.

(Sviridenko [unpublished], based on Guha and Khuller [1999] and Feige [1998])

Local Search as a general heuristic

Basic Framework:

- Define a neighbourhood graph on the feasible solutions.
- Start with any feasible solution x.
- If there is a neighbour y of x that is (significantly) better, set x := y and iterate.

Features:

- Quite successful for many practical (hard) problems
- Many variants of local search heuristics
- Typically no guarantees of running time and performance ratio.

Local Search in Combinatorial Optimization

Example: TSP

- Even simple 2-opt typically yields good solutions. Variants (chained Lin-Kernighan) with empirically less than 1% error
- ▶ Worst-case running time of *k*-opt is exponential for all *k*.
- Performance ratio $\Omega(n^{\frac{1}{2k}})$.

(Applegate et al. 2003, Chandra, Karloff, Tovey 1999)

Example: Facility Location

- Probably the first nontrivial problem where local search led to constant-factor approximation algorithms.
- (Korupolo, Plaxton and Rajamaran 2000, Arya et al. 2004) But: for metric UFL worse in theory (maybe also in practice)
- But. for metric of L worse in theory (maybe also in practice
- The only known technique to obtain a constant-factor approximation for CAPACITATED FACILITY LOCATION.

Capacitated Facility Location (CFL)

Instance:

- finite sets \mathcal{D} (clients) and \mathcal{F} (potential facilities);
- metric service costs $c_{ij} \in \mathbb{R}_+$ for $i \in \mathcal{F}$ and $j \in \mathcal{D}$;
- ▶ an opening cost $f_i \in \mathbb{R}_+$ for each facility $i \in \mathcal{F}$;
- a capacity $u_i \in \mathbb{R}_+$ for each facility $i \in \mathcal{F}$;
- ▶ a demand d_j for each client $j \in D$.

We look for:

- ▶ a subset S of facilities (called *open*) and
- ▶ an assignment $x : S \times D \to \mathbb{R}_+$ with $\sum_{i \in S} x_{ij} = d_j$ for $j \in D$ and $\sum_{j \in D} x_{ij} \le u_i$ for $i \in S$
- such that the sum of facility costs and service costs

$$\sum_{i \in S} \left(f_i + \sum_{j \in \mathcal{D}} c_{ij} x_{ij} \right)$$

is minimum.

Splittable or Unsplittable Demands

Assume that facilities with given capacities are open. Task: assign the clients to these facilities, respecting capacity constraints.

- Splittable (or uniform) demand: Hitchcock transportation problem.
- Unsplittable non-uniform demand: Generalizes bin packing.

Consequence: CFL with unsplittable demands has no approximation algorithm. It is strongly NP-hard to distinguish between instances with optimum cost 0 and ∞ .

Hence consider splittable demands only.

Universal Facility Location (UniFL)

Instance:

- finite sets \mathcal{D} (clients) and \mathcal{F} (potential facilities);
- metric service costs, i.e. a metric c on $\mathcal{D} \cup \mathcal{F}$;
- ▶ a demand $d_j \ge 0$ for each $j \in D$;
- ▶ for each $i \in \mathcal{F}$ a cost function $f_i : \mathbb{R}_+ \to \mathbb{R}_+ \cup \{\infty\}$, left-continuous and non-decreasing.

We look for:

• a function $x : \mathcal{F} \times \mathcal{D} \to \mathbb{R}_+$ with $\sum_{i \in \mathcal{F}} x_{ij} = d_j$ for all $j \in \mathcal{D}$ (a *feasible solution*), such that $c(x) := c_F(x) + c_S(x)$ is minimum, where

$$c_F(x) := \sum_{i \in \mathcal{F}} f_i\left(\sum_{j \in \mathcal{D}} x_{ij}\right)$$
 and $c_S(x) := \sum_{i \in \mathcal{F}} \sum_{j \in \mathcal{D}} c_{ij}x_{ij}.$

UniFL: Facility cost function given by an oracle

 $f_i(z)$: cost to install capacity z at facility i.

Given by an oracle that, for each $i\in\mathcal{F},\ u,c\in\mathbb{R}_+$ and $t\in\mathbb{R},$ computes $f_i(u)$ and

 $\max\{\delta \in \mathbb{R} : u + \delta \ge 0, f_i(u + \delta) - f_i(u) + c|\delta| \le t\}.$

Proposition

There always exists an optimum solution. (Mahdian and Pál 2003)

UniFL: important special cases

- ▶ UNCAPACITATED FACILITY LOCATION: $d_j = 1 \ (j \in D)$, and $f_i(0) = 0$ and $f_i(z) = t_i$ for some $t_i \in \mathbb{R}_+$ and all $z > 0 \ (i \in \mathcal{F})$.
- ▶ CAPACITATED FACILITY LOCATION: $f_i(0) = 0, f_i(z) = t_i \text{ for } 0 < z \le u_i \text{ and } f_i(z) = \infty \text{ for } z > u_i,$ where $u_i, t_i \in \mathbb{R}_+$ $(i \in \mathcal{F})$.
- ▶ SOFT-CAPACITATED FACILITY LOCATION: $d_j = 1 \ (j \in D)$, and $f_i(z) = \lceil \frac{z}{u_i} \rceil t_i$ for some $u_i \in \mathbb{N}$, $t_i \in \mathbb{R}_+$ and all $z \ge 0$ $(i \in \mathcal{F})$.

Simple local search operations

- ▶ ADD: open a facility (CFL); add capacity to a facility (UniFL).
- ▶ DROP: close a facility (CFL).
- ▶ SWAP: open one facility, close another one (CFL).

Even for CFL with non-uniform demands, these operations do not suffice:

When closing one facility, it may be necessary to open many other ones (and re-assign the demand along the edges of a star).

Pı	revious approximation	n algorithms fo	or CFL	and U	niFL
ſ	Kuehn, Hamburger 1963	add,drop,swap	CFL	—	
ſ	Korupolu, Plaxton, Raja-	add,drop,swap	CFL	8.001	uniform
	maran 1998				capacities
ſ	Chudak, Williamson	add,drop,swap	CFL	5.829	uniform
	1999				capacities
	Pál, Tardos, Wexler 2001	add,star	CFL	8.532	
ſ	Mahdian, Pál 2003	add,star	UniFL	7.873	
ſ	Zhang, Chen, Ye 2004	add,double-star	CFL	5.829	
ſ	Garg, Khandekar, Pandit	add,double-star	UniFL	5.829	not poly-
	2005				nomial!
ſ	Vygen 2005	add,comet	UniFL	6.702	
	All based on local search.	star	double star		met

Add Operation for UniFL

Let $t\in\mathcal{D}$ and $\delta>0.$ Replace current solution x by an optimum solution y of the transportation problem

$$\begin{split} \min \Biggl\{ c_{\mathcal{S}}(y) \quad \bigg| \quad y: \mathcal{F} \times \mathcal{D} \to \mathbb{R}_+, \ \sum_{i \in \mathcal{F}} y_{ij} = d_j \ (j \in \mathcal{D}), \\ \sum_{i \in \mathcal{D}} y_{ij} \leq \sum_{i \in \mathcal{D}} x_{ij} \ (i \in \mathcal{F} \setminus \{t\}), \ \sum_{i \in \mathcal{D}} y_{tj} \leq \sum_{i \in \mathcal{D}} x_{tj} + \delta \Biggr\}. \end{split}$$

We denote by

$$c^{x}(t,\delta) := c_{\mathcal{S}}(y) - c_{\mathcal{S}}(x) + f_{t}\left(\sum_{j \in \mathcal{D}} x_{tj} + \delta\right) - f_{t}\left(\sum_{j \in \mathcal{D}} x_{tj}\right)$$

the estimated cost (which is at least c(y) - c(x)).

How to find a profitable ADD operation Lemma Let $\epsilon > 0$ and $t \in \mathcal{F}$. Let x be a feasible solution. Then there is an algorithm with running time $O(|V|^3 \log |V|\epsilon^{-1})$ that \blacktriangleright finds a $\delta \in \mathbb{R}_+$ with $c^x(t, \delta) \leq -\epsilon c(x)$ \blacktriangleright or decides that no $\delta \in \mathbb{R}_+$ exists for which $c^x(t, \delta) \leq -2\epsilon c(x)$. (Mahdian, Pál 2003)

PIVOT Operation

Let x be a feasible solution. Let A be a graph with $V(A) = \mathcal{F}$ and

$$\delta \in \Delta_{\mathcal{A}}^{\times} := \left\{ \delta \in \mathbb{R}^{\mathcal{F}} \left| \sum_{j \in \mathcal{D}} x_{ij} + \delta_i \ge 0 \text{ for all } i \in \mathcal{F}, \sum_{i \in \mathcal{F}} \delta_i = 0 \right\}.$$

Then we consider the operation $\operatorname{PIVOT}(A, \delta)$, which means:

- Compute a minimum-cost (w.r.t. c) uncapacitated δ -flow in (A, c).
- ▶ W.I.o.g., the edges carrying flow form a forest.
- Scan these edges in topological order, reassigning clients according to flow values.
- This increases the cost of the solution by at most the cost of the flow plus

$$\sum_{i\in\mathcal{F}}f_i\left(\sum_{j\in\mathcal{D}}x_{ij}+\delta_i\right)-f_i\left(\sum_{j\in\mathcal{D}}x_{ij}\right).$$

How to find a profitable PIVOT operation

But: how to choose δ ?

- δ cannot be chosen almost optimally for the complete graph (unless P = NP).
- We show how to choose δ almost optimally if A is a forest.

Restrict attention to **PIVOT** on arborescences

Let A be an arborescence with $V(A) = \mathcal{F}$. Let x be a feasible solution.

For $\delta \in \Delta_A^{\mathsf{x}}$ define

$$c^{\mathsf{x}}_{\mathcal{A},i}(\delta) := f_i \left(\sum_{j \in \mathcal{D}} x_{ij} + \delta_i \right) - f_i \left(\sum_{j \in \mathcal{D}} x_{ij} \right) + \left| \sum_{j \in \mathcal{A}_i^+} \delta_j \right| c_{ip(i)}$$

for $i \in \mathcal{F}$ and

$$c^{x}(A,\delta):=\sum_{i\in\mathcal{F}}c^{x}_{A,i}(\delta).$$

Here A_i^+ denotes the set of vertices reachable from *i* in *A*, and p(i) is the predecessor of *i*.

How to find a profitable PIVOT for an arborescence Lemma Let $\epsilon > 0$. There is an algorithm with running time $O(|\mathcal{F}|^4 \epsilon^{-3})$ that • finds a $\delta \in \Delta_A^x$ with $c^x(A, \delta) \leq -\epsilon c(x)$ • or decides that no $\delta \in \Delta_A^x$ exists for which $c^x(A, \delta) \leq -2\epsilon c(x)$. (Vygen 2005)

Bounding the cost of a local optimum

Let $0 < \epsilon < 1$. Let x, x^* be feasible solutions to a given instance. Lemma

If
$$c^{\mathsf{x}}(t,\delta) \ge -\frac{\epsilon}{|\mathcal{F}|}c(\mathsf{x})$$
 for all $t \in \mathcal{F}$ and $\delta \in \mathbb{R}_+$, then

$$c_{\mathcal{S}}(x) \leq c_{\mathcal{F}}(x^*) + c_{\mathcal{S}}(x^*) + \epsilon c(x).$$

(Pál, Tardos and Wexler 2001)

If
$$c^{x}(A, \delta) \geq -\frac{\epsilon}{|\mathcal{F}|}c(x)$$
 for all stars and comets A and $\delta \in \Delta_{A}^{x}$, then

$$c_F(x) \le 4c_F(x^*) + 2c_S(x^*) + 2c_S(x) + \epsilon c(x).$$

(Vygen 2005)

The total cost of a local optimum

These two lemmata imply:

Theorem

If $c^{x}(t,\delta) > -\frac{\epsilon}{8|\mathcal{F}|}c(x)$ for $t \in \mathcal{F}$ and $\delta \in \mathbb{R}_{+}$ and $c^{x}(A,\delta) > -\frac{\epsilon}{8|\mathcal{F}|}c(x)$ for all stars and comets A and $\delta \in \Delta_{A}^{x}$, then

$$c(x) \leq (1 + \epsilon)(7c_F(x^*) + 5c_S(x^*)).$$

By scaling facility costs by $\frac{\sqrt{41}-5}{2}$ we get a polynomial-time $(\frac{\sqrt{41}+7}{2}+\epsilon)$ -approximation algorithm for UniFL.

How to bound the facility cost Let x be the current solution and x^* be an optimum solution. Let $b(i) := \sum_{j \in D} (x_{ij} - x_{ij}^*)$ $(i \in \mathcal{F})$. Let y be an optimum transshipment from $S := \{i \in \mathcal{F} : b(i) > 0\}$ to $T := \{i \in \mathcal{F} : b(i) < 0\}$. W.l.o.g., the edges where y is positive form a forest F. The cost of y is at most $c_S(x^*) + c_S(x)$. Using F and y, we will define a set of pivot operations on stars and comets, whose total estimated cost is at most $4c_F(x^*) - c_F(x) + 2c_S(x^*) + 2c_S(x)$. An operation (A, δ) closes $s \in S$ if $\delta_s = -b(s) < 0$, and it opens $t \in T$ if $0 < \delta_t \le -b(t)$. Over all operations to be defined, we will close each $s \in S$ once, open each $t \in T$ at most four times, and use an estimated routing cost at most twice the cost of y.









Complexity Results

(All the following results are by Maßberg and Vygen 2005)

Proposition

- There is no (1.5 ε)-approximation algorithm (for any ε > 0) unless P = NP.
- There is no (2 ε)-approximation algorithm (for any ε > 0) for any class of metrics where the Steiner tree problem cannot be solved exactly in polynomial time.
- There is a 2-approximation algorithm for geometric instances (similar to Arora's approximation scheme for the TSP).
 However, this is not practically efficient.

Lower bound: spanning forests

Let F_1 be a minimum spanning tree for (\mathcal{D}, c) . Let e_1, \ldots, e_{n-1} be the edges of F_1 so that $c(e_1) \ge \ldots \ge c(e_{n-1})$. Set $F_k := F_{k-1} \setminus \{e_{k-1}\}$ for $k = 2, \ldots, n$.

Lemma

 F_k is a minimum weight spanning forest in (\mathcal{D},c) with exactly k components.

Proof.

By induction on k. Trivial for k = 1. Let k > 1. Let F^* be a minimum weight k-spanning forest. Let $e \in F_{k-1}$ such that $F^* \cup \{e\}$ is a forest. Then

$$c(F_k) + c(e_{k-1}) = c(F_{k-1}) \le c(F^*) + c(e) \le c(F^*) + c(e_{k-1}).$$



Lower bound: number of facilities

Let t' be the smallest integer such that

$$\frac{1}{\alpha}c(F_{t'}) + d(\mathcal{D}) \leq t' \cdot u$$

Lemma

 t^\prime is a lower bound for the number of facilities of any solution.

Let t'' be an integer in $\{t', \ldots, n\}$ minimizing

$$\frac{1}{\alpha}c(F_{t''})+t''\cdot f.$$

Theorem $\frac{1}{\alpha}c(F_{t''}) + t'' \cdot f$ is a lower bound for the cost of an optimal solution.



Analysis of Algorithm A

Recall: $\frac{1}{\alpha}c(F_{t''}) + t'' \cdot f$ is a lower bound for the optimum. We set $L_r := \frac{1}{\alpha}c(F_{t''})$ and $L_f := t'' \cdot f$. Observe: $L_r + d(\mathcal{D}) \leq \frac{u}{f}L_f$.

The cost of the final solution is at most

$$c(F_{t''}) + t''f + \frac{2}{u} \Big(c(F_{t''}) + d(\mathcal{D}) \Big) f$$
$$= \alpha L_r + L_f + \frac{2f}{u} \big(\alpha L_r + d(\mathcal{D}) \big)$$
$$\leq \alpha L_r + L_f + 2\alpha L_f$$

Theorem

Algorithm A is a $(2\alpha + 1)$ -approximation algorithm.

Algorithm B

Define metric c' by $c'(v, w) := \min\{c(v, w), \frac{uf}{u+2f}\}.$

- 1. Compute a Steiner tree F for \mathcal{D} in (V, c') with some β -approximation algorithm.
- 2. Remove all edges e of F with $c(e) \ge \frac{uf}{u+2f}$.
- 3. Split up overloaded components of the remaining forest as in algorithm A.

Theorem

Algorithm B has perfomance ratio 3β .

Using the Robins-Zelikovsky Steiner tree approximation algorithm we get a 4.648-approximation algorithm.

With a more careful analysis of the Robins-Zelikovsky algorithm we can get a 4.099-approximation algorithm in $O(n^{2^{10000}})$ time.

Algorithm C

Define metric c'' by $c''(v, w) := \min\{c(v, w), \frac{uf}{u+f}\}$

- 1. Compute a tour F for \mathcal{D} in (V, c'') with some γ -approximation algorithm.
- 2. Remove the longest edge of F.
- 3. Remove all edges e of F with $c(e) \ge \frac{uf}{u+f}$.
- 4. Split up overloaded components of the remaining forest as in algorithm A.

Theorem

Algorithm C has perfomance ratio 3γ .

- Using Christofides' TSP approximation algorithm we get a
- 4.5-approximation algorithm in $O(n^3)$ time.

Comparison of the three approximation algorithms

- Algorithm A computes a minimum spanning tree.
- Algorithm B calls the Robins-Zelikovsky algorithm.
- Algorithm C calls Christofides' algorithm.
- Then each algorithm deletes expensive edges and splits up overloaded components.

В	general	4.099	$O(n^{2^{10000}})$
С	general	4.5	<i>O</i> (<i>n</i> ³)

Experimental Results

Algorithm A on six real-world instances:

	inst1	inst2	inst3	inst4	inst5	inst6
# terminals	3675	17140	45606	54831	109224	119461
MST length	13.72	60.35	134.24	183.37	260.36	314.48
t'	117	638	1475	2051	3116	3998
Lr	8.21	31.68	63.73	102.80	135.32	181.45
$L_r + L_f$	23.07	112.70	251.06	363.28	531.05	689.19
# facilities	161	947	2171	2922	4156	5525
service cost	12.08	54.23	101.57	159.93	234.34	279.93
total cost	32.52	174.50	377.29	531.03	762.15	981.61
gap (factor)	1.41	1.55	1.59	1.46	1.44	1.42
		•	•			

Igorithm A on fou	ur chinc co	maarod to	the provi	
auristic:	ir chips, co	mpared to	the previ	ously used
euristic.				
chip	Jens	Katrin	Bert	Alex
technology	180nm	130nm	130nm	130nm
# clocktrees	1	3	69	195
total # sinks	3805	137265	40298	189341
largest instance	375	119461	16260	35305
power (W, old)	0.100	0.329	0.306	2.097
power (W, new)	0.088	0.287	0.283	1.946
	11 10/	10.00/	7.5%	7 20/

Reduction of power consumption

Some Open Problems

- Close the gap between 1.46 and 1.52 for the approximability of UNCAPACITATED FACILITY LOCATION.
- Find better lower bounds than 1.46 for capacitated problems (such as CFL).
- ► Is Universal Facility Location really harder than CFL?
- ▶ Improve the approximation ratio for the problem with service capacities (in (ℝ², ℓ₁), with a practically efficient algorithm).
- In some real-world instances, there exists an interval graph on the terminals, and we have to partition this graph into cliques. Is there an approximation algorithm for the resulting problem?
- What other interesting problems combining facility location with network design, or routing, can be approximated?
- ► What about multi-stage extensions?

Further Reading

- J. Vygen. Approximation Algorithms for Facility Location Problems (lecture notes, with complete proofs and references). Can be downloaded at http://www.or.uni-bonn.de/~vygen
- B. Korte, J. Vygen. Combinatorial Optimization: Theory and Algorithms (Chapter 22). Springer, Berlin, third edition 2006. Also available in Japanese!
- J. Maßberg, J. Vygen. Approximation Algorithms for Network Design and Facility Location with Service Capacities. Proceedings of the 8th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2005); LNCS 3624, Springer, Berlin 2005, pp. 158–169

Algorithms for a Networked World

Magnús M. Halldórsson

Dept. of Computer Science, University of Iceland

Abstract:

The realization of Moore's law has ensured that computing ability has increased dramatically in our times. The law has held not only for processor power and quantity of internal and external memory, but also for the ability to communicate information. The resulting impact on essentially all spheres of society and life has been nothing less than astounding.

Along with increased connectivity, we are also seeing the introduction of a wide range of compact computing entities, possibly mobile and often non-statically connected into largely wireless networks. The explosion of the web and the internet as not only a source of information but also as a resource of computational intelligence, is poised to lead to a dramatic change in the way we view computation. The traditional view of an algorithm with full random access to its input, operating serially on a single processor, is on the retreat.

In comparison, it can be said that changes in CS theory are less dramatic. Surely, each year and each conference brings new topics, new subjects, new treatments, and new directions. Yet, we can also easily detect a great deal of consistency [one that is certainly comforting at times], and a measured pace of change. Are we theoreticians then by nature reactionaries? One of the theses of this talk is that the objects of study in CS theory are inherently fundamental and long-lasting, applying also to this Panopticon world of global and ubiquitous computing.

Yet, we cannot rest on our laurels, with self-satisfied smugness. We must find ways to treat the new means, ways, possibilities and limitations of computation in a systematic framework that continues to provide applied fields with rigorous guidance. The aim of the talk is to discuss some objectives, measures, and paradigms that address the changing nature of computing in a networked world. The concrete examples discussed will mostly relate to problems of coloring and packing, the topics of main focus of the speakers research.

This will by no means be a comprehensive overview - in fact, it is unlikely to be even a balanced introduction. Instead, the hope is that by posing some questions, some members of the audience will eventually be prompted to find some of the answers.

Algorithms for Sequence Manipulation and Related Problems D. T. Lee

Institute of Information Science Academia Sinica & & Department of Computer Science and Information Engineering National Taiwan University

2006 NHC Spring School Workshop March 2-3, 2006

Optimization on Sequences

Given a sequence $A = a_1, a_2, ..., a_n$, an optimization problem on sequences is to maximize or minimize some function, such as: sum of subsequence, density of subsequence, etc., with some constraints, such as: length, weight, etc.

3/92

4/92

Combinatorial Optimization

- n The combinatorial optimization typically deals with problems of *maximizing* or *minimizing* a function of one or many variables subject to a number of inequality constraints.
- n Consider two categories of problems on sequences:
 - n Optimization Problems
 - Range Search Problems

Range

Range Search Problems

A range (query) search problem is typically to report the subset S' to count the total number of elements of the subset S' of a set S contained in a query range Q subject to certain conditions.



n Given a sequence $A = a_1, a_2, ..., a_n$, and a range query, we want to *report* or *count* some subsequences of A contained in the query range satisfying certain conditions.









- n Sum of subsequence,
- n Density of subsequence
- n Selection of subsequence
- n subject to constraints on length, or weight of subsequences.







Applications in Bioinformatics

- Useful applications in bioinformatics including
- n finding tandem repeats, which are commonly used to map disease genes
- n locating DNA segments with rich CG content is a key step in gene finding and promoter prediction
- n low complexity filter, which is most commonly applied in sequence database search.

Technique Used – Lin *et al.*

- Lin, Jiang and Chao gave an O(n) time algorithm based on a clever technique called *left-negative decomposition*.
- J. of Computer and System Sciences '02

16/92

Max-Sum Segment Problem

Prior Results:

- n O(n) time for the special case $\ell = 1, u = n$ Gries [Science of Computer Programming'82]
- n O(n) time for the special case $\ell = 1$, u = nBentely. [Commun. ACM '84]
- n O(n) time. Lin, Jiang and Chao. [Journal of Computer and System Sciences '02]
- n O(n) time. Fan, Lee, Lu, Tsou, Wang, and Yao. [CIAA '03]

14/92

13/92

Lin et al. Algorithm ['02]

- n A sequence $A=a_1, a_2, ..., a_n$ is *left-negative* iff the sum of each proper prefix $a_1, a_2, ..., a_i$ is negative or zero for all $1 \le i \le n-1$.
- A partition of the sequence $A = A_1A_2...A_k$ is *minimal left-negative* if each A_i , $1 \le i \le k$, is *left-negative*, and, for each $1 \le i \le k-1$, the sum of A_i is positive.

Example – left-negative sequence

- n The sequence -4, 1, -2, 3 is left-negative.
- n The sequence -5, 3, 4, -1, 2, -6 is not left-negative.
- n The partition (-5, 3, 4), (-1, 2), (-6) is minimal left-negative.
- For every suffix of a sequence we can find a minimal left-negative partition.
- n (3)(4)(-1,2)(-6); (4)(-1,2)(-6); (-1,2)(-6), (2)(-6); (-6) are all possible minimal left-negative partitions

17/92





- a maximum sum segment must be at a block boundary.
- n (-5, (3), (4)), (-1, (2)), (-6)
- n Prefix sum sequence S= 0, -5, -2, 2, 1, 3, -3
- n Sum(1st) = 2, Sum(2nd) = 1, Sum(3rd) = -6
- n Max-Sum segment is a_2 , $a_3 a_4 a_5$ of sum $s_5 s_1 = 8$



Application: Max-Density Segment

- Finding the segment with the largest GC-ratio in a DNA sequence can be cast as a maximum-density segment problem.
- In Input sequence A corresponds to the given DNA sequence, where $a_i = 1$ if the corresponding nucleotide G or C; and $a_i = 0$ otherwise.
- Output feasible segment corresponds to the region with the largest GC-ratio.

21/92

Max-Density Segment Problem

Prior Results:

- ⁿ O(n ℓ) time algorithm for the special case u = nHuang. [Computer Appl. in the Biosciences '94]
- n O(n log ℓ) time algorithm for the special case u = n Lin, Jiang, and Chao. [J. Comp. and Syst. Sci. '02]
 n O(n log (u-ℓ)) time algorithm. Goldwasser, Kao, and
- Lu. [J. of Comp. and Syst. Sci. '03]
- n O(n) time algorithm. Kim. [IPL '03]– has a flaw
 - O(n) time algorithm. Chung and Lu. [SICOMP '04]

23/92

24/92

Application in Bioinformatics

- Maximum-Density Segment Problem arises from the investigation of non-uniformity of nucleotide composition within genomic sequences, which was first revealed through thermal melting and gradient centrifugation experiments.
- Researchers observed that the compositional heterogeneity is highly correlated to the GC content of the genomic sequences, and this motivates finding the segment with the largest GC-ratio.

Technique Used – Kim. IPL '03 3

- ⁿ Construct a point set in the plane $P = \{p_k | p_k = (k, s_k), k = 1, 2, ..., n, where s_k = a_1 + a_2 + ... + a_k$ is the prefix sum of sequence A.
- $\label{eq:product} \begin{array}{l} \mbox{n Construct lower hull of $P_{j} = \{p_{j-\ell}, \, p_{j-\ell+1}, \, \dots, \, p_{j-u}\}$} \\ \mbox{and find tangent segment t_{j} from p_{j} to P_{j}}. \end{array}$
- ⁿ The tangent segment of the maximum slope is the maximum-density segment of A.















n Input:

- n a sequence A of n real numbers $a_1, a_2, ..., a_n$
- n two nonnegative real numbers ℓ , u with $\ell \leq u$
- n a positive integer k.

n Output:

n the k segments such that their densities are the k largest over all $O((u-\ell) n)$ feasible segments.

33/92











Technique – Random Sampling

- Contract initial interval $[s_1, s_r] = (-\infty, \infty)$ into a smaller subinterval $[s_1', s_r']$ such that it contains the kth largest feasible segment s* and the subinterval $[s_1', s_r']$ contains at most $O(n^2/n^{1/2}) = O(n^{3/2})$ feasible segments.
- $\label{eq:sphere:sphe$
- If both steps are successful, output all the segments in $[s_1", s_r"]$ and find the solution segment with an appropriate rank, whose sum is s*, by using any *standard selection* algorithm.



















Range Search Problems

- n Single-shot mode query: O(n) query time is optimal for this query mode.
- n Repetitive mode query: Preprocessing allowed, and obtaining an o(n) time query time is the goal. A trade-off between storage and query time is expected.



51/92

49/92

Sum Range Search Problem single-shot mode

n Input:

- n a sequence A of n real numbers a_1, a_2, \ldots, a_n
- n two nonnegative real numbers ℓ , u with $\ell \leq u$
- n two real numbers s_1 , s_r with $s_1 \le s_r$.

n Output:

 ${\sf n}\,$ the segments over all $O((\textit{u-l}) \; n)$ feasible segments such that their sums are between s_l and $s_r.$

order-statistic tree 2
prefix sum x_i of the sequence A, x_i = a₁+...+a_i, i=1, 2, ..., n, and let x₀ = 0.
Let P = {x₀, x₁, ..., x_n}.
Let P_j = {x_{j-u}, x_{j-u+1}, ..., x_{j-l}}.
Maintain an order-statistic tree T(P_j) on P_j, by scanning the sequence of prefix sums.

Technique - maintaining an

Order-Statistic Tree

- An order-statistic tree is a balanced binary search tree with size information, size[z], stored in each node z of the tree, and containing the total number of nodes in the subtree rooted at z.
- n For an internal node z, size[z] = size[left[z]] + size[right[z]] + 1
- For a leaf node z, size[z]= 1.







Density Range Search Problem single-shot mode

n Input:

- $\ensuremath{\,{}^{n}}$ a sequence A of n real numbers $a_1, a_2, \, ..., \, a_n$
- n two nonnegative real numbers ℓ , u with $\ell \leq u$
- n two real numbers d_l , d_r with $d_l \le d_r$.

n Output:

n the segments over all $O((u-\ell) n)$ feasible segments such that their densities are between d_i and d_r

57/92



Density Range Search Problem single-shot mode

n Recent result:

- N We obtain an O(n log $(u-\ell)$ + h) time algorithm, where h is the output size.
- n Use problem transformation
- n Maintain a *priority search tree* T(P_i) on P_i.





- We can further transform this *geometric* slope range search problem into its dual problem, by mapping a point $p_i = (x_i, y_i)$ into its dual line ℓ_i : $y = x_i x - y_i$.
- For any two points p_i , and p_j , their corresponding dual lines ℓ_i , ℓ_j will intersect at the x-coordinate $x_{ij} = (y_j - y_i)/(x_j - x_i)$ which equals m(i, j).











Algorithm for Intersection Search

- The Intersection Search problem is now equivalent to an orthogonal range search of the form $R_j = [u_j, \infty) \times (-\infty, v_j]$ to report all the points of Q_j which lie in R_j for each j = 1, 2, ..., n.
- We use a data structure called *priority* search tree to support the above orthogonal range queries in logarithmic time.











- n a sequence A of n real numbers $a_1, a_2, ..., a_n$.
- two nonnegative real numbers ℓ , u with $\ell \le u$
- n two real numbers d_l , d_r with $d_l \le d_r$.
- n for an intervals $[d_l,\,d_r]$, reports the segments over all $O((\mathit{u}{-}\ell)\;n)$ feasible segments such that their densities are between d_l and d_r .

Density Range Search Problem repetitive mode Nork to Do:

n Try to preprocess A into a nice appropriate structure such that the query time is o(n + h).

Range **Minimum Search** Problem

n Input:

73/92

- ${\tt n}\,$ a sequence A of n real numbers $a_1,a_2,\,...,\,a_{{\tt n}.}$
- n two real numbers i, j, $i \leq j$.

n Online Query:

for each query interval [i, j], reports the index k with $i \le k \le j$ such that a_k achieves minimum.

75/92

Aggregate Range Search Problem

Given S and a range query Q, compute an aggregate function on the subset $S'=S \cap Q$ An aggregation function can be Min, Max, Sum, Count, Mean, Median (of S'), etc.



Range Minimum Search Problem

n Related work :

- O(n) preprocessing time and O(1) query time under the unit-cost RAM model. -- Gabow, Bentley, and Tarjan. STOC 1984
- n O(n) preprocessing time and O(1) query time under the unit-cost RAM model. -- Bender and Colton. In Proc. the 4th Latin American Symposium on Theoretical Informatics 2000

77/92

Problem Reduction

- n LCA problem reduces to ± 1 RMS problem
- n Observation: The LCA of nodes u and v in T is the lowest node encountered between the visits to u and to v during a depth first traversal of T, where the depths of the nodes in T differ by exactly one.

79/92

Technique Used ±1RMS problem: adjacent elements of the input sequence A differ by ±1 or -1 Least Common Ancestor (LCA) problem reduces to ±1RMS problem by depth first search traversal of input tree T of LCA. An (O(n log n), O(1))-time table-lookup algorithm for RMS Using the above algorithm on a smaller array A' obtained by partitioning A into 2n/logn blocks, each of size (log n)/2 and making use of the ±1 property, we can solve ±1RMS in (O(n), O(1))-time. RMS reduces to LCA building the Cartesian tree of A.



Problem Reduction

- RMS can be reduced to LCA by building a Cartesian tree of A.
- The root of a Cartesian tree is the minimum element of the array. The root is labeled with the position k of this minimum element.
- n The left and right children of the root are the roots of recursively constructed Cartesian trees of the left and right subarrays respectively.

81/92



Range **Maximum-Sum Segment** Search Problem -- definition

n Input:

- n a sequence A of n real numbers a_1, a_2, \ldots, a_n to be preprocessed.
- n real numbers i, j, i \leq j and k, l, k \leq l.

n Online Query:

n for two query intervals [i, j] and [k, l], reports A(x, y) with $i \le x \le j$ and $k \le y \le l$ that maximizes s(x, y). Dechnique Used
1. Let S = s₁, s₂, ..., s_n be the sequence where s_k = a₁+a₂+...+a_k is the prefix sum of sequence A.
1. Disjoint case: by min= RMinS(S, i, j), max= RMaxS(S, k, l), Ans.= s_{max} - s_{min}
1. Overlapping case: Divide into 3 possible cases and take minimum of the outputs of these three cases.
1. [i, k] and [k, l]: by RMinS(S, i, k), RMaxS(S, k, l)
1. [k, j] and [j, l]: by RMinS(S, k, j), a special case!

Range Maximum-Sum Segment Search Problem- A Special Case

n Input:

- n a sequence A of n real numbers $a_1, a_2, ..., a_n$.
- n two real numbers i, j, $i \leq j$.

n Online Query:

n for any query interval [i, j], reports A(x, y) with $i \le x \le y \le j$ such that A(x, y) is the maximum-sum segment of A(i, j).

85/92

Algorithm

- Let $C[\cdot]$ be the array of prefix sum of A.
- Define: left bound L[k] of A at index k to be the largest index *l* with $1 \le l \le k-1$ such that $C[l] \ge C[k]$, and L[k]=0, if no such *l* exists.
- Define: partner P[k] of A at index k to be the largest index p with $L[k]+1 \le p \le k$ that minimizes C[p-1].
- $\label{eq:alpha} \begin{array}{l} A(P[k],k) \text{ is a candidate segment of } A \text{ at index } k \\ \text{with sum } M[k] = s \ (P[k],k), \text{ for } 1 \leq k \leq n. \end{array}$

87/92

Range Maximum-Sum Segment Search Problem- A Special Case Related work : O(n) preprocessing time and O(1) query time

- ⁿ O(n) preprocessing time and O(1) query time under the unit-cost RAM model by Chen and Chao. ISAAC 2004
- n They solved RMSSS by using RMinS and RMaxS.

Algorithm of RMSSS(A, i, j) $r \leftarrow RMaxS(M, i, j), i.e. M[r]$ is maximum if P[r] < i then $p \leftarrow RMinS(C, i-1, r-1)+1$ $s \leftarrow RMaxS(C, i-1, r-1)+1$ if C[r]-C[p-1] < M(s) then output (P[s], s) else output (p, r) else output (P[r], r)

Range Maximum-Density Segment Search Problem

n Input:

n a sequence A of n real numbers $a_1, a_2, ..., a_n$

n Online Query:

n for two intervals [i, j] and [k, l], reports A(x, y) with $i \le x \le j$ and $k \le y \le l$ that maximizes d(x, y).

89/92

Generalization

n We can also generalize the above aggregation range search problems to other aggregation functions or consider the dynamic range query problems which support insertion, deletion, concatenation operations.

91/92

Range Maximum-Density Segment Search Problem- A Special Case ■ Input: ■ a sequence A of n real numbers a₁, a₂, ..., a_n. ■ two real numbers i, j, i ≤ j.

n Online Query:

n for any query interval [i, j], reports A(x, y) with $i \le x \le y \le j$ such that A(x, y) is the maximumdensity segment of A(i, j).


Dynamic Data Structures in Computational Geometry

Timothy M. Chan School of Computer Science University of Waterloo

Every student in computer science knows that we can maintain a set of n real numbers (in 1-d) to support insertions and deletions in $O(\log n)$ time and searches in $O(\log n)$ time, by his/her favorite kind of balanced search trees. The underlying question in this talk is: to what extent can this basic result hold in dimension beyond one? In other words, for what problems in computational geometry can we maintain solutions efficiently under insertions and deletions of data?

In the first part of the talk, we give a rough introduction to this fundamental and important topic called "dynamic computational geometry", which has been extensively studied for over three decades. We obviously do not have time to give a comprehensive survey (despite what the title says), but the purpose is to provide some examples to illustrate what types of problems have been studied, what types of results are possible, and what general techniques are available. Our focus is on fully dynamic algorithms that have guaranteed time bounds for arbitrary update sequences. We mention the following results:

- For orthogonal range searching or 2-d point location, dynamization with polylogarithmic update time is relatively easy, by directly applying standard balancing tricks to geometric search trees (such as range trees or segment trees); for these problems, research is thus all about removing logarithmic factors [16, 14, 10, 15]. The well-known dynamic 2-d convex hull problem [17, 5, 3] also fits in this category.
- Simple dynamization strategies (the "logarithmic method" and the "square-root method") are applicable to a general class of *decomposable search* problems [2].
- For some non-decomposable problems, like 2-d *smallest enclosing circle* and 3-d *linear programming*, dynamization can still be obtained by reducing them to decomposable problems through parametric search [13].
- A general technique by Eppstein [12] can solve dynamic max-max- or min-min-type problems like *bichromatic closest pair*. A key idea here is to handle deletions by re-insertions.
- For some other non-decomposable problems, like 2-d width [7] and rectangle connectivity [6], polylogarithmic dynamization is not known, but sublinear update time is possible.
- For even more challenging problems (like 2-d smallest enclosing rectangle), min-max-type problems (like 2-d discrete 1-center), and non-convex optimization problems (like 2-d largest empty circle), nothing is known at all, except in special cases (such as the insertion-only and offline settings) [8].

In the second part of the talk, we focus specifically on the dynamic 2-d *post office* problem, and discuss in more detail a new data structure recently discovered by this speaker [9], which is

the first polylogarithmic result known. The problem is to maintain a set of points in the Euclidean plane, to support exact nearest neighbor queries. (By a standard lifting transformation, it can also be viewed as a dynamic 3-d convex hull problem.) The previous solution by Agarwal and Matoušek [1] more than a decade ago has $O(n^{\varepsilon})$ amortized update time and $O(\log n)$ query time, or $O(\log^2 n)$ amortized update time and $O(n^{\varepsilon})$ query time, for any fixed $\varepsilon > 0$. The new solution has $O(\log^3 n)$ amortized insertion time, $O(\log^6 n)$ amortized deletion time, and $O(\log^2 n)$ query time. The techniques involved include removing "high-degree" elements (following Agarwal and Matoušek [1]), handling deletions by re-insertions (as in Eppstein [12]), and using geometric tools like "sampling" and "conflict lists" (as in [11, 4, 18]).

We conclude by mentioning applications of dynamic data structures in the design of geometric algorithms, and by listing some open problems in the area.

References

- P. K. Agarwal and J. Matoušek. Dynamic half-space range reporting and its applications. Algorithmica, 13:325-345, 1995. (Also in FOCS 1992.)
- J. Bentley and J. Saxe. Decomposable searching problems I: static-to-dynamic transformation. J. Algorithms, 1:301-358, 1980.
- [3] G. S. Brodal and R. Jacob. Dynamic planar convex hull. In FOCS 2002.
- [4] T. M. Chan. Random sampling, halfspace range reporting, and construction of $(\leq k)$ -levels in three dimensions. SIAM J. Comput., 30:561-575, 2001. (Also in FOCS 1998.)
- [5] T. M. Chan. Dynamic planar convex hull operations in near-logarithmic amortized time. J. ACM, 48:1-12, 2001. (Also in FOCS 1999.)
- [6] T. M. Chan. Dynamic subgraph connectivity with geometric applications. In STOC 2002.
- [7] T. M. Chan. A fully dynamic algorithm for planar width. Discrete Comput. Geom., 30:17-24, 2003. (Also in SoCG 2001.)
- [8] T. M. Chan. Semi-online maintenance of geometric optima and measures. SIAM J. Comput., 32:700-716, 2003. (Also in SODA 2002.)
- [9] T. M. Chan. Dynamic data structures for 3-d convex hulls and 2-d nearest neighbor queries. In SODA 2006.
- [10] Y.-J. Chiang and R. Tamassia. Dynamic algorithms in computational geometry. Proc. IEEE, 80:1412-1434, 1992.
- [11] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. Discrete Comput. Geom, 4:387-421, 1989.
- [12] D. Eppstein. Dynamic Euclidean minimum spanning trees and extrema of binary functions. Discrete Comput. Geom., 13:111-122, 1995. (Also in FOCS 1992.)
- [13] J. Matoušek. Linear optimization queries. J. Algorithms, 14:432-448, 1993.
- [14] K. Mehlhorn. Data Structures and Algorithms 3: Multi-Dimensional Searching and Computational Geometry. Springer-Verlag, 1984.
- [15] C. W. Mortensen. Fully-dynamic two dimensional orthogonal range and line segment intersection reporting in logarithmic time. In SODA 2003.
- [16] M. H. Overmars. The Design of Dynamic Data Structures. Springer-Verlag, 1983.
- [17] M. H. Overmars and J. van Leeuwen. Maintenance of configurations in the plane. J. Comput. Sys. Sci., 23:166-204, 1981. (Also in STOC 1980.)
- [18] E. Ramos. On range reporting, ray shooting, and k-level construction. In SoCG 1999.



Theme

Dynamic dota structures for searching in 1D:

balanced search trees

What about dynamic data structures in 2D (f beyond) ???

"Dynamic Computational Geometry"

w. guaranteed worst-case/amortized time

Settings:

fully dynamic (insert-only delete-only offline (all updates known) semi-online (delete times known) random updates kinetic

This Talk

Part I: Survey

Part II: A new result for dynamic 3D convex hull/ 2D post office [soon'06]







































Dynamic 3D Convex Hull/ 2D Post Office O(log³n) insert time (amort, road) O(log⁶n) delete time (amort, road) O(log²n) query time

Overview of Ideas

- handle insert by the logarithmic method (O(logn) subsets)
- move "high-degree stuff" to new subsets during pre-processing [also done by Agarwal, Matouiek, FOCS'92
- handle delete by re-inserts [inspired by Eppstein, FOCS'92]
- Use O(logn) somples & "conflict lists" [inspired by skip list & a static structure by Chan, FOCS'98]

Dual Problem Mainitain n halfspaces in 3D s.t. given vertical line q. find intersection at q i.e. mainitain n planes in 3D s.t. given vertical line q. find lowest plane at q. lower envelope



Preprocessing a Subset
For i=1.. logn
take sample of 2' planes

2' cells, each w.
~ n/2' conflicts
move planes that have > b conflicts
to new subset
Analysis:
planes moved ~
$$\frac{n}{b}$$
 logn
> only a fraction moved, for b ~ logn
> preproc. time O(nlogn)
=) insert time O(logn)







Open Problems - dynamic 3D CH / 2D Post office: O(logn) update & query ????? & O(n) space ?? - dynamic 2D paint location; O(logn) update & query ?? - dynamic 2D width: o(Un) ?? - dynamic 3D approx. width (or "core-sets" O(logn) ?? - dynamic 2D connectivity for line segments o(n) ?? - dynamic 2D conterpoint o(n) ?? : :







Drawing Metro Map





Introduction Model Drawing Metro Maps MIP & Experiments on traces with few crossings NP-Hardness

Mixed-Integer Programming

- Linear Programming: efficient optimization method for
 linear constraints and objective function,
 real-valued variables (domain R).
- Mixed-Integer Programming (MIP):
 - inxed-integer Programming (MIP).
 - allows also integer variables (domain Z),
 solution NP-hard in general.
- Still a practical method for many hard optimizat. problems.

Theorem (Nöllenburg & Wolff GD'05)

Alexander Wolff

Drawing Metro Maps MIP & Experiments ing trees with few crossings NP-Hardness

Example: Octilinearity and Relative Position





Spanni	Introduction Drawing Metro Maps ng trees with few crossings	Model MIP & Experiments NP-Hardness	
Example: Octil	inearity and	Relative Positio	n
2	Previous Secto	r	
prev 4	y(u) - y(v) -y(u) + y(v) x(u) - x(v)	$ \leq M(1 - \alpha_{\text{prev}}) \leq M(1 - \alpha_{\text{prev}}) \geq -M(1 - \alpha_{\text{prev}}) $	(u, v)) (u, v)) $(u, v)) + \ell_{uv}$
next		_ (p.o.(, ,, u
How does this v	work?		
Case 1: $\alpha_{prev}($	<i>u</i> , <i>v</i>) = 0		
	y(u) - y(v)	< <i>M</i>	
	-y(u)+y(v)	< <i>M</i>	
	x(u) - x(v)	$\geq \ell_{uv} - M$	(1) (1)
			(B) B
			- প্র
Alexander Wolff	13	42	Geometric Networks

Introduction Drawing Metro Maps Spanning trees with few crossings	Model MIP & Experiments NP-Hardness	
Example: Octilinearity and R	elative Position	
Previous Sector		
$\begin{array}{c} y(u) - y(v) \\ -y(u) + y(v) \\ x(u) - x(v) \end{array}$	$ \leq M(1 - \alpha_{prev}(u, v)) \\ \leq M(1 - \alpha_{prev}(u, v)) \\ \geq -M(1 - \alpha_{prev}(u, v)) + \ell_{uv} $	
How does this work?		
Case 2: $\alpha_{\text{prev}}(u, v) = 1$		
$y(u) - y(v) \le -y(u) + y(v) \le x(u) - x(v) \ge -x(v) = -x(v) \ge -x(v) = -$	≤ 0 ≤ 0 ≥ ℓ _{uv}	
		< ≥ <
Alexander Wolff 13 42	2 Geometric Net	works

Spanni	Drawing Metro Maps MIP & Experimen ng trees with few crossings NP-Hardness	nts
Example: Octil	inearity and Relative	Position
	Original Castor	
prev a rest	$\begin{array}{rcl} \text{Original Sector} \\ z_2(u) - z_2(v) &\leq & M_1 \\ -z_2(u) + z_2(v) &\leq & M_2 \\ z_1(u) - z_1(v) &\geq & -M_2 \end{array}$	$ \begin{array}{l} (1 - \alpha_{orig}(u, v)) \\ (1 - \alpha_{orig}(u, v)) \\ (1 - \alpha_{orig}(u, v)) + 2\ell_{uv} \end{array} $
prev 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2	Next Sector $x(u) - x(v) \leq M($ $-x(u) + x(v) \leq M($ $y(u) - y(v) \geq -M($	$1 - \alpha_{\text{next}}(u, v)) 1 - \alpha_{\text{next}}(u, v)) 1 - \alpha_{\text{next}}(u, v)) + \ell_{uv} $
Alexander Wolff	14 42	ି ଥି ୬୦.୯ Ccometric Networks





























Crossing-free spanning treesLet $k = \#$ crossings in G (crossing = pair of crossing edges). $\varphi(G) = 1 + \min \#$ drossings in spanning tree of G .Theorem (Knauer, Schramm, Spillner, Wolff ISAAC'05)If there is a poly-time $k^{1-\varepsilon}$ -approximation algo $A_{1-\varepsilon}$ of $\varphi(G)$ for some $\varepsilon \in (0, 1]$, then $\mathcal{P} = \mathcal{NP}$.Let $G' = \bigcirc $	Drawing Mero Maps Spanning trees with few crossings MIP Formulation
Let $k = \#$ crossings in G (crossing = pair of crossing edges). $\varphi(G) = 1 + \text{minimum } \#$ crossings in spanning tree of G . Theorem (Knauer, Schramm, Spillner, Wolff ISAAC'05) If there is a poly-time $k^{1-\varepsilon}$ -approximation algo $A_{1-\varepsilon}$ of $\varphi(G)$ for some $\varepsilon \in (0, 1]$, then $\mathcal{P} = \mathcal{NP}$. Let $G' = \bigcirc $	Crossing-free spanning trees
Theorem (Knauer, Schramm, Spillner, Wolff ISAAC'05)If there is a poly-time $k^{1-\varepsilon}$ -approximation algo $A_{1-\varepsilon}$ of $\varphi(G)$ for some $\varepsilon \in (0, 1]$, then $\mathcal{P} = \mathcal{NP}$.Let $G' = \bigcirc $	Let $k = $ #crossings in G (crossing = pair of crossing edges). $\varphi(G) = 1 + $ minimum #crossings in spanning tree of G .
If there is a poly-time $k^{1-\varepsilon}$ -approximation algo $A_{1-\varepsilon}$ of $\varphi(G)$ for some $\varepsilon \in (0, 1]$, then $\mathcal{P} = \mathcal{NP}$. Let $G' = \bigcirc $	Theorem (Knauer, Schramm, Spillner, Wolff ISAAC'05)
Let $G' = \bigcirc G \longrightarrow \bigcirc G \longrightarrow \bigcirc G = k^{1/\varepsilon - 1} \bigcirc G$ $\Rightarrow G'$ has $K = k \cdot k^{1/\varepsilon - 1} = k^{1/\varepsilon}$ crossings. $\boxed{\varphi(G) = 1 \Leftrightarrow \varphi(G') = 1} \qquad \forall s.$ $\downarrow A_{1-\varepsilon}(G') \qquad \qquad$	If there is a poly-time $k^{1-\varepsilon}$ -approximation algo $A_{1-\varepsilon}$ of $\varphi(G)$ for some $\varepsilon \in (0, 1]$, then $\mathcal{P} = \mathcal{NP}$.
$ \Rightarrow G' \text{ has } K = k \cdot k^{1/\varepsilon - 1} = k^{1/\varepsilon} \text{ crossings.} $ $ \hline \varphi(G) = 1 \Leftrightarrow \varphi(G') = 1 \text{vs.} \hline \varphi(G) \ge 2 \Leftrightarrow \varphi(G') \ge 1 + k^{1/\varepsilon - 1} \\ \downarrow A_{1-\varepsilon}(G') \downarrow A_{1-\varepsilon}(G') \\ \le K^{1-\varepsilon} = k^{1/\varepsilon - 1} \text{ cross.} \ge 1 + k^{1/\varepsilon - 1} \text{ cross. in sptree} $	Let $G' = \bigcirc $
$ \begin{array}{ c c c c c }\hline \varphi(G) = 1 \Leftrightarrow \varphi(G') = 1 & \text{vs.} & \hline \varphi(G) \ge 2 \Leftrightarrow \varphi(G') \ge 1 + k^{1/\varepsilon - 1} \\ \downarrow A_{1-\varepsilon}(G') & \downarrow A_{1-\varepsilon}(G') \\ \le K^{1-\varepsilon} = k^{1/\varepsilon - 1} \text{ cross.} & \ge 1 + k^{1/\varepsilon - 1} \text{ cross. in sptree} \end{array} $	$\Rightarrow G'$ has $K = k \cdot k^{1/\varepsilon - 1} = k^{1/\varepsilon}$ crossings.
$\leq K^{1-\varepsilon} = k^{1/\varepsilon-1} \text{ cross.} \qquad \geq 1 + k^{1/\varepsilon-1} \text{ cross. in sptree}$	$\boxed{\varphi(G) = 1 \Leftrightarrow \varphi(G') = 1} \text{VS.} \qquad \boxed{\varphi(G) \ge 2 \Leftrightarrow \varphi(G') \ge 1 + k^{1/\varepsilon - 1}} \\ + A = (G') \text{(G')}$
	$\leq K^{1-\varepsilon} = k^{1/\varepsilon-1} \text{ cross.} \qquad \geq 1 + k^{1/\varepsilon-1} \text{ cross. in sptree}$

Drawing Metro Maps Spanning trees with few crossings MIP Formulation

Fixed-parameter tractability

Aim

Solve (NP-) hard problem by restricting combinatorial explosion to parameter independent from input size.

Definition

A decision problem Π is *FPT w.r.t. a parameter k* if there is an algorithm that decides in $O(f(k) \cdot \text{poly}(|I|))$ time, whether an instance *I* of Π is a yes- or a no-instance. ($f : \mathbb{N} \to \mathbb{N}$ arbitrary!)

Example

k-VertexCover (Can we cover all edges of G with k vertices?) is FPT.



Introduction Drawing Metro Maps Spanning trees with few crossings	Non-Approximability Fixed-Parameter Algorithms MIP Formulation
ossingFreeSpanningTree	1
Theorem	
CrossingFreeSpanningTree is F w.r.t. #crossings in the input gra	PT ph.
i.e., can decide in $O(f(k) \cdot \text{poly}(k))$ vertices and k crossings has a c	n)) time, whether a graph with <i>n</i> crossing-free spanning tree.
Proof by enumeration of edges.	
$ \begin{array}{ll} X := \text{set of crossings} & \Rightarrow & E_X \\ G' := G \setminus E_X \end{array} $	$x := \bigcup X$ contains $\leq 2k$ edges.
For each subset H of E_X check: is H crossing-free? does G	$' \cup H$ span G?
Runtime: $O(2^{2n} \cdot \text{poly}) = O(4^{n} \cdot \text{poly})$	poly).

Note: G has c-f spanning tree \Leftrightarrow G has c-f spanning subgraph.

Introduction Approximability Drawing Metro Maps Spanning trees with few crossings MIP Formulation	Hitroduction Non-Approximel/281y Drawing Metro Maps Fixed-Parameter Algorithms Spanning trees with few crossings MIP Formulation
A faster algorithm	A yet faster algorithm
Theorem	Theorem
CrossingFreeSpanningTree can be decided in $O(3^k \cdot \text{poly})$ time.	CrossingFreeSpanningTree can be decided in $O(2^k \cdot poly)$ time.
Proof by enumeration of "crossing vectors".	Proof by enumeration of <i>maximal</i> crossing-free edge sets.
$X := \{\{e_1, f_1\}, \dots, \{e_k, f_k\}\} = $ set of crossings.	Let $e \in E_X \Rightarrow e$ crosses $c \ge 1$ other edges.
For each vector $x \in \{0, 1, 2\}^k$ $\begin{cases} e_i & \text{if } x_i = 1, \end{cases}$	If we put <i>e</i> into a crossing-free subset, we discard <i>c</i> crossings. If we don't put <i>e</i> into a c-f subset, we discard <i>c</i> crossings.
$H:=\bigcup_{i=1}^k \left\{ f_i \text{if } x_i=2, \right.$	$\Rightarrow E_X$ has $S(k) \le 2S(k-c) \le 2S(k-1) \le 2^k$ max. c-f subsets
Ø else.	For each such subset <i>H</i> check:
is <i>H</i> crossing-free? does $G' \cup H$ span G ?	does $G' \cup H$ span G ?
	2 0 0
Alexander Wolff 34 42 Coometric Network	s Alexander Wolff 35 42 Coometric Networks



Drawing Metro Maps es with few crossings MIP Formulation

A MIP for MinimumCrossingSpanningSubgraph
For each edge *e* of *G* let

$$y_e = \begin{cases} 1 & \text{if } e \text{ is in the desired spanning subgraph of } G, \\ 0 & \text{otherwise.} \end{cases}$$
We would like to have for each crossing $\{e, e'\} \in X$ a variable

$$x_{ee'} = \begin{cases} 1 & \text{if } y_e = y_{e'} = 1, \\ 0 & \text{otherwise.} \end{cases}$$
Then we could state as objective function

Then we could state as objective function

minimize
$$\sum_{\{e,e'\}\in X} x_{ee'}$$

Clearly $x_{ee'} = y_e \cdot y_{e'}$, but this is not linear!

exander Wolff

Introduction Drawing Metro Maps Spanning trees with few crossings	Nun-ryprovintiality Fixed-Parameter Algorithms MIP Formulation
Expressing x _{ee'}	
Reminder: we want $x_{ee'} = \begin{cases} 1 \\ 0 \end{cases}$	if $y_e = y_{e'} = 1$, otherwise.
We can also express this as xee	$y' = \min\{y_e, y_{e'}\}.$
So what about $x_{ee'} \le y_e$ and Recall our objective function: m	$egin{array}{lll} x_{ee'} \leq y_{e'} & (ext{and } x_{ee'} \geq 0)? \ \end{array}$ inimize $\sum_{\{e,e'\}\in X} x_{ee'} & :-(\end{array}$
Now here's the trick: $x_{ee'} \ge 0$	and $x_{ee'} \ge y_e + y_{e'} - 1$:-)
To do: must make sure that ea	dges <i>e</i> with $y_e = 1$ span <i>G</i> !
Idea: use flow to achieve cor	nnectivity! + skip one! + skip two!
Alexander Malff 20	40 Coomotión Maturatza



$$\sum_{(v,u)\in A} f_{(v,u)} = -1 + \sum_{(u,v)\in A} f_{(u,v)}$$

 \Rightarrow the set of arcs with non-zero flow is connected!

lexander Wolff

Htroduction Non-Approximability Drawing Metro Maps Fixed-Parameter Algorithms Spanning trees with few crossings MIP Formulation
From paths to connectivity
Choose arbitrary <i>center</i> vertex <i>s</i> . Go through all $t \neq s$.
Idea: establish simple <i>s</i> - <i>t</i> path π^t . Let $p_e^t = \begin{cases} 1 & \text{if } e \in \pi^t \\ 0 & \text{else.} \end{cases}$
Make sure π^t leaves <i>s</i> and reaches <i>t</i> :
$\sum_{e \text{ incident to } s} p_e^t = \sum_{e \text{ incident to } t} p_e^t = 1.$
Make sure π^t does not stop in any vertex v other than s and t :
$\{0,2\}$ $\ni \sum_{e \text{ incident to } v} p_e^t = 2d_e^t \text{ with } d_e^t \in \{0,1\}.$
Synchronize y_e with all variables of type p_e^t :
lexander Wolff 41 42 Coometric Networks



Geometric embeddings and graph expansion

James R. Lee

jrl@ias.edu Institute for Advanced Study (Princeton) and University of Washington (Seattle)

Abstract

Beginning in the early-90's, it became gradually apparent that techniques from high-dimensional geometry were highly relevant for a variety of computational tasks. At the center of this connection is the study of *graph partitioning problems* which involve breaking a graph into two or more large parts while minimizing the size of the "interface" between them. These problems are of central importance in numerous computational settings like data clustering, divide and conquer algorithms, and packet routing in networks. Additionally, the techniques involved have applications to areas like Markov chains, nearest-neighbor search, and learning theory.

We will discuss an approach to these problems based on *geometric embeddings*. The basic idea is to represent a graph G = (V, E) as a geometric object by mapping the vertices of G to points in a high-dimensional Euclidean space, and then to use geometric techniques to find good cuts in G—these are cuts which separate G into "large" pieces while minimizing the number of edges that are cut.

In the Sparsest Cut problem, the input is given as a graph G = (V, E) along with a subset of pairs of vertices $D \subseteq V \times V$ called *demands*. For a subset $S \subseteq V$, let $E(S, \bar{S})$ be the set of edges crossing from S to its complement \bar{S} , and let $D(S, \bar{S})$ be the set of demands crossing from S to \bar{S} . The sparsity of a cut (S, \bar{S}) is defined as the ratio

$$\Phi(S) = \frac{|E(S,\bar{S})|}{|D(S,\bar{S})|},$$

and the goal of the Sparsest Cut problem is to find the cut (S, \overline{S}) which minimizes $\Phi(S)$. Since solving this problem exactly is NP-hard, we will instead search for a set $S \subseteq V$ for which $\Phi(S)$ is *approximately optimal*, i.e. is within some factor C of the sparsest cut.

During the talk, I will first explain the relationship between geometric embeddings and the Sparsest Cut problem, and then I will discuss the current state-of-the-art techniques in constructing such embeddings. These employ a beautiful mix of semi-definite programming, high-dimensional convex geometry, probability theory, and combinatorics. A brief outline of the talk (with references) now follows.

The relation of geometric embeddings to the Sparsest Cut problem was first realized in the two papers [LLR95, AR98]. The basic approach follows.

1. Associate a *metric* to the graph G. This is a value d(u, v) for every pair of vertices $u, v \in V$ which is symmetric, i.e. d(u, v) = d(v, u) and satisfies the triangle inequality

$$d(u,v) \le d(u,w) + d(w,v)$$

for every triple of points $u, v, w \in V$. This metric is closely related to the structure of *multi-commodity flows in G*, and is the solution to a linear programming relaxation of the Sparsest Cut problem.

2. Given the metric d, we can think of (V, d) as a metric space. The second step is to *embed* this space into \mathbb{R}^n (where n = |V|) such that the embedding preserves the structure of (V, d). Such a mapping is called a *low-distortion embedding*. This is a map $f: V \to \mathbb{R}^n$ which satisfies

 $d(u,v) \le \|f(u) - f(v)\| \le C \cdot d(u,v) \qquad \text{for all } u, v \in V.$

The factor C is called the *distortion* of the map f, and our goal is to make C as small as possible. (Note that ||f(u) - f(v)|| is the Euclidean distance between f(u) and f(v).)

3. The final step is to use the geometric embedding to find a cut in G. One way to do this is by projecting the image of G (under the embedding f) onto a randomly oriented line. If we number the points of G from left to right on the line $1, 2, \ldots, n$ then there are n - 1 cuts of the form $(\{1, 2, \ldots, i\}, \{i + 1, i + 2, \ldots, n\})$. The final step of the algorithm returns the cut (S, \overline{S}) among these which has the smallest sparsity $\Phi(S)$.

Remarkably, it turns out that with high probability, the cut we return after step (3) will be within a factor C of optimal, where C is the distortion of the embedding from step (2)!

After discussing this general approach, we will examine the new techniques that achieve the best-known approximation factor for Sparsest Cut. These are based on semi-definite programming and novel high-dimensional arguments from [ARV04] (see also [Lee05, CGR05]), as well as new techniques for constructing low-distortion embeddings [KLMN05, Lee05, ALN05].

References

- [ALN05] S. Arora, J. R. Lee, and A. Naor. Euclidean distortion and the Sparsest Cut. In 37th Annual Symposium on the Theory of Computing, pages 553-562, 2005. Full version available at http://www.cs.berkeley.edu/~jrl.
- [AR98] Yonatan Aumann and Yuval Rabani. An O(log k) approximate min-cut max-flow theorem and approximation algorithm. SIAM J. Comput., 27(1):291–301 (electronic), 1998.
- [ARV04] Sanjeev Arora, Satish Rao, and Umesh Vazirani. Expander flows, geometric embeddings, and graph partitionings. In 36th Annual Symposium on the Theory of Computing, pages 222–231, 2004.
- [CGR05] S. Chawla, A. Gupta, and H. Räcke. Embeddings of negative-type metrics and an improved approximation to generalized sparsest cut. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 102–111, Vancouver, 2005.

- [KLMN05] R. Krauthgamer, J. R. Lee, M. Mendel, and A. Naor. Measured descent: A new embedding method for finite metrics. *Geom. Funct. Anal.*, 15(4):839–858, 2005.
- [Lee05] James R. Lee. On distance scales, embeddings, and efficient relaxations of the cut cone. In SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms, pages 92–101, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.
- [LLR95] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.

Distance Trisector curve and Voronoi diagram with neutral zone

Takeshi Tokuyama (Tohoku U) joint work with Tetsuo Asano (JAIST) and Jiri Matousek (Charles U)

Grand Challenge of NHC project

- Throw right on Computational Barriers
 - Identifying barriers (something looking impossible)
 - Breaking/proving/avoiding barriers
 - Making properties of barriers clearer
 - Well-known barriers in TCS
 - Computability
 - P vs NP, NP vs PSPACE etc
 - Approximation hardness, approximation ratio
 - Randomness (P vs PP, primarity checking, Yao's minmax principle)
 Lower bounds for online algorithms
 - Cryptographic barriers (Discrete log etc)
 - Cryptographic barriers (Discrete log etc)
 - Hopefully, find some more new barriers, because they will help progress of TCS

Classical barriers in math history

Let no one ignorant of Mathematics enter here. (plato)

- Incomputable geometric/arithmetic problems
 - Compute diagonal length of a unit square
 - Find a cube of volume 2
 - Dilemma of Pythagoreans
 - Draw a circle with the unit area
 It is easier to square the circle than to get round a mathematician.
 - Draw regular n-gon
 - Trisect a given angle



Tool to resolve a small but curious barrier.

Story of this talk

- In a computational geometric problem, we find a simple and natural geometric tool named *distance trisector curve*.
- A new transcendental curve??
- Some initial results have been obtained
- Existence and uniqueness of the trisector curves
- An algorithm to compute it.
- Use of the trisector curves.
- A new possibility in computational geometry
- Relaxed "computability" of geometric problems.

Voronoi diagram

- Subdivision of plane (space) into cells
 - $S = \{p_1, p_2, \dots p_n\}$ points in the plane
 - V(p_i) = { x : d(x, p_i) < d(x, p_j) for all j ≠ i}
 Voronoi cell: dominating region of p_i
- Great geometric structure with many applications
 - Mesh generation, Graphics,
 - Simulating economic/political equilibrium
 - Simulating biological cells / crystallization
- Efficient algorithm : O(n log n) time
- Many variants: VD of lines/discs/regions, higher dimension, non-Euclidean metric, power diagrams.





Motivation from nature

- We often see beautiful patterns with "objects" and "neutral zone" in natural structures
- Can we have a variation of Voronoi diagram to draw such a picture?



Voronoi diagram with neutral zone

- Classical Voronoi diagram: Partition the plane (space) into domination regions of input points.
- Idea: Partition the space (plane) into domination regions and a *neutral zone*
 - There may be many mathematical formulations
 - We adopt a natural formulation



Existence and uniqueness

- P=(0,1), Q=(0,-1)
- $y=f_1(x)$ is x-axis (bisector of PQ)
- $y = g_1(x)$:bisector parabola of P and x-axis
- y=f_j(x): bisector curve of P and y=-g_{j-1}(x)
- $y=g_i(x)$: bisector curve of P and $y=-f_i(x)$

Lemma. The trisecting curve C_P must be above $f_i(x)$ and below $g_i(x)$

Idea for computing F(x)

- If -t(x) is the x-value of the partner point,
- $t(x) < \beta x$ for a constant $\beta < \sqrt{3} 1 < 1$
- We can compute the Taylor expansion of F(x) and t(x) around zero to any precision.
 - F(x) and t(x) are computable if x < h for a small h.
- F(x) and x are computable from t(x), t(t(x)), F(t(x)), and F(t(t(x))).
- Guess z = t^k(x) < h, compute F(z) and F(t(z)), and recursively compute tⁱ(x) for i=k, k-1, k-2, ..., 0
- Determine correct value of z via binary search

Formulas to compute F(x)

- $x = \Phi(t(x), t(t(x)), F(t(x)), F(t(t(x))))$
- $F(x) = \psi(t(x), t(t(x)), F(t(x)), F(t(t(x))))$
- $\Phi(x,y,u,w) = x + y(x^2 + (1+u)^2)/2Q(x,y,u,w)$
- $\psi(x,y,u,w) = (2xyu + (1+u)(1+x^2-u^2))/2Q(x,y,u,w)$
- Q(x,y,u,w) = (1+u)(1+v)-xy

The above mentioned: STOC06, to appear

Voronoi edges

- Definition: If p is a boundary point on NV(p_i) and q is its nearest point among other region boundaries, we call q the partner point of p.
- Definition: The boundary of NV(p_i) is decomposed into curve segments each of which consists of points whose partner points are in a same region. The connected components are called Voronoi edges.

Questions

- N-Voronoi diagram always exists? – Yes
- Unique for a given point set? -Yes
- Efficient algorithm exist? - Yes, if we are given some oracles
- Really efficient in practice? – No.....

Given a set of regions $R_1, R_2, \dots R_n$ such that

$$\begin{aligned} & \mathsf{R}_{j} \text{ contains } \mathsf{p}_{j}, \text{ consider an operator } \boldsymbol{F} \\ & \boldsymbol{F}(\mathsf{R}_{1}, \mathsf{R}_{2}, \dots, \mathsf{R}_{n}) = (\mathsf{Q}_{1}, \mathsf{Q}_{2}, \dots, \mathsf{Q}_{n}) \\ & \text{where } \mathsf{Q}_{j} = \{ x: \ \mathsf{d}(x, \mathsf{p}_{j}) < \mathsf{d}(x, y) \text{ for } \ \mathsf{d}(x, y) \} \end{aligned}$$

 $\forall y \in U_{i \neq i}^{\prime} R_{i} \}$

Theorem. F has at least one fixed point

From Schauder-Kakutani's fixed point theorem

Corollary. N-Voronoi diagram is given as a fixed point of F

Uniquness

- Fixed point theorem does not assure uniqueness.
- Uniqueness is given in a constructive fashion
- Crystallization algorithm:
 - Growing radius of disks
 - Analogous to the space-sweep algorithm for computing a (classical) Voronoi diagram as a lower envelope of parabolic cylinders.

Curves in a N-Vonoroi diagram

Property 1: N-Voronoi diagram is a subset of an arrangement of curves in a curve family F

F consists of :

•Distance trisectors of pairs of input points are in F.

•Bisector curves of an input point and a curve (or point) C in F.

• Intersection points of curves in F.

Property 2: The number of applications of bisector operations is finite to obtain F

Analysis of crystallization algorithm

- For planar case, the algorithm shows uniqueness of N-Voronoi diagram
- Number of "structural changes" is O(n)
- Terminates in finite steps if we can draw
 - 1. bisector curve of a point and given point
 - 2. distance trisector curve
- Terminates in polynomial time if we assume that we can draw an "generalized" trisector curve

